

Rogue Apps, Hidden Web Tracking and Ubiquitous Sensors

Investigations into the Complexity of Modern Privacy Threats

vorgelegt von

M.Sc.

Jacob Leon Kröger

ORCID: 0000-0003-3559-8869

an der Fakultät IV - Elektrotechnik und Informatik

der Technischen Universität Berlin

zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften

- Dr. rer. nat. -

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr.-Ing. Uwe Nestmann

Gutachterin: Prof. Dr. Bettina Berendt

Gutachterin: Prof. Dr.-Ing. Ina Schieferdecker

Gutachter: Prof. Dr.-Ing. Carsten Trinitis

Gutachter: PD Dr. Sebastian Pape

Tag der wissenschaftlichen Aussprache: 9. Juni 2022

Berlin 2022

Zusammenfassung

Angesichts der heutzutage allgegenwärtigen Überwachung erscheinen Datenschutzbestrebungen häufig wie ein aussichtsloses Unterfangen. Um die Komplexität des Problems zu veranschaulichen und zur Erschließung realistischer Lösungswege beizutragen, befasst sich diese kumulative Dissertation mit Bedrohungen für unsere Privatsphäre, die von mobilen Apps, Web-Tracking und eingebetteten Sensoren ausgehen. Nachdem ein einleitendes Kapitel die Motivation für meine Forschung darlegt und theoretische Hintergründe beleuchtet, besteht die vorliegende Arbeit aus vier Teilen:

Teil I untersucht die Auswirkungen, die Sensoren in Alltagsgeräten auf unsere Privatsphäre haben. Zunächst geben drei Studien – in der Literatur erstmals in diesem Umfang – einen Überblick über die Vielfalt an persönlichen Informationen, die sich aus Eye-Tracking-Daten, Sprachaufnahmen und aus den Daten von Beschleunigungssensoren ableiten lassen. Anschließend untersucht eine Studie die Wahrnehmungen von Nutzer/-innen bezüglich der Möglichkeiten moderner Sprach- und Stimmanalyse. Das nachfolgende Kapitel befasst sich mit der Durchführbarkeit und Nachweisbarkeit von Smartphone-basierten Lauschangriffen, wobei Beschleunigungssensoren als möglicher Abhörkanal beleuchtet werden. Zuletzt präsentiert ein Kapitel Kategorien von persönlichen Informationen, die sich aus Videospiel-Daten (inkl. Sensordaten aus Gaming-Equipment) ableiten lassen.

Teil II konzentriert sich auf die Datenpraktiken mobiler Apps. Es wird eine vierjährige Undercover-Studie vorgestellt, bei der geprüft wurde, inwiefern App-Anbieter die in der EU vorgeschriebenen Transparenzpflichten einhalten. Obwohl das Gesetz Verbraucher/-innen das Recht auf Auskunft über ihre personenbezogenen Daten einräumt, zeigt die Studie, dass die Ausübung dieses Rechts in der Praxis mit erheblichen Hindernissen verbunden ist.

Teil III stellt zwei neue Ansätze vor, um Web-Tracking aufzudecken und erfahrbar zu machen. Zunächst wird ein Browser-Plugin präsentiert, das Browsersitzungen aufzeichnet, um Trainingsdaten für die automatische Erkennung von Web-Tracking zu sammeln. Künstliche Daten, die üblicherweise für diesen Zweck verwendet werden, haben diverse Nachteile, die durch die Verwendung realer Daten überwunden werden können. Anschließend erkundet ein Kapitel, wie Web-Tracking „sonifiziert“, d.h. für Nutzer/-innen durch Klänge und Melodien hörbar gemacht, werden kann. Zuletzt werden in Bezugnahme auf Teil I Vorschläge gemacht, wie das Spektrum der persönlichen Informationen, die sich aus verschiedenen Arten von Sensordaten ableiten lassen, ebenfalls digital registriert und interaktiv erfahrbar gemacht werden könnte.

Teil IV wirft einen kritischen Blick auf das Prinzip, dass Individuen per Einwilligung über die Verwendung ihrer persönlichen Daten entscheiden sollen („privacy self-management“). Basierend auf einer ganzheitlichen Untersuchung der Nachteile dieses Ansatzes wird argumentiert, dass die Selbstverwaltung persönlicher Daten in der Praxis nicht funktioniert und somit eine erhebliche Lücke im Datenschutzrecht darstellt.

Diese Dissertation unterstreicht die Notwendigkeit, neue Wege zu finden, um mit den exzessiven und obskuren Formen der Überwachung in unserer heutigen Gesellschaft umzugehen. Sie trägt zur wissenschaftlichen Literatur über mögliche Bedrohungen der Privatsphäre im Bereich der Verbraucherelektronik bei und zeigt das Versagen der aktuellen Gesetze zum Schutz unserer Privatsphäre auf. Die Dissertation schließt mit einer Diskussion von übergreifenden Themen. Zwar wird die Gefährlichkeit und Komplexität der untersuchten Eingriffe in unsere Privatsphäre hervorgehoben, doch wird eine allgemeine *privacy-is-dead*-Haltung entschieden zurückgewiesen. Vor dem Hintergrund der Forschungsergebnisse werden Politikempfehlungen und Ideen für zukünftige Forschung präsentiert.

Abstract

In the face of ubiquitous surveillance and people's loss of control over their personal data, informational privacy is widely perceived as irretrievably "dead". Illustrating the complexity of the problem and contributing to the search for realistic solutions, this cumulative dissertation deals with privacy threats posed by mobile apps, web tracking, and embedded sensors. The dissertation starts with an introductory chapter that establishes the motivation and introduces the theoretical background, followed by four parts:

Part I focuses on privacy threats posed by sensors embedded into consumer devices. First, three studies provide an overview of the rich variety of personal information that can be inferred from eye-tracking data, accelerometer data, and voice recordings. Second, a study on users' perceptions about the privacy impacts of voice and speech analysis is presented. Third, the feasibility and detectability of smartphone-based eavesdropping is investigated, addressing accelerometers as a possible eavesdropping channel. Fourth, the privacy-invading potential of video games and their associated sensor-equipped hardware is explored.

Part II focuses on data practices of mobile apps. An undercover investigation is presented, probing whether app vendors comply with transparency obligations prescribed by EU's General Data Protection Regulation. While the law grants consumers the right to access the personal data that companies hold about them, the study reveals severe obstacles to exercising this right in practice.

Part III presents two novel approaches for the detection and exposure of hidden web tracking. First, a browser extension is proposed that records internet browsing sessions to obtain training data for automated web-tracking detection. Artificial data, which is commonly used for this purpose, has severe drawbacks that can be overcome by using real-world browsing data. Second, methods are explored to "sonify" web-tracking activity, i.e., make it audible to internet users through indicative sounds and melodies. Furthermore, in reference to topics covered in Part I, suggestions are provided as to how the range of personal information that can be inferred from different types of sensor data could be recorded in a digital database in order to be presented in an interactive and updatable form.

Part IV sheds a critical light on the legal principle that people individually manage their privacy via notice and choice ("privacy self-management"), drawing on findings from the previous parts and related literature. Based on a holistic examination of its limitations, it is argued that privacy self-management does not function in practice, amounting to a major loophole in privacy law.

This dissertation emphasizes the need for new ways of dealing with the excessive and obscure forms of surveillance prevalent in modern life. It adds to the academic debate and scientific literature about possible privacy threats emerging from consumer electronics as well as to exposing the failure of current laws to protect our privacy. The dissertation concludes with a discussion of overarching themes. While emphasizing the seriousness and complexity of the privacy threats under investigation, a general privacy-is-dead attitude is firmly rejected. In light of the findings, policy recommendations and possible avenues for future research are presented.

Acknowledgements

Due to the inspiration and support I have received in both private and professional life, working on this dissertation has been a deeply fulfilling and joyful process.

First and foremost, I want to express my gratitude to Prof. Dr. Bettina Berendt and Prof. Dr.-Ing. Ina Schieferdecker for supervising my Ph.D. project. Thank you for your continuous trust and encouragement, for the many hours of undivided attention despite your busy schedules, and for your clear guidance, including during the unforeseen COVID-19 pandemic. To me and many others, you are invaluable leaders and role models.

I am also eternally indebted to Dr. Stefan Ullrich for his authentic and innovative way of leading our research group at the *Weizenbaum Institute for the Networked Society*. Stefan created a wonderful research environment, allowing everyone involved to follow their strengths and interests. Besides brilliant mentorship and intellectual input, Stefan provided valuable contacts across academia which significantly facilitated my work. This includes joint projects and stimulating exchanges with PD Dr. Sebastian Pape, Prof. Dr. Dominik Herrmann, Prof. Dr.-Ing. Carsten Trinitis, and Jens Lindemann, all of whom I would like to thank for their commitment, enthusiasm, and incredibly helpful feedback. I am grateful that Carsten and Sebastian have kindly agreed to be part of my Ph.D. jury.

It was a great honor to be among the first generation of doctoral researchers at the *Weizenbaum Institute*. My wholehearted thanks go to the expert consortium that developed the idea for the institute and successfully secured the first round of governmental funding in 2017. They have seized an excellent opportunity to support public interest research into the benefits and potential dangers of digital transformation. Independent research in this area is more urgently needed today than ever, and should by no means be taken for granted.

To all the wonderful people I had the pleasure of calling my colleagues and collaborators: Thank you so much for the educational, fun, exciting, productive, and deeply inspiring time we spent chatting, laughing, and – of course – working together. Beyond the individuals highlighted above, special mentions go to Philip Raschke, Otto Hans-Martin Lutz, Florian Müller, Milagros Miceli, Leon Gellrich, Saba Rebecca Brause, Andrea Hamm, Hans-Christian Gräfe, Rainer Rehak, and Jessica Percy Campbell.

I would also like to express my appreciation to the staff at the *Weizenbaum Institute* and at *TU Berlin* who generously supported my work by providing assistance with administrative affairs and public relations. Special thanks go to Evelyn Adams, Jana Peich, Filip Stiglmayer, Katharina Stefes, Johanna Hampf, Prof. Dr.-Ing. Uwe Nestmann, and Daniela Wroblewski.

My deepest heartfelt gratitude goes out to my beautiful family and friends. Their love and support have given me the strength to take on new challenges throughout my Ph.D. project and the confidence to stay true to my interests, even when this meant investigating and commenting on politically charged topics. Sophie and Ekkehard, Martin and Clarisse, Cora and Yéléna, *Oma* Jutta and *Großmama* Elisabeth, thank you for being my home and my light in these exciting times. You helped me to never lose hope and optimism despite the staggering and often alarming findings of my research. Yannis, Paul, Pablo, Toya, Soufiane, Iskender, Robin, Vincent, Valerio, Katja, Juli, Minu, George, Isabelle, Stefan, Olli, Simon, Talisa, and Marie – thank you all for your precious friendship, and for constantly inspiring me with your creativity, diverse perspectives, and colorful life experiences. Having you all in my life is the greatest blessing I could have ever asked for.

Among the people who make my heart shine, a special mention belongs to you, Maj. Knowing you has not only opened a new chapter in my life but also entirely new worlds to me. You have brought me closer to nature, mindfulness, and gratitude – to the things that really count in life. Our adventurous travels have been the perfect balance to my work as a doctoral researcher and are a source of unforgettable memories. Thank you for sharing your rare wisdom with me and others. The world needs more people like you.

Finally, in loving memory, I would like to dedicate this dissertation to *Opa* Ulrich Kröger, *Großpapa* Bernard Krebs, and my dear friend Lea Steinhardt – three beautiful souls that will continue to live in my heart forever.

Table of Contents

| | |
|--|--------------|
| Title Page | i |
| Zusammenfassung | iii |
| Abstract | v |
| List of Figures | xvii |
| List of Tables | xix |
| List of Included Papers | xxi |
| List of Abbreviations | xxiii |
| Chapter 1: Introduction | 1 |
| 1.1 Motivation: Privacy in Crisis | 1 |
| 1.1.1 Privacy Threats Posed by Ubiquitous Sensors | 3 |
| 1.2 Theoretical Background | 4 |
| 1.2.1 A Brief History of Privacy Discourse | 5 |
| 1.2.2 The Subject of Debate | 8 |
| 1.2.3 Information Privacy Research Areas | 10 |
| 1.2.3.1 The Types of Information Collected About Individuals | 10 |
| 1.2.3.2 Data Technologies and Applications | 11 |
| 1.2.3.3 Privacy Regulation | 12 |
| 1.2.3.4 People's Information Privacy Concerns | 12 |
| 1.2.3.5 Data Controllers' Information Privacy Practices | 13 |
| 1.2.3.6 Information Privacy Tools and Technologies | 13 |
| 1.3 Thesis Outline | 14 |
| I PRIVACY-INVADING POTENTIAL OF SENSOR DATA | 17 |
| Chapter 2: Voice Recordings, Eye Tracking, and Accelerometer Data | 19 |
| 2.1 Background and Motivation | 19 |
| 2.2 Research Scope and Limitations | 21 |
| P1 Paper 1: Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference | 25 |

TABLE OF CONTENTS

| | | |
|---------|--|----|
| P1-1 | Introduction | 25 |
| P1-2 | Inference of Personal Information from Voice Recordings | 26 |
| P1-2.1 | Speaker Recognition | 26 |
| P1-2.2 | Inference of Body Measures | 27 |
| P1-2.3 | Mood and Emotion Recognition | 28 |
| P1-2.4 | Inference of Age and Gender | 28 |
| P1-2.5 | Inference of Personality Traits | 29 |
| P1-2.6 | Deception Detection | 30 |
| P1-2.7 | Detection of Sleepiness and Intoxication | 30 |
| P1-2.8 | Accent Recognition | 30 |
| P1-2.9 | Speaker Pathology | 31 |
| P1-2.10 | Mental Health Assessment | 32 |
| P1-2.11 | Prediction of Interpersonal Perception | 32 |
| P1-2.12 | Inference of Socioeconomic Status | 33 |
| P1-2.13 | Classification of Acoustic Scenes and Events | 33 |
| P1-3 | Discussion and Implication | 34 |
| P1-4 | Conclusion | 36 |
| P1-5 | References | 36 |
| P2 | Paper 2: Privacy Implications of Accelerometer Data: A Review of Possible Inferences | 43 |
| P2-1 | Introduction | 43 |
| P2-2 | Possible Inferences | 43 |
| P2-2.1 | Activity and Behavior Tracking | 43 |
| P2-2.2 | Location Tracking | 44 |
| P2-2.3 | User Identification | 44 |
| P2-2.4 | Keystroke Logging | 45 |
| P2-2.5 | Inference of Health Parameters and Body Features | 45 |
| P2-2.6 | Inference of Demographics | 45 |
| P2-2.7 | Mood and Emotion Recognition | 45 |
| P2-2.8 | Inference of Personality Traits | 45 |
| P2-3 | Discussion and Implications | 46 |
| P2-4 | Conclusion | 47 |
| P2-5 | References | 47 |
| P3 | Paper 3: What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking | 51 |
| P3-1 | Introduction | 51 |
| P3-2 | Inference of Personal Information from Eye Tracking Data | 52 |
| P3-2.1 | Biometric Identification | 53 |
| P3-2.2 | Monitoring of Mental Workload and Cognitive Processes | 54 |
| P3-2.3 | Inference of Personality Traits | 54 |
| P3-2.4 | Inference of Cultural Affiliation and Ethnicity | 55 |
| P3-2.5 | Skill Assessment | 56 |
| P3-2.6 | Age and Gender Recognition | 56 |

| | | |
|---|--|-----------|
| | P3–2.7 Inference of Preferences and Aversions | 57 |
| | P3–2.8 Detection of Short- and Medium-Term User States | 57 |
| | P3–2.9 Health Assessment | 58 |
| | P3–3 Discussion and Implications | 59 |
| | P3–4 Conclusion | 61 |
| | P3–5 References | 62 |
| Chapter 3: User Knowledge and Perceptions about Sensor-based Inference | | |
| | Attacks | 67 |
| 3.1 | Background | 67 |
| P4 | Paper 4: Personal Information Inference from Voice Recordings: User Awareness and Privacy Concerns | 69 |
| P4–1 | Introduction | 69 |
| P4–2 | Related Work | 70 |
| | P4–2.1 Privacy Perceptions and Concerns about Audio Recording . . | 70 |
| | P4–2.2 Sensor-Based Inference of Personal Information | 70 |
| | P4–2.3 User Awareness and Perceptions about Sensor-Based Inference Attacks | 71 |
| P4–3 | Related Work | 72 |
| P4–4 | Research Methodology | 73 |
| | P4–4.1 Survey Instrument | 73 |
| | P4–4.2 Participant Recruitment | 74 |
| | P4–4.3 Sample Characteristics | 75 |
| | P4–4.2 Data Analysis | 75 |
| P4–5 | Results | 76 |
| | P4–5.1 How aware are people that personal information can be inferred from voice recordings? | 76 |
| | P4–5.2 How does the level of awareness differ across demographic groups? . | 76 |
| | P4–5.3 What concerns do people have about the inference of personal information from voice recordings? | 78 |
| | P4–5.4 How do people’s usage intentions for voice-enabled devices change when being informed on the topic? | 80 |
| P4–6 | Discussion and Implications | 81 |
| | P4–6.1 Consumer Education and Privacy-Enhancing Technologies . . | 82 |
| | P4–6.2 Regulatory Implications | 83 |
| P4–7 | Limitations | 83 |
| P4–8 | Conclusion | 84 |
| P4–9 | Acknowledgments | 84 |
| P4–10 | References | 85 |
| P4–11 | Appendix | 88 |
| | P4–11.A Appendix A: Survey Questionnaire | 88 |
| | P4–11.B Appendix B: Correlation Table | 90 |

Chapter 4: Special Cases and Application Areas of Sensor-based Inferences 91

| | | |
|-------|--|-----|
| 4.1 | Background: The Spies in Our Pockets | 91 |
| P5 | Paper 5: Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping | 95 |
| P5-1 | Introduction | 95 |
| P5-2 | Threat Model | 97 |
| P5-3 | Microphone-Based Eavesdropping | 98 |
| | P5-3.1 Microphone Access Permission | 98 |
| | P5-3.2 User Notifications and Visibility | 99 |
| P5-4 | Motion Sensor-Based Eavesdropping | 99 |
| | P5-4.1 Experimental Research Findings | 100 |
| | P5-4.2 Sampling Frequency Limits | 100 |
| | P5-4.3 Sensor Access Permissions and Energy Efficiency | 101 |
| P5-5 | Existing Mitigation and Detection Techniques | 101 |
| | P5-5.1 App Inspections Conducted by Ecosystem Providers | 101 |
| | P5-5.2 App Inspections Conducted by the Research Community | 102 |
| P5-6 | Ecosystem Providers as Potential Adversaries | 104 |
| P5-7 | Technical and Economic Feasibility | 104 |
| P5-8 | Unauthorized Access to Smartphones | 106 |
| P5-9 | Discussion | 107 |
| P5-10 | Conclusion | 108 |
| P5-11 | References | 109 |
| 4.2 | Background: Video Games and Privacy | 114 |
| P6 | Paper 6: Surveilling the Gamers: Privacy Impacts of the Video Game Industry | 115 |
| P6-1 | Introduction | 115 |
| P6-2 | Data Categories Collected by Video Game Companies | 117 |
| P6-3 | Inference of Personal Information from Gameplay Data | 118 |
| | P6-3.1 User Identification | 119 |
| | P6-3.2 Age and Gender Recognition | 119 |
| | P6-3.3 Emotion Recognition | 120 |
| | P6-3.4 Skill Assessment | 121 |
| | P6-3.5 Inference of Interests and Preferences | 121 |
| | P6-3.6 Inferences about Financial Status and Consumption Behavior | 122 |
| | P6-3.7 Inference of Personality Traits | 123 |
| P6-4 | Sensor-based Inference of Personal Information | 124 |
| P6-5 | Discussion and Implications | 125 |
| P6-6 | Limitations | 127 |
| P6-7 | Conclusion | 128 |
| P6-8 | References | 128 |

| | | |
|-------------------|---|------------|
| II | PRIVACY PRACTICES OF MOBILE APP VENDORS | 139 |
| Chapter 5: | GDPR Compliance of Mobile Apps | 141 |
| 5.1 | Background | 141 |
| P7 | Paper 7: How do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on iOS and Android Apps | 143 |
| P7-1 | Introduction | 143 |
| P7-2 | Related Work | 144 |
| P7-3 | Data Collection | 144 |
| P7-4 | Methodology for Interaction with the App Vendors | 145 |
| P7-5 | Analysis and Results | 145 |
| P7-5.1 | Overview of the Received Responses | 145 |
| P7-5.2 | Insufficient Responses | 146 |
| P7-5.3 | Data Sharing with Third Parties | 147 |
| P7-5.4 | Identity Verification and Security of Transmitted Data | 148 |
| P7-5.5 | Comparison by Vendor's Residence | 148 |
| P7-5.6 | Deceptive and Misleading Responses | 149 |
| P7-5.7 | Discontinued Apps and Accounts | 149 |
| P7-6 | Discussion | 150 |
| P7-7 | Ethical Considerations and Limitations | 151 |
| P7-8 | Conclusion | 151 |
| P7-9 | References | 152 |
| III | WEB TRACKING DETECTION AND SONIFICATION | 153 |
| Chapter 6: | Training Data for Automated Web Tracking Detection | 155 |
| 6.1 | Background | 155 |
| P8 | Paper 8: Towards Real-Time Web Tracking Detection with T.EX – The Transparency EXtension | 157 |
| P8-1 | Introduction | 157 |
| P8-2 | Objectives and Requirements | 159 |
| P8-3 | Limitations | 160 |
| P8-4 | Related Work | 162 |
| P8-5 | Implementation | 164 |
| P8-5.1 | HTTP and HTTPS Traffic Logging and Recording | 164 |
| P8-5.2 | Persistent Storage of Records | 165 |
| P8-5.3 | Encryption and Decryption of Chunks | 166 |
| P8-5.4 | Data Visualization | 166 |
| P8-6 | Evaluation | 167 |
| P8-7 | Conclusion and Outlook | 170 |
| P8-8 | References | 170 |

| | | |
|-------------------|---|------------|
| Chapter 7: | Making Web Tracking Audible | 173 |
| 7.1 | Background | 173 |
| P9 | Paper 9: Surfing in Sound: Sonification of Hidden Web Tracking | 175 |
| P9-1 | Introduction | 175 |
| P9-2 | Related Work | 175 |
| P9-3 | Framework Design | 176 |
| | P9-3.1 Implementation | 176 |
| | P9-3.2 Sound Design | 176 |
| | P9-3.3 Comparison to Existing Approaches | 177 |
| P9-4 | Evaluation | 177 |
| | P9-4.1 Study Design and Hypothesis | 177 |
| | P9-4.2 Results | 177 |
| P9-5 | Discussion | 178 |
| P9-6 | Future Research | 178 |
| P9-7 | Acknowledgment | 178 |
| P9-8 | References | 178 |
| Chapter 8: | Concept: Inference Mapping Tool | 179 |
| 8.1 | Background | 179 |
| 8.2 | Proposed Functionalities | 180 |
| | 8.2.1 Possible Extensions | 180 |
| 8.3 | Expected Benefits and Challenges | 182 |
| IV | A CRITICAL TAKE ON PRIVACY REGULATION | 185 |
| Chapter 9: | Challenging the Notice-and-Choice Approach to Privacy | 187 |
| 9.1 | Background and Motivation | 187 |
| P10 | Paper 10: The Myth of Individual Control: Mapping the Limitations of Privacy Self-Management | 191 |
| P10-1 | Introduction | 191 |
| P10-2 | Obstacles to Informed and Rational Privacy Choices | 193 |
| | P10-2.1 Time Constraints | 193 |
| | P10-2.2 Lack of Knowledge | 193 |
| | P10-2.3 Nature of Human Decision-Making | 194 |
| | P10-2.4 Information Overload | 194 |
| | P10-2.5 Lack of Transparency | 194 |
| | P10-2.6 Obstacles to Presenting Privacy Information | 195 |
| P10-3 | Obstacles to Voluntary Privacy Decisions | 195 |
| | P10-3.1 Dependence on Services | 195 |
| | P10-3.2 Nudging and Coercion | 196 |
| | P10-3.3 Financial Incentives | 196 |
| | P10-3.4 Non-negotiability of Privacy Practices | 196 |
| | P10-3.5 Uniformity of Privacy Practices | 197 |
| | P10-3.6 Social Norms | 197 |

| | |
|--|----------------|
| P10-4 Externalities of Privacy Choices | 197 |
| P10-5 Regulatory Blind Spots and Loopholes | 198 |
| P10-5.1 Vague and Undefined Legal Terminology | 198 |
| P10-5.2 Scope Limitations of Privacy Law | 199 |
| P10-6 Attempts to Fix Privacy Self-management | 199 |
| P10-6.1 Proposed Solutions | 200 |
| P10-6.2 Why Privacy Self-management Cannot Be Fixed | 200 |
| P10-7 Ways Forward | 201 |
| P10-8 Discussion and Conclusion | 202 |
| P10-9 References | 203 |
| DISCUSSION AND CONCLUSION | 209 |
| Chapter 10: Discussion | 209 |
| 10.1 Thesis Summary | 209 |
| 10.2 Public and Media Response | 213 |
| 10.3 Privacy is not Dead | 214 |
| 10.4 Technology Cannot Fix It | 216 |
| 10.5 Wide-Ranging Transparency as Ultima Ratio | 220 |
| 10.6 Recognizing Information Inferred about Individuals as Personal Data | 222 |
| 10.7 “Sensitive” vs. “Non-Sensitive” Data | 226 |
| 10.8 “Personal” vs. “Non-Personal” Data | 227 |
| 10.9 Call for a Stronger Risk-Based Regulation of Data Use | 229 |
| 10.10 A Note on the Independence of Internet Research | 234 |
| Chapter 11: Conclusion | 237 |
| 11.1 Directions for Future Development and Research | 237 |
| References | 241 |

List of Figures

| | | |
|-----------------|--|-----|
| Paper 1 | – Fig. 1: Overview of some sensitive attributes discernable from speech data . . . | 27 |
| Paper 2 | – Fig. 1: Classification of driving patterns based on streams of accelerometer data | 44 |
| | – Fig. 2: Map matching algorithm | 44 |
| | – Fig. 3: Overview of sensitive inferences that can be drawn from accelerometer data | 46 |
| Paper 3 | – Fig. 1: Overview of sensitive inferences that can be drawn from eye tracking data | 53 |
| Paper 4 | – Fig. 1: Distribution of participants’ level of awareness dependent on the inferred information | 76 |
| | – Fig. 2: Distribution of participants’ level of concern dependent on the inferred information | 80 |
| Paper 5 | – Fig. 1: A schematic and simplified overview of the threat model | 98 |
| Paper 6 | – Fig. 1: A classification of data types commonly collected by video games . . . | 118 |
| | – Fig. 2: Overview of sensitive inferences that can be drawn from eye-tracking data, voice recordings, and accelerometer data | 125 |
| Paper 7 | – Fig. 1: Popular apps dataset overview | 144 |
| | – Fig. 2: Frequency distribution of responses over time | 146 |
| | – Fig. 3: Evaluation of subject access request responses of all apps | 146 |
| Paper 8 | – Fig. 1: The user interface of the browser extension including a graph, a search feature, and further information on the third parties | 164 |
| | – Fig. 2: Records visualized on a timeline | 165 |
| | – Fig. 3: The results of the evaluation | 169 |
| Paper 9 | – Fig. 1: System overview | 176 |
| | – Fig. 2: Emotional content of the sound variations | 177 |
| Paper 10 | – Fig. 1: Overview of obstacles to the meaningful exercise of privacy self-management | 192 |
| Ch. 10 | – Fig. 10.1: Overview of the research papers included in this dissertation | 210 |

List of Tables

| | | |
|--------------------------|---|-----|
| Paper 4 – Tab. 1: | Participant demographics | 75 |
| – Tab. 2: | Regression results for AW_DEM | 77 |
| – Tab. 3: | Regression results for AW_STATE | 77 |
| – Tab. 4: | Regression results for AW_TRAIT | 77 |
| – Tab. 5: | Mean values for Grp-A and Grp-A.1 | 81 |
| – Tab. 6: | Spearman’s rank correlations between participant demographics and awareness for audio-based inferences | 90 |
| Paper 7 – Tab. 1: | Summary of results | 150 |
| Paper 9 – Tab. 1: | Sound variations: Means and standard deviations of emotional content scores | 178 |

List of Included Papers

Accepted and Published Papers

Kröger, J.L., Lutz, O.HM. and Raschke, P. (2020). Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference. In: Friedewald M., Önen M., Lievens E., Krenn S., Fricker S. (eds) *Privacy and Identity Management. Data for Better Living: AI and Privacy*. Springer, Cham, 242–258. Publisher’s version. https://doi.org/10.1007/978-3-030-42504-3_16

→ **Included in Chapter 2 of this thesis (p. 25)**

Kröger, J.L., Raschke, P. and Bhuiyan, T.R. (2019). Privacy Implications of Accelerometer Data: A Review of Possible Inferences. *Proceedings of the International Conference on Cryptography, Security and Privacy (ICCSP)*. ACM, New York, 81–87. Publisher’s version. <https://doi.org/10.1145/3309074.3309076>

→ **Included in Chapter 2 of this thesis (p. 43)**

Kröger, J.L., Lutz, O.HM. and Müller, F. (2020). What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking. In: Friedewald M., Önen M., Lievens E., Krenn S., Fricker S. (eds) *Privacy and Identity Management. Data for Better Living: AI and Privacy*. Springer, Cham, 226–241. Publisher’s version. https://doi.org/10.1007/978-3-030-42504-3_15

→ **Included in Chapter 2 of this thesis (p. 51)**

Kröger, J.L., Gellrich, L.K., Pape, S., Brause, S.R. and Ullrich, S. (2021). Personal Information Inference from Voice Recordings: User Awareness and Privacy Concerns. *Proceedings on Privacy Enhancing Technologies (PoPETs)*, 2022(1), Sciend, Warsaw, 6–27. Publisher’s version. <https://doi.org/10.2478/popets-2022-0002>

→ **Included in Chapter 3 of this thesis (p. 69)**

Kröger, J.L. and Raschke, P. (2019). Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping. In: Foley S. (eds) *Data and Applications Security and Privacy XXXIII (DBSec 2019)*. Springer, Cham, 102–120. Publisher’s version. https://doi.org/10.1007/978-3-030-22479-0_6

→ **Included in Chapter 4 of this thesis (p. 95)**

Kröger, J.L., Lindemann, J. and Herrmann, D. (2020). How do app vendors respond to subject access requests? A longitudinal privacy study on iOS and Android Apps. *ARES ’20: Proceedings of the 15th International Conference on Availability, Reliability and Security*. ACM, New York, 1–10. Publisher’s version. <https://doi.org/10.1145/3407023.3407057>

→ **Included in Chapter 5 of this thesis (p. 143)**

Raschke, P., Zickau, S., Kröger, J.L. and Küpper, A. (2019). Towards Real-Time Web Tracking Detection with T.EX – The Transparency EXtension. In: Naldi M., Italiano G., Rannenber K., Medina M., Bourka A. (eds) *Privacy Technologies and Policy. Annual Privacy Forum 2019*. Springer, Cham, 3–17. Publisher’s version. https://doi.org/10.1007/978-3-030-21752-5_1

→ **Included in Chapter 6 of this thesis (p. 157)**

Lutz, O.HM., Kröger, J.L., Schneiderbauer, M. and Hauswirth, M. (2019). Surfing In Sound: Sonification of Hidden Web Tracking. In: *Proceedings of the 25th International Conference on Auditory Display (ICAD 2019)*. Georgia Institute of Technology, Atlanta, 306–309. Publisher’s version. <https://doi.org/10.21785/icad2019.071>

→ **Included in Chapter 7 of this thesis (p. 175)**

Working Papers

Kröger, J.L., Raschke, P., Percy Campbell, J. and Ullrich, S. (2021). Surveilling the Gamers: Privacy Impacts of the Video Game Industry. Available at *SSRN*: <https://doi.org/10.2139/ssrn.3881279>

→ **Included in Chapter 4 of this thesis (p. 115)**

Kröger, J.L., Lutz, O.HM. and Ullrich, S. (2021). The myth of individual control: Mapping the limitations of privacy self-management. Available at *SSRN*: <https://doi.org/10.2139/ssrn.3881776>

→ **Included in Chapter 9 of this thesis (p. 191)**

List of Abbreviations

General

| | |
|-------------------|---|
| API | application programming interface |
| Art. 29 WP | Article 29 Working Party |
| CCPA | California Consumer Privacy Act |
| CIA | Central Intelligence Agency |
| CPU | central processing unit |
| ECG | electrocardiogram |
| EEG | electroencephalography |
| EU | European Union |
| GCHQ | Government Communications Headquarters |
| GDPR | General Data Protection Regulation |
| GPS | Global Positioning System |
| HTTP | Hypertext Transfer Protocol |
| HTTPS | Hypertext Transfer Protocol Secure |
| IT | information technology |
| IoT | Internet of Things |
| ISO | International Organization for Standardization |
| NGO | non-governmental organization |
| NSA | National Security Agency |
| OCEAN | “Big Five” traits (openness, conscientiousness, extroversion, agreeableness, neuroticism) |
| OS | operating system |
| OSC | Open Sound Control |
| PDS | personal data service |
| PIN | personal identification number |
| QoE | quality of experience |
| SMTPS | Simple Mail Transfer Protocol |
| T.EX | Transparency EXtension |
| TLS | Transport Layer Security |
| UK | United Kingdom |
| URL | uniform resource locator |
| VCVA | voice-controlled virtual assistant |

Paper 4: Subsets and Variables

| | |
|------------------|--|
| AW_DEM | level of awareness about DEM |
| AW_STATE | level of awareness about STATE |
| AW_TRAIT | level of awareness about TRAIT |
| CON_DEM | level of concern about DEM |
| CON_STATE | level of concern about STATE |
| CON_TRAIT | level of concern about TRAIT |
| DEM | inferences about user demographics |
| DVP | disposition to value privacy |
| EXP_CS | professional experience in computer science |
| EXP_DM | professional experience in data mining |
| EXP_DP | professional experience in data protection law |
| EXP_IS | professional experience in IT security |
| INNO | innovativeness |
| PA | privacy awareness |
| PPE | previous privacy experience |
| STATE | inferences about short- and medium-term states |
| TRAIT | inferences about physical & psychological traits |
| VCVA_UI | VCVA usage intention |

Paper 7: Subsets

| | |
|-------------------------|------------------------------------|
| DATA | responses containing personal data |
| DATA_X | cut set of DATA and subset X |
| EMS | vendors based in EU |
| GER | app vendors based in Germany |
| GONE | apps that not longer exist |
| NODATA | apps not collecting user data |
| NOTGONE | apps that could still be installed |
| OK | satisfactory request responses |
| R1 | data request round 1 (11/2015) |
| R2 | data request round 2 (03/2018) |
| R3 | data request round 3 (08/2019) |
| RESP | app vendors responding to request |
| RESP_X | cut set of RESP and subset X |
| WLD | vendors based outside the EU |

1

Introduction

1.1 Motivation: Privacy in Crisis

We live in a time where people have largely lost control over their personal data [1]. In the face of an increasingly diverse and complex universe of data collection and processing technologies, it has become virtually impossible for individuals to keep track of, let alone manage, the information that public and private sector organizations hold about them and what exactly the information is used for [2, 3]. Not only the amount of collected data, but also its accuracy and degree of intimate detail is steadily increasing, with modern methods of tracking and sensing reaching deep into all areas of private life [4, 5, 6].

In this world of ubiquitous surveillance, many people have adopted a fatalistic or even completely resigned attitude over the last decades [7, 8], regarding informational privacy as a “lost cause” [9, 10] or “completely and utterly dead” [11]. Commentators have long proclaimed the “end of privacy” [10, 12, 13], seeing us as already living in a “post-privacy” era [14, 15]. On the other hand, many researchers, activists, and politicians are still convinced of the importance of privacy protection and continue to lobby for more control over personal data [16, 17, 18, 19].

It is indisputable that there are major problems with the ways personal data is handled today, which include not only cybercrime issues [20] and excessive government surveillance [21] but also companies’ questionable data practices [19], a profound lack of transparency in data processing [22], legal loopholes [23], and the ineffective enforcement of data protection laws [24]. Admittedly, a sudden improvement of the situation does not seem likely. While facing this reality can make a pessimistic outlook seem justified, an all-out privacy-is-dead attitude can be dangerous and misleading in its generality. As will be argued in this thesis, there are indeed many big issues that cannot be solved or reversed, but also various possible options of tackling the privacy crisis and improving our situation. Clearly, democratic societies should use their regulatory power to bring technology and business practices in line with their core values, and not the other way around.

Helping to provide a scientific basis for shaping the digital transformation in a responsible manner, including identifying necessary framework conditions to uphold social self-determination, is the declared mission of the *Weizenbaum Institute for the Networked Society*,¹ where the research for this thesis was conducted.

Much of what is currently happening in the data economy is not widely desired among the general population. Behavioral targeting, for example, wherein search engines and websites serve ads based on people’s online behavior (e.g., time spent on web pages, links clicked on, searches made) is commonplace in the modern internet albeit widely opposed among its users. Already in 2012, when behavioral targeting capabilities were less developed than today, over two-thirds of the participants of a Pew Research Center survey expressed their disapproval of behavioral targeting [27]. In a more recent survey focusing on Facebook’s data practices, about half of the participants said that they are not comfortable when shown how the platform categorizes them [28]. When it comes to personal data, another questionable field of application is credit scoring – a system that, in its current way of functioning, “unfairly penalizes the poor and dramatically limits fair access to financial products at equitable prices” [29]. A survey from 2019 revealed that the vast majority of Americans (81%) feel “that the potential risks they face because of data collection by companies outweigh the benefits” [1].

Despite their justified concerns, people regularly consent to such data practices, e.g., by accepting companies’ privacy policies. Under current circumstances, individual declarations of consent to data processing should not be confused with genuine approval. As we will explore in-depth in this thesis, people’s privacy choices are typically irrational and/or involuntary due to corporate tricks, human limitations, and the complexities of modern data processing. One of the more obvious problems is: Given their time constraints and lack of expertise, it is unrealistic to expect ordinary consumers to understand, or even read, the privacy policies of all the companies they interact with in everyday life [2, 30]. And, as we will see, there are many more problems.

In the light of these problems, the question arises whether the legal principle of *notice and choice* (also referred to as “privacy self-management” [31]), which has dominated privacy law in Western countries for decades [32], really is a suitable approach to regulate personal data processing in the best interest of society. At the end of this thesis, we will arrive at a clear and well-founded response to this question.

Along the way, we will explore various technical and social aspects of modern data collection and processing. As detailed in the below outline (Chapter 1.3), this thesis focuses on privacy threats posed by mobile apps, web tracking, and sensors embedded into consumer devices. For obvious reasons given the expanse of these subjects, not all aspects of relevance can be covered in this thesis. However, the selected studies presented in this cumulative dissertation, which are all results of research collaborations, will exemplify and provide a deeper understanding of the actors, methods, and staggering complexities involved in privacy threats of the 21st century.

¹The *Weizenbaum Institute for the Networked Society*, also referred to as *The German Internet Institute*, conducts interdisciplinary and basic research into the transformation of society through digitalization [25, 26]. It is a joint project of six universities and research institutions in and around Berlin, funded by the German Federal Ministry of Education and Research (BMBF). The project is named after the late computer scientist and social critic Joseph Weizenbaum.

Based on these insights, I will highlight the futility of some proposed “solutions” to the privacy crisis while pointing out aspects that will be crucial for bringing about actual change and lifting us out of this crisis. In addition, as this thesis comprises a total of 10 research papers, numerous domain-specific research questions – all related to privacy issues with modern consumer electronics – will be addressed in the individual papers, as outlined in Chapter 1.3 and further explained in the corresponding thesis chapters.

While multiple empirical contributions are also made, some chapters of this thesis are of consolidating nature, answering research questions and providing classifications and analyses based on published patents and previous literature. These chapters, in my view, are at least of equal societal relevance and scientific value as the empirical findings presented. They, too, contribute to filling significant gaps identified in the literature – bringing much-needed structure, synthesis, and overview to areas where primary research is abundant but secondary research is still lacking. References to all of the papers contained in this thesis can be found in the above *List of Included Papers* (p. xxi).

As a particularly interesting example of the complex privacy threats that have been emerging and intensifying in recent years, this thesis places a special focus on the subject of so-called *sensor-based inference attacks*, which will be introduced in the following subchapter.

1.1.1 Privacy Threats Posed by Ubiquitous Sensors

Ever since the advent of smartphones, a variety of embedded sensors is constantly surrounding us – whether we are at work, in transit, or even within our own four walls. With visions of “smart homes”, “smart cities”, “connected healthcare”, and all sorts of new consumer electronics, the emerging Internet of Things (IoT) is predicted to further increase the number of sensors in our everyday environment by several orders of magnitude [33, 34].

New services and business models are enabled through the increasing pervasiveness and interconnection of sensing devices which promise to bring transformational improvements in many areas including health, safety, security, convenience, productivity, and sustainability. At the same time, with smart technologies reaching ever deeper into people’s lives, there is growing concern about potential privacy violations [35, 36, 37]. Through sensors in mobile and IoT devices, the extensive spying and internet surveillance practiced by companies, criminals, and governmental agencies could be further expanded into the physical world.

Despite the general awareness of electronic surveillance, there seems to be limited understanding of the privacy implications of specific devices and sensor types. While microphones, cameras, and GPS, for instance, are commonly perceived as privacy-sensitive [38, 39] and require the user’s explicit permission to be activated in current mobile operating systems [40], many inconspicuous sensors such as accelerometers, barometers, and gyroscopes are less well-understood in terms of their potential sensitivity, and also often less protected [41].

Various seemingly harmless sensors embedded in smartphones, tablets, smartwatches, smart home appliances, and other consumer electronics can be accessed by a wide range of possibly untrusted parties, including device manufacturers [42, 43] and service providers, such as [44] mobile app vendors [40] and even operators of visited websites [45], often without any clear notification to the user. With the help of advanced analysis methods, data from seemingly innocuous sensors can be exploited to infer highly sensitive information about the people

in their vicinity [4, 46, 47]. The analysis of data to illegitimately extract knowledge about a person is referred to as an “inference attack” [48]. For example, accelerometer data from smartphones and wrist-worn wearables alone may be sufficient to obtain information about a device holder’s location [46], physical activities [49], body features [50], gender [51], age [52], and emotional state [53].

With respect to the informational richness of sensor data and the increasing capabilities of inferential analytics, scholars and privacy professionals have argued that all IoT sensor data should be treated as personal data [54] and that “all our data will be health data one day” [55]. Blanke [56, p. 81] even states that “Inferences drawn from personal data have arguably become more dangerous to individual privacy than the vast collection and storage of the data itself.” In the same vein, the Article 29 Working Party (Art. 29 WP)² [57, p. 47] wrote: “More often than not, it is not the information collected in itself that is sensitive, but rather, the inferences that are drawn from it and the way in which those inferences are drawn, that could give cause for concern.”

Even sensors that are already widely regarded as privacy-intrusive, such as microphones, can capture and reveal much more information than commonly believed. Apart from the linguistic contents of speech, a speaker’s voice and speech characteristics in a voice recording can implicitly contain information about *their* biometric identity [58], emotions [59], mental and physical health [60, 61], age [62], gender [63], geographical origin [64], and personality [65] – and may thereby potentially reveal much more information than a speaker wishes and expects to communicate.

Proper privacy impact assessments and the selection of appropriate legal and technical protection measures require a fundamental grasp of the true richness and sensitivity of collected data. The need for such transparency in the age of ubiquitous computing becomes particularly evident in view of the many ways in which sensor data from consumer electronics can be misused. Examples include the secret tracking of media consumption habits, geo-location, and daily activities for purposes such as target advertising, adjusting health insurance premiums, or identifying less profitable customers to exclude them from certain services [19, 66, 67, 68].

Therefore, it is important to investigate whether the widely assumed innocuousness of certain types of sensor data matches reality. Although privacy violations through unexpected inferences can also occur in many other areas (e.g., social media data, public health data, census data, tax records), the focus of my research in this area lies specifically on sensor data from consumer electronics.

1.2 Theoretical Background

This thesis comprises ten research papers (see *List of Included Papers*, p. xxi). Information on the background, motivation, research questions, and related work are provided within the respective papers and corresponding introductory chapters. Additionally, to place the research topics and contributions in a larger context that goes beyond the individual studies, this

²The Art. 29 WP was an independent EU advisory body made up of the European Data Protection Supervisor and representatives from the European Commission and the data protection authorities of all EU member states, established according to Art. 29 of the Data Protection Directive 95/46/EC. In 2018, with the EU’s General Data Protection Regulation (GDPR) coming into effect, the Art. 29 WP was replaced by the European Data Protection Board (EDPB).

subchapter provides an introduction to the history (Ch. 1.2.1), central theme (Ch. 1.2.2), and different strands (Ch. 1.2.3) of privacy discourse.

1.2.1 A Brief History of Privacy Discourse

Certain human habits with potential relation to privacy preferences date back to prehistoric times (e.g., covering nakedness [69]). Legal and theoretical precursors of today's privacy laws and discourse have partially existed for several centuries (e.g., medical secret, seal of confession, bank secrecy, official secrecy, postal secrecy) [70].

Intellectual discussions on the subject of data privacy emerged at the end of the 19th century against the backdrop of new technical and social developments at the time (instantaneous photography, telegraphy, tabloid press). For example, in 1890 the U.S. lawyers Samuel Warren and Louis Brandeis [71] published a highly influential article suggesting that traditional mechanisms for protecting one's privacy were no longer sufficient. Until then, most forms of data collection had required some form of active participation of the data subject, such as sitting still in front of a camera due to long exposure times [72]. In view of the changing socio-technical environment, Warren and Brandeis derived the necessity of a "right to privacy", which they understand as an "instance of the enforcement of the more general right of the individual to be let alone" [71, p. 205]. In their view, the law should afford people the right to determine "to what extent [their] thoughts, sentiments, and emotions shall be communicated to others" [71, p. 198], unless this individual choice is overridden by societal interests (e.g., law enforcement, public administration, determining political candidates' fitness for office). While receiving broad attention, their assessment and demands did not directly lead to a change in U.S. law.

The beginning of modern privacy debate and research is dated to the mid-1960s to early 1970s [70, 73], shortly after the onset of the Computer Age. Early discussions were fuelled by a 1965 proposal of the U.S. government to set up a National Data Center to collect and store data on all citizens for administrative purposes [74], which was widely perceived as a threat to individual privacy and raised concerns about potential data misuses [75]. In the U.S. Senate and House of Representatives, several hearings on the topic of privacy were held in the 1960s in response to the development and commercialization of computers and automated data processing, stimulating both public and scientific discourse [70, 76].

One of the most influential theorists of this time was Alan F. Westin, whose understanding of privacy resembles the aforementioned view of Warren and Brandeis. In his 1967 book "Privacy and Freedom", Westin writes: "Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" [77, p. 7]. As will become apparent in the following, this conception is still influential today but can – especially in the context of today's technologized societies – be criticized as unrealizable, at least not in the form of truly informed privacy decisions (cf. Chapter 9).

Among the first popular books on how technology is used to undermine individual privacy is Vance Packard's 1964 "The Naked Society" [78]. As main causes for new privacy threats, Packard not only points to technical progress in the fields of data collection and processing, but also to existing efforts of forcing order upon an increasingly complex society, to the erosion

of civil liberties due to surveillance and persecution during the Cold War, and to companies' rising tendency to collect personal data for marketing purposes. In the same year, Myron Brenton published "The Privacy Invaders" [79], in which he warns against excessive forms of data collection and projects that privacy interests of individuals will increasingly be overridden by the interests of powerful data-hungry organizations.

In Europe, too, debates about privacy and the legal regulation of technical data collection systems began around this time and led, among other outcomes, to the adoption of the world's first data protection law in Hesse in 1970 [80] and to the adoption of the *German Federal Data Protection Act* (BDSG 1977) [81]. The BDSG 1977 afforded data subjects the right to request information on data stored about them, to have incorrect data corrected, and to block the use of or even erase personal data under certain conditions [81].

In the U.S. in 1973, guidelines for privacy protection in automated data systems were proposed by the government advisors Ware et al. [82]. These guidelines, also known as the *Fair Information Practices* (FIPs), include (1) that no personal data record-keeping system should be secret, (2) that data subjects should be able to find out what data is stored about them and how it is used, (3) that data subjects must be able to prevent their data from being used for undesired secondary purposes, (4) that there must be a way to correct faulty data records about themselves, and (5) that data controllers must assure that collected data is reliable for their intended purposes and take precautions to prevent data misuse [82].

The FIPs formed the basis for the *U.S. Privacy Act* [83], which was passed in 1974, imposing new rules on the collection, storage, use, and dissemination of personally identifiable information [84]. However, the law only applies to federal agencies – not to state governments or the private sector [85]. Initiatives to extend the law's scope to include the private sector did not succeed, as a landmark study conducted by the U.S. Privacy Protection Study Commission (PPSC) concluded that competition and self-regulation would be more effective means in the private sector than a law and a regulatory agency on the federal level [86].

While the purpose of German data protection law was initially to protect personal data from data processing not legitimized by law, the Federal Constitutional Court determined in its 1983 Census Act ("*Volkszählungsurteil*") [87, 88] that if data processing goes against the will of data subjects, it can violate their fundamental rights even if technically legal. In its ruling, the court derived a "right to informational self-determination" from the general right of personality, which is regarded as a groundbreaking decision in the evolution of privacy law [89].

Privacy concerns intensified over time with further development of new technologies for data collection and processing and new data-based applications and business models. In his 1987 article "Reviewing Privacy in an Information Society", Spiros Simitis noted that surveillance and excessive data collection about private individuals had become pervasive in everyday life [90]. In 1988, Priscilla M. Regan [91] argued that new technical developments – namely, the increasing quantity of stored data, qualitative changes in information processing due to computerization, the widespread introduction of personal computers, new methods of searching and matching data about individuals, and the increasing networking of systems – meant that existing approaches for protecting personal data had become insufficient.

With regard to the advancing possibilities of automated data processing, Oscar H. Gandy in the early 1990s introduced the term "panoptic sort", which refers to complex technologies

that process data from people’s “daily lives as citizens, employees, and consumers (...) to coordinate and control their access to the goods and services that define life in the modern capitalist economy” [92, p. 29] and to sort “individuals and groups (...) according to their presumed economic or political value” [92, p. 15f.].

To this day, despite various new privacy threats posed by economic and technological change, there is no singular law in the U.S. that covers all aspects of personal data processing across the public and private sector on the federal level [93], whereas European countries have pursued national and supranational forms of data protection regulation. The EU’s Data Protection Directive (*Directive 95/46/EC*), which was enacted in 1995 and has since been superseded (as will be described below), prescribed minimum standards for data protection that were implemented through national laws in all EU member states. The directive generally prohibited the processing of personal data, unless one of several conditions was met, such as the data subject giving consent or the processing being necessary for the performance of a contract. The interested reader will find in the literature detailed accounts of the development history of privacy law in the U.S. (e.g., Solove 2006 [94]) and in Europe (e.g., Fuster [95], Wuermeling 1995 [96]) – as well as comparisons and analyses of their complex interrelation (e.g., Rustad and Koenig 2019 [97]). For information on the evolution of privacy law in other parts of the world, see, for example, Gellman 2021 [32], Dowling 2009 [98], and Botha et al. 2017 [99].

The terrorist attacks at the break of the new millennium (2001 in New York, 2004 in Madrid, and 2005 in London) evoked a public debate about security versus freedom, and led to the introduction of new surveillance tools and laws (e.g., data retention, new forms of computer and network surveillance, state-operated spyware) [70]. Another event that significantly impacted the international privacy discourse occurred in 2013, when Edward Snowden, an employee and subcontractor of the *National Security Agency* (NSA), disclosed highly classified information about numerous global surveillance programs to the public, including cooperations between the NSA, the *Five Eyes Intelligence Alliance*, various telecommunication companies, and European governments [21]. Among other effects, these revelations reignited public debate on the topic of individual privacy [100]. However, on the political level the debate did not lead to a notable restriction of mass surveillance practices, but rather to their justification and legalization [70, 101].

In the private sector, too, data collection further expanded in the 21st century, driven by various business and technology trends, including social media, behavioral ad targeting, “big data”, the adoption of smartphones and mobile apps, wearables, smart home applications, and other connected devices [6, 102, 103, 104].

The EU’s *General Data Protection Regulation* (GDPR), with the aim of harmonizing and modernizing European data protection law [105], was adopted in April 2016 and came into effect on May 25, 2018. In contrast to the preceding Directive 95/46/EC, the GDPR applies directly EU-wide without the need for member states to transpose it into national law first. While the GDPR includes various new provisions and new responsibilities for data controllers (e.g., data portability, prompt notification of data breaches, appointment of data protection officers, increased fines) and is widely regarded as an important “legislative milestone” [106], “game changer” [107, 108, 109] or even “Copernican revolution” [110] in privacy law, on closer

examination it does not offer groundbreaking new approaches to privacy preservation [70, p. 206]. “Every day, people are confronted with misleading consent requests, uncontrolled tracking and surveillance in online advertising, and large tech firms’ uncanny knowledge of their intimate lives. The GDPR has had little impact“, observes Johnny Ryan, the Chief Policy & Industry Relations Officer of Brave Software, in a 2020 report [111]. In a similar vein, Jockum Hilden writes in his doctoral dissertation on “The Politics of Datafication” [112, p. 190]:

Scholars (...) have shown how the media advertising system has evolved from a fairly simple two-sided market comprising audiences and advertisers into a complex web of middlemen. Nothing in the GDPR inherently challenges this model of online surveillance, but it does create an added strain for the consumer-facing entry to obtain verifiable consent to online targeting. For this reason, the industry’s largest players with the most targeting power in the form of consumer-facing platforms are at an advantage because they can target their users directly through multiple avenues of communication.

Some of the key reasons why data protection laws around the world, including the GDPR, fail to meaningfully address and combat the problems modern societies have with personal data processing will be examined in Chapter 9. In the face of increasingly invasive forms of data collection and surveillance in everyday life and the absence of adequate countermeasures, many people have developed a resigned attitude towards privacy on the level of their individual decisions [7, 8, 113]. However, as Solove [114, p. 23f.] aptly writes:

[T]he fact that people share personal data doesn’t mean that they don’t care about privacy. In today’s Information Age, if people really wanted to keep all their information concealed, they’d have to live in a shack in the woods. The fact that people share data in an age where it is nearly impossible not to do so has little bearing on the value of privacy.

When asked about their preferences and concerns, most people *do care* about privacy protection and would prefer to be free from excessive surveillance [1, 115]. It remains an open question how this can be achieved in a meaningful manner. In politics and research, the quest for sensible answers to the on-going privacy crisis continues.

1.2.2 The Subject of Debate

An issue that privacy research has dealt with extensively is the fundamental question: *What is (and what is not) “privacy”?* As this dissertation focuses on data collection and processing technologies, it primarily addresses topics related to *information* privacy or *data* privacy, as opposed to other privacy concepts, such as *spatial* privacy or *decisional* privacy.³ But even

³With reference to Anita Allen’s work [116], Koops et al. [117] summarize: “[S]patial privacy refers to the privacy expectations in and around one’s home (...). A privacy intrusion here is, for example, the peeping tom invading the privacy of two people’s intimate life by looking through the bedroom window and taking photographs”, whereas “[d]ecisional privacy (...) is largely a protection against state intrusions against citizens’ right to make certain intimate choices regarding their lives and the way they choose to live, including choices about same-sex marriage or assisted suicide.”

for information privacy, there is currently no academic consensus on what that term exactly means, and there probably never will be.

Pohle [70] observes that, while at the beginning of the new millenium there was already an almost unmanageable number of definitions and theories about information privacy, what followed was not a consolidation of the discourse, but rather an excessive expansion. Based on a broad search of information privacy literature, Smith [73, p. 1002] concludes: “Many legal and social scholars (...) believe that general privacy – its conceptual understanding, rigorous definition, and the intensity of the individual and cultural beliefs it informs – is so dependent on the specific context that it is impossible to develop a one-size-fits-all conceptualization of general privacy.” Accordingly, Mulligan et al. [118] describe information privacy as an “essentially contested concept”.

As the views of Warren and Brandeis [71] and Westin [77] cited in Chapter 1.2.1 exemplify, traditional understandings of information privacy focus not only on the personal identifiability of data, but also on the personal preferences of the respective data subject. It is commonly considered a violation of information privacy if personal data is collected and/or used against the will of the data subject [31]. Correspondingly, in a review of contemporary privacy literature, Bélanger and Crossler [119] write: “Information privacy refers to the desire of individuals to control or have some influence over data about themselves.” This view is reflected in privacy laws around the world [31, 32, 99], which typically limit their scope of application to personally identifiable data (e.g., Art. 1 and 4(1) GDPR) and accept the data subject’s consent as a legal basis for data collection and processing (e.g., Art. 6(1)(a) and 9(2)(a) GDPR).

In fact, however, many of the problems that modern societies have with the collection and processing of information about people are not limited to the realm of personally identifiable data [31, 120, 121]. In Chapter 10.8, it will be exemplified how even the processing of de-identified data can have significant ramifications for individuals, groups, and society at large. Also, people’s ability to make informed decisions about the data collected on them is seriously impaired, if not completely obstructed, by a combination of human limitations and the obscurities and complexities of modern data practices, which raises doubts about data protection measures and privacy laws that heavily rely on individual control, i.e., the principle of notice-and-choice (cf. Chapter 9).

Hence, the scope of today’s academic privacy discourse is being extended beyond these conceptual boundaries. This includes the research conducted for this thesis, which is recognized as information privacy research but no longer limited to the narrow traditional understanding of the problem. Without confining the focus to personally identifiable data or non-consensual data use, this dissertation deals with questions related to what Westin [122, p. 1004] would call “data surveillance”, meaning “the collection, storage, exchange, and integration of comprehensive documentary information about individuals and groups through computers and other data-processing systems.”

Rather than focusing on interpersonal aspects of privacy (“social privacy” [123]), such as how information is shared among users of a social networking site, this thesis mainly focuses on “institutional privacy” [123], i.e., how information on people is collected and used by companies and other organizations. This is because the types of data sources and technologies under investigation (e.g., IoT devices, web trackers, inferential algorithms, mobile apps) are mostly

developed and provided by organizations, not by private individuals. Organizations are more likely to have the financial and technical resources required to operate these technologies, especially on a large scale. With references to patents, industry presentations, and commercial products, the focus of the studies contained in this thesis is also rather on companies than, for example, on organized crime groups or governmental agencies, although many of the findings could certainly be applied to organizations outside the private sector as well. Due to their broad applicability, the findings are not classified into a strict relational framework, such as “citizen-government”, “employee-employer”, or “consumer-business” [124].

While modern forms of data collection and processing have countless beneficial applications and hold great potential for society, it is also important to explore potential dangers and negative effects, and how these can be mitigated. Resting in the knowledge that possible benefits and applications are extensively being explored in business and research, this thesis focuses on threats arising from modern data practices. Here, based on the above considerations, the term “threat” not only refers to potential violations of individual privacy rights but also to the use of data against the general interest of society. The goal is to identify systemic problems and to underpin societal information privacy concern which, as Bélanger and Crossler [119, p. 1034] put it, “refers to the overall concerns citizens in societies taken as a whole have for the privacy of the information about them.”

1.2.3 Information Privacy Research Areas

The research in this dissertation spans across six areas, which were identified in literature reviews by Smith et al. [73] and Bélanger and Crossler [119] as material aspects of information privacy research, namely: (1) the types of information collected about individuals, (2) data technologies and applications, (3) privacy regulation, (4) people’s information privacy concerns, (5) data controllers’ information privacy practices, and (6) information privacy tools and technologies. These research areas may serve as a common conceptual framework for the diverse types of privacy studies included in this thesis. Descriptions of the individual areas and of how they relate to the research in this thesis will be provided in the following (Chapters 1.2.3.1 to 1.2.3.6).

1.2.3.1 The Types of Information Collected About Individuals

People’s privacy perceptions [47, 125] as well as legal and technical data protection measures [40, 126] often vary based on the types of data that are being collected and processed. For instance, basic demographic information like gender, marital status, occupation, and place of birth are widely perceived as less sensitive than location data, financial details or health-related information [127]. This thesis, in accordance with other scholars, will challenge the categorical distinction often drawn between “sensitive” and “non-sensitive” data, i.e., the idea that sensitivity is an inherent property of certain types of data (cf. Chapter 10.7). Data *uses* can be meaningfully recognized as harmless or dangerous based on their actual or expected consequences, but not the data itself. Nevertheless, to improve transparency and enable an informed public debate on modern data practices, it is important to investigate the breath of data that can be collected and inferred about individuals in modern technological settings. Here, Part I of this thesis makes an important contribution by exploring the rich spectrum

of personal information that can be inferred from recordings of sensors that are increasingly ubiquitous in everyday life, namely microphones, accelerometers, and eye-tracking cameras. The findings reported in Part I serve as arguments for the above-mentioned discussion: They highlight that algorithms can often derive intimate information from seemingly innocuous data, thus providing additional evidence that the “sensitivity” of data is highly contextual. For example, certain patterns in motion sensor data from mobile and wearable devices may provide cues to biometrically identify the user [128, 129], assess the user’s health [130, 131] or reveal the user’s location, even if GPS is disabled [68, 46]. While there has been extensive research on sensor-based inferences, the resulting privacy threats have rarely been structured and summarized in the literature and are often overlooked, even in privacy research. This is illustrated, for example, by the many privacy studies that focus on smart speakers but do not consider the aspect of algorithmic inferences from audio data (e.g., [132, 133, 134, 135]).

1.2.3.2 Data Technologies and Applications

Over time, the privacy discourse continuously adapts to new technological developments in the areas of data collection, storage, and processing. Faster computers, increasingly ubiquitous sensors, and advances in algorithmic data analysis, for example, enable new and more efficient ways of intruding people’s privacy. In a well-known “framework for analyzing privacy in modern societies”, Westin [124] distinguishes four contemporary stages of privacy development, based among other things⁴ on technological advance: *The Privacy Baseline* (1945-1960), characterized by “limited information-technology developments”; *The First Era* (1961-1979), characterized by “[a]dvances in physical, psychological, and data surveillance technologies” and large-scale, increasingly automated record systems, but still “high data processing and storage costs, and heavy software limitations”; *The Second Era* (1980-1989), characterized by “enhanced computer and telecommunications performance” as well as the adoption of workplace video display terminals (VDTs), the personal computer (PC) and distributed computing in data banks; *The Third Era* (1990-2002), characterized by the rise of the internet (e.g., emails, search engines, online shopping, online forums), the emerging ubiquity of wireless communication devices (e.g., cellphone), and the “development of data-mining software based on large data warehousing applications”. Westin published this framework in 2003 [124]. Other scholars have since updated the framework by conceptualizing a *Fourth Era* (2003-present), which involves a large range of new technology trends, including blockchain, autonomous AI, mobile and cloud computing, mobile apps, further internet fragmentation, connected sensors and the “Internet of Things”, “big data”, e-government, and social media [136, 137, 138]. As the title of this dissertation already suggests, the foci of the included studies are on data from increasingly ubiquitous sensors (Chapters 2, 3 and 4), mobile apps (Chapters 4 and 5), and web tracking (Chapters 6 and 7). Furthermore, as motivated in Chapter 1.1.1, the studies focusing on sensor data will put a special focus on information that can be extracted with the help of inferential analytics. In the above framework based on and extended from Westin’s [124] work, the technologies examined in this thesis can be ascribed to the third and fourth era.

⁴Besides technological advance, Westin [124] based his framework on legal and organizational developments and changing social climates. He considered factors, such as “public trust in government, business, and the non-profit sector and (...) general public comfort with the information collection and use activities of those organizations” as well as civil rights struggles, anti-war and other social-protest movements.

1.2.3.3 Privacy Regulation

While a fine-grained analysis of data protection laws is beyond the scope of this thesis, it is evaluated whether the reported findings could play a role in legal discourse (Chapter 9). Central problems that repeatedly surface in the studies conducted for this thesis are (1) the obscurity and staggering complexity of modern data processing, (2) the privacy-intrusion potential of seemingly harmless data, (3) companies' questionable data practices, (4) a widespread lack of understanding regarding these data practices and complexities and, as a result, (5) people often being tracked and profiled by companies against their will and without their awareness. These points provide additional evidence that, in today's socio-technical environment, most people are not able to manage their privacy via notice and choice in a meaningful and truly informed manner. Thus, they call into question the legal principle of privacy self-management, the criticism of which was sparked by groundbreaking work from Daniel J. Solove in 2013 [31] and remains a central issue in contemporary privacy discourse (e.g., [120, 139, 140, 141, 142]). Here, a gap in scientific literature was discovered, namely the lack of a publication providing a comprehensive summary and overview of the existing arguments against privacy self-management. The research presented in Chapter 9 is an attempt to fill this gap. From the considerations presented, it becomes apparent that a primary focus on individual consent as a legal basis for data processing is not useful. The focus should rather be placed on harms resulting from data use, and how these can be mitigated through regulation, e.g., by banning certain types of data collection and use irrespective of consent (cf. Chapter 10.9). Other points made in this thesis that fall into the realm of privacy regulation are that information inferred about individuals should be unambiguously covered and protected by privacy law (Chapter 10.6) and that strict distinctions between "sensitive" and "non-sensitive" data (Chapter 10.7) and between "personal" and "non-personal" data (Chapter 10.8) – which both often result in sparse protection afforded to the latter – should be challenged.

1.2.3.4 People's Information Privacy Concerns

Another important field of interest in information privacy research focuses on people's perceptions about the collection and use of their data and their level of concern about specific privacy threats. Some influential scholars in this field are Malhotra et al. [143], who proposed a widely used construct, scale, and causal model on internet users' information privacy concerns; Bellman et al. [144], who explored international differences in privacy perceptions; Dinev and Hart [145], who investigated the antecedents of people's internet privacy concerns; and Acquisti and Gross [146, 147], who explored people's privacy preferences and information sharing behavior on social media. While, as mentioned above, individual privacy decisions are typically neither truly free nor truly informed and should therefore be relied on less heavily in the regulation of data use (cf. Chapter 9), current legal frameworks, which largely build on the principle of privacy self-management, will not vanish overnight. Overhauling or replacing them will most likely be a lengthy endeavor. This means that individual privacy perceptions and behavior, even if uninformed, unfree, and often systematically distorted by cognitive biases, continue to have a major impact on how data is collected and used in practice. In this thesis, user perceptions will be investigated with regard to the privacy impacts of modern voice and speech analysis (Chapter 3). While previous studies have already dealt with people's

privacy perceptions regarding voice-based technology (e.g., always-listening devices [39, 148, 149], microphone-equipped devices getting hacked [135, 149, 150], reluctance to using voice assistants in public [133, 151]), the aspect of information inference from audio recordings via voice and speech analysis has received little attention in user studies so far.

1.2.3.5 Data Controllers’ Information Privacy Practices

Another sub-field of information privacy research explores “organizational actions regarding privacy protection or infringement, and various factors that affect these practices” [119, p. 1022]. While some commentators have warned against “[o]verbearing regulation [and] (...) [r]egulatory overreach” [152] and argued for more industry self-regulation in the realm of data protection [153, 154, 155], serious concerns have been raised about the trustworthiness and honesty of data controllers [6, 21, 102, 156]. These concerns have been fueled by a wide range of data misuse cases and scandals involving not only small organizations but also governmental agencies and major multinational corporations. Prime examples are the Facebook-Cambridge Analytica data scandal [157] and the global surveillance disclosures by Edward Snowden [21]. Data protection compliance and the lack thereof are of course not novel issues. Already in 1988, for example, Regan [91] observed that the rules prescribed by the U.S. Privacy Act of 1974 were not being complied with by various government entities. With regard to organizational misbehavior in the private sector, many scholars have concluded that industry self-regulation does not represent a sensible solution for regulating the collection and use of data about individuals (e.g., [156, 158, 159, 160, 161]). In this thesis, the topic of data controllers’ information privacy practices is addressed through a four-year undercover study focusing on the GDPR compliance of mobile apps (Chapter 5). Similar studies have already been conducted on various types of data controllers, including on smartphone app vendors and website owners [162], online tracking companies [163], CCTV operators [164], and various other private and public sector organizations [165, 166, 167, 168, 169, 170]. In contrast to existing work, the study presented in Chapter 5 follows a novel longitudinal approach, thus contributing a method for analyzing trends in data controller behavior over time. While it was hoped that the GDPR with its new rules and increased fines would have a significant positive impact on companies’ privacy practices [171, 172], the results of our study suggest that the level of compliance remained low even after the law came into effect. Of course, the discourse around data protection compliance should not be separated from the discourse around the meaningfulness of the underlying laws. While it is important to track the enforcement of existing laws, high levels of compliance alone do not help much if the laws that are complied with are misguided and senseless. As for some of the regulatory shortcomings and potential loopholes in existing privacy law, see the points raised above in Chapter 1.2.3.3.

1.2.3.6 Information Privacy Tools and Technologies

A wide variety of technical methods and tools have been developed to help prevent “unnecessary or unwanted processing of personal data” [173] and “to protect privacy by technically enforcing legal privacy principles” [174]. These technologies, which include, for example, tools for data minimization, anonymization, and encryption, are often subsumed under the collective term *privacy-enhancing technologies* (PETs) [175]. In the literature, there are extensive

overviews of existing PETs for specific areas of application, such as genome research [176], e-health [177], smart cities [178], and other Internet of Things applications [179] as well as generally for the field of “big data analytics” [180]. This thesis offers two contributions to the field of privacy tools and technologies, namely a novel browser extension for the collection of training data for the automated detection of web trackers (Chapter 6) and a software tool for the sonification of web tracking activities (Chapter 7). Besides the privacy design principle of “control” (i.e., the property of systems to provide choice over what information can be collected by whom), Bellotti and Abigail also identify the privacy design principle of “feedback” (i.e., the property of systems to inform users about details of data collection and use) [181]. The two aforementioned contributions in this thesis can primarily be attributed to the “feedback” category, although the collected web-tracking data could, of course, also serve as a basis for the development of control or defense mechanisms, such as anti-tracking tools. For reasons that were discussed in Chapter 1.2.3.3 and will be expanded upon in Chapter 9, the feedback provided by these tools will typically not suffice for users to make truly free and informed privacy choices. Nonetheless, making hidden forms of tracking and data collection experienceable is important to raise general awareness of the problem, increase accountability of data controllers, and inform public debate on modern types of privacy threats. In general, on the subject of PETs it should be noted: While technical means for privacy protection are clearly needed, they also have numerous limitations and should not be viewed as a standalone solution to the privacy crisis. This position will be further elaborated in Chapter 10.4.

1.3 Thesis Outline

The remainder of this dissertation is divided into four parts, which will be outlined below, followed by a general discussion. Each thesis part includes one or more of the publications introduced in the above *List of Included Papers* (p. xxi). Throughout the dissertation, when I write “we”/“us”/“our” in connection with a specific study (e.g., “we conducted a survey”, “our paper”), this will always refer to the co-authors of the respective publication and myself. The term “chapter” will refer to segments of the dissertation skeleton, whereas the term “section” will refer to segments of the included research papers. As already shown in the *Table of Contents*, cross-references to paper sections are prefixed with a matching index (e.g., “P5–1” for paper 5, section 1) to make them uniquely referable throughout the thesis.

Part I focuses on the various types of personal information that can be inferred from patterns and correlations in IoT and mobile sensor data. It contains three chapters. In Chapter 2, based on patents and experimental literature, we explore and illustrate the wealth of inferences that can be drawn from eye-tracking data, voice recordings, and accelerometer data, respectively. Leading into the chapter, I explain why these types of papers are important and why I chose to focus on these specific sensors. Building upon findings from the previous chapter, Chapter 3 presents results from a survey about users’ awareness and privacy concerns regarding personal information inference from voice recordings. Chapter 4 sheds light on two momentous privacy issues that are, in various ways, related to sensor-based inferences. Firstly, one paper investigates the privacy impacts of the video game industry, including a focus on sensor data collected through gaming equipment. Secondly, one paper addresses the

issue of “mobile eavesdropping”, i.e., the strange and frequently reported phenomenon that private conversations conducted in the presence of smartphones seemingly result in targeted online advertisements. Our paper reviews existing approaches to explain the phenomenon and examines the technical feasibility and detectability of mobile eavesdropping attacks, taking into account permission requirements, user notifications, sensor sampling frequencies, limited device resources, and existing security checks. Particular attention is paid to the role that inferential analytics could play in such attacks, with smartphone motion sensors being investigated as a possible eavesdropping channel.

Part II focuses on data practices of mobile apps. After Part I has shown that many companies with access to mobile and IoT sensor data (e.g., mobile app vendors) can analyze the data to infer sensitive information about users, this part takes a closer look at how transparently such service providers handle user data. In Chapter 5, an undercover investigation is presented that probes whether mobile app vendors comply with transparency obligations prescribed by EU’s General Data Protection Regulation. While the law grants consumers the right to access the personal data that companies hold about them, the study reveals severe obstacles to exercising this right in practice.

Part III broadens the focus beyond mobile apps and sensors by addressing another major type of modern user surveillance, namely web tracking. It deals with the detection and exposure of web-tracking activities hidden to ordinary internet users, before connecting these ideas back to the topic of sensor-based inferences covered in Part I. In Chapter 6, a novel browser extension called “T.EX – The Transparency EXtension” is proposed, which records network traffic during website visits in a privacy-preserving manner. An implementation is presented and evaluated for its performance. The real-world browsing sessions recorded by the extension can serve as a data basis for developing algorithms that automatically detect web trackers (e.g., in order to block trackers and protect users’ privacy). As yet, artificial data is often used for this purpose, which lacks in quality. Pioneering in a young and uncharted field of research, Chapter 7 explores ways in which web-tracking activity can be “sonified”, i.e., made audible through indicative melodies and sounds. When a connection to one of the leading tracking companies is established, this is indicated by a voice whispering the respective company name. Improving upon existing approaches on web tracking sonification, our proposed solution can monitor any network connection, including all applications, browsers, and devices. A small-scale experiment is included to test the effect of web-tracking sonification on users. To connect the topic of threat detection and presentation with the topic of sensor-based inference attacks covered previously in this thesis, Chapter 8 provides suggestions as to how the range of personal information that can be inferred from different types of sensor data could be visually presented.

Part IV zooms out to reflect on the the privacy threats examined in this thesis and explore regulatory implications. Drawing on findings from the previous parts and existing literature, Chapter 9 demonstrates that current legal frameworks for data protection blatantly fail to safeguard people’s privacy – and offers perspectives on what can be done about it. The focus is on privacy self-management (i.e., the legal principle that people individually manage

their privacy via consent), which is a cornerstone of data protection laws throughout the Western world. The chapter offers a summary and classification of the varied obstacles that render privacy self-management effectively useless in practice. The overview and analysis provided show that most of our privacy decisions are involuntary, irrational and/or legally circumventable by using tricks and legal loopholes. Additionally, the paper addresses the problem that privacy choices of individuals can have significant and unaccounted-for effects on other people and society at large. Based on these observations, the paper argues that other forms of government intervention are needed to meaningfully address the consequences of modern data processing.

Discussion and Conclusion. In the subsequent general discussion (Chapter 10), after a recapitulation of the main research findings (Ch. 10.1) and an overview of public and media responses (Ch. 10.2), I address several overarching themes. I start off by arguing against the widespread and misleading privacy-is-dead narrative (Ch. 10.3), before making the point that even the smartest privacy-enhancing technologies will not suffice to overcome our current privacy crisis (Ch. 10.4). I stress that, in the face of growing and fundamentally unpreventable privacy threats, the enforcement of radical transparency in data processing is often a powerful last resort, which should be used more extensively (Ch. 10.5). I advocate for unambiguously recognizing inferred data about individuals as falling within the material scope of privacy law (Ch. 10.6) and also generally question whether strict distinctions between “sensitive” and “non-sensitive” data (Ch. 10.7) and between “personal” and “non-personal” data (Ch. 10.8) really make sense. Finally, I make the case for stronger government interventions based on the harms and risks involved in data processing (Ch. 10.9), offer some remarks on academic freedom and current threats to the independence of internet research (Ch. 10.10), before ending with an outlook and conclusion, presenting various interesting avenues for future research (Ch. 11).

Please note: According to the requirements of the publishers and the Technische Universität Berlin, the papers included in this thesis are in the form and format in which they were originally published. The only change I have made is to remove the page numbers and partly also the running page headers from the papers to avoid confusion with the headers and page numbers of this thesis and thus provide more clarity for the reader. With kind permission, the two working papers included in this thesis were prepared using LaTeX preprint templates by *Elsevier* (Paper 6) and *SAGE Publishing* (Paper 10).

Part I

**PRIVACY-INVADING
POTENTIAL OF SENSOR DATA**

2

Voice Recordings, Eye Tracking, and Accelerometer Data

2.1 Background and Motivation

This chapter will present what are – to the best of my knowledge – the first⁵ research papers to broadly examine and illustrate the wealth of personal information that can be inferred from voice recordings (Paper 1), accelerometer readings (Paper 2), and eye-tracking data (Paper 3). While many other types of sensors, such as magnetometers, air quality sensors, and smart electric meters, can be exploited for inference attacks as well [4, 183], the focus on accelerometers, eye-tracking sensors, and microphones was chosen based on their expected privacy-intruding potential, their presence in experimental literature, and their role in current technology trends.

The accelerometer, for example, which measures acceleration forces, is one of the sensors most frequently accessed by mobile apps [40] and most widely built into wearable devices, such as smartphones, tablets, smartwatches, digital cameras, wearable fitness trackers, game controllers, and virtual reality headsets [184]. Microphones are similarly omnipresent in modern technology. Besides being widely integrated into mobile devices to enable voice memos, phone calls, and voice messages, they are increasingly used to communicate with voice-controlled virtual assistants [185]. The number of installed smart speakers is forecast to reach 640 million globally by 2024 [186]. Digital voice assistants are already used by two thirds of the adult population in the US [187] and one third to a half of the adult population in Germany [188, 189], approximately.

And while eye tracking can still be classified as an emerging technology, it is starting to be used in many areas, including personal computing, healthcare, automotive technology, and gaming [190], and it is generally seen as a promising and impactful technology that may

⁵For eye tracking data, there is one notable exception: Liebling and Preibusch’s “Privacy considerations for a pervasive eye tracking world” [182]. However, while highly valuable and insightful, this paper does not address some of the categories of possible inferences that are covered in our research presented in Paper 3.

considerably shape our lives in the medium- and long-term future [182] – especially through its applications in the domain of augmented and virtual reality [191, 192].

Although appearing similar at first glance, the three publications included in this chapter are very distinct from one another. In structure and content, a certain degree of overlap could not be avoided because, partially, similar information can be derived from accelerometer data, voice recordings, and eye-tracking data, and because similar implications had to be addressed in the discussion sections. Nonetheless, by addressing three fundamentally different types of sensors with different application areas and relevance to different industries and scientific communities, each of the papers provides a unique contribution. The distinctiveness of the individual papers becomes particularly evident when considering their reference lists: Among 284 sources cited in total, there are only three overlaps between Paper 1 and Paper 3, and not a single overlap between either of them and Paper 2.

Why is it important to explore and expose the wealth of inferences that can be drawn from certain types of sensor data – to fellow researchers, lawmakers, and the wider public? Here are the four main reasons that motivated our work:

- **Consumer education.** By examining the possibilities of sensor-based inference attacks, these types of research papers can underpin consumer education efforts, serving as a knowledge basis and inspiration. Given that we are constantly surrounded by sensors in everyday life, it is important for research to investigate and showcase the ways in which sensors can be (mis)used to intrude our privacy. Research – including our own (cf. Paper 4) – indicates that the general population, including even ICT professionals, is not sufficiently aware of the privacy risks posed by inferential analytics [47, 193, 194, 195]. The practical need for overview papers and consolidated information on the privacy implications of sensor data is also evidenced by existing educational campaigns in this field, such as the American Civil Liberties Union’s attempt at raising awareness about “The Privacy-Invasive Potential of Eye Tracking Technology” [196], and by the encouraging feedback and media response our papers have received (see Chapter 10.2).
- **Risk assessment and privacy safeguards.** Information about the richness and sensitivity of sensor data is crucial for holistic privacy impact assessments and the informed selection of appropriate technical and organizational privacy safeguards. For instance, it is important to avoid a false sense of security, such as when accelerometer data is wrongly perceived as “not particularly sensitive” [197] and, accordingly, not sufficiently protected from unauthorized access in mobile operating systems [40, 41]. Considering the higher level of protection afforded, for example, to cameras and GPS sensors, malicious parties could try to use less-protected sensors as substitutional data sources for gaining intimate insights about people [4]. To have collected data legally recognized as sensitive⁶ personal information (“special categories data”, Art. 9 GDPR), the claim that data contains sensitive information has to be substantiated [198]. In

⁶In Chapter 10.7, I challenge the legal distinction between “sensitive” and “non-sensitive” data. Ultimately, all sorts of data can be harmless or sensitive depending on a data controller’s means and intentions. In many harmful and discriminatory business practices, seemingly non-sensitive data can be used as a proxy for data legally recognized as sensitive [121]. However, as long as the legal distinction is being made, convincing arguments are needed to substantiate the harm-potential of seemingly harmless data types to ensure appropriate protection.

addition, as will be discussed in Chapter 10.6, with the legal status of inferred data under EU law being contested, solid arguments need to be provided in order for inferred data to receive any form of legal protection in the first place. In such argumentation processes, Papers 1, 2, and 3 can serve as supportive evidence for the examined types of sensor data. Furthermore, researchers have proposed various technical approaches to prevent the unwanted inference of personal information from sensor data (e.g., [199, 200, 201]). However, most of these approaches are still in the early stages of development and do not yet offer reliable protection [36, 202]. For direction and guidance in this emerging field, foundational research is urgently needed. Overview papers, such as the ones presented in this chapter, can help researchers in identifying sensor-related privacy threats and may thereby lower the risk of existing threats being overlooked.

- **Critical discourse on data protection law.** While technical privacy safeguards are in the making, dealing with the societal implications of inferential analytics is not only about preventing the unwanted collection and inference of personal information. By illustrating the privacy-invading potential of sensor data, our papers also underscore the importance of addressing the issues of ubiquitous surveillance and the increasingly transparent consumer on the regulatory level. They raise the question: How should we, as a society, best deal with a situation where powerful organizations can know virtually everything about us, and where this becomes increasingly difficult to prevent? Given the findings compiled in Paper 1, for example, it is not hard to imagine a future where companies may only need a few voice recordings (e.g., voice commands) to establish a fairly detailed profile about a person’s health and personality. With raw sensor data being a necessary requirement for many modern services, it will not be feasible to reliably avoid the disclosure of such data. Thus, while existing safeguards often focus on the aspect of data collection, there is a need for protective and countervailing measures further down the data lifecycle. In particular, legal analyses have shown that EU data protection law does not clearly and sufficiently address the privacy threats posed by inferential analytics [23, 56, 198, 203, 204], which shows that raising awareness for this issue remains crucially important (cf. Chapter 10.6).
- **Future research.** The information provided in our papers may also serve as a foundation for all sorts of further research on the societal impacts of sensor data, as illustrated by our own study on user awareness and privacy concerns about personal information inference from voice recordings (Paper 4), which was designed based on knowledge compiled in our literature review on the privacy impacts of voice and speech analysis (Paper 1).

2.2 Research Scope and Limitations

Papers 1, 2, and 3 present an astonishing variety of personal information that can be inferred from the respective type of sensor data. Critical readers may wonder whether the presented findings are realistic, given the state and limitations of current technology. In times where image-recognition algorithms still confuse abstract patterns with animals [205] and animals with raisin muffins and washing machines [206], one may ask: How is it possible that mere eye movements can reveal information about a person’s cultural background and personality traits?

And if it is true that mental health issues can be detected from voice recordings, why do people still need to see a psychologist to get a diagnosis? To address this quite understandable sort of skepticism, the following will briefly elaborate on the scope and limitations of our work.

First of all, to stay with the latter example, when our papers mention the sensor-based inference of mental health information, this does not necessarily mean that psychographic profiles inferred from sensor data match the comprehensiveness, validity, and reliability of judgements from human psychology experts. What is meant is rather – as will be detailed in the corresponding chapters – that some researchers have managed to detect certain mental health issues using the respective type of sensor data (e.g., schizophrenia and depression from eye-tracking data [207, 208] and voice recordings [61, 209]). This does not imply that the same can be achieved under all conditions and for every individual, but should primarily be seen as evidence that inferences in this area are possible in principle, representing a potential privacy threat that should be monitored and further investigated.

As with all types of predictive algorithms, various confounding factors (e.g., signal noise, environmental disturbances) can have an impact on the possibility and quality of drawn inferences, sometimes rendering the extraction of personal information from sensor data exceedingly difficult or even impossible. While inference methods tend to be accurate and reliable in “extreme cases”, they can show lower performance in “limit cases”, where the inferred attribute or characteristic of the data subject (e.g., intoxication) is less pronounced [210]. In addition, the cited studies are of course subject to their own limitations (e.g., small sample size, cost of the attack, error rates, controlled laboratory conditions), which will be reflected upon in our papers as well. While experimental research, patented systems, and existing commercial products prove that impressively accurate inferences are possible with all of the studied types of sensor data, the field of inferential analytics is clearly in its infancy, and it is by no means our intention to deny or trivialize the numerous remaining technical challenges and limitations.

While the approach of presenting inference attacks that are not yet reliably and universally applicable could be seen as too speculative or even alarmist, the early and thorough investigation of emerging threats associated with modern technologies can be very important for society. In defending their work on the potential future implications of brain-reading for people’s mental privacy, Mecacci and Haselager [211, p. 457] wrote:

[I]t would be unwise to wait with the assessment and discussion of potential implications (...) [until] the technology would be full-fledged. One shouldn’t delay the ethical discussion until it is too late (van de Poel and Royakkers 2011, p. 130). Societal debates take time too, and all too often technological (and economic) developments run ahead of proper societal evaluations to such an extent that it becomes extremely hard to correct them (consider e.g. the implications of internet tracking for privacy). Therefore, we suggest, one has no other option [than] to discuss the implications of technology under development[.]

These arguments equally apply to the research presented in this chapter, except that the sensor-based inferences covered are arguably more developed and more widely applied in practice than brain-reading technology – meaning that the discussion of their societal implications may be an even more pressing issue. Only a fraction of the cited inference approaches being used in

practice without appropriate safeguards would be sufficient to pose serious threats to consumer privacy. As a precaution, all threats, even potential ones, should be examined closely.

At the same time, it is important to maintain a critical perspective and question the advertised capabilities of predictive algorithms. Projects like *Calling Bullshit* [212, 213] and scholars like Kate Crawford [214], Cathy O’Neil [215], and Rashida Richardson [216] deal with the important question of how biases in data, methodological errors, deceptive claims, and overstated results can lead to inflated expectations and inappropriate applications of algorithmic systems. For example, highly controversial studies have claimed to predict criminal tendencies based on people’s facial features [217, 218]. These studies received criticism from a wide range of researchers and practitioners who view their results as dangerous and misleading [219]. One of the articles, which has since been retracted, claimed 97% accuracy in predicting criminality from portrait photographs [218]. Such systems, if applied by law enforcement in practice, could undermine the presumption of innocence and lead to serious and widespread discrimination [219]. With regard to such dangers, it needs to be stressed that the summaries of research results in Papers 1, 2, and 3 are not intended to obstruct critical scrutiny. Quite the contrary: By providing an overview of inference methods that are currently being developed in the academic and corporate world, the papers help to expose them to public scrutiny and catalyze scientific discourse. Both the actual capabilities of these methods and their mere development bear risks because any inferred information or assessment, whether accurate or not, can have real consequences for the data subject [5, 19].

After careful consideration, we decided against putting emphasis on the concrete accuracy results reported for individual inference methods cited in Papers 1, 2, and 3. This means that, apart from a few exceptions for illustrative purposes, the papers included in this chapter will not state numeric findings of cited studies. The following considerations led to this decision:

- **Algorithmic capabilities as trade secrets.** Since research conducted in the private sector is often subject to non-disclosure requirements [220, 221, 222], it is impossible to accurately and exhaustively portray the current state of the art in inferential analytics based on publicly available literature. An emphasis on numeric results could have misleadingly suggested otherwise. Data controllers (e.g., large tech corporations) can be much better equipped in terms of budget, resources, and technical expertise than the researchers whose published results we draw on in our papers, meaning that the risk of undesired inferences from sensor data is likely bigger in real life than it appears based on the specific reported findings we could have cited.
- **Pace of technological progress.** The reviewed literature shows that inference methods evolve quickly, which means that specific accuracies from current work are likely to be outdated very soon. Given new inference approaches and higher accuracy levels to be realized, reading outdated information on sensor-based inference attacks may generate a false sense of security among future readers – especially when considering the broad scope and synoptic nature of our privacy-focused overview papers.
- **Context dependence.** The inference accuracies stated in research publications typically refer to specific setups and scenarios (e.g., device models used, sensor sampling rates, experimental environments, sample demographics, instructions given to participants,

data analysis methods used). Thus, their numerical results cannot simply be generalized to all real-life situations and, therefore, have only limited significance – and may even cause confusion – in a broad overview paper.

- **Attack requirements.** Low inference accuracies do not necessarily imply low privacy risk. Inference algorithms do not need to be correct all of the time to be useful for intrusive profiling and tracking purposes. When companies use personal data for discriminatory scoring practices or to influence people’s behavior through micro-targeted communication, for example, their techniques are often based on approximate predictions and probabilities rather than factual knowledge [19, 121, 198]. The typical goal is to achieve a large enough number of correct hits to make a particular strategy successful or profitable at scale, rather than being exactly right in all cases. For hackers who are trying to infer passwords entered into smartphone touchscreens based on user hand motions, being correct in only 1% of the cases could already be lucrative [223]. Therefore, inference methods with limited accuracy deserve attention by privacy research and are a potential cause for concern, too. Citing a wide range of accuracy results in our papers, where space limitations did not allow for an in-depth discussion of this issue, could have been confusing or even misleading. Finally, it is also worth noting that even inference methods that are completely inaccurate and faulty are being applied in practice, which causes additional problems and discriminatory side effects [23, 224, 225]. Thus, for an assessment of societal impacts, it is relevant to know in which areas inferences are being used and explored, even if the achieved accuracies are still limited.

In view of the vast variety of existing inference methods and the pace of technological progress, the papers in this chapter are intended to be illustrative rather than exhaustive. Given the likelihood of existing inference methods further improving and new threats emerging, coupled with the lack of transparency resulting from companies’ non-disclosure policies, it is virtually impossible to determine what certain organizations are or will be capable of technologically, and to which extent these inference methods are really used in practice. The overviews provided in Papers 1, 2, and 3 are based on published findings from patents and peer-reviewed research but should by no means be interpreted as the upper bound of what is technically feasible.

All three projects were initiated and coordinated by me. I collaborated with Otto Hans-Martin Lutz, Florian Müller, and Towhidur Rahman Bhuiyan from the Weizenbaum Institute and Philip Raschke from Technische Universität Berlin, who all provided support in conducting the literature search, interpreting the results, and critically revising the manuscript.

Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference

Jacob Leon Kröger^{1,2(✉)}, Otto Hans-Martin Lutz^{1,2,3}, and Philip Raschke¹

¹ Technische Universität Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany
{kroeger, philip.raschke}@tu-berlin.de

² Weizenbaum Institute for the Networked Society, Berlin, Germany

³ Fraunhofer Institute for Open Communication Systems, Berlin, Germany

Abstract. Internet-connected devices, such as smartphones, smartwatches, and laptops, have become ubiquitous in modern life, reaching ever deeper into our private spheres. Among the sensors most commonly found in such devices are microphones. While various privacy concerns related to microphone-equipped devices have been raised and thoroughly discussed, the threat of unexpected inferences from audio data remains largely overlooked. Drawing from literature of diverse disciplines, this paper presents an overview of sensitive pieces of information that can, with the help of advanced data analysis methods, be derived from human speech and other acoustic elements in recorded audio. In addition to the linguistic content of speech, a speaker’s voice characteristics and manner of expression may implicitly contain a rich array of personal information, including cues to a speaker’s biometric identity, personality, physical traits, geographical origin, emotions, level of intoxication and sleepiness, age, gender, and health condition. Even a person’s socioeconomic status can be reflected in certain speech patterns. The findings compiled in this paper demonstrate that recent advances in voice and speech processing induce a new generation of privacy threats.

Keywords: Audio · Voice · Speech · Microphone · Privacy · Inference · Side channel

1 Introduction

Since the invention of the phonograph in the late 19th century, it has been technically possible to record and reproduce sounds. For a long time, this technology was exclusively used to capture pieces of audio, such as songs, audio tracks for movies, or voice memos, and for the telecommunication between humans. With recent advances in automatic speech recognition, it has also become possible and increasingly popular to interact via voice with computer systems [96].

Microphones are ubiquitous in modern life. They are present in a variety of electronic devices, including not only phones, headsets, intercoms, tablet computers, dictation machines and baby monitors, but also toys, household appliances, laptops, cameras, smartwatches, cars, remote controls, and smart speakers.

© The Author(s), Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference, published 2020 in M. Friedewald et al. (Eds.): Privacy and Identity 2019, pp. 226–241, Springer, reproduced with permission of Springer Nature.

https://doi.org/10.1007/978-3-030-42504-3_16

There is no question that microphone-equipped devices are useful and important in many areas. It is hard to imagine a future, or even a present, without them. However, as a growing proportion of audio recordings is disseminated through insecure communication networks and processed on remote servers out of the user’s control, the ubiquity of microphones may pose a serious threat to consumer privacy. Research and public debates have addressed this concern, with published reports looking into technical and legal aspects regarding data collection, processing, and storage, as well as access and deletion rights of the data subjects [18, 32, 96]. Yet, the recent privacy discourse has paid too little attention to the wealth of information that may unexpectedly be contained in audio recordings.

Certain characteristics of human speech can carry more information than the words themselves [94]. With the help of intelligent analysis methods, insights can not only be derived from a speaker’s accent, dialect, sociolect, lexical diversity, patterns of word use, speaking rate and rhythms, but also from acoustic properties of speech, such as intonation, pitch, perturbation, loudness, and formant frequencies. A range of statistics can be applied to extract hundreds or even thousands of utilizable speech parameters from just a short sequence of recorded audio [19, 80].

Based on literature of diverse scientific disciplines, including signal processing, psychology, neuroscience, affective computing, computational paralinguistics, speech communication science, phonetics, and biomedical engineering, Sect. 2 of this paper presents an overview of sensitive inferences that can be drawn from linguistic and acoustic patterns in audio data. Specifically, we cover inferences about a user’s biometric identity (Sect. 2.1), body measures (Sect. 2.2), moods and emotions (Sect. 2.3), age and gender (Sect. 2.4), personality traits (Sect. 2.5), intention to deceive (Sect. 2.6), sleepiness and intoxication (Sect. 2.7), native language (Sect. 2.8), physical health (Sect. 2.9), mental health (Sect. 2.10), impression made on other people (Sect. 2.11), and socioeconomic status (Sect. 2.12). Additionally, we examine information that can be extracted from the ambient noise and background sounds in a voice recording (Sect. 2.13). Section 3 provides a discussion of the presented findings with regard to their limitations and societal implications, followed by a conclusion in Sect. 4.

2 Inference of Personal Information from Voice Recordings

Based on experimental studies from the academic literature, this section presents existing approaches to infer information about recorded speakers and their context from speech, non-verbal human sounds, and environmental background sounds commonly found in audio recordings. Where available, published patents are also referenced to illustrate the current state of the art and point to potential real-world applications.

Figure 1 provides an introductory overview of the types of audio features and the categories of inferences discussed in this paper.

2.1 Speaker Recognition

Human voices are considered to be unique, like handwriting or fingerprints [100], allowing for the biometric identification of speakers from recorded speech [66]. This has been

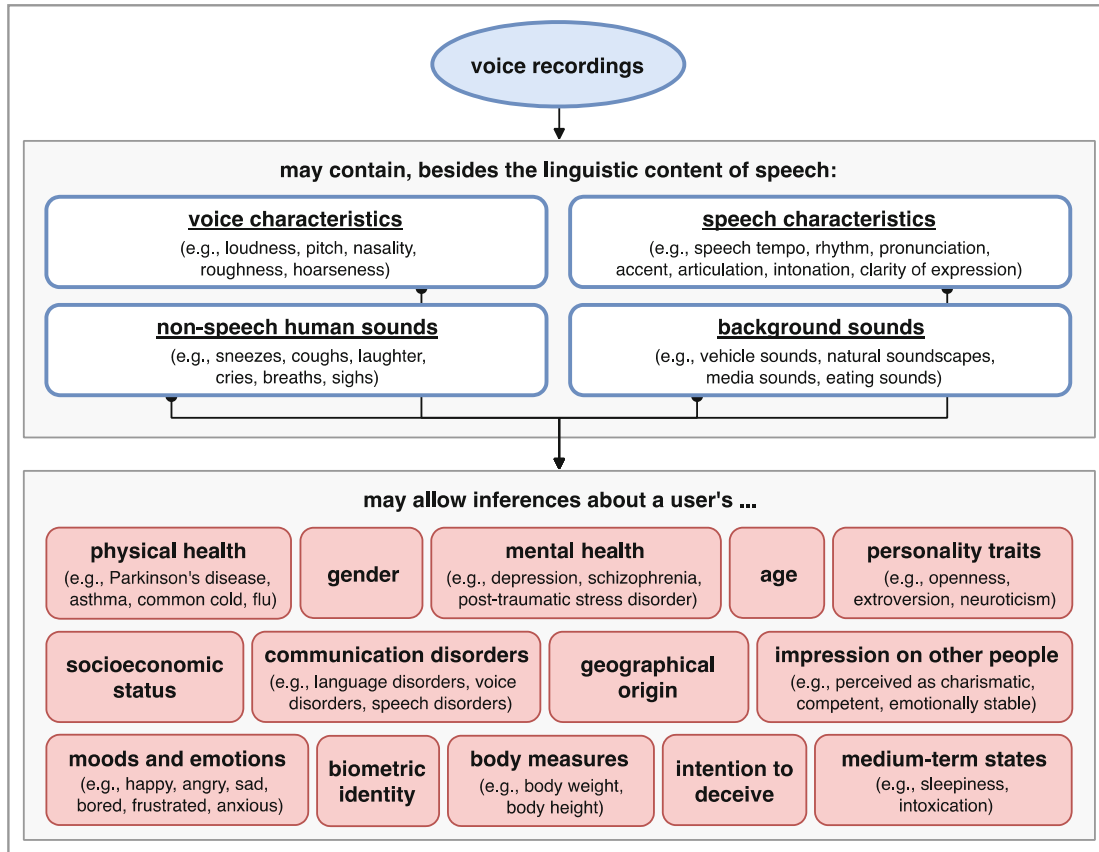


Fig. 1. Overview of some sensitive attributes discernable from speech data.

shown to be possible with speech recorded from a distance [71] and with multi-speaker recordings, even under adverse acoustic conditions (e.g., background noise, reverb) [66]. Voice recognition software has already been transferred into patents [50] and is being applied in practice, for example to verify the identity of telephone customers [40] or to recognize users of virtual assistants like Amazon Alexa [1].

Mirroring the privacy implications of facial recognition, voice fingerprinting could be used to automatically link the content and context of sound-containing media files to the identity of speakers for various tracking and profiling purposes.

2.2 Inference of Body Measures

Research has shown that human listeners can draw inferences about body characteristics of a speaker based solely on hearing the target's voice [42, 55, 69]. In [42], voice-based estimates of waist-to-hip ratio (WHR) of female speakers predicted the speaker's actual WHR, the estimated shoulder-to-hip ratio (SHR) of male speakers predicted the speaker's actual SHR measurements. In another study, human evaluators estimated the body height and weight of strangers from a voice recording almost as well as they did from a photograph [55].

Various attempts have been made to identify the acoustic voice features that enable such inferences [25, 29, 69]. In women, relationships were discovered between voice

parameters, such as subharmonics and frequency perturbation, and body features, including weight, height, body mass index, and body surface area [29]. Among men, individuals with larger body shape, particularly upper body musculature, are more likely to have low-pitched voices, and the degree of formant dispersion in male voices was found to correlate with body size (height and weight) and body shape (e.g., waist, chest, neck, and shoulder circumference) [25].

Although research on the speech-based assessment of body configuration is not as advanced as other inference methods covered in this paper, corresponding algorithms have already been developed. For instance, researchers were able to automatically estimate the body height of speakers based on voice features with an accuracy of 5.3 cm, surpassing human performance at this task [69].

Many people feel uncomfortable sharing their body measurements with strangers [12]. The researchers who developed the aforementioned approach for speech-based body height estimation suggest that their algorithm could be used for “applications related to automatic surveillance and profiling” [69], thereby highlighting just some of the privacy threats that may arise from such inference possibilities.

2.3 Mood and Emotion Recognition

There has been extensive research on the automatic identification of emotions from speech signals [21, 23, 53, 95, 99]. Even slight changes in a speaker’s mental state invoke physiological reactions, such as changes in the nervous system or changes in respiration and muscle tension, which in turn affect the voice production process [20]. Besides voice variations, it is possible to automatically detect non-speech sounds associated with certain emotional states, such as crying, laughing, and sighing [4, 23].

Some of the moods and emotions that can be recognized in voice recordings using computerized methods are happiness, anger, sadness, and neutrality [86], sincerity [37], stress [95], amusement, enthusiasm, friendliness, frustration, and impatience [35], compassion and sarcasm [53], boredom, anxiousness, serenity, and astonishment [99]. By analyzing recorded conversations, algorithms can also detect if there is an argument [23] or an awkward, assertive, friendly, or flirtatious mood [82] between speakers.

Automatic emotion recognition from speech can function under realistic noisy conditions [23, 95] as well as across different languages [21] and has long been delivering results that exceed human performance [53]. Audio-based affect sensing methods have already been patented [47, 77] and translated into commercial products, such as the voice analytics app *Moodies* [54].

Information about a person’s emotional state can be valuable and highly sensitive. For instance, Facebook’s ability to automatically track emotions was a necessary precondition for the company’s 2014 scandalous experiment in which the company observed and systematically manipulated mental states of over 600,000 users for opaque purposes [14].

2.4 Inference of Age and Gender

Numerous attempts have been made to uncover links between speech parameters and speaker demographics [26, 34, 48, 92]. A person’s gender, for instance, can be reflected

in voice onset time, articulation, and duration of vowels, which is due to various reasons, including differences in vocal fold anatomy, vocal tract dimensions, hormone levels, and sociophonetic factors [92]. It has also been shown that male and female speakers differ measurably in word use [26]. Like humans, computer algorithms can identify the sex of a speaker from a voice sample with high accuracy [48]. Precise classification results are achieved even under adverse conditions, such as loud background noise or emotional and intoxicated speech [34].

Just as the gender of humans is reflected in their anatomy, changes in the speech apparatus also occur with the aging process. During puberty, vocal cords are thickened and elongated, the larynx descends, and the vocal tract is lengthened [15]. In adults, age-related physiological changes continue to systematically transform speech parameters, such as pitch, formant frequencies, speech rate, and sound pressure [28, 84].

Automated approaches have been proposed to predict a target’s age range (e.g., child, adolescent, adult, senior) or actual year of birth based on such measures [28, 85]. In [85], researchers were able to estimate the age of male and female speakers with a mean absolute error of 4.7 years. Underlining the potential sensitivity of such inferred demographic information, unfair treatment based on age and sex are both among the most prevalent forms of discrimination [24].

2.5 Inference of Personality Traits

Abundant research has shown that it is possible to automatically assess a speaker’s character traits from recorded speech [3, 79, 80, 88]. Some of the markers commonly applied for this purpose are prosodic features, such as speaking rate, pitch, energy, and formants [68] and characteristics of linguistic expression [88].

Existing approaches mostly aim to evaluate speakers along the so-called “Big Five” personality traits (also referred to as the “OCEAN model”), comprising openness, conscientiousness, extroversion, agreeableness, and neuroticism [88]. The speech-based recognition of personality traits is possible both in binary form (high vs. low) and in the form of numerical scores [79]. High estimation accuracies have been achieved for all OCEAN traits [3, 80, 88].

Besides the Big Five, voice and word use parameters have been correlated with various other personality traits, such as gestural expressiveness, interpersonal awkwardness, fearfulness, and emotionality [26]. Even culture-specific attributes, such as the extent to which a speaker accepts authority and unequal power distribution, can be inferred from speech data [101].

It is well known that personality traits represent valuable information for customer profiling in various industries, including targeted advertising, insurance, and credit risk assessment – with potentially harmful effects for the data subjects [17, 18]. Some data analytics firms also offer tools to automatically rate job applicants and predict their likely performance based on vocal characteristics [18].

2.6 Deception Detection

Research has shown that the veracity of verbal statements can be assessed automatically [60, 107]. Among other speech cues, acoustic-prosodic features (e.g., formant frequencies, speech intensity) and lexical features (e.g., verb tense, use of negative emotion words) were found to be predictive of deceptive utterances [67]. Increased changes in speech parameters were observed when speakers are highly motivated to deceive [98].

Speech-based lie detection methods have become effective, surpassing human performance [60] and almost reaching the accuracy of methods based on brain activity monitoring [107]. There is potential to further improve the classification performance by incorporating information on the speaker's personality [2], some of which can be inferred from voice recordings as well (as we have discussed in Sect. 2.5).

The growing possibilities of deception detection may threaten a recorded speaker's ability to use lies as a means of sharing information selectively, which is considered to be a core aspect of privacy [63].

2.7 Detection of Sleepiness and Intoxication

Medium-term states that affect cognitive and physical performance, such as fatigue and intoxication, can have a measurable effect on a speaker's voice. Approaches exist to automatically detect sleepiness from speech [19, 89]. There is even evidence that certain speech cues, such as speech onset time, speaking rate, and vocal tract coordination, can be used as biomarkers for the separate assessment of cognitive fatigue [93] and physical fatigue [19].

Similar to sleepiness and fatigue, intoxication can also have various physiological effects, such as dehydration, changes in the elasticity of muscles, and reduced control over the vocal apparatus, leading to changes in speech parameters like pitch, jitter, shimmer, speech rate, speech energy, nasality, and clarity of pronunciation [5, 13]. Slurred speech is regarded as a hallmark effect of excessive alcohol consumption [19].

Based on such symptoms, intoxicated speech can be automatically detected with high accuracy [89]. For several years now, systems have been achieving results that are on par with human performance [13]. Besides alcohol, the consumption of other drugs such as \pm 3,4-methylenedioxymethamphetamine ("MDMA") can also be detected based on speech cues [7].

2.8 Accent Recognition

During childhood and adolescence, humans develop a characteristic speaking style which encompasses articulation, phoneme production, tongue movement, and other vocal tract phenomena and is mostly determined by a person's regional and social background [64]. Numerous approaches exist to automatically detect the geographical origin or first language of speakers based on their manner of pronunciation ("accent") [9, 45, 64].

Research has been done for discriminating accents within one language, such as regional Indian accents in spoken Hindi (e.g., Kashmiri, Manipuri, Bengali, neutral Hindi) [64] or accents within the English language (e.g., American, British, Australian, Scottish, Irish) [45], as well as for the recognition of foreign accents, such as Albanian,

Kurdish, Turkish, Arabic and Russian accent in Finnish [9] or Hindi, Russian, Italian, Thai, and Vietnamese accent in English [9, 39].

By means of automated speech analysis, it is not only possible to identify a person's country of origin but also to estimate his or her "degree of nativeness" on a continuous scale [33]. Non-native speakers can even be detected when they are very fluent in the spoken language and have lived in the respective host country for several years [62]. Experimental results show that existing accent recognition systems are effective and have long reached accuracies comparable to human performance [9, 39, 45, 62].

Native language and geographical origin can be sensitive pieces of personal information, which could be misused for the detection and discrimination of minorities. Unfair treatment based on national origin is a widespread form of discrimination [24].

2.9 Speaker Pathology

Through indicative sounds like coughs or sneezes and certain speech parameters, such as loudness, roughness, hoarseness, and nasality, voice recordings may contain rich information about a speaker's state of health [19, 20, 47]. Voice analysis has been described as "one of the most important research topics in biomedical electronics" [104].

Rather obviously, recorded speech may allow inferences about communication disorders, which can be divided into language disorders (e.g., dysphasia, underdevelopment of vocabulary or grammar), voice disorders (e.g., vocal fold paralysis, laryngeal cancer, tracheoesophageal substitute voice) and speech disorders (e.g., stuttering, cluttering) [19, 88].

But also conditions beyond the speech production can be detected from voice samples, including Huntington's disease [76], Parkinson's disease [19], amyotrophic lateral sclerosis [74], asthma [104], Alzheimer's disease [27], and respiratory tract infections caused by the common cold and flu [20]. The sound of a person's voice may even serve as an indicator of overall fitness and long-term health [78, 103].

Further, voice cues may reveal a speaker's smoking habit: A linear relationship has been observed between the number of cigarettes smoked per day and certain voice features, allowing for speech-based smoker detection in a relatively early stage of the habit (<10 years) [30]. Recorded human sounds can also be used for the automatic recognition of physical pain levels [61] and the detection of sleep disorders like obstructive sleep apnea [19].

Computerized methods for speech-based health assessment reach near-human performance in a variety of recognition and analysis tasks and have already been translated into patents [19, 47]. For example, Amazon has patented a system to analyze voice commands recorded by a smart speaker to assess the user's health [47].

The EU's General Data Protection Regulation classifies health-related data as a *special category of personal data* for which particular protection is warranted (Art. 9 GDPR). Among other discriminatory applications, such data may be used by insurance companies to adjust premiums of policyholders according to their state of health [18].

2.10 Mental Health Assessment

Speech abnormalities are a defining characteristic of various mental illnesses. A voice with little pitch variation, for example, is a common symptom in people suffering from schizophrenia or severe depression [36]. Other parameters that may reveal mental health issues include verbal fluency, intonation, loudness, speech tempo, semantic coherence, and speech complexity [8, 31, 36].

Depressive speech can be detected automatically with high accuracy based on voice cues, even under adverse recording conditions, such as low microphone quality, short utterances, and background environmental noise [19, 41]. Not only the detection, but also a severity assessment of depression is possible using a speech sample: In men and women, certain voice features were found to be highly predictive of their HAMD (Hamilton Depression Rating Scale) score, which is the most widely used diagnostic tool to measure a patient's degree of depression and suicide risk [36]. Researchers have even shown that it is possible to predict a future depression based on speech parameters, up to two years before the speaker meets diagnostic criteria [75].

Other mental disorders, such as schizophrenia [31], autism spectrum conditions [19], and post-traumatic stress disorder [102], can also be detected through voice and speech analysis. In some experiments, such methods have already surpassed the classification accuracy of traditional clinical interviews [8].

In common with a person's age, gender, physical health, and national origin, information about mental health problems can be very sensitive, often serving as a basis for discrimination [83].

2.11 Prediction of Interpersonal Perception

A person's voice and manner of expression have a considerable influence on how he or she is perceived by other people [44, 51, 88, 90]. In fact, a single spoken word is enough to obtain personality ratings that are highly consistent across independent listeners [10]. Research has also shown that personality assessments based solely on speech correlate strongly with whole person judgements [88]. Conversely, recorded speech may reveal how a speaker tends to be perceived by other people.

Studies have shown, for example, that fast talkers are perceived as more extroverted, dynamic, and competent [80], that individuals with higher-pitched voices are perceived as more open but less conscientious and emotionally stable [44], that specific intonation patterns increase a speaker's perceived trustworthiness and dominance [81], and that certain prosodic and lexical speech features correlate with observer ratings of charisma [88].

Researchers have also investigated the influence of speech parameters on the perception and treatment of speakers in specific contexts and areas of life. It was found, for instance, that voice cues of elementary school students significantly affect the judgements teachers make about their intelligence and character traits [90]. Similarly, certain speech characteristics of job candidates, including their use of filler words, fluency of speaking, and manner of expression, have been used to predict interviewer ratings for traits such as engagement, excitement, and friendliness [70]. Other studies show that voice plays an important role in the popularity of political candidates as it influences

their perceived competence, strength, physical prowess, and integrity [51]. According to [6], voters tend to prefer candidates with a deeper voice and greater pitch variability. The same phenomenon can be observed in the appointment of board members: CEOs with lower-pitched voices tend to manage larger companies, earn more, and enjoy longer tenures. In [65], a voice pitch decrease of 22.1 Hz was associated with \$187 thousand more in annual salary and a \$440 million increase in the size of the enterprise managed. On top of this, voice parameters also have a measurable influence on perceived attractiveness and mate choice [44].

Based on voice samples, it is possible to predict how strangers judge a speaker along certain personality traits – a technique referred to as “automatic personality perception” [88]. Considering that the impression people make on others often has a tangible impact on their possibilities and success in life [6, 51, 65, 90], it becomes clear how sensitive and revealing such information can be.

2.12 Inference of Socioeconomic Status

Certain speech characteristics may allow insights into a person’s socioeconomic status. There is ample evidence, for instance, that language abilities – including vocabulary, grammatical development, complexity of utterances, productive and receptive syntax – vary significantly between different social classes, starting in early childhood [38]. Therefore, people from distinct socioeconomic backgrounds can often be told apart based on their “entirely different modes of speech” [11]. Besides grammar and vocabulary, researchers found striking inter-class differences in the variety of perspectives utilized in communication and in the use of stylistic devices, observing that once the nature of the difference is grasped, it is “astonishing how quickly a characteristic organization of communication [can] be detected.” [87].

Not only language skills, but also the sound of a speaker’s voice may be used to draw inferences about his or her social standing. The menarcheal status of girls, for example, which can be derived from voice samples, is used by anthropologists to investigate living conditions and social inequalities in populations [15]. In certain contexts, voice cues, such as pitch and loudness, can even reveal a speaker’s hierarchical rank [52].

Based on existing research, it is difficult to say how precise speech-based methods for the assessment of socioeconomic status can become. However, differences between social classes certainly appear discriminative enough to allow for some forms of automatic classification.

2.13 Classification of Acoustic Scenes and Events

Aside from human speech, voice recordings often contain some form of ambient noise. By analyzing background sounds, it is possible to recognize the environment in which an audio sequence was recorded, including indoor environments (e.g., library, restaurant, grocery store, home, metro station, office), outdoor environments (e.g., beach, city center, forest, residential area, urban park), and transport modes (e.g., bus, car, train) [43, 97].

It is also possible to automatically detect and classify specific audio events, such as animal sounds (e.g., dog, cat, crow, crickets), natural sounds (e.g., rain, sea waves, wind, thunderstorm), urban sounds (e.g., church bells, fireworks, jackhammer), office

sounds (e.g., mouse click, keyboard typing, printer), bathroom sounds (e.g., showering, urination, defecation, brushing teeth), domestic sounds (e.g., clock tick, page turning, creaking door, keys placed on a table), and non-speech human sounds (e.g., crying, sneezing, breathing, coughing) [4, 16, 43, 97].

Algorithms can even recognize drinking and eating moments in audio recordings and the type of food a person is eating (e.g., soup, rice, apple, nectarine, banana, crisps, biscuits, gummi bears) [19, 91]. Commercial applications like *Shazam* further demonstrate that media sounds, such as songs and movie soundtracks, can be automatically identified and classified into their respective genre with high accuracy, even based on short snippets recorded in a noisy environment [49].

Through such inferences, ambient sounds in audio recordings may not only allow insights into a device holder's context and location, but also into his or her preferences and activities. Certain environments, such as places of worship or street protests, could potentially reveal a person's religious and political affiliations.

Sensitive information can even be extracted from ultrasonic audio signals inaudible to the human ear. An example that has received a lot of media attention recently is the use of so-called "ultrasonic beacons", i.e. high-pitched Morse signals which are secretly emitted by speakers installed in businesses and stores, or embedded in TV commercials and other broadcast content, allowing companies to unobtrusively track the location and media consumption habits of consumers. A growing number of mobile apps – several hundred already, some of them very popular – are using their microphone permission to scan ambient sound for such ultrasonic signals, often without properly informing the user about it [59].

3 Discussion and Implications

As illustrated in the previous section, sensitive inferences can be drawn from human speech and other sounds commonly found in recorded audio. Apart from the linguistic content of a voice recording, a speaker's patterns of word use, manner of pronunciation, and voice characteristics can implicitly contain information about his or her biometric identity, body features, gender, age, personality traits, mental and physical health condition, emotions, intention to deceive, degree of intoxication and sleepiness, geographical origin, and socioeconomic status.

While there is a rich and growing body of research to support the above statement, it has to be acknowledged that many of the studies cited in this paper achieved their classification results under ideal laboratory conditions (e.g., scripted speech, high quality microphones, close-capture recordings, no background noise) [10, 20, 30, 36, 55, 60, 70, 82, 94, 107], which may raise doubt about the generalizability of their inference methods. Also, while impressive accuracies have been reached, it should not be neglected that nearly all of the mentioned approaches still exhibit considerable error rates.

On the other hand, since methods for voice and speech analysis are often subject to non-disclosure agreements, the most advanced know-how arguably rests within the industry and is not publicly available. It can be assumed that numerous corporate and governmental actors with access to speech data from consumer devices possess much larger amounts of training data and more advanced technical capabilities than the researchers

cited in this paper. Amazon, for example, spent more than \$23 billion on research and development in 2017 alone, has sold more than 100 million Alexa-enabled devices and, according to the company’s latest annual report, “customers spoke to Alexa tens of billions more times in 2018 compared to 2017” [108]. Moreover, companies can link speech data with auxiliary datasets (e.g., social media data, browsing behavior, purchase histories) to draw other sensitive inferences [47] while the methods considered in this paper exclusively rely on human speech and other sounds commonly found in recorded audio. Looking forward, we expect the risk of unintended information disclosure from speech data to grow further with the continuing proliferation of microphone-equipped devices and the development of more efficient inference algorithms. Deep learning, for instance, still appears to offer significant improvement potential for automated voice analysis [3, 19].

While recognizing the above facts and developments as a substantial privacy threat, it is not our intention to deny the many advantages that speech applications offer in areas like public health, productivity, and convenience. Devices with voice control, for instance, improve the lives of people with physical disabilities and enhance safety in situations where touch-based user interfaces are dangerous to use, e.g., while driving a car. Similarly, the detection of health issues from voice samples (see Sect. 2.9) could help in treating illnesses more effectively and reduce healthcare costs.

But since inferred information can be misused in countless ways [17, 18], robust data protection mechanisms are needed in order to reap the benefits of voice and speech analysis in a socially acceptable manner. At the technical level, many approaches have been developed for privacy protection at different stages of the data life cycle, including operations over encrypted data, differential privacy, data anonymization, secure multi-party computation, and privacy-preserving data processing on edge devices [46, 72, 106]. Various privacy safeguards have been specifically designed or adjusted for audio mining applications. These include voice binarization, hashing techniques for speech data, fully homomorphic inference systems, differential private learning, the computation of audio data in separate entrusted units, and speaker de-identification by voice transformation [72, 73]. A comprehensive review of cryptography-based solutions for speech data is provided in [72]. Privacy risks can also be moderated by storing and processing only the audio data required for an application’s functionality. For example, where only the linguistic content is required, voice recordings can be converted to text in order to eliminate all voice-related information and thereby minimize the potential for undesired inferences.

In advocating data collection transparency and informational self-determination, the recent privacy discourse has put a focus on the recording mode of microphone-equipped devices, where a distinction can be made between “manually activated,” “speech activated,” and “always on” [32]. However, data scandals show that reporting modes cannot always be trusted [105]. And even where audio is only recorded and transmitted with a user’s explicit consent, sensitive inferences may unnoticeably be drawn from collected speech data, ultimately leaving the user without control over his or her privacy. Enabling the unrestricted screening of audio data for potentially revealing patterns and correlations, recordings are often available to providers of cloud-based services in unencrypted form – an example being voice-based virtual assistants [1, 22]. With personal data being

the foundation for highly profitable business models and strategic surveillance practices, it is certainly not unusual for speech data to be processed in an unauthorized or unexpected manner. This is well illustrated by recently exposed cases where Amazon, Google, and Apple ordered human contractors to listen to private voice recordings of their customers [22].

The findings compiled in this paper reveal a serious threat to consumer privacy and show that more research is needed into the societal implications of voice and speech processing. In addition to investigating the technical feasibility of inferences from speech data in more detail, future research should explore technical and legal countermeasures to the presented problem, including ways to enforce existing data protection laws more effectively. Of course, the problem of undesired inferences goes far beyond microphones and needs to be addressed for other data sources as well. For example, in recent work, we have also investigated the wealth of sensitive information that can be implicitly contained in data from air quality sensors, infrared motion detectors, smart meters [56], accelerometers [57], and eye tracking sensors [58]. It becomes apparent that sensors in many everyday electronic devices can reveal significantly more information than one would assume based on their advertised functionality. The crafting of solutions to either limit the immense amounts of knowledge and power this creates for certain organizations, or to at least avert negative consequences for society, will be an important challenge for privacy and civil rights advocates over the years to come.

4 Conclusion

Microphones are widely used in connected devices, where they have a large variety of possible applications. While recognizing the benefits of voice and speech analysis, this paper highlights the growing privacy threat of unexpected inferences from audio data. Besides the linguistic content, a voice recording can implicitly contain information about a speaker's identity, personality, body shape, mental and physical health, age, gender, emotions, geographical origin, and socioeconomic status – and may thereby potentially reveal much more information than a speaker wishes and expects to communicate.

Further research is required into the privacy implications of microphone-equipped devices, taking into account the evolving state of the art in data mining technology. As it is impossible, however, to meaningfully determine the limits of inference methods developed behind closed doors, voice recordings – even where the linguistic content does not seem rich and revealing – should be regarded and treated as highly sensitive by default. Since existing technical and legal countermeasures are limited and do not yet offer reliable protection against large-scale misuses of audio data and undesired inferences, more effective safeguards and means of enforcement are urgently needed. We hope that the knowledge compiled in this paper can serve as a basis for consumer education and will help lawmakers and fellow researchers in assessing the richness and potential sensitivity of speech data.

References

1. Amazon: Alexa and Alexa Device FAQs. <https://www.amazon.com/gp/help/customer/display.html?nodeId=201602230>. Accessed Nov 16 2019

2. An, G., et al.: Deep personality recognition for deception detection. In: INTERSPEECH, pp. 421–425 (2018). <https://doi.org/10.21437/Interspeech.2018-2269>
3. An, G., Levitan, R.: Lexical and acoustic deep learning model for personality recognition. In: INTERSPEECH, pp. 1761–1765 (2018)
4. Aytar, Y., et al.: SoundNet: learning sound representations from unlabeled video. In: Conference on Neural Information Processing Systems (NIPS), pp. 892–900 (2016)
5. Bae, S.-G., et al.: A judgment of intoxication using hybrid analysis with pitch contour compare in speech signal processing. *IJAER* **12**(10), 2342–2346 (2017)
6. Banai, B., et al.: Candidates' voice in political debates and the outcome of presidential elections. In: Psychology Days in Zadar, pp. 33–39. University of Zadar (2017)
7. Bedi, G., et al.: A Window into the intoxicated mind? Speech as an index of psychoactive drug effects. *Neuropsychopharmacology* **39**(10), 2340–2348 (2014)
8. Bedi, G., et al.: Automated analysis of free speech predicts psychosis onset in high-risk youths. *npj Schizophr.* **1**, 15030 (2015)
9. Behravan, H., et al.: i-vector modeling of speech attributes for automatic foreign accent recognition. *Trans. Audio Speech Lang. Process.* **24**(1), 29–41 (2016)
10. Belin, P., et al.: The sound of trustworthiness: acoustic-based modulation of perceived voice personality. *PLoS ONE* **12**(10), e0185651 (2017)
11. Bernstein, B.: Language and social class. *Br. J. Sociol.* **11**(3), 271–276 (1960)
12. Bindahman, S., et al.: 3D body scanning technology: privacy and ethical issues. In: Conference on Cyber Security, Cyber Warfare and Digital Forensic, pp. 150–154 (2012)
13. Bone, D., et al.: Intoxicated speech detection. *Comput. Speech Lang.* **28**(2), 375–391 (2014). <https://doi.org/10.1016/j.csl.2012.09.004>
14. Booth, R.: Facebook reveals news feed experiment to control emotions (2014). <https://www.theguardian.com/technology/2014/jun/29/facebook-users-emotions-news-feeds>
15. Bugdol, M.D., et al.: Prediction of menarcheal status of girls using voice features. *Comput. Biol. Med.* **100**, 296–304 (2018). <https://doi.org/10.1016/j.combiomed.2017.11.005>
16. Chen, J., Kam, A.H., Zhang, J., Liu, N., Shue, L.: Bathroom activity monitoring based on sound. In: Gellersen, H.-W., Want, R., Schmidt, A. (eds.) *Pervasive 2005*. LNCS, vol. 3468, pp. 47–61. Springer, Heidelberg (2005). https://doi.org/10.1007/11428572_4
17. Christl, W.: *How Companies Use Data Against People*. Cracked Labs, Vienna (2017)
18. Christl, W., Spiekermann, S.: *Networks of Control: A Report on Corporate Surveillance, Digital Tracking, Big Data & Privacy*. Facultas, Vienna (2016)
19. Cummins, N., et al.: Speech analysis for health: current state-of-the-art and the increasing impact of deep learning. *Methods* **151**, 41–54 (2018)
20. Cummins, N., et al.: “You sound ill, take the day off”: automatic recognition of speech affected by upper respiratory tract infection. In: *IEEE EMBC*, pp. 3806–3809 (2017)
21. Desplanques, B., Demuynck, K.: Cross-lingual speech emotion recognition through factor analysis. In: INTERSPEECH, pp. 3648–3652 (2018)
22. Drozdak, N., Turner, G.: Apple, Google, and Amazon May Have Violated Your Privacy by Reviewing Digital Assistant Commands. <https://fortune.com/2019/08/05/google-apple-amazon-digital-assistants/>. Accessed 03 Sept 2019
23. Dubey, H., et al.: BigEAR: inferring the ambient and emotional correlates from smartphone-based acoustic big data. In: *IEEE CHASE*, pp. 78–83 (2016)
24. EEOC: Charge Statistics. <https://www.eeoc.gov/eeoc/statistics/enforcement/charges.cfm>. Accessed 07 Nov 2019
25. Evans, S., et al.: Relationships between vocal characteristics and body size and shape in human males. *Biol. Psychol.* **72**(2), 160–163 (2006)
26. Fast, L.A., Funder, D.C.: Personality as manifest in word use: correlations with self-report, acquaintance report, and behavior. *J. Pers. Soc. Psychol.* **94**(2), 334–346 (2008)

27. Fraser, K.C., et al.: Linguistic features identify Alzheimer's disease in narrative speech. *J. Alzheimers Dis.* **49**(2), 407–422 (2015). <https://doi.org/10.3233/JAD-150520>
28. Ghahremani, P., et al.: End-to-end deep neural network age estimation. In: INTERSPEECH, pp. 277–281 (2018). <https://doi.org/10.21437/Interspeech.2018-2015>
29. González, J.: Correlations between speakers' body size and acoustic parameters of voice. *Percept. Mot. Skills* **105**(1), 215–220 (2007)
30. Gonzalez, J., Carpi, A.: Early effects of smoking on the voice: a multidimensional study. *Med. Sci. Monit.* **10**(12), CR649–CR656 (2004)
31. Gosztolya, G., et al.: Identifying schizophrenia based on temporal parameters in spontaneous speech. In: INTERSPEECH, pp. 3408–3412 (2018)
32. Gray, S.: Always On: Privacy Implications of Microphone-Enabled Devices. Future of Privacy Forum, Washington, DC (2016)
33. Grosz, T., et al.: Assessing the degree of nativeness and Parkinson's condition using Gaussian processes and deep rectifier neural networks. In: INTERSPEECH (2015)
34. Grzybowska, J., Ziółko, M.: I-vectors in gender recognition from telephone speech. In: Conference on Applications of Mathematics in Biology and Medicine (2015)
35. Haider, F., et al.: An active feature transformation method for attitude recognition of video bloggers. In: INTERSPEECH, pp. 431–435 (2018)
36. Hashim, N.W., et al.: Evaluation of voice acoustics as predictors of clinical depression scores. *J. Voice* **31**(2), 256.e1–256.e6 (2017). <https://doi.org/10.1016/j.jvoice.2016.06.006>
37. Herms, R.: Prediction of deception and sincerity from speech using automatic phone recognition-based features. In: INTERSPEECH, pp. 2036–2040 (2016)
38. Hoff, E.: How social contexts support and shape language development. *Dev. Rev.* **26**(1), 55–88 (2006). <https://doi.org/10.1016/j.dr.2005.11.002>
39. Honig, F., et al.: Islands of failure: employing word accent information for pronunciation quality assessment of English L2 learners. In: ISCA SLATE Workshop (2009)
40. HSBC: Welcome to Voice ID. <https://www.us.hsbc.com/customer-service/voice/>. Accessed 22 Oct 2019
41. Huang, Z., et al.: Depression detection from short utterances via diverse smartphones in natural environmental conditions. In: INTERSPEECH, pp. 3393–3397 (2018)
42. Hughes, S.M., et al.: Sex-specific body configurations can be estimated from voice samples. *J. Soc. Evol. Cult. Psychol.* **3**(4), 343–355 (2009). <https://doi.org/10.1037/h0099311>
43. IEEE AASP: Challenge results published. <http://www.cs.tut.fi/sgn/arg/dc2017/articles/challenge-results-published>. Accessed 22 Oct 2019
44. Imhof, M.: Listening to voices and judging people. *Int. J. List.* **24**(1), 19–33 (2010)
45. Jain, A., et al.: Improved accented speech recognition using accent embeddings and multi-task learning. In: INTERSPEECH, pp. 2454–2458 (2018)
46. Jain, P., et al.: Big data privacy: a technological perspective and review. *J. Big Data* **3**(1), 25 (2016)
47. Jin, H., Wang, S.: Voice-based determination of physical and emotional characteristics of users (2018). <https://patents.google.com/patent/US10096319B1/en?q=10096319>
48. Kabil, S.H., et al.: On learning to identify genders from raw speech signal using CNNs. In: INTERSPEECH, pp. 287–291 (2018)
49. Kaneshiro, B., et al.: Characterizing listener engagement with popular songs using large-scale music discovery data. *Front. Psychol.* **8**, 1–15 (2017)
50. Karpey, D., Pender, M.: Customer Identification Through Voice Biometrics (2016). <https://patents.google.com/patent/US9396730>
51. Klostad, C.A., et al.: Perceptions of competence, strength, and age influence voters to select leaders with lower-pitched voices. *PLoS ONE* **10**(8), e0133779 (2015)
52. Ko, S.J., et al.: The sound of power: conveying and detecting hierarchical rank through voice. *Psychol. Sci.* **26**(1), 3–14 (2015). <https://doi.org/10.1177/0956797614553009>

53. Koolagudi, S.G., Maity, S., Kumar, V.A., Chakrabarti, S., Rao, K.S.: IITKGP-SESC: speech database for emotion analysis. In: Ranka, S., et al. (eds.) IC3 2009. CCIS, vol. 40, pp. 485–492. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03547-0_46
54. Kotenko, J.: To infinity and Beyond Verbal (2013). <https://www.digitaltrends.com/social-media/exploring-beyond-verbal-the-technology-of-emotions-analytics/>
55. Krauss, R.M., et al.: Inferring speakers’ physical attributes from their voices. *J. Exp. Soc. Psychol.* **38**(6), 618–625 (2002)
56. Kröger, J.: Unexpected inferences from sensor data: a hidden privacy threat in the internet of things. In: Strous, L., Cerf, Vinton G. (eds.) IFIPIoT 2018. IAICT, vol. 548, pp. 147–159. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-15651-0_13
57. Kröger, J.L., et al.: Privacy implications of accelerometer data: a review of possible inferences. In: Proceedings of the 3rd International Conference on Cryptography, Security and Privacy (ICCSP). ACM, New York (2019). <https://doi.org/10.1145/3309074.3309076>
58. Kröger, J.L., Lutz, O.H.-M., Müller, F.: What does your gaze reveal about you? On the privacy implications of eye tracking. In: Friedewald, M., Önen, M., Lievens, E., Krenn, S., Fricker, S. (eds.) Privacy and Identity 2019. IFIP AICT, vol. 576, pp. 226–241. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-42504-3_15
59. Kröger, J.L., Raschke, P.: Is my phone listening in? On the feasibility and detectability of mobile eavesdropping. In: Foley, S.N. (ed.) DBSec 2019. LNCS, vol. 11559, pp. 102–120. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22479-0_6
60. Levitan, S.I., et al.: Acoustic-prosodic indicators of deception and trust in interview dialogues. In: INTERSPEECH, pp. 416–420 (2018)
61. Li, J.-L., et al.: Learning conditional acoustic latent representation with gender and age attributes for automatic pain level recognition. In: INTERSPEECH (2018)
62. Lopes, J., et al.: A nativeness classifier for TED Talks. In: ICASSP, pp. 5672–5675 (2011)
63. Magi, T.J.: Fourteen reasons privacy matters: a multidisciplinary review of scholarly literature. *Libr. Q. Inf. Community Policy* **81**(2), 187–209 (2011)
64. Malhotra, K., Khosla, A.: Automatic identification of gender & accent in spoken Hindi utterances with regional Indian accents. In: IEEE SLT Workshop, pp. 309–312 (2008)
65. Mayew, W.J., et al.: Voice pitch and the labor market success of male chief executive officers. *Evol. Hum. Behav.* **34**(4), 243–248 (2013)
66. McLaren, M., et al.: The 2016 speakers in the wild speaker recognition evaluation. In: INTERSPEECH, pp. 823–827 (2016). <https://doi.org/10.21437/Interspeech.2016-1137>
67. Mendels, G., et al.: Hybrid acoustic-lexical deep learning approach for deception detection. In: INTERSPEECH, pp. 1472–1476 (2017)
68. Mohammadi, G., et al.: The voice of personality: mapping nonverbal vocal behavior into trait attributions. In: Workshop on Social Signal Processing (SSPW), pp. 17–20 (2010)
69. Mporas, I., Ganchev, T.: Estimation of unknown speaker’s height from speech. *Int. J. Speech Technol.* **12**(4), 149–160 (2009). <https://doi.org/10.1007/s10772-010-9064-2>
70. Naim, I., et al.: Automated prediction and analysis of job interview performance. In: IEEE Conference on Automatic Face and Gesture Recognition, pp. 1–6 (2015)
71. Nandwana, M.K., et al.: Robust speaker recognition from distant speech under real reverberant environments using speaker embeddings. In: INTERSPEECH (2018)
72. Nautsch, A., et al.: Preserving privacy in speaker and speech characterisation. *Comput. Speech Lang.* **58**, 441–480 (2019). <https://doi.org/10.1016/j.csl.2019.06.001>
73. Nautsch, A., et al.: The GDPR & speech data: reflections of legal and technology communities, first steps towards a common understanding. In: INTERSPEECH, pp. 3695–3699 (2019). <https://doi.org/10.21437/Interspeech.2019-2647>
74. Norel, R., et al.: Detection of amyotrophic lateral sclerosis (ALS) via acoustic analysis. In: INTERSPEECH, pp. 377–381 (2018). <https://doi.org/10.1101/383414>

75. Ooi, K.E.B., et al.: Multichannel weighted speech classification system for prediction of major depression in adolescents. *IEEE Trans. Biomed. Eng.* **60**(2), 497–506 (2013)
76. Perez, M., et al.: Classification of huntington disease using acoustic and lexical features. In: *INTERSPEECH*, pp. 1898–1902 (2018)
77. Petrushin, V.A.: Detecting emotions using voice signal analysis (2007). <https://patents.google.com/patent/US7222075B2/en>
78. Pipitone, R.N., Gallup, G.G.: Women's voice attractiveness varies across the menstrual cycle. *Evol. Hum. Behav.* **29**(4), 268–274 (2008)
79. Polzehl, T., et al.: Automatically assessing personality from speech. In: *IEEE Conference on Semantic Computing (ICSC)*, pp. 134–140 (2010)
80. Polzehl, T.: *Personality in Speech*. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-09516-5>
81. Ponsot, E., et al.: Cracking the social code of speech prosody using reverse correlation. *Proc. Natl. Acad. Sci.* **115**(15), 3972–3977 (2018)
82. Ranganath, R., et al.: Detecting friendly, flirtatious, awkward, and assertive speech in speed-dates. *Comput. Speech Lang.* **27**(1), 89–115 (2013)
83. Reavley, N.J., Jorm, A.F.: Experiences of discrimination and positive treatment in people with mental health problems. *Aust. N. Z. J. Psychiatry* **49**(10), 906–913 (2015)
84. Reubold, U., et al.: Vocal aging effects on F0 and the first formant: a longitudinal analysis in adult speakers. *Speech Commun.* **52**(7–8), 638–651 (2010)
85. Sadjadi, S.O., et al.: Speaker age estimation on conversational telephone speech using senone posterior based i-vectors. In: *ICASSP*, pp. 5040–5044 (2016)
86. Sarma, M., et al.: Emotion identification from raw speech signals using DNNs. In: *INTERSPEECH*, pp. 3097–3101 (2018). <https://doi.org/10.21437/Interspeech.2018-1353>
87. Schatzman, L., Strauss, A.: Social class and modes of communication. *Am. J. Sociol.* **60**(4), 329–338 (1955). <https://doi.org/10.1086/221564>
88. Schuller, B., et al.: A survey on perceived speaker traits: personality, likability, pathology, and the first challenge. *Comput. Speech Lang.* **29**(1), 100–131 (2015)
89. Schuller, B., et al.: Medium-term speaker states - a review on intoxication, sleepiness and the first challenge. *Comput. Speech Lang.* **28**(2), 346–374 (2013)
90. Seligman, C.R., et al.: The effects of speech style and other attributes on teachers' attitudes toward pupils. *Lang. Soc.* **1**(1), 131 (1972)
91. Sim, J.M., et al.: Acoustic sensor based recognition of human activity in everyday life for smart home services. *Int. J. Distrib. Sens. Netw.* **11**(9), 679123 (2015)
92. Simpson, A.P.: Phonetic differences between male and female speech. *Lang. Linguist. Compass* **3**(2), 621–640 (2009). <https://doi.org/10.1111/j.1749-818X.2009.00125.x>
93. Sloboda, J., et al.: Vocal biomarkers for cognitive performance estimation in a working memory task. In: *INTERSPEECH*, pp. 1756–1760 (2018)
94. Soskin, W.F., Kauffman, P.E.: Judgment of emotion in word-free voice samples. *J. Commun.* **11**(2), 73–80 (1961). <https://doi.org/10.1111/j.1460-2466.1961.tb00331.x>
95. Stanek, M., Sigmund, M.: Psychological stress detection in speech using return-to-opening phase ratios in glottis. *Elektron Elektrotech.* **21**(5), 59–63 (2015)
96. Stanescu, C.G., Ievchuk, N.: Alexa, where is my private data? In: *Digitalization in Law*, pp. 237–247. Social Science Research Network, Rochester (2018)
97. Stowell, D., et al.: Detection and classification of acoustic scenes and events. *IEEE Trans. Multimed.* **17**(10), 1733–1746 (2015). <https://doi.org/10.1109/TMM.2015.2428998>
98. Streeter, L.A., et al.: Pitch changes during attempted deception. *J. Pers. Soc. Psychol.* **35**(5), 345–350 (1977). <https://doi.org/10.1037//0022-3514.35.5.345>
99. Swain, M., et al.: Databases, features and classifiers for speech emotion recognition: a review. *Int. J. Speech Technol.* **21**(1), 93–120 (2018)

100. Trilok, N.P., et al.: Establishing the uniqueness of the human voice for security applications. In: Proceedings of Student-Faculty Research Day, pp. 8.1–8.6. Pace University (2004)
101. Tsai, F.-S., et al.: Automatic assessment of individual culture attribute of power distance using a social context-enhanced prosodic network representation. In: INTERSPEECH, pp. 436–440 (2018). <https://doi.org/10.21437/Interspeech.2018-1523>
102. Vergyri, D., et al.: Speech-based assessment of PTSD in a military population using diverse feature classes. In: INTERSPEECH, pp. 3729–3733 (2015)
103. Vukovic, J., et al.: Women’s voice pitch is negatively correlated with health risk factors. *J. Evol. Psychol.* **8**(3), 217–225 (2010). <https://doi.org/10.1556/JEP.8.2010.3.2>
104. Walia, G.S., Sharma, R.K.: Level of asthma: mathematical formulation based on acoustic parameters. In: CASP, pp. 24–27 (2016). <https://doi.org/10.1109/CASP.2016.7746131>
105. Wolfson, S.: Amazon’s Alexa recorded private conversation and sent it to random contact (2018). <https://www.theguardian.com/technology/2018/may/24/amazon-alexa-recorded-conversation>
106. Zhao, J., et al.: Privacy-preserving machine learning based data analytics on edge devices. In: AAAI/ACM Conference on AI, Ethics, and Society (AIES), pp. 341–346 (2018)
107. Zhou, Y., et al.: Deception detecting from speech signal using relevance vector machine and non-linear dynamics features. *Neurocomputing* **151**, 1042–1052 (2015)
108. Annual Report 2018. Amazon.com, Inc., Seattle, Washington, USA (2019)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



Privacy Implications of Accelerometer Data: A Review of Possible Inferences

Jacob Leon Kröger
Technische Universität Berlin
Hardenbergstraße 32
10623 Berlin
+49 30 7001 41037
kroeger@tu-berlin.de

Philip Raschke
Technische Universität Berlin
Ernst-Reuter-Platz 7
10587 Berlin
+49 30 8353 58353
philip.raschke@tu-berlin.de

Towhidur Rahman Bhuiyan
Technische Universität Berlin
Hardenbergstraße 32
10623 Berlin
+49 30 7001 41001
t.bhuiyan@campus.tu-berlin.de

ABSTRACT

Accelerometers are sensors for measuring acceleration forces. They can be found embedded in many types of mobile devices, including tablet PCs, smartphones, and smartwatches. Some common uses of built-in accelerometers are automatic image stabilization, device orientation detection, and shake detection. In contrast to sensors like microphones and cameras, accelerometers are widely regarded as not privacy-intrusive. This sentiment is reflected in protection policies of current mobile operating systems, where third-party apps can access accelerometer data without requiring security permission. It has been shown in experiments, however, that seemingly innocuous sensors can be used as a side channel to infer highly sensitive information about people in their vicinity. Drawing from existing literature, we found that accelerometer data alone may be sufficient to obtain information about a device holder's location, activities, health condition, body features, gender, age, personality traits, and emotional state. Acceleration signals can even be used to uniquely identify a person based on biometric movement patterns and to reconstruct sequences of text entered into a device, including passwords. In the light of these possible inferences, we suggest that accelerometers should urgently be re-evaluated in terms of their privacy implications, along with corresponding adjustments to sensor protection mechanisms.

CCS Concepts

• Security and privacy

Keywords

Accelerometer, Sensor, Privacy, Side channel, Inference attack

1. INTRODUCTION

An accelerometer is an instrument for measuring acceleration forces caused by the movements and vibrations of an object, or by gravity. Today, all sorts of mobile devices, including smartphones, tablet PCs, smartwatches, digital cameras, wearable fitness trackers, game controllers, and virtual reality headsets, are equipped with built-in microelectromechanical accelerometers [1]. Studies even suggest that accelerometers are the most widely used sensor in wearable devices [2] and also the sensor that is most frequently accessed by mobile apps [3].

Among other common applications, acceleration signals are used for image stabilization in cameras, for measuring the orientation of a device relative to Earth's gravitational pull (e.g. to enable automatic display rotation between landscape and portrait mode), and for detecting user actions, such as moving or shaking a device.

While some sensors, such as microphones, cameras and GPS, are widely perceived as privacy-sensitive [4, 5] and require explicit user permission to be activated in current mobile operating systems [3], accelerometers are less well-understood in terms of their privacy implications, and also much less protected [6, 7]. Even scholarly literature has largely ignored potential issues in this field, with researchers describing accelerometer data as “not particularly sensitive” [8] or even “privacy preserving” [9].

Experimental studies have shown, however, that sensitive personal data can be inferred from accelerometer readings. This paper presents a non-exhaustive overview of possible inferences, drawing from multiple academic disciplines, including information science, psychology, health science, and computer science. According to our findings, accelerometers in mobile devices may reveal information about a user's activities (section 2.1), location (sect. 2.2), identity (sect. 2.3), device inputs (sect. 2.4), health condition and body features (sect. 2.5), age and gender (sect. 2.6), moods and emotions (sect. 2.7), and personality traits (sect. 2.8).

2. POSSIBLE INFERENCE

In this chapter, we present experimental studies from the scholarly literature in which sensitive information was successfully derived from accelerometer data. A visual overview is provided in Fig. 3, at the end of the chapter.

2.1 Activity and Behavior Tracking

A wide range of physical activity variables and behavior-related information can be derived from raw accelerometer data. Accelerometer-based pedometers (“step counters”), for instance, register the impacts produced by steps during motion and can estimate energy expenditure and distance walked [10]. In medical studies, wearable devices with embedded accelerometers are

widely used to assess the amount of sedentary time and physical activity among patients [11, 12].

Body-worn accelerometers have also been shown to enable real-time body posture and activity classification. High recognition accuracy has been achieved for basic physical activities, including running, walking, cycling, lying, climbing stairs, falling, sitting and standing [13–16], as well as for more complex activities, such as writing, reading, typing, painting, sorting paperwork or searching the internet [17].

Not only the type but also the duration of activities and temporal behavior patterns can be derived from acceleration signals [18, 19]. When worn during the night, mobile devices with built-in accelerometers may enable sleep-wake cycle monitoring, through variables such as sleep onset and offset, total sleep time and sleep intervals [20, 21], as well as the monitoring of sleep-related behaviors [11].

Accelerometers in handheld and wrist-worn devices can further be used to detect specific hand gestures [22], eating and drinking moments [23, 24], and smoking [25, 26]. Gait features of subjects, extracted from accelerometer data, can even reveal their level of intoxication. Researchers were able to distinguish “sober walk” from “intoxicated walk” [27] and to estimate blood alcohol content [28] as well as the number of drinks consumed [29] via accelerometry alone.

In [17], signals from a single body-worn accelerometer were used to detect if a subject is carrying a load. Accelerometer-based gait dynamics have also been used to estimate the weight of carried objects with robustness to variations in walking speeds, body types and walking conditions [30].




Figure 1: Classification of driving patterns based on streams of accelerometer data, from [31].

When located inside a car, motion sensors can be used to measure an operator’s driving behavior. In [31], Singh, Juneja and Kapoor identified events such as sudden breaking, sudden acceleration, right and left turns and lane changes from patterns in accelerometer data, as is illustrated in Fig. 1. From such information, researchers were able to detect aggressive or unsafe driving styles [32] and drunk driving patterns [33].

Based on indicative body movements and sound vibrations, both measured using accelerometers, researchers were able to derive

speech activity and social interactions of subjects [9, 34]. Even ways of reconstructing speech solely from recorded vibrations have been explored. AccelWord, developed in [35], can detect hotwords spoken by a user, utilizing accelerometer data from commercially available mobile devices. Patents have already been filed for a “method of detecting a user’s voice activity using an accelerometer” [36] and a “system that uses an accelerometer in a mobile device to detect hotwords” [37].

2.2 Location Tracking

It has been shown that accelerometers in mobile devices can be exploited for user localization and reconstruction of travel trajectories, even when other localization systems, such as GPS, are disabled. In [38], Han et al. were able to geographically track a person who is driving a car based solely on accelerometer readings from the subject’s smartphone. In their approach, they first calculate the vehicle’s approximate motion trajectory using three-axis acceleration measurements from an iPhone located inside the vehicle, and then map the derived trajectory to the shape of existing routes on a map. An example application of the algorithm is displayed in Fig. 2. Han et al. describe their results as “comparable to the typical accuracy for handheld global positioning systems.”

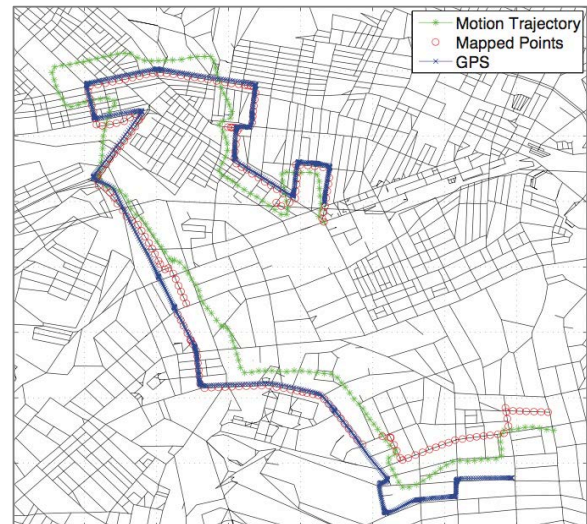


Figure 2: Map matching algorithm used in [38]. The green trail indicates the motion trajectory obtained from accelerometer data. The red trail indicates the inferred route. The blue trail indicates the actual route traveled (GPS data).

Hua, Shen and Zhong found that accelerometers in smartphones can also reveal the device’s location while the holder is using a metropolitan train system [39]. To achieve this, the researchers compare and match acceleration patterns with labeled training data to recognize specific station intervals through which the user travels. Results from experiments on a real metro line show that the accuracy of their approach could reach up to 89% and 92% if the metro ride is longer than 3 or 5 stations, respectively [39].

2.3 User Identification

Body movement patterns recorded by accelerometers in mobile devices have been demonstrated to be discriminative enough to differentiate between, or even uniquely identify, users. Various accelerometer-only approaches have been proposed to confirm the identity of a user based on biometric gait features [40, 41], hand gestures [42], or head movements [43]. Using accelerometer rea-

dings from smartphones, Kwapisz, Weiss and Moore were able to recognize individuals from a pool of 36 test subjects with 100% accuracy [44].

It has also been shown that, through aerial vibrations, accelerometers can be sensitive enough to capture sound, including human speech, in sufficient quality to distinguish between different speakers with high accuracy [35].

The location trajectory of a mobile device, which can be inferred from accelerometer data under certain conditions (as explained in section 2.2), may reveal a user's work and home addresses [45], and – in conjunction with white pages, employment directories, tax records, or other auxiliary datasets – a user's real identity [46].

Following an approach commonly referred to as *device fingerprinting*, users can further be told apart based on unique characteristics and features of their personal devices. Calibration errors in accelerometers, which are caused by imperfections in the manufacturing process, have been found sufficient to uniquely identify their encapsulating device [6, 47]. Such a “fingerprint” can be used, for instance, to track users across repeated website visits, even when private browsing is activated and other tracking technologies, such as canvas fingerprinting or cookies, are blocked [48].

2.4 Keystroke Logging

The input that users type into their devices through touchscreens and keyboards contains highly sensitive information such as text messages, personal notes, login credentials and transaction details.

Based on the observation that swipes, taps and keystrokes often correlate with distinctive hand movements of the user, it has been shown that inputs can be reconstructed using motion sensor data from handheld and wrist-worn devices [49–51]. Some researchers have exclusively used accelerometer data for such keystroke inference attacks. Aviv et al. demonstrated that accelerometers in smartphones can be exploited to infer tap- and gesture-based input, including PINs and graphical password patterns [52]. Based on the same type of data, Owusu et al. were able to obtain entire sequences of text entered through a phone's touchscreen [53].

Through examining the source code of other existing approaches, it has been found that even multi-sensor attacks solely use acceleration information for tap detection, leading to the conclusion that defense mechanisms against these kinds of side channel attacks should focus on accelerometers [54].

Not only does the above imply that accelerometer data could offer sensitive insights into a user's communication and transactions: Beltramelli and Risi even warn that a user's entire technological ecosystem could be compromised when passwords are leaked through embedded sensors in consumer electronics [55].

2.5 Inference of Health Parameters and Body Features

Body-worn accelerometers can be used to gain insight into a person's physical characteristics and health status. Using accelerometer data from smartphones, researchers were able to derive an approximation of the body weight and height of users [56, 57]. A strong correlation has been observed between accelerometer-determined physical activity and obesity [58].

Physical activity is generally recognized as a promoter and indicator of health [59]. A person's amount of physical activity can reveal sensitive information about latent chronic diseases and the person's degree of mobility [12] as well as about cognitive function and even risk of mortality [60]. As explained in section

2.1, a wide range of activity-related variables can be derived from accelerometer data, including energy expenditure, type of activity and temporal activity patterns. This association is increasingly put to use in health studies, where accelerometers are used to remotely assess the physical activity level of participants [61].

Another important factor in population health is the amount of sleep that people get. Sleep loss has been associated with developing serious illnesses, such as cardiovascular disease and diabetes, and even with increased all-cause mortality [62]. Numerous studies have shown that accelerometers in wearable devices can be used for evaluating sleep patterns [20], sleep fragmentation [63] and sleep efficiency [64]. Actigraphy, an accelerometer-based assessment method, has been described as an “essential tool in sleep research and sleep medicine” [20]. Experimental results from Pesonen and Kuula suggest that accelerometers in consumer-targeted wearables can be as effective for sleep monitoring as research-targeted devices [21].

Specialized accelerometers have been used to measure various other health parameters, including voice health [65], postural stability [12] and physiological sound [66].

2.6 Inference of Demographics

Estimates of demographic variables such as age and gender can be made based on data from body-worn accelerometers. It has long been demonstrated that adults and children differ in their smoothness of walking, which is reflected in accelerometer readings [67]. Menz, Lord and Fitzpatrick compared gait features between young and elder subjects using acceleration signals and discovered that younger subjects showed greater step length, higher velocity and smaller step timing variability [68]. Using data from accelerometers in smartphones, Davarci et al. were able to predict the age interval of test subjects with a success rate of 92.5% [69]. Their work is based on the observation that children and adults differ in the way they hold and touch smartphones.

Experimental results by Cho, Park and Kwon indicate that there are also gender-specific movement patterns [70]. In accordance, research has shown that it is possible to estimate the sex of individuals based on hip movements [56], gait features [71] and physical activity patterns [72], all derived from accelerometer data. An experiment also revealed that female gait patterns are significantly influenced by the heel height of their shoes [73]. Weiss and Lockhart emphasize that accelerometer-based gender recognition can work independently of a subject's weight and height [56]. Even acoustic vibrations caused by a person's voice and captured through a smartphone accelerometer can be used to classify speakers into male and female with high accuracy [35].

2.7 Mood and Emotion Recognition

The level of physical activity, which can be measured using body-worn accelerometers (see section 2.1), has been identified as a potential predictor of human emotions [74] and depressive moods [75]. Zhang et al. were able to recognize emotional states of test subjects (happy, neutral, and angry) with fair accuracy, relying only on accelerometer data from smart wristbands [76]. Accelerometers in smartphones have been used to detect stress levels [77] and arousal [78] in users. Also, Matic et al. found a positive association between accelerometer-derived speech activity and mood changes [9].

2.8 Inference of Personality Traits

Methods have been proposed for inferring preferences and other personality traits solely from body gestures and motion patterns. Englebienne and Hung used wearable accelerometers to estimate the motivations, interests and group affiliations of study

participants in scenarios of social interaction, based on their movements, body postures and expansiveness of gesturing [34].

A person's level of physical activity, which can also be measured using body-worn accelerometers (see section 2.1), has been shown to correlate with certain personality traits such as conscientiousness, neuroticism, openness, and extraversion [79]. Artese et al. evaluated the body movements of test subjects for seven days using accelerometer-based monitoring devices and found that agreeableness, conscientiousness and extraversion were positively and neuroticism negatively associated to more steps per day and other physical activity variables [80]. Examining correlates between the personality and physical activity of female college students, Wilson et al. discovered that neuroticism and the functioning of the behavioral inhibition system were both related to physical activity measures derived from accelerometer readings [81].

3. DISCUSSION AND IMPLICATIONS

As shown in the previous section, accelerometers in mobile devices can allow serious invasions of user privacy. Even when other sensors, such as cameras, microphones and GPS are turned off, accelerometer data can be sufficient to obtain information about a device holder's location, health condition, body features, age, gender, emotions and personality traits. Acceleration signals may even be used to uniquely identify a person based on biometric movement patterns and to reconstruct sequences of text entered into a device.

It has to be acknowledged that most experimental studies cited in this paper have substantial limitations. First, many approaches were only tested in controlled laboratory settings [14, 17, 24, 26, 32, 33, 35, 40, 41, 43, 53, 57]. For methods applied under real-life

conditions, considerable reductions in accuracy have been observed [9, 82]. Second, several of the presented methods require prior knowledge about the user or the user's context in order to function [39–44, 52]. Third, subjects in some of the experiments wore accelerometers attached to certain body parts, such as chest [9, 15], hip [40], waist [14], or even multiple body parts [24, 25, 64], whereas in reality, mobile devices are mostly worn around the wrist [23] or interchangeably in hands, bags, and pockets [83]. In light of these limitations, the real-world applicability of the presented methods can be questioned.

On the other hand, it may reasonably be assumed that at least some of the parties who regularly access accelerometer data from consumer devices (e.g. device manufacturers, service providers, app developers) possess larger sets of training data, more technical expertise and more financial resources than the researchers cited in this paper. Furthermore, data from other sensors and auxiliary data may be available to potential adversaries, improving their capability to draw sensitive inferences, while the methods considered in this paper solely rely on accelerometer data. Thus, our work represents only an initial and non-exhaustive exploration of the topic.

It would be enough if even one of the identified threats is realized, however, for user privacy to be seriously impacted. Also, it seems probable that the risk will continue to grow with further improvements of sensor technologies in terms of cost, size and accuracy, further advances in machine learning methods, and further proliferation of accelerometer-equipped mobile devices.

Given the widespread perception of accelerometers as non-intrusive, we call for an urgent reconsideration of their privacy implications, along with corresponding adjustments to technical

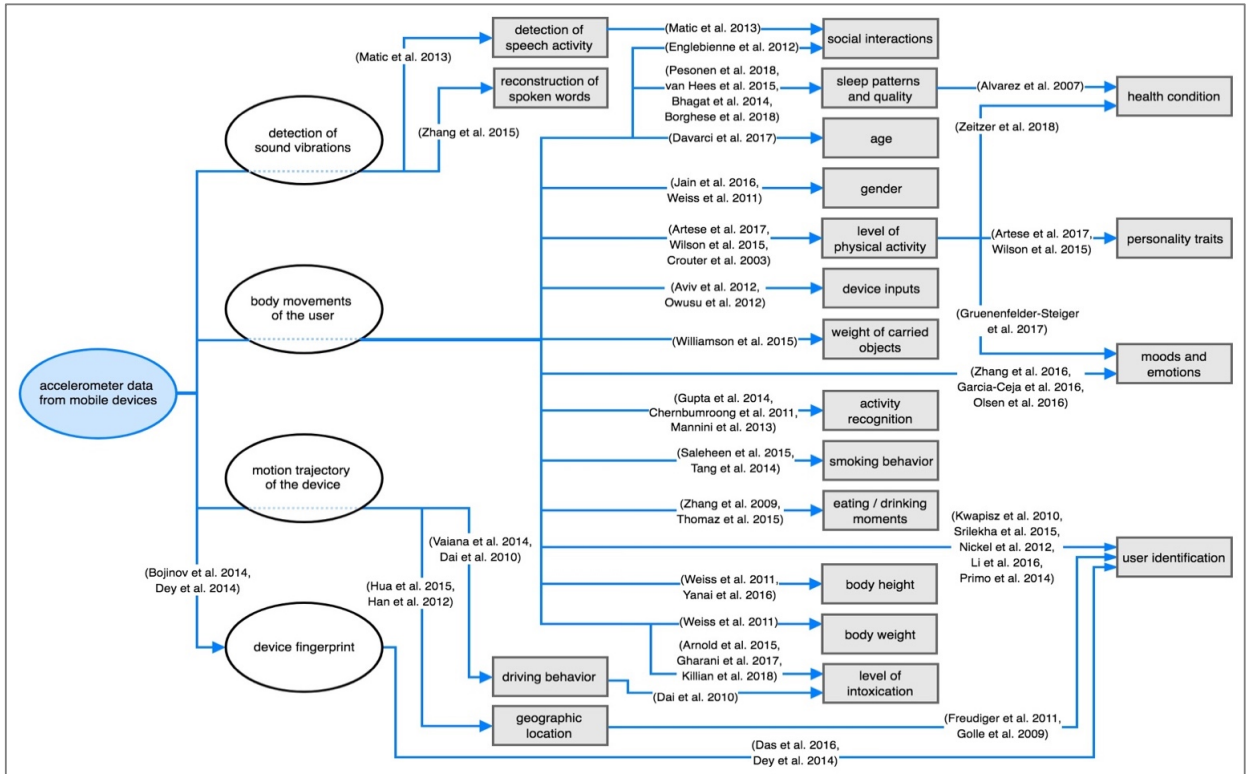


Figure 3: Overview of sensitive inferences that can be drawn from accelerometer data (according to the referenced studies).

and legal protection measures. In our opinion, the sensitivity of sensor data should generally be assessed in consideration of all inferences that could plausibly be drawn from it, and not based on the sensor's official purpose. Further research into the privacy-intrusion potential of accelerometers and other seemingly benign sensors is needed, taking into account state-of-the-art data mining techniques. As it is extremely difficult, however, to meaningfully determine the limits of continuously advancing inference methods, most sensors in mobile devices should be regarded and treated as highly sensitive by default.

4. CONCLUSION

Accelerometers are among the most widely used sensors in mobile devices, where they have a large variety of possible applications. They are commonly regarded as not privacy-intrusive and therefore often less access-restricted than other sensors, such as cameras and microphones. However, based on existing literature, we found that accelerometer data can enable serious privacy intrusions by allowing inferences about a device holder's location, identity, demographics, personality, health status, emotions, activities and body features.

Any trait or behavior of a user that results in characteristic movement patterns can potentially be detected through acceleration signals. Accelerometers are cheap, low in power consumption and often invisibly embedded into consumer devices. Thus, they represent a perfect surveillance tool as long as their data streams are not properly monitored and protected from potentially untrusted parties such as device manufacturers, service providers and app developers. In current mobile operating systems, third-party apps can access accelerometer data without requiring any permission or conscious participation from the user.

Although this paper conveys only a first impression of the privacy violations that could be enabled through accelerometers, the findings already are significant enough to express a warning to consumers who could be affected, as well as a call for action to the public and private actors who are entrusted with protecting user privacy in mobile devices.

5. REFERENCES

- [1] Wearable Devices That Have an Accelerometer: 2018. <https://vandrico.com/wearables/device-categories/components/accelerometer>. Accessed: 2018-06-06.
- [2] Richardson, S. and Mackinnon, D. 2017. *Left to their own Devices? Privacy Implications of Wearable Technology in Canadian Workplaces*. Surveillance Studies Centre.
- [3] Bai, X. et al. 2017. Sensor Guardian: prevent privacy inference on Android sensors. *EURASIP Journal on Information Security*. 2017, 1 (Dec. 2017). DOI:<https://doi.org/10.1186/s13635-017-0061-8>.
- [4] Hnat, T.W. et al. 2012. Doorjamb: unobtrusive room-level tracking of people in homes using doorway sensors. *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems* (2012), 309–322.
- [5] Klasnja, P. et al. 2009. Exploring privacy concerns about personal sensing. *International Conference on Pervasive Computing* (2009), 176–183.
- [6] Bojinov, H. et al. 2014. Mobile device identification via sensor fingerprinting. *arXiv:1408.1416*. (2014).
- [7] Xu, Z. and Zhu, S. 2015. SemaDroid: A Privacy-Aware Sensor Management Framework for Smartphones. *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy* (2015), 61–72.
- [8] Weiss, G.M. et al. 2016. Actitracker: A Smartphone-Based Activity Recognition System for Improving Health and Well-Being. *IEEE International Conference on Data Science and Advanced Analytics (DSAA)* (Oct. 2016), 682–688.
- [9] Matic, A. et al. 2013. Automatic Sensing of Speech Activity and Correlation with Mood Changes. *Pervasive and Mobile Sensing and Computing for Healthcare*. Springer Berlin Heidelberg, 195–205.
- [10] Crouter, S. et al. 2003. Validity of 10 Electronic Pedometers for Measuring Steps, Distance, and Energy Cost. *Med. Sci. Sport Exerc.* 35, (2003), 1455–60.
- [11] Migueles, J. et al. 2017. Accelerometer Data Collection and Processing Criteria to Assess Physical Activity and Other Outcomes: A Systematic Review and Practical Considerations. *Sports Medicine*. 47, (2017).
- [12] Yang, C.-C. and Hsu, Y.-L. 2010. A Review of Accelerometry-Based Wearable Motion Detectors for Physical Activity Monitoring. *Sensors*. 10, 8 (Aug. 2010), 7772–7788. DOI:<https://doi.org/10.3390/s100807772>.
- [13] Chernbumroong, S. et al. 2011. Activity classification using a single wrist-worn accelerometer. *2011 5th International Conference on Software, Knowledge Information, Industrial Management and Applications (SKIMA) Proceedings* (Sep. 2011), 1–6.
- [14] Gupta, P. and Dallas, T. 2014. Feature Selection and Activity Recognition System Using a Single Triaxial Accelerometer. *IEEE Transactions on Biomedical Engineering*. 61, 6 (Jun. 2014), 1780–1786. DOI:<https://doi.org/10.1109/TBME.2014.2307069>.
- [15] Khan, A.M. et al. 2010. A Triaxial Accelerometer-Based Physical-Activity Recognition via Augmented-Signal Features and a Hierarchical Recognizer. *IEEE Transactions on Information Technology in Biomedicine*. 14, 5 (Sep. 2010), 1166–1172. DOI:<https://doi.org/10.1109/TITB.2010.2051955>.
- [16] Lee, J.-V. et al. 2013. Smart Elderly Home Monitoring System with an Android Phone. *International Journal of Smart Home*. 7, 3 (2013), 16.
- [17] Mannini, A. et al. 2013. Activity recognition using a single accelerometer placed at the wrist or ankle. *Med. Sci. Sport Exerc.* 45, 11 (Nov. 2013), 2193–2203. DOI:<https://doi.org/10.1249/MSS.0b013e31829736d6>.
- [18] Tapia, E.M. 2008. *Using machine learning for real-time activity recognition and estimation of energy expenditure*. Massachusetts Institute of Technology.
- [19] Taraldsen, K. et al. 2012. Physical activity monitoring by use of accelerometer-based body-worn sensors in older adults: A systematic literature review of current knowledge and applications. *Maturitas*. 71, 1 (Jan. 2012), 13–19. DOI:<https://doi.org/10.1016/j.maturitas.2011.11.003>.
- [20] Bhagat, Y.A. et al. 2014. Clinical validation of a wrist actigraphy mobile health device for sleep efficiency analysis. *2014 IEEE Healthcare Innovation Conference (HIC)* (Oct. 2014), 56–59.
- [21] Pesonen, A.-K. and Kuula, L. 2018. The Validity of a New Consumer-Targeted Wrist Device in Sleep Measurement: An Overnight Comparison Against Polysomnography in Children and Adolescents. *J. Clin. Sleep Med.* 14, 04 (Apr. 2018), 585–591. DOI:<https://doi.org/10.5664/jcs.m.7050>.
- [22] Liu, J. et al. 2009. uWave: Accelerometer-based Personalized Gesture Recognition and Its Applications.

- Pervasive Mob. Comput.* 5, 6 (Dec. 2009), 657–675.
DOI:https://doi.org/10.1016/j.pmcj.2009.07.007.
- [23] Thomaz, E. et al. 2015. A Practical Approach for Recognizing Eating Moments with Wrist-Mounted Inertial Sensing. *UBICOMP*. 2015, (Sep. 2015), 1029–1040. DOI:https://doi.org/10.1145/2750858.2807545.
 - [24] Zhang, S. et al. 2009. Detection of Activities by Wireless Sensors for Daily Life Surveillance: Eating and Drinking. *Sensors*. 9, 3 (Mar. 2009), 1499–1517. DOI:https://doi.org/10.3390/s90301499.
 - [25] Saleheen, N. et al. 2015. puffMarker: A Multi-Sensor Approach for Pinpointing the Timing of First Lapse in Smoking Cessation. *UBICOMP*. 2015, (Sep. 2015), 999–1010.
 - [26] Tang, Q. 2014. Automated Detection of Puffing and Smoking with Wrist Accelerometers. *8th International Conference on Pervasive Computing Technologies for Healthcare* (2014).
 - [27] Killian, J. 2018. *Smartphone-Based Intelligent System: Training AI to Track Sobriety Using Smartphone Motion Sensors*. The Ohio State University.
 - [28] Gharani, P. et al. 2017. An Artificial Neural Network for Gait Analysis to Estimate Blood Alcohol Content Level. *Computing Research Repository*. (2017).
 - [29] Arnold, Z. et al. 2015. Smartphone Inference of Alcohol Consumption Levels from Gait. *2015 International Conference on Healthcare Informatics* (Oct. 2015), 417–426.
 - [30] Williamson, J.R. et al. 2015. Estimating load carriage from a body-worn accelerometer. *BSN* (Jun. 2015), 1–6.
 - [31] Singh, P. et al. 2013. Using Mobile Phone Sensors to Detect Driving Behavior. *Proceedings of the 3rd ACM Symposium on Computing for Development* (2013), 53:1–53:2.
 - [32] Vaiana, R. et al. 2014. Driving Behavior and Traffic Safety: An Acceleration-Based Safety Evaluation Procedure for Smartphones. *Modern Applied Science*. 8, (2014), 88–96.
 - [33] Dai, J. et al. 2010. Mobile phone based drunk driving detection. *2010 4th International Conference on Pervasive Computing Technologies for Healthcare* (Mar. 2010), 1–8.
 - [34] Englebienne, G. and Hung, H. 2012. Mining for motivation: using a single wearable accelerometer to detect people’s interests. *Proceedings of the 2nd ACM international workshop on Interactive multimedia on mobile and portable devices* (2012), 23.
 - [35] Zhang, L. et al. 2015. AccelWord: Energy Efficient Hotword Detection through Accelerometer. *MOBISYS* (2015), 301–315.
 - [36] Dusan, S.V. et al. 2016. System and method of detecting a user’s voice activity using an accelerometer. US9438985B2. Sep. 6, 2016.
 - [37] Mohapatra, P. et al. 2017. Energy-efficient, accelerometer-based hotword detection to launch a voice-control system. US20170316779A1. Nov. 2, 2017.
 - [38] Han, J. et al. 2012. ACComplix: Location inference using accelerometers on smartphones. *Fourth International Conference on Communication Systems and Networks* (Jan. 2012), 1–9.
 - [39] Hua, J. et al. 2015. We Can Track You If You Take the Metro: Tracking Metro Riders Using Accelerometers on Smartphones. *arXiv:1505.05958*. (May 2015).
 - [40] Nickel, C. et al. 2012. Authentication of Smartphone Users Based on the Way They Walk Using k-NN Algorithm. *Proceedings of the 2012 Eighth International Conference on Intelligent Information Hiding and Multimedia Signal Processing* (2012), 16–20.
 - [41] Primo, A. et al. 2014. Context-Aware Active Authentication Using Smartphone Accelerometer Measurements. *IEEE Conference on Computer Vision and Pattern Recognition Workshops* (Jun. 2014), 98–105.
 - [42] Srilekha, R. and Jayakumar, D. 2015. A Secure Screen Lock System for Android Smart Phones Using Accelerometer Sensor. *International Journal For Science Technology And Engineering*. 1, 10 (May 2015), 96–100.
 - [43] Li, S. et al. 2016. Whose move is it anyway? Authenticating smart wearable devices using unique head movement patterns. *IEEE International Conference on Pervasive Computing and Communications* (Mar. 2016), 1–9.
 - [44] Kwapisz, J.R. et al. 2010. Cell phone-based biometric identification. *IEEE International Conference on Biometrics: Theory, Applications and Systems* (Sep. 2010), 1–7.
 - [45] Freudiger, J. et al. 2011. Evaluating the privacy risk of location-based services. *International conference on financial cryptography and data security* (2011), 31–46.
 - [46] Golle, P. and Partridge, K. 2009. On the anonymity of home/work location pairs. *International Conference on Pervasive Computing* (2009), 390–397.
 - [47] Dey, S. et al. 2014. AccelPrint: Imperfections of Accelerometers Make Smartphones Trackable. *Network and Distributed System Security Symposium* (2014).
 - [48] Das, A. et al. 2016. Tracking Mobile Web Users Through Motion Sensors: Attacks and Defenses. *Network and Distributed System Security Symposium* (2016).
 - [49] Cai, L. and Chen, H. 2011. TouchLogger: Inferring Keystrokes on Touch Screen from Smartphone Motion. *HotSec*. 11, (2011), 9–9.
 - [50] Wang, H. et al. 2015. MoLe: Motion Leaks Through Smartwatch Sensors. *MOBICOM* (2015), 155–166.
 - [51] Xu, Z. et al. 2012. Taplogger: Inferring user inputs on smartphone touchscreens using on-board motion sensors. *Proceedings of the fifth ACM conference on Security and Privacy in Wireless and Mobile Networks* (2012), 113–124.
 - [52] Aviv, A.J. et al. 2012. Practicality of Accelerometer Side Channels on Smartphones. *Proceedings of the 28th Annual Computer Security Applications Conference* (2012), 41–50.
 - [53] Owusu, E. et al. 2012. ACCessory: Password Inference Using Accelerometers on Smartphones. *Proceedings of the Twelfth Workshop on Mobile Computing Systems* (2012), 9:1–9:6.
 - [54] Song, Y. et al. 2014. Two Novel Defenses against Motion-Based Keystroke Inference Attacks. *Computing Research Repository*. (2014).
 - [55] Beltramelli, T. and Risi, S. 2015. Deep-Spying: Spying using Smartwatch and Deep Learning. *arXiv:1512.05616*. (Dec. 2015).
 - [56] Weiss, G.M. and Lockhart, J.W. 2011. Identifying user traits by mining smart phone accelerometer data. *Proceedings of the Fifth International Workshop on Knowledge Discovery from Sensor Data* (2011).

- [57] Yanai, H.-F. and Enjyoji, A. 2016. Estimating Carrier's Height by Accelerometer Signals of a Smartphone. *HCI International* (Cham, 2016), 542–546.
- [58] Ferrari, G.L. de M. et al. 2017. Accelerometer-determined peak cadence and weight status in children from São Caetano do Sul, Brazil. *Ciência & Saúde Coletiva*. 22, (Nov. 2017), 3689–3698. DOI:https://doi.org/10.1590/1413-812320172211.21962015.
- [59] Warburton, D.E.R. et al. 2006. Health benefits of physical activity: the evidence. *Can. Med. Assoc. J.* 174, 6 (Mar. 2006), 801–809. DOI:https://doi.org/10.1503/cmaj.051351.
- [60] Zeitzer, J.M. et al. 2018. Daily Patterns of Accelerometer Activity Predict Changes in Sleep, Cognition, and Mortality in Older Men. *J. Gerontol.* 73, 5 (Apr. 2018), 682–687. DOI:https://doi.org/10.1093/gerona/glw250.
- [61] Chan, C.B. et al. 2012. Cross-sectional Relationship of Pedometer-Determined Ambulatory Activity to Indicators of Health. *Obesity Research*. 11, 12 (2012), 1563–1570. DOI:https://doi.org/10.1038/oby.2003.208.
- [62] Alvarez, G.G. and Ayas, N.T. 2007. The Impact of Daily Sleep Duration on Health: A Review of the Literature. *Progress in Cardiovascular Nursing*. 19, 2 (2007), 56–59. DOI:https://doi.org/10.1111/j.0889-7204.2004.02422.x.
- [63] Hees, V.T. van et al. 2015. A Novel, Open Access Method to Assess Sleep Duration Using a Wrist-Worn Accelerometer. *PLOS One*. 10, 11 (Nov. 2015). DOI:https://doi.org/10.1371/journal.pone.0142533.
- [64] Borghese, M. et al. 2018. Estimating sleep efficiency in 10-to-13-year-olds using a waist-worn accelerometer. *Sleep Health*. 4, 1 (Feb. 2018), 110–115. DOI:https://doi.org/10.1016/j.sleh.2017.09.006.
- [65] Lei, Z. et al. 2017. Supervised learning in voice type discrimination using neck-skin vibration signals: Preliminary results on single vowels. *J. Acoust. Soc. Am.* 141, 5 (May 2017), 3916–3916. DOI:https://doi.org/10.1121/1.4988844.
- [66] Liu, C. et al. 2015. A physiological sound sensing system using accelerometer based on flip-chip piezoelectric technology and asymmetrically gapped cantilever. *2015 IEEE 65th Electronic Components and Technology Conference (ECTC)* (May 2015), 1874–1877.
- [67] Smidt, G. et al. 1972. Accelerographic analysis of several types of walking. *Amer. J. Physical Med.* 50, (1972), 285–300.
- [68] Menz, H.B. et al. 2003. Age-related differences in walking stability. *Age and Ageing*. 32, 2 (Mar. 2003), 137–142. DOI:https://doi.org/10.1093/ageing/32.2.137.
- [69] Davarci, E. et al. 2017. Age group detection using smartphone motion sensors. *2017 25th European Signal Processing Conference* (Aug. 2017), 2201–2205.
- [70] Cho, S.H. et al. 2004. Gender differences in three dimensional gait analysis data from 98 healthy Korean adults. *Clin. Biomech.* 19, 2 (Feb. 2004), 145–152. DOI:https://doi.org/10.1016/j.clinbiomech.2003.10.003.
- [71] Jain, A. and Kanhangad, V. 2016. Investigating gender recognition in smartphones using accelerometer and gyroscope sensor readings. *International Conference on Computational Techniques in Information and Communication Technologies* (Mar. 2016), 597–602.
- [72] Jago, R. et al. 2005. Adolescent patterns of physical activity differences by gender, day, and time of day. *Amer. J. Physical Med.* 28, 5 (Jun. 2005), 447–452. DOI:https://doi.org/10.1016/j.amepre.2005.02.007.
- [73] MERRIFIELD, H.H. 1971. Female Gait Patterns in Shoes with Different Heel Heights. *Ergonomics*. 14, 3 (May 1971), 411–417. DOI:https://doi.org/10.1080/00140137108931260.
- [74] Kanning, M. and Schlicht, W. 2010. Be Active and Become Happy: An Ecological Momentary Assessment of Physical Activity and Mood. *Journal of Sport and Exercise Psychology*. 32, 2 (Apr. 2010), 253–261. DOI:https://doi.org/10.1123/jsep.32.2.253.
- [75] Gruenfelder-Steiger, A.E. et al. 2017. Physical Activity and Depressive Mood in the Daily Life of Older Adults. *The Journal of Gerontopsychology and Geriatric Psychiatry*. 30, 3 (Jan. 2017), 119–129. DOI:https://doi.org/10.1024/1662-9647/a000172.
- [76] Zhang, Z. et al. 2016. Emotion recognition based on customized smart bracelet with built-in accelerometer. *PeerJ*. 4, (Jul. 2016). DOI:https://doi.org/10.7717/peerj.2258.
- [77] Garcia-Ceja, E. et al. 2016. Automatic Stress Detection in Working Environments from Smartphones' Accelerometer Data: A First Step. *IEEE Journal of Biomedical and Health Informatics*. 20, 4 (Jul. 2016), 1053–1060. DOI:https://doi.org/10.1109/JBHI.2015.2446195.
- [78] Olsen, A.F. and Torresen, J. 2016. Smartphone accelerometer data used for detecting human emotions. *2016 3rd International Conference on Systems and Informatics (ICSAI)* (Nov. 2016), 410–415.
- [79] Wilson, K. and Dishman, R. 2014. Personality and Physical Activity: A Systematic Review and Meta-analysis. *Med. Sci. Sport Exer.* (May 2014), 473.
- [80] Artese, A. et al. 2017. Personality and Actigraphy-Measured Physical Activity in Older Adults. *Psychology and Aging*. 32, 2 (Mar. 2017), 131–138. DOI:https://doi.org/10.1037/pag0000158.
- [81] Wilson, K.E. et al. 2015. Personality Correlates of Physical Activity in College Women: *Medicine & Science in Sports & Exercise*. 47, 8 (Aug. 2015), 1691–1697. DOI:https://doi.org/10.1249/MSS.0000000000000570.
- [82] Foerster, F. et al. 1999. Detection of posture and motion by accelerometry: a validation study in ambulatory monitoring. *Computers in Human Behavior*. 15, 5 (Sep. 1999), 571–583. DOI:https://doi.org/10.1016/S0747-5632(99)00037-0.
- [83] Khan, A.M. et al. 2010. Accelerometer's position free human activity recognition using a hierarchical recognition model. *The 12th IEEE International Conference on e-Health Networking, Applications and Services* (Jul. 2010), 296–301.

What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking

Jacob Leon Kröger^{1,3}(✉), Otto Hans-Martin Lutz^{1,2,3}, and Florian Müller^{1,3}

¹ Technische Universität Berlin, Straße Des 17. Juni 135, 10623 Berlin, Germany
kroeger@tu-berlin.de

² Fraunhofer Institute for Open Communication Systems, Berlin, Germany

³ Weizenbaum Institute for the Networked Society, Berlin, Germany

Abstract. Technologies to measure gaze direction and pupil reactivity have become efficient, cheap, and compact and are finding increasing use in many fields, including gaming, marketing, driver safety, military, and healthcare. Besides offering numerous useful applications, the rapidly expanding technology raises serious privacy concerns. Through the lens of advanced data analytics, gaze patterns can reveal much more information than a user wishes and expects to give away. Drawing from a broad range of scientific disciplines, this paper provides a structured overview of personal data that can be inferred from recorded eye activities. Our analysis of the literature shows that eye tracking data may implicitly contain information about a user's biometric identity, gender, age, ethnicity, body weight, personality traits, drug consumption habits, emotional state, skills and abilities, fears, interests, and sexual preferences. Certain eye tracking measures may even reveal specific cognitive processes and can be used to diagnose various physical and mental health conditions. By portraying the richness and sensitivity of gaze data, this paper provides an important basis for consumer education, privacy impact assessments, and further research into the societal implications of eye tracking.

Keywords: Eye tracking · Gaze · Pupil · Iris · Vision · Privacy · Data mining · Inference

1 Introduction

Being an important part of visual perception and human behavior, eye movements have long been a subject of research interest. The first approaches to measure a person's gaze direction date back to the early 1900s [74]. Until recently, these technologies were severely limited by the cost of the equipment required, a lack of precision, and poor usability and were only used in very specific niches of research. Over the last few years, however, with rapid advances in sensor technology and data processing software, eye tracking solutions have become easy to use, lightweight, efficient, and affordable and found increasing adoption in many fields, including gaming, marketing, automotive technology, military, and healthcare [26].

While alternatives¹ exist, the most popular method today is video-based eye tracking, where mathematical models are used to calculate a person's gaze direction from video recordings, for example based on the shape and position of pupil and iris, or based on light reflection patterns in the eyes [2]. This method can not only be used in head-mounted devices, such as smart glasses and virtual reality headsets, but also through built-in cameras in laptops, tablets, and smartphones without requiring any additional hardware [45, 56]. With further improvements in cost and performance, eye tracking may soon be included as a standard feature in various consumer electronics, moving us towards a “pervasive eye tracking world” [58].

The many beneficial uses and enormous potentials of the rising technology have to be acknowledged and should be embraced. However, a more ubiquitous use of eye tracking will also raise serious privacy concerns – not only because gaze data may be collected and shared in non-transparent ways, but also because such data can unexpectedly contain a wealth of sensitive information about a user.

Drawing from a broad range of scientific disciplines, including neuroscience, human-computer interaction, medical informatics, affective computing, experimental economics, psychology, and cognitive science, this paper provides a structured overview and classification of sensitive pieces of information that can be disclosed by analyzing a person's eye activities. According to the reviewed literature, eye tracking data may reveal information about a user's biometric identity (Sect. 2.1), mental activities (Sect. 2.2), personality traits (Sect. 2.3), ethnic background (Sect. 2.4), skills and abilities (Sect. 2.5), age and gender (Sect. 2.6), personal preferences (Sect. 2.7), emotional state (Sect. 2.8), degree of sleepiness and intoxication (Sect. 2.8), and physical and mental health condition (Sect. 2.9). In order to take rapidly evolving technology trends and newly emerging privacy threats into account, we will consider not only proven and established approaches but also inference methods that are subject to ongoing research. Limitations of the presented methods and their practical applicability will be reflected upon in Sect. 3, followed by a conclusion in Sect. 4.

2 Inference of Personal Information from Eye Tracking Data

With reference to published research, filed patents, and existing commercial products, this section presents and categorizes personal information that can be inferred from eye tracking data. As a basis for potential inferences, eye tracking devices can record a large variety of gaze parameters.

Some of the most commonly measured eye movements are fixations, saccades, and smooth pursuit eye movements [85]. During a fixation, the eyes are relatively stable and focused on a specific position, allowing for information to be acquired and processed. Saccades are rapid eye movements from one fixation point to another, lasting 30 to 80 ms [87]. Smooth pursuit movements are performed when eyes are closely following a moving visual target. In addition to the spatial dispersion, duration, amplitude, acceleration, velocity, and chronological sequence of such eye movements, many eye trackers capture various other eye activities, including eye opening and closure (e.g., average distance

¹ For an overview of existing types of eye tracking, refer to [2].

between the eyelids, blink duration, blink frequency), ocular microtremors, pupil size, and pupil reactivity [19, 58]. Furthermore, most eye trackers videotape parts of the user’s face and may thereby capture additional information, such as the number and depth of wrinkles, and a user’s eye shape and iris texture [40]. Therefore, these parameters were also considered in our investigation into the richness and sensitivity of eye tracking data. Fig. 1 provides an introductory overview of common eye tracking measures and the categories of inferences discussed in this paper.

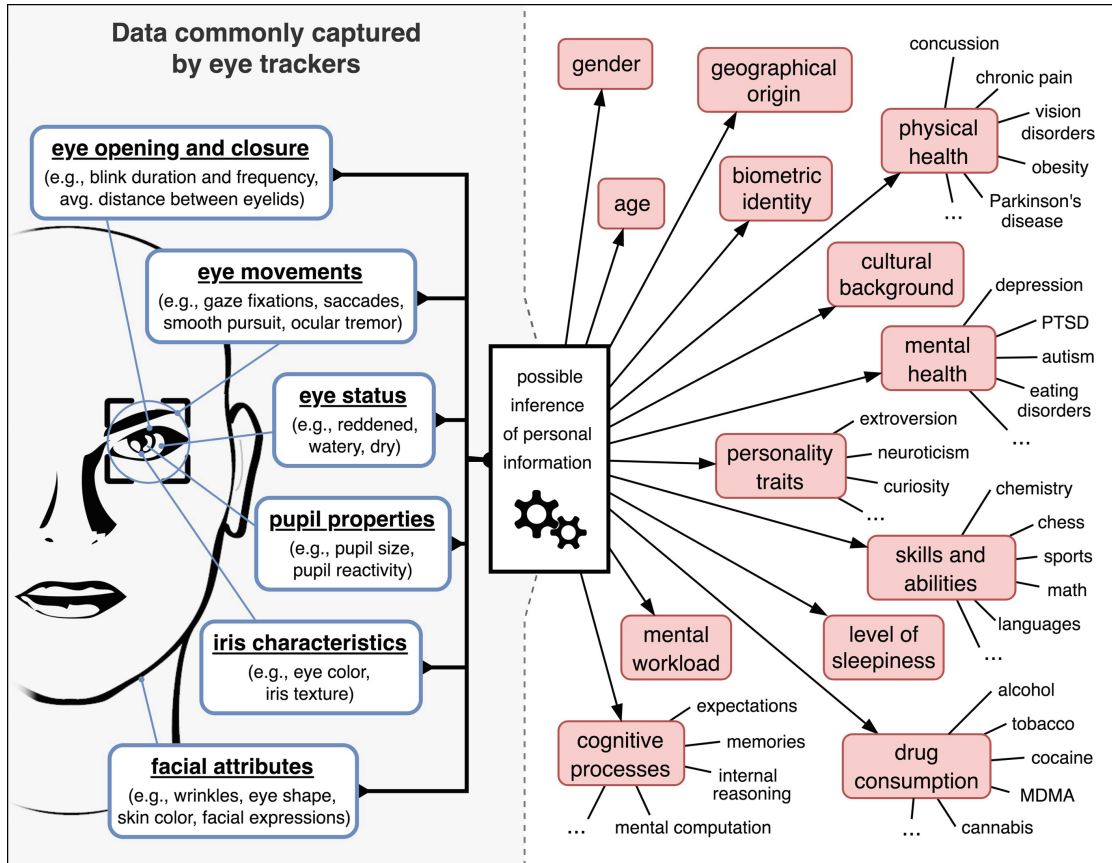


Fig. 1. Overview of sensitive inferences that can be drawn from eye tracking data.

2.1 Biometric Identification

Due to differences in physical oculomotor structure and brain functioning, certain gaze characteristics are unique for every individual, similar to fingerprints, and can thus be exploited for biometric identification [9, 74, 76]. Among other methods, people can be told apart based on distinct patterns of pupil reactivity and gaze velocity [9], or by comparing their eye movement trajectories when they focus on a moving target – even if the eye activity is only recorded through an ordinary smartphone camera [55].

Aside from such gaze-based measures, the complex textures and color patterns in a person’s iris are also suitable for biometric identification. This approach, called *iris*

recognition, is being used in a variety of real-world security and surveillance applications and has been recognized as “one of the most powerful techniques for biometric identification ever developed” [64]. Even though their iris scanning capability is usually not advertised, it should be understood that commodity eye trackers often record and process high-resolution images of the user’s iris, which can not only be used to uniquely identify the user but also to deceive iris-based authentication mechanisms and thereby steal the user’s identity [40].

In cases where a unique identification of an individual is not possible (e.g., because the person is not registered in the recognition system database), other attributes inferred from eye tracking data, such as age and gender (see Sect. 2.6), health condition (see Sect. 2.9), or ethnicity (see Sect. 2.4), can still help to classify the target person into a specific demographic group and thereby approximate the identity [74].

2.2 Monitoring of Mental Workload and Cognitive Processes

Certain patterns in eye movement, pupil dilation, and eye blinking have been recognized as reliable indicators of mental workload in people of any age [19, 63], sometimes offering higher accuracy than conventional methods like Electroencephalography [8]. Through eye tracking, it is also possible to distinguish a user’s moments of awareness from moments of distraction and mind wandering [31, 84].

Apart from detecting a user’s mental presence and measuring the mere intensity of cognitive processing, eye tracking can also provide insights into specific conscious and unconscious thought processes in a large variety of contexts. Among other mental tasks and activities, ocular measures have been used to study memory retrieval [19, 31], problem solving [31, 75], learning processes [44, 69], the formation of expectations [19, 27], internal reasoning [19], and mental computations [19, 31].

Eye tracking data can not only – to a certain extent – reveal what we remember, imagine, expect, and think about, but also our specific decision-making strategies [19, 28] and cognitive styles, i.e., individual differences in the way we acquire, process, and interpret information [72]. For example, people can be classified as field-dependent vs. field-independent (people of the latter type pay more attention to detail and exhibit a more analytical approach to processing visual information) [72], or as verbalizers vs. visualizers (people of the latter type can process visual information, such as images and diagrams, better than textual information) [44]. The gaze-based inference of such cognitive styles is feasible and can achieve high accuracies, as has been confirmed in a recent study by Raptis et al. [72].

Researchers from the field of cognitive science and experimental psychology have suggested that eye tracking data will not only be used for the real-time analysis but also for the prediction of human decisions and behavior [28].

2.3 Inference of Personality Traits

Experimental research has shown that it is possible to automatically infer personality traits from eye tracking data [34, 35, 42]. For example, gaze patterns captured during everyday tasks can be used to evaluate users along the so-called Big Five traits, namely openness to experience, conscientiousness, extroversion, agreeableness, and neuroticism

[34, 42]. The gaze-based assessment of personality traits is possible not only in binary form (high vs. low) but also in the form of ranges. In [35], for instance, eye movement analysis was used for the automatic recognition of different levels of curiosity.

Besides the Big Five traits and curiosity, gaze metrics were found to be associated with various other personality traits, including emotional intelligence [54], indecisiveness [36], the tendency to ruminate [21], trait anxiety [42], sexual compulsivity [87], boredom susceptibility [70], and general aggressiveness [6]. Eye tracking has even been used to investigate people’s attachment styles in interpersonal relationships (e.g., secure, withdrawn, fearful, enmeshed) [81].

Based on data from 428 study participants, Larsson et al. [53] also suggest that some personality traits, including tendermindedness, warmth, trust, and impulsiveness, are genetically linked to certain iris characteristics, offering – besides gaze behavior – another potential ocular biomarker to analyze people’s personalities.

2.4 Inference of Cultural Affiliation and Ethnicity

It is widely agreed that culture fundamentally shapes human cognitive processing and behavior [11]. Studies have shown that intercultural differences are reflected in certain gaze characteristics [12, 24, 41, 61]. For example, people of different cultural background were found to exhibit discriminative eye-movement patterns when seeking information on search engine results pages [61], when exploring complex visual scenes [12, 24], and when viewing videos of actors performing cultural activities [41]. Some cultural biases in visual processing are so pronounced that they can still be measured when external stimuli draw attention in an opposite manner to the respective bias [24].

Additionally, eye movements can reveal a person’s knowledge of certain cultural practices. For instance, in an eye tracking study by Green et al. [27], Chinese infants exclusively predicted the goal of eating actions performed by an actor with chopsticks, whereas European infants only anticipated that food would be brought to the mouth when eating actions were performed with Western cutlery, as indicated by their predictive gaze shifts towards the actor’s mouth.

Some studies have also investigated how people of different “race”² differ in their viewing behavior [25, 33, 88]. Apart from the fact that video-based eye trackers can directly record the eye color, eye shape, and skin color of a user, it has been observed in eye tracking studies that test subjects view “other-race faces” differently than faces of their “own race” in terms of the facial features scanned (e.g., initial focus and greater proportion of fixation time on the eyes vs. nose and mouth) [25, 88]. Furthermore, researchers have observed characteristic changes in pupil size, which are attributed to elevated cognitive effort during face recognition, when people look at “other-race faces” [88]. Such differences have been reported, for example, between “Black and White observers” [33] and between “Western Caucasian and East Asian observers” [25]

² The authors share the UNESCO’s position [60] that the classification of human populations into “races” is inadequate and obsolete. Nevertheless, it is important to monitor the state of research in this field, especially because any information indicative of a person’s ethnic background can serve as a basis for racist discrimination. All terms related to the concept of “race” in this paper are cited from external sources and do not reflect the authors’ views.

and could potentially allow inferences about the genetic and ethnic background of eye tracking users.

Eye tracking data may also allow inferences about a user's native language. For instance, considerable differences in eye movement patterns during reading can be observed between native and non-native speakers of English [39]. Eye tracking can even be used to determine which specific words are difficult to understand for a person [51]. Among other things, such information could help in estimating a subject's nationality or geographical origin.

2.5 Skill Assessment

Eye tracking has been used extensively in the study of human expertise and to discriminate between performance levels in a variety of areas [30, 31, 69, 75]. For example, gaze behavior can be analyzed to assess reading and listening comprehension skills [10, 92]. During a corresponding task or scenario, eye tracking can also be used to distinguish between experts and novices in chess [75], several sports [46], chemistry [69], mathematics [31], school teaching [14], and various medical skills, including surgery, nursing, anesthesia, and radiology [30].

Among other gaze characteristics, expertise is often associated with systematic eye movement patterns reflecting a specific task strategy [31], with the targeted inspection of important regions and task-relevant information [30, 75], and with more consistent gaze patterns over consecutive trials of a task [46].

In some fields, eye tracking has not only been used as a tool to discriminate between people of different skill levels, but also to predict people's task performance and learning curves [52, 69] and to examine specific learning disabilities, such as mathematical difficulties and dyslexia [31, 85].

2.6 Age and Gender Recognition

Just like physical shape, skin texture, and cognitive abilities, human eyes and visual behavior are fundamentally affected by the aging process [20, 36]. For example, eye tracking studies found age-related differences in people's visual explorativeness, pupil reactions to certain visual stimuli, and error rates in eye movement tasks [36, 42].

Furthermore, detailed frontal face images, which are typically required for video-based eye tracking, have already been used for automated age estimation, for instance based on wrinkles in the eye area [15]. Dynamic facial expressions, such as smiles, may also be analyzed to infer the age of test subjects [17]. Other parameters utilized for computerized age-group recognition include iris size and iris texture [20].

As with age, a person's gender can be reflected in certain eye tracking measures. For instance, studies found systematic gender differences in people's fixation distribution while viewing natural images (e.g., stills from romance films or wildlife documentaries) [68], during online shopping [38], when playing video games [42], and when viewing sexual stimuli [87]. Researchers have already used such differences in visual behavior to automatically classify the sex of test subjects [68].

2.7 Inference of Preferences and Aversions

Eye tracking is widely employed to investigate people's interests, likes, and dislikes. Spontaneous attention to specific objects in a visual scene (e.g., in terms of frequency, duration, and sequence of gaze fixations) is regarded as a natural indicator of interest [19, 74, 87]. For data presentation and analysis, gaze fixations are commonly aggregated into heat maps to quickly identify potential regions and objects of interest [74].

Besides the focus of visual attention, other eye parameters, such as pupil dilation and blink properties, can also be used to analyze a person's degree of interest and to distinguish between positive, neutral, and negative responses to visual stimuli [55]. Emotion detection from gaze data, which can assist in analyzing a user's interests and preferences [55, 83], will be discussed in Sect. 2.8.

Among other things, eye tracking has been used to examine preferences for certain types of gambling [65], mobile apps [56], activities of daily living [86], types of food [32], colors, geometric shapes, and product designs [3], pieces of clothing, animals, video game characters, and items of furniture [83]. Beyond mere interest, existing research even suggests that people's patterns of visual attention reflect their consumption and purchasing behavior [91].

Eye tracking has also been used extensively in the study of love and sexual desire. For example, researchers have analyzed pupillary responses and the allocation of visual attention to measure levels of sexual arousal and to investigate mating preferences towards specific facial characteristics, age groups, body shapes, body parts, and signs of social dominance [3, 87].

Apart from positive interests, visual attentional biases captured by eye trackers can also reflect a person's phobias and aversions (e.g., fear of spiders) [3, 37]. Some interests and preferences can already be inferred from eye tracking data with high accuracy [56, 73, 87] and several patents have been filed in this field [3, 83].

2.8 Detection of Short- and Medium-Term User States

Moods and Emotions. Eye tracking is increasingly used in the interdisciplinary field of affective computing, where systems are developed to automatically recognize human emotions based on physiological signals and behavioral cues [73, 83]. It has been shown that various ocular measures, including pupil size, blink properties, saccadic eye movements, and specific biases of visual attention, can contain information about a person's emotional state [4, 23, 55, 59].

Gaze data can reflect emotional arousal and the valence of emotions (positive, negative, neutral) [19, 55] as well as more specific affective states, such as happiness and enthusiasm [83], acute stress and worry [59], humorous moods and disgust [73], curiosity [4], distress, nervousness, and hostility [23], fear, anger, sadness, and surprise [55].

Eye tracking can not only be used to detect emotions with high accuracy [73] but also to estimate the intensity of emotions [55, 83]. Based on gaze parameters, existing methods can even distinguish whether a user's emotional response to a given stimulus is rational or purely instinctive [55].

Fatigue and Sleepiness. For over two decades, there have been approaches to automatically derive a person's level of sleepiness from certain ocular measures, such as blink rate, blink duration, average distance between the eyelids, fixation durations, and velocity of eye movements [57]. Recent studies have confirmed the suitability of eye tracking measures as indicators for sleepiness and fatigue [63, 89]. Corresponding methods have already been patented and achieve high accuracies – not only while the user is working on specific cognitive tasks, but also during everyday natural-viewing situations [57, 89].

Intoxication. The consumption of alcohol and other recreational drugs can have measurable effects on various eye and gaze properties, such as decreased accuracy and speed of saccades, changes in pupil size and reactivity, and an impaired ability to fixate on moving objects [29, 67, 85].

Apart from alcohol, significant abnormalities in oculomotor functioning were found in people under the influence of nicotine, 3,4-methylenedioxymethamphetamine (“MDMA”), and tetrahydrocannabinol (“THC”) [29, 70].

Researchers have demonstrated the ability to differentiate between drug-impaired and sober subjects with high accuracy based on eye tracking data [29]. The magnitude of some ocular effects is closely associated with the amount of drugs consumed [85] and certain effects can even be detected at non-intoxicating doses [77]. In addition to pupillary changes and eye movement impairments, an attentional bias towards drug-related visual stimuli has been observed among intoxicated test subjects [67].

Not only a state of intoxication, but also an acute state of drug deprivation and craving can have a distinct effect on certain eye tracking parameters [29, 70].

2.9 Health Assessment

Physical Health. Many diseases and medical conditions directly affect the eyes, or parts of the brain that are responsible for oculomotor function, and thereby cause gaze impairments [3, 19, 30]. Characteristic eye movement patterns were found, for example, in people suffering from concussion [43], fetal alcohol syndrome [3], irregular growth [3], chronic pain [22], neurocognitive impairment due to preterm birth [82], multiple sclerosis [3], Alzheimer's disease [30, 43], Tourette syndrome [19], Parkinson's disease [30], and various vision disorders (e.g., myopia, farsightedness, and blind spots) [3, 43].

As filed patents and published experimental studies show, eye movement analysis can be used to diagnose, monitor, prognose, and sometimes even predict various health disorders [30, 43] which can be subsumed under the umbrella term ETDCC (“Eye Tracking-Relevant Diseases, Conditions, and Characteristics”) [3].

Research has further demonstrated that certain patterns in gaze orientation and pupil reactivity to food-related stimuli (e.g., high vs. low calorie food images) can be indicative of overweight and obesity [32].

Mental Health. Abnormal eye movements can be used as behavioral biomarkers for the diagnosis of various mental health problems [1, 5, 29]. Oculomotor dysfunctions and gaze peculiarities are found, for example, in sufferers of anxiety disorder [29], depression [1], bipolar disorder [30], borderline personality disorder [6], schizophrenia [5], obsessive-compulsive disorder [13], binge-eating disorder [79], ADHD [7], mild cognitive impairment [30], autism [43], and posttraumatic stress disorder [66].

Some common symptoms of mental disorders are irregularities in blink rate and blink duration [19], abnormal stability and dispersal of gaze fixations during free viewing [5], unusual biases of visual attention [66], impaired smooth pursuit eye-movement performance [85], eye contact avoidance, and abnormal distance between the eyelids [1].

Certain mental illnesses, including depression and schizophrenia, can already be detected automatically via eye tracking [1, 5, 30] and corresponding methods have been filed as patents [43]. Besides the possibility of binary classification (suffering vs. not suffering), some ocular measures are associated with the severity of mental disorders [19]. Not only acute disorders can be reflected in gaze data, but also past mental health issues and even the personal risk of future outbreaks [71, 78]. For example, researchers have observed characteristic gaze patterns in previously depressed individuals [78] and found biases in visual attention that were predictive of future depression scores at a delay of more than two years [71].

Substance Use Disorders. Apart from acute states of intoxication (which we have discussed in Sect. 2.8), eye tracking data may contain information about a user’s longer-term drug consumption habits and addictions. Numerous eye tracking studies have reported a strong attentional bias towards drug-related visual cues in addicts of cocaine [16], alcohol [67], cannabis [90], and tobacco [18, 70].

Among other possible methods, such attentional biases can be detected by measuring how quickly, how often, and for how long a person’s eyes fixate on corresponding stimuli in comparison to neutral stimuli, or by testing the person’s ability to look away from drug-related stimuli on command [16, 18]. Significant biases have not only been observed in long-term addicts but also in habitual drug users without clinical symptoms of dependency [18, 67]. The strength of attentional biases towards drug-related visual cues was found to be correlated with scores on drug use scales, such as the Obsessive Compulsive Cocaine Scale [16] and with self-reported lifetime drug consumption [62]. Research has also shown that certain biases in visual attention can be predictive of craving and even relapse in drug addiction [16].

3 Discussion and Implications

As shown in the previous section, various kinds of sensitive inferences can be drawn from eye tracking data. Among other categories of personal data, recorded visual behavior can implicitly contain information about a person’s biometric identity, personality traits, ethnic background, age, gender, emotions, fears, preferences, skills and abilities, drug habits, levels of sleepiness and intoxication, and physical and mental health condition. To some extent, even distinct stages of cognitive information processing are discernable from gaze data. Thus, devices with eye tracking capability have the potential to implicitly capture much more information than a user wishes and expects to reveal. Some of the categories of personal information listed above constitute *special category data*, for which particular protection is prescribed by the EU’s General Data Protection Regulation (Art. 9 GDPR).

Of course, drawing reliable inferences from eye tracking data is not a trivial task. Many situational factors can influence eye properties and gaze behavior in complex

ways, making it difficult to measure the effect of a particular action, internal process, or personal characteristic of the user in isolation [55]. Seemingly identical ocular reactions can result from completely different causes. For example, an intensive gaze fixation on another person's face may indicate liking, aversion, confusion, recognition, and much more. Similarly, a sudden change in pupil size can be indicative of many different feelings or internal states, including physical pain, sexual arousal, interest, happiness, anger, or simply be a reaction to ambient events and conditions, such as noise or varying lighting [19, 55].

In spite of existing challenges and limitations, the reviewed literature demonstrates that there is considerable potential for inferences in many areas and that numerous research projects, patented systems, and even commercial products have already taken advantage of the richness of eye tracking data to draw inferences about individuals with high accuracy.

It should be acknowledged that many of the cited inference methods were only tested under controlled laboratory conditions and lack evaluation in real-world scenarios [4, 18, 27, 52, 65, 67, 69, 86, 88]. On the other hand, it may reasonably be assumed that some of the companies with access to eye tracking data from consumer devices (e.g., device manufacturers, ecosystem providers) possess larger sets of training data, more technical expertise, and more financial resources than the researchers cited in this paper. Facebook, for example, a pioneer in virtual reality and eye tracking technology, is also one of the wealthiest and most profitable companies in the world with a multi-billion dollar budget for research and development and a user base of over 2.3 billion people [93]. It seems probable that the threat of unintended information disclosure from gaze data will continue to grow with further improvements of eye tracking technology in terms of cost, size, and accuracy, further advances in analytical approaches, and the increasing use of eye tracking in various aspects of daily life.

In assessing the privacy implications of eye tracking, it is important to understand that, while consciously directed eye movements are possible, many aspects of ocular behavior are not under volitional control – especially not at the micro level [19, 55]. For instance, stimulus-driven glances, pupil dilation, ocular tremor, and spontaneous blinks mostly occur without conscious effort, similar to digestion and breathing. And even for those eye activities where volitional control is possible, maintaining it can quickly become physically and cognitively tiring [58] – and may also produce certain visible patterns by which such efforts can be detected. Hence, it can be very difficult or even impossible for eye tracking users to consciously prevent the leakage of personal information.

Though this paper focuses on privacy risks, we do not dispute the wide-ranging benefits of eye tracking. Quite the opposite: we believe that it is precisely the richness of gaze data and the possibility to draw insightful inferences from it that make the rising technology so valuable and useful. But to exploit this potential in a sustainable and socially acceptable manner, adequate privacy protection measures are needed.

Technical safeguards have been proposed to prevent the unintended disclosure of personal information in data mining, including specialized solutions for eye tracking data [58, 80]. These comprise the fuzzing of gaze data (i.e., inserting random noise into the

signal before passing it down the application chain) and the utilization of derived parameters (e.g., aggregated values instead of detailed eye fixation sequences) [58]. Experiments have already shown that approaches based on differential privacy can prevent certain inferences, such as user re-identification and gender recognition, while maintaining high performance in gaze-based applications [80]. In addition to approaches at the technical level, it should also be examined whether existing laws provide for sufficient transparency in the processing of gaze data and for proper protection against inference-based privacy breaches. The promises and limitations of existing technical and legal remedies are beyond the scope of this paper but deserve careful scrutiny and will be considered for future work.

Even though eye tracking is a demonstrative example, the threat of undesired inferences is of course much broader, encompassing countless other sensors and data sources in modern life [47]. In other recent work, we have examined sensitive inferences that can be drawn from voice recordings [49] and accelerometer data [48, 50], for instance. In our view, the vast possibilities of continuously advancing inference methods are clearly beyond the understanding of the ordinary consumer. Therefore, we consider it to be primarily the responsibility of technical experts, technology companies, and governmental agencies to inform consumers about potential consequences and protect them against such covert invasions of privacy. Also, since it is unlikely that companies will voluntarily refrain from using or selling personal information that can be extracted from already collected data, there should be strong regulatory incentives and controls.

4 Conclusion

While the widespread adoption of eye tracking holds the potential to improve our lives in many ways, the rising technology also poses a substantial threat to privacy. The overview provided in this paper illustrates that, through the lens of advanced data analytics, eye tracking data can contain a rich array of sensitive information, including cues to a user's biometric identity, gender, age, ethnicity, personality traits, drug consumption habits, moods and emotions, skills, preferences, cognitive processes, and physical and mental health condition. Since inference methods are often based on hidden patterns and correlations that are incomprehensible to ordinary consumers, it can be impossible for them to understand and control what information is revealed.

Although there is extensive literature on the analysis of eye tracking data, we believe that many possible inferences have not yet been investigated. Keeping track of the evolving possibilities of data mining methods in this field is therefore an important avenue for future research. This paper represents a crucial first step towards understanding the sensitivity of eye tracking data from a holistic perspective. The findings compiled herein are significant enough to warrant a warning to users whose privacy could be affected, as well as a call for action to the public and private actors entrusted with protecting user privacy in consumer electronics. Considering the rapid proliferation of eye tracking technology, existing technical and legal safeguards urgently need to be assessed regarding their ability to avert undesired inferences from gaze data, or to at least prevent the misuse of sensitive inferred information.

References

1. Alghowinem, S., et al.: Eye movement analysis for depression detection. In: IEEE International Conference on Image Processing, pp. 4220–4224 (2013)
2. Al-Rahayfeh, A., Faezipour, M.: Eye tracking and head movement detection: a state-of-art survey. *IEEE J. Transl. Eng. Health Med.* **1**, 2100212 (2013)
3. Avital, O.: Method and System of Using Eye Tracking to Evaluate Subjects (Patent No.: US20150282705A1) (2015)
4. Baranes, A., et al.: Eye movements reveal epistemic curiosity in human observers. *Vis. Res.* **117**, 81–90 (2015). <https://doi.org/10.1016/j.visres.2015.10.009>
5. Benson, P.J., et al.: Simple viewing tests can detect eye movement abnormalities that distinguish schizophrenia cases from controls with exceptional accuracy. *Biol. Psychiatry* **72**(9), 716–724 (2012). <https://doi.org/10.1016/j.biopsych.2012.04.019>
6. Bertsch, K., et al.: Interpersonal threat sensitivity in borderline personality disorder: an eye-tracking study. *J. Pers. Disord.* **31**(5), 647–670 (2017)
7. Blazey, R.N., et al.: ADHD Detection by Eye Saccades (Patent No.: US6652458B2) (2003)
8. Borys, M., et al.: An analysis of eye-tracking and electroencephalography data for cognitive load measurement during arithmetic tasks. In: 10th International Symposium on Advanced Topics in Electrical Engineering (ATEE), pp. 287–292 (2017)
9. Cantoni, V., et al.: Gaze-based biometrics: an introduction to forensic applications. *Pattern Recogn. Lett.* **113**, 54–57 (2018). <https://doi.org/10.1016/j.patrec.2016.12.006>
10. Chita-Tegmark, M., et al.: Eye-tracking measurements of language processing: developmental differences in children at high risk for ASD. *J. Autism Dev. Disord.* **45**(10), 3327–3338 (2015). <https://doi.org/10.1007/s10803-015-2495-5>
11. Chizari, S.: Exploring the role of culture in online searching behavior from cultural cognitive perspective: case study of American, Chinese and Iranian Graduate Students. In: iConference Proceedings. iSchools, Philadelphia (2016)
12. Chua, H.F., et al.: Cultural variation in eye movements during scene perception. *Proc. Natl. Acad. Sci.* **102**(35), 12629–12633 (2005). <https://doi.org/10.1073/pnas.0506162102>
13. Cludius, B., et al.: Attentional biases of vigilance and maintenance in obsessive-compulsive disorder: an eye-tracking study. *J. Obsessive Compuls. Relat. Disord.* **20**, 30–38 (2019). <https://doi.org/10.1016/j.jocrd.2017.12.007>
14. Cortina, K.S., et al.: Where low and high inference data converge: validation of CLASS assessment of mathematics instruction using mobile eye tracking with expert and novice teachers. *Int. J. Sci. Math. Educ.* **13**(2), 389–403 (2015)
15. Dehshibi, M.M., Bastanfard, A.: A new algorithm for age recognition from facial images. *Signal Process.* **90**(8), 2431–2444 (2010)
16. Dias, N.R., et al.: Anti-saccade error rates as a measure of attentional bias in cocaine dependent subjects. *Behav. Brain Res.* **292**, 493–499 (2015)
17. Dibeklioglu, H., et al.: A smile can reveal your age: enabling facial dynamics in age estimation. In: Proceedings of the 20th ACM International Conference on Multimedia, pp. 209–218. ACM Press, Nara (2012). <https://doi.org/10.1145/2393347.2393382>
18. DiGirolamo, G.J., et al.: Breakdowns of eye movement control toward smoking cues in young adult light smokers. *Addict. Behav.* **52**, 98–102 (2016)
19. Eckstein, M.K., et al.: Beyond eye gaze: what else can eyetracking reveal about cognition and cognitive development? *Dev. Cogn. Neurosci.* **25**, 69–91 (2017)
20. Erbilek, M., et al.: Age prediction from iris biometrics. In: 5th International Conference on Imaging for Crime Detection and Prevention (ICDP), pp. 1–5 (2013)
21. Fang, L., et al.: Attentional scope, rumination, and processing of emotional information: an eye-tracking study. *Emotion* **19**(7), 1259–1267 (2018)

22. Fashler, S.R., Katz, J.: Keeping an eye on pain: investigating visual attention biases in individuals with chronic pain using eye-tracking methodology. *J. Pain Res.* **9**, 551–561 (2016). <https://doi.org/10.2147/JPR.S104268>
23. Gere, A., et al.: Influence of mood on gazing behavior: preliminary evidences from an eye-tracking study. *Food Qual. Prefer.* **61**, 1–5 (2017)
24. Goh, J.O., et al.: Culture modulates eye-movements to visual novelty. *PLoS ONE* **4**(12), e8238 (2009). <https://doi.org/10.1371/journal.pone.0008238>
25. Goldinger, S.D., et al.: Deficits in cross-race face learning: insights from eye movements and pupillometry. *J. Exp. Psychol. Learn. Mem. Cogn.* **35**(5), 1105–1122 (2009)
26. Grand View Research: Global Eye Tracking Market Size By Type, Industry report. <https://www.grandviewresearch.com/industry-analysis/eye-tracking-market>. Accessed 25 Oct 2019
27. Green, D., et al.: Culture influences action understanding in infancy: prediction of actions performed with chopsticks and spoons in Chinese and Swedish infants. *Child Dev.* **87**(3), 736–746 (2016)
28. Guazzini, A., et al.: Cognitive dissonance and social influence effects on preference judgments: an eye tracking based system for their automatic assessment. *Int. J. Hum Comput Stud.* **73**, 12–18 (2015). <https://doi.org/10.1016/j.ijhcs.2014.08.003>
29. Hall, C.A., Chilcott, R.P.: Eyeing up the Future of the Pupillary Light Reflex in Neurodiagnostics. *Diagnostics* **8**(1), 1–20 (2018). <https://doi.org/10.3390/diagnostics8010019>
30. Harezlak, K., Kasproski, P.: Application of eye tracking in medicine: a survey, research issues and challenges. *Comput. Med. Imag. Graph.* **65**, 176–190 (2018)
31. Hartmann, M., Fischer, M.H.: Exploring the numerical mind by eye-tracking: a special issue. *Psychol. Res.* **80**(3), 325–333 (2016). <https://doi.org/10.1007/s00426-016-0759-0>
32. Hendrikse, J.J., et al.: Attentional biases for food cues in overweight and individuals with obesity: a systematic review of the literature. *Obes. Rev.* **16**(5), 424–432 (2015)
33. Hills, P.J., Pake, J.M.: Eye-tracking the own-race bias in face recognition: revealing the perceptual and socio-cognitive mechanisms. *Cognition* **129**(3), 586–597 (2013)
34. Hoppe, S., et al.: Eye movements during everyday behavior predict personality traits. *Front. Hum. Neurosci.* **12**, 1–8 (2018). <https://doi.org/10.3389/fnhum.2018.00105>
35. Hoppe, S., et al.: Recognition of curiosity using eye movement analysis. In: *International Conference on Pervasive and Ubiquitous Computing*, pp. 185–188 (2015)
36. Horsley, M. (ed.): *Current Trends in Eye Tracking Research*. Springer, Cham (2013). <https://doi.org/10.1007/978-3-319-02868-2>
37. Huijding, J., et al.: To look or not to look: an eye movement study of hypervigilance during change detection in high and low spider fearful students. *Emotion* **11**(3), 666–674 (2011). <https://doi.org/10.1037/a0022996>
38. Hwang, Y.M., Lee, K.C.: Using an eye-tracking approach to explore gender differences in visual attention and shopping attitudes in an online shopping environment. *Int. J. Hum. Comput. Interact.* **34**(1), 15–24 (2018)
39. Ito, A., et al.: Investigating the time-course of phonological prediction in native and non-native speakers of English: a visual world eye-tracking study. *J. Mem. Lang.* **98**, 1–11 (2018). <https://doi.org/10.1016/j.jml.2017.09.002>
40. John, B., et al.: EyeVEIL: degrading iris authentication in eye tracking headsets. In: *ACM Symposium on Eye Tracking Research & Applications (ETRA)*, pp. 1–5. ACM Press, Denver (2019). <https://doi.org/10.1145/3314111.3319816>
41. Kardan, O., et al.: Cultural and developmental influences on overt visual attention to videos. *Sci. Rep.* **7**(1), 11264 (2017). <https://doi.org/10.1038/s41598-017-11570-w>
42. Kaspar, K., König, P.: Emotions and personality traits as high-level factors in visual attention: a review. *Front. Hum. Neurosci.* **6**, 321 (2012)
43. Kempinski, Y.: *System and Method of Diagnosis Using Gaze and Eye Tracking* (Patent No.: US20160106315A1) (2016)

44. Koć-Januchta, M., et al.: Visualizers versus verbalizers: effects of cognitive style on learning with texts and pictures – an eye-tracking study. *Comput. Hum. Behav.* **68**, 170–179 (2017). <https://doi.org/10.1016/j.chb.2016.11.028>
45. Krafka, K., et al.: Eye tracking for everyone. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2176–2184. IEEE, Las Vegas (2016)
46. Kredel, R., et al.: Eye-tracking technology and the dynamics of natural gaze behavior in sports: a systematic review of 40 years of research. *Front. Psychol.* **8**, 1–15 (2017)
47. Kröger, J.: Unexpected inferences from sensor data: a hidden privacy threat in the Internet of Things. In: Strous, L., Cerf, V.G. (eds.) IFIP IoT 2018. IAICT, vol. 548, pp. 147–159. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-15651-0_13
48. Kröger, J.L., et al.: Privacy implications of accelerometer data: a review of possible inferences. In: Proceedings of the 3rd International Conference on Cryptography, Security and Privacy (ICCSP). ACM, New York (2019). <https://doi.org/10.1145/3309074.3309076>
49. Kröger, J.L., et al.: Privacy implications of voice and speech analysis - information disclosure by inference. In: Fricker, S., et al. (eds.) Privacy and Identity 2019. IFIP AICT, vol. 576, pp. 242–258. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-42504-3_16
50. Kröger, J.L., Raschke, P.: Is my phone listening in? On the feasibility and detectability of mobile eavesdropping. In: Foley, S.N. (ed.) DBSec 2019. LNCS, vol. 11559, pp. 102–120. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22479-0_6
51. Kunze, K., et al.: Towards inferring language expertise using eye tracking. In: CHI 2013 Extended Abstracts on Human Factors in Computing Systems, pp. 217–222. ACM Press, Paris (2013). <https://doi.org/10.1145/2468356.2468396>
52. Lallé, S., et al.: Prediction of users' learning curves for adaptation while using an information visualization. In: International Conference on Intelligent User Interfaces, pp. 357–368. ACM Press, Atlanta (2015)
53. Larsson, M., et al.: Associations between iris characteristics and personality in adulthood. *Biol. Psychol.* **75**(2), 165–175 (2007). <https://doi.org/10.1016/j.biopsycho.2007.01.007>
54. Lea, R.G., et al.: Trait emotional intelligence and attentional bias for positive emotion: an eye tracking study. *Pers. Individ. Differ.* **128**, 88–93 (2018)
55. Lemos, J.: System and Method for Determining Human Emotion by Analyzing Eye Properties (Patent No.: US20070066916A1) (2007)
56. Li, Y., et al.: Towards measuring and inferring user interest from gaze. In: International Conference on World Wide Web Companion, pp. 525–533. ACM Press, Perth (2017). <https://doi.org/10.1145/3041021.3054182>
57. Liang, C.-C., et al.: System for Monitoring Eyes for Detecting Sleep Behavior (Patent No.: US5570698A) (1996)
58. Liebling, D.J., Preibusch, S.: Privacy considerations for a pervasive eye tracking world. In: International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication, pp. 1169–1177. ACM Press, New York (2014)
59. Macatee, R.J., et al.: Attention bias towards negative emotional information and its relationship with daily worry in the context of acute stress: an eye-tracking study. *Behav. Res. Ther.* **90**, 96–110 (2017). <https://doi.org/10.1016/j.brat.2016.12.013>
60. Mader, G.: Declaration of Schlaining Against Racism, Violence and Discrimination. Austrian Commission for UNESCO, Vienna (1995)
61. Marcos, M.-C., et al.: Cultural differences on seeking information: an eye tracking study. In: CHI 2013: Workshop Many People, Many Eyes. ACM, Paris (2013)
62. Marks, K.R., et al.: Fixation time is a sensitive measure of cocaine cue attentional bias. *Addict. Abingdon Engl.* **109**(9), 1501–1508 (2014). <https://doi.org/10.1111/add.12635>
63. Martins, R., Carvalho, J.: Eye blinking as an indicator of fatigue and mental load—a systematic review. In: Arezes, P., et al. (eds.) Occupational Safety and Hygiene III, pp. 231–235. CRC Press (2015). <https://doi.org/10.1201/b18042-48>

64. Matey, J.R., et al.: Iris on the move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**(11), 1936–1947 (2006)
65. McGrath, D.S., et al.: The specificity of attentional biases by type of gambling: an eye-tracking study. *PLoS ONE* **13**(1), e0190614 (2018)
66. Milanak, M.E., et al.: PTSD symptoms and overt attention to contextualized emotional faces: evidence from eye tracking. *Psychiatry Res.* **269**, 408–413 (2018)
67. Miller, M.A., Fillmore, M.T.: Persistence of attentional bias toward alcohol-related stimuli in intoxicated social drinkers. *Drug Alcohol Depend.* **117**(2), 184–189 (2011)
68. Moss, F.J.M., et al.: Eye movements to natural images as a function of sex and personality. *PLoS ONE* **7**(11), e47870 (2012). <https://doi.org/10.1371/journal.pone.0047870>
69. Peterson, J., Pardos, Z., Rau, M., Swigart, A., Gerber, Colin, McKinsey, J.: Understanding student success in chemistry using gaze tracking and pupillometry. In: Conati, C., Heffernan, N., Mitrovic, A., Verdejo, M.F. (eds.) *AIED 2015. LNCS (LNAI)*, vol. 9112, pp. 358–366. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-19773-9_36
70. Pettiford, J., et al.: Increases in impulsivity following smoking abstinence are related to base-line nicotine intake and boredom susceptibility. *Addict. Behav.* **32**(10), 2351–2357 (2007). <https://doi.org/10.1016/j.addbeh.2007.02.004>
71. Price, R.B., et al.: From anxious youth to depressed adolescents: prospective prediction of 2-year depression symptoms via attentional bias measures. *J. Abnorm. Psychol.* **125**(2), 267–278 (2016). <https://doi.org/10.1037/abn0000127>
72. Raptis, G.E., et al.: Using eye gaze data and visual activities to infer human cognitive styles: method and feasibility studies. In: *Conference on User Modeling, Adaptation and Personalization (UMAP)*, pp. 164–173. ACM Press, Bratislava (2017)
73. Raudonis, V., et al.: Evaluation of human emotion from eye motions. *Int. J. Adv. Comput. Sci. Appl.* **4**(8), 79–84 (2013). <https://doi.org/10.14569/IJACSA.2013.040812>
74. Ravi, B.: *Privacy Issues in Virtual Reality: Eye Tracking Technology*. Bloomberg Law, Arlington County (2017)
75. Reingold, E., Sheridan, H.: Eye movements and visual expertise in chess and medicine. In: *Liversedge, S.P., Gilchrist, I.D., Everling, S. (eds.) The Oxford Handbook of Eye Movements*, pp. 528–550. Oxford University, Oxford (2011)
76. Rigas, I., et al.: Biometric recognition via eye movements: saccadic vigor and acceleration cues. *ACM Trans. Appl. Percept.* **13**(2), 1–21 (2016)
77. Roche, D.J.O., King, A.C.: Alcohol impairment of saccadic and smooth pursuit eye movements: impact of risk factors for alcohol dependence. *Psychopharmacology* **212**(1), 33–44 (2010). <https://doi.org/10.1007/s00213-010-1906-8>
78. Sears, C.R., et al.: Attention to emotional images in previously depressed individuals: an eye-tracking study. *Cogn. Ther. Res.* **35**(6), 517–528 (2011)
79. Sperling, I., et al.: Cognitive food processing in binge-eating disorder: an eye-tracking study. *Nutrients* **9**(8), 903 (2017). <https://doi.org/10.3390/nu9080903>
80. Steil, J., et al.: Privacy-aware eye tracking using differential privacy. In: *ACM Symposium on Eye Tracking Research & Applications*, pp. 1–9 (2019). <https://doi.org/10.1145/3314111.3319915>
81. Szymanska, M., et al.: How do adolescents regulate distress according to attachment style? A combined eye-tracking and neurophysiological approach. *Prog. Neuropsychopharmacol. Biol. Psychiatry* **89**, 39–47 (2019). <https://doi.org/10.1016/j.pnpbp.2018.08.019>
82. Telford, E.J., et al.: Preterm birth is associated with atypical social orienting in infancy detected using eye tracking. *J. Child Psychol. Psychiatry* **57**(7), 861–868 (2016)
83. Thieberger, G., et al.: *Utilizing Eye-tracking to Estimate Affective Response to a Token Instance of Interest* (Patent No.: US9569734B2) (2017)
84. Tobii: Tobii Pro wearable eye tracking for driver safety. <https://www.tobii.com/fields-of-use/psychology-and-neuroscience/customer-cases/audi-attitudes/>. Accessed 13 Sept 2019

85. Vidal, M., et al.: Wearable eye tracking for mental health monitoring. *Comput. Commun.* **35**(11), 1306–1311 (2012). <https://doi.org/10.1016/j.comcom.2011.11.002>
86. Wang, C.-Y., et al.: Multimedia recipe reading: predicting learning outcomes and diagnosing cooking interest using eye-tracking measures. *Comput. Hum. Behav.* **62**, 9–18 (2016)
87. Wenzlaff, F., et al.: Video-based eye tracking in sex research: a systematic literature review. *J. Sex Res.* **53**(8), 1008–1019 (2016)
88. Wu, E.X.W., et al.: Through the eyes of the own-race bias: eye-tracking and pupillometry during face recognition. *Soc. Neurosci.* **7**(2), 202–216 (2012)
89. Yamada, Y., Kobayashi, M.: Fatigue detection model for older adults using eye-tracking data gathered while watching video: evaluation against diverse fatiguing tasks. In: 2017 IEEE International Conference on Healthcare Informatics (ICHI), pp. 275–284 (2017). <https://doi.org/10.1109/ICHI.2017.74>
90. Yoon, J.H., et al.: Assessing attentional bias and inhibitory control in cannabis use disorder using an eye-tracking paradigm with personalized stimuli. *Exp. Clin. Psychopharmacol.* (2019). <https://doi.org/10.1037/pha0000274>
91. Zamani, H., et al.: Eye tracking application on emotion analysis for marketing strategy. *J. Telecommun. Electron. Comput. Eng.* **8**(11), 87–91 (2016)
92. Zhan, Z., et al.: Online Learners' reading ability detection based on eye-tracking sensors. *Sensors* **16**(9), 1457 (2016). <https://doi.org/10.3390/s16091457>
93. Fourth Quarter and Full Year 2018 Results. Facebook, Inc., Menlo Park, USA (2019)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



3

User Knowledge and Perceptions about Sensor-based Inference Attacks

3.1 Background

A significant body of research has explored people’s perceptions and awareness regarding the power of modern inferential analytics. Existing research in this area mostly focuses on inferences drawn from data collected by web trackers, search engines, and social networking sites (e.g., [226, 227, 228, 229, 230]). Several studies have also investigated people’s knowledge and concerns regarding the personal information that can be inferred from mobile and IoT sensor data (e.g., physiological sensors [104], smartphone motion sensors [47, 195], in-home sensing applications [193, 231]). Findings from prior research indicate that the privacy-invading potential and diversity of sensor-based inference attacks are widely unknown among the general public [47, 193, 195]. Accordingly, I have painted device users as vulnerable and unsuspecting victims of such attacks, for instance, by giving my first academic publication the title: “Unexpected Inferences from Sensor Data: A Hidden Privacy Threat in the Internet of Things” [4]. In Paper 1 of this thesis, it was concluded that, through the lens of data analytics, voice recordings “may (...) reveal much more information than a speaker wishes and expects to communicate.” Similarly, in Paper 3, it was stated: “[Eye-tracking] data can unexpectedly contain a wealth of sensitive information about a user. (...) Since inference methods are often based on hidden patterns and correlations (...), it can be impossible for [ordinary people] to understand and control what information is revealed.” But are such inferences really unexpected to most people? While previous research suggests so [47, 193, 195], the existing empirical base is sparse. As the above statements were made during our focus on the technical aspects of sensor-based inference attacks, before we thoroughly reviewed or conducted studies on people’s awareness and perceptions of inference attacks, they were mostly hypothetical.

Having obtained and published the results from Chapter 2, it was the logical next step to use some of these findings as a foundation for empirical research to answer questions like: How aware are people that personal information can be inferred from certain types of sensor data?

What perceptions and concerns do people have on this issue? And how do their attitudes and usage intentions towards specific devices or services change if they are informed about the possibilities of modern inferential analytics?

Specifically, after an in-depth scan of existing literature, the decision was made to conduct a large-scale user survey on the privacy impacts of voice and speech analysis, using knowledge compiled in Paper 1 as a basis for designing the questionnaire. The original idea had been to conduct a user study on accelerometer-based inference attacks but the focus then shifted to voice recordings, as the former had already been covered in previous work [47, 195] and microphones are similarly ubiquitous in modern life as accelerometers (cf. Chap. 2.1). I initiated and led the project, for which I collaborated with Saba Rebecca Brause and Dr. Stefan Ullrich from the Weizenbaum Institute, Leon Gellrich from Potsdam University, and Dr. Sebastian Pape from the Goethe University Frankfurt. My collaborators provided valuable help, mainly in the areas of data analysis and interpretation of results, but also with methodology design, data management, and the final editing of the manuscript. The resulting paper (Paper 4) was published in the journal *Proceedings on Privacy Enhancing Technologies* (PoPETs).

Jacob Leon Kröger*, Leon Gellrich, Sebastian Pape, Saba Rebecca Brause, and Stefan Ullrich

Personal information inference from voice recordings: User awareness and privacy concerns

Abstract: Through voice characteristics and manner of expression, even seemingly benign voice recordings can reveal sensitive attributes about a recorded speaker (e.g., geographical origin, health status, personality). We conducted a nationally representative survey in the UK ($n = 683$, 18–69 years) to investigate people’s awareness about the inferential power of voice and speech analysis. Our results show that – while awareness levels vary between different categories of inferred information – there is generally low awareness across all participant demographics, even among participants with professional experience in computer science, data mining, and IT security. For instance, only 18.7% of participants are at least somewhat aware that physical and mental health information can be inferred from voice recordings. Many participants have rarely (28.4%) or never (42.5%) even thought about the possibility of personal information being inferred from speech data. After a short educational video on the topic, participants express only moderate privacy concern. However, based on an analysis of open text responses, unconcerned reactions seem to be largely explained by knowledge gaps about possible data misuses. Watching the educational video lowered participants’ intention to use voice-enabled devices. In discussing the regulatory implications of our findings, we challenge the notion of “informed consent” to data processing. We also argue that inferences about individuals need to be legally recognized as personal data and protected accordingly.

Keywords: privacy, voice recording, speech, microphone, voice assistant, smart speaker, inference attack

DOI 10.2478/popets-2022-0002

Received 2021-05-31; revised 2021-09-15; accepted 2021-09-16.

***Corresponding Author: Jacob Leon Kröger:** Weizenbaum Institute for the Networked Society, Technische Universität Berlin, Germany, E-mail: kroeger@tu-berlin.de
Leon Gellrich: Universität Potsdam, Germany
Sebastian Pape: Goethe Universität, Frankfurt, Germany
Saba Rebecca Brause, Stefan Ullrich: Weizenbaum Institute for the Networked Society, TU Berlin, Germany

1 Introduction

Microphones are found everywhere in today’s technology-based society. They are embedded not only into phones, tablets, laptops, electronic toys, cameras, wearables, and car dashboards but also into intercoms, baby monitors, remote controls, smart speakers, and all sorts of other smart home devices. Recordings from these omnipresent microphones, which contain voice commands, memos, and private conversations, are often accessible to a range of different parties. For example, microphones in mobile devices are regularly accessed by device manufacturers, platform providers, installed apps, and even by third-party software components completely invisible to the ordinary user [75]. Private calls and voice messages can, if not properly encrypted, be intercepted by instant messaging providers, network operators, videoconferencing services, and other intermediaries [54, 81, 100]. Similarly, audio data uploaded to a cloud storage system, social networking site, or media-sharing platform can be accessible not only to the audience intended by the user but also to the respective service and/or infrastructure provider [94]. More recently, the soaring popularity of voice-enabled devices [77] and the COVID-19 pandemic with increased rates of virtual meetings and voice/video calls [20] have substantially added to the volume of speech data available to corporations.

Beyond their legitimate processing purposes, these organizations may use personal information extracted from voice recordings for malicious ends or pass it on to other parties with unknown intentions, potentially exposing users to the risk of discrimination, invasive advertising, extortion, identity theft, and other types of fraud and abuse [15, 22]. As shown in recent work, attackers can build a model of a victim’s voice using only a limited number of voice samples in order to fool voice-based user authentication systems or to mimic the victim in speech contexts (e.g., leaving fake voice messages or posting fake statements on the Internet) [66]. The use of speech data for unauthorized and unexpected purposes is not limited to hackers and organized crime groups, but also practiced by government agencies [81, 100] and technol-

ogy companies, including major multinational corporations [54, 88].

There is extensive research on the privacy implications of microphone-equipped devices, with numerous studies looking into user perceptions and concerns (cf. Sect. 2.1). However, existing work in this field has almost exclusively focused on linguistic content, i. e., *what* a person says in a voice recording. An equally serious, yet largely neglected, privacy threat is posed by the fact that sensitive personal information can be inferred from *how* a person speaks. In fact, certain patterns and correlations in recorded speech can be much more revealing than the words themselves (cf. Sect. 2.2). As we show in a recent literature review [41], a speaker’s voice characteristics and manner of expression can implicitly contain information about his or her geographical origin, age, personality, emotions, level of sleepiness, physical and mental health condition, and more. So far, little is known about users’ perceptions of the risks associated with these possible inferences (cf. Sect. 2.3).

Contribution. To examine privacy concerns related to voice recordings, we conducted a survey of 683 Internet users in the UK. Our survey is, to our knowledge, the first to focus specifically on the privacy impacts of audio-based inferences.

- Our results show a widespread lack of awareness about the inferential power of voice and speech analysis, with varying levels of awareness for different types of inferences. (Sect. 5.1)
- While we observe differences in awareness across demographic groups, these differences are not large. Even participants with professional experience in the ICT field score low on awareness. (Sect. 5.2)
- Participants’ reactions to audio-based inferences are quite evenly distributed between worried and unworried. An analysis of open text responses offers insight into their reasoning. (Sect. 5.3)
- Participants’ intention to use voice-controlled virtual assistants significantly decreased after a short educational video on the topic. (Sect. 5.4)

Outline. The paper proceeds as follows. First, Sect. 2 reviews related literature. Then, we describe our research goals (Sect. 3) and methodology (Sect. 4). In Sect. 5, we present the study results, a discussion of which follows in Sect. 6. We reflect on the limitations of our study in Sect. 7, before we conclude the paper in Sect. 8.

2 Related work

2.1 Privacy perceptions and concerns about audio recording

There are a variety of studies on users’ perceptions regarding the privacy impacts of microphone-equipped devices. Aspects that have been investigated include people’s mental models for a privacy/utility trade-off [12, 50, 52], concerns about always-listening devices [12, 13, 35, 52, 55, 102], user trust in providers of voice-based services [50, 52], concerns about voice data being shared with third parties or used for other purposes than stated [60, 91], and concerns around microphone-equipped devices getting hacked and other forms of unauthorized access [52, 55, 102].

Studies have also investigated how the context of use affects the acceptability of audio recordings (e. g., at work [35] or in public [64]). Moorthy and Vu [64], for example, investigated the usage of voice-controlled virtual assistants and found that participants preferred to use such devices in private locations to avoid drawing embarrassing attention and being overheard by strangers.

Another line of research focuses on users’ awareness and use of privacy controls for microphone-equipped devices, and the willingness to pay for security and privacy features [21, 50, 51, 60, 62]. Recent studies suggest that users are poorly informed about potential privacy risks and therefore not particularly concerned about embedded microphones in their devices [50, 51, 55, 61, 102].

2.2 Sensor-based inference of personal information

It has long been known that sensor data from consumer devices can be analyzed to extrapolate patterns and draw sensitive inferences about the user. The mining of data to illegitimately gain knowledge about a person is referred to as an “inference attack” [47]. Some published review articles summarize data categories that can be inferred from IoT and mobile sensor data [37, 87, 90], video game data [45], accelerometer data [44], and eye tracking data [40, 53].

In a recent literature review, we have examined the wealth of information that can be extracted from voice recordings [41]. Through the lens of advanced data analysis, the voice and manner of expression of a recorded speaker may reveal information about his or her geo-

graphical origin [5, 32], gender¹ and age [34, 83], mental health [4, 73], physical health [17, 18], level of intoxication [9, 86], moods and emotions [36, 95], sleepiness and fatigue [17, 86], and personality traits [78, 85].

For example, researchers have used speech errors and irregularities (e.g., number of false and unintelligible words, interrupts, hesitations) and rhythmicity features for distinguishing alcoholized from non-alcoholized speech [9, 86], used voice hoarseness and sounds like coughs and sniffles for the detection of sore throats and flu infections [18, 33], and used voice pitch variations and speech energy levels for automatic emotion recognition (e.g., anger, compassion, disgust, happiness, surprise) [33, 36, 95]. A large variety of features, including speaking rate, loudness, spectral features and characteristics of linguistic expression, has been applied for the inference of personality traits [78, 85].

While such algorithmic predictions are of course not always correct and can also be significantly impaired by situational factors (e.g., ambient noise, reverberation, microphone quality), remarkable accuracies have already been reported. Polzehl [78], for instance, reached 85.2% accuracy in classifying speakers into ten different personality classes along the Big Five traits (openness to experience, conscientiousness, extraversion, agreeableness, neuroticism). From voice energy features, Sadjadi et al. [83] estimated the age of speakers with a mean absolute error of 4.7 years. In computational paralinguistics challenges, researchers have achieved up to 91% accuracy in speech-based intoxication detection [86]. And, while mental health insights can also be derived from acoustic characteristics (e.g., monotone speech, glottal features) [73], Bedi et al. [4] analyzed semantic coherence and speech complexity to automatically detect emergent psychosis in test subjects, reaching the classification accuracy of traditional clinical interviews.

Findings from such experimental studies usually refer to specific experimental setups and demographic groups (e.g., “prediction of major depression in adolescents” [73]) and are subject to limitations (e.g., laboratory conditions, limited sample size). Therefore, while providing evidence that speech-based inferences in their respective field are possible in principle, their specific

findings are not necessarily generalizable to all people and all real-life situations.

It should be noted, on the other hand, that data controllers with access to speech data (e.g., large tech corporations) can be much better equipped in terms of budget, technical expertise and training data than the researchers cited above, meaning that the risk of undesired inferences from audio data is likely bigger in real life than it appears based on published results. Since data analysis methods are often subject to non-disclosure agreements, the most advanced know-how in voice and speech analysis arguably rests within the industry and is not publicly available [41]. It can also be assumed that the variety and effectiveness of audio-based inference attacks will further increase with the rising popularity of voice-enabled devices [77] and continuing advances in computing technologies and audio analysis methods (e.g., feature optimization [9] and deep learning approaches [17]). Therefore, despite the remaining technological challenges and limitations, such attacks pose a real and growing threat to consumer privacy that needs to be taken seriously and thoroughly investigated.

Existing products and features, such as voice-analytics tools for hiring assessment [96], call centers [57] and for illness detection based on smart speaker voice commands [33], indicate that companies intend to use speech data to draw sensitive inferences about users in practice.

2.3 User awareness and perceptions about sensor-based inference attacks

There is little research on users’ knowledge of the inferential power of mobile and IoT sensor data. A few previous studies have investigated people’s privacy concerns associated with inferences that can be drawn from physiological sensors (e.g., ECG and respiration) [80], smartphone motion sensors [16, 61], in-home sensing applications [14, 103], and IoT systems for companies (e.g., sensors for room occupancy monitoring and energy consumption tracking) [79]. Existing research indicates a widespread lack of awareness about sensor-based inference attacks [16, 61, 103]. Mehrnezhad et al. [61] found that smartphone users are not aware that various mobile sensors can be exploited to infer a personal identification number (PIN) typed on the touchscreen. In general, the researchers observed very low levels of understanding about the existence and functioning of embedded sensors. When presented with examples of personal data inference, users’ privacy concerns tend

¹ Note that people’s personal experience and internal understanding of their gender can vary from the gender socially assigned to them based on reproductive organs, anatomy, chromosomes, etc. Most existing work in the field under investigation focuses on inferring people’s self-reported gender.

to increase [16, 61, 80]. However, among the participants surveyed by Crager et al. [16], perceptions about inference attacks significantly varied between different demographic groups. In contrast to our study, none of these research efforts focuses specifically on audio data. And in existing studies on privacy concerns about voice recordings, threats posed by inferential analytics have been largely overlooked [50, 51, 55, 64].

3 Research goals

When the words spoken in a voice recording directly contain private information (e. g., expression of political opinion, credit card details, health-related information), the sensitivity of the recording is obvious. What has not been sufficiently explored in previous research, however, is people’s awareness of the wealth of information that can be inferred from a speaker’s voice characteristics and manner of expression (cf. Sect. 2). To fill the identified research gap, this study examines users’ awareness and concerns regarding audio-based inference attacks, the vulnerability of different demographic groups to this privacy threat, and whether informing people on this issue changes their intentions towards using microphone-equipped devices. In examining these issues, we focus on eight types of audio-based inferences that have emerged from our literature search (cf. Sect. 2.2). For analysis, we group these inferences into three clusters: inferences about demographics (**DEM**), short- and medium-term states (**STATE**), and physical and psychological traits (**TRAIT**). DEM includes geographical origin, gender and age. STATE includes level of intoxication, moods and emotions, sleepiness and fatigue. TRAIT includes mental health, physical health, and personality traits. Note that this paper makes no claim to be exhaustive. For instance, inferences could also be drawn from background sounds in voice recordings (e. g., media, urban or animal sounds) [92] or ultrasonic signals [43], which can be privacy-sensitive but were not included as they do not directly relate to users’ voice and speech characteristics. We pose the following research questions:

RQ-1: How aware are people that personal information can be inferred from voice recordings?

Since humans tend to perform quite well in estimating certain speaker attributes based on voice features

in everyday life (e. g., age, gender, emotions, intoxication) [49, 65, 84], they may better understand the possibility of automated inferences in these areas (compared to diagnosing a specific mental disorder based on speech features, for example). Human perception of voice identity is particularly related to a speaker’s age and gender [6, 48, 65]. Thus, we postulate:

H1: The level of awareness is higher for DEM inferences than for STATE inferences, and lowest for TRAIT inferences.

RQ-2: How does the level of awareness differ across demographic groups?

This question aims to identify at-risk populations by correlating awareness levels with participant demographics. All sorts of domain-related knowledge, technical understanding and privacy experience could assist in understanding the possibilities of modern data analytics. Advanced age, on the other hand, has been associated with lower degrees of ICT literacy [72]. We also explore the relationship between awareness and participants’ income, gender, and disposition to value privacy. For lack of clear indications in the literature, these were tested without directional hypothesis. We postulate:

H2.1: Awareness is positively correlated with previous privacy experience, general privacy awareness, innovativeness (i. e., a person’s tendency to be a technology pioneer), general level of education, and with professional experience in data protection law, computer science, data mining, and IT security.

H2.2: Awareness and age are negatively correlated.

H2.3: There are relationships between awareness and income, gender, and disposition to value privacy.

RQ-3: What concerns do people have about the inference of personal information from voice recordings?

Physical and psychological traits are more stable over time and may thus reveal more about a person’s life and character than temporary state variables. By contrast, basic demographic information is widely perceived as relatively non-sensitive [63] and is often already entered during sign-up to a digital service. Thus, we postulate:

H3: The level of concern is higher for TRAIT inferences than for STATE inferences, and lowest for DEM inferences.

RQ-4: How do people’s usage intentions for voice-enabled devices change when being informed on the topic?

For this question, while audio data can be recorded with any kind of microphone-equipped device, we decided to focus on voice-controlled virtual assistants (VCVAs). This choice was made with regard to the soaring popularity of services like Amazon Alexa and Apple’s Siri across the globe [77] and because user voice commands are usually available to the service providers in unencrypted form [54, 88], enabling them to screen the audio for revealing patterns and correlations. Based on previous research suggesting that privacy concerns are an important factor affecting the adoption of voice-enabled devices [29, 64], we postulate:

H4: The educational intervention will have a negative impact on VCVA usage intention.

4 Research methodology

Our four research questions were investigated by means of an online survey. After a brief overview of our study design, this section will provide detailed descriptions of our survey instrument (Sect. 4.1), participant recruitment process (Sect. 4.2), characteristics of our sample (Sect. 4.3) and methods used for data analysis (Sect. 4.4).

Participants’ awareness that personal information can be inferred from voice recordings (RQ-1) was studied based on self-reported measures. While alternative approaches exist (cf. Sect. 7), we asked about participants’ awareness after showing them a short educational video² on the topic, as inspired by Crager et al. [16]. To identify potential at-risk populations with particularly low levels of awareness (RQ-2), we included multiple demographic items in the survey and then correlated the results with participants’ reported levels of awareness. To explore participants’ concerns about personal information inference from voice recordings (RQ-3), we queried their reactions to the educational video through rating scales and one open text question.

To be able to test whether our video had an effect on participants’ interest in using voice-enabled devices (RQ-4), we decided to use two slightly different

questionnaires in our study (**Grp-A** and **Grp-B**). In Grp-A, participants were asked about their usage intention twice – once before, and once after the educational video, thus allowing a within-subject comparison to examine the effect of the intervention. However, repeating one question within a questionnaire may introduce a bias: Participants’ answers to the repeated question could be influenced by their previous response to the same question. Therefore, to create a control group, participants in Grp-B were asked about their usage intention only once (after the video). By comparing results from the post-intervention question in Grp-A and Grp-B using a Kolmogorov–Smirnov test, we checked whether repeating the question in Grp-A substantially affected participants’ responses. Besides underpinning the validity of the within-subject comparison among Grp-A participants, this approach allowed us to conduct a between-subject comparison between the pre-intervention question in Grp-A and the post-intervention question in Grp-B, providing additional insight into the impact of our educational video on participants’ intention to use voice-enabled devices.

4.1 Survey instrument

Both questionnaires, Grp-A and Grp-B, consist of one educational video² and 53 questions, including three attention checks. The video clip and all questions are exactly the same in Grp-A and Grp-B – only one question is repeated in Grp-A, as will be detailed below. The questionnaires were programmed with the software SoSci Survey (version 3.2.19) [93]. All responses capturing the intensity of feelings or level of agreement were measured on 5-point Likert-type scales. To allow for reproducibility, a copy of all survey items can be found in appendix A. It took participants a median of 8.5 minutes to complete our survey. The questionnaires contained the following:

- **9 privacy demographic items:** Using validated scales directly adapted from Xu et al. [101], participants were asked about their general privacy awareness (PA), disposition to value privacy (DVP), and previous privacy experience (PPE). Results from these items were used in answering RQ-2.
- **2 items on voice-controlled virtual assistants (VCVA):** To ensure a common level of understanding among participants, this section was started with a short textual definition of VCVA, including examples of common features. Participants were

² Video clip available here: https://youtu.be/Gr22YqS1_VA.

then asked (i) how often they use a VCVA in daily life, and (ii) to what extent they are interested in starting or continuing to use a VCVA. As explained in the beginning of Sect. 4, the position of question (ii) varied between Grp-A and Grp-B for the purpose of inter-group comparison and bias control. In questionnaire Grp-A, the question was posed before *and* after the educational video, in Grp-B it was posed *only once* after the video. Items from this block were used in answering RQ-4.

- **1 educational video²:** As a preparation for questions on this topic, participants were presented with a short informational video (1:44 minutes) about audio-based inferences. The video explains that, for certain functions, microphone-equipped devices typically transmit voice recordings to remote company servers, where the audio data can be analyzed to extract various kinds of personal information. Based on previous research (cf. Sect. 2.2), the video lists categories of data that could be derived from a speaker’s voice characteristics and manner of expression, namely *geographical origin, gender and age, mental health, physical health, level of intoxication, moods and emotions, sleepiness and fatigue, and personality traits*. To ensure that the audio track is audible and the video is watched through to the end, we included a preliminary sound check and displayed a code at the end of the video which was requested on the following page. If the code was not entered correctly, the survey was terminated and the corresponding participant was excluded from analysis.
- **20 items on awareness and concerns regarding audio-based inference attacks:** After the video, the participants were asked (i) whether they had been aware that such inferences are possible, (ii) how concerned they are about the possibility of such inferences, (iii) how often they have consciously thought about this issue before, (iv) how common they think it is for companies to draw such inferences from voice recordings, and (v) how concerned they are about individual categories of inferred information. Items (i) and (iii) were used in answering RQ-1, the other items were used in answering RQ-3.
- **10 technology demographic items:** Adapting a 9-item scale from Parasuraman and Colby [76], this section measures the participants’ level of innovativeness (INNO), i.e., the tendency to be a technology pioneer. Then, the participants were asked to select from a list all types of microphone-equipped devices they own. INNO was used in answering RQ-2,

the other question was used to provide descriptive sample statistics.

- **10 items on basic demographics and professional experience:** Participants were queried for their age, gender, net income, and level of education. Also, they were asked to specify their level of professional experience in the areas of data protection law, computer science, data mining, and IT security. These items were used in answering RQ-2.

Three attention checks were incorporated in the survey to screen for random responders and potential bot submissions (cf. questions 8, 16, 21 in appendix A). Before the actual online survey was conducted, we administered a pretest to a total of 58 participants using the crowdsourcing platform *Amazon Mechanical Turk* (<https://www.mturk.com/>). In this way, we were able to test our attention checks and refine the survey instruments, including a clarification of potentially ambiguous wording. Based on the pretest results, there were only minor adjustments.

Our survey instruments and research procedures were approved by the Ethics Committee at Goethe University Frankfurt.

4.2 Participant recruitment

To access a sample of UK adults, we used the services of the online market research firm *respondi AG* (<https://www.respondi.com/EN/>) which was carefully selected from a list of ten competing panel providers and fulfils the quality management system standards of ISO 20252 [26]. Although crowdsourcing platforms, such as *MTurk* and *Prolific*, offer several benefits in terms of cost efficiency, speed, and flexibility, we favored the option of hiring a panel company for several reasons. Above all, while recent studies have obtained high-quality results from crowdsourced samples [58, 82], there are widespread concerns about generalizability [11, 74, 98]. Significant differences between *MTurk* workers and general population estimates were found in family composition, political attitudes, and religiosity [11], level of education and health behavior [98], social engagement [58], and internet activity [82], to name a few examples.

According to our requirements, the sample for our study was designed to approximate the age and gender distribution of adults (18-69 years) from the latest UK census [24], which also resembles current population es-

Table 1. Participant demographics

| Age group | Grp-A (n = 349) | | Grp-B (n = 334) | |
|-----------|-----------------|-------------|-----------------|-------------|
| | male | female | male | female |
| 18-29 | 39 (11.2%) | 40 (11.5%) | 39 (11.7%) | 40 (12.0%) |
| 30-39 | 35 (10.0%) | 37 (10.6%) | 32 (9.6%) | 33 (9.9%) |
| 40-49 | 34 (9.7%) | 41 (11.7%) | 36 (10.8%) | 39 (11.7%) |
| 50-59 | 33 (9.5%) | 30 (8.6%) | 31 (9.3%) | 30 (9.0%) |
| 60-69 | 32 (9.2%) | 28 (8.0%) | 26 (7.8%) | 28 (8.4%) |
| Total | 173 (49.6%) | 176 (50.4%) | 164 (49.1%) | 170 (50.9%) |

timates from the UK Office for National Statistics [25]. Given a desired power of 95% and an estimated effect size of 0.3, an a priori power analysis revealed a required sample size of around 200 participants. Considering the explorative nature of the study and our available resources, we collected valid responses from $n = 683$ participants. Survey completers received a small compensation according to the terms of our panel provider. The survey was conducted between June 4 and July 1, 2020.

4.3 Sample characteristics

In total, 1,277 participants signed up for the survey. 588 responses were excluded for being incomplete, either because the participant had closed the questionnaire before answering all survey questions ($n=235$), or because the participant had failed to pass one of our attention checks ($n=353$). Additionally, six responses were eliminated due to obvious poor quality of their data, as assessed by independent raters. Our analysis is based on the remaining final sample of 683 participants. The age of participants ranges from 18 to 69 years ($\mu = 42.99$, $\sigma = 14.50$) with 50.7% being females. A breakdown of the age and gender distribution for both test groups is provided in Table 1. 99% of participants report to own at least one microphone-equipped device (95% smartphone, 79% laptop, 54% tablet, 36% smart speaker, 20% voice-enabled remote control, 14% in-vehicle voice control interface, 13% smartwatch). All participants are UK residents.

4.4 Data analysis

Statistical analysis. While all scales in our questionnaire are treated as parametric, we expected – due to the nature of the subject and based on existing literature – that results throughout the survey would be highly skewed (e.g., because related work indicates

low awareness levels for sensor-based inference attacks, cf. Sect. 2.3). After visually checking the histograms, Shapiro-Wilk tests confirmed that the survey results for the used scales are not normally distributed ($p < 0.001$). Thus, we used non-parametric tests for comparative analyses.

The Friedman test [23, p. 686ff] was used as a non-parametric alternative to a repeated-measures ANOVA. To test the difference between means of dependent variables, post-hoc Wilcoxon signed-rank tests with Bonferroni-corrected alpha [23, p. 914] were used as a non-parametric alternative to paired t-tests. To test the difference between means of *independent* variables, a Wilcoxon rank-sum test³ [23, p.655ff] was used as a non-parametric alternative to a two-sample t-test. For correlation analysis, since the commonly used Pearson correlation coefficient (Pearson’s r) requires normal distribution of the sample data when attempting to establish whether the correlation coefficient is significant [23, p. 219], we used Spearman’s rank correlation coefficient (Spearman’s ρ) [23, p. 223ff] instead. To obtain ordinal variables suitable for analysis, the variable income was clustered and the variable education was recoded (e.g., master’s degree above bachelor’s degree), as shown in the published dataset [38].

We further conducted regression analyses based on the Akaike information criterion (AIC), using forward selection and backward elimination procedures. The software environment R (version 4.0.0) was used for statistical data analysis.

Qualitative thematic analysis. Responses to the open text question (cf. appendix A, №11) were evaluated using a thematic analysis as proposed by Brown and Clarke [10], which is a method for identifying patterns of meaning within qualitative data.

After familiarizing himself with the data, a first rater systematically assigned descriptive and interpretative codes to all features in the data with potential relevance to the question posed. The resulting codebook was then used by a second rater to independently label and categorize the received responses, adding new codes where deemed appropriate. We used the Cohen’s Kappa coefficient to measure inter-rater reliability. All instances of discrepancy were discussed and resolved jointly by the two raters. The assigned codes were then used to identify frequent responses and overarching themes (cf. Sect. 5.3).

³ also known as Mann–Whitney U test

5 Results

We collected $n = 683$ complete and valid survey responses ($n = 349$ for Grp-A, $n = 334$ for Grp-B). In terms of age and gender distribution, Grp-A and Grp-B are approximately identical, both being nationally representative for the UK population between 18 and 69 years. In this section, we analyze the survey responses with respect to the research questions introduced in Sect. 3. We have released an annotated and sanitized dataset containing our results for all participants [38].

5.1 RQ-1. How aware are people that personal information can be inferred from voice recordings?

Our results presented in Fig. 1 indicate widespread unawareness of inferences that can be drawn from voice and speech parameters. Averaged over the eight types of inferences covered in our questionnaire, 67.6% of participants reported to be “not at all” or only “slightly” aware. We observed, however, that the level of awareness strongly differs between the individual inference categories. For example, while 48.2% of participants reported to be “somewhat”, “quite” or “very” aware about the possibility of inferring a speaker’s gender and age based on a voice recording, this figure drops to 18.7% for the inference of physical and mental health information.

For a statistical analysis of these differences, we compared the three clusters of inferences defined in Sect. 3, namely inferences about demographics (**DEM**), short- and medium-term states (**STATE**), and physical and psychological traits (**TRAIT**). A Friedman test [23, p. 686ff] yielded significant differences in awareness levels between DEM, STATE, and TRAIT ($\chi^2 = 425.53$, $p < 0.001$, Kendall’s $W = 0.312$). Post-hoc Wilcoxon [23, p. 667ff] signed-rank tests with Bonferroni-corrected alpha [23, p. 914] revealed that all three pair-wise comparisons are significant ($p < 0.001$).

In confirmation of hypothesis H1, the test results show that the level of awareness is higher for DEM inferences than for STATE inferences, and lowest for TRAIT inferences. We obtained a moderate effect size for the DEM-STATE comparison (0.342) and large effect sizes for the STATE-TRAIT (0.520) and DEM-TRAIT (0.707) comparisons. Post-hoc power analysis revealed that these tests had a very high power ($> 99\%$).

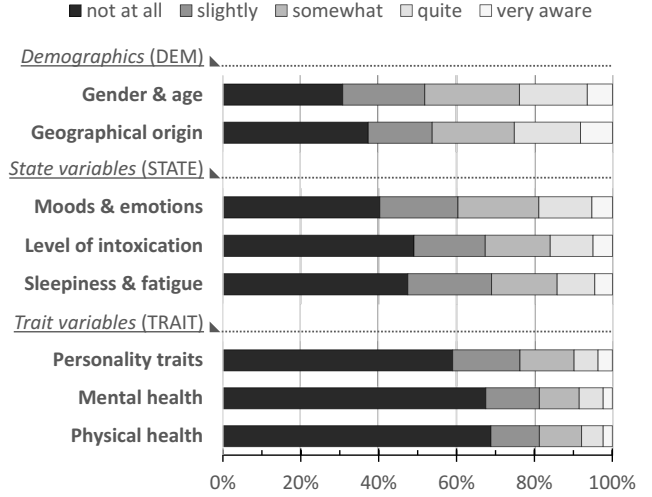


Fig. 1. Distribution of participants’ level of awareness dependent on the inferred information

Before taking the questionnaire, a large portion of participants has “never” (42.5%) or “rarely” (28.4%) consciously thought about the possibility of personal information being inferred from speech data. Only a small minority reports to have pondered on this issue “often” (7.0%) or “very often” (2.8%).

5.2 RQ-2. How does the level of awareness differ across demographic groups?

To explore statistical relationships between the awareness levels from RQ-1 and participant demographics, we first calculated three awareness scores for each participant: **AW_DEM** (avg. awareness about DEM inferences), **AW_STATE** (avg. awareness about STATE inferences), and **AW_TRAIT** (avg. awareness about TRAIT inferences).

We then tested for correlations between participant demographic attributes and **AW_DEM**, **AW_STATE**, and **AW_TRAIT** using Spearman’s rank correlation coefficient (Spearman’s ρ) [23, p. 223ff]. An overview of the correlation results, along with their Bonferroni-corrected significance levels, is provided in Table 6 in appendix B. Note that for all correlations for which we had a directional hypothesis (cf. Sect. 3), the confidence intervals are one-sided, meaning they end at 1.00 or -1.00 , respectively.

Supporting our hypotheses H2.1 and H2.2, **AW_DEM**, **AW_STATE**, and **AW_TRAIT** are negatively correlated with participant age and positively cor-

Table 2. Regression results for AW_DEM

| Coefficients | Estimate | Std. Error | t-value | p-value |
|--------------|----------|------------|---------|----------|
| (Intercept) | 1.72 | 0.34 | 5.15 | 3.85e-07 |
| age | -0.01 | 0.00 | -2.73 | 0.00653 |
| gender | -0.30 | 0.11 | -2.67 | 0.00798 |
| EXP_DM | 0.23 | 0.09 | 2.70 | 0.00712 |
| EXP_CS | 0.17 | 0.07 | 2.54 | 0.01132 |
| INNO | 0.13 | 0.07 | 1.97 | 0.04914 |
| PA | 0.11 | 0.07 | 1.41 | 0.15857 |

 Adjusted R²: 0.1532

Table 3. Regression results for AW_STATE

| Coefficients | Estimate | Std. Error | t-value | p-value |
|--------------|----------|------------|---------|---------|
| (Intercept) | 0.62 | 0.26 | 2.52 | 0.01208 |
| EXP_DM | 0.17 | 0.08 | 2.11 | 0.03482 |
| EXP_CS | 0.16 | 0.06 | 2.79 | 0.00558 |
| PA | 0.28 | 0.08 | 3.60 | 0.00036 |
| DVP | -0.12 | 0.06 | -1.90 | 0.05805 |
| PPE | 0.13 | 0.08 | 2.11 | 0.03546 |

 Adjusted R²: 0.1499

Table 4. Regression results for AW_TRAIT

| Coefficients | Estimate | Std. Error | t-value | p-value |
|--------------|----------|------------|---------|---------|
| (Intercept) | 0.58 | 0.27 | 2.19 | 0.02912 |
| age | -0.01 | 0.00 | -1.98 | 0.04854 |
| EXP_DM | 0.25 | 0.07 | 3.62 | 0.00033 |
| EXP_CS | 0.14 | 0.05 | 2.87 | 0.00431 |
| PA | 0.15 | 0.06 | 2.56 | 0.01070 |
| PPE | 0.10 | 0.06 | 1.49 | 0.13707 |

 Adjusted R²: 0.1954

AW_: average level of awareness about audio-based inference of demographics (DEM) / short- and medium-term states (STATE) / physical and psychological traits (TRAIT); **INNO**: innovativeness; **PA**: privacy awareness; **DVP**: disposition to value privacy; **PPE**: previous privacy experience; **EXP_**: professional experience in data mining (DM) / computer science (CS)

related with general level of education, degree of innovativeness (INNO), previous privacy experience (PPE), general privacy awareness (PA), and with professional experience in data protection law (EXP_DP), computer science (EXP_CS), data mining (EXP_DM), and IT security (EXP_IS). Across AW_DEM, AW_STATE, and AW_TRAIT, the most notable correlations were found with EXP_CS, EXP_DM, and EXP_IS. However, using guidelines from Dancey and Reidy [19] to interpret the results, even these are weak correlations.

The correlation test results further indicate that men tend to have slightly higher AW_TRAIT and AW_STATE (but not AW_DEM) than women. No significant correlation was found between participants' awareness and their disposition to value privacy (DVP) or level of income, contradicting hypothesis H2.3.

While the observed tendencies suggest that some population groups are somewhat less informed than others (and thus potentially more prone to unwittingly reveal sensitive information about themselves through speech data), there seems to be little awareness regarding audio-based inferences throughout all population segments. Even among those participants who reported "much" or "very much" EXP_CS ($n = 120$), EXP_IS ($n = 82$), or EXP_DM ($n = 60$), a large portion (49.4%, 45.9%, and 43.1%, respectively) stated to be "not at all" or only "slightly" aware (averaged over the eight types of inferences covered in our questionnaire).

For the continuous demographic variables (INNO, PA, DVP, PPE, EXP_DP, EXP_DM, EXP_CS, EXP_IS, gender, age), we additionally conducted regression analyses based on the Akaike information criterion (AIC) to test their predictive power on awareness. After forward selection and backward elimination procedures, we obtained models for AW_DEM, AW_STATE, and AW_TRAIT, with five to six predictors each and adjusted R² ranging from 15% to 19.5%. The results are shown in Tables 2, 3 and 4. For instance, an increase of EXP_DM by 1 point resulted in increased AW_DEM by 0.23, AW_STATE by 0.17, and AW_TRAIT by 0.25. The only predictors that were consistently significant across the three models were EXP_CS and EXP_DM, the other predictors were less relevant and varied between the models. The low R² results indicate that only a small portion of the variance in awareness can be explained by the demographic variables from our survey. This supports our above conclusion that all demographic groups under investigation are similarly vulnerable to audio-based inference attacks, i.e. similarly prone to disclose more information about themselves via speech data than they expect.

Regarding the internal consistency of constructs adapted from the literature (cf. Sect. 4.1), we obtained a good or excellent Cronbach's alpha for INNO (0.93), DVP (0.85), and PA (0.81), according to interpretation guidelines provided by George and Mallery [27]. For PPE, we obtained a Cronbach's alpha of 0.66, indicating a questionable internal consistency [27]. The lower Cronbach's alpha for PPE is mainly driven by the second item of the construct (cf. appendix A, № 3) which shows relatively low consistency with the other items of

PPE. Thus, the second item should receive special attention in future uses of the construct and may require improvement and re-validation.

5.3 RQ-3. What concerns do people have about the inference of personal information from voice recordings?

After watching a short educational video on the privacy implications of voice and speech analysis (cf. Sect. 4.1), participants were asked how worried they are about the possibility of personal information inference from speech data. Responses to this question were mixed. While 38.7% of participants reported to be “not at all” or “slightly” worried, a similar proportion (40.7%) stated to be “quite” or “very” worried.

The average participant believes that it is common rather than exceptional for companies to infer personal information from voice recordings. When asked to estimate the prevalence of this practice, 12.6% selected “somewhat” or “very” uncommon, 48.0% were “undecided”, and 38.4% selected “somewhat” or “very” common.

In an open text question, we asked the participants to provide a reasoning for their reported level of concern. Except for a handful of cases, all participants offered an intelligible response. The responses were evaluated by two raters using the thematic analysis method proposed by Brown and Clarke [10], as described in Sect. 4.4. The number of assigned codes per response varies between 1 and 6. The coding yielded a Cohen’s Kappa of 0.83, indicating a high degree of inter-rater reliability [59]. After instances of discrepancy were discussed and resolved jointly by the two raters, the assigned codes were used to identify overarching themes. Along with the complete dataset of our study, we have released the resulting codebook, including all identified themes [38].

In the following, we will summarize our findings regarding the most salient themes, namely participants’ *emotional reactions* (Sect. 5.3.1), *feared data misuses* (Sect. 5.3.2), *perceived benefits* and *inevitability* of voice-based technology (Sect. 5.3.3) as well as participants’ *knowledge gap’s and misconceptions* (Sect. 5.3.4). Additionally, findings regarding participants’ concerns towards specific types of inferences will be presented in Sect. 5.3.5. When quoting responses, we will either state the corresponding number of participants (**Ps**) or, if one individual participant is quoted, state the respective participant ID from the dataset (**P₁ to P₆₈₃**).

5.3.1 Emotional reactions

Approximately half of the open text responses illustrate or emphasize negative feelings. For instance, the inferential power of voice and speech analysis is perceived as “alarming” (2 Ps), “frightening” (2 Ps), “unnerving” (2 Ps), “unsettling” (2 Ps), “shocking” (2 Ps), “disturbing” (3 Ps), “uncomfortable” (4 Ps), “scary” (8 Ps), “concerning” (8 Ps), “Big Brother[ish]” (12 Ps), “worrying” (14 Ps), and “intrusive” (15 Ps). Partially, the negative reactions are quite strong, showing that confronting people with this issue can elicit “a great sense of helplessness” (P₃₆₁).

Participants are surprised at the variety of possible inferences, stating that they “hadn’t considered” (P₅₃₅), “didn’t realise” (3 Ps), had “no idea” (2 Ps) that “all this information could be revealed [...] by a voice recording” (P₃₇₄). In some cases, participants even express amazement about the possibilities of modern voice and speech analysis, e.g. describing them as “fascinating” (2 Ps) or “far beyond what I dreamed” (P₄₁₁). Others express confusion, e.g. by stating, “I don’t quite understand how they could possibly get this information from voice alone” (P₅₆). Participant P₄₄₁ concludes with the words: “The world is a lot cleverer than we realise.”

On the other hand, there are also participants who state to be completely indifferent about the the privacy implications of voice recordings. “If I like [a technology], I don’t care about side effects”, says P₃₉₉ and P₁₈₂ claims she “couldn’t care less what information people have on me.”

5.3.2 Feared data misuses

Participants express concern that microphone-equipped devices may collect more data than required for their functionality, and that the collected data might be used for “unrelated purposes” (P₄₆₀) without the user’s consent or awareness. While companies usually provide some form of privacy policy, it is objected that customers “rarely read them carefully or understand their implications fully” (P₆₇₈). Feared types of data processing and data misuse include targeted advertising to shape “political views/consumption habits” (P₅₂₄), data-based discrimination by insurances and employers (12 Ps) as well as “fraud” (2 Ps) and “identity theft” (P₂₆₀).

Further, there is concern that information extracted from voice recordings could end up in the “wrong hands” (6 Ps) by being passed on or leaked to third

parties, such as affiliate companies, hacker groups, or governmental agencies. Opposition is not only directed against criminal data use and governmental surveillance but also explicitly against “using very personal information for commercial purposes” (P₃₉).

Additional doubts and worries are expressed over the accuracy of inference algorithms. “A lot depends on how this information is interpreted”, says P₆₇₁. Inferences are feared to be “inaccurate and presumptive” (P₆₄₂) and “taken out of context” (P₆₄₄), leading to “assumptions being made that aren’t actually true for the individual” (P₂₆).

5.3.3 Perceived benefits and inevitability of voice-based technology

Despite their concerns, some participants perceive the disclosure of sensitive personal data as a necessary trade-off for using modern technology. “Unfortunately, I feel this is just the way the world is heading”, says P₅₅₉. Others agree: “companies have been collecting data for years” (P₄₇₆) and “there is little we can do about it” (P₆₆₄).

There are also responses specifically focusing on the beneficial uses of microphone-equipped technology, e. g., for creating “convenient products” (P₁₂₉), supporting “security and crime-fighting services” (P₁₃₇), targeting “adverts to sell me products/services that may assist” (P₃₀₃) or “alerting medical services if someone is in danger due to physical or mental health issues” (P₄₉₃). Voice control is perceived by some as “an evolutionary step in how we and our children will interact with devices” (P₅₂₉) which will “improve humanity” (P₄₁₃) and be used “for the greater good” (P₅₀₂). While they also see potential downsides, optimistic participants are confident that “the benefits far outweigh the negatives at the moment” (P₆₁₀) and that privacy loss is a “small price to pay for more convenient products” (P₁₂₉).

5.3.4 Knowledge gaps and misconceptions

It is striking that – although we specifically asked for their reasoning – none of the unworried participants provided a solid justification for their reported lack of privacy concern. Instead, their responses reveal potentially dangerous, yet understandable, misconceptions and false senses of security. For instance, unconvinced by our educational video, some participants do not believe that the presented audio-based inferences are tech-

nically feasible at all. While the sources and arguments compiled in Sect. 2.2 suggest otherwise, participants’ disbelief in a short educational video on a complex and unfamiliar topic is of course an understandable reaction.

Other participants do not see how data extracted from voice recordings could be used against their interest. “I really don’t care as I have no idea how this information could be used to my detriment”, states P₁₆₄. And P₄₄₇ asks: “Why would I be worried? I have nothing to hide.” The nothing-to-hide argument, which was put forth by many participants, has been criticized for its narrow view on privacy and for ignoring various threats that can arise from personal data being available to malicious or negligent parties [99].

Some participants explain that they are not worried because they do not own a voice-controlled device, such as a smart speaker. As exemplified in Sect. 1 and illustrated in our educational video, audio data (e. g., voice messages, voice memos, voice calls) can be recorded, analyzed and transmitted to remote servers by a wide variety of devices – not only by voice-controlled devices. Even living entirely without microphone-equipped devices would not guarantee protection against audio-based inference attacks, as a person’s voice can – intentionally or unintentionally – be recorded by other people’s devices (cf. Sect. 6.2). It should be noted, however, that our video focuses on direct user-device interaction and puts a slight emphasis on voice-controlled devices to prepare participants for questions related to RQ-4, which could be a source of misunderstanding.

Finally, a few participants base their sense of security on the assumption that their data will always be stored securely and only used sparingly and responsibly. For example, they doubt that any information extracted from voice recordings “would be used to identify me personally” (P₃₉₆), trusting that such data “would be in an anonymous format anyway” (P₁₁₀), whereas in reality this is often not the case. Others express confidence that their “privacy settings do the job” (P₂₁₁) and that companies would not use data “negatively against me” (2 Ps), “for truly bad purposes” (P₄₄₉) or “in any negative ways” (P₃₅₃). Participant P₆₅₄ states: “Given the [...] general consensus of privacy violations being bad for business, I don’t worry too much about inappropriate use.” We also received vague and confusing statements along this line, such as “It doesn’t trace it back to me personally as they will never meet me” (P₄₉₉) or “I assume the Internet has protection in place” (P₄₂₇).

In reality, however, companies can clearly leak or exploit personal data in harmful ways and commonly share such data with a range of third parties (cf. Sect. 1).

In light of the above observations, unworried reactions among participants appear to be largely explained by knowledge gaps.

5.3.5 Concerns towards specific types of inferences

We also asked the participants how concerned they would be if a company used voice recordings to infer specific types of information about them without their awareness. The results are shown in Fig. 2. As can be seen at first glance, the reported level of concern considerably varies between the information categories. For instance, while 60.8% of participants reported to be “quite” or “very” concerned about the disclosure of mental health information, only 31.3% of participants showed the same level of concern about inferences on their gender and age.

For a statistical analysis of these differences, we again compared the clusters defined in Sect. 3, namely inferences about demographics (DEM), short- and medium-term states (STATE), and physical and psychological traits (TRAIT). A Friedman test yielded significant differences in concern levels between these clusters ($\chi^2 = 386.87$, $p < 0.001$, Kendall’s $W = 0.283$). Post-hoc Wilcoxon signed-rank tests [23, p. 667ff] with Bonferroni-corrected alpha revealed that all three pairwise comparisons are significant ($p < 0.001$).

The test results confirm hypothesis H3 by showing that the level of concern is lowest for DEM inferences, followed by STATE inferences, and highest for TRAIT inferences. We obtained a moderate effect size for the DEM-STATE comparison (0.409) and large effect sizes for the DEM-TRAIT (0.643) and STATE-TRAIT (0.513) comparisons. Post-hoc power analysis revealed that these tests had a very high power ($> 99\%$).

Even in their response to the open text question (cf. Sect. 5.3), some participants have focused their concern on specific data categories (e.g., P₅₁₉: “Health [data] isn’t really something you want to be shared without consent”), while other types of inferred data, such as age, gender and level of intoxication, were rarely mentioned at all. Analogous to the knowledge gaps noted in Sect. 5.3.4, the variation in concern levels between different types of inferences could indicate a lack of awareness or understanding of how certain data categories can be misused.




Fig. 2. Distribution of participants’ level of concern dependent on the inferred information

5.4 RQ-4. How do people’s usage intentions for voice-enabled devices change when being informed on the topic?

To examine whether our educational intervention had an effect on the intention to use a voice-controlled virtual assistant (VCVA), we conducted a between- and a within-subject comparison, as explained in Sect. 4.

First, we compared responses from Grp-A ($n = 349$), where participants were asked about their VCVA usage intention (VCVA_UI) before the intervention, with responses from Grp-B ($n = 334$), where the same question was asked after the intervention. While a Wilcoxon rank-sum test [23, p.655ff] showed a significant difference between VCVA_UI in Grp-A and Grp-B in the direction predicted by hypothesis H4 ($p < 0.05$), i.e. lower VCVA_UI after the intervention, it only yielded a very small effect size (0.086). Given the small effect size, post-hoc power analysis also revealed low power (29%). The slightness of the result is not too surprising, as the expected effect may have been masked by those Grp-A participants who already reported no interest in using a VCVA before the intervention and thus left no room for a reduction in interest.

We then conducted a within-subject comparison in Grp-A. Participants in this group were asked about their VCVA_UI twice, once before and once after the intervention. Of course, repeating a question within a questionnaire may introduce a bias, as noted in Sect. 4. However, a Kolmogorov–Smirnov test revealed no significant difference ($p = 0.52$) between the distribu-

tion of responses to our post-intervention questions on VCVA_UI in Grp-A and Grp-B (despite the large sample sizes), which indicates that the two samples belong to the same distribution and that repeating the question in Grp-A did not substantially affect the distribution. Thus, we proceed with the within-subject comparison.

By comparing results from the pre- and post-intervention items on VCVA_UI within Grp-A, we analyze whether the intervention had a significant effect on VCVA_UI among Grp-A participants. To avoid the masking effect described above, we excluded those participants from the analysis who reported to be “not at all” or only “slightly” interested in using a VCVA before watching the educational video, leaving $n = 151$ participants for the analysis (Grp-A.1). Within Grp-A.1, a Wilcoxon signed-rank test [23, p. 667ff] not only yielded a significant decrease in VCVA_UI after the intervention ($p < 0.001$) but also a large effect size (0.590). Here, post-hoc power analysis revealed very high power ($> 99\%$).

These results suggest that for people with a medium or high interest in using such devices, information on the privacy implications of voice and speech analysis can have a strong negative effect on usage intentions, thus supporting hypothesis H4. Further research is required to investigate causes and motivational aspects behind these observations, and to determine whether these changes in reported usage intention sustain beyond the short term and actually translate into shifts in consumption behavior.

Table 5 shows VCVA_UI before and after the intervention, $\Delta VCVA_UI$ as well as awareness and concern levels, averaged over Grp-A and Grp-A.1. It can be seen that Grp-A and Grp-A.1 yield similar results for concern and awareness, indicating that Grp-A.1 is a representative subgroup of Grp-A. The similarity further indicates that those participants who reported to be “not at all” or only “slightly” interested in using a VCVA before watching the educational video (and thus excluded in Grp-A.1) showed little interest not because they are more concerned than the rest of Grp-A but, e.g., because they simply have no use for a VCVA.

In Grp-A.1, we additionally tested for correlations between the observed change in VCVA_UI ($\Delta VCVA_UI$) and the participants’ level of concern and awareness about the possibility of audio-based inference of demographics (DEM), short- and medium-term states (STATE), and physical and psychological traits (TRAIT). We found that concerns are positively correlated with $\Delta VCVA_UI$ ($p < 0.001$; Spearman’s $\rho = 0.38$ for DEM, 0.42 for STATE, 0.40 for

Table 5. Mean values for Grp-A and Grp-A.1

| | Grp-A (n = 349) | Grp-A.1 (n = 151) |
|-------------------|-----------------|-------------------|
| VCVA_UI (pre) | 2.47 | 3.97 |
| VCVA_UI (post) | 2.11 | 3.20 |
| $\Delta VCVA_UI$ | -0.36 (-14.6%) | -0.77 (-19.4%) |
| AW_DEM | 2.47 | 2.50 |
| AW_STATE | 2.08 | 2.03 |
| AW_TRAIT | 1.60 | 1.68 |
| CON_DEM | 2.62 | 2.53 |
| CON_STATE | 3.18 | 3.06 |
| CON_TRAIT | 3.55 | 3.40 |

VCVA_UI: VCVA usage intention before (pre) and after (post) the intervention; **$\Delta VCVA_UI$** : difference between VCVA_UI pre and post intervention; **AW_**: level of awareness about audio-based inference of demographics (DEM) / short- and medium-term states (STATE) / physical and psychological traits (TRAIT); **CON_**: level of concern about DEM / STATE / TRAIT inferences

TRAIT). The self-reported level of awareness about audio-based inferences was negatively correlated with $\Delta VCVA_UI$ (DEM: $p < 0.01$, Spearman’s $\rho = -0.23$; STATE: $p < 0.01$, Spearman’s $\rho = -0.21$; TRAIT: $p < 0.05$, Spearman’s $\rho = -0.21$). This means that VCVA_UI of participants with higher levels of prior knowledge about audio-based inferences tend to be less affected by our educational intervention. This is not surprising, as the intervention should logically have a higher educational impact on those participants who knew less on the topic beforehand.

6 Discussion and implications

The results of our survey reveal that there is widespread lack of understanding about the possibilities of voice and speech analysis. For instance, 81.3% of participants are not at all or only slightly aware that physical and mental health information can be inferred from a recorded speaker’s voice characteristics and manner of expression. Only 9.8% of participants have often or very often consciously thought about the possibility of personal information being inferred from voice and speech parameters.

Our results, offering a novel contribution by specifically focusing on voice recordings, are consistent with previous findings indicating a lack of awareness about inference attacks based on other types of sensor data [16, 61, 103]. While our analysis shows that awareness varies significantly depending on participants’ demographic

attributes, which was previously also shown to be the case for motion sensor-based inference attacks [16], the average level of awareness is low across all demographic groups – even among participants with professional experience in ICT (cf. Sect. 5.2).

In our sample, the degree of worry regarding audio-based inferences is quite evenly distributed between high and low. Results from analyzing open text responses suggest, however, that unconcerned reactions are largely explained by knowledge gaps about the risks that can arise from privacy intrusions (cf. Sect. 5.3.4). While some participants express worry about unauthorized data leakage to third parties, specific types of data misuse, and about being misrepresented by audio-based inferences (cf. Sect. 5.3.2), our results confirm previous findings about people’s unwarranted trust in companies’ data practices [50, 52, 102, 103] and a widespread nothing-to-hide mentality [50, 51, 91, 99, 103], potentially resulting in a false sense of security.

At the same time, the reported level of concern varies significantly between different categories of inferred data (e.g., high concern about inferred health information vs. low concern about inferred age – cf. Sect. 5.3.5), which may be due to individual preferences but perhaps also indicates a lack of understanding on how certain data categories can be used against the data subject’s interests.

Our educational intervention on the privacy implications of voice and speech analysis had a significant negative impact on participants’ intention to use voice-controlled virtual assistants. This result aligns with previous findings that users’ privacy concerns tend to increase when they are presented with examples of personal data inference [16, 61, 80].

6.1 Consumer education and privacy-enhancing technologies

While Internet-connected microphones have many beneficial applications (e.g., efficient human-computer interaction, assistance for physically disabled people, smart home convenience, driver safety), their increasing ubiquity in modern life calls for a debate on potential social ramifications. Besides the already omnipresent microphones in smartphones, laptops and other mobile devices, the number of installed smart speakers is forecast to reach 640 million globally by 2024 [70]. Nothing is fundamentally wrong with either microphone-equipped devices or speech data mining, but there is clearly a need for appropriate privacy safeguards.

Educating people on existing threats is an important starting point – not only to support informed purchase decisions but also to put critical pressure on the societal actors responsible for protecting consumer privacy in sensing devices.

With regard to data collection transparency in voice-controlled devices, there has been a focus on device recording modes, such as “speech-activated”, “manually activated”, and “always on” [12, 28]. In the face of recurring security breaches and privacy scandals, users have not only been advised to use the mute feature of their voice-controlled devices but also been encouraged to disconnect power supply or even purposely obfuscate audio signals to protect themselves against corporate and governmental eavesdropping [12].

However, while there are good reasons to be concerned about always-listening devices, it is important to understand that the mentioned safeguards – even if effectively applied in practice – will not prevent audio-based inference attacks (unless, of course, they permanently block the microphone and prevent any recording.) As discussed in this paper, voice and speech characteristics can unexpectedly carry sensitive personal information, which may later be extracted via advanced data analytics (cf. Sect. 2.2). Thus, even if a voice assistant is only consciously unmuted by a user to ask for the weather forecast, for instance, this can already lead to unwanted information leakage (e.g., based on sociolect, accent, intonation, pitch, loudness, or a hoarseness in the user’s voice).

To minimize privacy risks, voice recordings should preferably be encrypted before any upload or Internet transfer, and the data processing should take place as much as possible locally on the user’s device. In cases where the disclosure of speech data to service providers is unavoidable (e.g., because the data is necessary for service functionality or due to resource constraints of the end device), measures should be taken to prevent the illegitimate inference of personal information.

Some technical approaches that could help to defend against audio-based inference attacks are differential private learning, hashing techniques for speech data, fully homomorphic inference systems, and speaker de-identification by voice transformation [67, 68]. In recent work, for example, Aloufi et al. [2, 3] have proposed privacy-preserving intermediate layers to sanitize user voice input before sharing it with cloud service providers. These approaches, which are based on the automatic identification and obfuscation of sensitive features in speech data, have yielded promising evaluation results for certain use cases, such as protection against

unwanted emotion [2, 3] and gender recognition [3] while maintaining utility of the data for speech and speaker recognition.

Where possible without compromising the required functionality, voice recordings should also be transcribed to text in order to preserve task-relevant information while removing speaking speed, rhythm, voice characteristics, etc. and thus reduce the risk of inference attacks. In their proposed *Preech* system for privacy-preserving speech transcription, Ahmed et al. [1] apply voice transformation and the injection of noise to obfuscate users’ voice biometrics and thus prevent unauthorized identification and impersonation.

6.2 Regulatory implications

Considering that (i) existing technical solutions for protecting against sensor-based inference attacks have severe limitations [68, 87] and are still seen as “embryonic research topics” [69], (ii) companies obviously need strong incentives to apply privacy-enhancing technologies [30], (iii) many users are not willing to pay for privacy and their willingness to pay depends on the trust towards the provider of the privacy enhancing technology [31], and (iv) there is – as our study underscores – a very low level of risk awareness among users, adjustments in privacy regulation may be required as well.

To achieve a minimum level of transparency and oversight, inferences should at least be recognized as falling within the scope of data protection law. While the newly introduced California Consumer Privacy Act (CCPA), for example, specifically covers “inferences drawn” as part of its definition of personal information, most other data protection laws – including progressive ones, such as EU’s General Data Protection Regulation (GDPR) – do not sufficiently protect individuals against undesired inferences [8]. In a detailed legal analysis, Wachter and Mittelstadt [97] state that the GDPR “focuses primarily on mechanisms to manage the input side of processing. (...) [T]he few mechanisms in European data protection law that address the outputs of processing, including inferred and derived data, profiles, and decisions, are far weaker.”

Data protection law could, for example, make it mandatory for companies to provide comprehensive information on all types of inferences that they (attempt to) draw from collected personal data. Given that data mining algorithms are becoming an increasingly accurate and efficient access path to personal information, this could be a sensible measure.

For data subjects, being able to answer the question “who knows what about me?” is a necessary precondition for exercising other data protection rights (e.g., data rectification, erasure, restriction of processing) in an informed manner [39]. The widespread lack of understanding of how personal data can be collected, inferred, and misused calls into question the notion of “informed consent” and may warrant some form of paternalistic government intervention. As we argue in other recent work, people’s privacy choices are typically irrational, involuntary and/or easily circumventable [42].

Accordingly, various commentators have proposed a legal shift from the individualistic paradigm of notice and consent (“privacy self-management”) towards an increased focus on the ethical and social impacts of personal data use (e.g., [56, 71, 89, 97]). For instance, a general legal prohibition of using certain categories of personal data for ethically indefensible purposes based on the resulting harm potential could be helpful to protect consumers from consequences of their own unawareness.

Another argument against the self-management approach is that it ignores the various externalities that individual privacy choices have on other people and society at large [42]. In today’s interconnected world, people often share personal data of other people, giving rise to the notion of “interdependent privacy” [7]. Owners of microphone-equipped devices can become amateur data controllers without the data subject’s knowledge or consent. For example, someone might record a phone conversation and share it with a third party. Furthermore, a user’s device can record the voice, activities, etc. of persons in the vicinity (e.g., relatives, friends, visitors, bystanders), potentially scaling up the inference problem by a significant factor.

7 Limitations

While surveys are widely used in related empirical studies [16, 61, 79, 80], this form of data collection is subject to several potential limitations.

There is of course the risk of careless or random responding. We incorporated multiple attention checks into our survey to filter out low-quality responses. Only those respondents who passed all quality and attention checks were included in the analysis (cf. Sect. 4.3).

Furthermore, a self-reported survey captures subjective perceptions, which are prone to distortion. In particular, following an approach proposed by Crager

et al. [16], we asked participants for their awareness of audio-based inferences *after* showing them a short educational video on the topic. It may have been difficult for some participants to accurately recall what their level of knowledge was prior to watching the video.

A possible alternative would be to ask participants about awareness before showing the video (e.g., by asking how likely they think different types of inferences are). Even this approach, however, may evoke thoughts that participants would not have by themselves in everyday life. Moreover, participants with low levels of knowledge and skills may have a tendency to overestimate their abilities (a cognitive bias referred to as the Dunning-Kruger effect) [46]. Therefore, future work could build upon this study by using approaches that query the knowledge of participants more implicitly and objectively, instead of using self-reported measures of awareness.

Additionally, learning from our study’s limitations, follow-up studies should thoroughly test participants’ understanding of educational materials (e.g., in the form of a quiz). It is possible that participants did not understand everything in the video.

It is also possible that participants exaggerated certain responses in an attempt to present themselves in a more positive light, e.g., by stating that they are more familiar with technology than they actually are, by overstating their professional experience in some area, or by falsely claiming to be less interested in using a virtual assistant after our educational intervention. This effect may have been increased by asking participants about privacy attitudes at the beginning of the survey, priming them to think about privacy. We cannot exclude the possibility of a social-desirability bias but believe to have minimized the risk of occurrence through the neutral framing of our educational video and by informing participants in advance that the results of our online survey would be completely anonymous. Asking about privacy attitudes at the end would not have eliminated the issue of priming because, in this case, the privacy focus of the remaining survey may have influenced participants’ responses to these questions.

It should also be noted that our findings are only representative for the UK population, which is a typical WEIRD society (Western, Educated, Industrialized, Rich, Democratic). Replication studies in other contexts, such as in Asian or African countries, are required to establish cultural validity.

8 Conclusion

Microphones have become ubiquitous in modern life, embedded into mobile, wearable, and all sorts of smart home devices. While these devices provide useful functions, the increasing availability of private voice recordings to service providers, device manufacturers, app vendors, etc. has also become a major threat to consumer privacy. In this study, focusing on an issue that has received very little research attention to date, we investigated people’s awareness and privacy concerns about the wealth of personal information that can be inferred by analyzing a recorded speaker’s voice characteristics and manner of expression. Our results indicate a widespread lack of awareness about the possibilities of modern voice and speech analysis. Averaged over the eight types of inferences covered in our questionnaire, most participants reported to be “not at all” (50.0%) or only “slightly” (17.6%) aware. Even participants with professional experience in the ICT field scored low on awareness. Furthermore, while our results for participants’ level of concern about audio-based inference attacks do not show a clear tendency, many participants – judging from their text responses – seem to lack the background knowledge required to assess these threats in an informed manner. Overall, the findings of this study underscore that the complexities of modern data processing are beyond the comprehension of ordinary users – which calls into question the notion of “informed consent,” a cornerstone of most modern data protection laws, including EU’s GDPR. To prevent consent from being used as a loophole to excessively reap data from unwitting individuals, alternative and complementary technical, organizational, and regulatory safeguards urgently need to be developed. At the very least, inferred information relating to an individual should be classified as personal data by law, subject to corresponding protections and transparency rights. Results from our within- and between-subject comparisons suggest that education on data analytics may have an impact on smart device use, the mechanisms and implications of which are an interesting avenue for future research.

9 Acknowledgments

We thank the anonymous reviewers for their constructive feedback. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- [1] Shimaa Ahmed, Amrita Roy Chowdhury, Kassem Fawaz, and Parmesh Ramanathan. 2020. Preech: A system for privacy-preserving speech transcription. In *29th USENIX Security Symposium*. 2703–2720.
- [2] Ranya Aloufi, Hamed Haddadi, and David Boyle. 2019. Emotionless: privacy-preserving speech analysis for voice assistants. *preprint arXiv:1908.03632* (2019).
- [3] Ranya Aloufi, Hamed Haddadi, and David Boyle. 2020. Privacy-preserving Voice Analysis via Disentangled Representations. In *ACM SIGSAC Conference on Cloud Computing Security Workshop*. 1–14.
- [4] Gillinder Bedi et al. 2015. Automated analysis of free speech predicts psychosis onset in high-risk youths. *npj Schizophrenia* 1 (2015), 15030.
- [5] Hamid Behravan, Ville Hautamäki, Sabato Marco Siniscalchi, Tomi Kinnunen, and Chin-Hui Lee. 2015. I-Vector modeling of speech attributes for automatic foreign accent recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 24, 1 (2015), 29–41.
- [6] Pascal Belin, Shirley Fecteau, and Catherine Bedard. 2004. Thinking the voice: neural correlates of voice perception. *Trends in cognitive sciences* 8, 3 (2004), 129–135.
- [7] Gergely Biczók and Pern Hui Chia. 2013. Interdependent privacy: Let me share your data. In *Int. Conf. on Financial Cryptography and Data Security*. Springer, 338–353.
- [8] Jordan M Blanke. 2020. Protection for ‘Inferences Drawn’: A Comparison Between the General Data Protection Regulation and the California Consumer Privacy Act. *Global Privacy Law Review* 1, 2 (2020).
- [9] Daniel Bone, Ming Li, Matthew P Black, and Shrikanth S Narayanan. 2014. Intoxicated speech detection: A fusion framework with speaker-normalized hierarchical functionals and GMM supervectors. *Computer Speech & Language* 28, 2 (2014), 375–391.
- [10] Virginia Braun and Victoria Clarke. 2012. Thematic Analysis. In *APA Handbook of Research Methods in Psychology*. American Psychological Association, Washington, DC.
- [11] Jesse Chandler, Cheskie Rosenzweig, Aaron J Moss, Jonathan Robinson, and Leib Litman. 2019. Online panels in social science research: Expanding sampling methods beyond Mechanical Turk. *Behavior research methods* 51, 5 (2019), 2022–2038.
- [12] Varun Chandrasekar, Kassem Fawaz, Bilge Mutlu, and Suman Banerjee. 2018. Characterizing Privacy Perceptions of Voice Assistants: A Technology Probe Study. *arXiv:1812.00263 [cs]* (2018).
- [13] Chola Chhetri and Vivian Genaro Motti. 2019. Eliciting Privacy Concerns for Smart Home Devices from a User Centered Perspective. In *Information in Contemporary Society*, Natalie Greene Taylor et al. (Eds.). Springer, Cham, 91–101.
- [14] Eun Kyoung Choe, Sunny Consolvo, Jaeyeon Jung, Beverly Harrison, Shwetak N. Patel, and Julie A. Kientz. 2012. Investigating Receptiveness to Sensing and Inference in the Home Using Sensor Proxies. In *UbiComp*. ACM, 61–70.
- [15] Wolfie Christl. 2017. *How Companies Use Data Against People*. Cracked Labs, Vienna.
- [16] Kirsten Crager, Anindya Maiti, Murtuza Jadliwala, and Jibo He. 2017. Information Leakage through Mobile Motion Sensors: User Awareness and Concerns. In *EuroUSEC*. Internet Society, Paris, France.
- [17] Nicholas Cummins, Alice Baird, and Bjoern W Schuller. 2018. Speech analysis for health: Current state-of-the-art and the increasing impact of deep learning. *Methods* 151 (2018), 41–54.
- [18] Nicholas Cummins, Maximilian Schmitt, Shahin Amiriparian, Jarek Krajewski, and Björn Schuller. 2017. “You sound ill, take the day off”: Automatic recognition of speech affected by upper respiratory tract infection. In *Proceedings of the IEEE EMBC Conference*. 3806–3809.
- [19] Christine P Dancy and John Reidy. 2007. *Statistics without Maths for Psychology*. Pearson Education.
- [20] Evan DeFilippis, Stephen Michael Impink, Madison Singell, Jeffrey T Polzer, and Raffaella Sadun. 2020. *Collaborating during Coronavirus: The impact of COVID-19 on the nature of work*. National Bureau of Economic Research, Cambridge, MA.
- [21] Serge Egelman, Raghudeep Kannavara, and Richard Chow. 2015. Is This Thing On? Crowdsourcing Privacy Indicators for Ubiquitous Sensing Platforms. In *Proceedings of the CHI Conference*. ACM, New York, 1669–1678.
- [22] Federal Bureau of Investigation. 2020. 2019 Internet Crime Report. https://pdf.ic3.gov/2019_IC3Report.pdf
- [23] Andy Field, Jeremy Miles, and Zoë Field. 2012. *Discovering statistics using R*. Sage Publishing, Newbury Park, CA.
- [24] Office for National Statistics. 2013. The National Archives. <https://webarchive.nationalarchives.gov.uk/20160110194058/http://www.ons.gov.uk/ons/publications/re-reference-tables.html?edition=tcM%3A77-294277> (last accessed on 12 September 2021).
- [25] Office for National Statistics. 2019. Population estimates for the UK, England and Wales, Scotland and Northern Ireland: mid-2019. <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/bulletins/annualmidyearpopulationestimates/mid2019estimates>
- [26] International Organization for Standardization. 2019. ISO 20252:2019. <https://www.iso.org/obp/ui/#iso:std:iso:20252:ed-3:v1:en> (last accessed on 12 September 2021).
- [27] D George and P Mallery. 2003. *Reliability analysis. SPSS for Windows, step by step: a simple guide and reference*. Allyn & Bacon, Boston, MA.
- [28] Stacey Gray. 2016. Always On: Privacy Implications of Microphone-Enabled Devices. https://www.ftc.gov/system/files/documents/public_comments/2016/08/00003-128652.pdf
- [29] Sangyeal Han and Heetae Yang. 2018. Understanding adoption of intelligent personal assistants. *Industrial Management & Data Systems* 118, 3 (2018), 618–636.
- [30] David Harborth, Maren Braun, Akos Grosz, Sebastian Pape, and Kai Rannenberg. 2018. Anreize und Hemmnisse für die Implementierung von Privacy-Enhancing Technologies im Unternehmenskontext. In *Sicherheit 2018*. 29–41.
- [31] David Harborth, Xinyuan Cai, and Sebastian Pape. 2019. Why Do People Pay for Privacy?. In *ICT Systems Security and Privacy Protection – 34th IFIP TC 11 International Conference*. 253–267.

- [32] Abhinav Jain, Minali Upreti, and Preethi Jyothi. 2018. Improved Accented Speech Recognition Using Accent Embeddings and Multi-task Learning. In *Proc. Interspeech*. 2454–2458.
- [33] Huafeng Jin and Shuo Wang. 2018. Voice-based determination of physical and emotional characteristics of users. <https://patents.google.com/patent/US10096319B1/en> US Patent 10,096,319.
- [34] Selen Hande Kabil, Hannah Muckenhirn, et al. 2018. On Learning to Identify Genders from Raw Speech Signal Using CNNs. In *Proc. Interspeech*. 287–291.
- [35] Predrag Klasnja, Sunny Consolvo, Tanzeem Choudhury, Richard Beckwith, and Jeffrey Hightower. 2009. Exploring Privacy Concerns about Personal Sensing. In *International Conference on Pervasive Computing*. Springer, 176–183.
- [36] Shashidhar G Koolagudi, Sudhamay Maity, Vuppala Anil Kumar, Saswat Chakrabarti, and K Sreenivasa Rao. 2009. IITKGP-SESC: speech database for emotion analysis. In *International Conference on Contemporary Computing*. Springer, 485–492.
- [37] Jacob Kröger. 2019. Unexpected Inferences from Sensor Data: A Hidden Privacy Threat in the Internet of Things. In *Internet of Things. Information Processing in an Increasingly Connected World*, Leon Strous and Vinton G. Cerf (Eds.). Springer, Cham, 147–159.
- [38] Jacob Leon Kröger, Leon Gellrich, Sebastian Pape, Saba Rebecca Brause, and Stefan Ullrich. 2021. Response data - Survey on privacy impacts of voice & speech analysis. <http://dx.doi.org/10.14279/depositonce-12309.2>
- [39] Jacob Leon Kröger, Jens Lindemann, and Dominik Herrmann. 2020. How do app vendors respond to subject access requests? A longitudinal privacy study on iOS and Android Apps. In *International Conference on Availability, Reliability and Security*. 1–10.
- [40] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking. In *Privacy and Identity Management. Data for Better Living: AI and Privacy*, Samuel Fricker, Michael Friedewald, Stephan Krenn, Eva Lievens, and Melek Önen (Eds.). Springer, Cham, 226–241.
- [41] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Philip Raschke. 2020. Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference. In *Privacy and Identity Management. Data for Better Living: AI and Privacy*, Samuel Fricker et al. (Eds.). Springer, Cham, 242–258.
- [42] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Stefan Ullrich. 2021. The myth of individual control: Mapping the limitations of privacy self-management. <https://ssrn.com/abstract=3881776>. SSRN (2021).
- [43] Jacob Leon Kröger and Philip Raschke. 2019. Is my phone listening in? On the feasibility and detectability of mobile eavesdropping. In *IFIP Annual Conference on Data and Applications Security and Privacy*. Springer, 102–120.
- [44] Jacob Leon Kröger, Philip Raschke, and Towhidur Rahman Bhuiyan. 2019. Privacy Implications of Accelerometer Data: A Review of Possible Inferences. In *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy (ICCSPP)*. ACM, New York, NY, 81–87.
- [45] Jacob Leon Kröger, Philip Raschke, Jessica Percy Campbell, and Stefan Ullrich. 2021. Surveilling the Gamers: Privacy Impacts of the Video Game Industry. <https://ssrn.com/abstract=3881279>. SSRN (2021).
- [46] Justin Kruger and David Dunning. 1999. Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology* 77, 6 (1999), 1121.
- [47] John Krumm. 2007. Inference Attacks on Location Tracks. In *Pervasive Computing*, Anthony LaMarca et al. (Eds.). Springer, Cham, 127–143.
- [48] Norman J Lass, Karen R Hughes, Melanie D Bowyer, Lucille T Waters, and Victoria T Bourne. 1976. Speaker sex identification from voiced, whispered, and filtered isolated vowels. *J. Acoust. Soc. Am.* 59, 3 (1976), 675–678.
- [49] Marianne Latinus and Pascal Belin. 2011. Human voice perception. *Current Biology* 21, 4 (2011), R143–R145.
- [50] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-Seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction* 2 (2018), 1–31.
- [51] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. “Alexa, Stop Recording”: Mismatches between Smart Speaker Privacy Controls and User Needs. In *Symposium on Usable Privacy and Security*.
- [52] Yuting Liao, Jessica Vitak, Priya Kumar, Michael Zimmer, and Katherine Kritikos. 2019. Understanding the Role of Privacy and Trust in Intelligent Personal Assistant Adoption. In *Information in Contemporary Society*, Natalie Greene Taylor, Caitlin Christian-Lamb, Michelle H. Martin, and Bonnie Nardi (Eds.). Springer, Cham, 102–113.
- [53] Daniel J. Liebling and Sören Preibusch. 2014. Privacy Considerations for a Pervasive Eye Tracking World. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 1169–1177.
- [54] Natasha Lomas. 2019. Microsoft Tweaks Privacy Policy to Admit Humans Can Listen to Skype Translator and Cortana Audio. <https://social.techcrunch.com/2019/08/15/microsoft-tweaks-privacy-policy-to-admit-humans-can-listen-to-skype-translator-and-cortana-audio/> (last accessed on 12 September 2021).
- [55] Lydia Manikonda, Aditya Deotale, and Subbarao Kambhampati. 2018. What's up with Privacy?: User Preferences and Privacy Concerns in Intelligent Personal Assistants. In *AAAI/ACM Conference on AI, Ethics, and Society*. 229–235.
- [56] Alessandro Mantelero. 2018. AI and Big Data: A blueprint for a human rights, social and ethical impact assessment. *Computer Law & Security Review* 34, 4 (2018), 754–772.
- [57] Rob Matheson. 2016. Watch your tone. <https://news.mit.edu/2016/startup-cogito-voice-analytics-call-centers-ptsd-0120>. (last accessed on 12 September 2021).
- [58] Morgan N McCredie and Leslie C Morey. 2019. Who are the Turkers? A characterization of MTurk workers using the personality assessment inventory. *Assessment* 26, 5 (2019), 759–766.
- [59] Mary L McHugh. 2012. Interrater Reliability: The Kappa Statistic. *Biochemia Medica* 22, 3 (2012), 276–282.

- [60] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. 2017. Toys That Listen: A Study of Parents, Children, and Internet-Connected Toys. In *Proceedings of the CHI Conference*. 5197–5207.
- [61] Maryam Mehrnezhad, Ehsan Toreini, Siamak F. Shahandashti, and Feng Hao. 2018. Stealing PINs via Mobile Sensors: Actual Risk versus User Perception. *International Journal of Information Security* 17, 3 (2018), 291–313.
- [62] Cristina Mihale-Wilson, Jan Zibuschka, and Oliver Hinz. 2017. About User Preferences and Willingness to Pay for a Secure and Privacy Protective Ubiquitous Personal Assistant. In *Proceedings of the 25th European Conference on Information Systems (ECIS)*.
- [63] George R Milne, George Pettinico, Fatima M Hajjat, and Ereni Markos. 2017. Information sensitivity typology: Mapping the degree and type of risk consumers perceive in personal data sharing. *Journal of Consumer Affairs* 51, 1 (2017), 133–161.
- [64] Aarthi Easwara Moorthy and Kim-Phuong L. Vu. 2015. Privacy Concerns for Use of Voice Activated Personal Assistant in the Public Space. *International Journal of Human-Computer Interaction* 31, 4 (2015), 307–335.
- [65] Evelyne Moysse. 2014. Age estimation from faces and voices: a review. *Psychologica Belgica* 54, 3 (2014), 255–265. <http://dx.doi.org/10.5334/pb.aq>
- [66] Dibya Mukhopadhyay, Maliheh Shirvanian, and Nitesh Saxena. 2015. All your voices are belong to us: Stealing voices to fool humans and machines. In *European Symposium on Research in Computer Security*. Springer, 599–621.
- [67] Andreas Nautsch et al. 2019. Preserving privacy in speaker and speech characterisation. *Computer Speech & Language* 58 (2019), 441–480. <http://dx.doi.org/10.1016/j.csl.2019.06.001>
- [68] Andreas Nautsch, Catherine Jasserand, Els Kindt, Massimiliano Todisco, Isabel Trancoso, and Nicholas Evans. 2019. The GDPR & Speech Data: Reflections of Legal and Technology Communities, First Steps Towards a Common Understanding. *Proc. Interspeech* (2019), 3695–3699.
- [69] Andreas Nautsch, Jose Patino, Natalia Tomashenko, Junichi Yamagishi, Paul-Gauthier Noe, Jean-Francois Bonastre, Massimiliano Todisco, and Nicholas Evans. 2020. The Privacy ZEBRA: Zero Evidence Biometric Recognition Assessment. In *Proc. Interspeech*.
- [70] Evan Niu. 2020. Smart-Speaker Volumes Expected to Jump Next Year. <https://www.nasdaq.com/articles/smart-speaker-volumes-expected-to-jump-next-year-2020-10-23> (last accessed on 12 September 2021).
- [71] Data Ethics Commission of the Federal Government. 2019. *Opinion of the Data Ethics Commission*. German Federal Ministry of Justice and Consumer Protection, Berlin.
- [72] Tobias Olsson, Ulli Samuelsson, and Dino Viscovi. 2019. At risk of exclusion? Degrees of ICT access and literacy among senior citizens. *Information, Communication & Society* 22, 1 (2019), 55–72.
- [73] Kuan Ee Brian Ooi, Margaret Lech, and Nicholas B Allen. 2012. Multichannel weighted speech classification system for prediction of major depression in adolescents. *IEEE Transactions on Biomedical Engineering* 60, 2 (2012), 497–506. <http://dx.doi.org/10.1109/TBME.2012.2228646>
- [74] Yaakov Ophir, Itay Sisso, Christa SC Asterhan, Refael Tikochinski, and Roi Reichart. 2020. The Turker Blues: Hidden Factors Behind Increased Depression Rates Among Amazon’s Mechanical Turkers. *Clinical Psychological Science* 8, 1 (2020), 65–83.
- [75] Elleen Pan, Jingjing Ren, Martina Lindorfer, Christo Wilson, and David Chones. 2018. Panoptispy: Characterizing Audio and Video Exfiltration from Android Applications. In *Proceedings on Privacy Enhancing Technologies*. 33–50.
- [76] A Parasuraman and Charles L Colby. 2015. An Updated and Streamlined Technology Readiness Index: TRI 2.0. *Journal of Service Research* 18, 1 (2015), 59–74.
- [77] Sarah Perez. 2019. Report: Voice assistants in use to triple to 8 billion by 2023. <https://techcrunch.com/2019/02/12/report-voice-assistants-in-use-to-triple-to-8-billion-by-2023> (last accessed on 12 September 2021).
- [78] Tim Polzehl. 2016. *Personality in Speech: Assessment and Automatic Classification*. Springer, Cham.
- [79] Lesandro Ponciano, Pedro Barbosa, Francisco Brasileiro, Andrey Brito, and Nazareno Andrade. 2017. Designing for Pragmatists and Fundamentalists: Privacy Concerns and Attitudes on the Internet of Things. In *Brazilian Symposium on Human Factors in Computing Systems*.
- [80] Andrew Raji, Animikh Ghosh, Santosh Kumar, and Mani Srivastava. 2011. Privacy Risks Emerging from the Adoption of Innocuous Wearable Sensors in the Mobile Environment. In *Proceedings of the CHI Conference*. ACM.
- [81] Kari Rea. 2013. Glenn Greenwald: Low-Level NSA Analysts Have ‘Powerful and Invasive’ Search Tool. <http://abcnews.go.com/blogs/politics/2013/07/glenn-greenwald-low-level-nsa-analysts-have-powerful-and-invasive-search-tool>
- [82] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2019. How well do my results generalize? Comparing security and privacy survey results from MTurk, web, and telephone samples. In *2019 IEEE Symposium on Security and Privacy*. 1326–1343.
- [83] Seyed Omid Sadjadi, Sriram Ganapathy, and Jason W Pelecanos. 2016. Speaker age estimation on conversational telephone speech using senone posterior based i-vectors. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 5040–5044.
- [84] Florian Schiel. 2011. Perception of alcoholic intoxication in speech. In *Proc. Interspeech*. 3281–3284.
- [85] Björn Schuller, Stefan Steidl, Anton Batliner, Elmar Nöth, Alessandro Vinciarelli, Felix Burkhardt, Rob van Son, Felix Weninger, Florian Eyben, Tobias Bocklet, Gelareh Mohammadi, and Benjamin Weiss. 2015. A Survey on Perceived Speaker Traits: Personality, Likability, Pathology, and the First Challenge. *Computer Speech & Language* 29, 1 (2015), 100–131.
- [86] Björn Schuller, Stefan Steidl, Anton Batliner, Florian Schiel, Jarek Krajewski, Felix Weninger, and Florian Eyben. 2014. Medium-term speaker states — A review on intoxication, sleepiness and the first challenge. *Computer Speech & Language* 28, 2 (2014), 346–374.
- [87] Amit Kumar Sikder, Giuseppe Petracca, Hidayet Aksu, Trent Jaeger, and A. Selcuk Uluagac. 2018. A Survey on Sensor-Based Threats to Internet-of-Things (IoT) Devices and Applications. *arXiv:1802.02041 [cs]* (2018).

- [88] Dave Smith. 2019. MicrophoneGate: The World's Biggest Tech Companies Were Caught Sending Sensitive Audio from Customers to Human Contractors. Here's Where They Stand Now. <https://www.businessinsider.com/amazon-apple-google-microsoft-assistants-sent-audio-contractors-2019-8>
- [89] Daniel J Solove. 2013. Privacy self-management and the consent dilemma. *Harvard Law Review* 126 (2013), 1880.
- [90] Raphael Spreitzer, Veelasha Moonsamy, Thomas Korak, and Stefan Mangard. 2018. Systematic Classification of Side-Channel Attacks: A Case Study for Mobile Devices. *IEEE Communications Surveys & Tutorials* 20, 1 (2018), 465–488.
- [91] Anna Ståhlbröst, Annika Sällström, and Danilo Hollosi. 2014. Audio Monitoring in Smart Cities – an Information Privacy Perspective. In *12th International Conference e-Society*. 35–44.
- [92] Dan Stowell, Dimitrios Giannoulis, Emmanouil Benetos, Mathieu Lagrange, and Mark D Plumbley. 2015. Detection and classification of acoustic scenes and events. *IEEE Transactions on Multimedia* 17, 10 (2015), 1733–1746.
- [93] SoSci Survey. 2020. The Solution for Professional Online Questionnaires. <https://www.sosicisurvey.de/en/index>
- [94] Dan Svantesson and Roger Clarke. 2010. Privacy and consumer risks in cloud computing. *Computer Law & Security Review* 26, 4 (2010), 391–397.
- [95] Monorama Swain, Aurobinda Routray, and P. Kabisatpathy. 2018. Databases, Features and Classifiers for Speech Emotion Recognition: A Review. *International Journal of Speech Technology* 21, 1 (2018), 93–120.
- [96] VoiceSense. 2020. Speech Analysis as a Talent Recruiting and Retention Tool. <https://www.voicesense.com/solutions/talent-recruiting-and-retention>.
- [97] Sandra Wachter and Brent Mittelstadt. 2019. A right to reasonable inferences: Re-thinking data protection law in the age of big data and AI. *Colum. Bus. L. Rev.* (2019), 494–620.
- [98] Kelly Walters, Dimitri A Christakis, and Davene R Wright. 2018. Are Mechanical Turk worker samples representative of health status and health behaviors in the US? *PIOS One* 13, 6 (2018), e0198835.
- [99] Wikipedia. [n.d.]. Nothing to hide argument. https://en.wikipedia.org/wiki/Nothing_to_hide_argument (last accessed on 12 September 2021).
- [100] Wikipedia. [n.d.]. Telephone tapping. https://en.wikipedia.org/wiki/Telephone_tapping (last accessed on 12 September 2021).
- [101] Heng Xu, Tamara Dinev, Jeff Smith, and Paul Hart. 2011. Information privacy concerns: Linking individual perceptions with institutional privacy assurances. *Journal of the Association of Information Systems* 12, 12 (2011), 798–824.
- [102] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End user security and privacy concerns with smart homes. In *Symposium on Usable Privacy and Security* (Santa Clara, CA, USA). 65–80.
- [103] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Fearnster. 2018. User Perceptions of Smart Home IoT Privacy. *Proceedings of the ACM on Human-Computer Interaction* 2 (2018), 1–20.

A Survey questionnaire

1. What is your gender? (Female/Male/Other)
2. How old are you?
3. Privacy awareness⁴
 - I am aware of the privacy issues and practices in our society.⁵
 - I follow the news and developments about the privacy issues and privacy violations.⁵
 - I keep myself updated about privacy issues and the solutions that companies and the government employ to ensure our privacy.⁵
4. Disposition to value privacy⁴
 - Compared to others, I am more sensitive about the way companies handle my personal information.⁵
 - To me, it is the most important thing to keep my information privacy.⁵
 - Compared to others, I tend to be more concerned about threats to my information privacy.⁵
5. Previous privacy experience⁴
 - How often have you been a victim of what you felt was an improper invasion of privacy?⁶
 - How much have you heard or read during the past year about the use and potential misuse of the information collected from the Internet?⁶
 - How often have you experienced incidents where your personal information was used by a company without your authorization?⁶
6. Voice-controlled virtual assistant⁷

A virtual assistant is a software agent that can perform tasks based on voice commands, without the requirement for keyboard input. Some examples of commercially available virtual assistants are Apple's Siri, Amazon Alexa, Microsoft's Cortana, and Google Assistant. Among other capabilities, virtual assistants can set reminders, manage con-

⁴ The constructs Privacy Awareness (PA), Disposition to Value Privacy (DVP), and Previous Privacy Experience (PPE) are adapted from Xu et al. [101]

⁵ Item was measured on a 5-point Likert scale: Strongly disagree, Somewhat disagree, Neutral, Somewhat agree, Strongly agree

⁶ Item was measured on a 5-point Likert scale: Never, Rarely, Sometimes, Often, Very often

⁷ The order of items shown here reflects questionnaire Grp-A. In Grp-B, this component (№ 6, including the two questions) is positioned after the educational video, replacing question № 15. See Sect. 4.1 for explanation.

tacts, play music, take purchase orders, send messages and calls, provide weather reports, and manage smart home devices.

- How often do you use a voice-controlled virtual assistant in your daily life?⁶
 - Are you interested in starting or continuing to use a voice-controlled virtual assistant?⁸
7. Please attentively watch this video (1:44 minutes) about the privacy implications of voice data.⁹
 8. Please enter the code displayed in the video. (not case-sensitive)
 9. Before you watched the video, were you aware that these types of information can be inferred from voice recordings?
 - Geographical origin¹⁰
 - Gender and age¹⁰
 - Mental health information¹⁰
 - Physical health information¹⁰
 - Level of intoxication¹⁰
 - Moods and emotions¹⁰
 - Sleepiness and fatigue¹⁰
 - Personality traits¹⁰
 10. How worried are you about these possible inferences?¹¹
 11. Why do you feel this way? Please explain your reasoning in two or more sentences.
 12. Prior to this questionnaire, how often have you consciously thought about this issue when using a microphone-equipped device?⁶
 13. What do you think, how common is it for companies to draw such inferences from voice recordings?¹²
 14. Please rate how concerned you would be if a company used voice recordings to infer personal information about you without your awareness.
 - Geographical origin¹³
 - Gender and age¹³

8 Item was measured on a 5-point Likert scale: Not at all interested, Slightly interested, Somewhat interested, Quite interested, Very interested

9 Video clip available here: https://youtu.be/Gr22YqS1_VA.

10 Item was measured on a 5-point Likert scale: Not at all, Slightly, Somewhat, Quite well, Very well

11 Item was measured on a 5-point Likert scale: Not at all worried, Slightly worried, Somewhat worried, Quite worried, Very worried

12 Item was measured on a 5-point Likert scale: Very uncommon, Somewhat uncommon, Undecided, Somewhat common, very common

13 Item was measured on a 5-point Likert scale: Unconcerned, Slightly concerned, Somewhat concerned, Quite concerned, Very concerned

- Mental health information¹³
 - Physical health information¹³
 - Level of intoxication¹³
 - Moods and emotions¹³
 - Sleepiness and fatigue¹³
 - Personality traits¹³
15. Previously in this survey, you were asked about your interest in voice-controlled virtual assistants, such as Apple's Siri and Amazon Alexa. You are now asked about this a second time. Please answer based on your current thoughts and feelings, independent from your previous response.
 - Are you interested in starting or continuing to use a voice-controlled virtual assistant?⁸
 16. Please indicate the extent to which you agree or disagree with each statement.¹⁴
 - Other people come to me for advice on new technologies⁵
 - In general, I am among the first in my circle of friends to acquire new technology when it appears⁵
 - I can usually figure out new high-tech products and services without help from others⁵
 - I keep up with the latest technological developments in my areas of interest⁵
 - I enjoy the challenge of figuring out high-tech gadgets⁵
 - I find I have fewer problems than other people in making technology work for me⁵
 - I prefer to use the most advanced technology available⁵
 - Show that you are paying attention by skipping this row without making a tick⁵
 - I find new technologies to be mentally stimulating⁵
 - Learning about technology can be as rewarding as the technology itself⁵
 17. Do you own any devices that have a microphone? Select all that apply.
 - Phone/smartphone
 - Laptop
 - Tablet
 - Smartwatch
 - Camera
 - Smart speaker
 - Car with voice control interface
 - Voice-enabled remote control

14 This construct – Innovativeness (INNO) – is adapted from Parasuraman and Colby [76]

3. User Knowledge and Perceptions about Sensor-based Inference Attacks

- Other (please specify)
- 18. What is the highest level of education you have obtained?
 - Finished school with no qualifications
 - Still in secondary school
 - GCSE Level education (e. g., GCSE, O-Levels, Standards)
 - A-Level education (e. g., A, AS, S-Levels, Highers)
 - Some undergraduate education (i. e., university examinations but not completed degree)
 - Degree or Graduate education (e. g., BSc, BA)
 - Post-graduate education (e. g., MSc, MA)
 - Doctorate degree
 - Vocational education (e. g., NVQ, HNC, HND)
 - Other degree or qualification (please specify)
- 19. Do you have professional experience in the following areas?
 - Data protection law¹⁵
 - Computer science¹⁵
 - Data mining¹⁵
 - Information technology security (IT security)¹⁵
- 20. What is your monthly net income? (Net income is defined as your total income after tax and social security deductions.)
 - I do not have a personal income
 - Less than £250
 - £250 up to £500
 - £500 up to £1000
 - £1000 up to £1500
 - £1500 up to £2000
 - £2000 up to £3000
 - £3000 up to £4000
 - £4000 up to £5000
 - £5000 or more
 - Decline to answer
- 21. To show if we have expressed ourselves clearly enough, please tick the description that best reflects the topic of this study.
 - Health effects of urban air pollution
 - Privacy concerns related to voice recordings
 - Telecommunications in India
 - Professional music production
 - Health concerns about wireless device radiation
 - Landlord and tenant privacy rights

¹⁵ Item was measured on a 5-point Likert scale: No experience, Little experience, Some experience, Much experience, Very much experience

B Correlation table

Spearman’s rank correlations between participant demographics and participants awareness for audio-based inferences are shown in Table 6, along with Bonferroni-corrected significance levels.

Table 6. Spearman’s rank correlations between participant demographics and awareness for audio-based inferences

| | AW_DEM | 95% CI | AW_STATE | 95% CI | AW_TRAIT | 95% CI |
|-----------|----------|------------------|----------|------------------|----------|------------------|
| Age | −0.24*** | (−1.00 to −0.19) | −0.18*** | (−1.00 to −0.12) | −0.19*** | (−1.00 to −0.14) |
| Gender | 0.05 | (−0.03 to 0.12) | 0.11* | (0.03 to 0.19) | 0.11* | (0.04 to 0.18) |
| Income | 0.08 | (0.00 to 0.17) | 0.06 | (−0.03 to 0.14) | 0.04 | (−0.04 to 0.12) |
| Education | 0.29*** | (0.23 to 1.00) | 0.17*** | (0.10 to 1.00) | 0.14** | (0.07 to 1.00) |
| INNO | 0.24*** | (0.18 to 1.00) | 0.18*** | (0.12 to 1.00) | 0.21*** | (0.15 to 1.00) |
| PA | 0.18*** | (0.12 to 1.00) | 0.20*** | (0.14 to 1.00) | 0.19*** | (0.12 to 1.00) |
| DVP | 0.08 | (0.01 to 0.17) | 0.05 | (−0.02 to 0.13) | 0.09 | (0.03 to 0.17) |
| PPE | 0.19*** | (0.12 to 1.00) | 0.18*** | (0.11 to 1.00) | 0.18*** | (0.12 to 1.00) |
| EXP_DP | 0.20*** | (0.14 to 1.00) | 0.16*** | (0.10 to 1.00) | 0.14** | (0.07 to 1.00) |
| EXP_DM | 0.28*** | (0.22 to 1.00) | 0.26*** | (0.20 to 1.00) | 0.28*** | (0.22 to 1.00) |
| EXP_CS | 0.27*** | (0.21 to 1.00) | 0.22*** | (0.16 to 1.00) | 0.26*** | (0.20 to 1.00) |
| EXP_IS | 0.25*** | (0.18 to 1.00) | 0.22*** | (0.15 to 1.00) | 0.22*** | (0.14 to 1.00) |

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

AW_: average level of awareness about audio-based inference of demographics (DEM) / short- and medium-term states (STATE) / physical and psychological traits (TRAIT); **INNO**: innovativeness; **PA**: privacy awareness; **DVP**: disposition to value privacy; **PPE**: previous privacy experience; **EXP_**: professional experience in data protection law (DP) / data mining (DM) / computer science (CS) / IT security (IS); **CI**: confidence interval

4

Special Cases and Application Areas of Sensor-based Inferences

4.1 Background: The Spies in Our Pockets

In recent years, the topic of cellphone surveillance has been widely covered in the press (e.g., [232, 233, 234, 235] and on social media (e.g., [236, 237, 238, 239]). While cellphone surveillance has many faces and can serve many different purposes [240], there is one topic that stands out in terms of curiosity and received attention – namely, the suspicion that companies are secretly recording people’s private conversations in order to derive profiling information for targeted advertising. Many people reportedly receive ads for things they just talked about – and some of them are sure that this is no coincidence [241]. Surveys in several countries have shown that a significant portion of the adult population believes that smartphones listen to their conversations without being prompted [241, 242, 243], including 55% of Americans [242] and 61% of Britons [241].

Various commentators have dismissed these fears as baseless paranoia and conspiracy theories [244, 245, 246]. But can we be certain that they are? There is no question that wiretapping is possible in principle. When dealing with certain threats and criminal offenses, law enforcement agencies around the world regularly tap suspects’ phones (often requiring authorization by a court) [240]. Intelligence agencies also make use of such spying methods [247, 248, 249], which may also be the reason why Edward Snowden – with all his insight into the NSA’s hacking powers – famously asked reporters to put their phones in a fridge before giving them an interview [250].

Suspicious about eavesdropping operations run by companies targeting the general population, on the other hand, may seem rather far-fetched. How could such large-scale attacks be technically feasible and economically viable, especially when factoring in the damage to the responsible companies’ reputations in case their operations would ever be disclosed to the public? This sceptical perspective is well represented in public discourse and often reflected in news headlines like “No, Your Phone Can’t Hear You” [251], “Facebook’s Not Listening

Through Your Phone” [252], “Why phones that secretly listen to us are a myth” [253] or “Your phone is not recording your conversations” [254].

Most of these articles do acknowledge reports from people who received advertisements that eerily matched contents of previous private conversations. However, they argue that these situations are either the result of pure coincidence, or the result of sophisticated tracking and profiling algorithms based on people’s online behavior – not the result of secret eavesdropping. This reasoning is plausible, especially when taking into account human cognitive biases. It is not hard to imagine how people leap to the illusory conclusion that they must be eavesdropped upon when, in fact, they have long revealed their interests and preferences to companies through other sources.

At the same time, with regard to the obscurities and scandals of today’s data economy, it would not be unreasonable to examine the opposite perspective on this issue with an open mind. Given the central role and immense value of personal data in various industries, there is no doubt that seemingly far-fetched vulnerabilities, methods, and tricks are often exploited to obtain this resource. Just as dirty wars are fought over access to crude oil, gold ore, coltan, and other resources, no avenue is left unexplored to reap personal data, which – for this very reason – has been dubbed “the new oil” [255], “the oil of the 21st century” [256], “the new black gold” [257], and even been described by *The Economist* as “world’s most valuable resource” [258]. Also, there are many influential players in the data economy whose names are not commonly known [2, 31, 259] and who, accordingly, do not really have an image or reputation to lose.

It is well documented that many mobile apps ruthlessly spy on their users [66, 260, 261] and that many app vendors are not particularly law-abiding or transparent when it comes to their collection and use of personal data (cf. Paper 7). Even technical sub-components developed by arbitrary third parties from around the world can also secretly harvest personal data of mobile phone users, potentially without the knowledge of the host app itself [262, 263]. Why then should it be assumed that all these parties, despite their often dubious practices and frequent misbehavior, carefully stay away from the taboo of secretly recording and analyzing people’s conversations?

While it is quite possible that people’s purported eavesdropping experiences are merely the result of coincidence or conventional profiling methods, the phenomenon at least deserves a serious and objective examination. There is a wealth of research on privacy leaks in iOS and Android apps, including several studies specifically trying to detect stealthy audio exfiltration by means of analyzing network traffic or a device’s power consumption and memory usage (cf. Section P5–5.2). However, there is no trace of scientific consensus that would satisfactorily answer the questions raised above. In the academic literature, no evidence-backed statements could be located that directly address the practicability and detectability of commercially motivated large-scale eavesdropping operations against smartphone users. A thorough search did not reveal a single publication that provides a general introduction to this issue, let alone an overview of existing arguments supporting the different positions and theories surrounding it.

Shortly after I noted this research gap, during the preparation of a literature review on the categories of personal information that can be inferred from accelerometer data (Paper 2),

an interesting connection between this project and smartphone-based eavesdropping emerged: While still inconclusive, there is research suggesting it may be possible to reconstruct words spoken by users based on sound vibrations captured by smartphone accelerometers (cf. Section P5–5.4). Given that accelerometers are present in a broad range of mobile devices and wearables [184], are often less protected than microphones [40], are relatively inconspicuous when active (in terms of power consumption) [264], and are regularly accessed by device vendors [43, 265], mobile apps [40], and even visited websites [266], this seems like a connection of potential significance.

In collaboration with Philip Raschke, a doctoral researcher and IT security expert at TU Berlin’s Department of Telecommunication Systems, I started gathering and analyzing information on the “my phone is listening in on my conversations” phenomenon. Our research focus was on the technical feasibility and detectability of such an attack, including a consideration of the role that accelerometers could play in this. Existing arguments on the issue were collected, structured and then critically examined for plausibility. While I was the main person in charge of the study, Philip Raschke was involved in all stages of the project, from its conception over the analysis and interpretation of the literature to the editing and critical revision of the manuscript. The paper that resulted from our work (Paper 5) was presented in July 2019 at the *33rd Conference on Data and Applications Security and Privacy* (DBSec’ 19) in Charleston, South Carolina.

While our paper is by no means the end of the story, but rather a testament to the broad lack of transparency within mobile applications and operating systems, it is – to the best of my knowledge – the first academic publication to holistically address this topic, and has also received encouraging feedback [267, 268]. Wolfie Christl, for example, a well-known Austrian privacy researcher and activist, has described our paper as “[t]he most comprehensive take on the question I know of” [269]. Our findings regarding the technical feasibility of smartphone-based eavesdropping attacks were recently supported by an elaborate investigation of the German public-service broadcaster *Bayrischer Rundfunk* [270, 271].

Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping

Jacob Leon Kröger^{1,2(✉)} and Philip Raschke¹

¹ Technische Universität Berlin, Berlin, Germany

{kroeger, philip.raschke}@tu-berlin.de

² Weizenbaum Institute for the Networked Society, Berlin, Germany

Abstract. Besides various other privacy concerns with mobile devices, many people suspect their smartphones to be secretly eavesdropping on them. In particular, a large number of reports has emerged in recent years claiming that private conversations conducted in the presence of smartphones seemingly resulted in targeted online advertisements. These rumors have not only attracted media attention, but also the attention of regulatory authorities. With regard to explaining the phenomenon, opinions are divided both in public debate and in research. While one side dismisses the eavesdropping suspicions as unrealistic or even paranoid, many others are fully convinced of the allegations or at least consider them plausible. To help structure the ongoing controversy and dispel misconceptions that may have arisen, this paper provides a holistic overview of the issue, reviewing and analyzing existing arguments and explanatory approaches from both sides. Based on previous research and our own analysis, we challenge the widespread assumption that the spying fears have already been disproved. While confirming a lack of empirical evidence, we cannot rule out the possibility of sophisticated large-scale eavesdropping attacks being successful and remaining undetected. Taking into account existing access control mechanisms, detection methods, and other technical aspects, we point out remaining vulnerabilities and research gaps.

Keywords: Privacy · Smartphone · Eavesdropping · Spying · Listening · Microphone · Conversation · Advertisement

1 Introduction

Smartphones are powerful tools that make our lives easier in many ways. Since they are equipped with a variety of sensors, store large amounts of personal data and are carried throughout the day by many people, including in highly intimate places and situations, they also raise various privacy concerns.

One widespread fear is that smartphones could be turned into remote bugging devices. For years, countless reports have been circulating on the Internet from people who claim that things they talked about within earshot of their phone later appeared in targeted online advertisements, leading many to believe that their private conversations must have been secretly recorded and analyzed.

The reported suspicious ads range across many product and service categories, including clothing, consumer electronics, foods and beverages, cars, medicines, holiday destinations, sports equipment, pet care products, cosmetics, and home appliances – and while some of these ads were described as matching an overall discussion topic, others allegedly promoted a brand or even a very specific product mentioned in a preceding face-to-face conversation [6, 12]. Some people claim to have experienced the phenomenon frequently and that they have successfully reproduced it in private experiments. Interestingly, many of the purported witnesses emphasize that the advertised product or service seems not related to places they have visited, terms they have searched for online, or things they have mentioned in text messages, emails or social media [6, 40]. Furthermore, some reports explicitly rate it as unlikely that the respective advertisements were selected by conventional targeting algorithms, as they lay notably outside the range of advertising normally received and did sometimes not even appear to match the person’s consumer profile (e.g. in terms of interests, activities, age, gender, or relationship status) [6, 41].

Numerous popular media outlets have reported on these alleged eavesdropping attacks [3]. In a Forbes article, for instance, the US-based market research company Forrester reports that at least 20 employees in its own workforce have experienced the phenomenon for themselves [40]. The same holds true for one in five Australians, according to a recent survey [38]. Even the US House Committee on Energy and Commerce has started to investigate the issue by sending letters to Google and Apple inquiring about the ways in which iOS and Android devices record private conversations [77].

Many commentators, including tech bloggers, researchers and business leaders, on the other hand, view the fear that private companies could target their ads based on eavesdropped conversations as baseless and paranoid. The reputational risk, it is argued, would be far too high to make this a viable option [76]. With regard to CPU, battery and data storage limitations, former Facebook product manager Antonio García Martínez even considers the alleged eavesdropping scenario to be economically and technically unfeasible [51]. As an alternative explanation for suspiciously relevant ads, he points to the many established and well-documented methods that companies successfully use to track, profile and micro-target potential customers. Yet another possible explanation states that the frequently reported phenomenon is merely a product of chance, potentially paired with some form of confirmation bias [41]. Finally, some commentators also suggest that topics of private conversations are sometimes inspired by unconsciously processed advertisements, which may later cause the perception of being spied upon when the respective ad is encountered again [28].

Many views, theories and arguments have been put forward in attempt to explain the curious phenomenon, including experimental results and positions from the research community. However, a consensus has not yet been reached, not even regarding the fundamental technical feasibility of the alleged eavesdropping attacks. Therefore, this paper reviews, verifies and compares existing arguments from both sides of the discourse. Apart from providing a structured overview of the matter, conclusions about the feasibility and detectability of smartphone-based eavesdropping are drawn based on existing research and our own analysis.

In accordance with the reports found on the phenomenon, this paper will focus on smartphones – specifically, iOS and Android devices. Since smartphones are the most widespread consumer electronics device, and since iOS and Android together clearly dominate the mobile OS market [70], this choice seems justified to us. However, most of the considerations in this paper are applicable to other types of mobile devices and other operating systems as well.

The remainder of this paper is structured as follows. In Sect. 2, we describe the underlying threat model, distinguishing between three possible adversaries. Section 3 examines the possibility of using smartphone microphones for stealthy eavesdropping, expanding on aspects of security permissions and user notifications. Similarly, Sect. 4 considers smartphone motion sensors as a potential eavesdropping channel, taking into account sampling frequency limits enforced by mobile operating systems. Section 5 then looks into the effectiveness of existing mitigation and detection techniques developed by Google, Apple, and the global research community. In Sect. 6, the ecosystem providers themselves are considered as potential adversaries. Section 7 evaluates the technical and economic feasibility of large-scale eavesdropping attacks. After that, Sect. 8 examines ways in which governmental and criminal hackers can compromise the speech privacy of smartphone users. Finally, Sect. 9 provides a discussion of analysis results, followed by a conclusion in Sect. 10.

2 Threat Model

To target advertisements based on smartphone eavesdropping, an organization A, who is responsible for selecting the audience for certain online ads (either the advertiser itself or a contractor entrusted with this task, such as an advertising network¹), needs to somehow gain access to sensor data² from the corresponding mobile device, or to information derived from the sensor data.

Initially, speech is recorded through the smartphone by an actor B, which could be either (1) the operating system provider itself, e.g. Apple or Google, (2) non-system apps installed on the device, or (3) third-party libraries³ included in these apps. Potentially after some processing and filtering, which can happen locally on the device or on remote servers, actor B shares relevant information extracted from the recording – directly or through intermediaries – with organization A (unless A and B are one and the same actor, which is also possible).

Organization A then uses the received information to identify the smartphone owner as a suitable target for specific ads and sends a corresponding broadcast request to an ad publisher (organization A could also publish the ads itself if it has access to ad distribution channels). Finally, the publisher displays the ads on websites or apps – either on the smartphone through which the speech was recorded or on other devices

¹ Advertising networks are companies that match demand and supply of online ad space by connecting advertisers to ad publishers. They often hold extensive amounts of data on individual internet users to enable targeted advertising [17].

² “sensor data” can refer to either audio recordings or motion sensor data (see Sects. 3, 4).

³ The role and significance of third-party apps will be further explained in Sect. 3.1.

that can be linked⁴ to the smartphone owner, for example through logins, browsing behavior, or IP address matching. The websites and apps on which the advertisements appear do not reveal who recorded the smartphone owner's speech. Not even organization A necessarily understands how and by whom the received profiling information was initially collected. For illustration, Fig. 1 presents a simplified overview of the threat model.

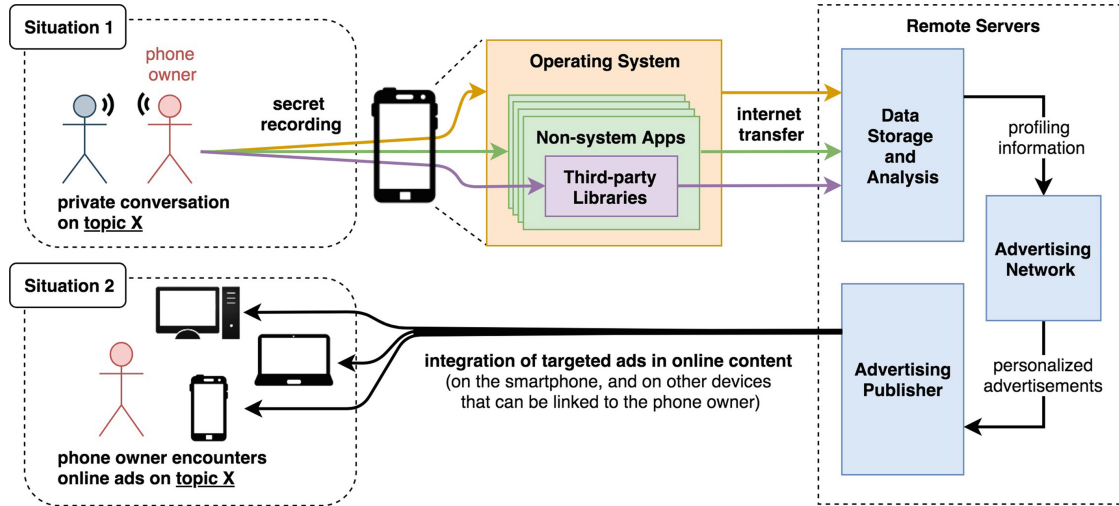


Fig. 1. A schematic and simplified overview of the threat model.

3 Microphone-Based Eavesdropping

Modern smartphones have the capability to tape any sort of ambient sound through built-in microphones, including private conversations, and to transmit sensitive data, such as the recording itself or information extracted from recorded speech, to remote servers over the Internet. Mobile apps installed on a phone could exploit these capabilities for secret eavesdropping. Aspects concerning app permissions and user notifications that could affect the feasibility and visibility of such an attack are examined in the following two subsections.

3.1 Microphone Access Permission

Before an app can access microphones in Android and iOS devices, permission has to be granted by the user. However, people tend to accept such requests blindly if they are interested in an app's functionality [10]. A survey of 308 Android users found that only 17% of respondents paid attention to permissions during app installation, and no more than 3% of the participants correctly answered the related comprehension questions [24].

⁴ For more information on cross-device tracking, refer to [65].

Encouraging app development at the expense of user privacy, current permission systems are much less strict than they were in early smartphones and have been criticized as “coarse grained and incomplete” [59]. Also, once a permission is granted, it is usually not transparent for users when and for which particular purpose data is being collected and to which servers it is being sent [62].

To include analytics and advertising capabilities, apps commonly make use of third-party libraries, i.e., code written by other companies. These libraries share multimedia permissions, such as microphone access, with their corresponding host app and are often granted direct Internet access [39]. Apart from the concern that third-party libraries are easily over-privileged, it is considered problematic that app developers often have limited or no understanding of the library code, which can also be changed dynamically at runtime [59]. Thus, not only users but also app developers themselves may be unaware of privacy leaks based on the abuse of granted permissions.

A large variety of existing apps has access to smartphone microphones. Examining over 17,000 popular Android apps, Pan et al. found that 43.8% ask for permission to record audio [59].

3.2 User Notifications and Visibility

Android and iOS apps with microphone permission can not only record audio at any time while they are active, i.e. running in the foreground, but also while they are in background mode, under certain conditions [7, 31]. Background apps have limited privileges and are often suspended to conserve the device’s limited resources. In cases, however, where they request the system to stay alive and continue recording while not in the foreground, there are ways to indicate this to the user.

In iOS, the status bar will automatically turn bright red when recording takes place in the background, allowing the user to immediately detect potentially unwanted microphone activity [8].

While the latest release of Android (version 9 Pie) implements similar measures [31], some older versions produce no visible indication when background apps access the microphone [10]. In this context, it might be worth noting that Android has been widely criticized for its slow update cycle, with hundreds of millions of devices running on massively outdated versions [56]. Also, quite obviously, notifications in the graphical user interface are only visible as long as the device’s screen is not turned off. And finally, some experimenters have already succeeded in circumventing the notification requirements for smartphone media recordings [69].

4 Motion Sensor-Based Eavesdropping

Adversaries might be able to eavesdrop on conversations through cell phones without accessing the microphone. Studies have shown that smartphone motion sensors – more specifically, accelerometers and gyroscopes – can be sensitive enough to pick up sound vibrations and possibly even reconstruct speech signals [36, 54, 79].

4.1 Experimental Research Findings

There are opposing views on whether non-acoustic smartphone sensors capture sounds at normal conversational loudness. While Anand and Saxena did not notice an apparent effect of live human speech on motion sensors in several test devices [3], other studies report very small but measurable effects of machine-rendered speech, significant enough to reconstruct spoken words or phrases [54, 79].

Using only smartphone gyroscopes, researchers from Israel’s defense technology group Rafael and Stanford University were able to capture acoustic signals rich enough to identify a speaker’s gender, distinguish between different speakers and, to some extent, track what was being said [54]. In a similar experiment, Zhang et al. demonstrated the feasibility of inferring spoken words from smartphone accelerometer readings in real-time, even in the presence of ambient noise and user mobility [79]. According to their evaluation, the achieved accuracies were comparable to microphone-based hotword detection applications such as Samsung S Voice and Google Now.

Both [79] and [54] have notable limitations. First of all, their algorithms were only able to detect a small set of predefined keywords instead of performing full speech recognition. Also, the speech in both experiments was produced by loudspeakers or phone speakers, which may result in acoustic properties different from live human speech. In [54], the playback device and the recording smartphone even shared a common surface, leading critics to suggest that the observed effect on sensor readings was not caused by aerial sound waves, but rather by direct surface vibrations [3]. Also, in contrast to Zhang et al., this approach only achieved low recognition accuracies, particularly for speaker-independent hotword detection. By their own admission, however, the authors of [54] are “security experts, not speech recognition experts” [32]. Therefore, the study should be regarded as an initial exploration rather than a perfect simulation of state-of-the-art spying techniques. With regard to the effectiveness of their approach, the researchers pointed out several possible directions for future improvement.

It might also be noteworthy that patents have already been filed for methods to capture acoustic signals through motion sensors, including a “method of detecting a user’s voice activity using an accelerometer” [21] and a “system that uses an accelerometer in a mobile device to detect hotwords” [55].

4.2 Sampling Frequency Limits

In order to limit energy consumption and because typical applications of smartphone motion sensors do not require highly sampled data, current mobile operating systems impose a cap on the sampling frequency of motion sensors, such as a maximum of 200 Hz for accelerometer readings in Android [3] and 100 Hz for gyroscopes in iOS [32]. For comparison, the fundamental frequency of the human speaking voice typically lies between 85 Hz and 155 Hz for men and 165 Hz and 255 Hz for women [79]. Thus, if at all, non-acoustic smartphone sensors can only capture a limited range of speech sounds, which presents a challenge to speech reconstruction attacks.

With the help of the aliasing effect explained in [54], however, it is possible to indirectly capture tones above the enforced frequency limits. Furthermore, experiments show that motion sensor signals from multiple co-located devices can be merged to obtain a signal with increased sampling frequency, significantly improving the effectiveness of speech reconstruction attacks [36]. Two or more smartphones that are located in proximity to each other and whose sensor readings are shared – directly or indirectly – with the same actor may therefore pose an increased threat to speech privacy.

It should also be noted that motion sensors in smartphones are usually capable of delivering much higher sampling frequencies (often up to 8 kHz) than the upper bounds prescribed by mobile operating systems [3]. Researchers already expressed concern that adversaries might be able to override and thereby exceed the software-based limits through patching applications or kernel drivers in mobile devices [3, 54].

4.3 Sensor Access Permissions and Energy Efficiency

While certain hardware components, such as camera, microphone and the GPS chip, are typically protected by permission mechanisms in mobile operating systems, motion sensors can be directly accessed by third-party apps in iOS and Android without any prior notification or request to the user [32, 45]. Thus, there is usually no way for smartphone owners to monitor, let alone control when and for what purposes data from built-in accelerometers and gyroscopes is collected. Even visited websites can often access smartphone motion sensors [32]. Exploiting accelerometers and gyroscopes to intrude user privacy is also much more energy-efficient and thus less conspicuous than recording via microphone [79].

5 Existing Mitigation and Detection Techniques

Many methods are applied by ecosystem providers and security researchers to screen mobile apps for vulnerabilities and malicious behavior. The following two subsections examine existing efforts with regard to their potential impact on the feasibility and detectability of mobile eavesdropping attacks.

5.1 App Inspections Conducted by Ecosystem Providers

Both iOS and Android apply a combination of static, dynamic and manual analysis to scan new and existing apps on their respective app market for potential security threats and to ensure that they operate as advertised [78]. Clearly, as the misbehavior of third-party apps can ultimately damage their own reputation, the platforms have strong incentives to detect and prevent abuse attempts.

Nevertheless, countless examples of initially undetected malware and privacy leaks have shown that the security screenings provided by Google and Apple are not always successful [19]. Google Play’s app inspection process has even been described as “fundamentally vulnerable” [29]. In a typical cat-and-mouse game, malicious apps evolve quickly to bypass newly implemented security measures [63], sometimes by

using “unbearably simple techniques” [29]. In Android devices from uncertified manufacturers, malware may even be pre-installed before shipment [14]. Significant vulnerabilities have also been found in official built-in apps. Apple’s FaceTime app, for example, allowed potential attackers to gain unauthorized access to iPhone cameras and microphones without any requirement of advanced hacking skills [15].

Leaving security loopholes aside, the existing security mechanisms do not guarantee privacy protection in terms of data minimization and transparency. Many mobile apps collect personal data with no apparent relevance to the advertised functionality [18, 62]. Even well-known apps like Uber have not been prevented from collecting sensitive user data that is not required for the service they offer [46].

There are also many documented cases of mobile apps using their microphone access in unexpected ways. An example that has received a lot of media attention recently is the use of so-called “ultrasonic beacons”, i.e. high-pitched Morse-style audio signals inaudible to the human ear which are secretly played in stores or embedded in TV commercials and other broadcast content in order to be able to unobtrusively track the location, activities and media consumption habits of consumers [10]. For this to work, the data subject needs to carry a receiving device that records and scans ambient sound for relevant ultrasonic signals and sends them back to the tracking network for automated comparison. A constantly growing number of mobile apps – several hundred already, some of them very popular – are using their microphone permission for exactly that purpose, often without properly informing the user about it [10, 47]. These apps, some of which are targeted at children and would not require audio recording for their core functionality, may even detect sounds while the phone is locked and carried in a pocket [47]. Even in cases where users are aware that their phone listens in, it is not clear to them what the audio stream is filtered for exactly and what information is being exfiltrated. Thus, the example of ultrasonic beacons shows how apps that have been approved into Apple’s App Store and Google Play can exploit their permissions for dubious and potentially unexpected tracking purposes.

Finally, it should not be overlooked that smartphone apps can also be obtained from various non-official sources, circumventing Apple’s and Google’s permission systems and auditing processes [62]. In Android, users are free in choosing the source of their applications [78]. Following a more restrictive policy, iOS only allows users to install apps downloaded from the official Apple App Store. However, kernel patches can be used to gain root access and remove software restrictions in iOS (“iOS jailbreaking”), which enables users to install apps from uncertified publishers [62].

5.2 App Inspections Conducted by the Research Community

In addition to the checks conducted by Google and Apple, mobile apps are being reviewed by a broad community of security and privacy researchers. A wide and constantly expanding range of manual and automated methods is applied for this purpose.

Pan et al., for instance, scanned 17,260 popular Android apps from different app markets for potential privacy leaks [59]. Through examining their media permissions, privacy policies and outgoing network flows, the researchers tried to identify apps that upload audio recordings to the Internet without explicitly informing the user about it.

While unveiling other serious forms of privacy violations, they found no evidence of such behavior. Based on these findings, the widely held suspicion of companies secretly eavesdropping on smartphone users was already portrayed as refuted in news headlines [34, 80].

However, the study comes with numerous limitations: Apart from considering only a small fraction of the over 2 million available Android apps, the researchers did not examine media exfiltration from app background activity, did not consider the use of privileged APIs, only tested a limited amount of each app’s functionalities for a short amount of time, used a controlled test environment with no real human interactions, did not consider iOS apps at all, and were not able to detect media that was intentionally obfuscated, encrypted at the application-layer, or sent over the network in non-standard encoding formats. Perhaps most importantly, Pan et al. were not able to rule out the scenario of apps transforming audio recordings into less detectable text transcripts or audio fingerprints before sending the information out. This would be a very realistic attack scenario. In fact, various popular apps are known to compress recorded audio in such a way [10, 33]. While all the choices that Pan et al. made regarding their experimental setup and methodology are completely understandable and were communicated transparently, the limitations do limit the significance of their findings. All in all, their approach would only uncover highly unsophisticated eavesdropping attempts.

Of course, many other researchers have also tried to detect privacy leaks in iOS and Android apps [62]. Besides analyzing decompiled code, permission requests and generated network traffic, other factors, such as battery power consumption and device memory usage, can also be monitored to detect suspicious app behavior [67]. Although some experts claim to have observed certain mobile apps recording and sending out audio with no apparent justification [58], the scientific community has not yet produced any hard evidence for large-scale eavesdropping through smartphone microphones.

Like the above-cited work by Pan et al., however, other existing methods to identify privacy threats in mobile devices also come with considerable limitations. Due to its closed-source nature, there is generally a lack of scalable tools for detecting malicious apps within iOS [19]. While, on the other hand, numerous efficient methods have been proposed for automatically scanning Android apps, none of these approaches is totally effective at detecting privacy leaks [59]. As with security checks of the official app stores (see Sect. 5.1), there is a wide range of possible obfuscation techniques and covert channels to circumvent detection mechanisms developed by the scientific community [10, 67]. Furthermore, many of the existing approaches do not indicate if detected data exfiltration activities are justified with regard to an app’s advertised functionality [62]. Yerukhimovich et al. even suggest that apps classified as safe or non-malicious are more likely to leak private information than typical “malware” [78].

Therefore, the fact that no evidence for large-scale mobile eavesdropping has been found so far should not be interpreted as an all-clear. It could only mean that it is difficult – under current circumstances perhaps even impossible – to detect such attacks effectively.

6 Ecosystem Providers as Potential Adversaries

Not only third-party apps but also mobile operating systems themselves can access privacy-sensitive smartphone data and transfer it over the Internet. It has been known for years that both, iOS and Android, do so extensively [5]. Examining the amount of data sent back to Google's and Apple's servers from test devices, a recent study found that iPhones – on average – received four requests per hour from their manufacturer during idle periods, and eighteen requests during periods of heavy use [68]. Leaving these numbers far behind, Android phones received forty hourly requests from Google when in idle state and ninety requests during heavy use. Of course, the number of requests per hour has only limited informational value. Data is often collected much more frequently, such as on a secondly basis or even constantly, to be later aggregated, compressed and sent out in data bundles [5].

While the establishment of network connections can be monitored, many aspects of data collection and processing in smartphones remain opaque. The source code of iOS is not made publicly available, and while Android is based on code from the Android Open Source Project, several of Google's proprietary apps and system components are closed-source as well [2]. Due to the resulting lack of transparency, it cannot be reliably ruled out that sensitive data is collected and processed without the will or knowledge of the smartphone owner – although, naturally, this would represent a considerable legal and reputational risk for the corresponding platform provider.

As an intermediary between applications and hardware resources, operating systems control the access to smartphone sensors, including microphones, accelerometers and gyroscopes, and can also decide whether or not sensor activity is indicated to the user on the device's screen. Other than with third-party apps, there is no superior authority in the system supervising the actions and decisions of iOS and Android. While external security experts can carry out inspections using similar methods as outlined in Sect. 5.2, they also face similar limitations. There is no reason to assume that operating systems refrain from using sophisticated obfuscation techniques to conceal their data collection practices. Additionally, being in control of the whole system, iOS and Android can access data on different levels of their respective software stack, which gives them more options for stealthy data exfiltration and could possibly impede detection.

7 Technical and Economic Feasibility

Even where adversaries manage to get around security measures and evade detection, it remains questionable whether a continuous and large-scale eavesdropping operation for the purpose of ad targeting would be technically feasible and economically viable. Based on estimations of CPU, battery, network transfer and data storage requirements, some commentators already stated their conclusion that such an operation would be far too expensive [51, 76] and may “strain even the resources of the NSA” [71]. Taking into account their underlying assumptions, these estimates appear valid. However, there are several ways in which smartphone-based eavesdropping could be made much more efficient and scalable, including:

- **Low quality audio recording.** To reduce the required data storage, processing power and energy consumption, adversaries could record audio at low bitrates. Speech signals do not even have to be intelligible to the human ear to be recognized and transcribed into text by algorithms [54].
- **Local pre-processing.** Some steps in the processing of recordings (e.g. transcription, extraction of audio features, data filtering, keyword matching, compression) can be performed locally on the device in order to transmit only the most relevant data to remote servers and thus reduce network traffic and required cloud storage.
- **Keyword detection instead of full speech recognition.** The amounts of processing power required for automatic speech recognition can be prohibitively high for local execution on mobile devices. A less CPU-intensive alternative to full speech recognition is keyword detection, where only a pre-defined vocabulary of spoken words is recognized. Such systems can even run on devices with much lower computational power than smartphones, such as 16-bit microcontrollers [25]. It has been argued that it would still be too taxing for mobile devices to listen out for the “millions or perhaps billions” of targetable keywords that could potentially be dropped in private conversations [51]. However, instead of listening for specific product and brand names, audio recordings can simply be scanned for trigger words that indicate a person’s interest, such as “love”, “enjoyed”, or “great”, in order to identify relevant snippets of the recording, which can then be analyzed in more depth. In fact, this very audio analysis method has already been patented, with the specific declared purpose of informing “targeted advertising and product recommendations” [22].
- **Selective recording.** Instead of recording continuously, an adversary could only record at selected moments using wake words or triggers based on time, location, user activity, sound level, and other context variables. This could significantly reduce the amount of required storage and network traffic [67].

Mobile apps that use all or some of the above techniques can be light enough to run smoothly on smartphones, as numerous commercial apps and research projects show [9, 10, 33, 58, 67].

But even if it is possible for companies to listen in on private conversations, some argue that this information might not be of much value to advertisers, since they would need to know a conversation’s context and speaker personalities very well in order to accurately infer personal preferences and purchase intentions from spoken phrases [51]. This argument is reasonable, but can equally be applied to many other profiling methods, including online tracking and location tracking, which are widely used nonetheless. Of course, where contextual information is sparse, such methods may lead to wrong conclusions about the respective data subject, possibly resulting in poor and inefficient ad targeting. However, this would not conflict with the above-mentioned reports of suspected eavesdropping: While the ads were perceived as inspired by topics raised in private conversations, they did not always reflect the purported witnesses’ actual needs and wants [6, 12].

From an outside perspective, it cannot be precisely determined how profitable certain types of personal data are for advertisers. It is therefore difficult, if not impossible, to draw up a meaningful cost-benefit calculation. However, it can generally

be assumed that private conversations contain a lot of valuable profiling information, especially when speakers express their interest in certain products or services. It is also worth mentioning that some of the world's largest companies earn a significant portion of their revenue through advertising – for Google and Facebook, this portion amounted to 85% and 98% in 2018, respectively [1, 23]. Profits from advertising can be considerably increased through effective targeting, which requires the collection of detailed personal information [68]. There is no doubt that smartphone sensor data can be very useful for this purpose. A recently filed patent describes, for example, how “local signals” from a mobile device, including motion sensor data and audio data from the microphone, can be analyzed to personalize a user's Facebook news feed [50].

8 Unauthorized Access to Smartphones

Although this is most likely no explanation for suspicious ad placement, it should be noted that there are many ways in which skilled computer experts or “hackers” can gain unauthorized access to mobile devices. The widespread use of smartphones makes them a particularly attractive hacking target [4].

Not only cyber criminals, but also law enforcement agencies and secret services invest heavily in their capabilities to exploit software flaws and other security vulnerabilities in consumer electronics [73]. It has been known for some time that intelligence agencies, such as NSA, GCHQ, and CIA, are equipped with tools to secretly compromise devices running iOS, Android and other mobile operating systems, enabling them “to move inside a system freely as if they owned it” [66, 75].

In addition to accessing sensitive data, such as geo-location, passwords, personal notes, contacts, and text messages, this includes the ability to turn on a phone's microphone without a user's consent or awareness [11]. With the help of specialized tools, smartphone microphones can even be tapped when the device is (or seems) switched off [73]. Such attacks can also be successful in high-security environments. In a recent case, for example, more than 100 Israeli servicemen had their phones infected with spyware that allowed unknown adversaries to control built-in cameras and microphones [57].

Besides the United States and some European nations, other developed countries, such as Russia, Israel and China, also have highly sophisticated spying technology at their disposal [75]. Less developed countries and other actors can buy digital eavesdropping tools from a flourishing industry of surveillance contractors at comparatively low prices [60]. That not only secret services but also law enforcement agencies in the US can be authorized to convert smartphones into “roving bugs” to listen in on private conversations has been confirmed in a 2012 court ruling [16]. Eavesdropping capabilities of criminal organizations should not be underestimated, either. According to a report by McAfee and the Center for Strategic and International Studies (CSIS), there are 20 to 30 cybercrime groups with “nation-state level” capacity in countries of the former Soviet Union alone [52].

9 Discussion

So far, despite significant research efforts, no evidence has been found to confirm the widespread suspicion that firms are secretly eavesdropping on smartphone users to inform ads. To the best of our knowledge, however, the opposite has not been proven either. While some threat scenarios (e.g. the constant transfer of uncompressed audio recordings into the cloud) can be ruled out based on existing security measures and considerations regarding an attack's visibility, cost and technical feasibility, there are still many security vulnerabilities and a fundamental lack of transparency that potentially leave room for more sophisticated attacks to be successful and remain undetected.

In comparison with the researchers cited in this paper, it can be assumed that certain companies have significantly more financial resources, more training data, and more technical expertise in areas such as signal processing, data compression, covert channels, and automatic speech recognition. This is – besides unresolved contradictions between cited studies and large remaining research gaps – another reason why existing work should not be seen as final and conclusive, but rather as an initial exploration of the issue.

While this paper focuses on smartphones, it should be noted that microphones and motion sensors are also present in a variety of other Internet-connected devices, including not only VR headsets, wearable fitness trackers and smartwatches, but also baby monitors, toys, remote controls, cars, household appliances, laptops, and smart speakers. Some of these devices may have weaker privacy safeguards than smartphones. For instance, they may not ask for user permission before turning on the microphone or may not impose a limit on sensor sampling frequencies. Numerous devices, including smart TVs [13], smart speakers [27], and connected toys [26], have already been suspected to spy on private conversations of their users. Certain smart home devices, such as home security alarms, may even contain a hidden microphone without disclosing it in the product specifications [44]. For these reasons, it is essential to also thoroughly examine non-smartphone devices when investigating suspicions of eavesdropping.

It is quite possible, at the same time, that the fears of advertising companies eavesdropping on private conversations are unfounded. Besides the widespread attribution to chance, one alternative approach to explaining strangely accurate advertisements points to all the established tracking technologies commonly employed by advertisers that do not depend on any phone sensors or microphones [51].

Drawing from credit card networks, healthcare providers, insurers, employers, public records, websites, mobile apps, and many other sources, certain multi-national corporations already hold billions of individual data points on consumers' location histories, browsing behaviors, religious and political affiliations, occupations, socioeconomic backgrounds, health conditions, personality traits, product preferences, and so on [17, 64]. Although their own search engines, social networks, email services, route planners, instant messengers, and media platforms already give them intimate insight into the lives of billions of people, advertising giants like Facebook and Google also intensively track user behavior on foreign websites and apps. Of the 17,260 apps examined in [59], for example, 48.22% share user data with Facebook in the

background. Through their analytics services and like buttons, Google and Facebook can track clicks and scrolls of Internet users on a vast number of websites [17].

The deep and potentially unexpected insights that result from such ubiquitous surveillance can be used for micro-targeted advertising and might thereby create an illusion of being eavesdropped upon, especially if the data subject is ill-informed about the pervasiveness and impressive possibilities of data linkage.

Even without being used for audio snooping, smartphones (in their current configuration) allow a large variety of actors to track private citizen in a much more efficient and detailed way than would ever have been possible in even the most repressive regimes and police states of the 20th century. At the bottom line, whether sensitive information is extracted from private conversations or collected from other sources does not make much difference to the possibilities of data exploitation and the entailing consequences for the data subject. Therefore, whether justified or not, the suspicions examined in this paper eventually lead to a very fundamental question: What degree of surveillance should be considered acceptable for commercial purposes like targeted advertising? Although this paper cannot offer an answer to this political question, it should not be forgotten that constant surveillance is by no means a technical necessity and that, by definition, democracies should design and regulate technology to primarily reflect the values of the public, not commercial interests.

Certainly, the fear of eavesdropping smartphones should never be portrayed as completely unfounded, as various criminal and governmental actors can gain unauthorized access to consumer electronics. Although such attacks are unlikely to result in targeted advertisement, they equally deprive the user of control over his or her privacy and might lead to other unpredictable harms and consequences. For example, digital spying tools have been used to infiltrate the smartphones of journalists [49] and human rights activists [60] for repressive purposes.

Finally, it should be recognized that – apart from the linguistic contents of speech – microphones and motion sensors may unexpectedly transmit a wealth of other sensitive information. Through the lens of advanced analytics, a voice recording can reveal a speaker’s identity [53], physical and mental health state [20, 37], and personality traits [61], for example. Accelerometer data from mobile devices may implicitly contain information about a user’s location [35], daily activities [48], eating, drinking and smoking habits [72, 74], degree of intoxication [30], gender, age, body features and emotional state [43] and can also be used to re-construct sequences of text entered into a device, including passwords [42].

10 Conclusion

After online advertisements seemingly adapted to topics raised in private face-to-face conversations, many people suspect companies to secretly listen in through their smartphones. This paper reviewed and analyzed existing approaches to explaining the phenomenon and examined the general feasibility and detectability of mobile eavesdropping attacks. While it is possible, on the one hand, that the strangely accurate ads were just a product of chance or conventional profiling methods, the spying fears were

not disproved so far, neither by device manufacturers and ecosystem providers nor by the research community.

In our threat model, we considered non-system mobile apps, third-party libraries, and ecosystem providers themselves as potential adversaries. Smartphone microphones and motion sensors were investigated as possible eavesdropping channels. Taking into account permission requirements, user notifications, sensor sampling frequencies, limited device resources, and existing security checks, we conclude that – under the current levels of data collection transparency in iOS and Android – sophisticated eavesdropping operations could potentially be run by either of the above-mentioned adversaries without being detected. At this time, no estimate can be made as to the probability and economic viability of such attacks.

References

1. Alphabet Inc.: Alphabet Announces Fourth Quarter and Fiscal Year 2018 Results (2019). https://abc.xyz/investor/static/pdf/2018Q4_alphabet_earnings_release.pdf?cache=adc3b38
2. Amadeo, R.: Google’s iron grip on Android: Controlling open source by any means necessary (2018). <https://arstechnica.com/gadgets/2018/07/googles-iron-grip-on-android-controlling-open-source-by-any-means-necessary/>
3. Anand, S.A., Saxena, N.: Speechless: analyzing the threat to speech privacy from smartphone motion sensors. In: 2018 IEEE Symposium on Security and Privacy, San Francisco, CA, pp. 1000–1017. IEEE (2018). <https://doi.org/10.1109/SP.2018.00004>
4. Aneja, L., Babbar, S.: Research trends in malware detection on Android devices. In: Panda, B., Sharma, S., Roy, N. (eds.) Data Science and Analytics. Communications in Computer and Information Science, vol. 799, pp. 629–642. Springer, Singapore (2018). https://doi.org/10.1007/978-981-10-8527-7_53
5. Angwin, J., Valentino-DeVries, J.: Apple, Google Collect User Data (2011). <https://www.wsj.com/articles/SB10001424052748703983704576277101723453610>
6. Anonymous: YouTube user demonstrating how Facebook listens to conversations to serve ads (2017). https://www.reddit.com/r/videos/comments/79i4cj/youtube_user_demonstrating_how_facebook_listens/
7. Apple: Background Execution. <https://developer.apple.com/library/archive/documentation/iPhone/Conceptual/iPhoneOSProgrammingGuide/BackgroundExecution/BackgroundExecution.html>
8. Apple: Record - iPhone User Guide. <https://help.apple.com/iphone/11/?lang=en#/iph4d2a39a3b>
9. Arcas, B.A., et al.: Now playing: continuous low-power music recognition. arXiv Comput. Res. Repos. abs/1711.10958 (2017). <http://arxiv.org/abs/1711.10958>
10. Arp, D., et al.: Privacy threats through ultrasonic side channels on mobile devices. In: 2017 IEEE European Symposium on Security and Privacy (EuroS&P), Paris, France, pp. 35–47. IEEE (2017). <https://doi.org/10.1109/EuroSP.2017.33>
11. Ball, J.: Angry Birds and “leaky” phone apps targeted by NSA and GCHQ for user data (2014). <https://www.theguardian.com/world/2014/jan/27/nsa-gchq-smartphone-app-angry-birds-personal-data>
12. BBC News Services: Is your phone listening in? Your stories (2017). <https://www.bbc.com/news/technology-41802282>

13. Beres, D.: How To Stop Your Smart TV From Eavesdropping On You (2015). https://www.huffpost.com/entry/your-samsung-tv-is-spying-on-you_n_6647762
14. Bocek, V., Chrysaidos, N.: Android devices ship with pre-installed malware (2018). <https://blog.avast.com/android-devices-ship-with-pre-installed-malware>
15. Bogost, I.: FaceTime Is Eroding Trust in Tech (2019). <https://www.theatlantic.com/technology/archive/2019/01/apple-facetime-bug-you-cant-escape/581554/>
16. Brown, A.J.: United States v. Oliva (United States Court of Appeals, D.C. No. 3:07-cr-00050-BR-1) (2012)
17. Christl, W.: Corporate Surveillance in Everyday Life. Cracked Labs, Vienna (2017)
18. Christl, W., Spiekermann, S.: Networks of Control: A Report on Corporate Surveillance, Digital Tracking, Big Data & Privacy. Facultas, Vienna (2016)
19. Cimitile, A., et al.: Machine learning meets iOS malware: identifying malicious applications on Apple environment. In: Proceedings of the 3rd International Conference on Information Systems Security and Privacy, Porto, Portugal, pp. 487–492. SciTePress (2017). <https://doi.org/10.5220/0006217304870492>
20. Cummins, N., et al.: Speech analysis for health: current state-of-the-art and the increasing impact of deep learning. *Methods* (2018). <https://doi.org/10.1016/j.ymeth.2018.07.007>
21. Dusan, S.V., et al.: System and Method of Detecting a User’s Voice Activity Using an Accelerometer (Patent No.: US9438985B2) (2014). <https://patents.google.com/patent/US9438985B2/en>
22. Edara, K.K.: Keyword Determinations from Voice Data (Patent No.: US20140337131A1) (2014). <https://patents.google.com/patent/US20140337131A1/en>
23. Facebook: Facebook Reports Fourth Quarter and Full Year 2018 Results. https://s21.q4cdn.com/399680738/files/doc_financials/2018/Q4/Q4-2018-Earnings-Release.pdf
24. Felt, A.P., et al.: Android permissions: user attention, comprehension, and behavior. In: Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS 2012), Washington, D.C. ACM Press (2012). <https://doi.org/10.1145/2335356.2335360>
25. Fourniols, J.-Y., et al.: An overview of basics speech recognition and autonomous approach for smart home IOT low power devices. *J. Signal Inf. Process.* **9**, 239–257. <https://doi.org/10.4236/jsip.2018.94015>
26. de Freytas-Tamura, K.: The Bright-Eyed Talking Doll That Just Might Be a Spy (2018). <https://www.nytimes.com/2017/02/17/technology/cayla-talking-doll-hackers.html>
27. Fussell, S.: Behind Every Robot Is a Human (2019). <https://www.theatlantic.com/technology/archive/2019/04/amazon-workers-eavesdrop-amazon-echo-clips/587110/>
28. Ganjoo, S.: Is Facebook secretly listening your conversations? New report says yes, security experts say no proof (2018). <https://www.indiatoday.in/technology/features/story/is-facebook-secretly-listening-your-conversations-new-report-says-yes-security-experts-say-no-proof-1255870-2018-06-09>
29. Gao, G., Chow, M.: Android Applications, Can You Trust Google Play on These. Tufts University (2016)
30. Gharani, P., et al.: An Artificial Neural Network for Gait Analysis to Estimate Blood Alcohol Content Level. *arXiv Comput. Res. Repos.* abs/1712.01691 (2017). <https://arxiv.org/abs/1712.01691>
31. Google: Android 9 Pie. <https://www.android.com/versions/pie-9-0/>
32. Greenberg, A.: The Gyroscopes in Your Phone Could Let Apps Eavesdrop on Conversations (2014). <https://www.wired.com/2014/08/gyroscope-listening-hack/>
33. Grosche, P., et al.: Audio content-based music retrieval. In: Müller, M., et al. (eds.) *Multimodal Music Processing. Dagstuhl Follow-Ups*. Dagstuhl Publishing, Wadern (2012)

34. Hale, J.L.: Does Your Smartphone Listen To You? A New Study Debunked This Common Conspiracy (2018). <https://www.bustle.com/p/does-your-smartphone-listen-to-you-a-new-study-debunked-this-common-conspiracy-9682413>
35. Han, J., et al.: ACComplice: location inference using accelerometers on smartphones. In: 2012 Fourth International Conference on Communication Systems and Networks (COMSNETS), pp. 1–9 (2012). <https://doi.org/10.1109/COMSNETS.2012.6151305>
36. Han, J., et al.: PitchIn: eavesdropping via intelligible speech reconstruction using non-acoustic sensor fusion. In: Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), pp. 181–192. ACM Press, Pittsburgh (2017). <https://doi.org/10.1145/3055031.3055088>
37. Hashim, N.W., et al.: Evaluation of voice acoustics as predictors of clinical depression scores. *J. Voice* **31**(2), 256.e1–256.e6 (2017). <https://doi.org/10.1016/j.jvoice.2016.06.006>
38. Hassan, B.: 1 in 5 Aussies convinced their smartphone is spying on them (2018). <https://www.finder.com.au/press-release-july-2018-1-in-5-aussies-convinced-their-smartphone-is-spying-on-them>
39. He, Y., et al.: Dynamic privacy leakage analysis of Android third-party libraries. In: 1st International Conference on Data Intelligence and Security (ICDIS), pp. 275–280 (2018). <https://doi.org/10.1109/ICDIS.2018.00051>
40. Khatibloo, F.: Is Facebook Listening (And So What If They Are)? (2017). <https://www.forbes.com/sites/forrester/2017/03/17/is-facebook-listening-and-so-what-if-they-are/>
41. Kleinman, Z.: Is your smartphone listening to you? (2016). <https://www.bbc.com/news/technology-35639549>
42. Kröger, J.: Unexpected inferences from sensor data: a hidden privacy threat in the internet of things. In: Strous, L., Cerf, V.G. (eds.) *Internet of Things. Information Processing in an Increasingly Connected World. IFIP Advances in Information and Communication Technology*, vol. 548, pp. 147–159. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-15651-0_13
43. Kröger, J.L., et al.: Privacy implications of accelerometer data: a review of possible inferences. In: Proceedings of the 3rd International Conference on Cryptography, Security and Privacy (ICCSP). ACM, New York (2019). <https://doi.org/10.1145/3309074.3309076>
44. Lee, D.: Google admits error over hidden microphone (2019). <https://www.bbc.com/news/technology-47303077>
45. Liu, X., et al.: Discovering and understanding Android sensor usage behaviors with data flow analysis. *World Wide Web* **21**(1), 105–126 (2018). <https://doi.org/10.1007/s11280-017-0446-0>
46. Lomas, N.: Uber to end controversial post-trip tracking as part of privacy drive (2017). <http://social.techcrunch.com/2017/08/29/uber-to-end-controversial-post-trip-tracking-as-part-of-privacy-drive/>
47. Maheshwari, S.: That Game on Your Phone May Be Tracking What You’re Watching on TV (2017). <https://www.nytimes.com/2017/12/28/business/media/alphonso-app-tracking.html>
48. Mannini, A., et al.: Activity recognition using a single accelerometer placed at the wrist or ankle. *Med. Sci. Sports Exerc.* **45**(11), 2193–2203 (2013). <https://doi.org/10.1249/MSS.0b013e31829736d6>
49. Marczak, B., et al.: Hacking Team and the Targeting of Ethiopian Journalists (2014). <https://citizenlab.ca/2014/02/hacking-team-targeting-ethiopian-journalists/>
50. Marra, C.J., et al.: Ranking of News Feed in a Mobile Device Based on Local Signals (Pub. No.: US20170351675A1) (2017). <https://patents.google.com/patent/US20170351675A1/en>
51. Martínez, A.G.: Facebook’s Not Listening Through Your Phone. It Doesn’t Have To (2017). <https://www.wired.com/story/facebooks-listening-smartphone-microphone/>

52. McAfee: Net Losses: Estimating the Global Cost of Cybercrime. Center for Strategic and International Studies (CSIS), Washington, D.C. (2014)
53. McLaren, M., et al.: The 2016 speakers in the wild speaker recognition evaluation. In: Proceedings of the 16th Annual Conference of the International Speech Communication Association (INTERSPEECH), pp. 823–827 (2016). <https://doi.org/10.21437/Interspeech.2016-1137>
54. Michalevsky, Y., et al.: Gyrophone: recognizing speech from gyroscope signals. In: Proceedings of the 23rd USENIX Security Symposium, pp. 1053–1067 (2014)
55. Mohapatra, P., et al.: Energy-efficient, Accelerometer-based Hotword Detection to Launch a Voice-control System. (Patent No.: US20170316779A1) (2017). <https://patents.google.com/patent/US20170316779A1/en>
56. Morris, I.: Android Is Still Failing Where Apple’s iOS Is Winning (2018). <https://www.forbes.com/sites/ianmorris/2018/04/13/android-is-still-failing-where-apples-ios-is-winning/>
57. Naor, I.: Breaking The Weakest Link Of The Strongest Chain (2017). <https://securelist.com/breaking-the-weakest-link-of-the-strongest-chain/77562/>
58. Nichols, S., Morgans, J.: Your Phone Is Listening and it’s Not Paranoia (2018). https://www.vice.com/en_uk/article/wjbzzy/your-phone-is-listening-and-its-not-paranoia
59. Pan, E., et al.: Panoptispy: Characterizing Audio and Video Exfiltration from Android Applications. Proc. Priv. Enhanc. Technol. **2018**(4), 33–50 (2018). <https://doi.org/10.1515/popets-2018-0030>
60. Perlroth, N.: Governments Turn to Commercial Spyware to Intimidate Dissidents (2017). <https://www.nytimes.com/2016/05/30/technology/governments-turn-to-commercial-spyware-to-intimidate-dissidents.html>
61. Polzehl, T.: Personality in Speech. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-09516-5>
62. Quattrone, A.: Inferring Sensitive Information from Seemingly Innocuous Smartphone Data. The University of Melbourne (2016)
63. Rahman, M., et al.: Search rank fraud and malware detection in Google Play. IEEE Trans. Knowl. Data Eng. **29**(6), 1329–1342 (2017). <https://doi.org/10.1109/TKDE.2017.2667658>
64. Ramirez, E., et al.: Data Brokers. A Call for Transparency and Accountability. Federal Trade Commission, Washington, D.C. (2014)
65. Ramirez, R., et al.: Cross-Device Tracking: An FTC Staff Report. Federal Trade Commission, Washington, D.C. (2017)
66. Rosenbach, M., et al.: iSpy: How the NSA Accesses Smartphone Data (2013). <http://www.spiegel.de/international/world/how-the-nsa-spies-on-smartphones-including-the-blackberry-a-921161.html>
67. Schlegel, R., et al.: Soundcomber: a stealthy and context-aware sound trojan for smartphones. In: Proceedings of the Network and Distributed System Security Symposium (NDSS) (2011)
68. Schmidt, D.C.: Google Data Collection. Digital Content Next, New York (2018)
69. Sidor, S.: Exploring limits of covert data collection on Android: apps can take photos with your phone without you knowing (2014). <http://www.ez.ai/2014/05/exploring-limits-of-covert-data.html>
70. Statista: Global mobile OS market share in sales to end users from 1st quarter 2009 to 2nd quarter 2018. <https://www.statista.com/statistics/266136/global-market-share-held-by-smartphone-operating-systems/>
71. Stern, J.: Facebook Really Is Spying on You, Just Not Through Your Phone’s Mic (2018). <https://www.wsj.com/articles/facebook-really-is-spying-on-you-just-not-through-your-phones-mic-1520448644>

72. Tang, Q., et al.: Automated detection of puffing and smoking with wrist accelerometers. In: Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare. pp. 80–87 (2014)
73. Taylor, P.: Edward Snowden interview: “Smartphones can be taken over” (2015). <https://www.bbc.com/news/uk-34444233>
74. Thomaz, E., et al.: A practical approach for recognizing eating moments with wrist-mounted inertial sensing. In: Proceedings of the ACM International Conference on Ubiquitous Computing, pp. 1029–1040. ACM Press (2015). <https://doi.org/10.1145/2750858.2807545>
75. Timberg, C., et al.: WikiLeaks: The CIA is using popular TVs, smartphones and cars to spy on their owners (2017). https://www.washingtonpost.com/news/the-switch/wp/2017/03/07/why-the-cia-is-using-your-tvs-smartphones-and-cars-for-spying/?noredirect=on&utm_term=.c162373021c3
76. Triggs, R.: No, your smartphone is not always listening to you (2018). <https://www.androidauthority.com/your-phone-is-not-listening-to-you-884028/>
77. Tsukayama, H., Romm, T.: Lawmakers press Apple and Google to explain how they track and listen to users (2018). <https://www.washingtonpost.com/technology/2018/07/09/lawmakers-press-apple-google-explain-how-they-track-listen-users/>
78. Yerukhimovich, A., et al.: Can smartphones and privacy coexist? Assessing technologies and regulations protecting personal data on Android and iOS devices. MIT Lincoln Laboratory, Lexington, MA (2016). <https://doi.org/10.7249/RR1393>
79. Zhang, L., et al.: AccelWord: energy efficient hotword detection through accelerometer. In: Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys), pp. 301–315. ACM Press (2015). <https://doi.org/10.1145/2742647.2742658>
80. No, Phones Aren’t Listening to Your Conversations, but May Be Recording In-App Videos: Study (2018). <https://www.justandroid.net/2018/07/05/no-phones-arent-listening-to-your-conversations-but-may-be-recording-in-app-videos-study/>

Open Access This chapter is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, duplication, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, a link is provided to the Creative Commons license and any changes made are indicated.

The images or other third party material in this chapter are included in the work’s Creative Commons license, unless indicated otherwise in the credit line; if such material is not included in the work’s Creative Commons license and the respective action is not permitted by statutory regulation, users will need to obtain permission from the license holder to duplicate, adapt or reproduce the material.



4.2 Background: Video Games and Privacy

After Chapters 2, 3, and 4.1, which all focused on specific types of sensors and devices, this chapter will address an entire industry well-suited to further illustrate and examine the privacy threats posed by inferential analytics, namely the video game industry. The focus on this industry was chosen with regard to its extensive collection of user data and its large-scale reliance on sensor-equipped devices.

A scan of the literature revealed that, compared to other types of digital services, video games have received little attention in privacy research thus far (cf. Section P6–1). While social networking sites, search engines, and ad networks are regularly mentioned and discussed by privacy scholars, video games – despite their huge popularity among the general population – are largely absent in the discourse.

Rather than a deep dive into one particular aspect of video games, the most notable gap in the literature was the scarcity of foundational work illustrating and discussing the potential privacy impacts of the video game industry in general. For instance, my literature search did not reveal a publication that provides a structured and comprehensive overview of the types of data collected by video games. It also yielded no published study that illustrates and classifies the wealth of personal information that can be inferred from gaming data. Paper 6 contributes to filling this gap. It is the result of a collaboration with Philip Raschke from Technische Universität Berlin; Jessica Percy Campbell, a Ph.D. candidate in the ‘Big Data Surveillance Project’ at the University of Victoria; and the head of our research group at the Weizenbaum Institute, Dr. Stefan Ullrich. Philipp Raschke and I conceived the project, which was then headed by me. All collaborators provided assistance in analyzing the literature and in editing and critically revising the manuscript.

Given my previous work, the paper examines not only inferences that can be drawn from regular gameplay data (e.g., an avatar’s movements and actions within the game environment) but also inferences from sensor data captured by gaming equipment. For the latter, we re-used and synthesized knowledge compiled in the above publications on voice recordings (Paper 1), accelerometer data (Paper 2), and eye-tracking data (Paper 3).

Surveilling the Gamers: Privacy Impacts of the Video Game Industry

Jacob Leon Kröger^{a,b,*}, Philip Raschke^b, Jessica Percy Campbell^c,
Stefan Ullrich^{a,b}

^aTechnische Universität Berlin, Straße des 17. Juni 135, Berlin, Germany

^bWeizenbaum Institute for the Networked Society, Hardenbergstraße 32, Berlin, Germany

^cUniversity of Victoria, David Turpin Building A316, Victoria BC V8P 5C2, Canada

Abstract

With many million users across all age groups and income levels, video games have become the world's leading entertainment industry. Behind the fun experience they provide, it goes largely unnoticed that modern game devices pose a serious threat to consumer privacy. To illustrate the industry's potential for illegitimate surveillance and user profiling, this paper offers a classification of data types commonly gathered by video games. Drawing from patents and literature of diverse disciplines, we also discuss how patterns and correlations in collected gameplay data may leak additional information in ways not easily understood or anticipated by the user. This includes inferences about a user's biometric identity, age and gender, emotions, skills, interests, consumption habits, and personality traits. Based on these findings, we argue that video games need to be brought into the focus of privacy research and discourse. Considering the granularity and enormous scale of the data collection taking place, this industry deserves the same level of scrutiny as other digital services, such as search engines, dating apps, or social media platforms. The knowledge compiled in this paper can serve as a basis for privacy impact assessments, consumer education, and further research into the societal impact of video games.

Keywords:

Video game, Privacy, Surveillance, Behavioral analysis, Data mining, Inference

1. Introduction

Playing video games is an extremely popular leisure activity. With annual revenues of over \$116 billion, video games are the world's leading entertainment medium, producing twice the revenue of digital music and cinema movies combined [1].

*Corresponding author: Jacob Leon Kröger, kroeger@tu-berlin.de

In order to adapt to varying preferences and requirements on the demand side, companies in this fast-growing industry have always been interested in collecting data about individual users and their gaming behavior. When the first video games were commercialized in the 1970s and 80s, these efforts were limited to traditional data collection methods, such as direct observation and videotaping of game play sessions, interviews, and questionnaires [2]. Soon after, with the advent of the World Wide Web and more powerful computers, it became technically possible to monitor users from a distance. Almost all of today's gaming devices are designed to transfer behavioral data to remote servers over the Internet [3]. The emergence of new business models, including free-to-play video games, microtransactions, and in-game advertising, have added to the industry's interest in personal data collection and user profiling. Over the past few years, all major game companies have invested substantially in their behavioral analytics capabilities [4, 5].

Throughout a user-game relationship, which can extend over months and years and thousands of hours of play time, “every single action taken, every decision made, every communication” can be recorded [6], sometimes with dozens of parameters being captured per second [7]. In conjunction with large user bases, which comprise up to hundreds of millions of players [8], this continuous gathering results in enormous amounts of high-dimensional user data. Due to technological trends, such as virtual reality, location-based gaming, physiological sensing, and affective computing, gaming also increasingly involves voice, facial, heart rate, skin response, GPS, eye tracking, and gesture recognition data [9].

While there are many legitimate processing purposes (e.g., game customization, discovery of bugs and usability issues, cheat detection, balanced team matching), gaming data can also be used for less noble ends. For example, knowledge about a player's psychological traits and vulnerabilities can be exploited for highly personalized persuasion and to increase the manipulative effect of targeted advertising [10] – not only to spur artificial demand for real-life goods and services or to sway political opinions and beliefs,¹ but also to make players spend more time in a game and purchase premium content [14, 15, 16]. Certain susceptible users, commonly referred to as “whales” in the gaming industry, can be induced to spend exorbitant sums of real money on virtual items or upgrades, often amounting to several hundred times the expenses of the average player [3]. Other possible types of data misuse include arbitrary mass surveillance, identity theft, and all sorts of discrimination [10, 17, 18]. There are methods for computing a “financial risk factor” from gameplay behavior, for instance, based on which a user may be denied a loan or a credit line extension [19], or methods to assess “essential qualities” based on gameplay data in order to determine

¹As “transformative learning tools” which often cover aspects of human history, economy, geography, culture, technology, and war [11], video games can be intentionally designed to propagandize populations and influence users' political leanings, functioning as an “interactive influence medium” [12] or “radicalising medium” [13].

a player’s suitability for certain jobs [20]. Besides the ever-present possibility of unintended data leaks, game companies regularly share user data with third parties, such as gaming networks, data brokers, middleware and analytics providers, government institutions, and advertising platforms [7, 9, 10, 21] who have their own intentions and may employ the knowledge unethically as well.

For an informed debate about these threats and to determine appropriate safeguards, an in-depth understanding of data collection and usage practices in the video game industry is crucial. Beyond Martinovic et al. [21], Moon [10], Russell et al. [9], and Whitson & Simon’s special issue of *Surveillance & Society* [12], there has been a lack of foundational research on the topic in recent years.

To provide a common basis of understanding for lawmakers, practitioners, and researchers of diverse backgrounds, this paper provides an overview and classification of the data categories commonly collected by video games (Sect. 2). Addressing an important issue that has been largely ignored in privacy research so far, we also explore how modern data analysis methods can be used to infer personal information from hidden patterns and correlations in collected gaming data (Sect. 3). Drawing from published patents and experimental studies, we found that in-game behavior can reveal information about a user’s biometric identity (Section 3.1), age and gender (Sect. 3.2), emotions (Sect. 3.3), skills (Sect. 3.4), interests (Sect. 3.5), consumption habits (Sect. 3.6), and personality traits (Sect. 3.7). The privacy implications of the sensors embedded in game devices will be the focus of Sect. 4. We then provide a discussion in Sect. 5 and a reflection on the limitations of our study in Sect. 6, before we conclude the paper in Sect. 7.

2. Data Categories Collected by Video Game Companies

Any interaction with a modern gaming system can be recorded in time-stamped log files, resulting in a history of all actions taken by the user and all player-related events happening in the game [15, 22]. This includes attributes and qualities, such as duration, frequency, direction, strength, speed, or accuracy of a player’s in-game actions.

Besides manual input, a range of sensors is increasingly being employed in gaming, e.g., to record a user’s voice, gestures, heart rate, facial expressions, or current geographical location (cf. Sect. 4). Gaming systems can also gather information about a user’s specific hardware and software setup and often use tracking technologies, such as identifiers, tags, and cookies [9, 10]. Additionally, many games seek permission to access data from other applications on the same device or from a user’s social media profile, such as documents, personal details, emails, or contact lists [9, 23]. An overview of all these data categories, along with specific examples, is provided in Fig. 1.

For storage and analysis, video games typically transmit their collected data to remote servers over the Internet – a process that is not traceable for the ordinary user and commonly referred to as “telemetry” [24] or “ex situ data

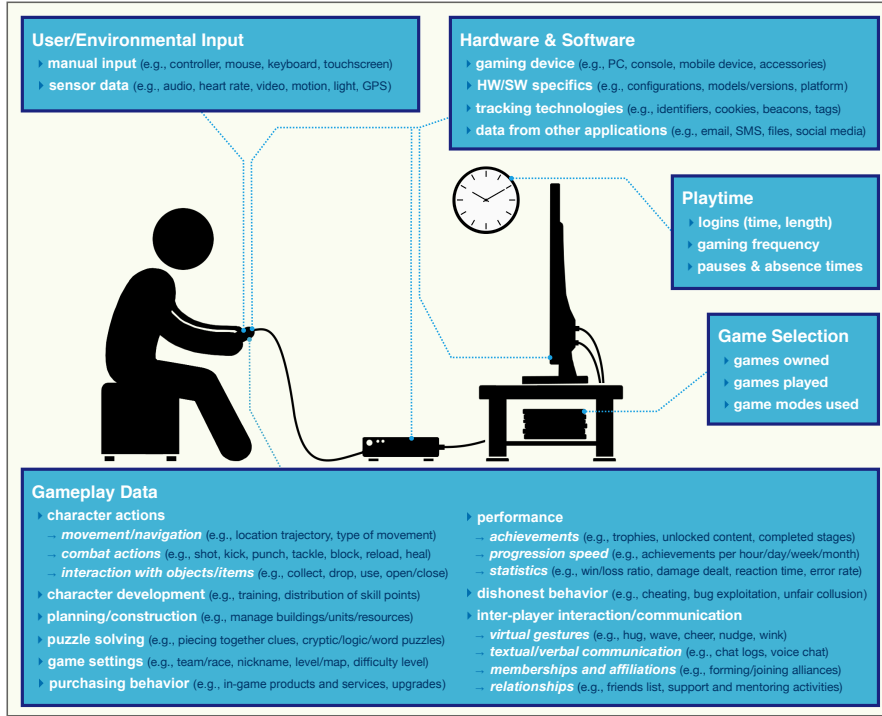


Figure 1: A classification of data types commonly collected by video games.

collection” [25]. Regarding the level of detail and granularity, it can be differentiated between shallow telemetry (i.e., collection of data on only a few behavioral variables) and deep telemetry (i.e., collection of data on all or a substantial fraction of the possible player behaviors) [26].

The specific metrics that are and can be captured naturally depend on the type of game and its underlying business model [7], but there is an overall trend towards more sophisticated video games which increases the variety of the data collected [27]. In many modern video games, “the level of granularity and completeness with which information is collected is unmatched by any real life experimental setup” [28]. One can distinguish between game-specific metrics (e.g., number of glowing bloatlflies killed in *Fallout 4*), genre-specific metrics (e.g., character progression in a role-playing game), and generic metrics which can be applied across game genres (e.g., total playtime) [7].

3. Inference of Personal Information from Gameplay Data

Apart from the information that is manually entered by the user or directly recorded (e.g., real name, birthdate, GPS location, chat logs), the data collected by video games can be mined for patterns and statistical relationships to

infer additional personal information. For this purpose, advanced data analysis methods are being applied in ways that may not be easily understood or anticipated by the user. With reference to published experimental research, this section presents and categorizes personal information that can be derived from gaming data – particularly from gameplay data, i.e., the player’s in-game behavior. In view of the vast variety of existing inference methods and the pace of technological progress, this overview is designed to be illustrative rather than comprehensive. Where available, relevant patents are referenced to exemplify corporate intentions and potential real-world applications.

3.1. User Identification

While even the tracking and profiling of anonymous users can be exploited for questionable purposes (e.g., invasive targeted advertising, price discrimination), a company’s ability to link gameplay behavior to a user’s real identity increases the potential for data misuse [21].

A user’s full name, email address, social media profile, postal address, credit card details, and other pieces of identity-related information are often entered during sign-up to a video game or gaming platform [7, 10]. Additionally, online game providers commonly employ tracking technologies, including web beacons, tags, browser fingerprinting, and cookies, enabling them to re-identify individual users and track their activities across different games, even when they are not logged in [9].

Users can also be identified based on characteristics of their individual playing style, such as their specific course of action in strategy games [21] or driving profile in racing games [29]. Besides, researchers found that cross-game tracking is often possible based solely on the analysis of player nicknames [30].

Apart from conventional tracking technologies and gameplay data, modern gaming devices increasingly capture potential biometric identifiers (e.g., voice, facial features, iris patterns, physical movements, body dimensions), as will be discussed in Sect. 4. Even typing rhythms and motion patterns recorded through basic input modalities like smartphone touchscreens [31], controller touch pads [32], computer mice [33], and physical keyboards [34] may be sufficient for user identification.

Due to the variety of existing user identification methods, it is extremely difficult, if not impossible, to guarantee for any gaming data that it is and will remain truly anonymous. Even in cases where the inference of a user’s real identity is not possible (e.g., due to limitations of the applied algorithm or because the target person is not registered in the recognition system database), other attributes derived from gaming data, such as age and gender (cf. Sect. 3.2), interests (cf. Sect. 3.5), socioeconomic status (cf. Sect. 3.6), and health condition (cf. Sect. 4) can still help to classify the target person into a specific demographic group and thereby approximate the identity.

3.2. Age and Gender Recognition

Just as name and address, video game players are often asked to provide their birthdate and gender during account registration [21]. Apart from that, there

are numerous approaches to infer such demographic attributes from playing behavior. For example, Likarish et al. [35] predicted the age of *World of Warcraft* players using 435 in-game features (e.g., character race, class, guild, level, faction, skills, achievements, and combat statistics), achieving a mean absolute error of ± 5 years for 53% of players. Using a similar set of features, Symborski et al. [36] inferred the self-reported gender identity of *Guild Wars* players with an accuracy of 83%.

Other game metrics that were identified as cues to age and/or gender include the tendency to play alone vs. joining multiplayer games [36, 37], the frequency of certain in-game activities (e.g., jumping, harvesting items, helping other players) [36, 37, 38], primary character gender and the number of selected male and female characters [21, 38], virtual-world language use (e.g., chat logs, character names) [36, 39, 40], the time spent playing certain game genres [21], and general play style (e.g., strategic player vs. social player) [36]. In addition to numerous approaches in the scientific literature, the inference of physical properties like gender and age from gameplay data has been incorporated in patents for over 15 years [41].

3.3. Emotion Recognition

Existing approaches for automatic emotion recognition are predominantly based on voice data [42], facial expressions or body language [43], or physiological data, such as heart rate and skin conductance [15], all of which are increasingly being captured by modern gaming technology (cf. Sect. 4). At the same time, however, there are many methods for deriving a person's affective state without using microphones, cameras, or biofeedback sensors. For example, it has long been proven possible to detect certain cognitive states of computer users, such as stress, by analyzing keyboard typing behavior [44] or cursor movements [45, 46]. Information about the emotions of a video game player can also be derived from playing characteristics, such as manner and direction of an avatar's movement, the types of weapons fired, objects destroyed, enemies killed, and items collected [47], or from a player's interaction with game dialogues, the frequency of game drop-outs, and overall performance metrics [46].

Behavioral patterns in interaction with a game can reflect a user's degree of engagement [46], level of motivation [45], emotional arousal and the valence of emotions (positive, negative, and neutral) [16] as well as more specific affective states such as fun, frustration, and the feeling of being challenged [47], distress, pride, shame, admiration, and reproach [48], anticipatory joy, hope, anxiety, anticipatory relief, and hopelessness [49], focus, curiosity and confusion [50], disappointment, boredom, interest, confidence, and satisfaction [46]. The affect detection model by Conati et al. in [48] even distinguishes between a player's emotions for the current state of the game, towards him/herself, and towards other characters in the game.

Emotion detection can be enhanced by incorporating information about the target's personality type [48, 50], some of which is inferable from gaming data as well (cf. Sect. 3.7). Attempts at analyzing the affective state of users based

on in-game behavior have been made by various game companies, including publishers of major commercial titles [15].

3.4. Skill Assessment

Video games are “problem solving spaces” [21] which usually require specific skills and abilities, such as strategic thinking, quick reflexes, aiming accuracy, multitasking, or eye-hand coordination [7, 24]. Two elements commonly found in video games are the repeated exposure of users to similar problems and the aspect of inter-player competition, which allow for multiple observations of a target behavior [51] and direct performance comparison between different players [24].

Extensive research efforts, including whole volumes dedicated to this topic, have established that the success of players in dealing with in-game tasks, puzzles, opponents, and other obstacles can be analyzed to assess their level of competence across a range of knowledge and skills [52, 53, 54]. In a widely used approach called “stealth assessment”, evaluation mechanisms are invisibly woven into a game’s environment to avoid the user being aware of the ongoing analysis [51, 55].

Some of the skills that can be assessed based on gameplay data are teamwork ability [19, 21, 54], language proficiency [19, 54, 56], financial investment skills [19], math fluency [24, 51, 54, 56], ICT skills [54, 57], creative problem solving [45, 51], spatial navigation [58], fine motor skills [51], metacognition and systems thinking [51], memory retention [10, 45, 59], cultural knowledge [41], and the understanding of specific science concepts, such as Newtonian mechanics [11, 60]. Approaches for game-based assessment can also allow to track a user’s cognitive development and learning trajectories over time [51, 61, 62, 63] and to examine specific gaps in knowledge [46, 63] or learning difficulties, such as reading problems and dyscalculia [61, 64].

With the technological and psychological foundations having long been in place, forms of stealth assessment are built into many of today’s commercial games [15, 51]. Patents in this field have existed for over ten years [19, 41].

3.5. Inference of Interests and Preferences

Since video gaming is a voluntary activity based on personal preferences, playing characteristics can allow insights into a player’s interests, likes, and dislikes [5, 37, 65]. Such inferences can not only be drawn from the type of gaming device used and the distribution of playtime across different games, but also from in-game behavior, such as the user’s allocation of budget to certain purposes (e.g., equipment, clothing, transportation), specific items collected and sold, targets of aggression, objectives pursued, modes of transportation used, team member selection, decisions made regarding character development, and patterns in social interaction with other players [19, 41, 66].

Among the user attributes that have been derived from gaming data are the proclivity for video gaming itself and the preference for certain games and game genres [15, 66, 67], game features (e.g., multiplayer vs. singleplayer mode) [67],

game design elements [65], and in-game activities (e.g., optimization, planning, trading, improvisation, imagining, co-operation) [68] as well as, for example, the preference for certain colors [41], car models [5, 19, 41], social relationships [69], sport and leisure activities [41], and types of financial investments [19]. Playing behavior may even reflect a user’s underlying basic desires or “life motives”, such as honor (i.e., the desire to obey traditional moral code), romance (i.e., the desire for courting and sex), acceptance (i.e., the desire for approval and to avoid criticism), independence (i.e., the desire for autonomy and self-reliance) [37], or the desire for social interaction and social achievement [65].

Emotion detection from gaming data, which includes the affective valence of a user’s reactions to specific stimuli inside the game (positive vs. negative) [16] and may thus assist in analyzing preferences and aversions, was discussed in Sect. 3.7 and will be addressed again – with a focus on sensor data – in Sect. 4.

3.6. Inferences about Financial Status and Consumption Behavior

Research has shown that economic behaviors of users in virtual worlds (e.g., collection and spending of in-game currency, trading of virtual goods and services, financial planning within a video game) resemble their real-world counterparts [21, 39], including even clandestine black-market activities [39]. Based on such correlations, game metrics can be indicative of a player’s financial status and consumption habits.

For example, a recently patented profiling method uses play traces to determine whether a user is frugal (e.g., indicated by saving in-game money even in the face of attractive spending options), fiscally responsible (e.g., indicated by investing carefully and focusing on strategically important purchases), or wasteful (e.g., indicated by taking financial risks, spending money quickly, and buying items not relevant to the goals of the game) [19]. The method also aims to evaluate whether a player is “trading-conscious”, i.e., fit for certain financial trading products, and to detect an “eagerness to go after new products or services” based on how players develop their in-game character.

Even non-financial aspects of a game can allow insights into a user’s money-management style. The above patent, for instance, proposes to assess a user’s level of frugality based on ammunition expenditure patterns in first-person shooter games (e.g., rate at which bullets are fired, percentage of hits, precision shots and controlled bursts vs. wasteful use of ammunition) or based on the user’s performance in driving games and flight simulators (e.g., aggressive driving, overspeed, crash frequency) [19].

Such links between gameplay and real-world spending behavior have also been reported in the scientific literature. Correlating the results of an online survey with log data from the popular sandbox video game *Minecraft*, for example, Canossa et al. [37] found that money-conscious players tend to build fewer sleeping accommodations for themselves and prefer to use cheap in-game materials, such as stone, sand, and iron instead of precious materials, such as diamond.

Besides in-game behavior and virtual consumption, it is common for game publishers to store actual payment information and purchase histories (e.g., when

the user buys games and upgrades online, or pays to unlock content) [9], which could increase their ability to estimate a user's economic proclivities [10]. Finally, even a user's set of gaming devices (e.g., cutting-edge game console, special equipment, high-end gaming PC) can be used as a cue to his or her financial standing [21].

3.7. Inference of Personality Traits

As with most aspects of human decision making and behavior (e.g., choice of literature, body language, leisure-time activities, decoration of personal space), personality traits play a central role in shaping how users respond to stimuli and experiences in virtual worlds [70, 71]. Therefore, even though players typically assume a fictional identity in video games – in terms of role (e.g., king, soldier, race driver), species (e.g., human, orc, elf), and other attributes (e.g., gender and age, body features, special abilities) – their individual playing styles often contain discernible traces of real-world personality [70, 72, 73].

By analyzing behavioral data from the multiplayer online games *Call of Duty* and *World of Warcraft*, Martinovic et al. [21] were able to assess certain character traits of players, including politeness (e.g., indicated by thanking and apologizing for in-game actions), leadership (e.g., indicated by remaining unchallenged in a leader role), defeatism (e.g., indicated by propensity to surrender early), disloyalty (e.g., indicated by tendency to betray own team mid-game), and punctuality (e.g., indicated by showing up in advance of scheduled games). A patent titled “Utilizing Gaming Behavior to Evaluate Player Traits” [19] comprises a method for inferring a player's untrustworthiness (e.g., indicated by dishonest behavior and cheating), aggressiveness (e.g., indicated by using excessive violence), goal orientation (e.g., indicated by actively pursuing specific tasks and objectives), patience (e.g., indicated by planning ahead and going after long-term goals), and risk aversion (e.g., indicated by avoiding unnecessary risks and challenges inside the game).

Other traits that have been correlated with and assessed based on game metrics include the tendency towards addiction [21], the disposition toward maximizing power versus security [68], tenacity and determination [74], self-confidence [19, 46], work ethic [73], and overall psychological maturity [19]. Furthermore, gameplay data has been used to evaluate users along the so-called Big Five personality factors, namely openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism [70, 72, 73, 75] and to evaluate subtypes of traits, e.g., different forms of curiosity (social curiosity, sensory curiosity, novelty-seeking curiosity, and explorative curiosity) [76].

It has been recognized by experts in the field that “most psychology experiments could be construed as video games” [77]. In fact, games can be specifically designed to expose certain personality traits. The first-person shooter game *America's Army*, for example, which was published by the U.S. Department of Defense as a platform for strategic communication and recruitment, can be used to specifically assess a player's “army values”, such as loyalty, duty, respect, selfless service, honor, integrity, and personal courage [51].

Experimental research has established that gameplay-personality associations can be sufficient to build a valid personality profile [70, 78]. The game-based recognition of personality traits is possible not only in binary form (high vs. low) but also in the form of numerical scores. In [73], for instance, the prediction result is presented using 100 possible scores along five personality dimensions. The effect size – i.e., the degree to which player personality is expressed in game behavior – was found to be “in line with those seen for professional, medical, and psychological applications of the MMPI, Big Five personality inventory, and Beck’s Hopelessness Scale” [73], which are standardized psychometric tests of adult personality and psychopathology. Researchers have even started to consider “whether a game is more suitable to predicting behaviors in a natural setting than a [conventional] personality test is” [78].

4. Sensor-based Inference of Personal Information

Modern game devices increasingly capture data from outside the game environment through a variety of embedded sensors. Some of the sensor-based technologies and features that are currently trending in the video game industry are eye tracking [10], emotion recognition [15], location-based gaming [79], physiological sensing [45, 74], body motion tracking [10, 15], and the combination of video gaming and physical exercise (“exergaming”) [80]. Some sensors are still being tested and explored for their applicability in video gaming (e.g., EEG, heart rate, skin conductance), whereas other sensors, such as cameras, microphones, GPS chips, and inertial motion sensors are already commonplace in off-the-shelf gaming devices. While sensors fulfill important functions and enable new forms of game interaction, they can unexpectedly reveal a large variety of sensitive personal information [42, 81, 82, 83, 84] and regularly collect data without the user’s knowledge [21, 81].

The precision of sensors found in gaming gear can be remarkable. Some commercially available video game systems (e.g., Xbox Kinect, Wii Balance Board) have already been confirmed as suitable instruments for diagnostic and functional assessment tasks in medical settings [85]. And in a way, even the most basic input devices can be seen as a proxy to gauge physiological measures because they implicitly capture characteristics of a user’s hand and body movements. For example, mouse clicks, keyboard keystrokes, and touchscreen taps can be analyzed to infer information about a user’s physical dexterity [19], state of health [4], emotions (cf. Sect. 3.3), and biometric identity (cf. Sect. 3.1). Illustrating the amount of detail obtainable from seemingly benign sensor data, there is a patented method that uses input from a simple touch pad to detect not only a user’s finger orientation and finger spacing, but also finger lengths and knuckle joint locations [32].

Naturally, this paper cannot cover in depth the whole diversity of sensors used in gaming. To exemplify the privacy implications, we decided to focus on three sensor types that are increasingly found in modern gaming devices and which we have thoroughly explored in previous work, namely accelerometers [81], microphones [42], and eye-tracking sensors [82]. Demonstrating the

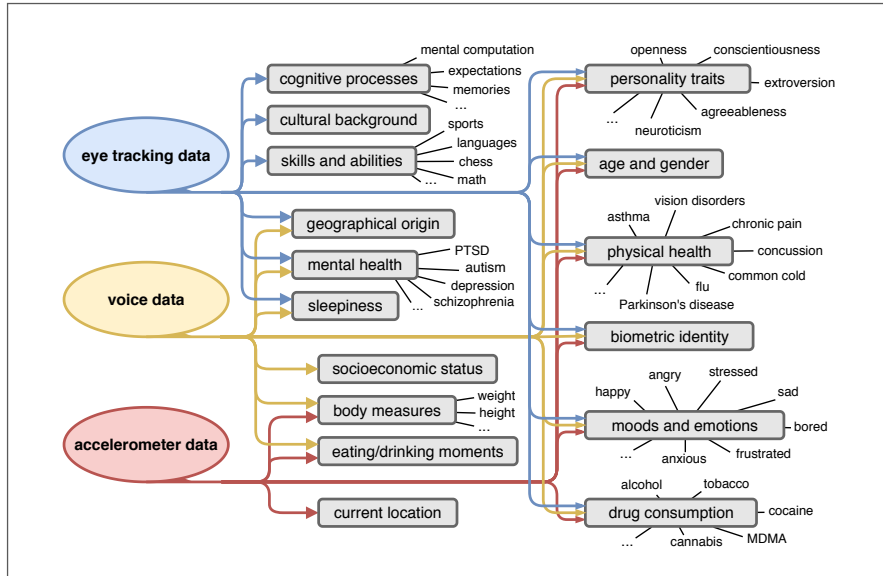


Figure 2: Overview of sensitive inferences that can be drawn from eye-tracking data [82], voice recordings [42], and accelerometer data [81].

richness and potential sensitivity of data from these sensors, Fig. 2 provides an overview of categories of personal information that can be inferred from accelerometer data, audio recordings, and eye-tracking data – representing a synthesis² of the findings from [42], [81], and [82]. For details and background information on these sensor types and inferences, please refer to the respective source. Several of the information categories shown in Fig. 2 did not yet appear in Sect. 3 (e.g., cognitive processes, cultural background, drug consumption, physical and mental health), suggesting that the growing use of sensors in entertainment electronics substantially increases the breadth of personal information discernible from video game data, going far beyond the information contained in traditional gameplay data.

5. Discussion and Implications

As we have explained and illustrated in this paper, video games can be used to collect a large variety of personal information about their users. With the help of advanced data analysis methods, patterns and correlations in gameplay and sensor data can be exploited to draw further inferences, e.g., about a user’s biometric identity, age and gender, emotions, skills, interests, socioeconomic

² From the three sources listed, we selected all inferences applicable to the context of video games.

status, consumption behavior, personality traits, physical and mental health condition, body measures, cultural and geographical background, drug habits, and cognitive processes – potentially much more information than a user wishes and expects to reveal. Considering the rapid developments in the area under investigation in recent years, as evident from the sources studied, we expect near-term discoveries of new threats and further improvements of existing inference methods in terms of speed and accuracy.

Not only the enormous volume and variety of the collected data, but also the high degree of experimental control held by the developers make game metrics so exceptionally sensitive and revealing. With respect to their unprecedented ability of placing millions of people in exactly replicated incentive environments, video games have been described by psychology and data science researchers as “rich natural laborator[ies]” [86], “ideal test bed[s] to collect and study data related to human behavior” [13], and “social engineering experiments that can generate a goldmine of behavioral data” [72]. Some forms of game-based assessment have already matched or even exceeded the accuracy of traditional psychological assessment methods, including self-reports [51, 64, 73, 87]. Above that, game mechanics can be intentionally designed to trigger the reactions and behaviors needed to analyze specific target attributes and qualities of the player [51, 88]. As a result, gaming data may allow intimate insights that are not obtainable from conventional sources of profiling information, such as a person’s browsing habits, loyalty card purchases, or credit history [19].

According to estimates, online video games may generate much more behavioral data than other Internet-based applications and services, including social media platforms [7, 77]. This seems plausible considering the typical amount and intensity of user-game interaction. Globally, gamers spend an average of over 28 hours each month playing, and around 20% of the population play more than 12 hours per week [89]. Also, unlike most other types of human-machine interaction, video gaming can be a deeply immersive experience integrating with a player’s sense of self [21, 46, 54] and may thus inhibit a rational consideration of potential risks and privacy implications.

As is clear from all the above, the video game industry urgently needs to be recognized and treated as a central issue in the discourse around consumer privacy, informational self-determination, and corporate surveillance – which is not currently the case. Users of video games deserve a high level of transparency around any collection, processing, and sharing of their personal data as well as effective protection measures against data access by unauthorized individuals or organizations.

In reality, these requirements are often far from being met. Many game publishers offer neither a sufficient explanation as to which of the data collected are really necessary for the functioning of the game, nor a simple way to opt out of non-essential data collection [7, 25]. Numerous companies in the video game industry, including market-leading players, have been involved in major data breaches [90], been criticized for being secretive about the data they collect and how this data is being used [7, 77, 91], and been accused of sharing personal data with third parties without a warrant and/or the user’s knowledge [21]. As

with many types of software, privacy policies of video games are widely written in ambiguous language and may omit important information [9].

Considering the complex plethora of data being collected by video game companies and the entertaining nature of their products, most users are likely neither motivated nor able to keep track of the ongoing data collection. Furthermore, as indicated by the persistence and prevalence of the nothing-to-hide argument [86], there still seems to be widespread ignorance about the serious risks that can arise from personal data being available to malicious or negligent parties. Thus, it can be questioned whether the doctrine of “informed consent”³ found at the core of even the most progressive data protection laws, such as EU’s GDPR [92], is appropriate and based on realistic assumptions, or whether more extensive forms of government intervention are needed to protect individuals from consequences of their own unawareness. Besides an obligation for companies to provide information on all personal data they collect *and infer*, this could mean restricting personal data usage for certain high-risk purposes irrespective of user consent [93].

In assessing the privacy impacts of video games and in the search for suitable protective measures, it should be considered that – while entertainment electronics appeal to people of all age groups – many game publishers market their products heavily towards minors who tend to be particularly unaware of privacy risks [9]. Furthermore, by putting players into a fictional environment without the immediate social context of real life, video games may give players a false sense of anonymity [10], making them “even more open to revealing their true self and thus [...] more vulnerable to prying eyes” [21].

6. Limitations

Being collected and applied in the service of corporate missions, gaming data and the algorithms used for data analysis are typically considered proprietary and not revealed to the public [5, 27]. Due to these confidentiality provisions, we cannot precisely assess the technological state of the art and current practices within the video game industry. Therefore, while being based on a broad range of valid empirical research, the overviews provided in Sect. 3 and 4 should be understood as an initial exploration of the respective issue, not as the upper bound of what is or may become technically feasible. It should further be noted that many of the cited inference methods were only tested on specific video game genres, individual games, or selected game components (e.g., [21, 29, 35, 36, 38, 75]), meaning that cross-game applicability of these methods remains largely unknown. Since researchers outside corporate laboratories rarely obtain direct access to large-scale user data collected by video game companies [27, 39], most of the cited studies also have relatively small sample sizes.

³In many jurisdictions, informed consent of the data subject is a legal basis for personal data processing. Under EU law, for example, a “freely given, specific, *informed* and unambiguous indication of the data subject’s wishes” is required for valid consent (Art. 4 GDPR). It can be questioned how often these legal requirements are really met in practice.

7. Conclusion

With users, developers, and the wider public being mainly focused on their features and entertainment value, it has been widely overlooked that video games constitute a substantial threat to consumer privacy. The overview provided in this paper illustrates that a wealth of potentially sensitive personal information can be collected and inferred from video game data. Our proposed data classification scheme is intended as a mutual reference point for readers of diverse backgrounds, and as a basic support tool for holistic privacy impact assessments. The example-rich sections on information inference will assist lawmakers, practitioners, and fellow researchers in further grasping the richness and potential sensitivity of gaming data.

Since the workings of data collection and data mining are completely invisible to ordinary video game users, it can be impossible for them to understand and control what information is revealed. Sophisticated surveillance and assessment mechanisms can be imperceptibly woven into the fabric of game environments and storylines. The immersive and distractive nature of video games may further impede a reasonable reflection on the staggering scope of the data harvesting taking place and on potential data misuses. Considering the immense and growing popularity of video gaming, consumer education in this field is urgently needed, along with effective technical and legal safeguards.

However, there still seems to be a long way to go. Various stakeholders of the video game industry are being criticized for a lack of transparency in data processing and have been involved in data scandals. Business models in the industry increasingly revolve around the harvesting and sharing of personal data. Under current circumstances, caution is definitely advisable. As with web browsing, users should not expect that their privacy will be protected or even respected when playing video games. At the same time, solving this problem cannot be left to the individual user. Only technology-savvy NGOs, research institutions, and governmental agencies are equipped to find sustainable solutions to this complex issue.

References

- [1] OppenheimerFunds, Investing in the Soaring Popularity of Gaming, 2018. URL: <https://www.reuters.com/sponsored/article/popularity-of-gaming>.
- [2] S. Santhosh, M. Vaden, Telemetry and Analytics Best Practices and Lessons Learned, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics: Maximizing the Value of Player Data*, Springer, London, 2013, pp. 85–109. URL: https://doi.org/10.1007/978-1-4471-4769-5_6.
- [3] T. V. Fields, Game Industry Metrics Terminology and Analytics Case Study, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics: Maximizing the Value of Player Data*, Springer, London, 2013, pp. 53–71. URL: https://doi.org/10.1007/978-1-4471-4769-5_4.

- [4] S. McCallum, J. Mackie, WebTics: A Web Based Telemetry and Metrics System for Small and Medium Games, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics: Maximizing the Value of Player Data*, Springer, London, 2013, pp. 169–193. URL: https://doi.org/10.1007/978-1-4471-4769-5_10.
- [5] R. Sifa, A. Drachen, C. Bauckhage, Profiling in Games: Understanding Behavior from Telemetry, in: K. Lakkaraju, G. Sukthankar, R. T. Wigand (Eds.), *Social Interactions in Virtual Worlds*, 1 ed., Cambridge University Press, 2018, pp. 337–374. URL: <https://doi.org/10.1017/9781316422823.014>.
- [6] S. Blickensderfer, N. A. Brown, Even the Games Have Eyes: Data Privacy and Gaming, 2019. URL: <https://www.natlawreview.com/article/even-games-have-eyes-data-privacy-and-gaming-podcast>.
- [7] A. Drachen, M. Seif El-Nasr, A. Canossa, Game Analytics – The Basics, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics*, Springer, London, 2013, pp. 13–40. URL: https://doi.org/10.1007/978-1-4471-4769-5_2.
- [8] List of most-played video games by player count, 2020. URL: https://en.wikipedia.org/w/index.php?title=List_of_most-played_video_games_by_player_count&oldid=976346186, page Version ID: 976346186.
- [9] N. C. Russell, J. R. Reidenberg, S. Moon, Privacy in Gaming, *Fordham Intellectual Property, Media & Entertainment Law Journal* (2020). URL: <https://doi.org/10.2139/ssrn.3147068>.
- [10] S. Moon, Privacy in Gaming and Virtual Reality Technologies: Review of Academic Literature 2012 – 2017, Technical Report, Center on Law and Information Policy (CLIP) at Fordham Law School, New York, NY, 2017. URL: https://www.fordham.edu/download/downloads/id/10331/privacy_in_gaming_and_virtual_reality_technologies_review_of_academic_literature_2012-2017.pdf.
- [11] V. J. Shute, F. Ke, Games, Learning, and Assessment, in: D. Ifenthaler, D. Eseryel, X. Ge (Eds.), *Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives*, Springer, New York, NY, 2012, pp. 43–58. URL: https://doi.org/10.1007/978-1-4614-3546-4_4.
- [12] J. R. Whitson, B. Simon, Game studies meets surveillance studies at the edge of digital culture: An introduction to a special issue on surveillance, games and play, *Surveillance & Society* 12 (2014) 309–319.
- [13] J. Ball, Xbox Live among game services targeted by US and UK spy agencies, 2013. URL: <https://www.theguardian.com/world/2013/dec/09/nsa-spies-online-games-world-warcraft-second-life>.

- [14] M. S. El-Nasr, *Game Analytics: Maximizing the Value of Player Data*, Springer, New York, 2013.
- [15] T.-H. D. Nguyen, Z. Chen, M. S. El-Nasr, Analytics-based ai techniques for a better gaming experience, in: S. Rabin (Ed.), *Game AI Pro 2: Collected Wisdom of Game AI Professionals*, volume 2, A K Peters/CRC Press, New York, NY, 2015, pp. 481–500.
- [16] G. N. Yannakakis, A. Paiva, Emotion in games, in: R. A. Calvo, S. D’Mello, J. Gratch, A. Kappas (Eds.), *The Oxford Handbook of Affective Computing*, Oxford University Press, New York, NY, 2014, pp. 459–471. Publisher: Oxford University Press.
- [17] W. Christl, *How Companies Use Data Against People*, Technical Report, Cracked Labs, Vienna, 2017.
- [18] Federal Bureau of Investigation, 2019 Internet Crime Report, 2020. URL: https://pdf.ic3.gov/2019_IC3Report.pdf.
- [19] E. J. Landers, B. G. Chun, S. B. Martin, M. H. Chang, A. Rao, T. H. Nguyen, D. S. Pelton, Utilizing gaming behavior to evaluate player traits, 2019. URL: <https://patents.google.com/patent/US10357713B1/en>.
- [20] Scoutible, 2020. URL: <https://www.scoutible.com>.
- [21] D. Martinovic, V. Ralevich, J. McDougall, M. Perklin, “You are what you play”: Breaching privacy and identifying users in online gaming, in: 2014 Twelfth Annual International Conference on Privacy, Security and Trust, IEEE, Toronto, 2014, pp. 31–39. URL: <https://doi.org/10.1109/PST.2014.6890921>.
- [22] A. Canossa, A. Drachen, Patterns of Play: Play-Personas in User-Centred Game Development, in: *Proceedings of the 2009 DiGRA International Conference*, Digital Games Research Association, London, 2009.
- [23] Pokémon GO Caught Millions of Players and Their Data, *Information Management* 50 (2016) 12.
- [24] C. S. Loh, Information Trails: In-Process Assessment of Game-Based Learning, in: D. Ifenthaler, D. Eseryel, X. Ge (Eds.), *Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives*, Springer, New York, NY, 2012, pp. 123–144. URL: https://doi.org/10.1007/978-1-4614-3546-4_8.
- [25] C. S. Loh, Y. Sheng, Measuring Expert Performance for Serious Games Analytics: From Data to Insights, in: C. S. Loh, Y. Sheng, D. Ifenthaler (Eds.), *Serious Games Analytics: Methodologies for Performance Measurement, Assessment, and Improvement*, *Advances in Game-Based Learning*, Springer, Cham, 2015, pp. 101–134. URL: https://doi.org/10.1007/978-3-319-05834-4_5.

- [26] A. Canossa, Interview with Nicholas Francis and Thomas Hagen from Unity Technologies, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics: Maximizing the Value of Player Data*, Springer, London, 2013, pp. 137–142. URL: https://doi.org/10.1007/978-1-4471-4769-5_8.
- [27] A. Drachen, C. Thureau, A Comparison of Methods for Player Clustering via Behavioral Telemetry, in: *Proceedings of the 8th International Conference on the Foundations of Digital Games*, Society for the Advancement of the Science of Digital Games, Chania, Crete, 2013.
- [28] K. J. Shim, N. Pathak, M. A. Ahmad, C. DeLong, Z. Borbora, A. Mahapatra, J. Srivastava, Analyzing human behavior from multiplayer online game logs: A knowledge discovery approach, *IEEE Intelligent Systems* 26 (2011) 85–89.
- [29] M. Kale, M. Bedekar, Driver Profiling Using Realistic Racing Games, in: *2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT)*, IEEE, Coimbatore, 2018, pp. 13–17. URL: <https://doi.org/10.1109/ICICCT.2018.8473154>.
- [30] M. A. Ahmad, C. Shen, J. Srivastava, N. Contractor (Eds.), *Predicting real world behaviors from virtual world data*, Springer, Cham, 2014. URL: <https://doi.org/10.1007/978-3-319-07142-8>.
- [31] T. Feng, Z. Liu, K.-A. Kwon, W. Shi, B. Carbunar, Y. Jiang, N. Nguyen, Continuous mobile authentication using touchscreen gestures, in: *2012 IEEE conference on technologies for homeland security (HST)*, IEEE, Waltham, MA, 2012, pp. 451–456.
- [32] B. Lukas, Q. Fu, S. Ranganathan, J. Karaoguz, T. W. Kwan, X. Yu, Hand-held gaming device that identifies user based upon input from touch sensitive panel, 2014. URL: <https://patents.google.com/patent/US8754746B2/en>.
- [33] B. Sayed, I. Traoré, I. Woungang, M. S. Obaidat, Biometric authentication using mouse gesture dynamics, *IEEE Systems Journal* 7 (2013) 262–274. URL: <https://doi.org/10.1109/JSYST.2012.2221932>, publisher: IEEE.
- [34] S. Maheshwary, S. Ganguly, V. Pudi, Deep secure: A fast and simple neural network based approach for user authentication and identification via keystroke dynamics, in: *IWAISe: First international workshop on artificial intelligence in security*, NUI Galway, Melbourne, 2017, pp. 59–67.
- [35] P. Likarish, O. Brdiczka, N. Yee, N. Ducheneaut, L. Nelson, Demographic Profiling from MMOG Gameplay, in: *11th Privacy Enhancing Technologies Symposium*, Springer, Waterloo, Canada, 2011, pp. 1–19.

- [36] C. Symborski, G. M. Jackson, M. Barton, G. Cranmer, B. Raines, M. M. Quinn, The Use of Social Science Methods to Predict Player Characteristics from Avatar Observations, in: M. A. Ahmad, C. Shen, J. Srivastava, N. Contractor (Eds.), Predicting Real World Behaviors from Virtual World Data, Springer, Cham, 2014, pp. 19–37. URL: https://doi.org/10.1007/978-3-319-07142-8_2, publisher: Springer, Cham.
- [37] A. Canossa, J. B. Martinez, J. Togelius, Give me a reason to dig Minecraft and psychology of motivation, in: 2013 IEEE Conference on Computational Intelligence in Games (CIG), IEEE, Niagara Falls, ON, 2013, pp. 1–8. URL: <https://doi.org/10.1109/CIG.2013.6633612>.
- [38] T. Kennedy, R. R. Ratan, K. Kapoor, N. Pathak, D. Williams, J. Srivastava, Predicting MMO Player Gender from In-Game Attributes Using Machine Learning Models, in: M. Ahmad, C. Shen, J. Srivastava, N. Contractor (Eds.), Predicting Real World Behaviors from Virtual World Data, Springer, Cham, 2014, pp. 69–84. URL: https://doi.org/10.1007/978-3-319-07142-8_5.
- [39] M. A. Ahmad, C. Shen, J. Srivastava, N. Contractor, On the Problem of Predicting Real World Characteristics from Virtual Worlds, in: M. A. Ahmad, C. Shen, J. Srivastava, N. Contractor (Eds.), Predicting Real World Behaviors from Virtual World Data, Springer Proceedings in Complexity, Springer, Cham, 2014, pp. 1–18. URL: https://doi.org/10.1007/978-3-319-07142-8_1.
- [40] A. Lawson, J. Murray, Identifying User Demographic Traits Through Virtual-World Language Use, in: M. A. Ahmad, C. Shen, J. Srivastava, N. Contractor (Eds.), Predicting Real World Behaviors from Virtual World Data, Springer Proceedings in Complexity, Springer, Cham, 2014, pp. 57–67. URL: https://doi.org/10.1007/978-3-319-07142-8_4.
- [41] D. Willis, Method and system for delivering advertising content to video games based on game events and gamer activity, 2006. URL: <https://patents.google.com/patent/US20060135232A1/en>.
- [42] J. L. Kröger, O. H.-M. Lutz, P. Raschke, Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference, in: M. Friedewald, M. Önen, E. Lievens, S. Krenn (Eds.), Privacy and Identity Management. Data for Better Living: AI and Privacy, Springer, Cham, 2020, pp. 242–258. URL: https://doi.org/10.1007/978-3-030-42504-3_16.
- [43] E. Novak, T. E. Johnson, Assessment of Student’s Emotions in Game-Based Learning, in: D. Ifenthaler, D. Eseryel, X. Ge (Eds.), Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives, Springer, New York, NY, 2012, pp. 379–399. URL: https://doi.org/10.1007/978-1-4614-3546-4_19.

- [44] L. M. Vizer, L. Zhou, A. Sears, Automated stress detection using keystroke and linguistic features: An exploratory study, *International Journal of Human-Computer Studies* 67 (2009) 870–886. URL: <https://doi.org/10.1016/j.ijhcs.2009.07.005>.
- [45] I. Ghergulescu, C. H. Muntean, Measurement and Analysis of Learner’s Motivation in Game-Based E-Learning, in: D. Ifenthaler, D. Eseryel, X. Ge (Eds.), *Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives*, Springer, New York, NY, 2012, pp. 355–378. URL: https://doi.org/10.1007/978-1-4614-3546-4_18.
- [46] E. E. Mattheiss, M. D. Kickmeier-Rust, C. M. Steiner, D. Albert, Approaches to Detect Discouraged Learners: Assessment of Motivation in Educational Computer Games, in: *Proceedings of eLearning Baltics (eLBa)*, Rostock, 2010, pp. 1–10.
- [47] N. Shaker, G. Yannakakis, J. Togelius, Towards Automatic Personalized Content Generation for Platform Games, in: *6th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE*, Stanford, CA, 2010, pp. 63–68.
- [48] C. Conati, H. Maclaren, Empirically Building and Evaluating a Probabilistic Model of User Affect, *User Modeling and User-Adapted Interaction* 19 (2009) 267–303. URL: <https://doi.org/10.1007/s11257-009-9062-8>.
- [49] K. Muñoz, P. M. Kevitt, T. Lunney, J. Noguez, L. Neri, An emotional student model for game-play adaptation, *Entertainment Computing* 2 (2011) 133–141. URL: <https://doi.org/10.1016/j.entcom.2010.12.006>.
- [50] J. Sabourin, B. Mott, J. C. Lester, Modeling Learner Affect with Theoretically Grounded Dynamic Bayesian Networks, in: S. D’Mello, A. Graesser, B. Schuller, J.-C. Martin (Eds.), *Affective Computing and Intelligent Interaction*, volume 6974, Springer, Berlin, 2011, pp. 286–295. URL: https://doi.org/10.1007/978-3-642-24600-5_32, series Title: *Lecture Notes in Computer Science*.
- [51] J. L. Plass, B. D. Homer, C. K. Kinzer, Y. K. Chang, J. Frye, W. Kaczetow, K. Isbister, K. Perlin, Metrics in Simulations and Games for Learning, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics: Maximizing the Value of Player Data*, Springer, London, 2013, pp. 697–729. URL: https://doi.org/10.1007/978-1-4471-4769-5_31.
- [52] D. Ifenthaler, D. Eseryel, X. Ge (Eds.), *Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives*, Springer, New York, NY, 2012. URL: <https://doi.org/10.1007/978-1-4614-3546-4>.
- [53] C. S. Loh, Y. Sheng, D. Ifenthaler (Eds.), *Serious games analytics: methodologies for performance measurement, assessment, and improvement*, *Advances in Game-Based Learning*, Springer, Cham, 2015. URL: <https://doi.org/10.1007/978-3-319-05834-4>.

- [54] M. C. Mayrath, J. Clarke-Midura, D. H. Robinson (Eds.), *Technology-Based Assessments for 21st Century Skills: Theoretical and Practical Implications from Modern Research*, Information Age Publishing, Charlotte, NC, 2012.
- [55] V. J. Shute, *Stealth Assessment in Computer-Based Games to Support Learning*, *Computer Games and Instruction* 55 (2011) 503–524.
- [56] A. Canossa, Interview with Simon Egenfeldt Nielsen from Serious Games Interactive, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics: Maximizing the Value of Player Data*, Springer, London, 2013, pp. 763–766. URL: https://doi.org/10.1007/978-1-4471-4769-5_33.
- [57] R. J. Mislevy, J. T. Behrens, K. E. Dicerbo, D. C. Frezzo, P. West, *Three Things Game Designers Need to Know About Assessment*, in: D. Ifenthaler, D. Eseryel, X. Ge (Eds.), *Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives*, Springer, New York, NY, 2012, pp. 59–81. URL: https://doi.org/10.1007/978-1-4614-3546-4_5.
- [58] F. Ke, V. Shute, *Design of Game-Based Stealth Assessment and Learning Support*, in: C. S. Loh, Y. Sheng, D. Ifenthaler (Eds.), *Serious Games Analytics*, Springer, Cham, 2015, pp. 301–318. URL: https://doi.org/10.1007/978-3-319-05834-4_13.
- [59] K. E. DiCerbo, M. Bertling, S. Stephenson, Y. Jia, R. J. Mislevy, M. Bauer, G. T. Jackson, *An Application of Exploratory Data Analysis in the Development of Game-Based Assessments*, in: C. S. Loh, Y. Sheng, D. Ifenthaler (Eds.), *Serious Games Analytics: Methodologies for Performance Measurement, Assessment, and Improvement*, *Advances in Game-Based Learning*, Springer, Cham, 2015, pp. 319–342. URL: https://doi.org/10.1007/978-3-319-05834-4_14.
- [60] E. Rowe, J. Asbell-Clarke, R. S. Baker, *Serious Games Analytics to Measure Implicit Science Learning*, in: C. S. Loh, Y. Sheng, D. Ifenthaler (Eds.), *Serious Games Analytics: Methodologies for Performance Measurement, Assessment, and Improvement*, *Advances in Game-Based Learning*, Springer, Cham, 2015, pp. 343–360. URL: https://doi.org/10.1007/978-3-319-05834-4_15.
- [61] B. Csapó, A. Lörincz, G. Molnár, *Innovative Assessment Technologies in Educational Games Designed for Young Students*, in: D. Ifenthaler, D. Eseryel, X. Ge (Eds.), *Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives*, Springer, New York, NY, 2012, pp. 235–254. URL: https://doi.org/10.1007/978-1-4614-3546-4_13.
- [62] D. Ifenthaler, D. Eseryel, X. Ge, *Assessment for Game-Based Learning*, in: D. Ifenthaler, D. Eseryel, X. Ge (Eds.), *Assessment in Game-Based Learning: Foundations, Innovations, and Perspectives*, Springer, New York, NY, 2012, pp. 1–8. URL: https://doi.org/10.1007/978-1-4614-3546-4_1.

- [63] R. D. Myers, T. W. Frick, Using Pattern Matching to Assess Gameplay, in: C. S. Loh, Y. Sheng, D. Ifenthaler (Eds.), *Serious Games Analytics: Methodologies for Performance Measurement, Assessment, and Improvement*, Advances in Game-Based Learning, Springer, Cham, 2015, pp. 435–458. URL: https://doi.org/10.1007/978-3-319-05834-4_19.
- [64] S. Klingler, T. Käser, A.-G. Busetto, B. Solenthaler, J. Kohn, M. von Aster, M. Gross, Stealth Assessment in ITS - A Study for Developmental Dyscalculia, in: A. Micarelli, J. Stamper, K. Panourgia (Eds.), *Intelligent Tutoring Systems*, volume 9684, Springer, Cham, 2016, pp. 79–89.
- [65] X. Li, C. Lu, J. Peltonen, Z. Zhang, A statistical analysis of Steam user profiles towards personalized gamification, in: *Proceedings of the 3rd International GamiFIN Conference*, CEUR-WS, Levi, 2019.
- [66] R. Sifa, C. Bauckhage, A. Drachen, Archetypal Game Recommender Systems, in: *Learning, Knowledge, Adaption (LMA) Conference*, CEUR-WS, Aachen, 2014, pp. 45–56.
- [67] J. Hamari, J. Tuunanen, Player Types: A Meta-synthesis, *Transactions of the Digital Games Research Association* 1 (2014) 29–53. URL: <https://doi.org/10.26503/todigra.v1i2.13>.
- [68] B. Cowley, D. Charles, Behavlets: a method for practical player modelling using psychology-based player traits and domain specific features, *User Modeling and User-Adapted Interaction* 26 (2016) 257–306. URL: <https://doi.org/10.1007/s11257-016-9170-1>.
- [69] P. P. Lai, S. Bai, D. Baack, K. Lee, Systems and methods for determining game level attributes based on player skill level prior to game play in the level, 2019. URL: <https://patents.google.com/patent/US10363487/en>.
- [70] P. Spronck, I. Balemans, Player Profiling with Fallout 3, in: *Proceedings of the AIIDE 2012 Conference*, AAAI Press, Palo Alto, CA, 2012, pp. 179–184.
- [71] V. Zeigler-Hill, S. Monica, The HEXACO model of personality and video game preferences, *Entertainment Computing* 11 (2015) 21–26. URL: <https://doi.org/10.1016/j.entcom.2015.08.001>.
- [72] N. Ducheneaut, N. Yee, Data Collection in Massively Multiplayer Online Games: Methods, Analytic Obstacles, and Case Studies, in: M. S. El-Nasr, A. Drachen, A. Canossa (Eds.), *Game Analytics: Maximizing the Value of Player Data*, Springer, London, 2013, pp. 641–664. URL: https://doi.org/10.1007/978-1-4471-4769-5_28.
- [73] S. Tekofsky, J. V. D. Herik, P. Spronck, A. Plaat, Psyops: Personality assessment through gaming behavior, in: *In Proceedings of the International*

- Conference on the Foundations of Digital Games, ACM, Chania, Crete, 2013, pp. 166–173.
- [74] J. R. Stafford, S. Osman, Automated video game rating, 2015. URL: <https://patents.google.com/patent/US9044675B2/en>.
- [75] N. Yee, N. Ducheneaut, L. Nelson, P. Likarish, Introverted Elves & Conscientious Gnomes: The Expression of Personality in World of Warcraft, in: Proceedings of the SIGCHI conference on human factors in computing systems, ACM, New York, NY, 2011, pp. 753–762.
- [76] M. Schaeckermann, G. Ribeiro, G. Wallner, S. Kriglstein, D. Johnson, A. Drachen, R. Sifa, L. E. Nacke, Curiously Motivated: Profiling Curiosity with Self-Reports and Behaviour Metrics in the Game "Destiny", in: Proceedings of the Annual Symposium on Computer-Human Interaction in Play - CHI PLAY '17, ACM, Amsterdam, 2017, pp. 143–156.
- [77] T. Upchurch, Google Stadia has kicked off a new age of gaming data harvesting, Wired UK (2019). URL: <https://www.wired.co.uk/article/google-stadia-data-harvesting>.
- [78] G. van Lankveld, P. Spronck, J. van den Herik, A. Arntz, Games as personality profiling tools, in: 2011 IEEE Conference on Computational Intelligence and Games (CIG'11), IEEE, Seoul, 2011, pp. 197–202. URL: <https://doi.org/10.1109/CIG.2011.6032007>.
- [79] D. Leorke, Location-Based Gaming: Play in Public Space, Palgrave Macmillan, Singapore, 2018. URL: <https://doi.org/10.1007/978-981-13-0683-9>.
- [80] E. Loos, Exergaming: Meaningful Play for Older Adults?, in: J. Zhou, G. Salvendy (Eds.), Human Aspects of IT for the Aged Population. Applications, Services and Contexts, Lecture Notes in Computer Science, Springer, Cham, 2017, pp. 254–265.
- [81] J. L. Kröger, P. Raschke, T. R. Bhuiyan, Privacy Implications of Accelerometer Data: A Review of Possible Inferences, in: Proceedings of the 3rd International Conference on Cryptography, Security and Privacy (ICCSP), ACM, New York, NY, 2019, pp. 81–87. doi:<https://doi.org/10.1145/3309074.3309076>.
- [82] J. L. Kröger, O. H.-M. Lutz, F. Müller, What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking, in: S. Fricker, M. Friedewald, S. Krenn, E. Lievens, M. Önen (Eds.), Privacy and Identity Management. Data for Better Living: AI and Privacy, IFIP Advances in Information and Communication Technology, Springer, Cham, 2019, pp. 226–241. URL: https://doi.org/10.1007/978-3-030-42504-3_15.

- [83] J. Kröger, Unexpected inferences from sensor data: a hidden privacy threat in the internet of things, in: IFIP International Internet of Things Conference, Springer, 2018, pp. 147–159. URL: https://doi.org/10.1007/978-3-030-15651-0_13.
- [84] J. L. Kröger, P. Raschke, Is my phone listening in? on the feasibility and detectability of mobile eavesdropping, in: IFIP Annual Conference on Data and Applications Security and Privacy, Springer, 2019, pp. 102–120. URL: https://doi.org/10.1007/978-3-030-22479-0_6.
- [85] J. Ruff, T. L. Wang, C. C. Quatman-Yates, L. S. Phieffer, C. E. Quatman, Commercially available gaming systems as clinical assessment tools to improve value in the orthopaedic setting: A systematic review, *Injury* 46 (2015) 178–183.
- [86] T. Casey, The Value of Deviance: Understanding Contextual Privacy, *Loyola University Chicago Law Journal* 51 (2019) 65–105.
- [87] D. J. Cornforth, M. T. P. Adam, Cluster Evaluation, Description, and Interpretation for Serious Games, in: C. S. Loh, Y. Sheng, D. Ifenthaler (Eds.), *Serious Games Analytics: Methodologies for Performance Measurement, Assessment, and Improvement*, Advances in Game-Based Learning, Springer, Cham, 2015, pp. 135–155.
- [88] J. T. Behrens, R. J. Mislevy, K. E. DiCerbo, R. Levy, An evidence centered design for learning and assessment in the digital world, Technical Report, National Center for Research on Evaluation, Standards, and Student Testing (CRESST), Los Angeles, CA, 2010.
- [89] Limelight Networks, The State of Online Gaming - 2019, 2019. URL: <https://www.limelight.com/resources/white-paper/state-of-online-gaming-2019/>.
- [90] List of data breaches, 2020. URL: https://en.wikipedia.org/w/index.php?title=List_of_data_breaches&oldid=980849764, page Version ID: 980849764.
- [91] J. L. Kröger, J. Lindemann, D. Herrmann, How do app vendors respond to subject access requests? a longitudinal privacy study on ios and android apps, in: International Conference on Availability, Reliability and Security, 2020, pp. 1–10. URL: <https://doi.org/10.1145/3407023.3407057>.
- [92] C. Utz, M. Degeling, S. Fahl, F. Schaub, T. Holz, (Un)informed Consent: Studying GDPR Consent Notices in the Field, in: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, ACM, London, 2019, pp. 973–990.
- [93] J. L. Kröger, O. H.-M. Lutz, S. Ullrich, The myth of individual control: Mapping the limitations of privacy self-management, *Social Science Research Network (SSRN)* (2021). URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3881776.

Part II

PRIVACY PRACTICES OF MOBILE APP VENDORS

5

GDPR Compliance of Mobile Apps

5.1 Background

After Part I has shown that many companies with access to mobile and IoT sensor data can analyze the data to infer detailed and intimate information about users, this chapter presents results from an investigation into how transparently such companies handle user data. Specifically, the study places the focus on mobile app vendors, as smartphones are increasingly becoming the primary device for many people [272]. In the top ten mobile markets worldwide, consumers spend an average of 4.8 hours a day on their smartphone [273]. It has been estimated that the average US consumer interacts with 46 different smartphone apps per month [274].

Paper 7 investigates – based on undercover field research – whether app vendors comply with transparency obligations prescribed by EU’s General Data Protection Regulation (GDPR). The law grants consumers the right to access the personal data that companies hold about them. Our study shows, however, that there are severe obstacles to exercising this right in practice. Based on our findings, we provide recommendations for regulatory action. While various previous studies exist in this field, it is the first time that such a study is conducted in a longitudinal fashion using a constant set of data controllers. The four-year study was initiated by Prof. Dr. Dominik Herrmann, head of the *Privacy and Security in Information Systems Group* at the University of Bamberg, and Jens Lindemann, research associate at the *Security and Privacy* department of the University of Hamburg.

When I joined the project, Dominik and Jens had already selected a sample of mobile apps and conducted two rounds of data collection. After expressing great interest in the subject, I was allowed to take the lead for the last two years of the study. During this time, I was the main person responsible for data collection and management, data analysis, and writing up the paper. The project has benefited greatly from the initiators’ experience in undercover field research. Dominik and Jens had already published a study in this specific area of research [162]. Our study was presented at the *15th International Conference on Availability, Reliability and Security* (ARES 2020), where it received the Best Paper Award.

How do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on iOS and Android Apps

Jacob Leon Kröger
kroeger@tu-berlin.de
TU Berlin, Weizenbaum Institute
Berlin, Germany

Jens Lindemann
lindemann@informatik.uni-hamburg.de
University of Hamburg
Hamburg, Germany

Dominik Herrmann
dh.psi@uni-bamberg.de
University of Bamberg
Bamberg, Germany

ABSTRACT

EU data protection laws grant consumers the right to access the personal data that companies hold about them. In a first-of-its-kind longitudinal study, we examine how service providers have complied with subject access requests over four years. In three iterations between 2015 and 2019, we sent subject access requests to vendors of 225 mobile apps popular in Germany. Throughout the iterations, 19 to 26 % of the vendors were unreachable or did not reply at all. Our subject access requests were fulfilled in 15 to 53 % of the cases, with an unexpected decline between the GDPR enforcement date and the end of our study. The remaining responses exhibit a long list of shortcomings, including severe violations of information security and data protection principles. Some responses even contained deceptive and misleading statements (7 to 13 %). Further, 9 % of the apps were discontinued and 27 % of the user accounts vanished during our study, mostly without proper notification about the consequences for our personal data. While we observe improvements for selected aspects over time, the results indicate that subject access request handling will be unsatisfactory as long as vendors accept such requests via email and process them manually.

CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy**; **Privacy protections**; • **Social and professional topics** → **Privacy policies**.

KEYWORDS

GDPR, compliance, subject access request, smartphone, mobile app, privacy

ACM Reference Format:

Jacob Leon Kröger, Jens Lindemann, and Dominik Herrmann. 2020. How do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on iOS and Android Apps. In *The 15th International Conference on Availability, Reliability and Security (ARES 2020)*, August 25–28, 2020, Virtual Event, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3407023.3407057>

© The authors 2020. This is the authors' version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 15th International Conference on Availability, Reliability and Security (ARES '20)*, Association for Computing Machinery, New York, NY, USA, pp. 1–10, <http://dx.doi.org/10.1145/3407023.3407057>.

1 INTRODUCTION

Many mobile apps collect personal data about their users and share it with third parties, such as analytics services and ad networks [3]. Given the increasing number of apps and the vast amount of data collected – often including data with no apparent relevance to the advertised functionality [30] – it is challenging for smartphone users to keep track of the data that app vendors hold about them.

As the preceding Directive 95/46/EC, the EU General Data Protection Regulation (GDPR [28]), which came into effect on 25 May 2018, affords individuals various rights over their personal data, including the rights to access, rectification, and erasure (Art. 15–17 GDPR). These rights allow smartphone users to demand transparency and regain control over the personal data collected by mobile apps. Numerous empirical investigations have revealed, however, that data subjects face various impediments in exercising their rights, with many data controllers completely refusing to comply with the law [1, 10, 24, 34, 37].

In this paper, we contribute to this line of research by presenting the results of a four-year undercover field study. While most existing work has investigated other types of public and private service providers, we focus on developers and vendors of mobile apps. This industry deserves scrutiny, as smartphones have become the primary device for many users [5].

Contribution. Until now, related studies (cf. Sect. 2) have evaluated data controller behavior based on one-time snapshot data. In contrast, we analyze how the behavior and compliance of a fixed set of app vendors change over time. Our longitudinal study includes observations both before and after the GDPR enforcement date. In three iterations between the years 2015 and 2019, we attempted to exercise the right of access with 225 vendors of mobile apps that were popular in Germany at the beginning of our study. Additionally, our first and our last round of inquiries included a question on third-party data sharing practices. In this paper, we examine in detail the obstacles we encountered as well as the veracity and completeness of the responses received. Secondly, we describe the measures that app vendors apply to verify the requester's identity and how they choose to protect the confidentiality of transmitted personal information. Thirdly, we shed light on the issue of *dissolution of personal data* that results from vendors going out of business, apps being discontinued, and stale user accounts being deleted without prior notice.

Outline. The remainder of this paper is structured as follows. First, we review related work in Sect. 2. In Sect. 3, we describe our data collection process. Then, we present our methodology for interaction with app vendors in Sect. 4. Section 5 summarizes the results of our study, a discussion of which is provided in Sect. 6. A

reflection on ethical aspects and limitations of our study follows in Sect. 7 before we conclude the paper in Sect. 8.

2 RELATED WORK

Numerous studies have examined how data controllers react to data subject requests in practice. In these studies, test requests for data access, erasure, and/or portability were either sent to specific types of organizations, such as CCTV operators [34], smartphone app vendors and website owners [10], online tracking companies [37], or to a broad range of public and private sector organizations [1, 2, 18, 19, 24, 42, 43].

In [24], investigations were carried out in ten different EU member states and, in addition to the quantitative evaluation, a detailed case-by-case assessment is provided for individual data controllers. With the exception of [1], [37] and [42], the above-referenced studies were conducted prior to the GDPR coming into force. Also, in contrast to our study, none of the existing publications offers a multi-year longitudinal evaluation over a constant set of data controllers. Since 2010, the French Association of Data Protection Officers¹ has published a yearly report on the right of access, including results from probing 150 to 200 service providers [1]. However, their list of examined data controllers is newly compiled each year, which means that behavioral changes and trends within individual organizations are not observed.

In line with our findings, previous studies report various anomalies, poor practices, and severe compliance issues on the side of the data controllers, resulting in low rates of satisfying responses to data subject requests [1, 10, 19, 24, 34, 37]. Besides a widespread unwillingness or inability to provide the requested data in time [2, 10, 34, 37], researchers have observed the use of inappropriate file formats for the transfer of personal data [42], instances of personal information leakage to impostors [7, 10], issues concerning the language and clarity of interactions [24], and unsafe procedures to authenticate data subjects [4, 25]. In some cases, researchers were not even able to locate the contact details of data controllers, rendering any request submission impossible from the outset [17, 24].

While privacy and security aspects of mobile apps have received a lot of research attention in recent years [3, 15, 21, 30, 31], little is yet known about the behavior of app vendors when it comes to fulfilling data subject requests, especially not with the GDPR being in effect.

3 DATA COLLECTION

Our objective was to obtain a comprehensive sample of apps from the iOS and Android app stores, considering a diverse set of app types and vendors. We did not want to focus on the most popular apps, because those might exhibit abnormal behavior as a result of being in the spotlight. On the other hand, we did not want to analyze outdated and niche apps with virtually no users.

To compile a suitable sample of popular mobile apps, we used market research information provided by AppAnnie (<https://www.appannie.com>). AppAnnie monitors the download counts of apps on the app stores of Google and Apple. Specifically, in August 2015,

¹ Official name in French: Association Française des Correspondants à la Protection des Données à Caractère Personnel (AFCDP)

Vendor's country of residence

| ? | rest EU | Germany | non-EU countries |
|---|---------|---------|------------------|
| 8 | 44 | 78 | 95 |

Operating system

| iOS | Android |
|-----|---------|
| 105 | 120 |

App categories (5 apps each)

| | |
|---|---|
| iOS: Business, Catalogue, Education, Entertainment, Finance, Food, Games, Health, Lifestyle, Medical, Music, Navigation, News, Photo, Productivity, Reference, Social Networking, Sports, Travel, Utility, Weather | Android: Books, Business, Comics, Communication, Education, Entertainment, Family, Finance, Games, Health, Lifestyle, Media, Medical, Music, News, Photography, Productivity, Shopping, Social, Sports, Tools, Transportation, Travel, Widgets |
|---|---|

Figure 1: Popular apps dataset overview; figures refer to number of apps out of 225 apps in total

we obtained AppAnnie's ranking lists containing the most popular apps in Germany. In current mobile operating systems, apps are assigned categories that describe their primary function or subject matter (e. g., Education, Music, or Health). For each of these app store categories, AppAnnie provides a ranked list (of varying length) that contains the most popular apps in the respective category according to the number of downloads within a fixed period. There were 24 categories for Android and 21 categories for iOS apps.

We randomly sampled apps from each of AppAnnie's category lists, subject to the following conditions. We picked a random app, installed it, and checked whether it offered users the possibility to create a personal account or to enter personal information in some other way. If an app did not meet this requirement, we picked another random app from the respective list. We also skipped an app if our sample already contained another app from the same vendor. We sampled apps until we had collected *five apps per category*, amounting to a total sample size of 225 apps (120 Android apps, 105 iOS apps). As a result, apps span a wide number of popularity ranks (avg. rank within a category: 134, std. deviation: 111, min. rank: 1, max. rank 407).

In the following, we characterize the dataset (see also Fig. 1). The largest proportion of the selected apps' vendors is based in Germany (35 %). Vendors located across other EU member states and outside of the EU account for 20 % and 42 %, respectively. For 3 % of the apps in our dataset, we could not find any information on the vendor's residence.

Like the German Federal Data Protection Act [26] and the now superseded EU Directive 95/46/EC [29], the GDPR can also apply to data controllers established in non-EU countries [40]. The GDPR explicitly states that it applies to organizations "not established in the Union, where the processing activities are related to the offering of goods or services, irrespective of whether a payment of the data subject is required, to such data subjects in the Union" (Art. 3 GDPR). While it may be difficult to exercise and enforce data protection rights against data controllers in foreign countries, we include app vendors based outside the EU in our study to compare their behavior with that of domestic vendors (see Sect. 5.5).

4 METHODOLOGY FOR INTERACTION WITH THE APP VENDORS

The apps in our sample were downloaded and installed on a smartphone with the corresponding operating system. Where it was possible, a new user account was created by signing up via email (76 %), Facebook (10 %), or Google (10 %); preference was given to signup via email if available. The remaining 4 % of the apps were populated with personal data without account registration.

For account creation and all other app interactions, we acted disguised as an ordinary consumer, always using the same identity of a volunteer (male, of legal age, German nationality). The app vendors were not informed in advance that they would be the subject of a study. We will reflect upon the necessity and ethical aspects of this covert approach in Sect. 7. Using the identity of a real person was necessary to overcome identification barriers (see Sect. 5.4).

After installation, we interacted with each app for about ten minutes, entering as much personal data as possible. Subsequently, over a period of four years, we issued three rounds of data subject requests to the app vendors (from now on referred to as R1, R2, and R3). **R1** requests were issued in November 2015, several weeks after the apps were installed. **R2** requests followed in March 2018, two months before the May 25 GDPR enforcement date. **R3** requests were sent in August 2019, more than one year after GDPR had come into effect.

The subject access requests were written in English or German (depending on the language of the app). We sent them to the vendors' email address specified on the respective page in the app marketplace. If no contact details were provided there, we searched for an email address on the vendor's website or, if it was the only option available on the website, submitted the request via a contact form. Ultimately, we were able to find a point of contact for all apps for R1. We updated our list of vendor email addresses once again before sending out our subject access requests in R2 and R3, respectively.

To account for changing external circumstances and to avoid being recognized as researchers, we deliberately used a different request text for each round of inquiries. While the right of access was the main focus of our study and, therefore, all three rounds contained data access requests, we additionally requested information about data sharing practices in R1 and R3. For R1, we chose a short and informal inquiry text, comprising only seven sentences. The texts for R2 and R3 were more elaborate. They included references to relevant data protection laws (GDPR, EU Directive 95/46/EC, and the German Federal Data Protection Act) as well as a warning that the responsible data protection authorities would be notified in the case of no response (which, in fact, we did not do).

A more formal approach was chosen for R2 and R3 because the introduction of the GDPR and the wide-ranging preparations for the new law aroused increased media attention and public awareness for data subject rights and data protection issues in general [9, 32]. In addition, since 2018, more and more self-help tools and templates for GDPR requests have become available through websites like *datarequests.org* and *gdpr.eu*. We, therefore, assumed that ordinary smartphone users were now better equipped and thus more likely

to make formal requests with legal references than they were in 2015, the year of R1.

To test how a reminder would affect the response rate, we sent a follow-up email to all vendors who had not replied to our request in R3. In the reminder (sent 85 days after the R3 request) we merely stated that we had not received a response to our previous message. Except where specifically indicated, the responses to this reminder will be ignored in the evaluation section, i. e., vendors who only responded *after* receiving a reminder will be counted as non-responders in R3.

5 ANALYSIS AND RESULTS

The received responses were subjected to a qualitative content analysis as proposed by Strauss and Corbin [35]. Following an open coding approach, we built and refined a codebook based on the data collected in R1. First, one coder went through the received responses, creating and applying an initial set of codes to all encountered findings. An independent second coder subsequently applied the resulting list of codes to the same set of responses. The initial coding yielded a Fuzzy Kappa value of 0.819, which indicates a high degree of agreement between the two coders.² Subsequently, the two coders consolidated the codebook, reassessed all responses, and resolved all conflicts. The emerging code categories and overarching themes were used to determine the foci of our paper, including the reachability and responsiveness of app vendors, the provision of the requested data, statements made on data sharing practices, occurring technical and communication problems as well as the security of data transmissions. A single coder applied the consolidated codebook to the responses in R2 and R3. The number of assigned codes per app and iteration varies between 1 and 9.

We have released a sanitized dataset containing our findings for every app [12]. A summary of our results is provided in Table 1, at the end of this section. In the published dataset and in the remainder of this paper, the apps in our sample are referred to using the pseudonyms **A1 to A225**. Where applicable, the corresponding round of requests is indicated through a subscript. **A5_{R1}**, for instance, designates the response from the vendor of App 5 to our first request. Section 5.1 provides an overview of our results, followed by an in-depth analysis of the obtained responses in Sects. 5.2–5.7. Unless specifically stated otherwise, percentage values refer to the whole sample of 225 apps. We will introduce capitalized labels to refer to particular subsets of apps.

5.1 Overview of the Received Responses

While the majority of the investigated app vendors reacted to our inquiries in some way (**RESP**: 78 %_{R1}, 81 %_{R2}, 74 %_{R3}), many did not respond to our requests (22 %_{R1}, 16 %_{R2}, 22 %_{R3}) and some were completely unreachable via email, yielding us only a delivery failure notification (0 %_{R1}, 3 %_{R2}, 4 %_{R3}). The reminder we sent after our last round of requests increased the response rate for R3 slightly from 74 % to 80 %, showing that exercising data protection rights can sometimes require perseverance.

² In contrast to Cohen's Kappa, Fuzzy Kappa by Kirilenko and Stepchenkova [11] can handle codebooks with codes that are not mutually exclusive, which matches our setting.

Out of all examined app vendors, only 14 %_{R1}, 44 %_{R2}, and 37 %_{R3} sent us a response that contained personal data (DATA). For reasons outlined in Sect. 5.2, however, many responses were unintelligible and useless. Overall, we received a proper usable export of personal data from merely 12 % of all vendors in R1, and from a significantly greater but still underwhelming proportion of 38 % in R2 and 28 % in R3. Some responses even contained deceptive or misleading statements (7 %_{R1}, 13 %_{R2}, 11 %_{R3}), which will be elaborated on in Sect. 5.6.

In each iteration of our study, the vast majority of responses arrived within five days (82 %_{R1}, 76 %_{R2}, 83 %_{R3}). Apart from a small number of exceptions, the rest arrived between day 6 and day 15 (14 %_{R1}, 22 %_{R2}, 11 %_{R3}). A minor group of outliers (2 %_{R1}, 0 %_{R2}, 2 %_{R3}) replied after more than 31 days, with one month being the time limit for responding to data protection rights requests as per Art. 12 GDPR. Before GDPR enforcement, neither the EU Directive 95/46/EC nor the German Federal Data Protection Act did prescribe such a specific time limit. However, according to legal commentaries by Wolff et al. [41, § 34 Rn. 104-106.1] and Müller-Glöße et al. [22, § 34 Rn. 1], the maximum appropriate response time ranged between two and four weeks. Figure 2 shows the frequency distribution of responses over time.

Out of all received responses (RESP), the proportion that fulfilled our subject access request (OK) by either containing a proper export of personal data or a credible statement that the data is no longer stored was 19 %_{R1}, 65 %_{R2}, and 55 %_{R3}. The others were recorded as incomplete responses. When taking all app vendors in our sample into account, including those who were unreachable or did not respond, the proportion of OK responses amounts to 15 %_{R1}, 53 %_{R2}, and 41 %_{R3}. An overview of our evaluation for all three rounds of subject access requests is provided in Fig. 3.

While we cannot precisely determine their individual influence, it can be assumed that both the introduction of the GDPR as well as the more formal and threatening tone of our inquiry in R2 and R3 (cf. Sect. 4) had an impact on the vendors' behavior. In [10], data provision inquiries yielded better response rates when they were written more formally and included explicit references to relevant data protection law.

In cases where data was provided (DATA), it was made available either by email (97 %_{R1}, 84 %_{R2}, 79 %_{R3}), by postal mail (0 %_{R1}, 9 %_{R2}, 2 %_{R3}), or through the app itself (3 %_{R1}, 7 %_{R2}, 19 %_{R3}). In most cases, personal data was provided in attachments in various file formats (namely *pdf*, *html*, *json*, *csv*, *jpg*, *png*, *docx*, and *txt*) and as plain text in the email body.

The following sections focus on particular aspects of the responses. First, we will consider deficiencies with format and content in Sect. 5.2. The degree to which the app vendors addressed our additional question on data sharing practices is examined in Sect. 5.3. After that, we consider the security of data in transit (Sect. 5.4) before we compare the responses in terms of their sender's residence (Sect. 5.5). Finally, we report on particularly intriguing aspects in Sects. 5.6 and 5.7, namely deceptive responses as well as the dissolution of personal data when vendors go out of business, apps are discontinued, and accounts are deleted unsolicitedly.




Figure 2: Frequency distribution of responses over time

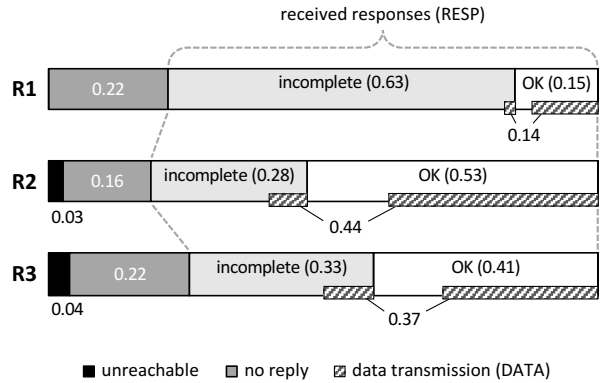


Figure 3: Evaluation of subject access request responses of all apps (n = 225)

5.2 Insufficient Responses

Failing to fulfil our subject access request, some of the received responses (RESP, $n = 175_{R1}$, 183_{R2} , 167_{R3}) contained only the labels of collected data, e. g., “birth date”, but not our actual data values, e. g., “1977-03-09” (22 %_{R1}, 6 %_{R2}, 5 %_{R3}). The frequency of this phenomenon strongly decreased over the course of our study, with a particularly large drop between R1 and R2.

Even when actual data was transmitted (DATA, $n = 32_{R1}, 100_{R2}, 84_{R3}$), it was often unintelligible due to serious formatting errors or obscure data labels (9 %_{R1}, 14 %_{R2}, 24 %_{R3}). In some of these cases,³ the data consisted of a continuous block of alphanumeric characters without headings, spaces, and line breaks. According to EU law, personal data should be made available “in a structured, commonly used and machine-readable format” (Art. 20 GDPR). Interestingly, despite the relatively short duration of app usage (10 to 15 minutes per app), we sometimes received data in large and confusing folder structures. The responses of A4_{R3} and A15_{R3}, for example, contained 17 folders and 27 individual files, respectively, but only a handful of relevant data points. Also, in several cases, data was only transferred as an image file (e. g., a screenshot of an internal database), which made it impossible to extract information via copy and paste (3 %_{R1}, 16 %_{R2}, 4 %_{R3}).

In other cases, technical or communication problems resulted in relevant information and personal data not being transmitted in the first place. These problems include dead download links and empty email attachments (A11_{R3}, A29_{R3}, A33_{R3}, A146_{R3}), unclear or incomprehensible instructions for how to obtain the data (A120_{R3}, A150_{R3}, A151_{R3}, A225_{R3}), references to inaccessible privacy functions inside the app (A157_{R3}), malfunctioning authentication mechanisms (A87_{R3}, A140_{R2}, A178_{R3}, A181_{R1}, A202_{R3}), and cases where the responsibility for privacy-related inquiries was referred back and forth between departments or affiliated companies (A21_{R1}, A25_{R1}, A118_{R3}). In most of these incidents, pointing out the problem to the app vendor did not help. For example, we exchanged 24 emails with A151_{R3} in an attempt to clarify a technical issue before finally giving up because the responses had become repetitive. These observations are in line with those by Ausloos et al. [2] who found that subject access requests often lead to endless sequences of emails without resulting in the requested transfer of personal data.

Furthermore, some of the investigated app vendors appear to lack essential resources to receive and process subject access requests in a professional manner. For example, from A30_{R3}, we received only a “recipient inbox full” error message. Other vendors explicitly stated that they lack the time, technical means, or personnel to compile and transmit the requested data (A105_{R2}, A172_{R3}), to reply in time (A50_{R3}, A83_{R3}), or to read emails in their inbox (A48_{R3}, A199_{R3}). While the many inadequate responses to our requests indicate widespread unawareness and ignorance about legal duties, some responses revealed a specific deficiency in legal knowledge. The person writing A112_{R2}, for example, did not know that a right of access already existed in EU law before the GDPR. Also, after the commencement of GDPR in May 2018, some vendors stopped offering a login function (A42_{R3}, A58_{R3}, A93_{R3}) or even discontinued the entire app, stating explicitly their inability to comply with the new law (A62_{R3}).

There were also language barriers. Although our requests were written in the corresponding app’s default language (English or German), two vendors replied in Spanish (A108_{R3}, A201_{R3}) and one always replied in Korean (A149). A31_{R3}, A51_{R3} and A174_{R3} changed the language of communication from English to German during

our email correspondence without prior announcement. Another vendor used the online translation service Google Translate to communicate with us, which led to ambiguities and confusion (A150_{R3}). A206_{R2} replied in bad English, and the responses A165_{R1} and A189_{R1} were utterly incomprehensible. Beyond that, some vendors used technical jargon, such as the term “SDK”⁴ (A210_{R3}). While jargon does not render a response useless per se, ordinary users may have difficulty understanding it.

Astonishingly, three vendors stated that data exports are only provided for paid subscriptions, not for users of their free version (A89_{R2}, A137_{R1}, A202_{R2}). And, in another grave violation of the most basic data protection principles, A100_{R2} sent us personal data of *another* user, including the person’s full name, contact details, login credentials, app usage logs, and even location data. In one case, we received the response to our request from a foreign email address with no apparent relation to the app in question (A77_{R3}).

Finally, some responses were impolite or contained careless mistakes. For instance, A161_{R2} declared – without any explanation – to have “no interest in having [us] as a customer.” A193_{R3} and A224_{R3} replied without a formal salutation, A202_{R3} misspelled the name in the salutation, and both A13_{R3} and A151_{R3} sent us the same emails multiple times for no apparent reason. In a strange mistake, the support agent writing A89_{R3} even addressed us with *their* name. Although these incidents were no severe obstacles to communication, they do indicate a lack of care and rigor in answering our requests.

5.3 Data Sharing with Third Parties

Along with the request for data access, in R1 and R3 we also requested information on third-party data sharing practices (cf. Sect. 4). Specifically, we asked the app vendors to name any third-party recipients to whom they had disclosed our personal data. Only a minority of the received responses (RESP, $n = 175_{R1}, 167_{R3}$) addressed this question at all (41 %_{R1}, 40 %_{R3}).

Out of this subset (DSANS, $n = 72_{R1}, 66_{R3}$), a decreasing proportion stated that personal data is never being shared with third parties (63 %_{R1}, 52 %_{R3}). The remaining responses vary widely in terms of the length and detail in which data sharing practices are explained. While some state specifically what data is made available to which partners for what purpose (56 %_{R1}, 37 %_{R3}), others list only a few generic reasons for data sharing (18 %_{R1}, 16 %_{R3}) or potential data recipients without naming the respective purpose (11 %_{R1}, 25 %_{R3}). A few vendors disclose only the categories of data being shared (11 %_{R1}, 6 %_{R3}) or even refrain from providing any specifics beyond mentioning the existence of data sharing (4 %_{R1}, 16 %_{R3}).

Overall, only 36 %_{R1} and 32 %_{R3} of the received responses (RESP) adequately addressed our question by either assuring that no data is being shared with third parties or by providing a list of data recipients. And there is still room for doubt, even in these cases. In particular, the proportion of apps that explicitly mention third-party tracking (only 19 %_{R1} and 14 %_{R3} of DSANS) is far below what one should expect based on the results of large-scale empirical studies on the topic. Binns et al. [3], for example, examined the

³ Out of ethical considerations, we refrain from publishing screenshots and excerpts of the responses (cf. Sect. 7).

⁴ SDK is an abbreviation of “software development kit” (i. e., a collection of software development tools)

prevalence of third-party trackers on 959,000 Android apps and found that 90.4 % included at least one tracker host. Therefore, it is well conceivable that some of the responding app vendors withheld relevant information on their data sharing practices.

5.4 Identity Verification and Security of Transmitted Data

Before responding to a subject access request, data controllers must carefully verify the identity of the requester. In a recent study, Martino et al. [7] demonstrated that unauthorized third parties could fairly easily obtain personal data through fake subject access requests sent from a different email address. In our study, most of the app vendors that sent us a copy of personal data (DATA, $n = 32_{R1}, 100_{R2}, 84_{R3}$) did not ask for prior identity verification (84 %_{R1}, 85 %_{R2}, 76 %_{R3}). Note, however, that we sent our inquiries using the email address that we had used for account registration, which provides a minimal form of authentication by itself.

Only a minority of app vendors in DATA demanded certain pieces of identity-related information. We were asked to provide our birth date, address, and customer number (13 %_{R1}, 2 %_{R2}, 0 %_{R3}) as well as a copy of a utility bill, ID card, or driving license (0 %_{R1}, 5 %_{R2}, 5 %_{R3}). A problem with this type of authentication mechanism is that it often requires those seeking to exercise their rights to provide the data controller with additional personal information. This practice has been recognized as “contrary to the spirit of data protection law” [25]. Implementing a rather unusual method for authentication, A119_{R2} demanded from us to “log into your social media account and post for public viewing a list of random characters.” Surprisingly, in three cases (A5_{R2}, A8_{R2}, A157_{R2}), personal data was disclosed to us before we had even responded to an authentication request.

In accordance with a GDPR best practice recommendation (from Recital 63 [28]), several app vendors did not directly send data to us. Instead, they provided remote access to the data through a self-service system within the password-protected app (3 %_{R1}, 7 %_{R2}, 19 %_{R3} of DATA). For this practice, a positive trend was observed over the three iterations of our study, with a particularly sharp increase between R2 and R3.

Some of the responders (RESP, $n = 175_{R1}, 183_{R2}, 167_{R3}$) failed to use transport-layer encryption via TLS (Transport Layer Security) in at least one of their emails. Transport-layer encryption is a baseline security measure, which has been deemed mandatory for the transmission of personal data via email [36]. For our analysis, we checked whether all relevant *Received* headers within emails contained signs of TLS being used (indicated by “SMTPS” or the statement of the negotiated cipher suite). Here, again, the worst result was recorded in R1, where 13 % of all responses (RESP) were not encrypted. This figure improved and remained stable over R2 and R3 (both 3 %). A21_{R2}, and A75_{R1} even transmitted the requested personal data without encryption. Besides, as indicated by email metadata (again, *Received* headers), many app vendors handle customer support via third-party services, e. g., ticket system software operated by Zendesk, Salesforce, Helpscout, Mandrill and others. As a result, even vendors that are themselves based in the EU may disclose personal data to third parties in other jurisdictions. In these

cases, exercising one’s right to access may disperse one’s personal data to additional data processors.

Where data was transferred (DATA), two other security mechanisms we observed were the expiration of download links after a few days or weeks (0 %_{R1}, 2 %_{R2}, 10 %_{R3}) and the protection of email attachments with passwords (6 %_{R1}, 3 %_{R2}, 12 %_{R3}). Except for one case where the corresponding password was communicated via telephone (A225_{R1}), the access keys were sent via email, sometimes in the same email as the password-protected data (A28_{R3}, A211_{R3}). Finally, for security reasons, several vendors chose to send the data by postal mail (0 %_{R1}, 9 %_{R2}, 2 %_{R3}) – sometimes even by registered mail (A16_{R2}, A81_{R2}, A107_{R2}, A225_{R2}).

Our correspondence with the vendors indicates that the employed authentication mechanisms are not the result of rigorous attacker modeling. Given the prevalence of TLS on mail servers, they provide little to no additional security against passive eavesdroppers. More importantly, they cannot protect against an adversary that has already compromised a user’s email account. Many pieces of information that have been requested from us for identity verification can be extracted from previously exchanged emails. Moreover, adversaries that control the email account of a user receive data and accompanying passwords when these are sent via email. Such adversaries could also compromise a user’s account on the app itself via the commonplace email-based password reset functionality, either to steal data directly from within the app or to change the postal address to which the data is sent.

5.5 Comparison by Vendor’s Residence

In the following, we explore behavioral differences between the app vendors from our sample that are based in Germany (GER, $n = 78$), in other EU member states (EMS, $n = 44$), and in the rest of the world (WLD, $n = 95$). Apps for which the vendor’s residence could not be determined ($n = 8$) are not considered in this analysis.

Across all three iterations of our study, the proportion of vendors responding to our inquiry was always somewhat higher among GER (82 %_{R1}, 88 %_{R2}, 78 %_{R3}) than among EMS (75 %_{R1}, 77 %_{R2}, 77 %_{R3}), and WLD (75 %_{R1}, 77 %_{R2}, 69 %_{R3}). We will refer to these responses as RESP_{GER} ($n = 64_{R1}, 69_{R2}, 61_{R3}$), RESP_{EMS} ($n = 33_{R1}, 34_{R2}, 34_{R3}$) and RESP_{WLD} ($n = 71_{R1}, 73_{R2}, 66_{R3}$). Throughout R1 and R2, the proportion of responses containing an export of personal data was highest among RESP_{GER} (30 %_{R1}, 71 %_{R2}, 61 %_{R3}). RESP_{EMS} caught up towards the end of our study (21 %_{R1}, 62 %_{R2}, 62 %_{R3}), leaving the rate of data transmissions among RESP_{WLD} far behind (7 %_{R1}, 38 %_{R2}, 38 %_{R3}). We name these data containing responses DATA_{GER} ($n = 19_{R1}, 49_{R2}, 37_{R3}$), DATA_{EMS} ($n = 7_{R1}, 21_{R2}, 21_{R3}$) and DATA_{WLD} ($n = 5_{R1}, 28_{R2}, 25_{R3}$).

Where sufficient for comparison, only the percentage figures from the most recent inquiry (R3) will be shown below. If the previous rounds of requests yielded deviating results, we will explicitly mention them.

In terms of OK responses (a usable export of personal data or a credible statement that no personal data is being stored), the three geographical subsets followed a similar development. The results show a sharp improvement between R1 and R2 followed by deterioration in R3 (GER: 51 %_{R3}, EMS: 45 %_{R3}, WLD: 31 %_{R3}). In all three runs, GER had the highest rate of OK responses; the lowest rate

was consistently found among WLD. Contributing to this discrepancy, Germany-based vendors scored better on the structure and intelligibility of transmitted data. Due to a lack of these qualities, 24 %_{R3} and 33 %_{R3} of the data exports provided by DATA_{WLD} and DATA_{EMS} were useless, respectively. Among DATA_{GER}, this figure was relatively low (19 %_{R3}).

And while the geographical subsets had comparable response rates to our question on data sharing practices in R1 (see Sect. 5.3), GER achieved by far the best result in our last round of requests (49 %_{R3}), followed by EMS (25 %_{R3}) and then WLD (17 %_{R3}).

Averaging over all responses received in our three rounds of inquiries, the vendors residing in Germany took longer to respond (5.5 days) than the vendors based in other EU member states (3.2 days) and outside of the EU (2.9 days). However, measured against legally prescribed time constraints (see Sect. 5.1), all of these average durations are acceptable.

The Germany-based vendors exhibited substantially inferior performance in the area of requester authentication. Compared to the already surprisingly small portion of vendors who conducted some form of identity check among DATA_{EMS} (48 %_{R3}) and DATA_{WLD} (28 %_{R3}), authentication requests among DATA_{GER} were much sparser (8 %_{R3}). In our first inquiry, this figure had been 21 %_{R1} for DATA_{GER} and 0 %_{R1} for DATA_{WLD}, but subsequently rose sharply for DATA_{WLD} and, contrary to our expectations, decreased considerably for DATA_{GER}. Before taking the sudden lead in R3, DATA_{EMS} had much weaker authentication results (14 %_{R1}, 10 %_{R2}).

At the end of our study, more apps from GER (10 %) and especially EMS (14 %) had been discontinued than from WLD (7 %). The proportion of disappearing user accounts (see Sect. 5.7) was distributed fairly evenly across the three groups.

5.6 Deceptive and Misleading Responses

Many of the responding app vendors (RESP, $n = 175_{R1}, 183_{R2}, 167_{R3}$) made misleading or demonstrably false statements. Several claimed that the app in question or our user account no longer existed – although we were still able to install, log in to, and use the app on our test devices (6 %_{R1}, 11 %_{R2}, 6 %_{R3}). Some vendors even contradicted themselves. For example, A34_{R3} repeatedly picked up specific points from our email request while pretending that the same message had never arrived. A188_{R3} claimed that no data related to our account had been collected. A213_{R3} stated that, as a matter of principle, copies of personal data would never be provided for security reasons. Both companies, however, had already sent us a copy of stored personal data a few days before. Moreover, some app vendors falsely claimed that they had already replied to our request earlier (A127_{R3}) or falsely promised to get back to us within a specified time (A2_{R3}, A33_{R3}, A59_{R2}, A102_{R2}, A118_{R3}, A126_{R3}).

Besides the capacity limits mentioned in Sect. 5.2, one reason for these observed inconsistencies and contradictions seems to be a lack of standardized processes for dealing with subject access requests. Our inquiries were sometimes processed and answered simultaneously by multiple different departments or employees (A4_{R3}, A21_{R2}, A37_{R3}, A106_{R3}, A153_{R3}, A188_{R3}, A191_{R3}, A197_{R3}, A213_{R3}). Additionally, in some cases where a data export was provided, we noticed that the export was not complete, i. e., some pieces

of personal data stored in the app had been omitted for unknown reasons (A47_{R2}, A158_{R1}, A173_{R1}).

We also suspected that some vendors merely *pretended* to be poorly reachable when they received subject access requests – while others actually *had* insufficient resources to process incoming emails. To confirm this hypothesis, we tested how the vendors that failed to respond to our requests reacted to non-privacy related inquiries. Using another (different) fake identity, we emailed the vendors who had not replied in R1 and R3, expressing interest in promoting their apps on a personal blog (R1) or YouTube channel (R3). Out of the group of initial non-responders, 31 %_{R1} and 22 %_{R3} replied to these dummy requests, many of them within a few hours, proving that their email inbox was in fact being monitored.

5.7 Discontinued Apps and Accounts

Manual checks revealed that prior to R2, 4 % of the initial 225 apps no longer existed. 9 % were gone prior to R3. Some of the corresponding vendors (**GONE**, $n = 10_{R2}, 21_{R3}$) still responded to our requests, informing us that their app had been discontinued (30 %_{R2}, 14 %_{R3}) or providing us with a copy of personal data that was still being stored (10 %_{R2}, 0 %_{R3}). The majority of them, however, did not address our request in their response or remained silent (50 %_{R2}, 57 %_{R3}) and some were not reachable via email anymore (10 %_{R2}, 29 %_{R3}).

Even among the vendors of the apps that could still be installed (**NOTGONE**, $n = 225_{R1}, 215_{R2}, 204_{R3}$), several responded that none of our submitted personal data is stored on their servers (8 %_{R1}, 24 %_{R2}, 17 %_{R3}). While some of these vendors (**NODATA**, $n = 17_{R1}, 51_{R2}, 34_{R3}$) explained that they only process data locally on the user's device without ever having direct access to it (24 %_{R1}, 25 %_{R2}, 24 %_{R3}), most of them simply stated that a matching record did not exist in their database *anymore* (76 %_{R1}, 75 %_{R2}, 76 %_{R3}).

Based on manual login attempts, 27 % of all user accounts that we had initially created were no longer accessible at the end of our study, including discontinued apps (**GONE**). In a few astounding cases, our user account was deleted or deactivated in direct response to our inquiry (A12_{R3}, A93_{R2}, A133_{R2}, A161_{R2}, A209_{R2}) – either with unconvincing excuses or with no justification at all. A12_{R3}, for example, cited “credit card problems” as the reason, although we had never connected a credit card with any of the investigated apps.

All account discontinuations and erasures of personal data were carried out without any request or consent from our side. It is striking that we found no provisions in the corresponding privacy policies or terms of service that would explain the deletions. Also, apart from two vendors who gave notice that our account would be discontinued due to inactivity between R1 and R2 (A84, A106), we received no account deletion notification outside the responses to our requests. Neither a terminated app nor a deleted account necessarily implies, of course, that all of the user's data was, in fact, entirely deleted by the app vendor and potential affiliates. This is well illustrated by a few cases where the vendor transferred our personal data despite having discontinued the app or our user account (A93_{R2}, A128_{R2}, A133_{R2}).

Among the vendors who provided an export of personal data (**DATA**, $n = 32_{R1}, 100_{R2}, 84_{R3}$), many included in their response an

unsolicited offer to delete our user account and the related data – especially in R2 and R3 (9 %_{R1}, 39 %_{R2}, 37 %_{R3}). It seems that the loss of a single user is often perceived as advantageous over bearing the effort and potential legal trouble associated with responding to data subject requests.

Table 1: Summary of results; figures refer to number of apps out of 225 apps in total

| | R1 (2015) | R2 (2018) | R3 (2019) |
|---|--------------|--------------|--------------|
| Request Response (Sect. 5.1) | | | |
| Vendor unreachable (delivery failure) | 1 | 6 | 9 |
| Vendor did not respond | 49 | 36 | 49 |
| Vendor responded (RESP) | 175 | 183 | 167 |
| Response sufficient (OK) | 33 | 119 | 92 |
| Response contains data (DATA) | 32 | 100 | 84 |
| Response deceptive or misleading | 13 | 24 | 19 |
| Average response time [in days] | 3.7 | 3.8 | 4.0 |
| Data Sharing Practices (Sect. 5.3) | | | |
| Response addresses data sharing (DSANS) | 72 | n/a | 66 |
| Response states that data is shared | 27 | n/a | 32 |
| Response states that data is not shared | 45 | n/a | 34 |
| Sufficient information on data sharing | 63 | n/a | 54 |
| Requester Authentication (Sect. 5.4) | | | |
| Authentication mechanism in place | 5 | 15 | 20 |
| App login required | 1 | 7 | 16 |
| Identification document required | 0 | 5 | 4 |
| Identity-related information required | 4 | 2 | 0 |
| Other authentication mechanism | 0 | 1 | 0 |
| Security of Transmitted Data (Sect. 5.4) | | | |
| Data file(s) protected with password | 2 | 3 | 10 |
| Download link has expiration date | 0 | 2 | 8 |
| Data export via postal mail | 0 | 9 | 2 |
| TLS encryption failure | 23 | 6 | 5 |
| Existence of App & Account (Sect. 5.7) | | | |
| Discontinued apps (aggregated) | 0 | 10 | 21 |
| Discontinued user accounts (aggregated) | 0 | 49 | 61 |

6 DISCUSSION

The results of our four-year undercover study reveal severe obstacles to exercising the right of access in the mobile app space. We found that the documented problems largely persist over time and have not substantially improved since 2014, when Herrmann and Lindemann obtained success rates that are comparable with the ones we obtained in 2019 [10].

In certain areas, we did observe positive trends. Between 2015 and 2019, the investigated vendors seem to have increased their willingness and ability to disclose stored personal information to data subjects upon request. Out of all examined app vendors, the proportion that managed to provide a proper export of personal data rose from 12 %_{R1} to 28 %_{R3}. Overall, the rate of acceptable responses (OK) to our subject access requests went up from 15 %_{R1} to 41 %_{R3} (see Sect. 5.1).

There was also a crucial improvement in the protection of private messages and transmitted personal information. Specifically, among all received replies (RESP, $n = 175_{R1}, 183_{R2}, 167_{R3}$), we observed

a reduction of unencrypted email correspondence from 13 %_{R1} to 3 %_{R3}. And we noticed a growing tendency among the vendors to not send data exports directly through email or postal mail but to make them available via a portal in the password-protected app (see Sect. 5.4). Among all vendors who sent us a copy of personal data (DATA, $n = 32_{R1}, 100_{R2}, 84_{R3}$), the prevalence of this practice increased from 3 %_{R1} to 19 %_{R3}.

Even with these improvements, however, getting access to one’s data is still a frustrating endeavor most of the time. While we observed substantial improvements between 2015 (R1) and 2018 (R2), most of the identified problems (e.g., low overall response rate, incomplete responses, deceptive and misleading statements, unintelligible data exports) remained static or even worsened between R2 and our last round of inquiries in 2019 (R3). Besides the impact of the emerging GDPR, the significant improvements between R1 and R2 can probably be attributed to a considerable extent to the less formally written request in R1 (see Sect. 4), reflecting previous findings from [10]. Although formal and informal requests have the same legal validity, our results suggest that successfully exercising data protection rights in practice may require a certain manner of expression and a degree of legal expertise, potentially going beyond the capacity of average smartphone users.

Why has there been no further improvement between R2 and R3? One explanation may be that app vendors have realized that the mostly underfunded and poorly organized supervisory authorities have insufficient resources to enforce the regulation and penalize misconduct [33, 38, 39]. As a result, there is little incentive for vendors to handle subject access requests professionally. The handling of our requests varied significantly, indicating a lack of well-established best practices in the mobile app industry.

We observed numerous alarming cases of disregard for the most basic information security and data protection principles. Among other failures, 76 to 85 % of the data exports (DATA) were provided without any attempt at verifying the requester’s identity. Where authentication mechanisms were in place, they were sometimes accidentally circumvented by the app vendors themselves. We also received exports of personal data via non-encrypted channels and were even given access to personal data not belonging to us (see Sect. 5.4).

Furthermore, we observed that many of the investigated apps were discontinued (9 %), with some vendors becoming completely unreachable during our study (4 %). The latter case can create unsettling uncertainty for data subjects, leaving them in doubt about the fate of their data. Given this “personal data dissolution,” users cannot get answers to the essential question, “Who knows what about me?” anymore. Answering this question, in turn, is a natural prerequisite for exercising the remaining ARCO rights (access, rectification, cancellation, opposition) in an informed manner, and thus an essential basis for informational self-determination [17]. While the case of unreachable and disappearing data controllers illustrates the problem well and potentially exacerbates it, any non-reaction or incomplete response to a subject access request can, of course, have the same consequence.

Personal data dissolution does not only result from vendors going out of business or apps being discontinued. Another source is user accounts that disappear without prior notice: 27 % of the initial accounts were gone by the end of our study. Of course, blocking or

deleting accounts after long periods of user inactivity is a legitimate concern of app providers. The practice also aligns with the GDPR’s storage limitation principle [8] and is often in the interest of the data subject by helping to prevent the accumulation of unused “zombie accounts” [23]. However, it seems evident that affected users should be informed in advance, be offered a final copy of their data, and be asked for explicit consent if the data is retained for further use. With few exceptions, these basic rules were not followed by the respective vendors in our study. In the vast majority of cases, we did not even receive an account deletion notification.

Overall, our results show very clearly that current processes for receiving and dealing with data subject requests have plenty of room for improvement. Existing research-based suggestions and recommendations for data controllers need to be compiled into actionable guidelines and distributed in a form that makes them digestible for small- and medium-sized organizations, such as app vendors. This includes, in particular, guidance on how to authenticate data subjects safely [4, 7, 25], how to transfer personal data [42], and how to facilitate the submission of requests [2, 10, 17]. It should be a key objective to replace the error-prone manual processing of data subject requests with automated and standardized interfaces for obtaining personal data and other privacy-related information. To incentivize the rapid and broad-scale adoption of such approaches, industry-specific legal requirements along these lines could be helpful.

Examining a fixed group of data controllers over a sustained period of time has proven to be suitable for this type of investigation and should be used more widely. The basic issues and challenges to exercising data protection rights in practice are now well understood. More longitudinal research, however, is needed to monitor whether progress is being made, to assess the impact of legal measures, as well as to track and refine emerging best practices.

Future studies and discussions on the privacy behavior of mobile apps should take into account that the stored personal data is not necessarily limited to recorded actions and inputs of the user. Patterns and correlations in collected data may leak additional information in a way that is not easily understood or anticipated by the user [13, 14, 16]. To achieve an adequate level of transparency, it would be necessary for users to be informed about such forms of data acquisition as well.

7 ETHICAL CONSIDERATIONS AND LIMITATIONS

As explained in Sect. 4, we interacted with the app vendors disguised as an ordinary app user. One crucial ethical aspect of undercover field studies is that resources of the studied subjects are consumed without their consent – namely, in our case, the working time of employees of the examined app vendors. However, as in related studies [1, 2, 10, 24, 37, 42], the covert approach was necessary to test the behavior of the vendors under realistic conditions and to prevent a research participation effect [20]. It can be assumed that employees of the investigated companies would have put more care and diligence into answering our inquiries if they had known about the nature of our study. For this reason, we consider undercover field research the only viable approach to investigate problems related to exercising data protection rights in practice.

It is not the purpose of this paper to name and shame individual companies. Therefore, we only referred to apps from our sample using pseudonyms. Moreover, we have decided not to inform the investigated companies about the study we have conducted to protect the responsible employees from negative consequences. For the same reason, after careful consideration and given the many inadequate responses to our inquiries, we also decided to publish neither the collected postal and email correspondence nor the received data records – not even in pseudonymized form. There would always be a residual risk that a vendor might recognize itself. In the interest of reproducibility and traceability, we did, however, release our *results dataset*, which consists of a sanitized and comprehensive table containing all findings for every app (represented by the numeric ID used in this paper) [12].

Our study has several limitations that need to be recognized. First, we want to stress that we have studied a particular sample of mobile apps. Our results, therefore, cannot be generalized to the whole population of apps, nor to all popular apps. Also, as explained in Sect. 4, we used a different request text for each iteration of our study, varying in length, formality, and in terms of the laws that were cited. Thus, the results obtained from R1, R2, and R3 are not unconditionally comparable.

Furthermore, the sending of three repeated requests, with long pauses between them, after only a few minutes of user activity on each app in the beginning, presumably deviates strongly from normal user behavior. This activity pattern may have aroused curiosity or suspicion among the investigated companies and potentially affected their response to our inquiries. Simulating an active user on 225 apps over a period of four years was not feasible with our available resources. On the other hand, this approach allowed us to measure the extent of *personal data dissolution* (cf. Sect. 6).

Other surrounding factors, such as the increasing public interest in data protection issues [9, 32] and several smartphone-related privacy scandals [6, 27], could also have affected the compliance with subject access requests in the mobile app industry. Therefore, despite the longitudinal nature of our study, we cannot measure the impact of the newly enacted GDPR in isolation.

8 CONCLUSION

In this longitudinal study, we investigated the obstacles faced by data subjects in exercising their right to access with mobile app vendors. Our results indicate positive trends in selected areas. The overall situation, however, is still as unsatisfactory today – with GDPR in force – as it was at the beginning of our study in 2015. Even in the second iteration of our study, in which the reactions to our subject access request were the most promising, we received an acceptable response from only 53 % of the examined vendors. In the first and third iteration, this figure was 15 % and 41 %, respectively. Response rates to questions we asked about third-party data sharing practices were similarly disappointing.

Besides a general lack of responsiveness, the observed problems range from malfunctioning download links and authentication mechanisms over confusing data labels and file structures to impoliteness, incomprehensible language, and even serious cases of carelessness and data leakage. It is evident from our results that there are no well-established and standardized processes for subject

access requests in the mobile app industry. Moreover, we found that many vendors lack the motivation to respond adequately. Many of the responses we received were not only completely insufficient, but also deceptive or misleading. Equally worrisome are cases of unsolicited dissolution of personal data, for instance, due to the apparently widespread practice of deleting stale accounts without prior notice.

With regard to the sensitive personal data that is regularly collected by mobile apps, this deficient and stagnating status quo is hardly tolerable. What could help to improve the situation is a combination of random compliance checks by authorities, coupled with better support for data controllers through industry-specific guidelines and best practices. In particular, there should be mandatory standard interfaces for providing data exports and other privacy-related information to data subjects, obviating the need for the manual processing of GDPR requests.

REFERENCES

- [1] Association Française des Correspondants à la protection des Données à caractère Personnel. 2020. Données personnelles - Index AFCDP 2020 du Droit d'accès. <https://afcdp.net/index-du-droit-d-acces/>
- [2] Jef Ausloos and Pierre Dewitte. 2018. Shattering One-Way Mirrors. Data Subject Access Rights in Practice. *Data Subject Access Rights in Practice (January 20, 2018)*. International Data Privacy Law 8, 1 (2018), 4–28.
- [3] Reuben Binns, Ulrik Lyngs, Max Van Kleek, Jun Zhao, Timothy Libert, and Nigel Shadbolt. 2018. Third party tracking in the mobile ecosystem. In *Proceedings of the 10th ACM Conference on Web Science*. 23–31.
- [4] Coline Boniface, Imane Fouad, Nataliia Bielova, Cédric Lauradoux, and Cristiana Santos. 2019. Security Analysis of Subject Access Request Procedures. In *Annual Privacy Forum*. Springer, 182–209.
- [5] Christina Bröhl, Peter Rasche, Janina Jablonski, Sabine Theis, Matthias Wille, and Alexander Mertens. 2018. Desktop PC, tablet PC, or smartphone? An analysis of use preferences in daily activities for different technology generations of a worldwide sample. In *International Conference on Human Aspects of IT for the Aged Population*. Springer, 3–20.
- [6] Catalin Cimpanu. 2019. Another Facebook privacy scandal, this time involving its mobile analytics SDK. *ZDNet* (Feb. 2019). <https://www.zdnet.com/article/another-facebook-privacy-scandal-this-time-involving-its-mobile-analytics-sdk/>
- [7] Mariano Di Martino, Pieter Robyns, Winnie Weyts, Peter Quax, Wim Lamotte, and Ken Andries. 2019. Personal Information Leakage by Abusing the GDPR “Right of Access”. In *Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security (SOUPS’19)*. USENIX Association, USA, 371–386.
- [8] Majid Hatamian. 2020. Engineering Privacy in Smartphone Apps: A Technical Guideline Catalog for App Developers. *IEEE Access* 8 (2020), 35429–35445.
- [9] Alex Hern. 2018. What is GDPR and how will it affect you? *The Guardian* (May 2018). <https://www.theguardian.com/technology/2018/may/21/what-is-gdpr-and-how-will-it-affect-you>
- [10] Dominik Herrmann and Jens Lindemann. 2016. Obtaining personal data and asking for erasure: do app vendors and website owners honour your privacy rights?. In *GI Sicherheit 2016*. Gesellschaft für Informatik e.V., Bonn, 149–160. arXiv:1602.01804
- [11] Andrei P. Kirilenko and Svetlana Stepchenkova. 2016. Inter-Coder Agreement in One-to-Many Classification: Fuzzy Kappa. *PLOS ONE* 11, 3 (03 2016), 1–14.
- [12] Jacob Leon Kröger, Jens Lindemann, and Dominik Herrmann. 2020. Subject Access Request response data - 105 iOS and 120 Android apps. <https://doi.org/10.14279/depositonce-10338>
- [13] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What does your gaze reveal about you? On the privacy implications of eye tracking. In *Privacy and Identity Management*. Springer, 226–241.
- [14] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Philip Raschke. 2020. Privacy Implications of Voice and Speech Analysis–Information Disclosure by Inference. In *Privacy and Identity Management*. Springer, 242–258.
- [15] Jacob Leon Kröger and Philip Raschke. 2019. Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping. In *Data and Applications Security and Privacy XXXIII*. Springer International Publishing, Cham, 102–120.
- [16] Jacob Leon Kröger, Philip Raschke, and Towhidur Rahman Bhuiyan. 2019. Privacy implications of accelerometer data: a review of possible inferences. In *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy*. 81–87.
- [17] Xavier Duncan L’Hoiry and Clive Norris. 2015. The honest data protection officer’s guide to enable citizens to exercise their subject access rights: lessons from a ten-country European study. *International Data Privacy Law* 5, 3 (2015), 190–204.
- [18] Xavier L’Hoiry and Clive Norris. 2017. Exercising Access Rights in the United Kingdom. In *The Unaccountable State of Surveillance*. Springer, 359–404.
- [19] René Mahieu, Hadi Asghari, and Michel van Eeten. 2018. Collectively exercising the right of access: individual effort, societal effect. *Internet Policy Review* 7, 3 (2018).
- [20] Jim McCambridge, John Witton, and Diana R Elbourne. 2014. Systematic review of the Hawthorne effect: new concepts are needed to study research participation effects. *Journal of Clinical Epidemiology* 67, 3 (2014), 267–277.
- [21] Nurul Momen, Majid Hatamian, and Lothar Fritsch. 2019. Did App Privacy Improve After the GDPR? *IEEE Security & Privacy* 17, 6 (2019), 10–20.
- [22] Rudi Müller-Glöße, Ulrich Preis, and Ingrid Schmidt (Eds.). 2016. *Erfurter Kommentar zum Arbeitsrecht (ErfKoArbR)* (16 ed.). Beck, München.
- [23] David Nield. 2019. How To Clear Out Your Zombie Apps and Online Accounts. *Wired* (May 2019). <https://www.wired.com/story/delete-old-apps-accounts-online/>
- [24] Clive Norris and Xavier L’Hoiry. 2017. Exercising Citizen Rights Under Surveillance Regimes in Europe–Meta-analysis of a Ten Country Study. In *The Unaccountable State of Surveillance*. Springer, 405–455.
- [25] Chris Norval, Heleen Janssen, Jennifer Cobbe, and Jatinder Singh. 2018. Reclaiming data: Overcoming app identification barriers for exercising data protection rights. In *ACM International Joint Conference and International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. 921–930.
- [26] German Federal Office of Justice. 2019. Federal Data Protection Act (BDSG). https://www.gesetze-im-internet.de/englisch_bdsge/englisch_bdsge.html#p0014
- [27] Kate O’Flaherty. 2019. Huawei Security Scandal: Everything You Need to Know. *Forbes* (Feb. 2019). <https://www.forbes.com/sites/kateoflaherty/2019/02/26/huawei-security-scandal-everything-you-need-to-know/>
- [28] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2016.119.01.0001.01.ENG
- [29] Article 29 Data Protection Working Party. 2010. Opinion 8/2010 on applicable law (0836-02/10/EN). *WP 179* (2010).
- [30] Anthony Quattrone, Lars Kulik, Egemen Tanin, Kotagiri Ramamohanarao, and Tao Gu. 2015. PrivacyPalisade: Evaluating app permissions and building privacy into smartphones. In *2015 10th International Conference on Information, Communications and Signal Processing (ICICSP)*. IEEE.
- [31] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. 2018. Apps, Trackers, Privacy, and Regulators: A Global Study of the Mobile Tracking Ecosystem. In *Proceedings 2018 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA.
- [32] Adam Satariano. 2018. G.D.P.R., a New Privacy Law, Makes Europe World’s Leading Tech Watchdog. *The New York Times* (May 2018). <https://www.nytimes.com/2018/05/24/technology/europe-gdpr-privacy.html>
- [33] Adam Satariano. 2020. Europe’s Privacy Law Hasn’t Shown Its Teeth, Frustrating Advocates. *New York Times* (April 2020). <https://www.nytimes.com/2020/04/27/technology/GDPR-privacy-law-europe.html>
- [34] Keith Spiller. 2016. Experiences of accessing CCTV data: The urban topologies of subject access requests. *Urban Studies* 53, 13 (2016), 2885–2900.
- [35] Anselm Strauss and Juliet Corbin. 1990. *Basics of Qualitative Research*. Sage Publications.
- [36] Jörg Thoma. 2014. Datenschutzbeauftragter mahnt mangelnde Verschlüsselung an. *Golem* (Sept. 2014). <https://www.golem.de/news/bayern-datenschutzbeauftragter-mahnt-mangelnde-verschluesselung-an-1409-109260.html>
- [37] Tobias Urban, Dennis Tatang, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2019. A Study on Subject Data Access in Online Advertising after the GDPR. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*. Springer, 61–79.
- [38] Siddharth Venkataramakrishnan. 2020. GDPR accused of being toothless because of lack of resources. *Financial Times* (April 2020). <https://www.ft.com/content/a915ae62-034e-4b13-b787-4b0ac2aaff7e>
- [39] Nicholas Vinocur. 2019. ‘We have a huge problem’: European tech regulator desponds over lack of enforcement. *Politico* (Dec. 2019). <https://www.politico.com/news/2019/12/27/europe-gdpr-technology-regulation-089605>
- [40] Paul Voigt and Axel von dem Bussche. 2017. Scope of Application of the GDPR. In *The EU General Data Protection Regulation (GDPR)*. Springer, 9–30.
- [41] Heinrich Amadeus Wolff and Stefan Brink (Eds.). 2018. *Datenschutzrecht in Bund und Ländern* (23 ed.). Beck, München.
- [42] Janis Wong and Tristan Henderson. 2019. The right to data portability in practice: exploring the implications of the technologically neutral GDPR. *International Data Privacy Law* 9, 3 (2019), 173–191.
- [43] Nils Zurawski. 2017. Exercising Access Rights in Germany. In *The Unaccountable State of Surveillance*. Springer, 109–133.

Part III

WEB TRACKING DETECTION AND SONIFICATION

6

Training Data for Automated Web Tracking Detection

6.1 Background

Besides the data collection in video games and mobile apps, another major type of modern user surveillance is *web tracking*, which refers to website operators and third parties storing information about the browsing activities of internet users. As with the sensor data discussed in Part I, diverse types of personal information (e.g., interests and demographics) can be inferred by analyzing web tracking data and linking the browsing activities of individual users across multiple pages [227]. In today's internet, web tracking is extremely prevalent. Based on an analysis of billions of web pages, Schelter and Kunegis [275] found that, depending on their content, 60 to 90% were actively tracking their visitors. While there are various benefits and legitimate reasons for web tracking (e.g., enable website functions, support free content, improve website usability), such surveillance of online activity entails serious privacy concerns.

Not only is it problematic that, due to the ubiquity of web tracking, it is difficult to effectively evade it and stay private while browsing online, but information collected through web tracking can also be used for many potentially harmful purposes, such as price discrimination, personalization of search results, determination of insurance coverage, assessment of financial credibility, government surveillance, and even identity theft [276]. Furthermore, there is a widespread lack of transparency, with companies often providing incomplete and/or misleading information about their web tracking activities [227, 277].

Of particular concern with regard to non-transparency and potential data misuse is third-party tracking, which is a form of web tracking performed by actors other than the actual operator of a visited website [278]. Here, the website operator (the “first party”) grants access to their users' activity data to other companies (the “third parties”) to incorporate certain services and functions (e.g., targeted advertising, web analytics, social media gadgets) into the website [279]. As Binns et al. write: “Third-party tracking potentially raises greater privacy concerns than first-party data collection because it can often link records from multiple

websites (...) to a single user. This can provide a more complete picture of an individual, from where they shop, to their social networks and likely political opinions” [279, p. 1f.]. Third-party tracking can also be present on highly sensitive websites that may reveal intimate details about the user, such as gay dating platforms or information pages on prostate cancer [280]. On many websites, users are being tracked by dozens of third parties at the same time [281]. In this way, third-party trackers can collect enormous amounts of personal information about individual users.

In this chapter, a paper is presented that contributes to the detection of web-tracking activity hidden to ordinary internet users. The paper proposes a novel browser extension called “T.EX – The Transparency EXtension”, which records browsing sessions in a secure and privacy-preserving manner. An implementation is presented and evaluated for its performance. The real-world browsing data collected by the extension can be used to feed machine learning algorithms for the automated detection of web trackers. As yet, artificial data is often used for this purpose, which has considerable drawbacks.

It needs to be acknowledged that web-tracking detection, even if done flawlessly, has its limitations as a method for privacy preservation. As the studies in Part I have exemplified, companies’ data practices can be highly obscure and inferences drawn from collected data can be varied and complex. Therefore, even if users had complete transparency as to when their browsing activities are tracked by whom, it would still not be apparent what possible inferences can be drawn about them and for which purposes this information is ultimately used. This general problem, which is of course not limited to the domain of web tracking, cannot be solved on the technical level alone but will also require regulatory changes (cf. Chapters 10.5, 10.6, and 10.9). Nonetheless, given that web tracking activities are mostly invisible to ordinary internet users, an effective and reliable detection approach is an essential step towards improved transparency and an important basis for critical public discourse on the matter, to which we hope to contribute with our proposed browser extension “T.EX”.

The project was conducted in collaboration with researchers from *Telekom Innovation Laboratories* and the faculty *Electrical Engineering and Computer Science* of Technische Universität Berlin, headed by the computer science Ph.D. candidate Philip Raschke. I contributed to the scanning of related work and identification of the research gap, to the conceptualization of the browser extension (incl. the statement of objectives) as well as to the testing and evaluation of the browser extension and the identification of its limitations. Our paper was presented at the *Annual Privacy Forum* (APF) 2019 in Rome, Italy, organized by the *European Union Agency for Cybersecurity* (ENISA).

Towards Real-Time Web Tracking Detection with T.EX - The Transparency EXtension

Philip Raschke^(✉), Sebastian Zickau, Jacob Leon Kröger, and Axel Küpper

Service-centric Networking, Weizenbaum-Institut, Telekom Innovation Laboratories,
Technische Universität Berlin, Berlin, Germany

{philip.raschke,sebastian.zickau,kroeger,axel.kuepper}@tu-berlin.de

Abstract. Targeted advertising is an inherent part of the modern Web as we know it. For this purpose, personal data is collected at large scale to optimize and personalize displayed advertisements to increase the probability that we click them. Anonymity and privacy are also important aspects of the World Wide Web since its beginning. Activists and developers relentlessly release tools that promise to protect us from Web tracking. Besides extensive blacklists to block Web trackers, researchers used machine learning techniques in the past years to automatically detect Web trackers. However, for this purpose often artificial data is used, which lacks in quality.

Due to its sensitivity and the manual effort to collect it, real user data is avoided. Therefore, we present T.EX - The Transparency EXtension, which aims to record a browsing session in a secure and privacy-preserving manner. We define requirements and objectives, which are used for the design of the tool. An implementation is presented, which is evaluated for its performance. The evaluation shows that our implementation can be used for the collection of data to feed machine learning algorithms.

Keywords: Web tracking · Browsing behavior · Data privacy ·
Browser extension · Data quality · Machine-learning ·
Classification algorithm

1 Introduction

There is no doubt that our Web browsing behavior is very sensitive. The websites we visit and the content we consume reveal information about our personality, our preferences, orientations, and habits. We give away our physical addresses, our phone numbers, and bank account information to use services or order goods. Simultaneously, the majority of websites nowadays integrates content from multiple external sources or third parties. Consequently, when visiting a website (also referred as first party) these third parties are given notice about our visit the moment our browser requests the external content. While our

© Springer Nature Switzerland AG 2019, Towards Real-Time Web Tracking Detection with T.EX - The Transparency EXtension, published 2019 in APF 2019, LNCS 11498, Springer, reproduced with permission of Springer Nature.
https://doi.org/10.1007/978-3-030-21752-5_1

physical address, phone number, or bank account information is not disclosed to these third parties, a link to the website we visited is.

The reasons for websites to integrate external content are manifold. Services embed images, audio, or videos without having to host or being allowed to host the content on an own server. But also many third-party scripts are integrated for various reasons. They are in particular critical, since their integration enables the execution of third-party code on the user's machine. There, they can access and gather information of the device and send it to a server where it is aggregated and analyzed. This way, a malicious third party can track every mouse movement, every key stroke, and every change of the scroll position of a user on a *different* website even without his or her awareness.

While on paper this sounds like a severe data security and privacy threat, this technique is widely used in the field of targeted advertising and Web analytics to track user behavior across multiple websites. In fact, Web trackers are an inherent part of the modern Web, because of their economic value for content providers and publishers. Websites display advertisements provided by ad exchanges or advertising networks in exchange for a payment per view or click. This way, each user of a website generates revenue.

While there is a variety of browser extensions that promise to tackle the issue, they are mostly blacklist-based, i.e. manual effort is required to identify trackers, which are then blocked (often by the domain name). This has four major disadvantages: (i) trackers can easily change their domain name, (ii) websites may offer relevant content or services, while also tracking user behavior (Amazon, Google, etc.), (iii) blacklists can be wrong, not complete, or outdated, and (iv) blocking requests to domains might create errors that prevent access to the desired content of the first party. The latter also occurs in the opposite causal direction, i.e. first parties block users from their content, if they block requests to third parties. Another conceptual flaw of blacklists is that they are not transparent themselves by providing little to no information on the third party in question and why it is blocked or not.

Consequently, an automated approach to detect Web trackers is desirable. This is a classification problem, which can be solved with machine learning techniques. However, machine learning approaches require rather large amounts of training data, which ideally is real data. However, researchers in this field often use bots to generate this data by crawling the Alexa.com top K websites. While this method produces large amounts of data rather quickly, it has a major drawback: it is not *real* data. These bots open the website, wait until it is finished loading, and then open the next in the list. These bots cannot log into websites like Facebook or Twitter, which even have implemented countermeasures for artificial users of their services. Even worse, the front page of these services are very limited and only offer a login form. It can be safely assumed that most of the third-party communication takes place after the login. By using bots, tracking of user interactions like moving the mouse, pressing a key, or scrolling is completely neglected.

For this reason, we present *T.EX: Transparency EXtension (T.EX)*, a secure browser extension to enable client-side recording, storage, and analysis of individual browsing behavior. With this tool researchers can generate data sets with real users in a secure, privacy-preserving, and user-friendly way. In this paper, we define requirements concerning security, privacy, and usability and explain how they were met. In addition, the extension provides data visualization capabilities allowing (experienced) users to assess their browsing behavior and the third-party communication involved in it.

The remainder of this paper is structured as follows: Sect. 2 defines the objectives and requirements of the tool. Section 3 elaborates on the limitations and the derived design decisions. Section 4 gives an overview of related work and assesses whether suitable solutions already exist. Section 5 presents the implementation of the tool. In Sect. 6, we evaluate the tool with regard to the specified objectives. Finally, a conclusion is given including an outlook.

2 Objectives and Requirements

As stated above, the main objective of the tool is to enable the generation of *real* user data in a secure, privacy-preserving, and user-friendly manner by allowing users to record browsing sessions. On this basis, we derive the following objectives:

- Obj1** The tool needs to be able to monitor Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) traffic, including header information, parameters, and the body.
- Obj2** An accurate differentiation between first and third party must be realized. The first party should not be identified only by its host name but rather by the actual page (HTTP path) the user visited.
- Obj3** The network traffic must be persistently stored for a certain amount of time. This data must be securely (i.e. encrypted) stored on the user's device, so no other (malicious) software on the user's machine can access it.
- Obj4** The extraction of data must be in a privacy-preserving manner, i.e. only relevant data should be collected. Furthermore, no external servers must be involved.
- Obj5** The user must be able to completely delete the data at any time. There should be a means to prove the erasure of the data.
- Obj6** Furthermore, the user must be able to export the data in a machine-readable format.
- Obj7** The user must be able to disable the recording of network traffic at any time. Ideally, the user can be given a guarantee or proof that the recording is stopped.
- Obj8** Usage of the tool should be user-friendly to the extent that the perceived Quality of Experience (QoE) is not impacted by it.

Obj9 The tool must offer data visualization capabilities so that users can review the recorded data before they export it. A search function enables users to check if any sensitive information are contained within the data set.

3 Limitations

Unfortunately, the above defined objectives cannot be realized without constraints. In this section, we infer limitations from these objectives and elaborate on consequent design decisions for the tool.

In order to realize Obj1, HTTP and HTTPS traffic needs to be intercepted. Obviously, this is a severe data security risk and infringement of the user's privacy. For this reason, the collected and recorded data must remain on the user's device (see Obj4). However, intercepting HTTPS traffic on the network layer is not possible without aggressive intervention. A *man-in-the-middle* attack could be used in order to intercept the encrypted traffic, but this would put the user's overall data security at risk.

Fortunately, we can rely on capabilities offered by browser vendors. Experienced users or system administrators have the expertise to obtain the data using the browser's developer tools like Google Chrome's *DevTools* or the *Inspector* of Firefox. However, the data, that is logged there, is separated from other browser sessions (tabs). Consequently, for a holistic view, an aggregation of the data is required. The user would need to open the corresponding tool before the begin of each browsing session in each tab. The log is cleared with every new page the user visits, so a checkbox needs to be ticked to persist the log (in each tab). To export the recorded data, only Firefox' *Inspector* offers a complete export of the data, while Chrome's *DevTools* only offer an option to export one request at a time. Collecting data using this method is cumbersome and error-prone, which violates Obj8. Further inspection of this method also revealed that Obj2 is violated, since the exported data either does not contain the first party (Chrome) or only gives the host name of it (Firefox).

Clearly, a more sophisticated method is required. Luckily, HTTP and HTTPS traffic can be logged using Chrome's or Firefox' extension Application Programming Interface (API). So, Obj1 can be best implemented in a browser extension. In fact, we found no alternative approach to realize Obj1 without aggressively interfering with the user's device. Using the extension API also allows us to identify the first party including the HTTP path (see Obj2). Besides an *initiator* field in the traffic log, it is possible to map a request to a certain open and active tab of which the URL can be used.

To persistently store the data like stated in Obj3, a sophisticated database like *MySQL* or *MongoDB* would be ideal, however this would require users to install additional software on their device (violation of Obj8) or to transmit the data to an external server (violation of Obj2). Browser extensions are able to store data in the so-called *local storage*, which offers limited storage capabilities. The local storage is a key-value store, thus complex queries cannot be

easily expressed. Furthermore, the local storage is not encrypted, thus malicious software on the user’s device could easily gain access to it. Therefore, encryption must be implemented within the browser extension. However, inconvenient key-pair generation and management must be avoided in order to not violate Obj8.

In order to realize a collection of data in a privacy-preserving manner (Obj4), only the outgoing traffic is recorded. This way, we follow a data minimization approach. The HTTP response, besides the actual content the user consumes, contains cookies and identifiers that are assigned to the user and which are used for subsequent requests. By neglecting the HTTP response, we miss these assignments. However, we assume the preserved privacy is of higher value than the benefit gained from the HTTP responses. Moreover, it is not sure whether the accuracy of a classification algorithm to detect Web trackers would be increased if the HTTP response is taken into consideration. It would be interesting to investigate this in a separate study.

Since the HTTP body is used to transmit sensitive data like passwords, messages, photos or videos, recording it can be highly sensitive. Therefore, it is not recorded by default but the user is able to enable this feature at own risk. The reason why we do not completely exclude it, like we do with the HTTP response, is that we could observe Web trackers using it for passing identifiers to their servers.

The local storage can be cleared at any time; therefore, the user is given a button to trigger the erasure of all data (Obj5). Moreover, the local storage is file-based, i.e. its content can be found in plain text in files on the user’s machine. Thus, to ensure the erasure of all personal data, the user can additionally delete the corresponding files. The path to these files is static, it can be given to the user so he or she can find it.

To export the data in a machine-readable format (Obj6) the whole local storage must be queried, requests must be decrypted, and saved to a dedicated file. Since data in the local storage is in JSON format, it is reasonable to export it as such. Due to the diverse structure of the recorded data, an export in CSV is rather unhandy.

Disabling the recording (Obj7) can be realized with a set of means: by implementing blacklists (or whitelists), by offering a button to start and stop recording at any time, or by disabling the extension completely. The latter is undoubtedly the safest and easiest way to guarantee that the recording is disabled. Blacklists or whitelists determine on which websites recording should be disabled or enabled respectively. This approach, however, requires users to invest some effort for pre-configuration, which might violate Obj8. A button to start and stop recording is rather easy to implement, but offers no advantage compared to enabling or disabling the extension, since this can be triggered with one click as well.

To achieve Obj8, all other objectives must be realized by involving as less user effort as possible. This means that the usage of the extension itself is realized in a user-friendly manner. But furthermore, the usage of the extension should not impact the perceived QoE while browsing the Web, i.e. websites should not

take longer to load or that CPU and memory consumption drastically increase so that other applications are affected.

The visualization of the data (Obj9) can be done in the browser using Hyper-text Markup Language (HTML), Cascading Style Sheets (CSS), JavaScript, and Scalable Vector Graphics (SVG). To highlight the communication flows, we chose a graph representation of the data. A search function is provided to users allowing them to query the data for personal information they do not want to be included in a resulting data set, which is further processed.

4 Related Work

Trackers enjoy a long presence in the history of the Web. In fact, they exist almost as long as the Web itself. Lerner et al. [11] proved the presence of Web trackers in 1996 by examining and analyzing the Web Archive. The Internet, as a distributed system, is built upon interconnections of nodes, thus, third parties are conceptually nothing to despise. However, for the precise personalization of displayed advertisements, personal data is required, which is often collected without a user's awareness using Web tracking techniques. One could argue that the most severe issue with third-party content is not its presence but users' unawareness of it. A study by Thode et al. [14] shows that users' expectations regarding third-party tracking heavily differ from reality. With the General Data Protection Regulation (GDPR) [7] coming into effect in May 2018, this circumstance becomes problematic, since it requires the processing of personal data to be transparent.

Bujlow et al. [2] published a sophisticated survey on all known Web tracking techniques to date. Most modern and often more accurate methods mostly rely on third-party scripts that are executed on the user's device to obtain a set of data items to generate a so-called *browser fingerprint*, which is sufficient to uniquely identify the user among other users.

Today Web trackers are subject to extensive studies due to the threat they impose on our data privacy. A very sophisticated study was conducted by Englehardt et al. [6] in 2016, who aimed to measure and analyze the extent of third-party presence on one million websites. Therefore, they designed and developed the tool OpenWPM to measure and record HTTP traffic. Yet, OpenWPM uses Selenium to crawl the top one million websites, which is a framework to simulate and automate user interactions. Thus, their measured data is not real user data. Regardless of the data quality, they found third-party scripts present on nearly all considered websites. Their results further show that only few third parties are present on a high number of first parties. This is clear evidence for data monopolies of the most prominent Web trackers. However, this circumstance is also an advantage: one has to identify and block the few most prominent third parties only to effectively protect oneself from Web tracking on the most popular websites at least. This is one of the reasons why the blacklist-based approach is so popular: it is very effective.

There are many browser extensions for all major browsers that follow this approach. Their promise is to protect users from unintended and unauthorized

third-party information disclosure. Browser extensions like Ghostery [10], Ultra-Block - Privacy Protection & Adblocker [16], Crumble [4], or Privacy Badger [13] are very popular tools with millions of users. However, only Privacy Badger tries to identify Web trackers based on their prominence in addition to blacklists. Privacy Badger blocks a third party if its presence is observed on three distinct first parties. An additional challenge of these browser extensions is to maintain the same level of user-perceived QoE after the extension has been installed. From a user’s perspective, blocking third-party requests is very beneficial, since loading times are decreased and computing resources are spared, as a study of Kontaxis and Chew [8] confirms.

However, the above presented browser extensions give little to no information on the tracking third party itself nor technical details about the process of data exposure. However, there are browser extensions that give more information: uMatrix [17] and uBO-Scope [15]. The extension uMatrix provides the user with insights on the type of HTTP requests issued to the corresponding third parties. While, to our knowledge, the extension uBO-Scope is the only one that accurately gives information on the extent of presence of a specific third party during the current browsing session. A high presence of a third party is indicated with red in the extension’s pop-up window.

Nonetheless, all the above presented browser extensions rather aim to identify and block tracking activities than serving as tool to assess data flows to third parties. They offer limited data visualization capabilities and no recording options, which makes it difficult to analyze or further process the measured data. The browser extension closest to the objectives of T.EX is Firefox’ Lightbeam [9], which has strong visualization features (Obj9), but fails to give more insights on the communication that has taken place and the third parties itself (Obj1). Lightbeam allows to export the recorded data in machine-readable format (Obj6), yet the exported information does not include the first party with its HTTP path (Obj2).

The idea to use machine-learning techniques to identify Web trackers was proposed by Bau et al. [1] in 2013. They elaborate on useful data sources and how to obtain labeled training sets. Following the paper’s position, there were several publications of researchers in the following years describing supervised or unsupervised classification of Web tracking activities. In 2014, Metwalley et al. [12] present an unsupervised approach that leads to successful results. Their algorithm is able to detect 34 Web trackers that have never been documented before. Similar results are achieved by Wu et al. [18] in 2016. They use a supervised approach and detect 35 new tracking parties. Despite their successful revelation of new Web trackers, both research groups use crawlers to generate the data with which they feed their machine-learning algorithms.

The importance of proper data quality is highlighted by the publication of Yu et al. [19], who achieve remarkable results with regard to accuracy and performance of detecting Web trackers. The authors are a research group from the Cliqz browser development team, which is a German browser vendor of the same-named browser Cliqz [3]. Through their product, they were able to use browsing

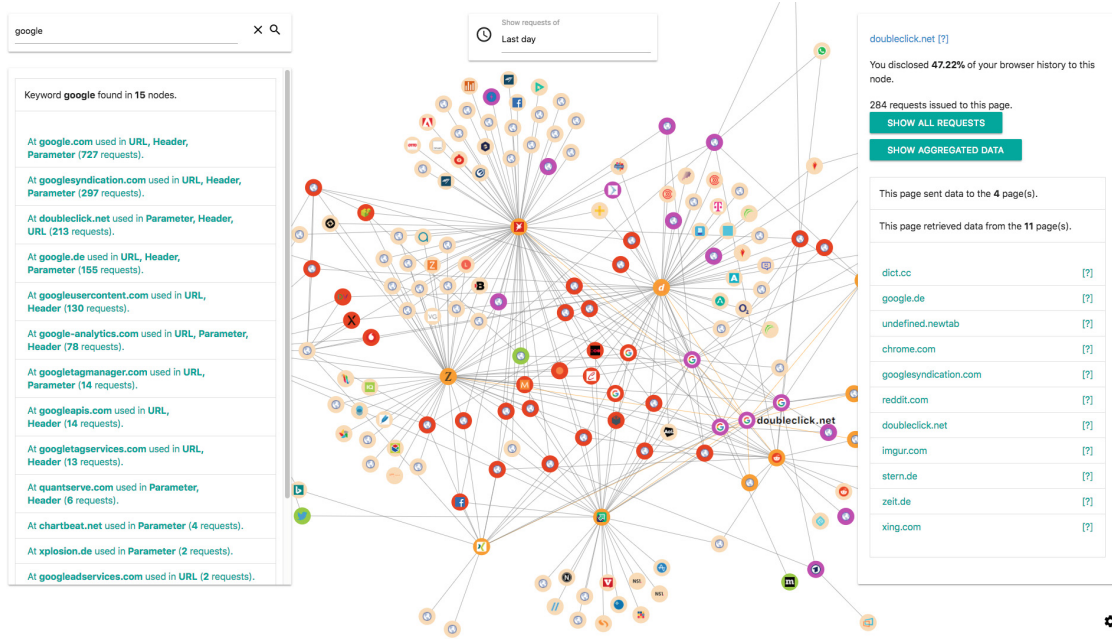


Fig. 1. The user interface of the browser extension including a graph, a search feature, and further information on the third parties. Highly connected nodes are colored red to indicate third parties with high extent of presence on other websites. (Color figure online)

data of 200.000 users for their algorithm. This way, they were able to outperform their commercial competitor Disconnect.me [5], which is also used by Firefox.

5 Implementation

This chapter presents the implementation of T.EX and explains how the individual objectives were realized. T.EX has been implemented for Google Chrome, however it is planned to port the implementation to Mozilla Firefox. Since the offered browser extension APIs of the two browser vendors are based on the WebExtension APIs, it can be expected that most of the code can be reused for the implementation of a Firefox extension.

5.1 HTTP and HTTPS Traffic Logging and Recording

To intercept and log HTTP and HTTPS traffic, the interface *webRequest* is used. Chrome and Firefox emit an event *onBeforeRequest* before a request is issued. Extensions can subscribe to the event by adding a listener to it. Both browsers provide extensions with valuable information on the issued request, including all necessary information on the target t of the request, search parameters S , request headers H , form data F and even data in the request body B . Interestingly, determining the source s of a request requires more effort in Google Chrome. While Firefox emits the initiator of a request in the *originUrl* field, Chrome only

Fig. 2. Records visualized on a timeline enabling users to investigate requests initiated by a certain website to a certain third party. By selecting an event, users can see the corresponding record including all recorded data.

gives information on the source in an optional field called *initiator*. To retrieve the source even if the field is not set, a query of open tabs with the *tabId* is required. A logged event is called *record* *r*, which is defined as follows:

5.2 Persistent Storage of Records

For this reason, two strategies are implemented: the aggregation of records into chunks and the writing of chunks into the local storage in a defined interval i . This way, the local storage is less demanded and the work load is evenly distributed over time. However, these strategies raise the question of appropriate keys that can be used for the chunks, so that they can be queried later in time.

To enable this, we implement a chain of chunks C , i.e. each chunk c is pointing to the last chunk and the key of the most recent chunk is stored in a global field called *currentId*. Each chunk retrieves a timestamp ts , which is used as key for the chunk.

$$c \in C := (ts, lastId, R_{[ts-i, ts]}) \quad (4)$$

$$currentId = ts \quad (5)$$

Eventually, this implementation enables queries of chunks in a certain time range. Moreover, this implementation allows the erasure of old chunks after a predefined time. Given that the local storage by default is limited to 5.24 megabytes, this feature is crucial. Both Chrome and Firefox have the extra permission *unlimitedStorage*. Extensions that ask for the privilege are allowed to store more data. Nonetheless, an implementation that does not rely on the permission is desirable.

5.3 Encryption and Decryption of Chunks

Since the local storage resides on the user's machine unencrypted, encryption needs to be implemented in order to ensure data security. Otherwise, a malicious application on the user's device could gain access to this data and gain valuable information like passwords, the browser history, email addresses, bank account information and suchlike. Without encryption, T.EX would rather constitute a severe risk than contribute to improved data security and privacy.

To implement encryption, the user is prompted to generate a key pair (*pubKey* and *privKey*) after the installation of the browser extension. This requires the user to enter a password *pwd*. The generated private key is encrypted with the entered password using the Advanced Encryption Standard (AES). The generated public key and the encrypted private key *encPrivKey* are then stored in the local storage.

To encrypt chunks, a random key *aesKey* is generated that serves as symmetric key for the encryption. This random key is used for the whole browsing session until the browser is closed. This key is encrypted with the public key so that only the private key can decrypt it. This encrypted symmetric key *encAesKey* is stored along with the encrypted chunk in the local storage. To decrypt chunks, the user is prompted to enter the password to decrypt the private key, which is then used to decrypt the symmetric key to eventually retrieve the chunks.

5.4 Data Visualization

As it can be seen in Fig. 1, data flows are represented by a graph $G := (V, E)$, which illustrates connections between visited websites (green-colored nodes) and involved third parties (beige or red-colored nodes). Red-colored nodes are highly connected nodes that retrieve data from various websites and Web applications. For the coloring, a rather simple rule-based approach was used for the beginning. However, it is planned to extend the coloring function at a later point in time.

Algorithm 1. Set-up and encryption of chunks

```

1:  $privKey, pubKey \leftarrow generateKeyPair()$ 
2:  $pwd \leftarrow$  user-entered password
3:  $encPrivKey \leftarrow encrypt(privKey, pwd)$ 
4:  $save(encPrivKey, pubKey)$ 
5:  $c = (ts, lastId, R_{[ts-i, ts]})$ 
6:  $aesKey \leftarrow generateRandomKey()$  for each session
7:  $encAesKey \leftarrow encrypt(aesKey, pubKey)$ 
8:  $c' \leftarrow (ts, lastId, encrypt(R_{[ts-i, ts]}, encAesKey), encAesKey)$ 
9:  $save(c')$ 
    
```

Algorithm 2. Decryption of chunks

```

1:  $encPrivKey \leftarrow$  load from local storage
2:  $pwd \leftarrow$  password prompt
3:  $privKey \leftarrow decrypt(encPrivKey, pwd)$ 
4:  $c' \leftarrow$  load from local storage
5:  $aesKey \leftarrow decrypt(c'_{encAesKey}, privKey)$ 
6:  $c \leftarrow (ts, lastId, decrypt(R_{[ts-i, ts]}, aesKey))$ 
    
```

A more gradient color function is currently researched to highlight only the Web trackers in the graph.

$$G := (V, E) \tag{6}$$

$$V := \{r_s, r_t | r \in R\} \tag{7}$$

$$E := \{(r_s, r_t) | r \in R\} \tag{8}$$

Users can search for keywords that might appear in URLs, headers, or parameters. Purple-colored nodes (as seen in Fig. 1) are nodes that contain the keyword in the record. By clicking on a node the user is able to retrieve more information on the corresponding node such as to which nodes data has been sent to or from which nodes data was retrieved. For further investigation of the occurred communication, the user can investigate requests to or from one node, which are visualized on a timeline. By selecting an entry on the timeline the record is visualized (see Fig. 2).

6 Evaluation

The aim of this section is to evaluate whether the usage of T.EX implies an unneglectable impact on the user-perceived QoE while browsing the Web. Therefore, we investigate whether the loading time of a website noticeably increases, when using T.EX. We measure loading times by recording key events: *onDOMContentLoaded* and *onCompleted*. Both events occur strictly sequential, i.e. the *DOMContentLoaded*, which indicates that the Document Object Model (DOM) is fully built, always occurs before *DOMContentCompleted*, which indicates that

also all referenced resources are fully loaded and initialized. From a user’s perspective, the first event occurs close to the moment when the user is able to see the website. In contrast to the latter, which is triggered when the loading indicator of the browser disappears.

Analogously, we measure the resource consumption (i.e. CPU and memory usage) during a website request and loading in order to learn the impact of the browser extension on hardware resources. For this purpose, we request and compute CPU and memory usage in a determined interval (so-called tick each 50 ms). Besides CPU and memory usage, we further evaluate the disk space consumption of T.EX on a general level to find out how fast the extension reserves disk space for its purpose.

As stated above, we open websites with and without T.EX activated. We additionally repeat the procedure with a different, comparable browser extension activated in order to be able to assess the performance of T.EX in comparison with other extensions. For this purpose we identified Privacy Badger as good candidate, since it uses the same APIs to analyze traffic in real-time. However, we know that Privacy Badger decreases loading times of websites, while we expect T.EX to increase loading times. This is due to Privacy Badger preventing HTTP requests from occurring, thus saving time to load, while T.EX logs, processes, and stores HTTP requests. For both hardware resources are used. With this evaluation procedure we aim to put the increased hardware usage of T.EX into perspective.

As appropriate websites for the test, we use the German news site *spiegel.de* and the front page of *google.de*, which differ in the amount of third-party content they integrate. While accessing *google.de* triggers *only* 23 requests, which only request content from Google servers, requesting *spiegel.de* involves more than 400 requests to more than 50 third parties. We expect hardware usage and loading times to increase linearly with the number of involved requests, thus we selected two websites that are rather bipolar in that respect. The experiment was conducted on a machine with an Intel Core i7 (2.2 GHz quad-core) and 16 GB memory. The machine was connected to the Internet via a 1 Gbit Ethernet connection. The experiments were repeated three times each to detect anomalies.

The results of the experiment are depicted in Fig. 3. The rows represent the corresponding runs without T.EX activated (top), with T.EX activated (middle), and with Privacy Badger activated (bottom). In each run the CPU usage (left column), memory usage (middle column), and loading times (right column) were measured.

By comparing the individual results displayed in the first column, an increase of CPU usage is clearly observable. The CPU is working much closer to capacity and maintains this level during the whole time the website is loaded. The reason for the CPU demand of T.EX is found in the steady encryption of records in the background. Thus, disabling the encryption would gain performance, yet would constitute a violation of the extension’s main objectives. Additional CPU capacity is used, since requests are preprocessed before they are stored in the local storage. This preprocessing could be executed at a later point in time, for

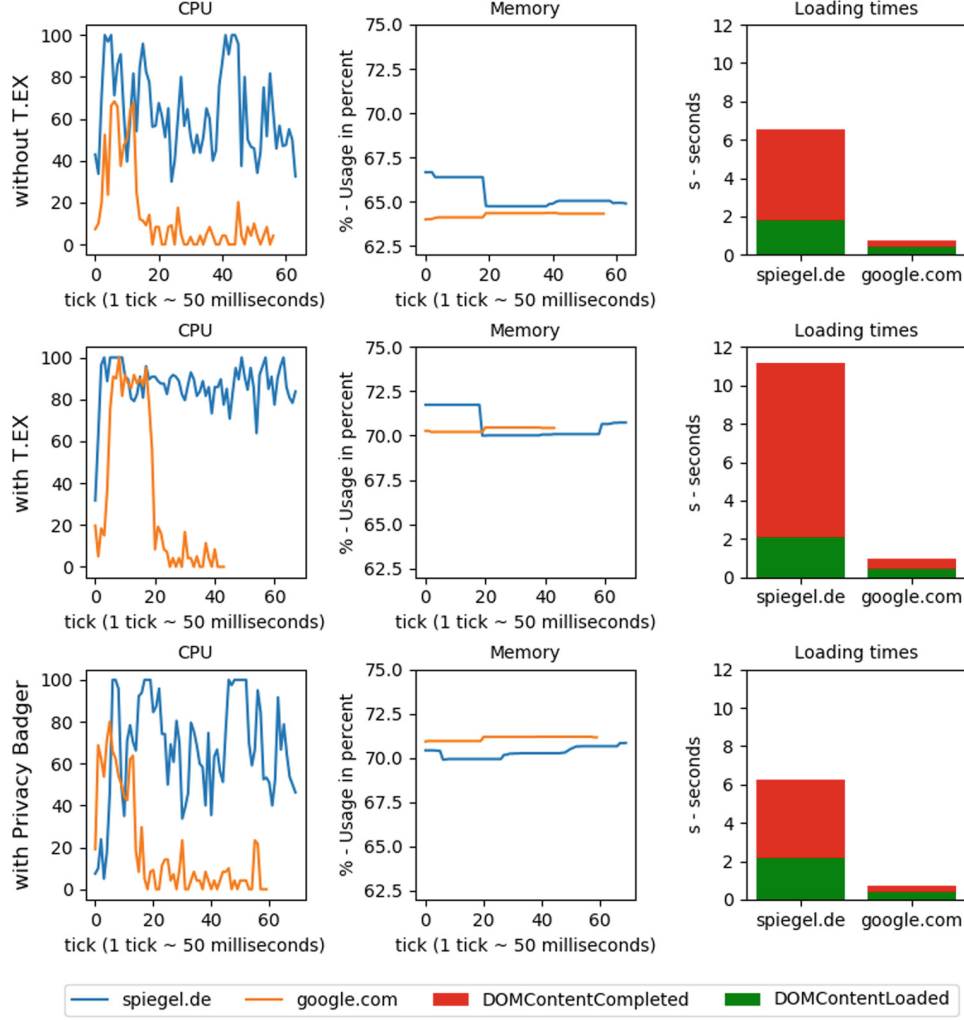


Fig. 3. The results of the evaluation: the first column shows the CPU usage, the second column the memory usage, and the third column the loading times. The first row represents the measurements without T.EX activated, the second row with enabled T.EX, and the last row with Privacy Badger activated.

example, when the browser is in the idle state for a certain amount of time, i.e. the browser is currently not used by the user.

The memory consumption is rather consistent with our expectation: the usage is increased fairly but not excessively. Comparable browser extensions like Privacy Badger that perform similar tasks show the same level of memory consumption. The perceived QoE should not be affected to much by this circumstance. In contrast to the loading times, which seem to be strongly affected by the usage of T.EX. When comparing the third column in Fig. 3, it is noticeable that the loading time is drastically increased, when T.EX was activated. This does not apply on the *DOMContentLoaded* event, but on the *DOMContentLoaded* event. Note that the page is usable much earlier, so that the user can already interact with it, before the DOM content is fully loaded. Yet the performance of T.EX with regard to loading times requires improvement. It is also noteworthy

that the performance for the loading times of *google.de* are comparable to the performance achieved in the other runs. Consequently, the drastic increase of the loading time occurs on websites with massive third-party involvement. An exponential increase relative to the number of involved third parties could be ruled out.

Finally, we aim to investigate the disk space consumption. While it can be measured easily by simply checking how big the local storage files are, it is rather difficult to define a rule to estimate the storage usage. In general, it heavily depends on the usage and browsing behavior of the user. In a dedicated three-hour lasting session, we were able to collect 80 megabyte of data, while on a different machine that is exclusively used during office hours (then extensively), we collected almost 700 megabyte in a single month. Nonetheless, it must be stated that the storage requirements imposed by the usage of T.EX exceed the requirements of other browser extensions. Therefore, users of T.EX must be aware that the recording of browsing sessions is storage intensive.

7 Conclusion and Outlook

This paper presents T.EX a browser extension to provide transparency to experienced users or system administrators, who want to record and analyze communication flows to external third parties while browsing the Web. Therefore, objectives and requirements have been defined and their implementation has been presented. T.EX will serve as tool to conduct measurements and obtain real user data in a secure and privacy-preserving manner, which might contribute to more accurate machine learning models to identify Web trackers and tracking activities in real-time. We evaluated T.EX by measuring its impact on the performance to derive consequences on the user-perceived QoE. Our results show that T.EX achieves performance, which is comparable to other privacy browser extensions like Privacy Badger. However, it has an impact on the loading times of certain websites that cannot be neglected. The issue will be investigated in future works. Furthermore, we will use T.EX to collect data that will be used to identify trackers and their tracking activities.

Acknowledgments. Supported by the European Union’s Horizon 2020 research and innovation programme under grant 731601.

References

1. Bau, J., Mayer, J., Paskov, H., Mitchell, J.: A promising direction for web tracking countermeasures. In: Workshop on Web 2.0 Security and Privacy (2013)
2. Bujlow, T., Carela-Espanol, V., Lee, B.R., Barlet-Ros, P.: A survey on web tracking: mechanisms, implications, and defenses. *Proc. IEEE* **105**(8), 1476–1510 (2017)
3. Cliqz - Der sichere Browser mit integrierter Schnell-Suche. <https://cliqz.com/>. Accessed 4 Feb 2019

4. Crumble - Online Privacy, Stop Tracking. <https://chrome.google.com/webstore/detail/crumble-online-privacy/icpfjjckgkocbkdaodapelofhgjncoh>. Accessed 4 Feb 2019
5. Disconnect. <https://disconnect.me/>. Accessed 4 Feb 2019
6. Englehardt, S., Narayanan, A.: Online tracking. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security - CCS 2016, no. 1, pp. 1388–1401 (2016)
7. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ L 119, pp. 1–88, 4 May 2016
8. Kontaxis, G., Chew, M.: Tracking protection in Firefox for privacy and performance. In: IEEE Web 2.0 Security & Privacy, June 2015
9. Firefox Lightbeam - Add-ons for Firefox. <https://addons.mozilla.org/de/firefox/addon/lightbeam/>. Accessed 4 Feb 2019
10. Ghostery Makes the Web Cleaner, Faster and Safer! <https://www.ghostery.com>. Accessed 4 Feb 2019
11. Lerner, A., Simpson, A. K., Kohno, T., Roesner, F.: Internet Jones and the raiders of the lost trackers: an archaeological study of web tracking from 1996 to 2016. In: Usenix Security (2016)
12. Metwalley, H., Traverso, S., Mellia, M.: Unsupervised detection of web trackers. In: IEEE Global Communications Conference (GLOBECOM), pp. 1–6 (2015)
13. Privacy Badger - Electronic Frontier Foundation. <https://www.eff.org/privacybadger>. Accessed 4 Feb 2019
14. Thode, W., Griesbaum, J., Mandl, T.: I would have never allowed it: user perception of third-party tracking and implications for display advertising. Re:inventing information science in the networked society. In: Proceedings of the 14th International Symposium on Information Science (ISI 2015), Zadar, Croatia, 19th–21st May 2015, vol. 66, pp. 445–456 (2015)
15. BO-Scope: a tool to measure over time your own exposure to third parties on the web. <https://github.com/gorhill/uBO-Scope>. Accessed 4 Feb 2019
16. UltraBlock - Block Ads, Trackers and Third Party Cookies. <https://ultrablock.org/>. Accessed 4 Feb 2019
17. Matrix: point and click matrix to filter net requests according to source, destination and type. <https://github.com/gorhill/uMatrix>. Accessed 4 Feb 2019
18. Wu, Q., Liu, Q., Zhang, Y., Liu, P., Wen, G.: A machine learning approach for detecting third-party trackers on the web. In: Askoxylakis, I., Ioannidis, S., Katsikas, S., Meadows, C. (eds.) ESORICS 2016. LNCS, vol. 9878, pp. 238–258. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45744-4_12
19. Yu, Z., Macbeth, S., Modi, K., Pujol, J. M.: Tracking the trackers. In: Proceedings of the 25th International Conference on World Wide Web - WWW 2016, pp. 121–132 (2016)

7

Making Web Tracking Audible

7.1 Background

There has been extensive research on methods to uncover and defend against unwanted web tracking [276, 282], including our contribution presented in Paper 8. But how can real-time information about on-going web tracking best be conveyed to internet users for purposes such as user studies, data collection transparency, and privacy education?

While global internet usage is steadily increasing [283], exposing more and more people to online surveillance, studies show that many internet users are highly unaware about the pervasiveness and possibilities of web tracking [227, 284]. Although consumer education alone will not suffice to ensure fair information processing, insights and transparency are necessary to foster a basic awareness of the problems associated with ubiquitous surveillance and support an informed public debate on privacy matters (cf. Chapter 9.1). As explained in Chapter 6.1, companies often try to obfuscate their web tracking activities and the resulting data can be used against the data subjects in various ways, which further underlines the importance of improving transparency. One obvious approach for exposing hidden web tracking activity to internet users is visualization [285], as is also well illustrated in Fig. 1 of Paper 8. But other types of sensory input may also be worth exploring – not only to assist visually impaired people, but also because learners vary in their preference for different modalities (e.g., visual vs. auditory) [286] and because visual attention is limited during browsing sessions as users are typically busy looking at web content.

Entering a young and largely untouched field of research, the paper included in this chapter examines ways in which web-tracking activity can be “sonified” for users, i.e., made audible through indicative melodies, sounds, and whispering voices. Compared to the few existing approaches on web-tracking sonification in non-professional settings [287, 288], we added new features and also conducted a small experiment to test the effect of our sonification approach on users’ perceptions.

The project was a collaboration of researchers from the Weizenbaum Institute, the Fraunhofer Institute for Open Communication Systems, Technische Universität Berlin, and Humboldt-

Universität Berlin, headed by the human-computer interaction expert and doctoral researcher Otto Hans-Martin Lutz. I contributed to researching related work and identifying the research gap, to conceptualizing the sonification tool, to the tool’s testing and evaluation, and to the writing and critical revision of the manuscript. The resulting paper was published in the proceedings of the *International Conference on Auditory Display* (ICAD 2019).

It may be worth noting that, although not included in this thesis, we have continued our work in the area of privacy-related sonification, as exemplified by our recent publication “That Password Doesn’t Sound Right” [289], where we propose a novel approach for interactive password strength sonification.

SURFING IN SOUND: SONIFICATION OF HIDDEN WEB TRACKING

Otto Hans-Martin Lutz^{abc}, Jacob Leon Kröger^{ac}, Manuel Schneiderbauer^{ad} and Manfred Hauswirth^{abc}

^a Weizenbaum Institute for the Networked Society, Berlin

^b Fraunhofer FOKUS

Kaiserin-Augusta-Allee 31, 10589 Berlin, Germany

{otto.lutz, manfred.hauswirth}@fokus.fraunhofer.de

^c Technische Universität Berlin ^d Humboldt-Universität Berlin

ABSTRACT

Web tracking is found on 90 % of common websites. It allows online behavioral analysis which can reveal insights to sensitive personal data of an individual. Most users are not aware of the amount of web tracking happening in the background. This paper contributes a sonification-based approach to raise user awareness by conveying information on web tracking through sound while the user is browsing the web.

We present a framework for live web tracking analysis, conversion to Open Sound Control events and sonification. The amount of web tracking is disclosed by sound each time data is exchanged with a web tracking host. When a connection to one of the most prevalent tracking companies is established, this is additionally indicated by a voice whispering the company name. Compared to existing approaches on web tracking sonification, we add the capability to monitor any network connection, including all browsers, applications and devices.

An initial user study with 12 participants showed empirical support for our main hypothesis: exposure to our sonification significantly raises web tracking awareness.

1. INTRODUCTION

Web tracking collects information about a particular user's activity on the World Wide Web. It is widely used, with some form of web tracking found on 90 % of common websites, and on 60 % of websites with highly privacy-critical content [1]. Although complex and extremely diverse, the ecosystem of web trackers is dominated by a small number of companies, notably by Google, Facebook and Amazon, who are inconspicuously present as third-party data collectors on many websites [2]. Recent empirical results suggest that third-party scripts owned by Google alone are present in about 80% of web traffic of the top 600 websites, and are used in a tracking context in about 40 % [3].

Since a person's browsing behavior reveals insights into his or her personality, habits and sensitive aspects such as financial and

medical situation or political views, web tracking may constitute a serious privacy threat [4]. Even though web tracking is seen unfavorably by the majority of internet users due to privacy concerns [5], they do not understand the full extent, the methods and possibilities of online behavioral tracking [6].

Web tracking is invisible to the user by design. Studies show that there is no sufficient awareness of web tracking [7]. We use sonification of clandestine web traffic to tracking providers as a means of raising awareness for online privacy issues. If visualization is used instead for the same objective, users must divert their visual attention from their primary task (surfing the web). Using the auditory domain, we can simultaneously communicate information in a different modality, which provides additional attention and workload resources [8]. Furthermore, sonification is suitable to present temporal data in real-time and can be shaped to convey emotional content [9, p.11, p.92].

Our contribution is a sonification-based approach to raise user awareness of web tracking which extends the possibilities of existing approaches like *Soundbeam* by Hutchins et al. [10]. We describe a framework for live web tracking analysis and conversion to OSC¹ events, which can be used to monitor web tracking on any network connection – across all kinds of browsers, apps and devices. We discuss our system, sonification and sound design. Finally, we present results of an initial user study with 12 participants. We found empirical support for our main hypothesis: exposure to the sonification significantly raised web tracking awareness.

2. RELATED WORK

There is a comprehensive body of work on using sonification for network traffic monitoring to achieve higher situational awareness in a network operations center (e.g., [11, 12], systematic overview in [13]). In this context, users are network security specialists which use the auditory modality as supplementary resource to achieve their objectives in pattern, anomaly and intrusion detection. The scope of our approach, however, focuses on the average user who, in contrast to network operations professionals, is often unaware of the extent of web tracking [6]. Here, awareness refers to a general consciousness on the prevalence of web tracking. Sonification of web tracking can increase this awareness as it

This work has been funded (in part) by the Federal Ministry of Education and Research of Germany (BMBF) under grant no. 16DII111 ("Deutsches Internet-Institut").



This work is licensed under Creative Commons Attribution Non Commercial 4.0 International License. The full terms of the License are available at <http://creativecommons.org/licenses/by-nc/4.0>

¹Open Sound Control, a network-based protocol for sound and media control: <http://opensoundcontrol.org>

provides immediate auditory feedback to the user while he or she is browsing the internet.

Soundbeam [10] sonifies third-party connections extracted by Mozilla Lightbeam, a plug-in for the Mozilla Firefox browser. It sends data on intentionally visited websites and unintentionally visited third-parties (e.g., analytics or advertisement providers) to the SuperCollider synthesis engine via OSC. Soundbeam is designed for ensemble performance. Several users can run the software on different computers in the same network. When user B encounters a third-party element that has been identified by user A before, it is sonified for both users. This is intended to “highlight both the ubiquitousness and interconnectedness of tracking” [10].

Another related project is an earcon-based sonification of internet security threats for vision-impaired users [14]. Here, warning sounds that convey their intended meanings with little-to-no user training (e.g., casting a fishing reel to warn about a phishing attack) were used to notify users about security threats while browsing on a screen reader.

3. FRAMEWORK DESIGN

Our software runs in the background while the user is browsing the web. The framework comprises four stages: (1) monitoring network traffic, (2) filtering for connections to known web trackers, (3) extracting different kinds of tracking-related events, and (4) sending these events to the sound generator via OSC (see Figure 1).




Figure 1: System overview

In the prototyping phase, we used Ableton Live [16], with a Max for Live OSC receiver for sound synthesis. We aim to switch to cross-platform (Linux supported) open source software in the future.

3.1. Implementation

In order to be able to intercept any network connection, we use Python to create several instances of TShark processes, a text-based version of the network protocol analyzer Wireshark [15]. These processes listen to the traffic of the selected network connection. They are configured with filter lists of web tracker IP addresses, so only traffic to these addresses is analyzed in the following steps.

Tracker identification: Connections to tracker services are detected by tracker identification lists available from different sources (e.g., whotracks.me [17], easyList [18], or generated from Mozilla Lightbeam). Each list has benefits and disadvantages.

For our prototype, we used a semi-automated approach, accessing all Alexa Top 50 Websites International and Germany [19] with Mozilla Lightbeam running in the background and exporting the list of third-parties accessed. When testing the lists by browsing random websites, this semi-automatically generated list caught more third-party connections than the whotracks.me list. On the other hand, the whotracks.me list supplies a differentiation between different categories of third-parties (e.g., advertising, analytics, content delivery networks), which can provide a clearer picture of the intentions behind the third-party connection. We aim to systematically compare different tracker lists in the future.

Event separation: We configured TShark to listen to ports 80 (HTTP) and 443 (HTTPS) of the IP addresses generated from the tracker lists. We spawned separate TShark processes: a) monitoring establishment of a connection (SYN events) and b) monitoring data transferred to trackers (GET / TLS application data events). We further filter the SYN events by connections to the top 10 most prevalent trackers to further accentuate these acoustically (see Section 3.2).

All these events are stored in buffers and then sent out via OSC. As sound events which happen in close temporal proximity are not discernible anymore (precedence effect) [20], we send out the buffered events with a short pause in-between. In a heuristic pre-test, a pause of 70 ms turned out to provide the best balance between discerning single events and an overall coherent impression.

3.2. Sound design

The overall purpose of our approach is raising awareness, creating interest and stimulating thought on the topic of web tracking. The auditory representation is designed to show the amount of web tracking in the background, raise interest and convey some degree of danger in order to feature the associated privacy concerns. Not only the amount of tracking is important, but the fact that a group of very few companies are present on most websites. Therefore, we aim to disclose the oligopoly of these companies as well.

When a connection to one of the top 10 tracking companies is established, we present an audio recording of the company’s name in a whispered manner. Reverb is added to the whispers to intensify the spatial and suspicious impression, as a reference to the intrusion on privacy. Some of the companies are well known to users (e.g., Google, Facebook), others are less known (e.g., ComScore, Criteo). The whispered names are supposed to stimulate questions about these companies as well.

Each data transfer event is presented with a short sound event. The following sound variations were designed for comparison regarding users’ perception in terms of interest, curiosity, danger, and fear. We aimed to design our sounds in a way to reflect either power or fragility to convey both the power of tracking and the hidden, brittle quality it has as well. The powerful and fragile sounds were designed both in a musical and an abstract sound variation. Their numbers correspond to the sequence used in evaluation.

1. powerful and musical: low cello and tuba
2. fragile and abstract: granular synthesis
3. powerful and abstract V1: deep bleeps
4. fragile and musical: piccolo flute and violine
5. powerful and abstract V2: like V1, added delay

A video containing both an impression of the sonic experience with our system while surfing and examples of all sound variations can be found at <http://s.fhg.de/SonificationICAD2019>.

3.3. Comparison to existing approaches

Our approach of monitoring the internet traffic itself instead of relying on the Lightbeam browser plugin extends the capabilities of Soundbeam by:

- supporting all browsers and combinations of ad / tracking blocker plug-ins.
- supporting monitoring of any physical or virtual network connection on the host computer. This enables monitoring traffic generated not only by web browsing but by apps as well.
- supporting monitoring the traffic of any device (e.g., laptop, smartphone), if we open and monitor an ad-hoc wireless network that this device connects to.
- usage and comparison of different tracker blocking lists.
- conveying the name of the tracking company by whispers.

As we have no means of identifying which addresses or links the user wants to visit, our approach does not support differentiation between intentional website visits and third-party connections. Therefore, the quality of the tracker identification list is an essential factor for a reliable result.

For now, we do not support ensemble performance as we currently aim to make an individual user aware of the tracking he or she personally is subjected to. To create a multi-user experience, the capability for sending OSC events to different computers in the network can be added to our framework.

4. EVALUATION

4.1. Study design and hypothesis

We conducted an initial user study with 12 participants (6 male, 5 female, 1 no gender stated) with an age range between 23 and 36 years, mean age was 28.9 years. In a within-subjects design, we presented the recordings of five different sound variations in a classroom setting. Each recording represented the sonification of accessing the same website. It showed the actual sonic experience while surfing, consisting of several single bleeps occurring shortly after each other. Whispering of the tracker names was muted in order to set focus on the tonal quality of the sonified events. After each sound variation, participants filled out a questionnaire regarding the perceived emotional qualities of the respective sonic experience. We asked participants to rate their overall auditory impression of the sound playback (as if visiting a website), not the single sound elements. At the end, we presented all sound variations again and asked participants to state their favorite.

For the emotional qualities of the sounds, we asked participants to rank each sound between the following poles on a four-point likert scale. For statistical analysis, we assigned the numbers (-2, -1, 1, 2) to the scale items.

- innocent (-2) to dangerous (2)
- relaxing (-2) to frightening (2)
- boring (-2) to interesting (2)
- indifferent (-2) to curious (2)

As we designed the system to raise awareness, our main hypothesis is that the awareness regarding web tracking gets higher after exposure to the sonification. We assessed awareness before and after the sonification experience each with a five-point likert scale (low, rather low, medium, rather high, high).

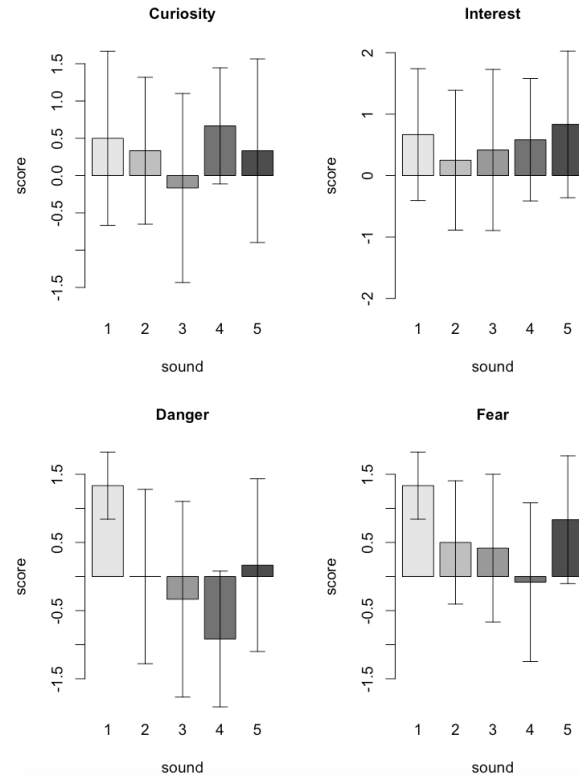


Figure 2: Emotional content of the sound variations. Error bars in plot: +/- one standard deviation

4.2. Results

As Shapiro-Wilk normality tests showed that normal distributions cannot be assumed in our sample, we performed a one-sided Wilcoxon signed rank test with continuity correction (see [21, p. 977]) to assess the differences between awareness scores prior to and after exposition to the sonification. The test results support our main hypothesis: Awareness levels were significantly higher after exposure to the sonification than before ($mean_{before} = 0.75$, $mean_{after} = 1.25$, $p = 0.024$, $r = -0.652$).

Results on the emotional qualities curiosity, interest, danger and fear were less distinct and not significant (see Table 1 and Figure 2). Hence, all statements on the emotional qualities of the sounds are descriptive only. For sound 1 (low cello and tuba), danger and fear ratings were both high in mean and with a smaller standard deviation compared to the other sounds. Interestingly, sound 4 (piccolo flute and violin) was perceived least dangerous, but raised the most curiosity. Sound 1 was stated most often as favorite (five times), followed by sounds 4 and 5 (three times each).

| Sound variation: | 1 | 2 | 3 | 4 | 5 |
|------------------|-------|-------|--------|--------|-------|
| mean(curiosity) | 0.500 | 0.333 | -0.167 | 0.667 | 0.333 |
| sd(curiosity) | 1.168 | 0.985 | 1.267 | 0.778 | 1.231 |
| mean(interest) | 0.667 | 0.250 | 0.417 | 0.583 | 0.833 |
| sd(interest) | 1.073 | 1.138 | 1.311 | 0.996 | 1.193 |
| mean(danger) | 1.333 | 0 | -0.333 | -0.917 | 0.167 |
| sd(danger) | 0.492 | 1.279 | 1.435 | 0.996 | 1.267 |
| mean(fear) | 1.333 | 0.500 | 0.417 | -0.083 | 0.833 |
| sd(fear) | 0.492 | 0.905 | 1.084 | 1.165 | 0.937 |

Table 1: Sound variations: Means and standard deviations of emotional content scores

5. DISCUSSION

The initial user study has limitations: Most notably, as the sounds were presented in a classroom setting, a sequence effect is expected. Future studies will benefit from individual presentation via headphones and randomisation of the sound variations. Adjectives of the emotional quality poles were not selected from standardized test batteries on emotional content. Additionally, the sample size of 12 participants was quite small. Nevertheless, some effect of the sonification experience on web tracking awareness could be shown.

6. FUTURE RESEARCH

As our initial results are encouraging, we will continue and extend our work in the following ways: First, we aim to set it up in a way that supports connecting a user's own device (laptop, smartphone) to a special wireless network we provide and monitor. By this, we allow users to explore the tracking sounds of their own browser or app configuration. We are also looking into porting the framework to a small computer like the Raspberry Pi [22]. This can ease the usage of our system in installations in public. Then, we plan to conduct a larger user study that assesses the impact of our approach to web tracking awareness in the field.

Future research questions regarding sound design are manifold: We aim to disclose not only the amount of web tracking, but the oligopoly of the tracking companies as well. So far, we approached this with the tracker name whispering when connecting initially. In future, we want to design signature sounds for each company, so the corresponding single events can be linked to these companies. Another significant step is moving on from producing the sounds in Ableton Live to a model-based sonification. Additionally, incorporating the spatial domain can help conveying tracker parameters by placement in the virtual room.

7. ACKNOWLEDGMENT

The authors want to thank Jan Maria Kopankiewicz for his support with implementation.

8. REFERENCES

- [1] S. Schelter and J. Kunegis, "On the Ubiquity of Web Tracking: Insights from a Billion-Page Web Crawl," pp. 53–66, 2016. [Online]. Available: <http://arxiv.org/abs/1607.07403>
- [2] S. Macbeth, "Tracking the Trackers: Analysing the global tracking landscape with GhostRank," Cliqz GmbH, Tech. Rep., 2017.
- [3] A. Karaj, S. Macbeth, R. Berson, and J. M. Pujol, "WhoTracks.Me: Monitoring the online tracking landscape at scale," pp. 1–15, 2018. [Online]. Available: <http://arxiv.org/abs/1804.08959>
- [4] A. Acquisti, L. Brandimarte, and G. Loewenstein, "Privacy and human behavior in the age of information," *Science*, vol. 347, no. 6221, pp. 509–514, 2015.
- [5] K. Purcell, J. Brenner, and L. Rainie, "Search engine use 2012," *Search*, 2012.
- [6] T. Bujlow, V. Carela-Espanol, B. R. Lee, and P. Barlet-Ros, "A Survey on Web Tracking: Mechanisms, Implications, and Defenses," *Proceedings of the IEEE*, vol. 105, no. 8, pp. 1476–1510, 2017.
- [7] W. Thode, J. Griesbaum, and T. Mandl, "I would have never allowed it": User Perception of Third-party Tracking and Implications for Display Advertising," in *Proc. International Symposium on Information Science*, 2015.
- [8] C. D. Wickens, "Multiple resources and mental workload," *Human factors*, vol. 50, no. 3, pp. 449–55, 2008.
- [9] T. Hermann, A. Hunt, and J. G. Neuhoff, *The Sonification Handbook*, 1st ed. Berlin: Logos Publishing House, 2011.
- [10] C. Hutchins, H. Ballweg, S. Knotts, J. Hummel, and A. Roberts, "Soundbeam: A Platform for Sonyfing Web Tracking," *Proceedings of the International Conference on New Interfaces for Musical Expression*, pp. 497–498, 2014.
- [11] M. Ballora, N. Giacobe, and D. Hall, "Songs of cyberspace: an update on sonifications of network traffic to support situational awareness," *Proc. SPIE Defense + Commercial Sensing*, vol. 8064, pp. 1–6, 2011.
- [12] M. Debashi and P. Vickers, "Sonification of network traffic flow for monitoring and situational awareness," *PLoS ONE*, vol. 13, no. 4, pp. 1–31, 2018.
- [13] L. Axon, S. Creese, M. Goldsmith, and J. R. C. Nurse, "Reflecting on the Use of Sonification for Network Monitoring," *Proc. SECURWARE 2016*, pp. 254–261, 2016.
- [14] A. Siami Namin, R. Hewett, K. S. Jones, and R. Pogrud, "Sonifying Internet Security Threats," *Proc. 2016 Conference on Human Factors in Computing Systems Extended Abstracts*, pp. 2306–2313, 2016.
- [15] <https://www.wireshark.org>, [Accessed 20.04.2019].
- [16] <https://www.ableton.com>, [Accessed 20.04.2019].
- [17] <https://github.com/cliqz-oss/whotracks.me>, [Accessed 20.04.2019].
- [18] <https://github.com/easylist>, [Accessed 20.04.2019].
- [19] <https://www.alexa.com/topsites>, [Accessed 20.04.2019].
- [20] H. Wallach, E. B. Newman, and M. R. Rosenzweig, "A Precedence Effect in Sound Localization," *The Journal of the Acoustical Society of America*, vol. 21, p. 468, 1949.
- [21] A. Field, J. Miles, and Z. Field, *Discovering Statistics Using R*. SAGE Publications, 2012.
- [22] <https://raspberrypi.org>, [Accessed 20.04.2019].

8

Concept: Inference Mapping Tool

8.1 Background

The collaborative research projects and tools that were presented in Papers 8 and 9 focus on web tracking. With regard to the topic covered in Part I, the question arose whether similar approaches could also be used to present the wealth of personal information that can be inferred from certain types of sensor data.

As the overview graphics provided in Paper 1 (Fig. 1), Paper 2 (Fig. 3), Paper 3 (Fig. 1), and Paper 6 (Fig. 2) show, it is possible to present such information in a structured and easy to digest manner. However, the static nature of these figures has several drawbacks. Due to the pace of technological development, such literature-based summaries can quickly become outdated and potentially generate a false sense of security by overlooking inferences that may become possible at a later stage. Furthermore, each of these overviews focuses on one particular type of sensor data, leaving aside the additional inferences that could be drawn by linking it with data from other sources. Also, due to the limited space and scope of the individual papers and the reasons provided in Chapter 2.2, these overviews leave out various details, including sensor properties (e.g., sampling rate), the device models and classification algorithms used, sample characteristics, and other specifics about the experimental setups.

Given all this, it could be desirable to present categories of inferable information in a more detailed, dynamic, and updatable form, using digital databases and visualization tools. Such an “Inference Mapping Tool” was not implemented as part of this Ph.D. project due to time constraints, leaving it as an interesting avenue for future development and research. In the following, based on the experience gained by examining the privacy impacts of sensors using experimental research and patents (cf. Part I), some recommendations will be provided to facilitate a future implementation. Ideas for basic functionalities (Ch. 8.2) and possible extensions (Ch. 8.2.1) will be presented, followed by a discussion of expected benefits and potential challenges (Ch. 8.3).

8.2 Proposed Functionalities

A tool for exploring the privacy-invading potential of sensor data could be realized in its simplest form through a specialized literature database with a search function. When users search for a certain type of sensor, they would be presented with a list of relevant experimental studies and patents related to the inference of personal information from that sensor's data. To ensure usability and help users in handling the sheer amount of information, the key findings of each publication should be displayed in a concise and standardized manner and search hits should be grouped, such as based on the categories of inferred information covered (e.g., personality traits, mental health, emotions). The search results could not only be presented in the form of lists and static overview graphics, such as the ones provided in Paper 1 (Fig. 1), Paper 2 (Fig. 3), and Paper 3 (Fig. 1), but also in the form of interactive diagrams, similar to what we proposed for web tracking data in Paper 8.

8.2.1 Possible Extensions

- **Host devices.** An advanced version of the tool could allow users to not only search for or select individual sensors, but also entire devices equipped with multiple sensors, such as smartphones, tablets, smartwatches, and fitness trackers – potentially even specific device models (e.g., Samsung Galaxy S16, iPhone 13). The tool would then present the sensors built into the selected device and the variety of personal information that can be inferred from the different types of sensor data. As exemplified in Papers 1, 2, and 3, sensors such as accelerometers, microphones, and eye-tracking cameras can be found embedded into a wide range of devices.
- **Sensor fusion.** If corresponding research results are available, the tool could also show categories of personal information that can be inferred by combining data from different sensors. These sensors could be situated within one device or across multiple devices. A highly advanced version of the tool may even allow the user to simulate an entire technological ecosystem (e.g., office room, home environment), highlighting the various types of inferences that could be drawn based on the sensor-equipped devices present. While such a tool would without question be very complex and costly to build, it could serve as an important protective instrument in high-risk areas (e.g., privacy protection for celebrities, high-ranking officers or heads of state). Not only can sensor fusion make it possible to infer additional information [290] but it can also make existing attacks more efficient. For example, it has been demonstrated that the accuracy of accelerometer-based inference attacks, such as the ones presented in Paper 2, can be improved by incorporating data from other sensors (e.g., gyroscope and magnetometer) [291].
- **Device positioning.** The tool could also provide users with the option to specifically search for inferences based on how the respective device is carried or where it is located (e.g., in hands of the user, on wrist, carried in pants pocket or bag, lying around in user's vicinity). For instance, it is possible to detect a user's eating [292] and smoking activity [293] using accelerometer data from a wrist-worn device, whereas, to

my knowledge, this type of inference is not possible by only using accelerometer data from a smartphone carried in a pocket.

- **Study details.** For all cited experimental studies, besides the key results (e.g., achieved accuracies), the tool could offer an optional summary of study details available, including experimental setting (real world vs. lab), sample size, used methods, computational cost of inferences, sensor properties (e.g., sampling rate), and the study’s limitations.
- **Non-sensor data.** While this dissertation has put a focus on sensor-based inference attacks (cf. Part I), the threat of undesired inferences is of course much broader, encompassing countless other data sources in modern life [48, 294, 295, 296]. Thus, the tool could also, for example, provide an overview of personal information categories that can be inferred from people’s browsing history [227], Facebook ‘likes’ [294, 297], playing strategies in video games [298] or online purchasing habits [295].
- **Data access.** Based on user permission systems and other authorization measures in place, the tool could show under which circumstances which data sources can be accessed by which parties (e.g., device manufacturers, platform providers, mobile apps, visited websites). Where operating systems do not provide enough transparency to clearly tell when resources are or can be accessed by certain parties (including by the system itself), it should – as a precaution – be assumed that these resources are regularly accessed.
- **Risk score.** Based on possible inferences identified from the literature, the tool could provide some form of privacy risk score for certain sensors, and potentially for entire devices and technical ecosystems. Such a score could indicate the estimated maturity of existing inference attacks (e.g., based on details of published studies, such as sample size, results, and limitations) and thus the likelihood that certain types of information can be inferred in a given situation. This approach could be extended into an app or browser extension informing users in real-time about their current risk of unknowingly disclosing personal information via inference. Such warnings, however, should always come with the important disclaimer that, due to non-disclosure requirements, some companies likely have far greater inference capabilities than what is known from published research, as I have discussed in Chapter 2.2. Beyond the mere risk of unwanted information disclosure, an advanced version of the tool could – based on context – incorporate possible types of data misuse and the severity of estimated consequences into the risk score calculation.
- **Inference chains.** Logically, pieces of inferred information can, in turn, be analyzed and linked to infer additional information. Thus, the tool could introduce the concept of “inference chains”, representing a concatenation of inferences and possible sub-inferences. For instance, people’s daily motion trajectories and location traces, which can be discernible from accelerometer data [46, 68], can in turn be used to infer people’s home and work locations, other points of interest, and identities [48, 299].
- **Collaborative design.** To distribute the work of continuously updating the tool, it could be designed as a collaborative platform so that, upon verifying their qualification and academic affiliations, contributors from around the world can enter new research

findings through a standardized online form. Of course, as with any type of open collaboration, there would need to be some form of quality control.

8.3 Expected Benefits and Challenges

Besides other benefits of presenting information related to inference attacks to the public in a structured and appealing manner (cf. Chapter 2.1), the continuously updated tool could also facilitate fellow scholars in searching for research gaps in the experimental literature. By clicking through devices and sensors and exploring the categories of personal information inferable from various types of data, they may discover areas where inference attacks appear plausible but have not yet been sufficiently tested in published research. Or, by looking at the details of studies in certain areas, they could also identify specific methods and approaches that have not yet been attempted (or sufficiently replicated).

Critics may of course argue that such a tool would serve data-hungry organizations as an inspiration for how to stealthily extract intimate personal information from collected data. To address this criticism, let me reiterate that – given the increasing financial value of personal data [256, 257] and large research budgets of many corporate actors – there arguably is much more advanced know-how on inferential analytics within the private sector already than what is visible in published research. Providing the research community and the wider public at least with some insight into these technical capabilities is important. Also, the possibility of inferences is of course not a bad thing per se, as I have briefly exemplified in Sections P1–3 and P2–1. The crucial question is how our society chooses to use and regulate these possibilities. And to arrive at a reasonable and informed response to this question, wide-ranging understanding and transparency on this issue are needed.

Here are some foreseeable challenges the developers of an Inference Mapping Tool would likely have to deal with:

- **Algorithmic complexity.** Given the infinite amount of possible ways to link and analyze data, it will never be possible to provide a final and exhaustive list of all possible inferences for a specific type of data (cf. Chapter 2.2). This inevitable limitation should always be prominently stated to avoid misinterpretation.
- **Information overload.** Evaluating the relevant literature can be an extremely time-consuming task, and the number of experimental publications in the area is set to further increase steeply in the coming years. For this reason, as proposed above, it should be considered whether the tool can be designed in a collaborative manner.
- **Methodological differences.** As I have observed during my research for the publications included in Chapter 2, there is a great deal of methodological heterogeneity and lack of standardization among existing studies on the sensor-based inference of personal information. One reason for this is that inference results for one sensor can usually be found across multiple scientific disciplines with different conventions (e.g., preferred methods, reporting standards, accuracy measures). Another reason is that many research areas under investigation are still in their infancy. Also, individual studies

often greatly differ in their experimental setups. If study results are to be meaningfully linked and compared, ways need to be found to overcome or deal with this heterogeneity.

- **Non-disclosure.** Corporate research efforts and data analytics operations are typically kept confidential. Thus, without access to companies' internal information, the platform would rely on public sources, which may understate the privacy threats resulting from companies' real analytical capabilities. This limitation needs to be clearly stated.
- **Threat dilution.** Compiling and chaining together possible inferences may eventually lead to the realization that intimate personal information can be inferred from almost all types of mobile and IoT sensor data, as Schneble, Elger, and Shaw [55] predict by stating that "all our data will be health data one day". While, at first glance, this might call into question the purpose and usefulness the tool itself, I contend that such a finding would be highly valuable. It would support the thesis that there is no inherent distinction between "sensitive" and "non-sensitive" data categories, which still is a common assumption in legal and technical data protection, as I will further discuss in Chapter 10.7.
- **Acquisition of funding.** As a basis for realizing the concept outlined above and to ensure real-world impact, appropriate funding would be required – not only for the development of the tool itself, but preferably also for the knowledge transfer into academia and the wider society as well as for the continued operation and maintenance of the tool (including content moderation and quality control, cf. point "Collaborative design" in Chapter 8.2.1). Unless funding from governmental and/or non-profit organizations can be secured, the challenge would be to make the funding of the project attractive to private-sector funders without compromising scientific neutrality and the fundamental transparency-enhancing mission of the project (cf. Chapter 10.10). While companies that process sensor data as part of their business model (e.g., app developers, mobile platform providers) may have a potential conflict of interest, funding could for example come from companies in the field of data protection and IT security. In any case, to ensure the greatest possible impact and stimulate public discourse, it would be desirable for the tool's functions to be available not only to the funders and their affiliates but to the general public as much as possible.

Part IV

A CRITICAL TAKE ON PRIVACY
REGULATION

9

Challenging the Notice-and-Choice Approach to Privacy

9.1 Background and Motivation

“Managing one’s privacy is a vast, complex, and never-ending project that does not scale; it becomes virtually impossible to do comprehensively.”

– Daniel J. Solove [114]

The *Weizenbaum Institute for the Networked Society*, where this thesis was conducted, encourages to see research results through the lens of civil society. The mission of the institute is to help provide a scientific basis for shaping the digital transformation in a responsible manner, including the identification of necessary framework conditions to uphold individual and social self-determination [300]. Central problems that repeatedly surfaced in the investigations in this dissertation so far are (1) the obscurity and staggering complexity of modern data processing, (2) the privacy-intrusion potential of seemingly harmless data, (3) companies’ questionable data practices, (4) a widespread lack of understanding regarding these data practices and complexities and, as a result, (5) people often being tracked and profiled by companies against their will and without their awareness.

As exemplified in Sections P1–3, P3–3, and P4–6.1, there are technical solution approaches to defend against identified privacy threats posed by inferential analytics. However, these approaches are as yet severely limited, often only offering protection against a few specific inferences while overlooking others (cf. Chapter 10.4). Another obstacle on the path to meaningful privacy protection is that, while companies usually make every effort to defend against potential external attackers and prevent data leaks to the outside world, they often lack incentives to deploy self-limiting data protection measures [301, 302]. In highly competitive industries and in times where personal data is known as the “new black gold” [257], why would profit-oriented companies voluntarily refrain from extracting knowledge from already collected data, especially when they can do so secretly, i.e., without reporting inferences drawn to data subjects or the public?

Given people’s widespread lack of knowledge and understanding about data privacy, would better consumer education be a solution? Of course, education on privacy matters is crucial, not least for raising awareness about the societal problems and injustices associated with personal data use. Fostering a sense of urgency around the dire state of privacy and the possible consequences of data misuse is particularly important in the face of widespread indifference, apathy, and a baseless but persistent nothing-to-hide attitude [7, 8, 303]. However, it would clearly be an unrealistic goal to educate the general public in minute detail about the countless forms of data collection and utilization prevalent in today’s technologized and data-centric society. Even for privacy researchers and experts, who spend hours per day reading, thinking about, and discussing privacy matters, it is difficult to see through the thicket of modern data processing – which is partly due to the time constraints of everyday life and limited memory capacity, but also due to the obscurity and overwhelming complexity of the corresponding technologies. The observation that the capabilities of inferential analytics are severely underestimated even by technical professionals (see, for example, Section P2–1 for accelerometer data and Section P4–5.2 for voice recordings) underlines the futility of expecting thorough understanding from ordinary users.

Inference attacks can be based on highly complex and unintuitive patterns and correlations, meaning that users cannot reliably anticipate what information will be inferable from the data collected about them. Picking up an example from Paper 4, even if a voice assistant is only used by a user to ask for the weather forecast, this can already lead to unexpected information leakage (e.g., based on the user’s sociolect, accent, intonation, pitch, loudness, or voice hoarseness). As a result of such obscurities, it has become impossible for individual consumers to meaningfully manage the information that organizations hold about them. Thus, logically, they cannot be expected to “defend themselves” against data misuse or give truly “informed consent” to data processing. As in the field of medical practice, where it is the responsibility of physicians to provide their patients with sufficient information to make informed decisions regarding proposed treatments [304], data protection law puts information obligations on the data controllers (cf. Art. 12–14 GDPR). However, it is highly doubtful whether the latter can even be fulfilled without exceeding the amount of text a single person can realistically process [30, 31].

This calls into question the prevalent legal paradigm of *notice and choice*, also referred to as *privacy self-management* [31], which mandates that people individually manage their privacy via consent. The very notion of *informational self-determination*, originally coined by the German constitutional court [87, 88] and formally defined as “the authority of the individual to decide [themselves] (...) when and within what limits information about [their] private life should be communicated to others” [305] seems to be based on wrong assumptions and may therefore – despite all good intentions – be an unachievable and misleading concept.

The problem applies not only to indications of privacy preferences (e.g., declarations of consent, privacy settings, user permissions) but also to the choice and use of privacy-enhancing tools and services (e.g., anti-tracking extensions; encryption tools; privacy-friendly web browsers, search engines, and instant messaging apps). While using such tools and services can be highly advisable and may offer significant protection against unwanted tracking and surveillance [306, 307], they, too, have diverse limitations (cf. Chapter 10.4), can introduce

new privacy threats (e.g., by bringing in additional third parties and requiring the disclosure of additional data, cf. Section P10–6.2), and also logically require some level of understanding for proper selection and operation, which adds to the already overwhelming problem of information overload (cf. Section P10–2.4). Therefore, while the development and provision of privacy-enhancing technologies to end users is certainly laudable, these technologies do not present a sufficient solution to the on-going privacy crisis and will not obviate the need for regulatory-level adjustments.

The main problem is not that people are too lazy or careless to adopt technical safeguards (although low adoption is also a recognized problem [308, 309]) but that these technologies will simply not solve some of the fundamental shortcomings of privacy self-management (cf. Section P10–6.2). These shortcomings are systemic and, accordingly, require systemic responses. As they cannot be comprehensively solved on an individual level, blame for their continued existence should not be put on individuals. Instead, with regard to human limitations and the complexities of modern data processing, privacy laws should focus on the anticipatory prevention of harms (cf. Chapter 10.9) and rely to a much lesser extent on data subjects’ supposedly “free” and “informed” decisions.

However, as long as privacy choices of individuals are falsely painted and perceived as expressions of their freedom and autonomy, there will be little resistance to the status quo. Widespread opposition is needed because, despite its deficiencies, the self-management paradigm has been deeply enshrined into privacy law over decades [31, 32] and alternative regulatory approaches are as yet relatively vague and undefined and will require significant effort to be further developed, implemented, and tested.

Although not yet sufficiently reflected in law, critical perspectives on privacy self-management are well-represented in the literature, with authors describing the approach as “dysfunctional” [139], “failed (...), impractical” [140], “destined to fail” [310], not “fit for purpose” [141], and as a “market failure” [142], “a fundamental dilemma” [31], and a deceptive “neoliberal technique of power” [120]. In 2013, the renowned data protection lawyer Eduardo Ustaran [311] concluded: “Yes, consent is dead. Further, continuing to give it a central role is dangerous”. Or, as professor Viktor Mayer-Schönberger put it: “The naked truth is that informational self-determination has turned into a formality devoid of meaning and import” [312].

All sorts of interesting and compelling arguments have been brought forth in opposition of privacy self-management. Among other topics, these arguments deal with the intricacies of data collection and processing [2, 310], people’s dependence on certain services [2, 120], and people’s inability to read, let alone understand the privacy policies of all their used services due to complex language and the enormous amount of time that would be required to do so [31, 313]. For instance, *New York Times* journalists examined privacy policies of major tech and media platforms and concluded that they are “verbose and full of legal jargon” – short: an “incomprehensible disaster” [314]. Taken together, arguments scattered across existing

literature add up to a strong case against privacy self-management, as it is commonly practiced throughout the Western world.⁷

However, while there are numerous publications that deal with one or several of these arguments in great detail, a thorough literature search did not reveal a publication that provides a comprehensive summary and overview of the existing arguments. Such an overview is urgently needed, especially given the lack of knowledge transfer into politics, as indicated by the law's heavy, continued, and – in the political realm – widely unquestioned reliance on privacy self-management. The persistence of the legal paradigm is in the interest of powerful corporate actors who benefit from the status quo, using the myth of “informed consent” as a justification for the excessive collection and questionable use of personal data. Organized efforts to keep this paradigm alive in its current dysfunctional form need to be confronted with a broad and well-structured barrage of arguments against it.

To respond to the aforementioned research gap, existing arguments from the literature are categorized and synthesized in Paper 10. I initiated and headed the project, which was conducted in collaboration with two colleagues from the Weizenbaum Institute: Otto Hans-Martin Lutz and Dr. Stefan Ullrich. They provided valuable assistance in structuring, analyzing, and interpreting relevant literature.

⁷By putting the focus of criticism on the Western world, I do not intend to imply that data protection efforts in other parts of the world are more meaningful or effective. Rather, other parts of the world are excluded from this analysis due to varying cultural and political factors that make a direct comparison difficult. Such an endeavor, while certainly interesting, is beyond the scope of this thesis.

The myth of individual control: Mapping the limitations of privacy self-management

Jacob Leon Kröger^{1,2}, Otto Hans-Martin Lutz^{2,3} and Stefan Ullrich^{1,2}

Abstract

Despite years of heavy criticism, privacy self-management (i.e., the principle that people individually manage their privacy via notice and choice) remains the standard of privacy protection throughout the Western world. Building on previous research, this paper provides an overview and classification of the manifold obstacles that render privacy self-management largely useless in practice. People's privacy choices are typically irrational, involuntary and/or circumventable due to human limitations, corporate tricks, legal loopholes and the complexities of modern data processing. While, at first glance, privacy self-management appears to embody fundamental values such as liberalism and individualism, in reality it rather deprives than strengthens people's autonomy. Moreover, the self-management approach ignores the consequences that individual privacy choices have on other people and society at large. Regarding future research, we argue that the focus should not be on whether privacy self-management can be fixed by making it more user-friendly or efficient – it cannot. The concept is based on fundamentally wrong assumptions. To meaningfully address the potentials and dangers of personal data processing in the 21st century, a shift away from relying so heavily on individual control is inevitable. We discuss potential ways forward, stressing the need for government intervention to regulate the social impact of personal data processing.

Keywords

Personal data, privacy self-management, informed consent, notice-and-choice, GDPR, fair information practices

1 Introduction

In 1973, the US government advisors Ware et al. famously proposed guidelines for privacy protection in automated data systems, the *Fair Information Practices* (FIPs). Two core provisions of the FIPs are that individuals should (1) have access to information on the data collected about them and (2) be able to decide for which purposes their personal data may be used. These principles – now commonly referred to as “notice and choice” or, in conjunction, “privacy self-management” – shaped not only data protection procedures in the US but were also widely adopted across the world, including in Canada, Australia and Europe (Gellman 2021).

Individual control over personal data is often portrayed as a reflection of fundamental rights and values, such as autonomy and human dignity (van Ooijen and Vrabec 2019), as is well illustrated by the establishment of the “Basic Right on Informational Self-Determination” in a landmark decision of Germany's constitutional court (1983). Under the paradigm of privacy self-management, unless explicitly prohibited by law, all forms of personal data processing can be authorized via consent (Solove

2012). Thus, considerable power and responsibility is vested in the individual data subject.

Most recently, the EU has continued its reliance on privacy self-management by making it one of the central principles of the new *General Data Protection Regulation* (cf. Art. 5 & 6(1)(a) GDPR). While the law also places new responsibilities on data controllers (e.g., data portability, state-of-the-art security measures, appointment of data protection officers), “the need for individual control seems to be addressed more explicitly and with greater prudence [in the GDPR] compared to earlier regulations”, as van Ooijen and Vrabec (2019, p. 92) observe.

¹Technische Universität Berlin, Germany

²Weizenbaum Institute for the Networked Society, Berlin, Germany

³Fraunhofer Institute for Open Communication Systems, Berlin, Germany

Corresponding author:

Jacob Leon Kröger, Weizenbaum Institute for the Networked Society, Hardenbergstraße 32, 10623 Berlin, Germany.

Email: kroeger@tu-berlin.de

*This paper was prepared using sagej.cls [Version: 2017/01/17 v1.20]
by courtesy of SAGE Publications (www.sagepub.com).*

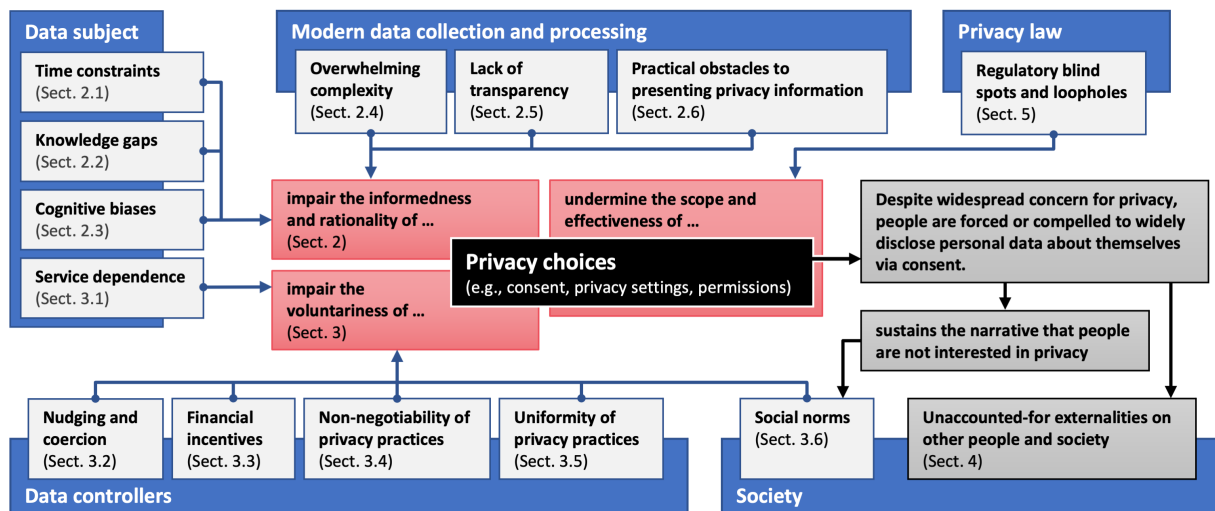


Figure 1. Overview of the obstacles to the meaningful exercise of privacy self-management covered in this paper.

Throughout the Western world,* privacy self-management remains the dominant approach to privacy regulation (Gellman 2021). This is despite the fact that the notice-and-choice model largely fails to avert dangerous and reprehensible data practices (Christl 2017b; Madge 2018; Zuboff 2019) and has been criticized as fundamentally dysfunctional (Hull 2015; Lehtiniemi and Kortessniemi 2017; Obar 2015; Rubinstein 2012; Scherf 2015; Solove 2012). To visualize how time-consuming it would be to read the privacy policies of large tech corporations, the artist Dima Yarovsky has printed them out and hung them on the wall of a gallery, stating that “Ticking the box, ‘I have read and agree to the Terms,’ is the biggest lie on the web today” (Schwab 2018). In his well-known article “Privacy Self-Management and the Consent Dilemma”, professor Daniel J. Solove (2012) explains in detail how people’s time constraints and cognitive biases obstruct the meaningful exercise of privacy self-management, and why the costs and benefits of personal data disclosure are better addressed at a societal level than through isolated personal choices.

Solove’s work on privacy self-management has strongly inspired academic discourse on the matter, including this paper. His critique of the notice-and-choice paradigm receives broad support, for example from the fields of law (Rothchild 2018; Rubinstein 2012; Scherf 2015), behavioral economics (Acquisti and Grossklags 2007; Acquisti et al. 2015; van Ooijen and Vrabec 2019) communication studies (Baruh and Popescu 2015) and philosophy (Hull 2015), with an increasing focus on the challenges posed by the rise of big data (Baruh and Popescu 2015; Hull 2015; Lehtiniemi and Kortessniemi 2017; Rubinstein 2012).

To underpin the debate going forward, this paper provides a structured overview of arguments that scholars have brought forth in opposition of privacy self-management. These arguments concern the informedness and rationality (Sect. 2), the voluntariness (Sect. 3) and the unaccounted-for externalities (Sect. 4) of individual privacy choices. Additionally, we point out loopholes in privacy law that undermine the effectiveness of privacy self-management (Sect. 5). While our legal analysis focuses primarily on the GDPR, which is presently regarded as the most comprehensive and influential privacy regulation worldwide (Miglicco 2018; Zarsky 2016), the essence of the arguments is generally applicable to privacy laws that embrace the notice-and-choice paradigm.

An introductory overview of the arguments covered in this paper is provided in Figure 1. As can be seen there, existing obstacles to privacy self-management are rooted in limitations on the data subject side, characteristics of modern data processing, data controller strategies, social norms and regulatory loopholes. We also address existing approaches to improve the usability and effectiveness of privacy self-management (Sect. 6.1) and show that some of the most critical problems are inherent to the paradigm and thus fundamentally unsolvable (Sect. 6.2). Then, we reflect upon the identified weaknesses of privacy self-management and suggest potential ways forward (Sect. 7), before drawing a conclusion in Sect. 8.

*The focus on the Western world is not meant to imply that data protection efforts in non-Western countries are more meaningful or effective. Rather, other parts of the world are excluded from this analysis due to varying cultural and political factors that make a direct comparison difficult. Such an endeavor, while certainly interesting, is beyond the scope of this paper.

2 Obstacles to Informed and Rational Privacy Choices

In this section, we summarize existing obstacles to the informedness and rationality of privacy choices, which can be ascribed to human limitations and the nature of modern data processing (see Figure 1). Specifically, this section covers people's time constraints (Sect. 2.1), lack of knowledge (Sect. 2.2) and cognitive biases (Sect. 2.3), the complexity (Sect. 2.4) and obscurity (Sect. 2.5) of today's data economy, and practical obstacles to presenting privacy information (Sect. 2.6).

2.1 Time Constraints

There is a vast discrepancy between the time required for the meaningful exercise of privacy self-management and people's time constraints (Obar 2015; Rothchild 2018). In order to make an informed choice, data subjects need to invest time in (i) gathering all relevant information, (ii) carefully examining the information, (iii) estimating costs and benefits of data disclosure based on the information and (iv) determining whether the expected consequences are compatible with their preferences (van Ooijen and Vrabec 2019). The section "Information Overload" (2.4) below will give an idea of how tedious this process – even step (i) alone – can be. Additionally, data subjects have to repeat the process if they want to compare the terms of competing service providers. As it is not unusual for companies to frequently modify their privacy policies, studying them all just once would still not suffice (Solove 2012).

Mcdonald and Cranor (2008) estimated that an average Internet user would need more than six full working weeks (244 hours) to read the privacy policies of every website visited in a one-year period, which would result in \$781 billion of lost productivity in the US alone. Notably, the study was conducted in 2008, since when the amount of Internet traffic has increased more than tenfold (Cisco Systems 2009, 2020). The study also exclusively focuses on web browsing and does not consider time required for policy re-reading and policy comparisons between alternative service providers. Therefore, while the outcome is astonishing, it can be regarded as a highly conservative estimate and most certainly understates the effort that would be required today to read the privacy policies of all services used by the average consumer (utilities, insurances, financial services, mobile apps, etc.).

Given the time constraints of everyday life, it is unrealistic to expect data subjects to read through thousands of pages of privacy policies. "There is no reason to think this is currently happening, or will ever happen", to say it in the words of Beales and Muris (2008, p. 114). Rothchild (2018, p. 559) even argues that, considering their limited

time budgets, it is the most rational choice for individual data subjects to leave privacy policies unread ("rational inattention"). This argument gains further weight when taking into account the many deficiencies of the notice-and-choice framework presented in this paper, which will most likely result in the time invested being fruitless anyway.

2.2 Lack of Knowledge

Making truly informed privacy choices in a modern, technology-based society would require a significant amount of economic, technical and legal background knowledge. Privacy self-management therefore builds on the concept of a knowledgeable and privacy-conscious consumer (Baruh and Popescu 2015; Solove 2012). However, a growing body of research suggests that average consumers do not possess nearly enough knowledge to make sound judgements and decisions about disclosing personal data (Acquisti and Grossklags 2007; Hull 2015; Liu and Gummadi 2011; Scherf 2015; Weinshel et al. 2019).

There are widespread knowledge gaps and misconceptions about key concepts of data protection law (Auxier et al. 2019; Solove 2012), the amount of data disclosed in everyday life (Hull 2015), the economic value of personal data (Acquisti and Grossklags 2007; Hull 2015; Scherf 2015), the meaning of online privacy settings (Felt et al. 2012; Liu and Gummadi 2011) and the possibilities of modern tracking and data processing technologies (Weinshel et al. 2019), including big data analytics (Baruh and Popescu 2015).

For example, only 37% of the Facebook users examined by Liu and Gummadi (2011) had selected privacy settings that match their expectations, with most participants revealing more information about themselves than assumed. Similarly, 74% of the Facebook users interviewed by Hitlin and Rainie (2019) did not know that the social media platform maintains a list of their interest and traits. In a survey by Felt et al. (2012), merely 3% of the participants correctly answered comprehension questions about mobile app privacy permissions. When confronted with detailed descriptions of web tracking and algorithmic inferences, people are often surprised by companies' technical capabilities and the extent of data collection (Weinshel et al. 2019). In a recent Pew Research Center survey, only 6% of participants reported that they completely understand what companies do with their personal data (Auxier et al. 2019). Also, roughly two-thirds of the participants said they have very little or no understanding of current data protection regulations.

Given this endemic lack of knowledge, it is unreasonable to expect ordinary consumers to understand what exactly they are consenting to when accepting privacy policies.

2.3 Nature of Human Decision-making

According to the logic of the notice-and-choice paradigm, people make privacy decisions based on a rational cost-benefit analysis (Hull 2015; van Ooijen and Vrabec 2019). Empirical findings from behavioral science and psychology show, however, that “people’s actual ability to make such informed and rational decisions does not even come close to the vision contemplated by privacy self-management” (Solove 2012, p. 1883). Even if complete information was available to data subjects (an unrealistic assumption, of course, considering the complexities described in Sect. 2.4), the problem remains that human perception and decision-making is heavily influenced by cognitive biases, often resulting in remarkably irrational and contradictory decisions (Acquisti and Grossklags 2007; Acquisti et al. 2015).

As van Ooijen and Vrabec (2019, p. 95) observe, “the more information individuals have access to about what happens to their data, the less information they are able to filter, process, and weigh to make decisions that are in line with their own privacy preferences.” In situations of information overload, the bounded rationality innate to humankind “makes us rely on simplified mental models, approximate strategies, and heuristics” (Acquisti and Grossklags 2007, p. 6). Humans also tend to ignore long-term consequences of their decisions in favor of short-term benefits and immediate gratification (Acquisti and Grossklags 2007). For example, when interested in a mobile app’s functionality, many smartphone users do not pay attention to the permissions granted during installation (Felt et al. 2012; Wottrich et al. 2018).

Some other cognitive biases that have been highlighted for their impact on privacy self-management are the availability heuristic (i.e., tendency to assess familiar dangers as riskier than unfamiliar ones), the valence effect (i.e., tendency to overestimate the likelihood of favorable events), overconfidence (i.e., tendency to overestimate one’s knowledge and capabilities) and the status quo bias (i.e., preference for things to stay the same) (Acquisti and Grossklags 2007; Solove 2012). The various pitfalls and systematic errors in human decision-making clearly obstruct rational privacy choices – even more so in the face of incomplete information and uncertainty.

Of course, strictly speaking, there is hardly any area of life in which people make purely rational decisions. It is very common to rely on simple heuristics and ignore large parts of the available information – and in many contexts, this can actually lead to more accurate judgements than logic and statistical models (Gigerenzer and Gaissmaier 2011). However, with regard to the complexities involved (see Sect. 2.4) and the serious and far-reaching ramifications (see Sect. 4), it is highly doubtful that people’s individual gut decisions and heuristics will

suffice to regulate personal data processing in their own and society’s best interest.

2.4 Information Overload

Over the last decades, technological innovations and new data-driven business models have bred an incredibly complex ecosystem of data collection and utilization. Globally, an estimated 59 zettabytes of data were created in 2020 and this figure is predicted to triple until 2025 – resulting in so much data that, if it was stored on DVDs, the stack of DVDs could circle Earth 222 times (Reinsel et al. 2018). Apart from the sheer amount of data, there are also “too many entities that collect, use, and disclose people’s data for the rational person to handle” (Solove 2012, p. 1881). Personal data is collected in various ways and forms – today, most of it is not manually entered or uploaded by data subjects but passively collected in the background (e.g., generated by an Internet user’s browsing activity) (Mehmood et al. 2016), making it even harder for data subjects to keep track of the data they disclose about themselves.

Given these overwhelming complexities, it is considered impracticable for people in today’s modern world to grasp, let alone to manage (i) what, when, by whom and how their personal data is collected (Mendes and Vilela 2017; Rothchild 2018), (ii) where and for how long the data is stored (Hartzog 2018), (iii) what inferences can be drawn from the data (Baruh and Popescu 2015; Kröger et al. 2021b), (iv) what exactly the data is used for (Rothchild 2018), (v) which parties receive the data (Hartzog 2018), and (vi) which positive and negative consequences may arise from all the above (Rothchild 2018; van Ooijen and Vrabec 2019). Modern data processing algorithms can be so sophisticated that even their developers do not fully understand how they work (Rainie and Anderson 2017). The consequences for data subjects can be gradual, cumulative or dispersed over time and thus virtually impossible to trace back to isolated data transactions (Le Métayer 2016; Solove 2012).

Although privacy policies already represent a coarse abstraction of organizations’ actual data practices, they are often lengthy and difficult to comprehend, even for well-educated people (Beales and Muris 2008; Fabian et al. 2017; van Ooijen and Vrabec 2019). Considering the mind-boggling depths and complexities of modern data processing, it can be assumed that sufficiently informed privacy decisions are an extremely rare occurrence.

2.5 Lack of Transparency

The invisible character of digital data poses another challenge for the informed exercise of privacy self-management. In modern devices, details about data collection and transfer are often hidden from the user

(Mehmood et al. 2016; Kröger and Raschke 2019) as are the logic and inner workings of data processing algorithms (van Ooijen and Vrabec 2019). Personal data is commonly duplicated and widely shared across the data economy, along what Madge (2018) calls “invisible data chains”. Some of the largest players involved, such as advertising networks and data brokers, act largely outside the public eye and almost never interact directly with data subjects (Rothchild 2018; Solove 2012).

Privacy policies do not provide much help in clearing up the obscurities, as they are often filled with opaque legal jargon (Beales and Muris 2008; Fabian et al. 2017) and can only provide a rough sketch of the underlying complexities (see Sect. 2.4). Certain details about data processing are even considered trade secrets and may be kept confidential, potentially overriding transparency obligations of data protection law (Fischer 2020; Madge 2018).

Privacy self-management typically requires people to decide about consenting to data processing before their data is initially collected (Solove 2012). However, relevant information might not be available at this point. While data subjects can choose to retract or alter their consent at a later stage (e.g., by adjusting privacy settings), the same problem applies to these situations – because here, too, potentially relevant events lie in the uncertain future. It can be unknown, for example, what exact outputs the data processing will produce (e.g., prediction, score, recommendation). Also, there is what Solove (2012) calls the “problem of aggregation”: Collected data can be linked and analyzed to infer previously undisclosed private information in unforeseen ways (Kröger 2018; Kröger et al. 2019a,b, 2020, 2021a). Regarding this issue, the transparency prescribed by law is often limited (see Sect. 5.2).

In an effort to enhance transparency and allow data subjects to manage their privacy in an “informed” manner, many modern data protection laws grant people the right to access the personal data that organizations store about them (e.g., Art. 15 GDPR). However, in practice, attempts to exercise such access rights are often ignored or responded to in a deceptive and insufficient manner (Kröger et al. 2020).

In sum, such obscurities can make it impossible for data subjects to exercise meaningful control over their personal data and to check whether promises made in privacy policies are kept, leading to many implications of data protection law being only theoretical.

2.6 Obstacles to Presenting Privacy Information

Modern technological environments often present an obstacle to the notice-and-choice paradigm (Rothchild

2018). For example, audio data, photos, videos and smart home sensor data can all incidentally contain information about people who are in the vicinity of the recording device but have not consented to the data collection. The affected individuals, who could be persons known to the device owner (e.g., colleagues, house guests, family members, flatmates) or simply strangers in public, are deprived of the control suggested by privacy self-management.

When data is collected by a device that has a very small display or no display at all (e.g., fitness tracker, smart speaker, home surveillance camera), it may even be difficult to deliver privacy information to device owners themselves. As Wachter (2018a, p. 4) observes: “[F]eatures of the IoT, which create numerous privacy risks, are frequently designed to go unnoticed by users in order to provide a ‘seamless’ experience. (...) The seamless implementation of these techniques can cause users to forget that their data is constantly being collected.”

3 Obstacles to Voluntary Privacy Decisions

In this section, we outline existing obstacles to the voluntariness of privacy choices. These can be divided into people’s dependence on certain services (Sect. 3.1), various data controller strategies (Sect. 3.2 to 3.5) and social norms (Sect. 3.6). A full overview is provided in Figure 1.

3.1 Dependence on Services

People have no other choice than submitting to a company’s privacy practices if they strongly depend on its services, e.g., for access to job opportunities, apartment offers, online communities, dating options or certain products. For example, in today’s interconnected world, a complete withdrawal from social media can mean a severe loss in social capital (Hull 2015) or even “digital suicide” (Baruh and Popescu 2015, p. 17) and is therefore clearly not a viable option for most people. Similarly, people depend on payment methods, such as credit cards, to effectively participate in the online marketplace (Baruh and Popescu 2015). Without access to these types of critical services, people essentially become “hermits in self-exile from the online world” (Rothchild 2018, p. 559). People with certain medical conditions may also depend on modern health services, e.g. mobile health apps, which often exhibit dubious privacy practices (Dehling et al. 2015). Besides practical needs, another reason for a user’s dependence on a service can be addiction (e.g., social media addiction, smartphone addiction, gaming addiction).

Service dependence is particularly problematic in monopolized industries or when a service is only offered by one company (e.g., a specific video game), so that the user cannot choose between alternative providers and

is thus subjected to one company's privacy conditions. Additionally, the problem can be exacerbated by social marginalization. For example, the trend of disconnecting from social or online media ("digital detox") can be a luxury only available to people in privileged positions in society (Ward 2017). Thus, the issue of service dependence can contribute to already existing systemic inequalities.

3.2 Nudging and Coercion

Under the notice-and-choice paradigm, data controllers typically have the power to "shape the playing field that guides individual decisions" (Schwartz 2003, p. 2082). For example, they can decide where, when and how often privacy notices are displayed and how exactly they are worded and designed (e.g., visual appearance of "accept" vs. "decline" button). Research has shown that the way privacy choices are framed and presented has a considerable impact on consumer decision-making (Acquisti et al. 2015; Acquisti and Grossklags 2007; Solove 2012). As Solove (2012, p. 1887) puts it: "[P]rivacy preferences are not developed in the abstract but in context."

The possibility of activating or suppressing people's privacy concerns through contextual cues is what Acquisti et al. (2015, p. 509) call "the malleability of privacy preferences". User experience tactics in mobile apps and websites specifically designed to subtly mislead and trick users into doing things they do not want to do are known as "dark patterns" (Gray et al. 2018). In the following, we present some examples.

A) Use of Default Settings. One approach to nudge privacy choices in a particular direction is to present the desired choice as the default option. Many people accept default settings without further review (van Ooijen and Vrabec 2019), not only to save time but also because default settings are often unconsciously perceived as recommendations (McKenzie et al. 2006).

B) Illusion of Control. People tend to take greater risks and disclose more about themselves when they feel in control of a situation, even if the sense of control is illusory (Solove 2012). Thus, to increase consent rates, data controllers can use confidence-inspiring communication (e.g., wording, design) that elicits – perhaps deceitfully – the subjective experience of control and self-determination.

C) Concealment & Obstruction. Once people have given their consent, data controllers can try to make it difficult for them to retract or modify their decision. For example, privacy options can be hidden in a complex settings menu, be labeled in an ambiguous manner, be overwhelmingly numerous and fine-grained or require some prerequisite steps to be selected. A specific example would be a website that does not offer a global opt-out of third-party

data sharing, but where users are instead required to visit external websites of ad publishers and follow their varying opt-out procedures (Hull 2015).

D) Annoyance. Another method of getting unwilling people to agree to data practices is to show the respective request to them repeatedly and/or in a distracting manner (e.g., cookie banners, prompts, pop-ups). When this happens, data subjects may have no other choice than accepting the conditions if they want to use the service undisturbed (Solove 2012).

3.3 Financial Incentives

People's economic circumstances can compel them to accept intrusive privacy practices when they are coupled with some form of financial incentive. There are various ways for companies to incentivize the disclosure of personal data, including loyalty and rewards programs, free vs. paid subscriptions (e.g., online news, music streaming, mobile apps) or discounts in exchange for a sign-up (e.g., newsletter, user account). Another incentive model that has gained traction in recent years is the dynamic adjustment of insurance premiums based on sensor-based measurements (e.g., car insurance rates based on driving behavior, health insurance based on fitness tracker data) (Spender et al. 2019).

For socially disadvantaged individuals, securing such financial rewards may not just be tempting but a financial necessity. Indigent single parents, low-income earners or unemployed people, for instance, might see no other option than consenting to the constant monitoring of their driving style and physical activities in order to receive insurance discounts, even if such tracking goes way beyond their level of comfort.

There have been attempts by legislators to address this issue. The California Consumer Privacy Act (CCPA), for example, prohibits financial incentives related to the collection and retention of personal information when "unjust, unreasonable, coercive, or usurious in nature" (§1798.125(b)(4) CCPA). However, it needs to be understood that – as highlighted above – the coercive effect of financial incentives may not only be driven by their specific scope and conditions but also simply by people's economic neediness. It may thus be impossible to offer financial incentives for personal data disclosure without coercive side effects.

3.4 Non-negotiability of Privacy Practices

The privacy options available to data subjects are usually dictated by the respective data controller, often leaving little room for customization (Lehtiniemi and Kortessniemi 2017; Schwartz 2003; Utz et al. 2019). Some companies even follow a "take-it-or-leave-it" approach by strictly

conditioning access to their services on the disclosure of sensitive personal data (Madge 2018; Solove 2012), which may not fulfill people's context-dependent privacy needs (Baruh and Popescu 2015). Not only giving but also retracting consent can be designed as an all-or-nothing decision for data subjects, forcing them to cease using a service altogether if they want to disclose less information about themselves (Lehtiniemi and Kortessniemi 2017).

While Art. 7(4) GDPR does not recognize consent as "freely given" if a service is conditioned on the acceptance of unnecessary data processing, it can be difficult for users to assess whether this is the case. Since the monetization of personal data has become an economic backbone of entire industries (Christl 2017b; Zuboff 2019), it can be expected that companies will be innovative in stretching the definition of "necessary" and will continue exploring ways to force through their desired terms – with the help of nudging if required (see Sect. 3.2).

Given the weak negotiating position of data subjects and the incessant demands for personal data disclosure – essentially as a form of "payment" for services (Scherf 2015) – privacy self-management often doesn't reflect the preferences people would express if given a meaningful choice (Lehtiniemi and Kortessniemi 2017; Utz et al. 2019).

3.5 Uniformity of Privacy Practices

In theory, having the choice between alternative service providers may help customers to find an offer that reasonably matches their privacy preferences. In reality, however, competing service providers often exhibit very similar data practices (Bischoff 2017; Dehling et al. 2015; Rothchild 2018). Comparisons between tech giants (e.g., device manufacturers, social networking sites, e-commerce platforms, streaming media services) have shown that they all engage in comprehensive user tracking and profiling (Bischoff 2017; Rothchild 2018). This problem also exists in highly sensitive industries, such as health information technology. A comparison of 17,979 health-related apps by Dehling et al. (2015) revealed, for example, that – except for a small minority of 4.37% – all examined apps pose a notable risk of privacy infringement. Based on an analysis of privacy policies of websites and mobile apps, Rothchild (2018, p. 71) concludes that "notice-and-choice cannot work in the context of Internet-enabled data flows because the uniformity of privacy practices leaves little or no room for the exercise of choice."

3.6 Social Norms

There is popular belief that privacy protection is "primarily an antiquated roadblock on the path to greater innovation" (Hull 2015, p. 97) and only relevant for people who have "something to hide" (Marwick and Hargittai 2019, p. 1706). Although these narratives are misleading, they

can shape public opinion and frame privacy-conscious people as technophobic, paranoid or even suspicious. Such social norms and judgements can strongly influence human decision-making, including privacy choices (Lehtiniemi and Kortessniemi 2017).

Large corporations whose profits rely on the extensive collection of personal data can use their influence on public discourse to shift privacy norms in their favor, which could potentially affect consumer privacy behavior as well as the beliefs and values of their own employees responsible for developing data-based services. During his time as Google's CEO, Eric Schmidt notoriously stated: "If you have something that you don't want anyone to know, maybe you shouldn't be doing it in the first place" (Esguerra 2009). Facebook's founder and CEO Mark Zuckerberg publicly argued that privacy is no longer a "social norm" (Johnson 2010).

In the age of virtual communities and dating platforms, social pressure toward data disclosure can also come from the mainstream consensus that people need to maintain an online identity (Hull 2015). The influence of norms on people's privacy choices may not only be fueled by stigmatization and pressure from peers, but also by the possible sanctioning of norm violations by data controllers. As Solove (2006, p. 493) puts it: "Because of its inhibitory effects, surveillance is a tool of social control, enhancing the power of social norms". For instance, data subjects could be classified into an unfavorable profiling category as a consequence of choosing restrictive privacy settings or refusing to share detailed information about themselves.

Finally, the manifold obstacles that render privacy self-management useless in practice (cf. Figure 1) produce a situation characterized by perceived loss of control (Auxier et al. 2019) and feelings of fatalism and apathy (Hargittai and Marwick 2016), where – due to the futility of trying to manage one's own privacy – carelessness effectively becomes the normal (or most "rational") choice (Rothchild 2018).

4 Externalities of Privacy Choices

Privacy self-management assumes that people decide about disclosing personal data based on their subjective costs and benefits (Lehtiniemi and Kortessniemi 2017). This individualistic focus fails to account for the manifold effects that privacy choices have on other people and society at large (Hull 2015; Solove 2012). The result is a collective action problem where people supposedly act in their own interest, but in doing so harm each other and, eventually, themselves. The following examples illustrate this issue.

A) Discrimination. Personal data is often used by companies (e.g., insurances, employers, banks) to assess people and treat them differently (Kröger et al. 2021;

Zuboff 2019). However, an often overlooked aspect is that data from one person can help to evaluate (and potentially discriminate) others, e.g., when the assessment is made based on group classifications, when the data feeds algorithms that are being used to analyze other people or when information provided by one person helps to judge others based on the exclusion principle. As Mühlhoff (2021, p. 19) puts it, big data-driven discrimination requires “the million data points from the majority of ‘normal people’ who feel they have ‘nothing to hide’ for its algorithms to learn what ‘normal’ (translate: ‘privileged’) means so that the system can discriminate against allegedly non-normal, dangerous, ill, ... persons.”

B) Data-driven Persuasion. A major danger of the big data era is seen in the possibility of manipulating people’s desires, opinions and behavior through highly personalized communication (Zuboff 2019). Insights into a person’s life and perception of reality (e.g., through learning about his or her character traits, preferences, fears and political attitudes) can be used to micro-target and psychologically tailor messages to better resonate with the individual receiver. Online ads, for example, are often served in real-time, not only based on previous profiling, but also based on the receiver’s current location and what he or she is doing or searching for online (Christl 2017b). Leaked documents revealed that Facebook can identify when internet users feel “worthless”, “insecure” and “need a confidence boost”, amounting to what has been described as “a perfect model of what buttons you can push in a particular person” (Lewis 2017). Such methods are not only used in commercial advertising but also by political campaigns (Christl 2017b). Thus, when serving the purpose of data-driven persuasion, personal data disclosure can threaten electoral freedom and the foundations of liberal democracy (Zuboff 2019).

C) Free Thought and Expression. The possibility of avoiding judgmental eyes and ears has been recognized as a precondition for the “uninhibited exploration of ideas” (Scherf 2015, p. 45). Richards (2008, p. 387) writes: “Intellectual privacy is vital to a robust culture of free expression, as it safeguards the integrity of our intellectual activities by shielding them from the unwanted gaze or interference of others.” Conversely, exposing oneself to continuous monitoring and tracking can drive self-censorship, anticipatory conformity and inhibition (Solove 2006). Such chilling effects do not only constrain individual freedom but also stifle pluralism, creativity and innovation on a collective level and thus impoverish entire societies (Hull 2015; Solove 2012). Solove (2012, p. 1893) asserts that privacy self-management, due to its individualistic focus, “does not prevent, redress, or even consider infringements on those social values.”

5 Regulatory Blind Spots and Loopholes

In this section, we describe how vagueness and exceptions in privacy law undermine the effectiveness of privacy self-management, using the GDPR as a reference. Specifically, we look at undefined legal terminology (Sect. 5.1) and the scope limitations of privacy law (Sect. 5.2). Our observations show that, even if privacy self-management was a reasonable approach to privacy regulation, some of its key provisions can be circumvented by data controllers under current EU legislation.

5.1 Vague and Undefined Legal Terminology

The legal provisions that enshrine the principle of privacy self-management in EU data protection law have been criticized for being vague and containing too many exemptions (Blanke 2020; Madge 2018; Wachter 2018b; Wachter and Mittelstadt 2019). For example, it remains unclear under which conditions data controllers have to fulfill a person’s data erasure request (van Ooijen and Vrabec 2019) and whether they are obligated to inform third-party data recipients about the erasure (Wachter and Mittelstadt 2019), to what extent personal data can be used for other purposes than stated in a company’s privacy policy (Zarsky 2016) and in which cases data controllers have to notify data subjects about an occurred data breach (Wachter 2018b). It is also not clear what data subjects can expect from exercising their so-called “right of explanation” (van Ooijen and Vrabec 2019, p. 96), whether such a right actually exists in the GDPR (Wachter et al. 2017) and whether data controllers need to inform people that their data is used for profiling purposes (Madge 2018).

The underlying problem is that various important terms and concepts in the GDPR are broadly described or left completely undefined. For instance, Art. 22 GDPR grants data subjects the right “not to be subject to a decision based solely on automated processing, including profiling”. The lack of further specification may represent “a loophole in which nominal human involvement in a computer-driven decision-making process renders the provisions inapplicable” (Wachter 2018b).

According to Art. 5(1)(b) and 6(4) GDPR, subsequent processing of personal data may be permitted without the data subject’s consent if the processing is “compatible” with the purpose for which the data was initially collected. Zarsky (2016) has discussed the uncertainties around this notion and how they may be taken advantage of by big data analytics firms. Art. 19. GDPR requires data controllers to communicate any rectification or erasure of personal data to the recipients to whom the personal data has been disclosed, unless it involves “disproportionate effort”. Here again, the GDPR does not specify how to assess the level of proportionality (Wachter and Mittelstadt 2019). Similarly, data controllers are only required to communicate a data

breach to the affected individuals if the breach is “likely to result in a high risk to [their] rights and freedoms” (Art. 34 GDPR), whereas the law does not specify the notion of “high risk” (Wachter 2018b).

According to another vague and highly controversial provision, data processing can be lawful without the data subject’s consent if it is necessary to pursue purposes of “legitimate interest” to the data controller or a third party (cf. Art. 6(1)(f) GDPR). Based on examples provided in GDPR recital 47, legitimate interests can range from critical security purposes, such as fraud prevention, to purely commercial purposes, such as direct marketing. While the GDPR stipulates that any such “legitimate interest” must be balanced against the interests and fundamental rights of the data subject, the weighing of interests is usually conducted by the data controller without any supervision or transparency obligation (Madge 2018). The European Parliament (2021) has expressed concern “that ‘legitimate interest’ is very often abusively mentioned as a legal ground for processing”. Given the GDPR’s strict consent requirements, shifting from consent to another legal basis for data processing is increasingly becoming an attractive option for companies. As Madge (2018) observes: “Organisations handling personal data, particularly those that are in the business of marketing, are in general revising their data protection procedures to use the claim of legitimate interests instead of consent.”

It should be noted that EU data protection law has been heavily influenced by corporate lobbyists (Madge 2018; Corporate Europe Observatory 2019; Zarsky 2016). While the GDPR’s general wording is to some extent important to ensure adaptability to a fast-changing technology landscape, the lack of clear definitions can amount to legal loopholes, leaving important issues open for interpretation by data controllers and courts. It has been argued that certain clauses of the GDPR “hollow it out to the extent that the exceptions themselves become a rule” (Brkan 2019, p. 23).

5.2 Scope Limitations of Privacy Law

The effectiveness of privacy self-management is also limited by the fact that, by means of transformation, personal data can be taken outside the scope of data protection law without removing its harm potential.

For example, none of the data subject rights granted by EU data protection law apply to anonymous information (Art. 2 and 4(1) GDPR). Thus, as soon as companies have anonymized collected data, which commonly happens in big data analytics (Mehmood et al. 2016), subsequent processing will not require consent. While there are good reasons to incentivize data anonymization (Finck and Pallas 2020), there are two problems with leaving anonymous data unregulated.

First, there are often ways to link seemingly anonymous data back to individuals (Blanke 2020; Le Métayer 2016; Kasperbauer 2020). Narayanan and Felten (2014, p. 6) contend that “[m]ost ‘anonymized’ datasets require no more skill than programming and basic statistics to deanonymize.” Advances in data analytics continue to reduce the reliability of existing anonymization methods (Mehmood et al. 2016). While, in principle, de-anonymized information is of course considered personal data subject to the GDPR (cf. recital 50 GDPR), data anonymization – or the appearance of it – can be purposefully used by data controllers to avoid the restrictions of data protection law (Wachter and Mittelstadt 2019).

Second, even securely anonymized data can be used in harmful ways. Some of the most prevalent uses of personal data, such as behavioral profiling, price discrimination and ad targeting, do not necessarily require real names attached to the data to have detrimental effects on individuals (Christl 2017b; Data Ethics Commission 2019; Wachter 2019). Also, data collected from one person (even if anonymized) can harm other individuals, population groups and society at large, as was illustrated in Sect. 4.

Information inferred about an individual based on collected personal data can also fall outside the scope of data protection law (Blanke 2020; Fischer 2020; Wachter and Mittelstadt 2019). Analyses of EU privacy law and jurisdiction have shown that “compared to other types of personal data, inferences are effectively ‘economy class’ personal data in the [GDPR]. Data subjects’ rights (...) are significantly curtailed for inferences” (Wachter and Mittelstadt 2019, p. 494f.). The insufficient regulation of inferred information has been recognized as a “significant loophole of the GDPR” (Skiljic 2021). Nothing in the wording of the law inherently prevents inferred information about individuals from being protected by the GDPR, but protection is also not guaranteed as there is substantial room for interpretation (Fischer 2020; Ufert 2020). While the California Consumer Privacy Act (CCPA), for example, explicitly includes “inferences drawn” in its scope of application, there is no such specific mention in the GDPR, which creates uncertainty and leaves the question of applicability to the courts (Blanke 2020; Ufert 2020).

6 Attempts to Fix Privacy Self-management

In response to the blatant failure of privacy self-management, there are often calls for more consumer education (Dehling et al. 2015; McMahon et al. 2020; Weinshel et al. 2019). While these calls are reasonable, they are unlikely to be realized on the scale required and may also create the false impression that a lack of education

is the only problem standing in the way of a functioning notice-and-choice regime.

There are also numerous approaches to help data subjects make their privacy choices in a more convenient and efficient manner, a brief introduction of which will be given in Sect. 6.1. However, these approaches will not be able to fix the fundamental dysfunctions of privacy self-management, as will be explained in Sect. 6.2.

6.1 Proposed Solutions

One approach to improve privacy self-management is the development of simpler and more intuitive ways of presenting privacy policies – for example, by shortening text and using plainer language (Fabian et al. 2017; Solove 2012), presenting privacy information through standardized icons (Information Commissioner’s Office 2017; van Ooijen and Vrabec 2019) or using other innovative forms of presentation, such as explanatory videos and cartoons (Information Commissioner’s Office 2017), interactive visualizations (Reeder et al. 2008) or question-answering chatbots (Polis 2021). Solutions have also been proposed to help people compare privacy practices of different service providers, including comparison tables (Reeder et al. 2008; Rothchild 2018), privacy seals and certification (Information Commissioner’s Office 2017) and even crowdsourcing-based approaches where people “analyse privacy policies and warn their peers about unacceptable terms” (Le Métayer 2016, p. 417).

Additionally, diverse software tools are developed to provide users more oversight and control over the data they disclose about themselves. These range from domain-specific tools, such as anti-tracking browser extensions (Le Métayer 2016; Scherf 2015) to comprehensive personal data services (PDSes) providing users with a central point of control over all sorts of personal information (e.g., browsing history, fitness data, credit card purchases, music streamed) (Data Ethics Commission 2019; Obar 2015; Rubinstein 2012). There are cloud-based PDSes and solutions that store the data locally on the user’s own device, both from non-profit organizations and commercial providers. Among other features, PDSes can allow users to protect their data through encryption (Information Commissioner’s Office 2017; Le Métayer 2016), manage their privacy preferences for different data controllers through a single user interface (Drozd and Kirrane 2020; Lehtiniemi and Kortessniemi 2017; Le Métayer 2016) and facilitate the exercise of data protection rights (e.g., data access, deletion, rectification) (Datenanfragen.de 2021; Le Métayer 2016). Furthermore, PDSes can offer decision support for the selection of privacy settings – based either on the user’s stated preferences or on preferences inferred from the user’s behavior (an approach known as “adaptive privacy settings”) (Baruh and Popescu 2015, p. 19).

With respect to people’s limited time, attention and technical understanding, legal provisions have also emerged to ensure that data-based services uphold and protect privacy rights of data subjects by default without requiring their manual input, e.g., by means of information security measures, pseudonymization and data minimization (cf. “Data protection by design and by default”, Art. 25 GDPR).

6.2 Why Privacy Self-management Cannot Be Fixed

While the proposed solutions can help mitigate some of the deficiencies of privacy self-management, they in turn face their own diverse set of challenges. Apart from practical hurdles, such as the challenge for PDSes to manage privacy preferences across a large number of potentially uncooperative data controllers (Raschke et al. 2018) and a widespread lack of adoption (Scherf 2015), solution approaches typically focus on fixing a particular problem of privacy self-management, whereas other problems are ignored or even worsened.

For example, while more fine-grained consent models might be desirable from the perspective of customizability, they add to the already overwhelming complexity of privacy choices, creating “greater risks of confusion” (Solove 2012, p. 1885). On the other hand, standardized privacy icons save people time but are “very generalized and simplistic (...) [and] do not provide comprehensive knowledge about data collection practices” (van Ooijen and Vrabec 2019, p. 98). While automated tools could be used to compare privacy policies of different service providers, they will not change the fact that the terms and data practices of many providers do not differ substantially (see Sect. 3.5), that privacy choices are shaped by cognitive biases (see Sect. 2.3) and that people often depend on a specific service and thus cannot simply switch to alternatives (see Sect. 3.1).

Another problem is that PDSes are provided by third parties who pursue their own (e.g., financial) interests and “may take advantage of the power vested in the intermediary positions” (Lehtiniemi and Kortessniemi 2017, p. 1). This issue is particularly critical if a PDS processes personal data. Adaptive privacy systems, for example, can require sensitive personal information for the prediction of the user’s privacy preferences, such as locational, temporal, behavioral and/or social media data (Baruh and Popescu 2015). Furthermore, the informed use of privacy tools obviously requires data subjects to be familiar with their specific functions and limitations, which may further exacerbate the issue of information overload (see Sect. 2.4).

Importantly, approaches that enhance the usability and effectiveness of privacy self-management cannot overcome the collective action dilemma outlined in Sect. 4. Whether privacy icons, fine-grained consent models or personal data services – all are situated in a process that revolves around

the choices and self-interest of individual data subjects. The harm these choices may inflict on other people and society at large are usually ignored.

Finally, while data protection measures, such as data minimization and pseudonymization (cf. "Data protection by design and by default", Art. 25 GDPR), can certainly reduce the risk of unnecessary data exposure, the legal paradigm of privacy self-management allows to circumvent such measures by way of "informed" and "freely given" consent of the data subject (cf. Art. 4(11) GDPR) – although, in reality, indications of consent rarely fulfil these criteria, as we have argued in this paper.

7 Ways Forward

It has been convincingly argued that, in order to avert the dangers of excessive paternalism, the element of individual control should not be eliminated from privacy protection entirely (Solove 2012). Future research needs to thoroughly examine in which areas the individualistic approach is really tenable and practicable, considering the obstacles summarized in this paper.

The prerequisite for meaningful progress on the legal front is that the manifold limitations of privacy self-management are recognized and treated as such by legislators. This also means admitting that our privacy laws – including recent and much-anticipated ones, such as the GDPR and the California Consumer Privacy Act – are based on wrong assumptions and therefore not truly fit for purpose. While sticking with privacy self-management may be the path of least effort in the short run, this policy ignores the long-term consequences of uninformed and involuntary privacy choices, which can be severe not only for individuals but also for society at large (see Sect. 4).

Instead of leaving privacy protection purely up to individual choice, many commentators have proposed that the use of personal data should be controlled and authorized in a more institutionalized form, based on social impact assessments (e.g., Beales and Muris (2008); Data Ethics Commission (2019); European Commission (2021); Mühlhoff (2021); Rothchild (2018); Scherf (2015)).

Beales and Muris (2008, p. 118), for example, call for regulation approaches based on "the consequences of information use and misuse". Kasperbauer (2020, p. 770) contends that, "[a]s we lose control of our data, we will need better tools to defend against and penalise those who use our data against us." The German Data Ethics Commission (2019) proposes that algorithmic systems should face additional requirements (e.g., transparency obligations, audits, real-time supervision) or even a complete ban based on their potential for harm. Similarly, publicly appointed experts could determine that the use of certain data categories for certain purposes (e.g., personalized pricing, credit/insurance scoring, targeted advertising) should be

prohibited because the expected societal benefits do not outweigh the costs and risks involved. Solove (2012, p. 1903) proposes "hard boundaries that block particularly troublesome practices". Under Barack Obama, even the Executive Office of the US President suggested examining "whether a greater focus on how data is used and reused would be a more productive basis for managing privacy rights in a big data environment" (Podesta et al. 2014, p. 61).

Impact assessments are already an established tool in the field of data protection (e.g., "data protection impact assessment", Art. 35 GDPR) but, thus far, they are mainly conducted by the data controllers themselves, which can obviously entail significant conflicts of interest. Two recent EU initiatives towards a stronger risk-based government regulation of data use are the proposed *AI Act*, under which artificial intelligence applications could face special requirements or even a legal prohibition based on their estimated harm potential (European Commission 2021), and the proposed *Digital Services Act*, which could include restrictions on profiling-based advertising (Stolton 2021). These initiatives are interesting, but their real impact remains to be seen and depends, of course, on the final legislative outcome and the rigor of enforcement.

One thing is clear however: The range of regulatory possibilities is far from exhausted. Not only the use, but even the collection or inference of certain types of personal information could be illegalized if deemed ethically indefensible (Rothchild 2018). The underlying risk-benefit calculations should consider social, economic, psychological and physical consequences for individuals (Mühlhoff 2021) and entire population groups targeted by algorithmic profiling (e.g., based on gender, income, ethnicity, religious affiliation, political views) (Taylor 2016). Instead of only focusing on harms that have already occurred, the Data Ethics Commission (2019, p. 16) argues that "[p]ossible future cumulative effects, network effects and effects of scale, technological developments and changing actor constellations must be taken into account when gauging the potential impact" of personal data processing.

Further research is required to determine how exactly such risk-based approaches could be enshrined in law and implemented in practice. This research must pay due consideration to potential challenges, such as resistance from corporate and governmental stakeholders (Madge 2018; Scherf 2015; Zarsky 2016), the difficulties of estimating and weighing the consequences of personal data processing (Hull 2015; Scherf 2015) and the hurdle of breaking with the current legal framework, in which, for example, prohibitions of data processing are rather uncommon (Wachter 2018a). Furthermore, it will be a key challenge to strike a sensible balance between

individual autonomy, meaningful consumer protection and the prevention of overly paternalistic regulation (Solove 2012).

8 Discussion and Conclusion

The arguments presented in this paper show that, while privacy self-management may function in very simple and idealized cases, it clearly does not represent a sufficient solution for the informed authorization of data processing across today's technology-based society.

Various problems, ranging from people's limited knowledge, time constraints and cognitive biases over obscurities in data processing to various forms of external influence on privacy decisions, call into question the notion of "informed" and "freely given" consent, as prescribed, for example, in Art. 4(11) GDPR. Additionally, there is a fundamental mismatch between the individualistic focus of the notice-and-choice paradigm and the societal impacts of individual privacy decisions. Lastly, the vagueness and scope of relevant legal provisions leave numerous loopholes that further dilute the applicability and effectiveness of privacy self-management, allowing data controllers to undermine data subject rights. Even if privacy self-management *was* a reasonable approach to privacy regulation (which is clearly not the case), the latter fact alone should suffice to recognize the notice-and-choice provisions of the GDPR, and many similar laws across the world, as fundamentally unfit for purpose.

By structuring existing arguments against privacy self-management, our paper supports and builds upon previous work in this direction (e.g., Baruh and Popescu (2015); Hull (2015); Rothchild (2018); Scherf (2015); Solove (2012)), underscoring the conclusion that "being tasked with doing work beyond its capabilities (...) this paradigm alone cannot serve as the centerpiece of a viable privacy regulatory regime" (Solove 2012, p. 1880). While, at first glance, privacy self-management appears to embody Western values of liberalism and individualism, in practice it rather deprives than strengthens people's autonomy. The claim that privacy should be a matter of individual control has even been described as a neoliberal technique of power (Baruh and Popescu 2015; Hull 2015). Hull (2015, p. 89-98) writes very aptly:

[P]rivacy self-management isn't about protecting people's privacy; it's about inculcating the idea that privacy is an individual, commodified good that can be traded for other market goods (...) [It] forces privacy into the market, obstructs the functioning of other, more social, understandings of privacy, and occludes the various ways that individuals attempt to resist adopting the market-based

view of themselves and their privacy (...) [P]rivacy self-management functions as a technology of neoliberal governance, by inculcating the belief that subjectivity and ethical behavior are matters primarily of individual risk management coupled with individual responsibility for poorly-managed risks. (...) [It] obscures a social struggle, repackaging it as a well-functioning market.

The present situation, where the vast majority of consumers perceive a loss of control over the data that companies collect about them (Auxier et al. 2019) and where widely opposed data practices fuel entire industries (Christl 2017a; Zuboff 2019), is a consequence and testament of the dismal failure of privacy self-management. The fact that many people surrender the attempt to manage their privacy under these conditions (Auxier et al. 2019; Felt et al. 2012) should by no means be taken as evidence that people are fundamentally uninterested in the fair and well-regulated processing of personal data.

Existing ideas for improving privacy self-management offer symptomatic treatment but fail to holistically address the root problems of the paradigm (see Sect. 6). Although seemingly reasonable in the context of today's policy landscape, such "solutions" may foster unwarranted hope and false trust in a broken system while distracting from actual, transformative solutions that go beyond the limitations of privacy self-management. Based on existing literature, we outlined potential ways out of this stagnant situation, stressing the need for regulatory approaches that focus on the consequences of personal data use. Most scholarly suggestions in this area are still vague and hypothetical. To arrive at actionable policy recommendations, further research on this issue is needed.

Finally, on a more holistic note, it is critical to understand that many of the dangers associated with personal data processing are merely a reflection and continuation of long-standing socioeconomic disparities. Modern forms of data-based discrimination, such as credit/insurance scoring, often disadvantage people that are already marginalized by society (Christl 2017a; Zuboff 2019). Were wealth and life opportunities more equally distributed, not only would the risk of marginalization be lower but it is also likely that – through a higher average level of education – people would be more resilient to other abuses of their data, such as personalized persuasion and manipulation attempts.

While addressing the immediate ramifications of personal data processing is vital, the strategic focus should be on tackling the underlying reasons that cause these problems to arise in the first place. Thus, the threats posed by modern data processing are yet another reminder that our society urgently needs to take on the pressing issue of social injustice.

References

- Acquisti A, Brandimarte L and Loewenstein G (2015) Privacy and human behavior in the age of information. *Science* 347: 509–514.
- Acquisti A and Grossklags J (2007) What Can Behavioral Economics Teach Us About Privacy? *Digital Privacy: Theory, Technologies and Practices* 18: 363–377.
- Auxier B et al. (2019) Americans and Privacy. Technical report, Pew Research Center. URL <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/>.
- Baruh L and Popescu M (2015) Big data analytics and the limits of privacy self-management. *New Media & Society* 19: 579–596.
- Beales JH and Muris TJ (2008) Choice or Consequences: Protecting Privacy in Commercial Information. *U. Chi. L. Rev.* 75: 109–135.
- Bischoff P (2017) Comparing the privacy policy of internet giants side-by-side. URL <https://www.comparitech.com/blog/vpn-privacy/we-compared-the-privacy-policies-of-internet-giants-side-by-side/>. (accessed 12 March 2021).
- Blanke JM (2020) Protection for ‘inferences drawn’. *Global Privacy Law Review* 1(2): 81–92.
- Brkan M (2019) Do algorithms rule the world? *Int. J. Law Inf. Technol.* 27(2): 91–121.
- Christl W (2017a) Corporate surveillance in everyday life. Technical report, Cracked Labs, Vienna.
- Christl W (2017b) How companies use data against people. Technical report, Cracked Labs, Vienna.
- Cisco Systems (2009) Visual networking index: 2008–2013. Technical report. URL https://www.cisco.com/c/dam/global/pt_br/assets/docs/whitepaper/VNI_06-09.pdf.
- Cisco Systems (2020) Annual internet report (2018–2023). Technical report. URL <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>.
- Corporate Europe Observatory (2019) Captured states: when EU governments are a channel for corporate interests. Technical report, Brussels. URL https://corporateeurope.org/sites/default/files/ceo-captured-states-final_0.pdf.
- Court GC (1983) Census Act, BVerfGE 65, 1. URL <https://freiheitsfoo.de/files/2013/10/Census-Act.pdf>. (accessed 20 July 2021).
- Data Ethics Commission (2019) Opinion of the Data Ethics Commission. Technical report. URL https://datenethikkommission.de/wp-content/uploads/DEK_Gutachten_engl_bf_200121.pdf.
- Datenanfragende (2021) Request generator. URL <https://www.datarequests.org>. (accessed 27 February 2021).
- Dehling T, Gao F, Schneider S and Sunyaev A (2015) Exploring the Far Side of Mobile Health. *JMIR mHealth and uHealth* 3(1): e3672.
- Drozd O and Kirrane S (2020) Privacy CURE: Consent Comprehension Made Easy. In: Hölbl M, Rannenberg K and Welzer T (eds.) *ICT Systems Security and Privacy Protection*. Springer, pp. 124–139.
- Esguerra R (2009) Google CEO eric schmidt dismisses the importance of privacy. URL <https://www.eff.org/deeplinks/2009/12/google-ceo-eric-schmidt-dismisses-privacy>. (accessed 20 March 2021).
- European Commission (2021) Proposal for a Regulation of the European Parliament of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts. URL <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52021PC0206&from=EN>.
- European Parliament (2021) Resolution of 25 March 2021 on the Commission evaluation report on the implementation of the General Data Protection Regulation – 2021/c 494/11. *Official Journal of the European Union* URL <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52021IP0111>.
- Fabian B, Ermakova T and Lentz T (2017) Large-scale readability analysis of privacy policies. In: *International Conference on Web Intelligence*. pp. 18–25.
- Felt AP et al. (2012) Android permissions: User attention, comprehension, and behavior. In: *Symposium on Usable Privacy and Security*.
- Finck M and Pallas F (2020) They who must not be identified – distinguishing personal from non-personal data under the GDPR. *International Data Privacy Law* 10: 11–35.
- Fischer C (2020) *The legal protection against inferences drawn by AI under the GDPR*. LL.M. thesis (Law and Technology), Law School, Tilburg University. URL <http://arno.uvt.nl/show.cgi?fid=151926>.
- Gellman R (2021) Fair Information Practices: A Basic History - Version 2.20. *SSRN Electronic Journal*.
- Gigerenzer G and Gaissmaier W (2011) Heuristic decision making. *Annual Review of Psychology* 62: 451–482.
- Gray CM et al. (2018) The dark (patterns) side of UX design. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. pp. 1–14.
- Hargittai E and Marwick A (2016) “What can I really do?” Explaining the privacy paradox with online apathy. *International journal of communication* 10: 3737–3757.

- Hartzog W (2018) The case against idealising control. *European Data Protection Law Review* 4(4): 423–432.
- Hitlin P and Rainie L (2019) Facebook Algorithms and Personal Data. Technical report, Pew Research Center, Washington, D.C. URL <https://www.pewresearch.org/internet/2019/01/16/facebook-algorithms-and-personal-data/>.
- Hull G (2015) Successful failure: What Foucault can teach us about privacy self-management in a world of Facebook and big data. *Ethics and Information Technology* 17(2): 89–101.
- Information Commissioner's Office (2017) Big data, artificial intelligence, machine learning and data protection. Technical report. URL <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>.
- Johnson B (2010) Privacy no longer a social norm, says Facebook founder. URL <https://www.theguardian.com/technology/2010/jan/11/facebook-privacy>. (accessed 10 March 2021).
- Kasperbauer TJ (2020) Protecting health privacy even when privacy is lost. *Journal of Medical Ethics* 46(11): 768–772.
- Kröger J (2018) Unexpected inferences from sensor data: a hidden privacy threat in the Internet of Things. In: *IFIP International Internet of Things Conference*. Springer, pp. 147–159.
- Kröger JL, Lindemann J and Herrmann D (2020) How do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on iOS and Android Apps. In: *International Conference on Availability, Reliability and Security*.
- Kröger JL, Miceli M and Müller F (2021) How data can be used against people: A classification of personal data misuses. *SSRN Electronic Journal* URL <https://ssrn.com/abstract=3887097>.
- Kröger JL and Raschke P (2019) Is my phone listening in? On the feasibility and detectability of mobile eavesdropping. In: *IFIP Annual Conference on Data and Applications Security and Privacy*. Springer, pp. 102–120.
- Kröger JL, Lutz OHM and Müller F (2019a) What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking. In: Friedewald M et al. (eds.) *Privacy and Identity Management. Data for Better Living: AI and Privacy*. Cham: Springer, pp. 226–241.
- Kröger JL, Lutz OHM and Raschke P (2020) Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference. In: Friedewald M et al. (eds.) *Privacy and Identity Management. Data for Better Living: AI and Privacy*. Cham: Springer, pp. 242–258.
- Kröger JL, Raschke P and Bhuiyan TR (2019b) Privacy Implications of Accelerometer Data: A Review of Possible Inferences. In: *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy (ICCSPP)*. ACM, pp. 81–87.
- Kröger JL, Raschke P, Campbell JP and Ullrich S (2021a) Surveilling the Gamers: Privacy Impacts of the Video Game Industry. *SSRN Electronic Journal* URL <https://ssrn.com/abstract=3881279>.
- Kröger JL et al. (2021b) Personal Information Inference from Voice Recordings: User Awareness and Privacy Concerns. *Proceedings on Privacy Enhancing Technologies (forthcoming)* 2022(1).
- Le Métayer D (2016) Whom to trust? Using technology to enforce privacy. In: Wright D and De Hert P (eds.) *Enforcing Privacy*. Springer, pp. 395–437.
- Lehtiniemi T and Kortetniemi Y (2017) Can the obstacles to privacy self-management be overcome? Exploring the consent intermediary approach. *Big Data & Society* 4(2): 1–11.
- Lewis P (2017) Our minds can be hijacked. URL <https://www.theguardian.com/technology/2017/oct/05/smartphone-addiction-silicon-valley-dystopia>. (accessed 16 March 2021).
- Liu Y and Gummadi KP (2011) Analyzing Facebook Privacy Settings: User Expectations vs. Reality. In: *ACM Conference on Internet Measurement*. pp. 61–70.
- Madge R (2018) Five loopholes in the GDPR. URL <https://medium.com/mydata/five-loopholes-in-the-gdpr-367443c4248b>. (accessed 12 March 2021).
- Marwick A and Hargittai E (2019) Nothing to hide, nothing to lose? *Information, Communication & Society* 22(12): 1697–1713.
- Mcdonald AM and Cranor LF (2008) The Cost of Reading Privacy Policies. *Journal of Law and Policy for the Information Society* : 543–568.
- McKenzie CR, Liersch MJ and Finkelstein SR (2006) Recommendations implicit in policy defaults. *Psychological Science* 17(5): 414–420.
- McMahon A, Buyx A and Prainsack B (2020) Big Data Governance Needs More Collective Responsibility. *Medical Law Review* 46(1): 155–182.
- Mehmood A, Natgunanathan I, Xiang Y, Hua G and Guo S (2016) Protection of big data privacy. *IEEE Access* 4: 1821–1834.
- Mendes R and Vilela JP (2017) Privacy-preserving data mining. *IEEE Access* 5: 10562–10582.
- Miglicco G (2018) GDPR is here and it is time to get serious. *Computer Fraud & Security* 2018(9).
- Mühlhoff R (2021) Predictive Privacy: Towards an Applied Ethics of Data Analytics. *Ethics. Inf. Technol.* 23: 675–690.
- Narayanan A and Felten EW (2014) No silver bullet: De-identification still doesn't work. Technical report, Princeton University.
- Obar JA (2015) Big Data and The Phantom Public. *Big Data & Society* 2(2): 1–16.

- Podesta J et al. (2014) Big data: Seizing opportunities, preserving values. Technical report. URL https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf.
- Polisis (2021) AI-powered Privacy Policies. URL <https://pribot.org/>. (accessed 9 March 2021).
- Rainie L and Anderson J (2017) Code-dependent: Pros and cons of the algorithm age. Technical report, Pew Research Center. URL <https://www.pewresearch.org/internet/2017/02/08/code-dependent-pros-and-cons-of-the-algorithm-age/>.
- Raschke P et al. (2018) Designing a GDPR-Compliant and Usable Privacy Dashboard. In: Hansen M et al. (eds.) *Privacy and Identity Management*. Springer, pp. 221–236.
- Reeder RW et al. (2008) A User Study of the Expandable Grid Applied to P3P Privacy Policy Visualization. In: *ACM Workshop on Privacy in the Electronic Society*. pp. 45–54.
- Reinsel D, Gantz J and Rydning J (2018) The Digitization of the World from Edge to Core. Technical report, International Data Corporation. URL <https://www.seagate.com/files/www-content/our-story/trends/files/dataage-idc-report-final.pdf>.
- Richards NM (2008) Intellectual Privacy. *Texas Law Review* 87(2).
- Rothchild JA (2018) Against Notice and Choice: The Manifest Failure of the Proceduralist Paradigm to Protect Privacy Online (or Anywhere Else). *Cleveland State Law Review* 66: 559–648.
- Rubinstein I (2012) Big Data: The End of Privacy or a New Beginning? *International Data Privacy Law* : 12–56.
- Scherf R (2015) Yes I agree*: Assessing the failure of privacy “Self-Management” and its regulatory reforms. *Public Policy & Governance* 6(2): 37–54.
- Schwab K (2018) Printing out the privacy policies of Facebook, Snap, and others. URL <https://www.fastcompany.com/90171107/printing-out-the-privacy-policies-of-facebook-snap-and-others>. (accessed 19 July 2021).
- Schwartz PM (2003) Property, privacy, and personal data. *Harv. L. Rev.* 117: 2055–2128.
- Skiljic A (2021) The status quo of health data inferences. URL <https://iapp.org/news/a/the-status-quo-of-health-data-inferences/>. (accessed 5 April 2021).
- Solove DJ (2006) A Taxonomy of Privacy. *University of Pennsylvania Law Review* 154(3): 477–564.
- Solove DJ (2012) Privacy Self-Management and the Consent Dilemma. *Harvard Law Review* 126: 25.
- Spender A et al. (2019) Wearables and the Internet of Things: Considerations for the Life and Health Insurance Industry. *British Actuarial Journal* 24: e22.
- Stolton S (2021) Prohibit targeted advertising in Digital Services Act, EU data watchdog says. URL <https://www.euractiv.com/section/digital/news/prohibit-targeted-advertising-in-digital-services-act-eu-data-watchdog-says/>. (accessed 20 March 2021).
- Taylor L (2016) Safety in numbers? Group privacy and big data analytics in the developing world. *Group Privacy: New Challenges of Data Technologies* 126: 13–36.
- Ufert F (2020) AI Regulation Through the Lens of Fundamental Rights. *European Papers* 5(2): 1087–1097. URL <https://www.europeanpapers.eu/it/europeanforum/ai-regulation-through-the-lens-of-fundamental-rights>.
- Utz C et al. (2019) (Un)informed Consent: Studying GDPR Consent Notices in the Field. In: *ACM Conference on Computer and Communications Security*. pp. 973–990.
- van Ooijen I and Vrabec HU (2019) Does the GDPR Enhance Consumers’ Control over Personal Data? *Journal of Consumer Policy* 42(1): 91–107.
- Wachter S (2018a) The GDPR and the Internet of Things: A three-step transparency model. *Law, Innovation and Technology* 10(2): 266–294.
- Wachter S (2018b) Normative challenges of identification in the Internet of Things. *Comput. Law Secur. Rev.* 34(3): 436–449.
- Wachter S (2019) Data protection in the age of Big Data. *Nature Electronics* 2(1): 6–7.
- Wachter S and Mittelstadt B (2019) A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI. *Colum. Bus. L. Rev.* (2): 494–620.
- Wachter S, Mittelstadt B and Floridi L (2017) Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law* 7(2): 76–99.
- Ward M (2017) Your digital detox isn’t as radical as you think. URL <https://www.washingtonpost.com/news/made-by-history/wp/2017/07/30/your-digital-detox-isnt-as-radical-as-you-think/>. (accessed 20 July 2021).
- Ware W et al. (1973) Records, computers, and the rights of citizens. Technical report. URL <https://aspe.hhs.gov/report/records-computers-and-rights-citizens>.
- Weinshel B et al. (2019) Oh, the Places You’ve Been! User Reactions to Longitudinal Transparency About Third-Party Web Tracking and Inferencing. In: *ACM Conference on Computer and Communications Security*. pp. 149–166.
- Wottrich VM et al. (2018) The privacy trade-off for mobile app downloads. *Decision support systems* 106: 44–52.
- Zarsky TZ (2016) Incompatible: The GDPR in the Age of Big Data. *Seton Hall Law Review* 47(4): 995–1020.
- Zuboff S (2019) *The age of surveillance capitalism*. Profile books.

DISCUSSION AND CONCLUSION

10

Discussion

This chapter will first summarize the contributions of this thesis by breaking down the topics covered and main findings reported in the individual parts and studies (Chapter 10.1). Then, it will provide a brief overview of the public and media response our research has received so far (Chapter 10.2), followed by a general discussion, including numerous drawn conclusions and recommendations (Chapters 10.3 to 10.9). Finally, Chapter 10.10 will offer some reflections on the independence of internet research.

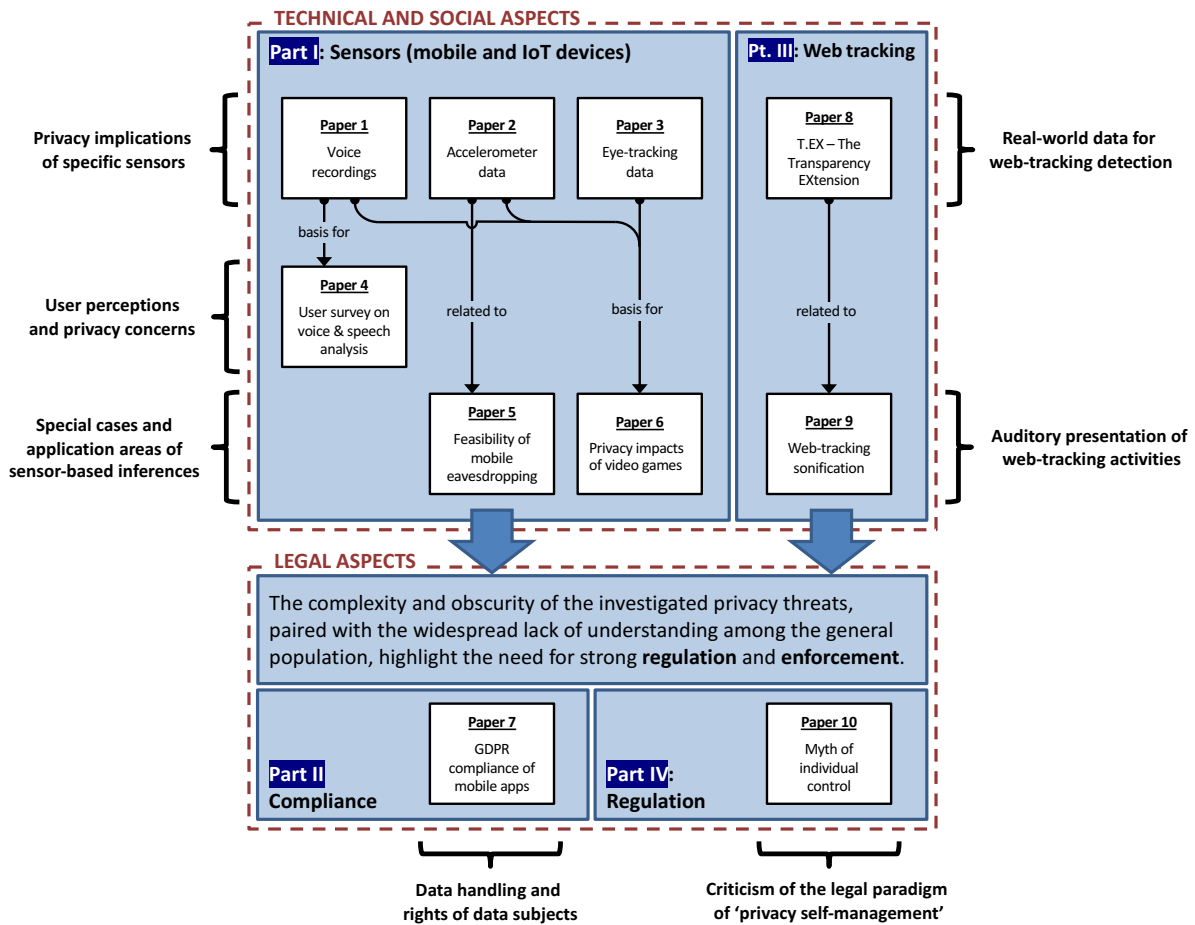
10.1 Thesis Summary

In the following, the four parts of this thesis will be summarized. For better orientation, Fig. 10.1 provides a visual overview of the research papers included in this thesis.

Part I

In Part I, we have looked at the rich variety of personal information that can be inferred from patterns and correlations in IoT and mobile sensor data. Specifically, based on relevant patents, available commercial products, and experimental literature, three papers have explored and illustrated the privacy-invading potential of inferences drawn from voice recordings (Paper 1), accelerometer readings (Paper 2), and eye-tracking data (Paper 3), respectively. In Chapter 2, it was explained why these types of papers are important and why the focus was placed on these specific sensors. The three above papers show that the sensors under investigation can allow serious invasions of user privacy, even where their recordings do not appear sensitive at first glance. While based on prior work, the papers offer a novel contribution, as existing knowledge about the informational richness of individual types of sensor data has rarely been structured and consolidated, especially not from a privacy perspective. Our findings highlight that, for meaningful and holistic privacy impact assessments of sensors, one should consider the inferences that can plausibly be drawn from their data, rather than considering only a sensor's official purpose. Further, it was explained why companies' non-disclosure policies make it impossible to exactly determine the applications, capabilities, and limitations of state-of-the-art

Figure 10.1: Overview of the research papers included in this dissertation.



inference methods used in practice. Considering the extensive financial resources and technical know-how of certain companies, the real threat level may go far beyond what is visible in published research.

To address the data subject’s perspective, Paper 4 presented results from a survey about user awareness and privacy concerns regarding personal information inference from voice recordings, building upon findings from Paper 1. A few previous studies have already found a widespread lack of awareness about sensor-based inference attacks [47, 193, 195]. Our study confirms these findings and provides a novel contribution by focusing specifically on voice recordings. In accordance with findings by Crager et al. [195], the measured levels of awareness varied depending on participants’ demographic attributes – but only slightly. Even participants with professional experience in the ICT field scored low on awareness. While some participants expressed worry about unauthorized data leakage to third parties, specific types of data misuse, and about being misrepresented by audio-based inferences, they overall expressed only moderate privacy concern. However, open text responses indicated that unconcerned reactions are largely explained by unwarranted trust in companies’ data practices and knowledge gaps about possible data misuses. The reported level of concern varied strongly between different categories of inferred information (e.g., higher concern about health information than about inferred age). Showing participants a short educational video on the privacy impacts of voice and speech analysis had a significant negative impact on their intention to use voice-controlled virtual assistants, which aligns with prior studies

showing that people’s privacy concerns tend to grow when presented with examples of personal data inference [47, 104, 227].

Then, Paper 5 explored the issue (or potential illusion) of smartphone eavesdropping, highlighting an endemic lack of transparency in the collection and processing of mobile sensor data and investigating both microphones and accelerometers as possible eavesdropping channels. For the first time in academic literature, our paper provides a comprehensive overview of the issue, including a description of the threat model and a summary and analysis of existing arguments and explanatory approaches. Non-system mobile apps, third-party libraries, and ecosystem providers were considered as possible adversaries. We provided insights on the technical feasibility of such attacks and explained why existing mitigation and detection techniques applied by ecosystem providers and the international research community are not sufficient to reliably rule them out. At the same time, while disputing the widely held opinion that the spying fears are completely unrealistic, we stressed that hard evidence for such attacks is thus far lacking and that there are many alternative explanations for eerily accurate ads (e.g., conventional tracking methods, profiling information from various sources being combined, inference algorithms, pure coincidence, people’s cognitive biases in memory and attention).

Paper 6 investigated the privacy impacts of the video game industry, including a focus on the sensors found in gaming equipment. Apart from offering a classification of data types commonly collected by video games and a detailed discussion of personal data categories that can be inferred from gaming data, the paper discussed social and regulatory implications of the presented findings.

Part II

In Part II, an undercover field study was presented that examines – in a first-of-its-kind longitudinal approach – how transparently mobile app vendors handle user data (Paper 7). Specifically, the study investigated in three iterations between 2015 and 2019 whether a set of 225 companies complied with transparency obligations prescribed by the GDPR. In theory, the law grants consumers the right to access the personal data that companies hold about them. However, our undercover research revealed severe obstacles to exercising this right in practice. 19 to 26% of the vendors did not reply to our requests or were completely unreachable. And the responses that we obtained exhibit a long list of shortcomings, including deceptive statements and other severe violations of information security and data protection principles. Among other problems, 76 to 85% of data exports were provided without a verification of the requester’s identity. And where verification mechanisms were in place, they were sometimes circumvented by the app vendors themselves. We were also provided with personal data via non-encrypted channels and received access to personal data not belonging to our user accounts. While positive trends were observed in certain areas, most of the observed problems persisted over time (e.g., low overall response rate, incomplete responses, deceptive and misleading statements, unintelligible data exports). All in all, responses were satisfactory in only 15 to 53% of the cases, with a surprising decline of response quality between the GDPR enforcement date and the end of our study. Furthermore, 27% of our user accounts vanished during the study, mostly without proper notification. The rate of satisfactory responses we obtained in 2019 is comparable with results Herrmann and Lindemann obtained in 2014 [162],

which further indicates that progress in the privacy compliance of mobile apps is stagnating. Based on our findings, we provided specific recommendations for regulatory action.

Part III

In Part III, two papers were presented that contribute to the detection and/or exposure of web-tracking activity largely hidden to ordinary internet users. Paper 8 introduced “T.EX – The Transparency EXtension”, a novel browser plug-in which records browsing sessions in a secure and privacy-preserving manner. The real-world browsing data collected by the plug-in can be used to feed algorithms to enable the automated detection of web trackers. Automated detection is needed to overcome the shortcomings of the conventional blacklist-based approach to blocking web trackers (e.g., manual effort required, trackers can change domain name, tracking and relevant services can come from the same domain, manually-created blacklists often contain errors). We also highlight the downsides of using artificial data in the training of algorithms for automated web-tracking detection (e.g., bots can easily be fooled; artificial data does not include human interactions like moving the mouse, scrolling, or pressing a key; bots usually do not log into services and, thus, miss some of the most important third-party communication). In the paper, we define specific requirements and objectives for the design of our browser plug-in, then describe the implementation, and finally discuss limitations and evaluate the performance of the tool in terms of website loading times and memory and CPU usage.

Then, Paper 9 explored ways in which web-tracking activity can be “sonified”, i.e., made audible through indicative melodies and sounds. Compared to existing approaches, our proposed framework for web tracking sonification can monitor any network connection, independent of browser, application, and device. The paper describes how a prototype was developed, including our approach to tracker identification, data transfer event separation, and sound design. An initial user study, which is reported in the paper, showed different emotional reactions based on the type of sound played, and that exposing users to sonification significantly increased their web-tracking awareness.

With regard to the privacy implications of inferential analytics covered in Part I, the question arose whether similar presentation approaches as in Paper 8 could also be used to illustrate the wealth of sensitive information that can be inferred from certain types of sensor data. While a corresponding software tool was not developed due to time constraints, suggestions were provided for potential future implementations in Chapter 8, comprising basic functionalities, possible extensions as well as challenges to be expected for such a project.

Part IV

After the previous parts of this dissertation dealt with the astonishing privacy impacts of sensor data, companies’ obscure and unethical data practices, users’ privacy perceptions and limited knowledge regarding these issues, and attempts to improve transparency, Part IV zoomed out to the big picture by addressing the overall state of informational privacy in our society and the deficiencies of current data protection law.

Specifically, Paper 10 deals with the legal principle of privacy self-management, which has shaped privacy law in Western countries for decades and also plays a central role in EU’s

GDPR. Based on previous literature, the paper offers an overview and classification of the various limitations of privacy self-management that undermine the principle’s usefulness in managing the dangers and potentials of modern data processing. These limitations include obstacles to informed and rational privacy choices (Section P10–2), obstacles to voluntary privacy choices (Section P10–3), unaccounted-for externalities of individual privacy choices on other people and society at large (Section P10–4), and several regulatory blind spots and loopholes (Section P10–5). The result is a situation where most data subjects feel confused and powerless, perceiving a loss of control over their data [1], while large industries live on highly questionable data practices and benefit from people’s disorientation and confusion [6, 120, 314].

In theory, the legal paradigm of privacy self-management grants people the right to only authorize the collection and processing of their personal data when they feel comfortable with it. In reality, however, the level of loss of control is so immense that it leads to widespread feelings of apathy and complete resignation [7, 8, 303]. Considering the issues structured and summarized in Fig. 1 of Paper 10, individual control over personal data is a misleading notion or, as we put it in our paper: a “myth”. In reality, such control does not exist, at least not to a meaningful or sufficient extent. As argued in Section P10–6, even if widely alleged, the problems presented will not be solved by individual data subjects making better “educated” and “informed” privacy choices. We have outlined perspectives on what can and should be done about the failure of privacy self-management in Section P10–7. Note, however, that most existing ideas in this area are still hypothetical. To arrive at more specific, actionable policy recommendations, further research on this issue is needed (cf. Chapter 11).

10.2 Public and Media Response

The papers presented in this thesis have received encouraging feedback and attention in the press and across social media platforms. They were mentioned in hundreds of tweets around the globe [315, 316, 317], including by researchers and best-selling authors, such as Cory Doctorow [318], Steve Stewart-Williams [319, 320], and Jordan Peterson [321]; by famous journalists and activists, such as Wolfie Christl [269] and M. Serdar Kuzuloğlu [322]; and by influential technologists, such as Maderas [323] and Joe Biden’s former cybersecurity expert Jacqueline Singh [324].

At the time of writing, the papers have jointly received over 200,000 downloads from SpringerLink [222, 325, 326, 327], Georgia Tech Library [328], SSRN [329, 330], Sciendo [194], and ACM Digital Library [331, 332]; numerous citations by domain experts [333]; and various mentions across Facebook [315, 316], blogs [35, 334, 335, 336, 337, 338, 339, 340], podcasts [341, 342, 343], YouTube channels [344, 345], and in news papers and magazines [346, 347, 348, 349, 350, 351, 352, 353, 354, 355, 356, 357, 358, 359, 360, 361]. For instance, our paper on the privacy implications of eye tracking (Paper 3) was not only widely discussed on social media [316] but also covered by the award-winning podcast *Curiosity Daily* [343], the *VICE* magazine [346], the *TIME* magazine [360], *POLITICO* [361] and the *Boston Globe* [362] and has already been referenced by the *Electronic Frontier Foundation* (EFF) [363], the *IEEE Global Initiative on Ethics of Extended Reality* [364], the *Bipartisan Policy Center* (BPC) [365] and the *Information*

Technology & Innovation Foundation (ITIF) in Washington, D.C. [366], which has been ranked as the world’s most authoritative science and technology policy think tank [367]. As the online magazine *NEO.LIFE* summarizes, the paper “has been reverberating throughout the industry, consumer advocate organizations, and academia” [368]. Our investigation into the privacy compliance of mobile app vendors (Paper 7) received the Best Paper Award at the *ARES 2020 Conference*.

I gave interviews about my research to leading national news outlets, such as *Deutschlandfunk* [369, 370] and *Bayerischer Rundfunk* [270] in Germany, *El País* in Spain [371] and *la Repubblica* in Italy [372]. Furthermore, my co-authors and I have summarized some of our main research findings in Twitter threads, for which we also received positive feedback (e.g., [373, 374, 375, 376]).

I am delighted about the broad response in new as well as conventional media and throughout the general public, which confirms the social relevance of our research. It is satisfying to see that the way we have presented our findings not only appeals to fellow scholars, but also seems to make them digestible and understandable for laypeople despite the complexity of the topics under investigation.

10.3 Privacy is not Dead

Given the state of informational privacy in today’s legal and socio-technical environment, where people are being tracked and profiled wherever they go and whatever they do, feelings of resignation and powerlessness are more than understandable. As most people have lost control over their personal data [1] and powerful organizations continue to exploit our data for ethically dubious, reprehensible, and often widely opposed purposes, such as behavioral targeting [27, 377], discriminatory credit scoring [6, 29] and the unsolicited inference of intimate attributes [28, 378], privacy has been declared a “lost cause” [9, 10] and “completely and utterly dead” [11].

There is no question that there are many problems with the ways personal data is being handled today and, indeed, a sudden improvement of the situation should not be expected. In the face of this reality, resignation may appear as the only logical reaction or consequence. Depending on the perspective taken, some of the issues covered in this dissertation, such as the wealth of sensitive inferences that can be drawn from seemingly innocuous sensor data (cf. Part I) or the extent of hidden web tracking (cf. Part III), may even contribute to a fatalistic sentiment by illustrating the opacity and overwhelming complexities of modern data processing. In addition, there are narratives portraying privacy as “primarily an antiquated roadblock on the path to greater innovation” [120], as we have discussed in Section P10–3.6.

Under these circumstances, more than ever, it is crucial for people to understand (1) why privacy protection matters, both for individuals and society at large, and (2) that privacy is by no means “dead”. The general assumption that privacy has ceased to exist or to be relevant, and that any effort of protecting it is therefore senseless, is a misguided narrative that helps those organizations that want to collect and use personal data without being held accountable. What we need is a more differentiated view on this issue. People from the *I-have-nothing-to-hide* camp who doubt the relevance of privacy protection altogether [8, 303] deserve to be

educated about the basics of personal data use and misuse. Some thought-provoking reflection questions should suffice to convince most people that they would not, in fact, like to disclose anything about themselves to the entire world [379, 380]. This becomes particularly apparent when considering that nearly everyone keeps secrets (e.g., about drug/alcohol abuse, sexually transmitted diseases, infidelity, mental health, religious beliefs, cheating at work or in school, illegal activities) [381]. According to David H. Flaherty, even people convinced that privacy does not matter “cannot withstand even a few minutes’ questioning about intimate aspects of their lives without capitulating to the intrusiveness of certain subject matters” [382].

And the functions and value of informational privacy go far beyond the mere avoidance of embarrassment and the guarding of our deepest personal secrets. Data protection is not ultimately about the protection of secrets and other personal data, but about protecting individual data subjects and the wider society from the varied and serious harms that can result from personal data collection and use. To provide illustration, in collaboration with Milagros Miceli from the Weizenbaum Institute and Florian Müller from the University of Kassel I have published a holistic overview and classification of the ways in which personal data can be used against people [383].

Those who do understand that informational privacy may serve important protective functions but simply believe that it is unattainable under today’s circumstances and irretrievably “dead” should also carefully check their assumptions. Yes, it is true and obvious that some economical, social, and technological developments of the past decades will be impossible to reverse, and that there is not much of privacy left to save in certain areas. For example, unless choked off by natural resource constraints or some other external crisis, it can reasonably be expected that data collection and processing technologies will continue to become more and more efficient and pervasive, thereby exacerbating many privacy threats, such as illegitimate data de-anonymization and the inference problem addressed in Part I. It is also true that computer systems are fundamentally insecure, and that hacker attacks and other unauthorized intrusions and data leaks should be expected to remain a regular occurrence. In fact, as Wachter [384] underscores, promising otherwise would be deceptive and may generate a false sense of security. And yes, it can also be suspected that the increasing inter-connectedness and technologization of society will bring about more and more products and services that require the collection of granular personal information. It is important to recognize such trends and the manifold dangers that come with them.

At the same time, however, even in the face of widespread surveillance by powerful criminal, corporate, and governmental organizations, informational privacy is *not* an all-or-nothing issue. While various privacy-eroding trends are irreversible and only set to accelerate, there are still many aspects and facets of privacy that can and should be protected and/or reclaimed. Identifying these strategic fields of action, where there still is a realistic prospect of improvement, is a critical task for privacy researchers, policymakers, and civil society watchdogs. In areas where dangerous trends and technological advancements can no longer be prevented or reverted, the question needs to be addressed how we want to deal with these realities on a regulatory level.

For instance: How should we, as a society, deal with the fact that companies gain more and more insights into our personal lives, often much more than we would want or expect?

Of course, the goal must not necessarily be a reduction of the amount of data collected and stored (which would also seem rather unrealistic given current technological developments) but improvement measures can also revolve around transparency and accountability in the handling of personal data or, for example, a better regulation of its use. The focus should be on adverse effects on people and society and how these can be prevented, or remediated.

Extrapolating from the current dysfunctional frameworks of data governance, a pessimistic (or even fatalistic) outlook seems justified, indeed. While the GDPR's general principles stipulate that personal data should be processed "fairly and in a transparent manner" (Art. 5(1)(a) GDPR) and also mandates the consideration of "risk to the rights and freedoms of natural persons" (Art. 35 GDPR), the reality of today's data economy is characterized by intrusive and dangerous forms of mass surveillance that strongly deviate from popular understandings of well-regulated and fair data processing (cf. Chapter 1).

However, laws are not engraved in stone; and there is also considerable room for improvement with regard to their implementation and enforcement in practice. While an extensive exploration of regulatory options for action is beyond the scope of this thesis, the following subchapters will address some implications and recommendations that emerge from our findings and the research reviewed during my doctoral studies. First, it will be highlighted that privacy-enhancing technologies alone will not suffice to overcome our current privacy crisis (Chapter 10.4) and that, in the face of growing and fundamentally unpreventable privacy threats, the enforcement of radical transparency in data processing is a powerful last resort, which should be used more widely (Chapter 10.5). I will then advocate for information inferred about people to be unambiguously covered and protected by privacy law (Chapter 10.6) and generally challenge strict distinctions between "sensitive" and "non-sensitive" data (Chapter 10.7) and between "personal" and "non-personal" data (Chapter 10.8). Addressing the shortcomings of privacy self-management, Chapter 10.9 will make the case for stronger government interventions based on the harms and risks involved in data processing. Finally, Chapter 10.10 will offer some remarks on academic freedom and current threats to the independence of internet research.

10.4 Technology Cannot Fix It

As for the dire state of data privacy around the world (cf. Chapter 1.1), technological innovations are often demanded and proposed as solutions to the problem. For instance, a new data management system devised by Tim Berners-Lee and his colleagues is expected by some to "save the internet" [385] or "fix the web" [386]. Swanson [387], a senior fellow of the American Enterprise Institute, argues that "Europe And California Get It Wrong; Technology Is The Solution To Digital Privacy" [387]. New technology trends, such as the blockchain [388], differential privacy [389, 390], smart personal information management systems [385, 386], or decentralized computing [391, 392] are often celebrated as the ultimate solution to the problem. This is well illustrated by headlines like "AI could help solve the privacy problems it has created" [393], "AI has a privacy problem, but these techniques could fix it" [389], "How blockchain could solve the internet privacy problem" [394] or "Concerns around data privacy are rising, and blockchain is the solution" [395]. Besides such major technology trends, there

are countless other, sometimes highly specialized, technical defense mechanisms, as we have exemplified in Sections P1–3, P3–3, and P4–6.1.

While people’s clinging to desperate hopes is understandable considering the current situation and while technology-based efforts at protecting people’s data are laudable and urgently needed, it can be dangerous to overestimate the protective potential of pure technology. While opaque privacy threats should be over- rather than underestimated as a precaution (cf. Chapter 2.2), the opposite applies to defense mechanisms, as overstating their abilities may lead to a false sense of security. As Le Métayer [396] puts it:

[G]reat care should always be taken not to over-emphasise the level of protection provided by technology. Protection is never absolute or irrevocable: a technology (e.g., a cryptography software) that may reasonably be perceived as secure at a given time may become unsafe later on because new attacks have been discovered; similarly, a dataset that is supposed to be anonymised using the best available techniques at a given time could possibly be de-anonymised later on, either because new auxiliary data is published or because more powerful data analysis techniques are available. The worst would therefore be to convey a misleading impression of complete protection that would lead individuals to care less about their privacy: technology can play a key role in enforcing privacy but should not be used in isolation or seen as a convenient way to forget about privacy.

When assessing the impact of technical privacy safeguards, there are some issues that need to be born in mind related to both their functional limitations and the way they are applied in practice:

- **Data controllers’ reluctance to apply self-limiting privacy safeguards.** As it represents an increasingly valuable resource in many industries [6, 256, 257], it is obvious that most companies will not refrain from collecting and processing personal data unless they have a strong incentive to do so. Accordingly, while companies are usually committed to protect their collected data against external attackers (e.g., hacker groups) and unauthorized access (e.g., by other users of the same service), they will typically not make collected data inaccessible or unusable to themselves through the use of privacy-enhancing technologies. For instance, it seems unlikely that a company would voluntarily forgo inferences it could draw from already collected data, e.g. by distorting or encrypting the data. Of course, there can be exceptions, such as when data controllers have no use for certain data anyway or want to establish an image as a privacy-friendly company. However, without strong incentives or well-enforced regulatory obligations, it should not be assumed that data controllers protect intimate personal information from their own prying eyes.
- **The consent loophole.** Even where certain data protection measures are prescribed by law (e.g., encryption, data minimization, storage limitation), these requirements can be stretched and circumvented by collecting the data subject’s consent to extensive data collection and processing [2, 310, 397]. Unfortunately, as explained in Paper 10, ordinary people cannot be expected to make free and truly informed privacy decisions in today’s

complex socio-technical environment. Logically, this not only applies to providing consent to data collection and processing, but also to deciding for or against optional privacy protections. Under EU’s current legal framework, meaningful privacy protections often depend on users’ willingness to pay for them, which is notoriously limited [398, 399].

- **Technical challenges and limitations.** Privacy-enhancing technologies often face severe hurdles, such as computational complexity, extensive power consumption, lack of usability, and technical immaturity [36, 400, 401]. While methods for privacy preservation in data mining and statistical databases have been under development for decades [402, 403, 404], there remain open challenges that impede many real-world applications [405, 406]. Existing solutions to protect against sensor-based inference attacks, such as the ones covered in Part I, have even been described as “embryonic research topics” [202]. Truly effective privacy safeguards, on the other hand, often compromise the usability or functionality of the respective product or service, giving rise to a privacy-utility trade-off [407, 408]. For instance, the privacy-invading potential of sensor data (e.g., eye-tracking data, voice recordings) can be drastically limited by inserting random noise into the signal [409, 410]. However, if done heavily – which could be required to defend against unknown attacks (cf. Chapter 2.2) – this may deprive the data of any real utility. The latter statement should not be misunderstood as an argument against the use of privacy-enhancing technologies. In fact, with regard to the many ways in which data can be weaponized against people [383], using technology to purposely reduce data availability or granularity can be necessary to protect societal interests and people’s freedoms and fundamental rights. However, the problems outlined above indicate how limited the effectiveness of mere technological safeguards is, especially in the face of quickly advancing privacy threats, such as sensor-based inference attacks.
- **Limited scope of protection.** In discussions around privacy-enhancing technologies, the primary focus is usually on *what* they protect and *how*. One aspect that is at least as important, however, is what they *do not* protect. As explained in Section P4–6.1, for example, there are various privacy protections popularly proposed for voice-enabled devices (e.g., mute feature, disconnect power supply, use of ultrasonic noise to distort the recording) that may improve speech privacy by preventing unwanted recording, but will not reliably protect against audio-based inference attacks. Similarly, a seemingly obvious solution to the accelerometer-based privacy threats described in Paper 2 would be for mobile operating systems to give users more control and transparency, e.g., by asking them every time a mobile app or visited website wants to access accelerometer data. Again, unfortunately, this “solution” will not provide reliable protection against inference attacks. As inferences are often based on complex patterns and algorithms, ordinary users cannot be expected to understand what information is indirectly revealed once they have granted access to the sensor [47, 195]. More examples for the limited scope and effect of privacy-enhancing technologies are provided in Section P10–6.2. While technical approaches have been proposed to render specific inferences impossible, these approaches typically focus on certain categories of inferred information, leaving many other categories out of consideration [199, 200, 201, 409, 410]. Also, even with protections advancing further, given the pace of technological progress and lack of transparency, it is

impossible for safeguards to guarantee that all possible inferences are covered. Strictly speaking, as explained in Section P10–5.2, not even anonymization provides complete protection against potential harms resulting from data disclosure. It is reasonable to assume that, even where many technical safeguards are used in combination, large blind spots and loopholes will remain.

Taken together, the above arguments underscore that technology alone will not solve the privacy threats caused by technology. To be precise, the problems under investigation in this dissertation are of course not only of technological origin. As *The Guardian* states in its “view on internet privacy” [411]:

It would be a mistake to see these problems as primarily technological because that would suggest that their solutions would be technological, too. In fact, the preservation of personal privacy (...) is a political and social task as much as it is one for the very few experts who understand the ramifications of mathematical magics like public key cryptography. Technological solutions will only work within a legal and political context, and the real threats to privacy come not from vulnerable widgets but weak laws (...) and feeble oversight.

There is no question that privacy-enhancing technologies play a crucial part in data protection, and it is important that laws make it mandatory for data controllers to deploy state-of-the-art technical privacy safeguards (e.g., Art. 25 and 32 GDPR). Considering the financial temptations of exploiting personal data and the widespread lack of data protection compliance among companies (cf. Paper 7), there need to be strong regulatory incentives and controls to achieve this. A greater focus should be put on technologies that reliably protect data from undue access and exploitation by the data controllers themselves.

But all this will not suffice. Technical safeguards cannot be expected to be more beneficial than the values and objectives guiding their development. And many existing problems associated with the processing of personal data are not primarily of technical nature, but involve fundamental social, ethical, and political questions yet to be answered [6, 102, 215]. Regarding the threat of inference attacks addressed in Part I, Mühlhoff contends that “no technical solution will resolve the fundamental threat to human dignity and autonomy that arises when aggregate inferences are turned into individual predictions” [121]. Technical solutions are important but clearly not enough. To tap the full potential of personal data processing while meaningfully protecting people against the resulting dangers, innovative solutions are needed on the regulatory level as well. Especially, lawmakers will need to address the question of how to deal with the failure of the notice-and-choice approach (cf. Paper 10). Some possible directions were touched on in Section P10–6.2 and will be discussed in Chapter 10.9. Also, where neither law nor technology can offer reliable protection, a drastic increase in the transparency of business and data practices may be needed as a precautionary measure, as will be further argued in Chapter 10.5.

10.5 Wide-Ranging Transparency as Ultima Ratio

To secure a minimum level of accountability and control under the given complexities of data processing, regulation should strive for the greatest possible transparency at all stages of the data processing lifecycle. As many technological trends and privacy threats cannot be averted *per se*, transparency and accountability are in some respect means of last resort and should be seen as strategic tools of utmost importance.

The research for this thesis has repeatedly been restricted by technical obscurities and the limited public availability of information, as becomes particularly evident in our investigation into mobile eavesdropping (Paper 4) but also throughout our literature-based investigations into the privacy-invading potential of certain types of sensor data (Papers 1, 2, and 3). Here are some examples of the lack of transparency prevalent in today's data economy:

- While such information is sometimes partly revealed through patents, company presentations, and privacy policies, EU's current regulation does not unequivocally require data controllers to comprehensively inform data subjects or the public about the inferences they (attempt to) draw from collected personal data (cf. Chapter 10.6). This is despite inferences becoming an increasingly accurate and efficient path to access intimate personal information (cf. Part I).
- While user permission requests are now commonplace (e.g., to authorize an app to access a smartphone's microphone), the user interfaces of most mobile devices do not provide detailed logs of when certain resources (e.g., accelerometer) are accessed by the system or mobile apps – and when resources are entirely blocked from access.
- Closed-source systems and obfuscated code make it difficult, if not impossible, even for experts, to understand the inner workings and trace the actions of operating systems and individual applications (what specific data they collect, when and where they send it, how they process it, etc.).
- Even service providers and data controllers themselves often cannot observe all stages of data processing [412]. For instance, as explained in Section P5-3.1, software components embedded into mobile apps which are provided by external companies ("third-party libraries") inherit privacy permissions and may exfiltrate and use personal data in ways not even transparent to the host app.
- The GDPR does not require data controllers to provide the names of third-party data recipients (categories of recipients can suffice, cf. Art. 13 and 15 GDPR). In the case of categories, recipients cannot be checked for scandals and reputation. Note that some of today's biggest privacy threats are posed by companies that most people have never heard of or directly interacted with (e.g., ad networks, data brokers) [2, 259].
- The GDPR's "legitimate interest" clause is increasingly used by data controllers as an alternative legal basis for data processing instead of asking for consent [22, 413]. The European Parliament has expressed concern "that 'legitimate interest' is very often abusively mentioned as a legal ground for processing" [414]. While high-profile

investigations have shown that claims of legitimate interest can be legally challenged with significant consequences for data controllers [415, 416], the weighing of interests involves subjectivity and is typically conducted by the data controller without supervision or transparency obligation [22, 417].

- According to Art. 35 GDPR, data controllers have to conduct a data protection impact assessment (DPIA) when data processing “is likely to result in a high risk to the rights and freedoms of natural persons”. However, although recommended by the Article 29 Working Party, the GDPR does not require data controllers to publish the methods and results of their data protection impact assessments [384].
- While many contemporary privacy laws afford people the right to demand a copy of the personal data that companies store about them (e.g., Art. 15 GDPR), exercising this right is often met with silence or responded to in a deceptive and insufficient manner (cf. Paper 7).
- Although cookie banners suggest improved control over the extent of online tracking, extensive and hidden web tracking activity is still ubiquitous on the internet and many websites also display deceptive information in this regard [418].

While Art. 5(1)(a) GDPR stipulates as a general principle that “Personal data shall be processed (...) in a transparent manner in relation to the data subject”, the above examples show that there is a lack of sensible context-specific transparency obligations, and also a severe lack of effective enforcement of existing laws. Wide-ranging transparency is important not only to expose misconduct, but also to subject data practices to public scrutiny, enabling an informed debate on what types of data use should be considered appropriate and acceptable [419]. While ordinary people often do not understand enough of the subject matter and do not have the time to really do much with additional published details (cf. Sections P10–2.1 and P10–2.2), it would be valuable to society if not only supervisory authorities but also journalists, privacy researchers, and civil society watchdogs had more comprehensive information access and insight. Such actors could be granted special rights to information to enable public welfare purposes, such as conducting systematic analyses to uncover discriminatory business practices (cf. Chapter 10.9).

Specifically for the issue of inferential analytics addressed in this thesis, much more transparency is clearly possible than currently realized. For example, Fischer [203] demands that data controllers should provide information on “who aims to draw inferences, who drew existing inferences, who is processing drawn inferences for what purposes and what the inferences state about the individual.” Further, Clarke [420] demands: “Where actions are taken based on inferences drawn from data analytics, [we need to] ensure that the rationale for the decisions is transparent to people affected by them, and that mechanisms exist whereby stakeholders can access information about, and if appropriate complain about and dispute interpretations, inferences, decisions and actions.” As will be further explained in Chapter 10.6, this is not currently guaranteed by EU legislation.

Of course, when contemplating legal transparency obligations, it needs to be born in mind that companies have a legitimate concern to protect their intellectual property and trade

secrets. Even though closed-source systems and non-disclosure agreements can be a major impediment to data protection efforts and privacy research (cf. Sections P1–3, P5–6, and P6–6), companies may have good reasons for choosing these approaches. Regulatory solutions will need to strike a balance between companies’ right to secrecy on the one hand and people’s right to privacy, transparency, and freedom from harm on the other.

It should also be understood that transparency does not resolve the complexity of modern data collection and processing (cf. Section P10–2.4). For instance, studies have shown that even when software is open source, a thorough review by experts is often not feasible due to the sheer amount of code [421, 422, 423]. As Pfleeger [422] states: “Calls for open source code require a workforce capable of and willing to review code for undesirable features.” Transparency-enhancing legal requirements should therefore be linked to specific objectives and always be checked for feasibility.

In any case, data processing will never be completely transparent. There will always be methods to cover up reprehensible data practices. For effective legal oversight and enforcement, it will therefore remain crucial to encourage, facilitate, and protect whistleblowers (e.g., through secure hotlines, legal advice, financial support, identity protection). Recent cases, such as the 2018 Facebook-Cambridge Analytica data scandal [424], have demonstrated how important internal insights can be for the revelation and prosecution of personal data misuse. Despite recent progress on the legal front (e.g., EU’s 2019 directive for whistleblower protection [425]), various problems remain. Around the world, numerous countries still offer no comprehensive level of whistleblower protection [426]. And while the numbers may have somewhat improved since then, EU’s Directorate-General for Justice and Consumers [427] stated in 2018 that only “15% of citizens know about existing rules on whistleblower protection”. Also, in a recent legal analysis, Vigjilence Abazi concluded that EU’s new whistleblower directive “is an important legal development but (...) only in the early stages towards meaningful protection, rather than a ‘game changer’ for whistleblowers in the EU” [428, p. 642]. Furthermore, regarding the law’s real impact, she notes: “[W]hether the Directive will attain the expected high standards of protection depends, *inter alia*, on the transposition of the rules into national law, the enforcement of the Directive’s protections and the embeddedness of the rules in organisational culture” [428, p. 641].

10.6 Recognizing Information Inferred about Individuals as Personal Data

By providing an overview of categories of personal information that can be inferred from certain types of mobile and IoT sensor data, this thesis has highlighted the privacy-invading potential of modern inferential analytics (cf. Chapter 2). There is no doubt that inferences drawn from patterns and correlations in collected data are becoming an increasingly efficient way to access intimate information and assemble detailed profiles about unsuspecting individuals. Jordan M. Blanke even contends that inferences drawn from collected data have “become more dangerous to individual privacy than the vast collection and storage of the data itself” [56, p. 81]. With regard to the risk of spurious correlations and incorrect inferences, the Article 29 Working

Party⁸ has pointed out that it is “crucial that data subjects/consumers are able to correct or update” data inferred about them [429, p. 47].

There have been legal efforts to put inferred personal information under the scope of privacy law. For instance, the California Consumer Privacy Act (CCPA), which was introduced in 2018, specifically covers “inferences drawn” as part of its definition of personal information [430]. Under EU’s current regulation, however, the situation is less clear and essentially leaves the question of applicability to the courts. Fabienne Ufert observes that “the concept of personal data is not exhaustively defined [in the GDPR] (...). It is especially unclear if inferences drawn from personal data – something that AI is particularly good at – form part of the concept” [204, p. 1097]. The GDPR’s ambiguous position regarding inferred data has been recognized as a legal loophole [22, 23]. Most recently, legal expert Alina Skiljic criticised that the GDPR “does not properly define, regulate or refer to [inferred data.] (...) [I]t is inappropriate that not even a single article in the GDPR is dedicated to inferences specifically” [23]. In a comprehensive legal analysis of the GDPR, Wachter and Mittelstadt [198, p. 498ff.] write:

The legal status of inferences is heavily disputed in legal scholarship, and marked by inconsistencies and contradictions within and between the views of the Article 29 Working Party and the European Court of Justice (ECJ). (...) Compared to other types of personal data, inferences are effectively ‘economy class’ personal data in the [GDPR]. Data subjects’ rights (...) are significantly curtailed for inferences. (...) [EU] data protection law focuses primarily on mechanisms to manage the input side of processing. (...) [T]he few mechanisms in European data protection law that address the outputs of processing, including inferred and derived data, profiles, and decisions, are far weaker.

For drawn inferences to constitute personal data within the meaning of the GDPR and thus fall within the law’s scope of application, four conditions have to be fulfilled: According to Art. 4(1) GDPR, the inferences must be (1) *information* that (2) *relates to* an (3) *identified or identifiable* (4) *natural person* [431]. These conditions may seem straightforward and easy to satisfy, and they can be, but in many cases the legal reality is more complex. With regard to the notion of “personal data” under EU data protection law, Lorenzo Dalla Corte states: “Despite the crucial importance of [this] notion, the boundaries of the concept are often blurry. Ascertaining whether data is personal frequently depends on each individual processing’s concrete context and characteristics. As a result of the contextual and relative character of the notion of personal data (...) much is left to the discretion of the interpreter” [431, p. 1].

A legal analysis by Celin Fischer confirms that, in principle, the definitions in the GDPR provide for the possibility of treating inferences as personal data and obliging data controllers to inform people about “intentions of drawing inferences about them, (...) the existence of inferred data about them and (...) the purposes the inferences are intended to be used for” [203, p. 65]. However, whether this really is mandated by the law ultimately depends on a case-by-case assessment [23, 432]. Fischer points out that “each inference, when considering its state of being personal data, needs to be considered within the concrete processing instance” [203, p. 40f.]. Due to the GDPR’s vague wording, the outcomes of such examinations – and, thus,

⁸For an explanation of Article 29 Working Party, see *supra* note 2 in Chapter 1.1.1.

the law’s protective effect against intrusive inferences – heavily depend on interpretations and jurisprudence [198, 23, 203]. As Mittelstadt et al. [433, p. 14] put it: “The GDPR can be a toothless or a powerful mechanism to protect data subjects dependent upon its eventual legal interpretation: the wording of the regulation allows either to be true.”

Some of the main arguments brought forth against recognizing inferences as personal data are that they are not “provided by⁹ the data subject” [434] and usually consist of probabilistic assumptions, not verifiable knowledge [198, 435]. While certain inferences can of course be verified (e.g., by investigating through other channels whether a person’s predicted income or age range is correct), other inferences are subjective and inherently unverifiable, at least not in the present (e.g., data subject being classified as “untrustworthy”, “likely to commit a crime” or “aggressive driver”) [198]. It is important to understand, however, that any information or assessment about a person, verifiable or unverifiable, correct or incorrect, can have real consequences for the data subject [5, 19]. This is what matters in the end. While inference methods are becoming more and more accurate and efficient with technological progress, they do not need to be 100% accurate for many attacks and intrusive profiling purposes, as I have pointed out in Chapter 2.2. Also, completely inaccurate methods may be used nonetheless, causing additional discriminatory side-effects. Thus, in line with Kamann and Braun [436] and Wachter and Mittelstadt [198], I argue that inferences about individuals should fall within the scope of data protection law irrespective of verifiability because, in either case, they may have an impact on people’s lives.

Aside from how inferred information should be handled and communicated, another important question is: To what extent does the drawing of inferences constitute a data processing purpose that needs to be made transparent under EU law? According to Art. 5(1)(b) GDPR, personal data shall be “collected for specified, explicit (...) purposes and not further processed in a manner that is incompatible with those purposes”. Art. 13(1)(c) further demands that the intended purposes of data processing need to be clearly stated at the time when personal data is collected from a data subject. One could expect that the intended application of inferential analytics would need to be stated as such a purpose. However, again, the GDPR does not explain how such situations should be dealt with exactly, e.g., in terms of how detailed the provided information should be [203, p. 65]. It also depends on jurisdiction whether the process of drawing inferences can be described as “compatible” with previously stated purposes, which under the GDPR obviates the need for further legitimization (cf. Art. 5(1)(b) and Recital 50 GDPR). For the interested reader, Joe O’Callaghan [432] offers a more detailed insight into how the GDPR’s purpose limitation principle relates to inferred data. Naturally, at the time of data collection, it can be impossible to disclose in full detail “the processing of inferred data that has yet to be generated or determined (...) particularly where the insights are unforeseen or unforeseeable” [432, p. 10].

Also, the GDPR’s principle of data portability, which grants people the right to receive the data that controllers hold about them “in a structured, commonly used and machine-readable format” will likely not cover inferred data as the principle only applies to data that was provided to a controller by the data subject (Art. 20(1) GDPR).

⁹Observed data, such as recorded browsing behavior or location data, is also legally defined as (indirectly or passively) “provided by” the data subject [198, p. 516].

In conclusion, despite being celebrated as one of the most progressive privacy laws in the world [171, 172, 437], the GDPR does not reliably protect individuals against the privacy threats posed by inferential analytics. This lack of protection can be exploited by data controllers, as illustrated by a complaint filed by the consumer watchdog Privacy International [259] warning that “many companies (...) seem to work under the assumption that derived, inferred and predicted don’t count as personal data, even if they are linked to unique identifiers or used to target individuals.”

While the GDPR offers possible ways of dealing with this issue, improvements in regulatory certainty would be advisable. To avoid complicated case-by-case decisions and enable an efficient widespread application of the law in practice, also with regard to the varying interpretations across EU member states [438, 439], clear and specific wording should be chosen that leaves as little room as possible for weak interpretations that render the GDPR “toothless”. As Fischer [203, p. 65] states: “Only with enough detail, will the provisions of the GDPR be sufficient to mitigate the risks, inaccurate inferences can pose for individuals.” Besides added detail and clarity, sensible regulatory measures could ensure:

- that data subjects’ ARCO rights (access, rectification, cancellation, and opposition) equally and unequivocally apply to inferred data, as long as it pertains to an identifiable individual,
- that data controllers have to publish comprehensive¹⁰ information about all types of inferences that they are drawing – and attempting to draw – about individuals, and
- that inferred data is taken into account when examining potential harms resulting from data use and in the identification of appropriate countermeasures (see Chapter 10.9). In high-risk scenarios, this may include legal prohibitions of using certain types of inferences for certain purposes, or prohibitions of drawing certain inferences in the first place.

It should be taken into consideration that besides vagueness and loopholes in existing data protection laws, trade secrets (cf. Sections P1–3 and P6–6) and the complexity of modern data processing systems (cf. Section P10–2.4) present additional obstacles to ensuring transparency and proper regulation of inferential analytics. Fischer, for example, recognizes that her “analysis of applying the GDPR to inferences drawn (...) is done without taking into account some important barriers. Intellectual property rights and trade secrets, among others, will likely play an important role in regard to the impact, the GDPR will have on the challenges posed by inferences” [203, p. 69f.]. With respect to inferred data, Ufert writes: “[T]here exists the risk that controllers use [the] complexity and the autonomy of AI as an excuse to circumvent their information [obligations]” [204, p. 1095]. While certainly not easy to solve, future regulatory approaches should put a focus on this problem, e.g., by putting strict legal boundaries on ethically indefensible data uses and excessive information inference (cf. Chapter 10.9) and by improving enforcement and accountability through transparency-enhancing policies and robust whistleblower protection (cf. Chapter 10.5).

¹⁰It is often not clear what information will be inferable from collected data and thus not always possible to disclose all relevant details at the time of data collection (cf. Section P10–2.5). Therefore, processing purposes should constantly be updated, as has been proposed, for example, by UK’s Information Commissioner’s Office: “As your processing purposes become clearer, update your privacy information and actively communicate this to people” [440].

10.7 “Sensitive” vs. “Non-Sensitive” Data

Both technical privacy safeguards and data protection law often vary in their level of protection between “sensitive” personal data and data rated as less sensitive or even non-sensitive. Art. 9 GDPR, for example, affords particular protection to the processing of “special categories of personal data”, such as information related to a person’s ethnic origin, genetics or political and religious beliefs. Similarly, mobile operating systems typically classify the access to certain types of data and system resources, such as the microphone or the user’s contact list and GPS location, as sensitive, thus imposing extra safeguards (e.g., always requiring user permission), whereas other data sources in mobile devices have received less protection [4, 40]. With regard to publicized data scandals and the more obvious and well-known forms of data misuse (e.g., identity theft, discriminatory scoring), it is very understandable that certain data categories are widely viewed as more sensitive than others [47, 125, 127]. However, the assumption that such a distinction is universally applicable is a dangerous fallacy [70, 441].

Depending on data controllers’ possibilities and intentions, even seemingly non-sensitive data can serve for nefarious ends and result in considerable harm for the data subject (and potentially others, cf. Section P10–4). For instance, as discussed in Part I of this thesis, even data types widely believed to be completely innocuous can, through the lens of data analytics, implicitly reveal various categories of intimate personal information. Even basic input modalities like keyboard keystrokes, mouse clicks, and touchscreen taps can be sufficient to not only biometrically identify users, but also infer information about their state of health, level of stress, and physical dexterity (cf. Paper 6). Serving as another example, the wide range of personal information inferable from mobile and IoT accelerometer data was illustrated in Paper 2, including a user’s level of intoxication, daily activities, smoking habits, driving behavior, approximate location, and much more.

Considering the higher level of protection afforded to cameras, microphones, GPS, and other “privacy-sensitive” sensors, malicious parties could try to use less-protected sensors, such as accelerometers, as substitutional data sources for gaining intimate insights about people [4]. And the same holds true, of course, for all sorts of seemingly benign non-sensor data (e.g., telephone metadata [296]). In general, while the wording “sensitive data” can be a useful and appropriate rhetorical device to underscore imminent threats of abuse, sensitivity is not a static property of data but is situational and can vary, for example, based on time, online/offline context, and the organizations involved in the data processing [441]. As illustrated above, especially the classification of data as “non-sensitive” is problematic because it may lead to possible threats being overlooked.

In conclusion, as Wachter [442] states, “[o]utdated, ineffective and fluid categorisations of data as (...) sensitive or non-sensitive must be abandoned.” Wachter and Mittelstadt [198] not only criticize the significant differences in the level of legal protection for “sensitive” vs. “non-sensitive” data under the GDPR, but also the law’s ambiguous wording: While inferred data can theoretically be recognized as “sensitive data”, this requires vaguely defined conditions to be fulfilled and ultimately depends on court decisions. As will be further argued in Chapter 10.9, the level of technical protection and governmental intervention in data processing

should be based on the context-dependent level of expected harm, not on the mere category of collected data.

10.8 “Personal” vs. “Non-Personal” Data

As with the above observations regarding “sensitive” vs. “non-sensitive” data, the strict and consequential legal distinction between “personal” and “non-personal” also needs to be questioned. The scope of data protection laws is typically limited to personal data, i.e., information “relating to an identified or identifiable natural person” (Art. 4 (1) GDPR). Thus, the data subject rights and data controller obligations prescribed by the law do not apply to anonymized data (Art. 2 GDPR). While the de-identification of data has important benefits in terms of privacy protection [443], leaving anonymous data completely unregulated is problematic for two reasons.

First, there is extensive evidence in the literature that seemingly anonymous data can often be linked back to individuals [396, 444, 445]. Ohm described this issue as the “surprising failure of anonymization”, warning that it is often possible to “‘reidentify’ or ‘deanonymize’ individuals hidden in anonymized data with astonishing ease” [446, p. 1701]. Narayanan and Felten even contend that for certain types of data, there simply is no reliable form of anonymization available and also “no evidence that it’s meaningfully achievable [in the future]” [447, p. 1]. The pervasive interlinkage of data sources and advancing capabilities of modern computers and data analytics methods continue to increase the efficiency of de-anonymization attacks [4, 445, 448, 449]. For instance, as richly illustrated in Part I, machine learning algorithms can be used to biometrically identify people based on patterns and correlations in sensor data, such as as a user’s voice and speech characteristics in audio recordings (Paper 1), recognizable patterns in accelerometer data (e.g., gait features, hand gestures, head movements) (Paper 2) or iris textures and eye-movement behavior captured by eye trackers (Paper 3).

Second, as we have discussed in Sections P10–4 and P10–5.2, “anonymous” data can be used for various tracking and profiling purposes which may inflict harm on individuals, population groups, and society at large. For example, companies can build detailed user profiles without real names attached to the data, and then use this information for behavioral ad targeting, the personalized tailoring of persuasive (e.g., political) messages, price discrimination, or exclusion from certain services based on a customer’s estimated risk or profitability [19, 198, 383, 450]. For such purposes, being able to distinguish between “anonymous” individuals can be sufficient. In a recent publication, I provide examples of how mobile and IoT sensor data can be used for the tracking and profiling of unidentified users [4, p. 154]. For example, small calibration errors in sensors originating from imperfections in the manufacturing process can be used to tell the devices of different users apart (“sensor fingerprinting”), without the need to know their real identity [266].

Furthermore, “anonymous” data from many people can feed machine learning algorithms that are used to infer sensitive information about individuals based on pattern recognition [121]. As Hildebrandt and Gutwirth state, sensitive inferences are often not drawn “from the personal data of the categorised person but inferred from a large amount of often anonymised data of

many other people” [451, p. 11], which is just another example of how anonymous data can be weaponized and have an impact on people’s lives.

In sum, this means that data “anonymization”, which regularly happens in modern data processing [445], can be a means of evading the restrictions of data protection laws [19, 198] but does not effectively remove the data’s potential to harm individuals and groups. Thus, de-identified data must not be ignored when assessing potential ramifications of data collection and use, and when designing corresponding organizational, technical, and regulatory safeguards. Of course, modern privacy laws such as the GDPR take into account the risk of de-anonymization. For instance, the GDPR recital 26 [452] states:

To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly. (...) [A]ccount should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments. The principles of data protection should therefore not apply (...) to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable.

This means that, at least in theory, data that can easily be de-anonymized does not legally count as anonymous data. Accordingly, some of the examples of “anonymous” data mentioned above would likely not satisfy the GDPR’s requirements for anonymous data and – if challenged – corresponding claims by data controllers may not withstand legal scrutiny.

Now, if the law already addresses the risk of de-anonymization, where is the problem? Apart from the general problem that the GDPR lacks effective enforcement [24, 453, 454] and, therefore, it can be assumed that many classifications of data as “anonymous” will simply not be challenged in practice, the points raised in this subchapter lead to the following two open questions: (1) With regard to the limits of anonymization, could it be that the very notion and legal recognition of “anonymous data” is unrealistic in most cases and, thus, largely misleading? Based on recent research on this issue, Alex Hern from *The Guardian* concludes that “successfully anonymising data is practically impossible for any complex dataset” [455]. (2) In cases where data is successfully declared as “anonymous” and thereby taken outside the scope of data protection law, how can regulation nevertheless ensure sufficient transparency in data processing and prevent negative ramifications?

Of course, as long as “anonymous” data is not de-anonymized, it does not make sense to demand data subject rights for that data because there is no identified data subject in the first place – and thus nobody whom these rights could be ascribed to. Hence, there is certainly a *raison d’être* for laws that specifically apply to personally identifiable data. But, as will be further elaborated in Chapter 10.9, the focus of political and scientific discourse should be on the various ways in which data can be weaponized and how these threats can be prevented or mitigated, rather than zeroing in on “personal data” per se.

To address the threat of de-anonymization, it might be helpful, as the German Data Ethics Commission [450] has suggested, to make the linking of allegedly “anonymous” data to real

identities a criminal and punishable offense. But even if this idea gains momentum, it should not be forgotten that, depending on the use, even anonymous data can cause considerable harm.

10.9 Call for a Stronger Risk-Based Regulation of Data Use

Paper 10 provided an overview of the numerous obstacles that render the legal principle of privacy self-management useless in practice, showing that people's privacy choices are typically irrational (Section P10-2), involuntary (Section P10-3) and/or circumventable (Section P10-5). The paper also highlighted the problem that the individualistic approach disregards the consequences that privacy choices of individuals have on other, seemingly uninvolved people and the wider society (Section P10-4). These observations make it apparent that data protection regimes which heavily rely on privacy self-management, such as EU's GDPR, are not truly fit for purpose.

From the perspective of individual freedom and empowerment, it is tempting to argue in favor of strengthening people's informational self-determination (it should be acknowledged that I have previously done so as well, e.g., [4, p. 155f.]). Upon closer examination, however, this approach has severe limitations and is not a sufficient solution to the ongoing privacy crisis. To meaningfully address the dangers of data processing in an increasingly complex and technologized world, authorization for many types of data use should not (or, at least, significantly less) be obtained via consent of individual data subjects but rather allocated or withheld by some collective entity based on risk assessments (cf. Section P10-7). Additionally, in some areas a "fiduciary duty" requirement could be introduced, i.e., a rule that data must exclusively be used in the data subject's interest and not for any further purposes [456]. Of course, a purely objective assessment of risks and benefits is not achievable, and there will always be the possibility that dangers are overlooked or misjudged. But since the task of assessment is clearly beyond the resources and capacities of individual data subjects, and since the use of one person's data can ultimately have an impact on other people as well, a public and expert-driven form of decision-making seems more promising and sensible than decisions at the individual level.

The GDPR already stipulates that when "legitimate interest" is being claimed as a legal basis for data processing (Art. 6(1)(f) GDPR) and generally when data processing is "likely to result in a high risk to the rights and freedoms of natural persons" (Art. 35(1) GDPR), data controllers must perform an assessment of potential consequences ("balancing test" [457] or "data protection impact assessment" [458], respectively). While such rules clearly are important, it can be doubted whether they suffice to effectively avert harm – especially in the way they are currently implemented. It is problematic that the assessment results are usually not publicly available (cf. Chapter 10.5) and that the rules may often be circumvented through meaningless "rubber-stamping" procedures [204, 414]. An even more fundamental limitation lies in the fact that these assessments are not typically carried out by independent experts and elected representatives of the people (or their appointees), but by the data controllers themselves, which comes with significant conflicts of interest. It is evident that current frameworks of data governance, including these rules and the way they are enforced, fail to align corporate

data practices with the will of the general population and fundamental values of our society (cf. Chapter 1.1).

The above call for more institutionalized forms of risk assessment on a societal and political level does not mean that the element of individual control should be eliminated from privacy protection entirely. A certain degree of choice and control over the availability of intimate personal information to others is required to satisfy basic human needs [459, 460]. Unless overwritten by urgent societal interests (e.g., law enforcement, public health, public safety), hiding one’s information and retreating into privacy should always be permitted based on personal preferences. However, for data controllers, depending on the intended purpose and the level of risk involved, there should be legal limits to the amount and types of data they are allowed to collect, infer, and use. As O’Callaghan proposes: “It may be, for example, that systems are designed to have certain ‘blind spots’ (...) both where it comes to the collection and surveillance of source data and also to the inferences that are generated from such source data” [432, p. 15]. For some readers, the idea of blocking certain types of data collection and processing irrespective of user consent may appear revolutionary or even dangerous because, after all, the notice-and-choice approach has been the legal standard of privacy protection across the Western world for decades and is widely seen as an embodiment of fundamental values, such as autonomy and human dignity (cf. Section P10-1). Accordingly, challenging this legal framework may be perceived as an attack on people’s rights and freedoms.

The truth is, however, that under the complex circumstances of modern life, privacy self-management does not strengthen or even preserve people’s autonomy with regard to their data. Quite the opposite: Despite the notion of “freely given” and “informed” consent (cf. Art. 4 GDPR), people have completely lost control over their data [1]. And, given the arguments provided in Section P10-2, it seems likely that people’s loss of control and the amount and granularity of data collected about them are much vaster than commonly assumed. Beneath the widespread lack of understanding and awareness, beneath the confusion and obscurity, organizations can hide questionable data practices. As Litman-Navarro [397] states, many privacy policies only serve to “opaquely establish companies’ justifications for collecting and selling your data. The data market has become the engine of the internet, and these privacy policies we agree to but don’t fully understand help fuel it.” Although groundbreaking in some sense, EU’s GDPR, which is built upon the principle of privacy self-management, will certainly not solve the most pressing issues related to personal data collection and processing we as a society are currently facing.

While there are many other reasons for this (cf. Fig. 1 in Paper 10), the topic of inference attacks discussed in Part I of this thesis is an illustrative example of the overwhelming technological complexities we face in modern life. The fact that seemingly benign data can be sufficient to infer intimate information about people (cf. Papers 1 and 2) undermines the individual oversight and control suggested by privacy self-management. User studies [47, 193, 195], including our own study presented in Paper 4, have confirmed that most people’s understanding of sensor-based inference attacks is very limited.

As an example of why privacy choices are often not only uninformed but also unfree, I would like to elaborate on the issue of nudging and manipulation outlined in Section P10–3.2. The astonishing possibilities and inherent dangers of nudging methods are amplified by recent

advances in data analytics and machine learning. With the help of modern technologies, it is possible to predict people's behavior (e.g., online browsing behavior) with increasing accuracy, to automatically compare the effectiveness of different marketing and communication approaches in real-time (e.g., A/B testing), and to automatically refine and micro-personalize persuasion techniques based on individual behavior and preferences [102, 215, 377]. These capabilities, which are widely employed for improving response rates in online advertising, can of course also be applied by companies to influence people's privacy choices and achieve high opt-in rates [31, 310, 461]. Thus, amassed information about the behavior, preferences, mental functioning, and psychological vulnerabilities of individuals can be used to elicit consent and extract even more personal data, resulting in a vicious cycle of privacy loss.

In view of all the above and the arguments summarized in Fig. 1 of Paper 10, individuals are not capable of managing the wide-ranging risks associated with modern data collection and processing. To overcome this problem, novel forms of government intervention are needed. In describing the regime of privacy self-management, Solove [31, p. 1880ff.] writes:

Privacy self-management takes refuge in consent. It attempts to be neutral about substance – whether certain forms of collecting, using, or disclosing personal data are good or bad – and instead focuses on whether people consent to various privacy practices. (...) Privacy self-management cedes substantial responsibility for preserving privacy to individuals, and it assumes that the primary harm to be redressed is nonconsensual data collection, use, or disclosure. (...) Any way forward will require the law to make difficult substantive decisions. (...) Under privacy self-management, most forms of data collection, use, or disclosure are acceptable if consensual. Consent often becomes a convenient way to reach outcomes without confronting the central values at stake. To move forward, this kind of neutrality cannot be sustained. (...) Privacy law has said far too little about the appropriate forms of collection, use, and disclosure of data. (...) [T]he law must take a stronger stance about substance.

In conclusion, data practices should be politicized and recognized as ethical issues of societal proportions. The primary focus of corresponding regulation that replaces and/or complements the notice-and-choice approach should be on how the data is used and the (anticipated) resulting consequences. Some proponents of this opinion and preliminary ideas of how it could be realized were presented in Section P10–7.

One example for such a risk-based approach that is currently being contemplated by lawmakers is the *Artificial Intelligence Act* (AI Act) proposed by the European Commission in April 2021 [462]. The proposal envisions that AI applications should be classified based on their harm potential. High-risk applications would face new mandatory checks and requirements (e.g., assessment by an independent third party). Certain particularly harmful AI practices could even be prohibited, for example: “practices that have a significant potential to manipulate persons through subliminal techniques beyond their consciousness (...) [.] AI-based social scoring for general purposes done by public authorities (...) [and] the use of ‘real time’ remote biometric identification systems in publicly accessible spaces for the purpose of law enforcement“ [462]. While moving in a promising direction, the proposal has raised concerns and criticism [463, 464,

465]. For instance, there is concern that legal harmonization between EU member states could lead to systems “being scrutinized less thoroughly” under the new law [463]. It has also been demanded that the law should not exclusively focus on the vaguely defined technology area of “artificial intelligence” but more generally apply to “automated / algorithmic decision-making systems” and that the risk assessment process should be inclusive and transparent, involving not only technical experts but also civil society actors and independent observers [463]. Further, civil rights advocates urge that the law should include a complete ban on automated systems to predict criminal behavior (“predictive policing”) to mitigate risks of discrimination [464, 466]. Griff Ferris, the Legal and Policy Officer at criminal justice watchdog *Fair Trials*, states: “The only way to protect people from these harms and other fundamental rights infringements is to prohibit their use” [467].

The consistent implementation of a risk-based regulation of data use might require fundamental changes and restructuring in certain industries and business models (e.g., advertising, insurance, credit scoring). As Cory Doctorow argues in his book *How to Destroy Surveillance Capitalism*: “[I]f we’re to have any hope of destroying surveillance capitalism, we’re going to have to destroy the monopolies that currently comprise the commercial web as we know it. Only by breaking apart the tech giants that totally control our online experiences can we hope to return to a more open and free web – one where predatory data-harvesting is not a founding principle” [468]. Privacy-friendly alternatives to surveillance-based business models (e.g., search engines like *DuckDuckGo.com* and *Startpage.com*, web browsers like *Firefox* and *Brave*, instant messengers like *Threema* and *Signal*) could be incentivized or even become the legally prescribed minimum standard in terms of data protection.

There have recently been promising developments in the political sphere. In the EU, a coalition of political leaders, companies, and civil society organisations (“Tracking-Free Ads Coalition”) [469] and the European Data Protection Supervisor [470] are currently pursuing stricter rules concerning tracking-based online advertising. Several restrictions for behavioral advertising are being incorporated into the proposal for the EU’s prospective *Digital Services Act* [471, 472]. Furthermore, the Belgian Data Protection Authority has recently ruled that a mechanism widely used by companies to collect users’ consent to data processing violates multiple provisions of the GDPR [473]. This consent popup system, the so-called *Transparency and Consent Framework* (TCF), was portrayed as a “a cross-industry best practice standard” for years [474], relied on by tech giants such as Google and Amazon and present “on 80% of the European internet” [475]. The ruling from the Belgian supervisory body states that, among other problems, consent obtained via the TCF is “currently not given in a sufficiently specific, informed and granular manner” [473, p. 115]. It is encouraging that this problem is at least to a certain extent being recognized and addressed in law enforcement already. At the same time, on closer inspection, the above-cited judgement of the Belgian Data Protection Authority equally applies to almost all privacy decisions we make in everyday life (cf. Paper 10) and it would be misguided to assume that a few adjustments to systems like the TCF (e.g., more granular information) would magically lead to free and informed privacy decisions.

A recent example of political advocacy beyond EU borders is Amnesty International’s [476] call for a legal ban on “surveillance advertising”. Similarly, it has been demanded that “[d]ata trading for ad revenue must be regulated like finance, aviation, medicine, and power” [477].

For a thorough risk-benefit analysis across industries, it should be carefully evaluated how much added value the use of detailed personal information really brings for certain products and services – not primarily in terms of companies’ profit, but in terms of general welfare.

It is a recognized problem that many services collect more data than required for their advertised functionality [2, 66, 262, 478, 479]. Therefore, where possible, it should be investigated and defined which types of data are necessary for the provision and development of beneficial products and services, thus setting industry-specific standards for data minimization and allowable data collection. EU privacy law already demands that data controllers take into account the technological state of the art when choosing data protection measures (Art. 25(1) GDPR) and that “by default, only personal data which are necessary for each specific purpose of the processing are processed” (Art. 25(2) GDPR). However, such provisions concerning data protection “by default” are regularly being circumvented by collecting data subjects’ supposedly “informed” consent [480, 481] and through other legal loopholes (Section P10–5). To avoid this, special permissions from supervisory bodies could be made a requirement to expand data collection beyond the generally defined limits (e.g., when additional data is needed to offer special features or for the purpose of exploratory innovation). In these cases, too, the use of data should be subject to clear rules and permission should depend on estimated benefits and harms for individuals and society as a whole.

This thesis cannot offer a comprehensive and detailed picture of what a future data governance framework may look like. This important task will be left as an avenue for future research (cf. Chapter 11). It already seems clear, however, that to effectively deal with the dangers posed by technology trends, such as “big data” and machine learning, data protection law will need to be interlinked and work in conjunction with anti-discrimination law [482, 483] as well as with competition and consumer protection law, as has been pointed out by Wojciech Wiewiórowski, the current European Data Protection Supervisor [484].

To address harms that are caused by algorithms (e.g., through discriminatory scoring or recommender systems) but are not easily detected in individual cases, group-level advocacy and legal representation have been proposed as a remedy (“group privacy” [485, 486]). In practice, such an approach could be implemented in the form of collective action lawsuits [483]. Even if there is no single clearly identifiable victim, discriminatory practices can be uncovered through systematic analyses conducted, for example, by journalists and civil society organizations [487, 488, 489]. Accordingly, it has been demanded that the law should afford watchdog organizations a right to collective action against such harms [483].

In addition to legislative changes, governmental authorities should also assume a more proactive role in the area of compliance monitoring. To take more of this burden off the shoulders of individual data subjects and civil society organizations, random compliance checks by authorities (e.g., undercover investigations using fake user accounts, cf. Paper 7) might be a sensible measure – for which, of course, supervisory bodies need to be equipped with corresponding human and material resources, which is not currently the case [24, 450, 453].

10.10 A Note on the Independence of Internet Research

During my doctoral studies at the Weizenbaum Institute for the Networked Society, I enjoyed the invaluable privilege of true scientific freedom. Whilst of course being part of an institute and a research group with their own research agendas, I was always completely free in choosing the specific foci of my research projects by myself or, if I wanted to, in consultation with prospective collaborators. This freedom can be attributed to the Weizenbaum Institute's public funding [490] and the liberal management style shared by my superiors, Prof. Dr.-Ing. Ina Schieferdecker and Prof. Dr. Bettina Berendt, and the leader of our research group, Dr. Stefan Ullrich. I am very grateful for the trust and confidence they placed in me.

The granted freedom not only allowed me to choose research topics that naturally drew my interest, but also allowed me to take uninhibited critical perspectives, to address politically charged topics, and also protected my research results from external influence, such as corporate interests. Of course, the work of publicly funded institutes is also subject to evaluation: The continuation and funding depends on the opinion of evaluators and the goodwill of relevant policy makers, who theoretically may attempt to exert undue influence on the research being carried out (e.g., by expressing overly specific expectations regarding research outcomes). However, during my time at the Weizenbaum Institute, I did not experience the slightest such attempt to influence the work of our research group.

Recent journalistic investigations have shown that Europe's leading tech policy institutes significantly depend on funding from tech giants like Google, Facebook, Microsoft, and Amazon [491]. This includes research into AI "ethics", competition in digital markets, and companies' privacy practices. Clarke, Willimas, and Swindells state that, "While this funding tends to come with guarantees of academic independence, [it] creates an ethical quandary where the subject of research is also often the primary funder of it" [491]. The power of corporations to intimidate and suppress critical research focusing on their own business is also well illustrated by a recent case where an Instagram monitoring project from AlgorithmWatch was shut down after threats from Facebook [492] and by an incidence where Facebook shut down the accounts of researchers at New York University who have been among the company's biggest critics [493].

Considering the excessive economic power and questionable business practices of Big Tech companies and their enormous influence on many aspects of society, policymakers urgently need to step up and protect the independence of public interest research in this field, as has been demanded in an open letter to EU lawmakers, signed by a broad coalition of civil society organizations [494]. The letter, which was initiated by AlgorithmWatch [419], warns:

[L]arge platforms continue to suppress public interest research by scientists, civil society watchdogs, and journalists. (...) Facebook – as one of the largest platforms – has repeatedly restricted researchers' access to data, hindering public interest research not only in the United States but also in Europe[.] (...) Facebook clearly misuses its power to quash public interest research and therefore prevents an evidence-based debate about the impact platforms have on democratic processes and fundamental rights. (...) In your position as representatives of the European citizenry, we ask you to ensure that Terms of Service cannot be weaponized

against individuals or organizations that attempt to hold large platforms to account. Financial power must not be the currency that governs our public sphere. (...) Only if we understand how our public sphere is influenced by platforms' algorithmic choices can we take measures towards ensuring they do not undermine individuals' autonomy, freedom, and the collective good.

An expansion of public funding could be a crucial building block in safeguarding the independence of internet research, along with the constant and thorough scrutiny of purported guarantees of academic independence in both public and private funding. Furthermore, as proposed in Chapter 10.5, wide-ranging transparency obligations are needed, including laws that ensure access to platform data – not only for researchers but also for journalists and independent civil society organizations [494].

11

Conclusion

In this thesis, we have looked at various facets of modern-day privacy threats. We have explored the plethora of sensitive personal information that can unexpectedly be inferred from sensor data collected by everyday consumer devices. And we have learned that many people, including even ICT professionals, are largely unaware of these types of privacy threats. We have seen that many mobile apps do not comply with basic data protection rights prescribed by the GDPR, and have explored new ways to detect hidden web tracking and make it visible – or even audible. Furthermore, we have seen that current privacy laws based on the principle of privacy self-management are deeply dysfunctional for a multitude of reasons, and clearly not sufficient to regulate data processing in today’s complex socio-technical environment.

By answering their respective research questions, the studies included in this thesis provide numerous empirical and theoretical contributions to the scientific discourse. The findings highlight the severity and complexity of the privacy threats under investigation. At the same time, a general privacy-is-dead attitude is rejected. It is too early for complete resignation, and many things can still be changed and drastically improved – especially on the regulatory level. In the individual studies and in the closing discussion, I have addressed societal implications of our research findings and shared some ideas of how we could respond to the privacy challenges we are collectively facing.

11.1 Directions for Future Development and Research

In the course of the collaboration projects and studies incorporated in this thesis, various interesting research gaps were identified. Due to time constraints, some of these were left untouched. There remain many open questions and paths to be explored. This includes the development and evaluation of technical tools. In particular, I see the following exciting avenues for future development and research:

- **User studies on the privacy impacts of sensors.** Building upon findings from our literature review on the privacy implications of voice and speech analysis (Paper 1), we conducted a survey about users’ awareness and privacy concerns regarding personal

information inference from voice recordings (Paper 4). Similar studies could be conducted for other types of sensor data, such as for accelerometer data (based on the knowledge compiled in Paper 2) and eye-tracking data (based on Paper 3). While there have been previous studies in this field (cf. Section P4–2.3), many inference categories identified in our review papers and many sensors have not yet been covered. For scholars interested in this line of research, it could be helpful to consider the limitations and lessons learnt we have summarized in Section P4–7.

- **Longitudinal checks for data protection compliance.** In Paper 7, we used fake accounts over the course of several years to test repeatedly whether mobile app vendors complied with data subject rights prescribed by law. This novel approach has proven fruitful for this type of investigation and should be used more widely, also on other types of data controllers. The basic problems in this area are now well understood and documented [162, 163, 495]. However, further longitudinal research is needed to assess the impact of legal measures and monitor whether actual progress is being made, as well as to observe and help refine emerging best practices.
- **Monitoring the state of the art in inferential analytics.** As analytical capabilities are quickly advancing [496, 497, 498], a continuous monitoring of these capabilities with respect to their privacy implications is needed. As far as possible, privacy research should monitor the inferential power of state-of-the-art data mining and machine learning techniques. Here, both experimental research and consolidating secondary research remain important. At the same time, it needs to be acknowledged that – due to non-disclosure policies – it is impossible to exactly determine the limits of continuously advancing inference attacks based on publicly available information (cf. Sections P1–3 and P6–6).
- **Exploring solutions to plug the loophole of privacy self-management.** Since the notice-and-choice approach to data protection is deeply dysfunctional (as we have explained in Paper 10), new regulatory approaches to replace and/or complement the old paradigm need to be devised and tested in practice. As argued in Chapter 10.9, one promising legal direction could be risk-based restrictions of data collection and use. It seems clear, however, that there is no simple solution to this problem. Addressing the complexities involved will arguably require highly interdisciplinary research (incl. law, economics, computer science, sociology, and behavioral science) and a combination of different legal approaches, including strict prohibitions of ethically indefensible data processing (cf. Section P10–7), increased transparency obligations (cf. Chapter 10.5) and, where harms cannot be avoided, innovative ways of harm compensation. As mentioned and exemplified in Section P10–7, future research on this issue should closely examine the numerous potential obstacles that regulatory alternatives to privacy self-management will likely be faced with, and ways how these can be overcome.
- **Automatic web tracking detection.** Further research is needed to build effective classification models and identify suitable features and characteristics to accurately distinguish tracker from non-tracker web traffic [499, 500, 501]. As we have explained in Paper 8, real browsing data can – and should – be used to train algorithms for the

automatic detection of web-tracking activity, as it offers several important advantages over artificial data. Real browsing sessions can be recorded in a privacy-preserving manner, for example, by using our novel browser extension *T.EX – The Transparency EXtension* [327].

- **Web tracking sonification.** After the preliminary evaluation of our sonification approach yielded encouraging results in terms of increasing users’ web tracking awareness (Paper 9), we intended to conduct larger user studies to measure the effect in the field. However, due to other projects and time constraints, we were not able to pursue this further, thus leaving this undertaking for future research.
- **Development and testing of inference mapping tool.** For reasons outlined in Chapter 8.1, it would be desirable to present categories of personal information that can be inferred from certain types of data (e.g., certain sensor signals) in an interactive and updatable form. Suggestions were provided on how such an “Inference Mapping Tool” could look like in terms of basic functionalities (Chapter 8.2) and possible extensions (Chapter 8.2.1), followed by a discussion of potential challenges to be born in mind during implementation and testing (Chapter 8.3).
- **Impact of privacy education on technology use.** Findings from our within- and between-subject comparisons in Paper 4 suggest that education on the privacy-invading potential of data analytics may have an impact on people’s intention to use smart devices and services, such as voice-controlled virtual assistants. Further research is encouraged to examine motivational aspects, causes, and implications of these observations. It would be interesting to see whether changes in self-reported usage intention actually translate into shifts in consumption and device usage behavior.

References

- [1] Auxier, B. et al. *Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information*. 2019. URL: <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/>.
- [2] Rothchild, J. A. “Against Notice and Choice: The Manifest Failure of the Proceduralist Paradigm to Protect Privacy Online (or Anywhere Else)”. In: *Cleveland State Law Review* 66 (2018), pp. 559–648. URL: <https://heinonline.org/HOL/LandingPage?handle=hein.journals/clevslr66&div=22>.
- [3] Mendes, R. and Vilela, J. P. “Privacy-preserving data mining”. In: *IEEE Access: Practical Innovations, Open Solutions* 5 (2017), pp. 10562–10582. DOI: 10.1109/ACCESS.2017.2706947.
- [4] Kröger, J. “Unexpected Inferences from Sensor Data: A Hidden Privacy Threat in the Internet of Things”. In: *Internet of Things. Information Processing in an Increasingly Connected World*. Ed. by L. Strous and V. G. Cerf. Cham: Springer, 2019, pp. 147–159.
- [5] Christl, W. *Networks of Control*. Vienna: Cracked Labs, 2016. ISBN: 978-3-7089-1473-2. URL: https://crackedlabs.org/dl/Christl_Spiekermann_Networks_Of_Control.pdf.
- [6] Christl, W. *Corporate Surveillance in Everyday Life*. Vienna: Cracked Labs, 2017. URL: https://crackedlabs.org/dl/CrackedLabs_Christl_CorporateSurveillance.pdf.
- [7] Marwick, A. and Hargittai, E. “Nothing to hide, nothing to lose?” In: *Information, Communication & Society* 22.12 (2019). Publisher: Taylor & Francis, pp. 1697–1713. DOI: 10.1080/1369118X.2018.1450432.
- [8] Hargittai, E. and Marwick, A. ““What can I really do?” Explaining the privacy paradox with online apathy”. In: *International Journal of Communication* 10 (2016), pp. 3737–3757.
- [9] De Datta, R. “The truth about privacy”. In: *Forbes Magazine* (Oct. 2020). URL: <https://www.forbes.com/sites/forbestechcouncil/2020/10/21/the-truth-about-privacy/> (visited on 09/20/2021).
- [10] Burt, A. and Geer, D. “Opinion | The End of Privacy”. In: (2017). URL: <https://www.nytimes.com/2017/10/05/opinion/privacy-rights-security-breaches.html> (visited on 10/05/2021).
- [11] Morgan, J. “Privacy Is Completely And Utterly Dead, And We Killed It”. In: (2014). URL: <https://www.forbes.com/sites/jacobmorgan/2014/08/19/privacy-is-completely-and-utterly-dead-and-we-killed-it/> (visited on 10/05/2021).
- [12] Rubenfeld, J. “The End of Privacy”. In: *Stanford Law Review* 61 (2008), pp. 101–162. URL: <https://heinonline.org/HOL/LandingPage?handle=hein.journals/stflr61&div=6>.
- [13] Sykes, C. J. *The End of Privacy: The Attack on Personal Rights at Home, at Work, On-Line, and in Court*. St. Martin’s Publishing Group, Oct. 1999. ISBN: 978-0-312-26830-5.
- [14] Posey, J. *Studying Digital Marketing in a Post-Privacy Era*. University of Pennsylvania. 2017. URL: <https://penntoday.upenn.edu/spotlights/studying-digital-marketing-post-privacy-era> (visited on 10/05/2021).
- [15] Heller, C. *Post-Privacy: Prima leben ohne Privatsphäre*. German. CH Beck, 2011. ISBN: 978-3-406-62223-6.
- [16] Electronic Frontier Foundation. *About EFF*. URL: <https://www EFF.org/about> (visited on 10/05/2021).

REFERENCES

- [17] Privacy International. *About Us*. URL: <https://privacyinternational.org/about> (visited on 10/05/2021).
- [18] United States Pirate Party. *About*. URL: <https://uspirates.org/about/> (visited on 10/05/2021).
- [19] Christl, W. *How companies use data against people*. Vienna: Cracked Labs, 2017. URL: https://crackedlabs.org/dl/CrackedLabs_Christl_DataAgainstPeople.pdf.
- [20] Federal Bureau of Investigation. *2020 Internet Crime Report*. 2021. URL: https://www.ic3.gov/Media/PDF/AnnualReport/2020_IC3Report.pdf (visited on 03/09/2022).
- [21] Snowden, E. *Permanent Record*. New York: Picador Paper, 2020. ISBN: 978-1-250-77290-9.
- [22] Madge, R. *Five loopholes in the GDPR*. 2018. URL: <https://medium.com/mydata/five-loopholes-in-the-gdpr-367443c4248b> (visited on 10/05/2021).
- [23] Skiljic, A. *The status quo of health data inferences*. 2021. URL: <https://iapp.org/news/a/the-status-quo-of-health-data-inferences/> (visited on 02/18/2022).
- [24] Venkataramakrishnan, S. “GDPR accused of being toothless because of lack of resources”. In: *Financial Times* (Apr. 2020). URL: <https://www.ft.com/content/a915ae62-034e-4b13-b787-4b0ac2aaff7e> (visited on 11/16/2021).
- [25] Weizenbaum Institute for the Networked Society. *Das Institut*. German. URL: <https://www.weizenbaum-institut.de/das-institut/> (visited on 10/07/2021).
- [26] Weizenbaum Institute for the Networked Society. *Mission Statement*. URL: <https://www.weizenbaum-institut.de/en/the-institute/mission-statement/> (visited on 10/07/2021).
- [27] Heimlich, R. *Internet Users Don't like Targeted Ads*. 2012. URL: <https://www.pewresearch.org/fact-tank/2012/03/13/internet-users-dont-like-targeted-ads/> (visited on 10/05/2021).
- [28] Hitlin, P. and Rainie, L. *Facebook Algorithms and Personal Data*. Tech. rep. Washington, D.C.: Pew Research Center, 2019. URL: <https://www.pewresearch.org/internet/2019/01/16/facebook-algorithms-and-personal-data/>.
- [29] Rainie, L., Janna Anderson, and Vogels, E. A. *Worries about developments in AI*. 2021. URL: <https://www.pewresearch.org/internet/2021/06/16/1-worries-about-developments-in-ai/> (visited on 10/06/2021).
- [30] Mcdonald, A. M. and Cranor, L. F. “The Cost of Reading Privacy Policies”. In: *Journal of Law and Policy for the Information Society* (2008), pp. 543–568.
- [31] Solove, D. J. “Privacy Self-Management and the Consent Dilemma”. In: *Harvard Law Review* 126 (2012), p. 25. URL: https://scholarship.law.gwu.edu/cgi/viewcontent.cgi?article=2093&context=faculty_publications.
- [32] Gellman, R. “Fair Information Practices: A Basic History - Version 2.21”. In: *SSRN Electronic Journal* (2021). DOI: 10.2139/ssrn.2415020.
- [33] Safaei, B. et al. “Reliability side-effects in Internet of Things application layer protocols”. In: *International Conference on System Reliability and Safety (ICSRS)*. 2017, pp. 207–212. DOI: 10.1109/ICSRS.2017.8272822.
- [34] Amy Nordrum. *Popular Internet of Things Forecast of 50 Billion Devices by 2020 Is Outdated*. 2016. URL: <https://spectrum.ieee.org/popular-internet-of-things-forecast-of-50-billion-devices-by-2020-is-outdated> (visited on 10/06/2021).
- [35] Higginbotham, S. *A warning about sensors and surveillance*. 2021. URL: <https://staceyoniot.com/a-warning-about-sensors-and-surveillance/> (visited on 10/29/2021).
- [36] Sikder, A. K. et al. “A survey on sensor-based threats to internet-of-things (iot) devices and applications”. In: *Preprint arXiv:1802.02041* (2018). URL: <https://arxiv.org/abs/1802.02041>.
- [37] Spreitzer, R. et al. “Systematic classification of side-channel attacks: A case study for mobile devices”. In: *IEEE Communications Surveys & Tutorials* 20.1 (2017), pp. 465–488. DOI: 10.1109/COMST.2017.2779824.

- [38] Hnat, T. W. et al. “Doorjamb: unobtrusive room-level tracking of people in homes using doorway sensors”. In: *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*. New York, USA: Association for Computing Machinery, 2012, pp. 309–322. DOI: 10.1145/2426656.2426687.
- [39] Klasnja, P. et al. “Exploring Privacy Concerns about Personal Sensing”. In: *International Conference on Pervasive Computing*. Ed. by H. Tokuda et al. Springer, 2009, pp. 176–183. DOI: 10.1007/978-3-642-01516-8_13.
- [40] Bai, X., Yin, J., and Wang, Y.-P. “Sensor Guardian: prevent privacy inference on Android sensors”. In: *EURASIP Journal on Information Security* 2017.10 (June 2017). DOI: 10.1186/s13635-017-0061-8. (Visited on 10/06/2021).
- [41] Xu, Z. and Zhu, S. “SemaDroid: A Privacy-Aware Sensor Management Framework for Smartphones”. In: *Proceedings of the ACM Conference on Data and Application Security and Privacy*. 2015, pp. 61–72. DOI: 10.1145/2699026.2699114.
- [42] Notra, S. et al. “An experimental study of security and privacy risks with emerging household appliances”. In: *IEEE Conference on Communications and Network Security*. 2014, pp. 79–84. DOI: 10.1109/CNS.2014.6997469.
- [43] Capritto, A. *The complete guide to Apple’s Health app*. 2019. URL: <https://www.cnet.com/health/the-complete-guide-to-apples-health-app/> (visited on 11/03/2021).
- [44] Yang, L., Ting, K., and Srivastava, M. B. “Inferring occupancy from opportunistically available sensor data”. In: *IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 2014, pp. 60–68. DOI: 10.1109/PerCom.2014.6813945.
- [45] Bojinov, H. et al. “Mobile Device Identification via Sensor Fingerprinting”. In: *arXiv:1408.1416 [cs]* (2014). URL: <http://arxiv.org/abs/1408.1416> (visited on 10/06/2021).
- [46] Hua, J., Shen, Z., and Zhong, S. “We Can Track You if You Take the Metro: Tracking Metro Riders Using Accelerometers on Smartphones”. In: *IEEE Transactions on Information Forensics and Security* 12.2 (Feb. 2017), pp. 286–297. DOI: 10.1109/TIFS.2016.2611489.
- [47] Mehrnezhad, M. et al. “Stealing PINs via mobile sensors: Actual risk versus user perception”. In: *International Journal of Information Security* 17.3 (2018), pp. 291–313. DOI: 10.1007/s10207-017-0369-x.
- [48] Krumm, J. “Inference attacks on location tracks”. In: *International Conference on Pervasive Computing*. 2007, pp. 127–143. DOI: 10.1007/978-3-540-72037-9_8.
- [49] Gupta, P. and Dallas, T. “Feature Selection and Activity Recognition System Using a Single Triaxial Accelerometer”. In: *IEEE Transactions on Biomedical Engineering* 61.6 (June 2014), pp. 1780–1786. DOI: 10.1109/TBME.2014.2307069.
- [50] Weiss, G. M. and Lockhart, J. W. “Identifying user traits by mining smart phone accelerometer data”. In: *Proceedings of the International Workshop on Knowledge Discovery from Sensor Data*. 2011, pp. 61–69. DOI: 10.1145/2003653.2003660.
- [51] Jain, A. and Kanhangad, V. “Investigating gender recognition in smartphones using accelerometer and gyroscope sensor readings”. In: *International Conference on Computational Techniques in Information and Communication Technologies (ICCTICT)*. 2016, pp. 597–602. DOI: 10.1109/ICCTICT.2016.7514649.
- [52] Davarci, E. et al. “Age group detection using smartphone motion sensors”. In: *25th European Signal Processing Conference (EUSIPCO)*. Aug. 2017, pp. 2201–2205. DOI: 10.23919/EUSIPCO.2017.8081600.
- [53] Zhang, Z. et al. “Emotion recognition based on customized smart bracelet with built-in accelerometer”. In: *PeerJ* 4 (2016), e2258. DOI: 10.7717/peerj.2258. URL: <https://peerj.com/articles/2258>.
- [54] Kohnstamm, J. and Madhub, D. “Mauritius Declaration on the Internet of Things”. In: 2014. URL: https://edps.europa.eu/sites/edp/files/publication/14-10-14_mauritius_declaration_en.pdf.

REFERENCES

- [55] Schneble, C. O., Elger, B. S., and Shaw, D. M. “All Our Data Will Be Health Data One Day: The Need for Universal Data Protection and Comprehensive Consent”. In: *Journal of Medical Internet Research* 22.5 (2020). Publisher: JMIR Publications, e16879. DOI: 10.2196/16879. URL: <https://www.jmir.org/2020/5/e16879>.
- [56] Blanke, J. M. “Protection for ‘Inferences drawn’: A comparison between the general data protection regulation and the california consumer privacy act”. In: *Global Privacy Law Review* 1.2 (2020), pp. 81–92. URL: <https://kluwerlawonline.com/journalarticle/Global+Privacy+Law+Review/1.2/GPLR2020080>.
- [57] Article 29 Working Party. *Opinion 03/2013 on Purpose Limitation*. Tech. rep. 2013. URL: https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf (visited on 10/07/2021).
- [58] McLaren, M. et al. “The 2016 Speakers in the Wild Speaker Recognition Evaluation”. In: *INTER-SPEECH Proceedings*. 2016, pp. 823–827. DOI: 10.21437/Interspeech.2016-1137.
- [59] Swain, M., Routray, A., and Kabisatpathy, P. “Databases, features and classifiers for speech emotion recognition: a review”. In: *International Journal of Speech Technology* 21.1 (2018), pp. 93–120. DOI: 10.1007/s10772-018-9491-z.
- [60] Cummins, N., Baird, A., and Schuller, B. W. “Speech analysis for health: Current state-of-the-art and the increasing impact of deep learning”. en. In: *Methods. Health Informatics and Translational Data Analytics* 151 (2018), pp. 41–54. DOI: 10.1016/j.ymeth.2018.07.007.
- [61] Bedi, G. et al. “Automated analysis of free speech predicts psychosis onset in high-risk youths”. In: *npj Schizophrenia* 1 (2015). Publisher: Nature Publishing Group Article number: 15030. DOI: 10.1038/npjrschz.2015.30.
- [62] Sadjadi, S. O., Ganapathy, S., and Pelecanos, J. W. “Speaker age estimation on conversational telephone speech using senone posterior based i-vectors”. In: 2016, pp. 5040–5044. DOI: 10.1109/ICASSP.2016.7472637.
- [63] Kabil, S. H., Muckenhirn, H., et al. “On learning to identify genders from raw speech signal using CNNs”. In: *INTER-SPEECH Proceedings*. 2018, pp. 287–291. DOI: 10.21437/Interspeech.2018-1240.
- [64] Behravan, H. et al. “i-Vector Modeling of Speech Attributes for Automatic Foreign Accent Recognition”. In: *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 24.1 (2016), pp. 29–41. DOI: 10.1109/TASLP.2015.2489558.
- [65] Polzehl, T. *Personality in Speech: Assessment and Automatic Classification*. Springer, 2015. ISBN: 978-3-319-09516-5. DOI: 10.1007/978-3-319-09516-5.
- [66] Arp, D. et al. “Privacy threats through ultrasonic side channels on mobile devices”. In: *IEEE european symposium on security and privacy (EuroS&P)*. 2017, pp. 35–47. DOI: 10.1109/EuroSP.2017.33.
- [67] Greveler, U. et al. “Multimedia content identification through smart meter power usage profiles”. In: *Proceedings of the international conference on information and knowledge engineering (IKE)*. Athens, 2012.
- [68] Han, J. et al. “Accomplice: Location inference using accelerometers on smartphones”. In: *Fourth International conference on Communication Systems and Networks (COMSNETS)*. 2012. DOI: 10.1109/COMSNETS.2012.6151305.
- [69] Grömer, K. *The Art of Prehistoric Textile Making: The development of craft traditions and clothing in Central Europe*. Naturhistorisches Museum Wien, 2016. ISBN: 978-3-902421-94-4. URL: https://doi.org/10.26530/oapen_604250.
- [70] Pohle, J. “Datenschutz und Technikgestaltung: Geschichte und Theorie des Datenschutzes aus informatischer Sicht und Folgerungen für die Technikgestaltung”. German. Doctoral thesis. Humboldt-Universität zu Berlin, 2018. URL: <https://doi.org/10.18452/19136>.
- [71] Warren, S. and Brandeis, L. “The Right to Privacy”. In: *Harvard Law Review* 4.5 (1890), pp. 193–220.
- [72] Benjamin, W. “A Short History of Photography”. In: *Screen* 13.1 (1972), pp. 5–26. DOI: 10.1093/screen/13.1.5.

-
- [73] Smith, H. J., Dinev, T., and Xu, H. "Information Privacy Research: An Interdisciplinary Review". In: *MIS Quarterly* 35.4 (2011). Publisher: Management Information Systems Research Center, University of Minnesota, pp. 989–1015. DOI: 10.2307/41409970.
 - [74] Dunn, E. S. "The Idea of a National Data Center and the Issue of Personal Privacy". In: *The American Statistician* 21.1 (1967). Publisher: Taylor & Francis, pp. 21–27. DOI: 10.1080/00031305.1967.10481787.
 - [75] Kraus, R. "Statistical Déjà Vu: The National Data Center Proposal of 1965 and Its Descendants". In: *Journal of Privacy and Confidentiality* 5.1 (2013). DOI: 10.29012/jpc.v5i1.624.
 - [76] Kamlah, R. *Right of privacy: Das allgemeine Persönlichkeitsrecht in amerikanischer Sicht, unter Berücksichtigung neuer technologischer Entwicklungen*. German. Erlanger Juristische Abhandlungen 4. Köln: Heymann, 1969.
 - [77] Westin, A. F. *Privacy and Freedom*. New York: Atheneum, 1967.
 - [78] Packard, V. *The Naked Society*. New York: David McKay Company, 1964.
 - [79] Brenton, M. *The Privacy Invaders*. New York: Coward-McCann, 1964.
 - [80] Cate, F. H. "The EU data protection directive, information privacy, and the public interest". In: *Iowa Law Review* 80 (1994). Publisher: HeinOnline, p. 431. URL: <https://heinonline.org/HOL/LandingPage?handle=hein.journals/ilr80&div=25>.
 - [81] Riccardi, J. L. "The German Federal Data Protection Act of 1977: Protecting the Right to Privacy?" In: *Boston College International and Comparative Law Review* 6.1 (1983). URL: <https://lawdigitalcommons.bc.edu/iclr/vol6/iss1/8>.
 - [82] Ware, W. et al. *Records, computers, and the rights of citizens*. Tech. rep. Washington, D.C.: U.S. Department of Health, Education & Welfare, 1973. URL: <https://aspe.hhs.gov/report/records-computers-and-rights-citizens>.
 - [83] Mulligan, D. K. "The Enduring Importance of Transparency". In: *IEEE Security Privacy* 12.3 (2014), pp. 61–65. DOI: 10.1109/MSP.2014.58.
 - [84] Beverage, J. "The Privacy Act of 1974: An Overview". In: *Duke Law Journal* 1976.2 (1976). Publisher: Duke University School of Law, pp. 301–329. DOI: 10.2307/1371980. URL: <https://www.jstor.org/stable/1371980>.
 - [85] Borgesius, F. Z., Gray, J., and Eechoud, M. van. "Open data, privacy, and fair information principles: Towards a balancing framework". In: *Berkeley Technology Law Journal* 30.3 (2015). Publisher: JSTOR, pp. 2073–2131. URL: <https://www.jstor.org/stable/26377585>.
 - [86] Privacy Protection Study Commission. *Personal privacy in an information society*. Washington, D.C.: U.S. Government Printing Office, 1977.
 - [87] German Constitutional Court. *Volkszählung, BVerfGE 65, 1*. German. 1983. URL: <https://servat.unibe.ch/tools/DfrInfo?Command=ShowPrintText&Name=bv065001>.
 - [88] Konrad-Adenauer-Stiftung. *English Translation of essential parts of the German "Volkszählungsurteil" from 15 December 1983*. 2013. URL: <https://freiheitsfoo.de/files/2013/10/Census-Act.pdf>.
 - [89] Hornung, G. and Schnabel, C. "Data protection in Germany I: The population census decision and the right to informational self-determination". In: *Computer Law & Security Review* 25.1 (2009), pp. 84–88. DOI: 10.1016/j.clsr.2008.11.002.
 - [90] Simitis, S. "Reviewing Privacy in an Information Society". In: *University of Pennsylvania Law Review* 135.3 (1987), p. 707. DOI: 10.2307/3312079.
 - [91] Regan, P. M. "From Paper Dossiers to Electronic Dossiers: Gaps in the Privacy Act of 1974". In: *Office Technology and People* 4.3 (1988), pp. 279–296. DOI: 10.1108/eb022661.
 - [92] Gandy Jr, O. H. *The panoptic sort: A political economy of personal information*. Boulder, CO, USA: Westview Press, 1993. ISBN: 0-8133-1657-X. URL: <https://eric.ed.gov/?id=ED377817>.

REFERENCES

- [93] Klosowski, T. *The State of Consumer Data Privacy Laws in the US (And Why It Matters)*. Sept. 2021. URL: <https://www.nytimes.com/wirecutter/blog/state-of-privacy-laws-in-us/> (visited on 01/26/2022).
- [94] Solove, D. J. *A Brief History of Information Privacy Law*. Tech. rep. 215. Washington, D.C.: George Washington University. URL: <http://ssrn.com/abstract=914271>.
- [95] Fuster, G. G. *The emergence of personal data protection as a fundamental right of the EU*. Vol. 16. Law, Governance and Technology Series. Springer Science & Business, 2014. ISBN: 3-319-05023-0.
- [96] Wuermeling, U. U. "Harmonization of European Union Privacy law". In: *John Marshall Journal of Computer and Information Law* 14.3 (1996), p. 411. URL: <https://www.jcil.org/journal/articles/308.html>.
- [97] Rustad, M. L. and Koenig, T. H. "Towards a Global Data Privacy Standard". In: *Florida Law Review* 71.2 (2019), p. 91. URL: <https://scholarship.law.ufl.edu/flr/vol71/iss2/3>.
- [98] Dowling Jr, D. C. *International data protection and privacy law*. Publisher: Practising Law Institute. 2009. URL: https://intellicentrics.ca/wp-content/uploads/dlm_uploads/2014/09/article_intldataprotectionandprivacylaw_v5-1.pdf (visited on 01/13/2022).
- [99] Botha, J. et al. "A High-Level Comparison between the South African Protection of Personal Information Act and International Data Protection Laws". In: 2017, pp. 57–66.
- [100] Murata, K., Adams, A. A., and Lara Palma, A. M. "Following Snowden: a cross-cultural study on the social impact of Snowden's revelations". In: *Journal of Information, Communication and Ethics in Society* 15.3 (2017), pp. 183–196. DOI: 10.1108/JICES-12-2016-0047.
- [101] Meister, A. *Wir veröffentlichen den Gesetzentwurf zur BND-Reform - Große Koalition will Geheimdienst-Überwachung legalisieren*. German. June 2016. URL: <https://netzpolitik.org/2016/wir-veroeffentlichen-den-gesetzentwurf-zur-bnd-reform-grosse-koalition-will-geheimdienst-ueberwachung-legalisieren/> (visited on 01/26/2022).
- [102] Zuboff, S. *The Age of Surveillance Capitalism - The Fight for a Human Future at the New Frontier of Power*. Profile Books, 2019. ISBN: 978-1-78125-685-5. URL: <https://profilebooks.com/work/the-age-of-surveillance-capitalism/>.
- [103] Lin, H. and Bergmann, N. W. "IoT privacy and security challenges for smart home environments". In: *Information* 7.3 (2016). Publisher: Multidisciplinary Digital Publishing Institute, p. 44. DOI: 10.3390/info7030044.
- [104] Raij, A. et al. "Privacy risks emerging from the adoption of innocuous wearable sensors in the mobile environment". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2011, pp. 11–20. DOI: 10.1145/1978942.1978945.
- [105] Bundesministerium des Innern und für Heimat. *Häufig nachgefragt: Datenschutz-Grundverordnung*. German. URL: http://www.bmi.bund.de/SharedDocs/faqs/DE/themen/it-digitalpolitik/datenschutz/datenschutzgrundvo-liste.html;jsessionid=60591ADF5C1123600E22CA2092E50701.2_cid373?nn=9393768 (visited on 01/13/2022).
- [106] Thacker, N. *GDPR – A legislative milestone for a digital age*. Tech. rep. Austin, Texas: Forcepoint, 2017. URL: https://www.forcepoint.com/sites/default/files/resources/files/guide_gdpr_en.pdf.
- [107] Surya, S. *GDPR : The game changer in data privacy*. June 2018. URL: <https://medium.com/@sundaramsidharth/gdpr-the-game-changer-in-data-privacy-4277d38c17a> (visited on 01/13/2022).
- [108] Joshi, M. "General data protection regulation will be a game-changer but businesses will have to reshape their strategy-Business News , Firstpost". In: *Firstpost* (Nov. 2017). URL: <https://www.firstpost.com/business/general-data-protection-regulation-will-be-a-game-changer-but-businesses-will-have-to-reshape-their-strategy-4220539.html> (visited on 01/13/2022).
- [109] Dubey, S. *GDPR - Understanding The Game Changer*. May 2018. URL: <https://www.moengage.com/blog/gdpr-understanding-the-game-changer/> (visited on 01/13/2022).

- [110] Kuner, C. "The European Commission's proposed data protection regulation: A copernican revolution in European data protection law". In: *Bloomberg BNA Privacy and Security Law Report* 6.2012 (2012). Publisher: The Bureau of National Affairs, pp. 1–15.
- [111] Brave Software. *Europe's governments are failing the GDPR*. Tech. rep. 2020. URL: <https://brave.com/static-assets/files/Brave-2020-DPA-Report.pdf> (visited on 01/13/2022).
- [112] Hildén, J. "The Politics of Datafication: The influence of lobbyists on the EU's data protection reform and its consequences for the legitimacy of the General Data Protection Regulation". ISBN: 9789515134103. PhD thesis. University of Helsinki, 2019. URL: <https://helda.helsinki.fi/handle/10138/305981>.
- [113] Wirth, J., Maier, C., and Laumer, S. "The influence of resignation on the privacy calculus in the context of social networking sites: an empirical analysis". In: *European Conference on Information Systems (ECIS) Proceedings*. 2018. URL: https://aisel.aisnet.org/ecis2018_rp/161.
- [114] Solove, D. J. "The Myth of the Privacy Paradox". In: *George Washington Law Review* 89.1 (2021). URL: <https://www.gwlr.org/wp-content/uploads/2021/01/89-Geo.-Wash.-L.-Rev.-1.pdf>.
- [115] Madden, M. and Rainie, L. *Americans' Views About Data Collection and Security*. 2015. URL: <https://www.pewresearch.org/internet/2015/05/20/americans-views-about-data-collection-and-security/> (visited on 01/26/2022).
- [116] Allen, A. *Unpopular Privacy: What Must We Hide?* Oxford University Press, Nov. 2011. ISBN: 978-0-19-992088-4.
- [117] Koops, B.-J. et al. "A Typology of Privacy". In: *University of Pennsylvania Journal of International Law* 38.2 (2017), pp. 483–575.
- [118] Mulligan, D. K., Koopman, C., and Doty, N. "Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374.2083 (2016). DOI: 10.1098/rsta.2016.0118.
- [119] Belanger, F. and Crossler, R. "Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems". In: *Management Information Systems Quarterly* 35.4 (2011), pp. 1017–1041. URL: <https://aisel.aisnet.org/misq/vol35/iss4/12>.
- [120] Hull, G. "Successful failure: What Foucault can teach us about privacy self-management in a world of Facebook and big data". In: *Ethics and Information Technology* 17.2 (2015), pp. 89–101. DOI: 10.1007/s10676-015-9363-z.
- [121] Mühlhoff, R. "Predictive privacy: towards an applied ethics of data analytics". In: *Ethics and Information Technology* 23 (2021). Publisher: Springer, pp. 675–690. DOI: 10.1007/s10676-021-09606-x.
- [122] Westin, A. F. "Science, privacy, and freedom: Issues and proposals for the 1970's. Part I – The current impact of surveillance on privacy". In: *Columbia Law Review* 66.6 (1966), pp. 1003–1050. DOI: 10.2307/1120997.
- [123] Raynes-Goldie, K. "Aliases, creeping, and wall cleaning: Understanding privacy in the age of Facebook". In: *First Monday* 15.1 (2010). DOI: 10.5210/fm.v15i1.2775.
- [124] Westin, A. F. "Social and Political Dimensions of Privacy". In: *Journal of Social Issues* 59.2 (2003), pp. 431–453. DOI: 10.1111/1540-4560.00072.
- [125] Schomakers, E.-M. et al. "Internet users' perceptions of information sensitivity—insights from Germany". In: *International Journal of Information Management* 46 (2019). Publisher: Elsevier, pp. 142–150. DOI: 10.1016/j.ijinfomgt.2018.11.018.
- [126] Zarsky, T. Z. "Incompatible: the GDPR in the age of big data". In: *Seton Hall Law Review* 47.4 (2017). Publisher: HeinOnline, p. 995.
- [127] Milne, G. R. et al. "Information Sensitivity Typology: Mapping the Degree and Type of Risk Consumers Perceive in Personal Data Sharing". In: *Journal of Consumer Affairs* 51.1 (2017), pp. 133–161. DOI: 10.1111/joca.12111.
- [128] Nickel, C., Wirtl, T., and Busch, C. "Authentication of smartphone users based on the way they walk using k-nn algorithm". In: *International Conference on Intelligent Information Hiding and Multimedia Signal Processing*. IEEE, 2012, pp. 16–20. DOI: 10.1109/IIH-MSP.2012.11.

REFERENCES

- [129] Primo, A. et al. "Context-aware active authentication using smartphone accelerometer measurements". In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 2014, pp. 98–105. DOI: 10.1109/CVPRW.2014.20.
- [130] Cornet, V. P. and Holden, R. J. "Systematic review of smartphone-based passive sensing for health and wellbeing". In: *Journal of Biomedical Informatics* 77 (2018), pp. 120–132. DOI: 10.1016/j.jbi.2017.12.008.
- [131] Saleheen, N. et al. "puffMarker: a multi-sensor approach for pinpointing the timing of first lapse in smoking cessation". In: *Proceedings of the ACM international joint conference on pervasive and ubiquitous computing*. 2015, pp. 999–1010. DOI: 10.1145/2750858.2806897.
- [132] Lau, J., Zimmerman, B., and Schaub, F. "Alexa, stop recording": Mismatches between smart speaker privacy controls and user needs". In: *Poster at the 14th Symposium on Usable Privacy and Security (SOUPS)*. 2018. URL: <https://www.usenix.org/sites/default/files/soups2018posters-lau.pdf>.
- [133] Easwara Moorthy, A. and Vu, K.-P. L. "Privacy concerns for use of voice activated personal assistant in the public space". In: *International Journal of Human-Computer Interaction* 31.4 (2015). Publisher: Taylor & Francis, pp. 307–335.
- [134] Lau, J., Zimmerman, B., and Schaub, F. "Alexa, are you listening? Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers". In: *Proceedings of the ACM on Human-Computer Interaction*. ACM New York, NY, USA, 2018, pp. 1–31. DOI: 10.1145/3274371.
- [135] Manikonda, L., Deotale, A., and Kambhampati, S. "What's up with privacy? User preferences and privacy concerns in intelligent personal assistants". In: *Proceedings of the 2018 AAAI/ACM conference on AI, ethics, and society*. 2018, pp. 229–235. DOI: 10.1145/3278721.3278773.
- [136] Jozani, M. et al. "Privacy concerns and benefits of engagement with social media-enabled apps: A privacy calculus perspective". In: *Computers in Human Behavior* 107 (2020), p. 106260. DOI: 10.1016/j.chb.2020.106260.
- [137] Kandeh, A. T., Botha, R. A., and Futchet, L. A. "Enforcement of the Protection of Personal Information (POPI) Act: Perspective of data management professionals". In: *South African Journal of Information Management* 20.1 (2018). Publisher: AOSIS Publishing. DOI: 10.4102/sajim.v20i1.917.
- [138] Buck, C. "App-Privacy As An Abstract Value – Approaching Contingent Valuation For Investigating The Willingness To Pay For App Privacy". In: *International Conference on Electronic Business (ICEB) Proceedings*. 2015. URL: <https://aisel.aisnet.org/iceb2015/64>.
- [139] Rubinstein, I. "Big Data: The End of Privacy or a New Beginning?" In: *International Data Privacy Law* 3.2 (2013). Publisher: Oxford University Press, pp. 74–87. DOI: 10.1093/idpl/ips036.
- [140] Obar, J. A. "Big Data and The Phantom Public". In: *Big Data & Society* 2.2 (2015). Publisher: SAGE Publications, pp. 1–16. DOI: 10.1177/2053951715608876.
- [141] Razon, A. K. "Towards Financial Inclusion Through Digital Financial Services: Examining the Impact of the 'Notice and Consent' Privacy Mechanism". In: *Case Western Reserve Journal of Law, Technology and the Internet* 11.2 (2020), pp. 50–83. URL: <https://scholarlycommons.law.case.edu/jolti/vol11/iss1/3/>.
- [142] Scherf, R. "Yes I agree*: Assessing the failure of privacy "Self-Management" and its regulatory reforms". In: *Public Policy & Governance Review* 6.2 (2015). Publisher: University of Toronto, pp. 37–54. URL: <https://munkschool.utoronto.ca/wp-content/uploads/2015/04/ppgr-6-2-final-website.pdf#page=37>.
- [143] Malhotra, N. K., Kim, S. S., and Agarwal, J. "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model". In: *Information Systems Research* 15.4 (2004), pp. 336–355. DOI: 10.1287/isre.1040.0032. URL: <https://pubsonline.informs.org/doi/abs/10.1287/isre.1040.0032>.
- [144] Bellman, S. et al. "International Differences in Information Privacy Concerns: A Global Survey of Consumers". In: *The Information Society* 20.5 (2004). Publisher: Routledge, pp. 313–324. DOI: 10.1080/01972240490507956.

- [145] Dinev, T. and Hart, P. "Internet privacy concerns and their antecedents - measurement validity and a regression model". In: *Behaviour & Information Technology* 23.6 (2004), pp. 413–422. DOI: 10.1080/01449290410001715723. (Visited on 01/17/2022).
- [146] Acquisti, A. and Gross, R. "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook". In: *International Workshop on Privacy Enhancing Technologies*. Ed. by G. Danezis and P. Golle. Lecture Notes in Computer Science. Springer, 2006, pp. 36–58. DOI: 10.1007/11957454_3.
- [147] Gross, R. and Acquisti, A. "Information revelation and privacy in online social networks". In: *Proceedings of the ACM Workshop on Privacy in the Electronic Society*. 2005, pp. 71–80. DOI: 10.1145/1102199.1102214.
- [148] Chhetri, C. and Motti, V. G. "Eliciting privacy concerns for smart home devices from a user centered perspective". In: *International Conference on Information*. Springer, 2019, pp. 91–101. DOI: 10.1007/978-3-030-15742-5_8.
- [149] Liao, Y. et al. "Understanding the role of privacy and trust in intelligent personal assistant adoption". In: *iConference: International Conference on Information*. Springer, 2019, pp. 102–113. DOI: 10.1007/978-3-030-15742-5_9.
- [150] Zeng, E., Mare, S., and Roesner, F. "End user security and privacy concerns with smart homes". In: *Symposium on Usable Privacy and Security*. 2017, pp. 65–80.
- [151] Sweeney, M. and Davis, E. "Alexa, Are You Listening? An Exploration of Smart Voice Assistant Use and Privacy in Libraries". In: *Information Technology and Libraries* 39.4 (Jan. 2021). DOI: 10.6017/ital.v39i4.12363.
- [152] Swanson, B. *The best solution to digital privacy challenges? More technology, not European-style regulation*. Sept. 2018. URL: <https://www.aei.org/technology-and-innovation/the-best-solution-to-digital-privacy-challenges-more-technology-not-european-style-regulation/> (visited on 01/19/2022).
- [153] Greenley-Giudici, A. *When Self-Regulation Works, Your Privacy Is In Good Hands | TrustArc*. July 2012. URL: <https://trustarc.com/blog/2012/07/27/when-self-regulation-works-your-privacy-is-in-good-hands/>, %20<https://trustarc.com/blog/2012/07/27/when-self-regulation-works-your-privacy-is-in-good-hands/> (visited on 01/27/2022).
- [154] Davenport, T. H. "Should the U.S. Adopt European-Style Data-Privacy Protections?" In: *Wall Street Journal* (Mar. 2013). ISSN: 0099-9660. URL: <https://online.wsj.com/article/SB10001424127887324338604578328393797127094.html> (visited on 01/27/2022).
- [155] Bowie, N. E. and Jamal, K. "Privacy Rights on the Internet: Self-Regulation or Government Regulation?" In: *Business Ethics Quarterly* 16.3 (2006). Publisher: Cambridge University Press, pp. 323–342. URL: <https://www.jstor.org/stable/3857919>.
- [156] Gellman, R. and Dixon, P. "Many failures: A brief history of privacy self-regulation in the united states". In: *World Privacy Forum*. Lake Oswego, USA, 2011, pp. 1–29. URL: <http://www.worldprivacyforum.org/www/wprivacyforum/pdf/WPFselfregulationhistory.pdf>.
- [157] Wong, J. C. "The Cambridge Analytica scandal changed the world – but it didn't change Facebook". In: *The Guardian* (Mar. 2019). ISSN: 0261-3077. URL: <https://www.theguardian.com/technology/2019/mar/17/the-cambridge-analytica-scandal-changed-the-world-but-it-didnt-change-facebook> (visited on 01/27/2022).
- [158] Hoofnagle, C. J. "Privacy self regulation: A decade of disappointment". In: *Consumer Protection in the Age of the 'Information Economy'*. Ed. by J. K. Winn. Routledge, 2016. ISBN: 1-317-16120-3.
- [159] Budnitz, M. E. "Privacy Protection for Consumer Transactions in Electronic Commerce: Why Self-Regulation Is Inadequate". In: *South Carolina Law Review* 49.4 (1998), pp. 847–886. URL: <https://heinonline.org/HOL/Page?handle=hein.journals/sclr49&id=863&div=&collection=>.
- [160] Culnan, M. J. "Protecting Privacy Online: Is Self-Regulation Working?" In: *Journal of Public Policy & Marketing* 19.1 (2000). Publisher: SAGE Publications, pp. 20–26. DOI: 10.1509/jppm.19.1.20.16944.

REFERENCES

- [161] Gellman, R. and Dixon, P. “Failures of Privacy Self-Regulation in the United States”. In: *Enforcing Privacy: Regulatory, Legal and Technological Approaches*. Ed. by D. Wright and P. De Hert. Cham: Springer, 2016, pp. 53–77. ISBN: 978-3-319-25047-2. URL: https://doi.org/10.1007/978-3-319-25047-2_3.
- [162] Herrmann, D. and Lindemann, J. “Obtaining personal data and asking for erasure: Do app vendors and website owners honour your privacy rights?” In: *GI Sicherheit 2016*. Bonn: Gesellschaft für Informatik e.V., Apr. 2016, pp. 149–160.
- [163] Urban, T. et al. “A study on subject data access in online advertising after the GDPR”. In: *Data Privacy Management, Cryptocurrencies and Blockchain Technology*. Ed. by C. Pérez-Solà et al. Springer, 2019, pp. 61–79. URL: https://doi.org/10.1007/978-3-030-31500-9_5.
- [164] Spiller, K. “Experiences of accessing CCTV data: The urban topologies of subject access requests”. In: *Urban Studies* 53.13 (2016). Publisher: Sage Publications, pp. 2885–2900. DOI: 10.1177/0042098015597640.
- [165] Norris, C. and L’Hoiry, X. “Exercising citizen rights under surveillance regimes in Europe—Meta-analysis of a ten country study”. In: *The Unaccountable State of Surveillance*. Ed. by C. Norris et al. Springer, 2017, pp. 405–455. URL: https://doi.org/10.1007/978-3-319-47573-8_14.
- [166] Association Française des Correspondants à la protection des Données à caractère Personnel. *Données personnelles - Index AFCDP 2020 du Droit d’accès*. French. 2020. URL: <https://afcdp.net/index-du-droit-d-acces/> (visited on 04/09/2020).
- [167] Ausloos, J. and Dewitte, P. “Shattering one-way mirrors. Data subject access rights in practice”. In: *International Data Privacy Law* 8.1 (2018), pp. 4–28. DOI: 10.1093/idpl/ipy001.
- [168] Wong, J. and Henderson, T. “The right to data portability in practice: exploring the implications of the technologically neutral GDPR”. In: *International Data Privacy Law* 9.3 (2019). Publisher: Oxford University Press, pp. 173–191. DOI: 10.1093/idpl/ipz008.
- [169] Zurawski, N. “Exercising access rights in Germany”. In: *The Unaccountable State of Surveillance*. Ed. by C. Norris et al. Springer, 2017, pp. 109–133. URL: https://doi.org/10.1007/978-3-319-47573-8_6.
- [170] L’Hoiry, X. and Norris, C. “Exercising Access Rights in the United Kingdom”. In: *The Unaccountable State of Surveillance*. Ed. by C. Norris et al. Springer, 2017, pp. 359–404. URL: https://doi.org/10.1007/978-3-319-47573-8_13.
- [171] Albrecht, J. “How the GDPR Will Change the World”. In: *European Data Protection Law Review* 2.3 (2016), pp. 287–289. DOI: 10.21552/EDPL/2016/3/4.
- [172] Tiku, N. “How Europe’s New Privacy Law Will Change the Web, and More”. In: *Wired* (2018). ISSN: 1059-1028. URL: <https://www.wired.com/story/europes-new-privacy-law-will-change-the-web-and-more/> (visited on 11/12/2021).
- [173] Van Blarckom, G., Borking, J. J., and Olk, J. E. *Handbook of privacy and privacy-enhancing technologies*. The Hague: College Bescherming Persoonsgegevens, 2003. ISBN: 90-74087-33-7.
- [174] Fischer-Hübner, S. and Berthold, S. “Privacy-Enhancing Technologies”. In: *Computer and Information Security Handbook*. Ed. by J. R. Vacca. Boston: Morgan Kaufmann, Jan. 2017, pp. 759–778. ISBN: 978-0-12-803843-7. URL: <https://doi.org/10.1016/B978-0-12-803843-7.00053-3>.
- [175] Goldberg, I., Wagner, D., and Brewer, E. “Privacy-enhancing technologies for the internet”. In: *Proceedings IEEE COMPCON 97. Digest of Papers*. 1997, pp. 103–109. DOI: 10.1109/COMPCON.1997.584680.
- [176] Mittos, A., Malin, B., and De Cristofaro, E. “Systematizing Genome Privacy Research: A Privacy-Enhancing Technologies Perspective”. In: *Proceedings on Privacy Enhancing Technologies (PoPETs)* 2019.1 (2019), pp. 87–107. DOI: 10.2478/popets-2019-0006.
- [177] Becher, S. et al. “Big Picture on Privacy Enhancing Technologies in e-Health: A Holistic Personal Privacy Workflow”. In: *Information* 11.7 (July 2020). Publisher: Multidisciplinary Digital Publishing Institute, p. 356. DOI: 10.3390/info11070356.

- [178] Curzon, J., Almeahadi, A., and El-Khatib, K. “A survey of privacy enhancing technologies for smart cities”. In: *Pervasive and Mobile Computing* 55 (2019), pp. 76–95. DOI: 10.1016/j.pmcj.2019.03.001.
- [179] Malina, L. et al. “A Privacy-Enhancing Framework for Internet of Things Services”. In: *Network and System Security*. Ed. by J. K. Liu and X. Huang. Cham: Springer, 2019, pp. 77–97. ISBN: 978-3-030-36937-8. DOI: 10.1007/978-3-030-36938-5_5.
- [180] European Union Agency for Network And Cybersecurity. *Privacy by design in big data: an overview of privacy enhancing technologies in the era of big data analytics*. Publications Office of the European Union, 2016. ISBN: 978-92-9204-160-1. URL: <https://doi.org/10.2824/641480>.
- [181] Bellotti, V. and Sellen, A. “Design for Privacy in Ubiquitous Computing Environments”. In: *Proceedings of the European Conference on Computer-Supported Cooperative Work*. Ed. by G. de Michelis, C. Simone, and K. Schmidt. Springer, 1993, pp. 77–92. URL: https://doi.org/10.1007/978-94-011-2094-4_6.
- [182] Liebling, D. J. and Preibusch, S. “Privacy considerations for a pervasive eye tracking world”. In: *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct publication*. 2014, pp. 1169–1177. DOI: 10.1145/2638728.2641688.
- [183] Matyunin, N. et al. “MagneticSpy: Exploiting Magnetometer in Mobile Devices for Website and Application Fingerprinting”. In: *Proceedings of the 18th ACM Workshop on Privacy in the Electronic Society*. 2019, pp. 135–149. DOI: 10.1145/3338498.3358650.
- [184] Vandrico Solutions. *List of Wearables*. 2021. URL: <https://vandrico.com/wearables/list> (visited on 11/01/2021).
- [185] Hoy, M. B. “Alexa, Siri, Cortana, and more: an introduction to voice assistants”. In: *Medical Reference Services Quarterly* 37.1 (2018). Publisher: Taylor & Francis, pp. 81–88. DOI: 10.1080/02763869.2018.1404391.
- [186] Niu, E. *Smart-Speaker Volumes Expected to Jump Next Year*. 2020. URL: <https://www.nasdaq.com/articles/smart-speaker-volumes-expected-to-jump-next-year-2020-10-23> (visited on 11/01/2021).
- [187] National Public Radio. *The Smart Audio Report*. Tech. rep. Washington, D.C., 2020. URL: https://www.nationalpublicmedia.com/uploads/2020/04/The-Smart-Audio-Report_Spring-2020.pdf (visited on 11/01/2021).
- [188] OMD. *OMD Germany launch 'The Age of Voice 3.0' report*. 2021. URL: <https://www.ond.com/news/what-potential-does-voice-have-in-brand-management/> (visited on 11/01/2021).
- [189] Zweites Deutsches Fernsehen. *Postbank-Studie: Sprachassistenten sehr beliebt*. German. 2019. URL: <https://www.zdf.de/uri/c3b7d25c-d255-47f9-b067-ec8aa41edee2> (visited on 11/01/2021).
- [190] Tobii Tech. *Markets - Eye Tracking Opportunities*. URL: <https://tech.tobii.com/markets/> (visited on 11/01/2021).
- [191] Pupil Labs. *VR & AR - Eye tracking*. 2021. URL: <https://pupil-labs.com/products/vr-ar/> (visited on 11/01/2021).
- [192] Tobii Tech. *Why eye tracking will be standard on all VR headsets*. 2020. URL: <https://blog.tobii.com/why-eye-tracking-will-be-standard-on-all-vr-headsets> (visited on 11/01/2021).
- [193] Zheng, S. et al. “User perceptions of smart home IoT privacy”. In: *Proceedings of the ACM on Human-Computer Interaction* (2018), pp. 1–20. DOI: 10.1145/3274469.
- [194] Kröger, J. L. et al. “Personal Information Inference from Voice Recordings: User Awareness and Privacy Concerns”. In: *Proceedings on Privacy Enhancing Technologies (PoPETs)* 2022.1 (2022), pp. 6–27. DOI: 10.2478/popets-2022-0002.
- [195] Crager, K. et al. “Information Leakage through Mobile Motion Sensors: User Awareness and Concerns”. In: *Proceedings of the European Workshop on Usable Security (EuroUSEC)*. Internet Society, 2017. DOI: 10.14722/eurosec.2017.23013.
- [196] Jay, S. *The Privacy-Invasive Potential of Eye Tracking Technology*. 2013. URL: <https://www.aclu.org/blog/national-security/privacy-and-surveillance/privacy-invading-potential-eye-tracking-technology> (visited on 11/01/2021).

REFERENCES

- [197] Weiss, G. M. et al. “Actitracker: a smartphone-based activity recognition system for improving health and well-being”. In: *IEEE international conference on data science and advanced analytics (DSAA)*. 2016, pp. 682–688. DOI: 10.1109/DSAA.2016.89.
- [198] Wachter, S. and Mittelstadt, B. “A right to reasonable inferences: Re-thinking data protection law in the age of big data and AI”. In: *Columbia Business Law Review* (2019). Publisher: HeinOnline, pp. 494–620. URL: <https://heinonline.org/HOL/LandingPage?handle=hein.journals/colb2019&div=15>.
- [199] Aloufi, R., Haddadi, H., and Boyle, D. “Emotionless: privacy-preserving speech analysis for voice assistants”. In: *eprint arXiv:1908.03632* (2019). URL: <https://arxiv.org/abs/1908.03632>.
- [200] Aloufi, R., Haddadi, H., and Boyle, D. “Privacy-preserving voice analysis via disentangled representations”. In: *ACM SIGSAC Conference on Cloud Computing Security Workshop*. 2020, pp. 1–14.
- [201] Steil, J. et al. “Privacy-aware eye tracking using differential privacy”. In: *Proceedings of the ACM symposium on eye tracking research & applications*. 2019. DOI: 10.1145/3314111.3319915.
- [202] Nautsch, A. et al. “The privacy ZEBRA: Zero evidence biometric recognition assessment”. In: *Proc. Interspeech*. 2020, pp. 1698–1702. DOI: 10.21437/Interspeech.2020-1815.
- [203] Fischer, C. “The legal protection against inferences drawn by AI under the GDPR”. LL.M. thesis (Law and Technology). Law School, Tilburg University, 2020. URL: <http://arno.uvt.nl/show.cgi?fid=151926>.
- [204] Ufert, F. “AI Regulation Through the Lens of Fundamental Rights: How Well Does the GDPR Address the Challenges Posed by AI?” In: *European Papers* 5.2 (2020), pp. 1087–1097. DOI: 10.15166/2499-8249/394. URL: <https://www.europeanpapers.eu/it/europeanforum/ai-regulation-through-the-lens-of-fundamental-rights>.
- [205] Heaven, D. *Why deep-learning AIs are so easy to fool*. 2019. URL: <https://nature.com/articles/d41586-019-03013-5> (visited on 09/20/2021).
- [206] Watson, T. *Top artificial intelligence fails in image and facial recognition*. 2019. URL: <https://medium.com/swlh/top-artificial-intelligence-fails-in-image-and-facial-recognition-dfc1527b2295> (visited on 09/20/2021).
- [207] Alghowinem, S. et al. “Eye movement analysis for depression detection”. In: *IEEE International Conference on Image Processing*. 2013, pp. 4220–4224.
- [208] Benson, P. J. et al. “Simple viewing tests can detect eye movement abnormalities that distinguish schizophrenia cases from controls with exceptional accuracy”. In: *Biological Psychiatry* 72.9 (2012). Publisher: Elsevier, pp. 716–724. DOI: 10.1016/j.biopsych.2012.04.019.
- [209] Ooi, K. E. B., Lech, M., and Allen, N. B. “Multichannel weighted speech classification system for prediction of major depression in adolescents”. In: *IEEE Transactions on Biomedical Engineering* 60.2 (2012), pp. 497–506. DOI: 10.1109/TBME.2012.2228646.
- [210] Schuller, B. et al. “Medium-term speaker states - A review on intoxication, sleepiness and the first challenge”. In: *Computer Speech & Language* 28.2 (2014). Publisher: Elsevier, pp. 346–374. DOI: 10.1016/j.cs1.2012.12.002.
- [211] Mecacci, G. and Haselager, P. “Identifying criteria for the evaluation of the implications of brain reading for mental privacy”. In: *Science and Engineering Ethics* 25.2 (2019). Publisher: Springer, pp. 443–461. DOI: 10.1007/s11948-017-0003-3.
- [212] Bergstrom, C. T. and West, J. D. *Calling bullshit: the art of skepticism in a data-driven world*. New York: Random House, 2020. ISBN: 978-0-525-50918-9.
- [213] Calling Bullshit. *Calling Bullshit - Data Reasoning in a Digital World*. URL: <https://www.callingbullshit.org/> (visited on 02/14/2022).
- [214] Solon, O. “Artificial intelligence is ripe for abuse, tech researcher warns: ‘a fascist’s dream’”. In: *The Guardian* (Mar. 2017). ISSN: 0261-3077. URL: <https://www.theguardian.com/technology/2017/mar/13/artificial-intelligence-ai-abuses-fascism-donald-trump>.

- [215] O’Neil, C. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown, Sept. 2016. ISBN: 978-0-553-41881-1.
- [216] Richardson, R., Schultz, J. M., and Crawford, K. “Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice”. In: *New York University Law Review* 94 (2019). Publisher: HeinOnline. URL: <https://heinonline.org/HOL/LandingPage?handle=hein.journals/nyulro94&div=3>.
- [217] Wu, X. and Zhang, X. “Responses to Critiques on Machine Learning of Criminality Perceptions (Addendum of arXiv:1611.04135)”. In: *arXiv:1611.04135 [cs]* (May 2017). URL: <http://arxiv.org/abs/1611.04135> (visited on 02/14/2022).
- [218] Hashemi, M. and Hall, M. “RETRACTED ARTICLE: Criminal tendency detection from facial images and the gender bias effect”. In: *Journal of Big Data* 7.1 (2020), p. 2. DOI: 10.1186/s40537-019-0282-4.
- [219] Coalition for Critical Technology. *Abolish the #TechToPrisonPipeline*. Sept. 2021. URL: <https://medium.com/@CoalitionForCriticalTechnology/abolish-the-techtoprisonpipeline-9b5b14366b16> (visited on 02/14/2022).
- [220] Drachen, A. and Thureau, C. “A Comparison of Methods for Player Clustering via Behavioral Telemetry”. In: *Proceedings of the 8th International Conference on the Foundations of Digital Games*. Chania, Crete: Society for the Advancement of the Science of Digital Games, 2013.
- [221] Sifa, R., Drachen, A., and Bauckhage, C. “Profiling in Games: Understanding Behavior from Telemetry”. In: *Social Interactions in Virtual Worlds*. Ed. by K. Lakkaraju, G. Sukthankar, and R. T. Wigand. Cambridge University Press, 2018, pp. 337–374. URL: <https://doi.org/10.1017/9781316422823.014>.
- [222] Kröger, J. L., Lutz, O. H.-M., and Raschke, P. “Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference”. In: *Privacy and Identity Management. Data for Better Living: AI and Privacy*. Ed. by M. Friedewald et al. Cham: Springer, Mar. 2020, pp. 242–258. URL: https://doi.org/10.1007/978-3-030-42504-3_16.
- [223] Owusu, E. et al. “Accessory: password inference using accelerometers on smartphones”. In: *Proceedings of the 12th workshop on mobile computing systems & applications*. 2012, pp. 1–6. DOI: 10.1145/2162081.2162095.
- [224] Lohr, S. *Sizing up big data, broadening beyond the internet*. 2013. URL: <https://bits.blogs.nytimes.com/2013/06/19/sizing-up-big-data-broadening-beyond-the-internet/> (visited on 09/20/2021).
- [225] Osoba, O. A. and Welser IV, W. *An intelligence in our image: The risks of bias and errors in artificial intelligence*. Santa Monica, USA: RAND Corporation, 2017.
- [226] Wills, C. E. and Zeljkovic, M. “A personalized approach to web privacy: awareness, attitudes and actions”. In: *Information Management & Computer Security* 19.1 (2011), pp. 53–73. DOI: 10.1108/09685221111115863.
- [227] Weinshel, B. et al. “Oh, the Places You’ve Been! User Reactions to Longitudinal Transparency About Third-Party Web Tracking and Inferencing”. In: *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2019, pp. 149–166. DOI: 10.1145/3319535.3363200.
- [228] Rader, E., Hautea, S., and Munasinghe, A. ““I Have a Narrow Thought Process”: Constraints on Explanations Connecting Inferences and Self-Perceptions”. In: 2020, pp. 457–488. ISBN: 978-1-939133-16-8. URL: <https://www.usenix.org/conference/soups2020/presentation/rader>.
- [229] Schaub, F. et al. “Watching Them Watching Me: Browser Extensions Impact on User Privacy Awareness and Concern”. In: *Workshop on Usable Security*. San Diego, CA: Internet Society, 2016. DOI: 10.14722/usec.2016.23017.
- [230] Petkos, G., Papadopoulos, S., and Kompatsiaris, Y. “PScore: a framework for enhancing privacy awareness in online social networks”. In: *International Conference on Availability, Reliability and Security*. IEEE, 2015, pp. 592–600. DOI: 10.1109/ARES.2015.80.
- [231] Choe, E. K. et al. “Investigating receptiveness to sensing and inference in the home using sensor proxies”. In: *Proceedings of the ACM Conference on Ubiquitous Computing*. 2012, pp. 61–70. DOI: 10.1145/2370216.2370226.

REFERENCES

- [232] Nichols, S. *Your Phone Is Listening and it's Not Paranoia*. June 2018. URL: <https://www.vice.com/en/article/wjbzzy/your-phone-is-listening-and-its-not-paranoia> (visited on 02/12/2022).
- [233] Kleinman, Z. "Is your smartphone listening to you?" In: *BBC News* (Mar. 2016). URL: <https://www.bbc.com/news/technology-35639549> (visited on 02/12/2022).
- [234] Bond, C. *Is Your Phone Recording Your Conversations? The Answer Might Surprise You*. 2021. URL: https://www.huffpost.com/entry/smartphone-devices-recording-privacy_1_5d570869e4b056fafd0bbc09 (visited on 02/12/2022).
- [235] Pettijohn, N. "Of Course Your Phone Is Listening To You". In: *Forbes* (2019). URL: <https://www.forbes.com/sites/nathanpettijohn/2019/09/03/of-course-your-phone-is-listening-to-you/> (visited on 02/12/2022).
- [236] Fearnow, B. "Watch: Viral Video 'Proves' Facebook Listening To Your Real-Life Conversations". In: *International Business Times* (Oct. 2017). URL: <https://www.ibtimes.com/watch-viral-video-proves-facebook-listening-your-real-life-conversations-2608077> (visited on 02/12/2022).
- [237] the_ticklemONSTER (Reddit user). *Phone listening in to my conversations*. Reddit Post. Dec. 2017. URL: www.reddit.com/r/GalaxyS8/comments/7mb72k/phone_listening_in_to_my_conversations/ (visited on 02/12/2022).
- [238] Comrade Cat (Twitter user). *Duuuuuuude over the last couple months my phone has done so much shady ... [Tweet]*. Oct. 2021. URL: <https://twitter.com/CatJova/status/1454361194987327489> (visited on 02/12/2022).
- [239] Ceelz, W. *Hey guys are the phones listening to us forreal because I asked a ... [Tweet]*. Feb. 2022. URL: <https://twitter.com/NavyCeelz/status/1489427715254112257> (visited on 02/12/2022).
- [240] Wikipedia. *Telephone tapping*. URL: https://en.wikipedia.org/wiki/Telephone_tapping (visited on 09/20/2021).
- [241] Nolsoe, E. *Is my phone listening to my conversations? Britons believe the answer is yes*. 2021. URL: <https://yougov.co.uk/topics/technology/articles-reports/2021/07/16/my-phone-listening-my-conversations-britons-believ> (visited on 02/12/2022).
- [242] Koetsier, J. *55% Of Americans Say Smartphones Spy On Conversations To Customize Ads*. 2019. URL: <https://www.forbes.com/sites/johnkoetsier/2019/05/31/55-of-americans-say-smartphones-spy-on-conversations-to-customize-ads/> (visited on 02/12/2022).
- [243] Nationwide News. "Is my smartphone spying on me?" In: (July 2018). URL: <https://www.news.com.au/technology/gadgets/mobile-phones/one-in-five-people-think-their-smartphone-is-eavesdropping-on-them/news-story/313c66fa2df5e4f08549c083dcf8d3b3> (visited on 02/12/2022).
- [244] Kelly, M. *Your phone isn't listening to you, researchers say, but it may be watching everything you do*. July 2018. URL: <https://www.theverge.com/2018/7/3/17531698/conspiracy-theory-facebook-android-phone-listening> (visited on 02/12/2022).
- [245] Johnson, E. *Your phone is not secretly spying on your conversations. It doesn't need to*. July 2018. URL: <https://www.vox.com/2018/7/20/17594074/phone-spying-wiretap-microphone-smartphone-northeastern-dave-choffnes-christo-wilson-kara-swisher> (visited on 02/12/2022).
- [246] Tidy, J. "Why phones that secretly listen to us are a myth". In: *BBC News* (Sept. 2019). URL: <https://www.bbc.com/news/technology-49585682> (visited on 02/12/2022).
- [247] Timberg, C., Dwoskin, E., and Nakashima, E. "WikiLeaks: The CIA is using popular TVs, smartphones and cars to spy on their owners". In: *Washington Post* (Mar. 2017). ISSN: 0190-8286. URL: <https://www.washingtonpost.com/news/the-switch/wp/2017/03/07/why-the-cia-is-using-your-tvs-smartphones-and-cars-for-spying/> (visited on 11/03/2021).
- [248] Taylor, P. "Edward Snowden interview: 'Smartphones can be taken over'". In: *BBC News* (Oct. 2015). URL: <https://www.bbc.com/news/uk-34444233> (visited on 11/03/2021).

- [249] Rosenbach, M., Poitras, L., and Stark, H. “How the NSA Spies on Smartphones Including the BlackBerry”. In: *Der Spiegel* (Sept. 2013). ISSN: 2195-1349. URL: <https://www.spiegel.de/international/world/how-the-nsa-spies-on-smartphones-including-the-blackberry-a-921161.html> (visited on 11/03/2021).
- [250] Naughton, J. “The WhatsApp spyware story tells us that nothing is secure”. In: *The Guardian* (2013). URL: <https://www.theguardian.com/commentisfree/2019/may/19/whatsapp-spyware-story-tells-us-nothing-is-secure-pegasus-nso> (visited on 09/20/2021).
- [251] BuzzFeed Daily. *No, your phone can't hear you — But here's how it gets your data anyway*. 2021. URL: <https://www.buzzfeed.com/daily/phone-cant-hear-you-still-gets-data> (visited on 09/20/2021).
- [252] Martínez, A. G. “Facebook’s not listening through your phone. It doesn’t have to”. In: *The Guardian* (2017). URL: <https://www.theguardian.com/commentisfree/2019/may/19/whatsapp-spyware-story-tells-us-nothing-is-secure-pegasus-nso>.
- [253] Joe Tidy. *Why phones that secretly listen to us are a myth*. 2019. URL: <https://www.bbc.com/news/technology-49585682> (visited on 09/20/2021).
- [254] Purtill, C. *Your phone isn't really spying on your conversations - the truth might be even creepier*. 2019. URL: <https://qz.com/1609356/your-phone-is-not-recording-your-conversations/> (visited on 09/20/2021).
- [255] Hirsch, D. D. “The glass house effect: Big Data, the new oil, and the power of analogy”. In: *Maine Law Review* 66.2 (2014). Publisher: HeinOnline, pp. 373–395.
- [256] Lange, C. *Why is Personal Data the Oil of the 21st Century*. Mar. 2018. URL: <https://medium.com/pdata-token/why-is-personal-data-the-oil-of-the-21st-century-819a65e769a0> (visited on 09/20/2021).
- [257] García-Gasco Romero, M. “Personal Data: The new Black Gold”. In: *Security in the Global Commons and Beyond*. Ed. by J. M. Ramírez and B. Bauzá-Abril. Springer, 2021, pp. 171–182. ISBN: 978-3-030-67973-6.
- [258] The Economist. *The world's most valuable resource is no longer oil, but data*. May 2017. URL: <https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data> (visited on 09/20/2021).
- [259] Privacy International. *Why we've filed complaints against companies that most people have never heard of – and what needs to happen next*. 2018. URL: <http://privacyinternational.org/advocacy/2434/why-weve-filed-complaints-against-companies-most-people-have-never-heard-and-what> (visited on 09/24/2021).
- [260] Maheshwari, S. “That Game on Your Phone May Be Tracking What You’re Watching on TV”. In: *The New York Times* (Dec. 2017). URL: <https://www.nytimes.com/2017/12/28/business/media/alphonso-app-tracking.html> (visited on 09/21/2021).
- [261] Harkin, D., Molnar, A., and Vowles, E. “The commodification of mobile phone surveillance: An analysis of the consumer spyware industry”. In: *Crime, Media, Culture* 16.1 (2020). Publisher: SAGE Publications, pp. 33–60. DOI: 10.1177/1741659018820562.
- [262] Pan, E. et al. “Panoptispy: Characterizing audio and video exfiltration from android applications.” In: *Proceedings on Privacy Enhancing Technologies (PoPETs)* 2018.4 (2018), pp. 33–50. DOI: 10.1515/popets-2018-0030.
- [263] He, Y. et al. “Dynamic privacy leakage analysis of Android third-party libraries”. In: *Journal of Information Security and Applications* 46 (2019). Publisher: Elsevier, pp. 259–270. DOI: 10.1016/j.jisa.2019.03.014.
- [264] Zhang, L. et al. “Accelword: Energy efficient hotword detection through accelerometer”. In: *Proceedings of the Annual International Conference on Mobile Systems, Applications, and Services*. 2015, pp. 301–315. DOI: 10.1145/2742647.2742658.

REFERENCES

- [265] Stolyar, B. “A Google Fit update turns Pixel phone cameras into health trackers”. In: *Mashable* (Feb. 2021). URL: <https://mashable.com/article/google-fit-pixel-update-heart-rate-respiratory-rate> (visited on 11/03/2021).
- [266] Das, A., Borisov, N., and Caesar, M. “Tracking mobile web users through motion sensors: Attacks and defenses”. In: *Network and Distributed System Security (NDSS) Symposium*. 2016. DOI: 10.14722/ndss.2016.23390.
- [267] Altmetric. *Attention for chapter 6: Is my phone listening in? On the feasibility and detectability of mobile eavesdropping*. URL: <https://www.altmetric.com/details/63157232/chapter/63158112> (visited on 09/20/2021).
- [268] PlumX Metrics. *Metrics Details - Is my phone listening in? On the feasibility and detectability of mobile eavesdropping*. URL: https://plu.mx/plum/a/?doi=10.1007/978-3-030-22479-0_6 (visited on 09/20/2021).
- [269] Christl, W. *Are companies secretly eavesdropping on smartphone users to inform ads? ... [Tweet]*. Oct. 2019. URL: <https://web.archive.org/web/20220315110721/https://twitter.com/WolfieChristl/status/1182631705301131269> (visited on 02/12/2022).
- [270] PULS Reportage. *Können iPhone und Android Handys heimlich mithören? Wir programmieren Apps, um das zu beweisen!* Feb. 2022. URL: <https://www.youtube.com/watch?v=rX2tK-qSVpk> (visited on 02/10/2022).
- [271] Ciesielski, R. *Is your phone listening to your conversations?* Feb. 2022. URL: <https://medium.com/br-next/is-your-phone-listening-to-your-conversations-5182bc8ed45> (visited on 02/14/2022).
- [272] Anderson, M. *Mobile Technology and Home Broadband 2019*. 2019. URL: <https://www.pewresearch.org/internet/2019/06/13/mobile-technology-and-home-broadband-2019/> (visited on 02/15/2022).
- [273] Takahashi, D. *App Annie: Smartphone users spent 3.8 trillion hours on mobile in 2021*. Jan. 2022. URL: <https://venturebeat.com/2022/01/12/app-annie-smartphone-users-spent-3-8-trillion-hours-on-mobile-in-2021/> (visited on 02/15/2022).
- [274] Chan, S. *U.S. Consumers Used an Average of 46 Apps Each Month in the First Half of 2021*. 2021. URL: <https://sensortower.com/blog/apps-used-per-us-smartphone> (visited on 02/15/2022).
- [275] Schelter, S. and Kunegis, J. “On the Ubiquity of Web Tracking: Insights from a Billion-Page Web Crawl”. In: *The Journal of Web Science* 4 (2018). DOI: 10.1561/106.00000014.
- [276] Bujlow, T. et al. “A Survey on Web Tracking: Mechanisms, Implications, and Defenses”. In: *Proceedings of the IEEE* 105.8 (2017), pp. 1476–1510. DOI: 10.1109/JPROC.2016.2637878. URL: <http://ieeexplore.ieee.org/document/7872467/> (visited on 11/04/2021).
- [277] Andreou, A. et al. “Investigating Ad Transparency Mechanisms in Social Media: A Case Study of Facebook’s Explanations”. In: *Proceedings 2018 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2018. DOI: 10.14722/ndss.2018.23191. URL: https://www.ndss-symposium.org/wp-content/uploads/2018/02/ndss2018_10-1_Andreou_paper.pdf (visited on 11/04/2021).
- [278] Mayer, J. R. and Mitchell, J. C. “Third-party web tracking: Policy and technology”. In: *IEEE Symposium on Security and Privacy*. tex.organization: IEEE. 2012, pp. 413–427. DOI: 10.1109/SP.2012.47.
- [279] Binns, R. et al. “Measuring Third-party Tracker Power across Web and Mobile”. In: *ACM Transactions on Internet Technology* 18.4 (Aug. 2018), 52:1–52:22. DOI: 10.1145/3176246.
- [280] Acar, G. et al. *Facebook Tracking Through Social Plug-ins*. Technical report prepared for the Belgian Privacy Commission. Katholieke Universiteit Leuven, 2015.
- [281] Deußer, C., Passmann, S., and Strufe, T. “Browsing Unicity: On the Limits of Anonymizing Web Tracking Data”. In: *2020 IEEE Symposium on Security and Privacy (SP)*. 2020, pp. 777–790. DOI: 10.1109/SP40000.2020.00018.

- [282] Sanchez-Rola, I. et al. "The web is watching you: A comprehensive review of web-tracking techniques and countermeasures". In: *Logic Journal of the IGPL* 25.1 (2017). Publisher: Oxford University Press, pp. 18–29. DOI: 10.1093/jigpal/jzw041.
- [283] Roser, M., Ritchie, H., and Ortiz-Ospina, E. *Internet*. July 2015. URL: <https://ourworldindata.org/internet> (visited on 02/15/2022).
- [284] Larsson, S. and Jensen-Urstad, A. "Notified But Unaware: Third-Party Tracking Online". In: *Critical Analysis of Law* 8.1 (2021), pp. 101–120. URL: <https://cal.library.utoronto.ca/index.php/cal/article/view/36282>.
- [285] Takano, Y. et al. "Mindyourprivacy: Design and implementation of a visualization system for third-party web tracking". In: *Annual International Conference on Privacy, Security and Trust*. tex.organization: IEEE. 2014, pp. 48–56. DOI: 10.1109/PST.2014.6890923.
- [286] Lehmann, J. and Seufert, T. "The Interaction Between Text Modality and the Learner's Modality Preference Influences Comprehension and Cognitive Load". In: *Frontiers in Psychology* 10.2820 (2020). DOI: 10.3389/fpsyg.2019.02820.
- [287] Hutchins, C. et al. "Soundbeam: A platform for sonifying web tracking". In: *Proceedings of the International Conference on New Interfaces for Musical Expression*. 2014, pp. 497–498.
- [288] Namin, A. S. et al. "Sonifying internet security threats". In: *Proceedings of the CHI conference extended abstracts on human factors in computing systems*. 2016, pp. 2306–2313. DOI: 10.1145/2851581.2892363.
- [289] Lutz, O. H.-M. et al. "That password doesn't sound right: interactive password strength sonification". In: *Proceedings of the 15th International Conference on Audio Mostly*. 2020, pp. 206–213. DOI: 10.1145/3411109.3412299.
- [290] Ding, W. et al. "A survey on data fusion in internet of things: Towards secure and privacy-preserving fusion". In: *Information Fusion* 51 (Nov. 2019), pp. 129–144. DOI: 10.1016/j.inffus.2018.12.001. URL: <https://www.sciencedirect.com/science/article/pii/S1566253518304731>.
- [291] Javed, A. R. et al. "AlphaLogger: detecting motion-based side-channel attack using smartphone keystrokes". In: *Journal of Ambient Intelligence and Humanized Computing* (2020). DOI: 10.1007/s12652-020-01770-0.
- [292] Thomaz, E., Essa, I., and Abowd, G. D. "A practical approach for recognizing eating moments with wrist-mounted inertial sensing". In: *Proceedings of the ACM international joint conference on pervasive and ubiquitous computing*. 2015, pp. 1029–1040. DOI: 10.1145/2750858.2807545.
- [293] Tang, Q. et al. "Automated detection of puffing and smoking with wrist accelerometers". In: *International Conference on Pervasive Computing Technologies for Healthcare*. 2014, pp. 80–87. DOI: 10.4108/icst.pervasivehealth.2014.254978.
- [294] Tinker, B. *How Facebook 'likes' predict race, religion and sexual orientation*. 2018. URL: <https://www.cnn.com/2018/04/10/health/facebook-likes-psychographics/index.html> (visited on 11/09/2021).
- [295] Calandrino, J. A. et al. "'You Might Also Like:' Privacy Risks of Collaborative Filtering". In: *IEEE Symposium on Security and Privacy*. Oakland, CA, USA: IEEE, 2011, pp. 231–246. DOI: 10.1109/SP.2011.40. URL: <http://ieeexplore.ieee.org/document/5958032/>.
- [296] Mayer, J., Mutchler, P., and Mitchell, J. C. "Evaluating the privacy properties of telephone metadata". In: *Proceedings of the National Academy of Sciences* 113.20 (2016), pp. 5536–5541. DOI: 10.1073/pnas.1508081113.
- [297] Cadwalladr, C. and Graham-Harrison, E. "How Cambridge Analytica turned Facebook 'likes' into a lucrative political tool". In: *The Guardian* (2018). URL: <https://www.theguardian.com/technology/2018/mar/17/facebook-cambridge-analytica-kogan-data-algorithm> (visited on 11/09/2021).
- [298] Martinovic, D. et al. "'You are what you play': Breaching privacy and identifying users in online gaming". In: *International Conference on Privacy, Security and Trust*. tex.organization: IEEE. 2014, pp. 31–39. DOI: 10.1109/PST.2014.6890921.

REFERENCES

- [299] Freudiger, J., Shokri, R., and Hubaux, J.-P. “Evaluating the privacy risk of location-based services”. In: *International Conference on Financial Cryptography and Data Security*. 2011, pp. 31–46. DOI: 10.1007/978-3-642-27576-0_3.
- [300] Weizenbaum Institute for the Networked Society. *Mission Statement*. URL: <https://www.weizenbaum-institut.de/en/the-institute/mission-statement/> (visited on 02/10/2022).
- [301] Harborth, D. et al. “Anreize und hemmnisse für die implementierung von privacy-enhancing technologies im unternehmenskontext”. German. In: *SICHERHEIT 2018*. Gesellschaft für Informatik e.V., 2018.
- [302] Goold, B. J. “Building it in: the role of privacy enhancing technologies (PETs) in the regulation of surveillance and data collection”. In: *New Directions in Surveillance Privacy*. Ed. by B. J. Goold and D. Neyland. Routledge, 2013, pp. 41–61. ISBN: 1-134-04599-9.
- [303] Solove, D. J. “I’ve got nothing to hide and other misunderstandings of privacy”. In: *San Diego Law Review* 44 (2007). Publisher: HeinOnline, p. 745. URL: <https://heinonline.org/HOL/LandingPage?handle=hein.journals/sanlr44&div=40>.
- [304] Murray, B. “Informed Consent: What Must a Physician Disclose to a Patient?”. In: *AMA Journal of Ethics* 14.7 (July 2012). Publisher: American Medical Association, pp. 563–566. DOI: 10.1001/virtualmentor.2012.14.7.hlaw1-1207.
- [305] Rouvroy, A. and Poulet, Y. “The Right to Informational Self-Determination and the Value of Self-Development: Reassessing the Importance of Privacy for Democracy”. In: *Reinventing Data Protection?* Ed. by S. Gutwirth et al. Springer, 2009, pp. 45–76. ISBN: 978-1-4020-9497-2. URL: http://link.springer.com/10.1007/978-1-4020-9498-9_2.
- [306] Hodge, R. *Signal, WhatsApp and Telegram: Here’s which secure messaging app you should use*. 2022. URL: <https://www.cnet.com/tech/services-and-software/signal-whatsapp-and-telegram-heres-which-secure-messaging-app-you-should-use/> (visited on 02/22/2022).
- [307] Mozilla Foundation. *Seven of the best browsers in direct comparison*. URL: <https://www.mozilla.org/en-US/firefox/browsers/compare/> (visited on 02/22/2022).
- [308] Vemou, K. and Karyda, M. “A Classification of Factors Influencing Low Adoption of PETs Among SNS Users”. In: *Trust, Privacy, and Security in Digital Business*. Ed. by S. Furnell, C. Lambrinoudakis, and J. Lopez. Springer, 2013, pp. 74–84. DOI: 10.1007/978-3-642-40343-9_7.
- [309] Zou, Y. et al. “Examining the Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices”. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 2020. DOI: 10.1145/3313831.3376570.
- [310] Baruh, L. and Popescu, M. “Big data analytics and the limits of privacy self-management”. In: *New Media & Society* 19 (2015), pp. 579–596. DOI: 10.1177/1461444815614001.
- [311] Ustaran, E. *Yes, Consent Is Dead. Further, Continuing To Give It A Central Role Is Dangerous*. 2013. URL: <https://iapp.org/news/a/yes-consent-is-dead-further-continuing-to-give-it-a-central-role-is-danger/> (visited on 02/18/2022).
- [312] Pfeifle, S. *Keynote: Forget Notice and Choice, Let’s Regulate Use*. 2013. URL: <https://iapp.org/news/a/keynote-forget-notice-and-choice-lets-regulate-use/> (visited on 03/12/2022).
- [313] Fabian, B., Ermakova, T., and Lentz, T. “Large-scale readability analysis of privacy policies”. In: *Proceedings of the International Conference on Web Intelligence*. 2017, pp. 18–25. DOI: 10.1145/3106426.3106427.
- [314] Litman-Navarro, K. “We Read 150 Privacy Policies. They Were an Incomprehensible Disaster.” In: *The New York Times* (June 2019). URL: <https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html,%20https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html> (visited on 10/08/2021).
- [315] PlumX Metrics. *Citation Data for DOI 10.1007/978-3-030-22479-0_6*. URL: https://plu.mx/plum/a/?doi=10.1007/978-3-030-22479-0_6 (visited on 11/01/2021).
- [316] PlumX Metrics. *Citation Data for DOI 10.1007/978-3-030-42504-3_15*. URL: https://plu.mx/plum/a/?doi=10.1007/978-3-030-42504-3_15 (visited on 11/01/2021).

-
- [317] PlumX Metrics. *Citation Data for DOI 10.1145/3309074.3309076*. URL: <https://plu.mx/plum/a/?doi=10.1145/3309074.3309076> (visited on 11/01/2021).
 - [318] Doctorow, C. *Adding sensors to our computers revolutionized them ... [Tweet]*. July 2021. URL: <https://twitter.com/doctorow/status/1420816801743597569> (visited on 11/01/2021).
 - [319] Stewart-Williams, S. *Accelerometer data from smartphones etc. can reveal people's location ... [Tweet]*. 2021. URL: <https://twitter.com/SteveStuWill/status/1432372540484243461> (visited on 11/01/2021).
 - [320] Stewart-Williams, S. *Visual behaviour can reveal people's sex, age, ethnicity, personality traits ... [Tweet]*. Aug. 2020. URL: <https://twitter.com/SteveStuWill/status/1432372540484243461> (visited on 11/01/2021).
 - [321] Peterson, J. *What computers will soon learn about you and how ... [Tweet]*. URL: <https://twitter.com/jordanbpeterson/status/1376030447902326786> (visited on 11/01/2021).
 - [322] Kuzuloğlu, M. S. *Göz tanıma teknolojisi zannettiğinizden çok daha fazla bilgi toplayabiliyor ... [Tweet]*. Turkish. Mar. 2021. URL: <https://twitter.com/SteveStuWill/status/1432372540484243461> (visited on 11/01/2021).
 - [323] Maderas. *Now Eye Tracking, data analytics vs. your eyes gaze direction ... [Tweet]*. Apr. 2020. URL: <https://twitter.com/hackermaderas/status/1254823763046268928> (visited on 11/01/2021).
 - [324] Singh, J. *Accelerometer data alone is sufficient to obtain a device holder's ... [Tweet]*. May 2021. URL: <https://web.archive.org/web/20220315110536/https://twitter.com/hackingbutlegal/status/1390205634658852864> (visited on 11/01/2021).
 - [325] Kröger, J. L. and Raschke, P. "Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping". In: *Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec)*. Cham: Springer, 2019, pp. 102–120. DOI: 10.1007/978-3-030-22479-0_6.
 - [326] Kröger, J. L., Lutz, O. H.-M., and Müller, F. "What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking". In: *Privacy and Identity Management. Data for Better Living: AI and Privacy*. Ed. by M. Friedewald et al. Cham: Springer, 2020, pp. 226–241. URL: https://doi.org/10.1007/978-3-030-42504-3_15.
 - [327] Raschke, P. et al. "Towards real-time web tracking detection with T.EX - The transparency EXtension". In: *Annual Privacy Forum*. Springer, 2019, pp. 3–17. DOI: 10.1007/978-3-030-21752-5_1.
 - [328] Lutz, O. H.-M. et al. "Surfing in Sound: Sonification of hidden web tracking". In: *International conference on auditory display (ICAD)*. Georgia Institute of Technology, 2019, pp. 306–309. DOI: 10.21785/icad2019.071.
 - [329] Kröger, J. L., Lutz, O. H.-M., and Ullrich, S. "The myth of individual control: Mapping the limitations of privacy self-management". In: *SSRN Electronic Journal* (2021). DOI: 10.2139/ssrn.3881776. URL: <https://ssrn.com/abstract=3881776>.
 - [330] Kröger, J. L. et al. "Surveilling the Gamers: Privacy Impacts of the Video Game Industry". In: *SSRN Electronic Journal* (2021). DOI: 10.2139/ssrn.3881279. URL: <https://ssrn.com/abstract=3881279>.
 - [331] Kröger, J. L., Raschke, P., and Bhuiyan, T. R. "Privacy Implications of Accelerometer Data: A Review of Possible Inferences". In: *Proceedings of the International Conference on Cryptography, Security and Privacy*. ACM, 2019, pp. 81–87. DOI: 10.1145/3309074.3309076.
 - [332] Kröger, J. L., Lindemann, J., and Herrmann, D. "How do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on iOS and Android Apps". In: *Proceedings of the 15th International Conference on Availability, Reliability and Security*. ACM, 2020, pp. 1–10. DOI: 10.1145/3407023.3407057.
 - [333] Google Scholar. *Jacob Leon Kröger*. 2021. URL: https://scholar.google.com/citations?user=DAG_00EAAAAJ&hl=en (visited on 11/01/2021).
 - [334] Voßberg, R. *Datenschutz bei Apps - Das große Schweigen*. German. 2021. URL: <https://www.ferchau.com/de/de/blog/02-03-2021-das-grosse-schweigen-4937> (visited on 11/01/2021).

REFERENCES

- [335] Mysk, T. *iPhone Apps Can Tell Many Things About You Through the Accelerometer*, Mysk. 2021. URL: <https://www.mysk.blog/2021/10/24/accelerometer-ios/> (visited on 11/01/2021).
- [336] Doctorow, C. *Tracking you with accelerometer signatures*. 2021. URL: <https://doctorow.medium.com/tracking-you-with-accelerometer-signatures-5e2e762de5fd> (visited on 11/01/2021).
- [337] Wikrent, T. *Week-end Wrap – Political Economy*. 2021. URL: <https://real-economics.blogspot.com/2021/08/week-end-wrap-political-economy-august.html> (visited on 11/01/2021).
- [338] Fortuna, A. *Some thoughts about detectability and feasibility of mobile eavesdropping on smartphones*. 2020. URL: <https://www.andreafortuna.org/2020/05/07/some-thoughts-about-detectability-and-feasibility-of-mobile-eavesdropping-on-smartphones/> (visited on 11/01/2021).
- [339] Bar-Zeev, A. *Advertising is like Vegas*. 2021. URL: <https://avibarzeev.medium.com/advertising-is-like-vegas-f603e601d993> (visited on 11/01/2021).
- [340] Doctorow, C. *Pluralistic: 29 Jul 2021*. 2021. URL: <https://pluralistic.net/2021/07/29/impunity-corrodes/> (visited on 11/06/2021).
- [341] Gray, B. C. *Conspiracy Theorist About Issues of Privacy, ft. Carolyn Hoessler*. URL: <https://yougotthis.trubox.ca/podcast/episode-23-conspiracy-theorist-about-issues-of-privacy-ft-carolyn-hoessler/> (visited on 11/01/2021).
- [342] Techlore. *Surveillance Report 10*. URL: <https://surveillance-report.castos.com/episodes/surveillance-report-10> (visited on 11/01/2021).
- [343] Gough, C. and Hamer, A. *How Technology Is Enabling Archeological Discoveries and the Privacy Risks of Eye-Tracking Tech*. URL: <https://www.curiositydaily.com/how-technology-is-enabling-archeological-discoveries-w-elizabeth-sawchuk-and-mary-prendergast-and-the-privacy-risks-of-eye-tracking-tech/> (visited on 11/01/2021).
- [344] Matteo311 (Youtube channel). *New VR News - Eye Privacy Concerns | The Latest VR games*. Oct. 2020. URL: <https://www.youtube.com/watch?v=SPEK09D5MJ8> (visited on 11/01/2021).
- [345] Eric For President Clips (Youtube channel). *NO ONE is Ready For THIS In VR...* Aug. 2021. URL: <https://www.youtube.com/watch?v=q7APNV2LodU> (visited on 11/01/2021).
- [346] Rose, J. *Eye-Tracking Tech Is Another Reason the Metaverse Will Suck*. Mar. 2022. URL: <https://www.vice.com/en/article/93b8v8/eye-tracking-tech-is-another-reason-the-metaverse-will-suck> (visited on 03/12/2022).
- [347] Mitteldeutscher Rundfunk. *Spion in der Hosentasche: Schon die App-Nutzung verrät unsere Identität | MDR.DE*. German. 2022. URL: <https://www.mdr.de/wissen/smartphone-handy-spion-ab hoeren-daten-aktivitaet-identitaet-apps-100.html> (visited on 03/15/2022).
- [348] Zorz, Z. *The state of GDPR compliance in the mobile app space*. Aug. 2020. URL: <https://www.helpnetsecurity.com/2020/08/26/gdpr-compliance-mobile-app-space/> (visited on 11/01/2021).
- [349] Petereit, D. *Apps: Forscher empfehlen, Nutzerprofile mit falschen Daten anzulegen*. German. URL: <https://t3n.de/news/apps-forscher-empfehlen-fake-profile-1330150/> (visited on 11/01/2021).
- [350] Stellmach, V. *Falsche Angaben: Datenschützer und Forscher raten, bei Apps zu lügen*. German. Oct. 2020. URL: <https://www.basichthinking.de/blog/2020/10/28/falsche-angaben-datenschutz/> (visited on 11/01/2021).
- [351] Der Neue Wiesentbote. *Bamberger Informatiker im Undercover-Einsatz*. German. Oct. 2020. URL: <https://www.wiesentbote.de/2020/10/22/bamberger-informatiker-im-undercover-einsatz/> (visited on 11/01/2021).
- [352] Achter, P. *Informatiker im Undercover-Einsatz*. German. Oct. 2020. URL: <https://nachrichten.idw-online.de/2020/10/21/informatiker-im-undercover-einsatz/> (visited on 11/01/2021).
- [353] Homburg 1. *Informatiker im Undercover-Einsatz: Geben App-Anbieter persönliche Nutzerdaten auf Anfrage heraus?* German. Oct. 2020. URL: <https://homburg1.de/informatiker-im-undercover-einsatz-geben-app-anbieter-persoenliche-nutzerdaten-auf-anfrage-heraus-107887/> (visited on 11/01/2021).

- [354] Soyez, F. *Mythe ou réalité: votre téléphone vous espionne-t-il vraiment ?* French. URL: <https://www.cnetfrance.fr/news/mythe-ou-realite-votre-telephone-vous-espionne-t-il-vraiment-39896041.htm> (visited on 11/01/2021).
- [355] Blain, L. *Eye tracking can reveal an unbelievable amount of information about you*. Mar. 2021. URL: <https://newatlas.com/science/science/eye-tracking-privacy/> (visited on 11/01/2021).
- [356] Back, E. *Vos yeux révèlent quasiment tous les secrets sur votre identité*. French. URL: <https://www.futura-sciences.com/tech/actualites/technologie-vos-yeux-revelent-quasiment-tous-secrets-votre-identite-86553/> (visited on 11/01/2021).
- [357] Kardoudi, O. *El movimiento de tus ojos puede revelar tus secretos más íntimos*. Spanish. Apr. 2021. URL: https://www.elconfidencial.com/tecnologia/novaceno/2021-04-07/movimiento-ojos-revela-secretos-intimos_3024448/ (visited on 11/01/2021).
- [358] Lavoie, E. *We Should Worry About Virtual Reality Sex*. Dec. 2020. URL: https://www.realclearscience.com/articles/2020/12/09/we_should_worry_about_virtual_reality_sex_652506.html (visited on 11/01/2021).
- [359] Puiu, T. “Microphone-enabled smart devices are a huge privacy concern, but most of us aren’t aware of it”. In: *ZME Science* (Nov. 2021). URL: <https://www.zmescience.com/science/microphone-voice-recs-privacy-0931321/> (visited on 11/23/2021).
- [360] Wheeler, T. “If the Metaverse Is Left Unregulated, Companies Will Track Your Gaze and Emotions”. In: *Time magazine* (2022). URL: <https://time.com/6188956/metaverse-is-left-unregulated-companies-will-track-gaze-emotions/> (visited on 07/12/2022).
- [361] Scott, M. “Digital Bridge: Policing the metaverse — Kids’ online safety — Digital tax next steps”. In: *POLITICO* (2022). URL: <https://www.politico.eu/newsletter/digital-bridge/policing-the-metaverse-kids-online-safety-digital-tax-next-steps/> (visited on 07/12/2022).
- [362] Selinger, E. “Facebook’s next privacy nightmare will be a sight to see - The Boston Globe”. In: *Boston Globe* (2021). URL: <https://www.bostonglobe.com/2021/11/12/opinion/facebooks-next-privacy-nightmare-will-be-sight-see/> (visited on 11/15/2021).
- [363] Rodriguez, K., Opsahl, K., and Leufer, D. *Virtual Worlds, Real People: Human Rights in the Metaverse*. Dec. 2021. URL: <https://www.eff.org/deeplinks/2021/12/virtual-worlds-real-people-human-rights-metaverse> (visited on 01/05/2022).
- [364] McGill, M. *The IEEE Global Initiative on Ethics of Extended Reality (XR) Report - Extended Reality (XR) and the Erosion of Anonymity and Privacy*. Tech. rep. 2021, p. 24. URL: <https://standards.ieee.org/wp-content/uploads/import/governance/iccom/extended-reality-anonymity-privacy.pdf>.
- [365] Soroushian, J. and Neschke, S. *Thinking Ahead About XR: Privacy and Security in an Immersive World | Bipartisan Policy Center*. 2021. URL: <https://bipartisanpolicy.org/blog/thinking-ahead-about-xr-privacy-and-security-in-an-immersive-world/> (visited on 12/08/2021).
- [366] Dick, E. *Balancing user privacy and innovation in augmented and virtual reality*. Washington, D.C.: Information Technology and Innovation Foundation, 2021. URL: <https://itif.org/publications/2021/03/04/balancing-user-privacy-and-innovation-augmented-and-virtual-reality>.
- [367] McGann, J. G. *TTCS Global Go To Think Tank Index Reports*. Tech. rep. Publisher: University of Pennsylvania. Think Tanks and Civil Societies Program, 2019. URL: https://repository.upenn.edu/think_tanks/16.
- [368] Pontoniere, P. *Biometrics on the Rise*. Jan. 2022. URL: <https://neo.life/2022/01/biometrics-on-the-rise/> (visited on 01/21/2022).
- [369] Kloiber, M. *App-Anbieter reagieren nur mangelhaft auf Auskunftsbegehren*. 2020. URL: <https://www.deutschlandfunk.de/app-anbieter-reagieren-nur-mangelhaft-auf-auskunftsbegehren-interv-jacob-kroeger-dlf-a7e75d73-100.html> (visited on 11/01/2021).

REFERENCES

- [370] Sickert, T. and Kogel, D. *Lauschangriff für personalisierte Werbung - Vom eigenen Smartphone ausspioniert*. German. 2021. URL: <https://web.archive.org/web/20220220003950/https://www.deutschlandfunkkultur.de/spionage-durch-eigenes-smartphone-100.html> (visited on 02/20/2022).
- [371] Pérez, M. H. “‘Gamers’ y cazaballenas: así impactan los videojuegos en la privacidad de sus usuarios”. Spanish. In: *El País* (2021). URL: <https://elpais.com/tecnologia/2021-08-02/gamers-y-cazaballenas-asi-impactan-los-videojuegos-en-la-privacidad-de-sus-usuarios.html> (visited on 11/01/2021).
- [372] Aluffi, G. “Gli smartphone ci spiano davvero?” Italian. In: *la Repubblica* (Sept. 2021). URL: https://www.repubblica.it/cronaca/2021/09/30/news/telefonini_spia-320120446/ (visited on 11/01/2021).
- [373] Kröger, J. L. *What can companies learn about you by analyzing how you hold and move your mobile devices ... [Tweet]*. July 2021. URL: https://twitter.com/JL_Kroger/status/1420681035617116163 (visited on 11/01/2021).
- [374] Kröger, J. L. *I’ll summarize our study on the privacy implications of eye tracking ... [Tweet]*. May 2021. URL: https://twitter.com/JL_Kroger/status/1392789775569018881 (visited on 11/01/2021).
- [375] Kröger, J. L. *Are apps listening to people’s conversations to improve ad targeting? ... [Tweet]*. Aug. 2021. URL: https://twitter.com/JL_Kroger/status/1445333142869577733 (visited on 11/01/2021).
- [376] Herrmann, D. *I’ll summarize our study on app vendors’ responses to Subject Access Requests ... [Tweet]*. Aug. 2021. URL: <https://twitter.com/herdom/status/1298921164295876608> (visited on 11/01/2021).
- [377] Matz, S. C. et al. “Psychological targeting as an effective approach to digital mass persuasion”. In: *Proceedings of the National Academy of Sciences* 114.48 (2017), pp. 12714–12719. DOI: 10.1073/pnas.1710966114.
- [378] Cabañas, J. G. et al. “Does Facebook use sensitive data for advertising purposes?” In: *Communications of the ACM* 64.1 (2020), pp. 62–69. DOI: 10.1145/3426361.
- [379] Solove, D. J. “Why privacy matters even if you have ‘nothing to hide’”. In: *Chronicle of Higher Education* (2011). URL: http://www.woldww.net/classes/Information_Ethics/Solove-ChronicleArticle-NothingToHide.pdf.
- [380] Amnesty International. *7 reasons why ‘I’ve got nothing to hide’ is the wrong response to mass surveillance*. 2015. URL: <https://www.amnesty.org/en/latest/campaigns/2015/04/7-reasons-why-ive-got-nothing-to-hide-is-the-wrong-response-to-mass-surveillance/> (visited on 11/10/2021).
- [381] Slepian, M. L. and Kirby, J. N. “To Whom Do We Confide Our Secrets?” In: *Personality and Social Psychology Bulletin* 44.7 (2018), pp. 1008–1023. DOI: 10.1177/0146167218756032.
- [382] Flaherty, D. H. “Visions of Privacy: Past, Present and Future”. In: *Visions of Privacy: Policy Choices for the Digital Age*. Ed. by C. J. Bennett and R. A. Grant. University of Toronto Press, 1999, pp. 19–38. ISBN: 978-1-4426-8310-5.
- [383] Kröger, J. L., Miceli, M., and Müller, F. “How data can be used against people: A classification of personal data misuses”. In: *SSRN Electronic Journal* (2021). DOI: 10.2139/ssrn.3887097. URL: <https://ssrn.com/abstract=3887097>.
- [384] Wachter, S. “The GDPR and the Internet of Things: a three-step transparency model”. In: *Law, Innovation and Technology* 10.2 (2018). Publisher: Taylor & Francis, pp. 266–294. DOI: 10.1080/17579961.2018.1527479.
- [385] Verdegem, P. “Tim Berners-Lee’s plan to save the internet: give us back control of our data”. In: *The Conversation* (2021). URL: <http://theconversation.com/tim-berners-lees-plan-to-save-the-internet-give-us-back-control-of-our-data-154130> (visited on 10/15/2021).
- [386] Orphanides, C. “How Tim Berners-Lee’s Inrupt project plans to fix the web”. In: *Wired UK* (2019). URL: <https://www.wired.co.uk/article/inrupt-tim-berners-lee> (visited on 10/15/2021).

-
- [387] Swanson, B. “Europe And California Get It Wrong; Technology Is The Solution To Digital Privacy”. In: *Forbes* (2018). URL: <https://www.forbes.com/sites/washingtonbytes/2018/09/25/europe-and-california-get-it-wrong-technology-is-the-solution-to-digital-privacy/> (visited on 10/15/2021).
 - [388] Dorri, A. et al. “Blockchain: A distributed solution to automotive security and privacy”. In: *IEEE Communications Magazine* 55.12 (2017), pp. 119–125.
 - [389] Wiggers, K. *AI has a privacy problem, but these techniques could fix it*. Dec. 2019. URL: <https://venturebeat.com/2019/12/21/ai-has-a-privacy-problem-but-these-techniques-could-fix-it/> (visited on 10/15/2021).
 - [390] Greenberg, A. “How One of Apple’s Key Privacy Safeguards Falls Short”. In: *Wired* (2017). URL: <https://www.wired.com/story/apple-differential-privacy-shortcomings/> (visited on 10/15/2021).
 - [391] Yeung, C.-m. A. et al. “Decentralization: The Future of Online Social Networking”. In: *W3C Workshop on the Future of Social Networking Position Papers*. 2009. URL: <https://www.w3.org/2008/09/msnws/papers/decentralization.pdf>.
 - [392] Waterhouse, S. “These two technologies could supercharge our privacy on the internet”. In: *Fast Company* (2021). URL: <https://www.fastcompany.com/90602972/encryption-and-decentralization-privacy> (visited on 10/15/2021).
 - [393] Gangopadhyay, A. and Chen, Z. *AI could help solve the privacy problems it has created*. 2020. URL: <http://theconversation.com/ai-could-help-solve-the-privacy-problems-it-has-created-130510> (visited on 10/15/2021).
 - [394] Mearian, L. “How blockchain could solve the internet privacy problem”. In: *Computerworld* (Apr. 2018). URL: <https://www.computerworld.com/article/3267930/how-blockchain-could-solve-the-internet-privacy-problem.html> (visited on 10/15/2021).
 - [395] Fauvre-Willis, A. *Concerns around data privacy are rising, and blockchain is the solution*. 2021. URL: <https://cointelegraph.com/news/concerns-around-data-privacy-are-rising-and-blockchain-is-the-solution> (visited on 10/15/2021).
 - [396] Le Métayer, D. “Whom to trust? Using technology to enforce privacy”. In: *Enforcing Privacy*. Ed. by D. Wright and P. De Hert. Springer, 2016, pp. 395–437. URL: https://doi.org/10.1007/978-3-319-25047-2_17.
 - [397] Litman-Navarro, K. “We Read 150 Privacy Policies. They Were an Incomprehensible Disaster”. In: *The New York Times* (June 2019). URL: <https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html,%20https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html> (visited on 10/26/2021).
 - [398] Mihale-Wilson, C., Zibuschka, J., and Hinz, O. “About User Preferences and Willingness to Pay for a Secure and Privacy Protective Ubiquitous Personal Assistant”. In: *European Conference on Information Systems (ECIS) Proceedings*. 2017. URL: https://aisel.aisnet.org/ecis2017_rp/3.
 - [399] Beresford, A. R., Kübler, D., and Preibusch, S. “Unwillingness to pay for privacy: A field experiment”. In: *Economics Letters* 117.1 (2012), pp. 25–27. DOI: 10.1016/j.econlet.2012.04.077.
 - [400] Agrawal, N. et al. “Exploring Design and Governance Challenges in the Development of Privacy-Preserving Computation”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (May 2021), pp. 1–13. DOI: 10.1145/3411764.3445677. (Visited on 11/11/2021).
 - [401] Leon, P. et al. “Why Johnny can’t opt out: a usability evaluation of tools to limit online behavioral advertising”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2012, pp. 589–598. DOI: 10.1145/2207676.2207759.
 - [402] Beck, L. L. “A security mechanism for statistical database”. In: *ACM Transactions on Database Systems (TODS)* 5.3 (1980). Publisher: ACM New York, NY, USA, pp. 316–3338. DOI: 10.1145/320613.320617.
 - [403] Hinke, T. H., Delugach, H. S., and Wolf, R. P. “Protecting databases from inference attacks”. In: *Computers & Security* 16.8 (1997), pp. 687–708. DOI: 10.1016/S0167-4048(97)87607-9.

REFERENCES

- [404] Domingo-Ferrer, J. “A Survey of Inference Control Methods for Privacy-Preserving Data Mining”. In: *Privacy-Preserving Data Mining: Models and Algorithms*. Ed. by C. C. Aggarwal and P. S. Yu. Advances in Database Systems. Springer, 2008, pp. 53–80. ISBN: 978-0-387-70992-5. URL: https://doi.org/10.1007/978-0-387-70992-5_3.
- [405] Hegde, A. et al. “SoK: Efficient Privacy-preserving Clustering”. In: *Proceedings on Privacy Enhancing Technologies* 2021.4 (2021), pp. 225–248. DOI: 10.2478/popets-2021-0068. URL: <https://www.sciendo.com/article/10.2478/popets-2021-0068>.
- [406] Koo, J., Kang, G., and Kim, Y.-G. “Security and Privacy in Big Data Life Cycle: A Survey and Open Challenges”. In: *Sustainability* 12.24 (2020), p. 10571. DOI: 10.3390/su122410571. URL: <https://www.mdpi.com/2071-1050/12/24/10571> (visited on 02/20/2022).
- [407] Butler, D. J. et al. “The Privacy-Utility Tradeoff for Remotely Teleoperated Robots”. In: *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. Portland, USA: ACM, 2015, pp. 27–34. DOI: 10.1145/2696454.2696484. URL: <https://dl.acm.org/doi/10.1145/2696454.2696484>.
- [408] Zhang, X.-Y. et al. “Privacy-Functionality Trade-Off: A privacy-preserving multi-channel smart metering system”. In: *Energies* 13.12 (2020). Publisher: Multidisciplinary Digital Publishing Institute, p. 3221. DOI: 10.3390/en13123221.
- [409] Liu, A. et al. “Differential privacy for eye-tracking data”. In: *Proceedings of the ACM Symposium on Eye Tracking Research & Applications*. June 2019. DOI: 10.1145/3314111.3319823. URL: <https://dl.acm.org/doi/10.1145/3314111.3319823>.
- [410] Ahmed, S. et al. “Preech: A system for privacy-preserving speech transcription”. In: *29th USENIX Security Symposium*. 2020, pp. 2703–2720.
- [411] The Guardian. “The Guardian view on internet privacy: technology can’t fix it | Editorial”. In: (Jan. 2017). ISSN: 0261-3077. URL: <https://www.theguardian.com/commentisfree/2017/jan/13/the-guardian-view-on-internet-privacy-technology-cant-fix-it>.
- [412] Ooijen, I. van and Vrabec, H. U. “Does the GDPR enhance consumers’ control over personal data? An analysis from a behavioural perspective”. In: *Journal of Consumer Policy* 42.1 (2019). Publisher: Springer, pp. 91–107. DOI: 10.1007/s10603-018-9399-7.
- [413] IAB Europe. *GDPR Guidance: Legitimate Interests Assessments (LIA) for Digital Advertising*. 2021. URL: <https://web.archive.org/web/20210331063930/https://iabeurope.eu/blog/gdpr-guidance-legitimate-interests-assessments-lia-for-digital-advertising/> (visited on 02/20/2022).
- [414] European Parliament. *European Parliament resolution of 25 March 2021 on the Commission evaluation report on the implementation of the General Data Protection Regulation two years after its application (2020/2717(RSP))*. Tech. rep. Brussels, 2021. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52021IP0111>.
- [415] Boardman, R. *Irish Data Protection Commission WhatsApp decision: what do you need to know?* 2021. URL: <http://www.twobirds.com/en/news/articles/2021/uk/irish-data-protection-commission-whatsapp-decision> (visited on 02/20/2022).
- [416] Hale, W. C. P. et al. *Belgian Data Protection Authority Rules IAB Cookie Consent Framework Violates the GDPR*. 2022. URL: <https://www.lexology.com/library/detail.aspx?g=51643f4e-a723-4d77-be00-d682f2797c7d> (visited on 02/20/2022).
- [417] Information Commissioner’s Office. *How do we apply legitimate interests in practice?* URL: <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/legitimate-interests/how-do-we-apply-legitimate-interests-in-practice/> (visited on 02/20/2022).
- [418] Sanchez-Rola, I. et al. “Can I Opt Out Yet? GDPR and the Global Illusion of Cookie Control”. In: *Proceedings of the ACM Asia conference on computer and communications security*. 2019, pp. 340–351. DOI: 10.1145/3321705.3329806.

- [419] AlgorithmWatch. *Use the DSA to Stop Platforms from Suppressing Public Interest Research*. 2021. URL: https://algorithmwatch.org/en/wp-content/uploads/2021/09/Open_Letter_to_European_Lawmakers_15_sept_21.pdf (visited on 02/22/2022).
- [420] Clarke, R. “Guidelines for the responsible application of data analytics”. In: *Computer Law & Security Review* 34.3 (2018). Publisher: Elsevier, pp. 467–476. DOI: 10.1016/j.clsr.2017.11.002.
- [421] Schmidt, A. “Secrecy versus openness: Internet security and the limits of open source and peer production”. Doctoral thesis. TU Delft, 2014. URL: <https://repository.tudelft.nl/islandora/object/uuid%3Aecf237ed-7131-4455-917f-11e55e03df0d>.
- [422] Pfleeger, C. P. “Looking into Software Transparency”. In: *IEEE Security Privacy* 14.1 (Jan. 2016), pp. 31–36. DOI: 10.1109/MSP.2016.5.
- [423] Hansen, M., Köhntopp, K., and Pfitzmann, A. “The Open Source approach - opportunities and limitations with respect to security and privacy”. In: *Computers & Security* 21.5 (2002). Publisher: Elsevier, pp. 461–471. DOI: 10.1016/S0167-4048(02)00516-3.
- [424] Cadwalladr, C. and Graham-Harrison, E. “Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach”. en-GB. In: *The Guardian* (Mar. 2018). URL: <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election> (visited on 11/11/2021).
- [425] *EU Directive 2019/1937*. Nov. 2019. URL: <http://data.europa.eu/eli/dir/2019/1937/oj/eng> (visited on 10/14/2021).
- [426] Sharma, J. P., Kanojia, S., and Sachdeva, S. “Comparison of Whistle-blower Protection Mechanism of Select Countries”. In: *Indian Journal of Corporate Governance* 11.1 (2018). Publisher: SAGE Publications, pp. 45–68. DOI: 10.1177/0974686218769198.
- [427] Directorate-General for Justice and Consumers. *Whistleblower protection*. 2018. URL: https://ec.europa.eu/info/sites/default/files/placeholder_11.pdf (visited on 10/14/2021).
- [428] Abazi, V. “The European Union Whistleblower Directive: A ‘Game Changer’ for Whistleblowing Protection?” In: *Industrial Law Journal* 49.4 (2020), pp. 640–656. ISSN: 0305-9332. DOI: 10.1093/indlaw/dwaa023. URL: <https://doi.org/10.1093/indlaw/dwaa023> (visited on 10/14/2021).
- [429] Article 29 Working Party. *Opinion 03/2013 on purpose limitation*. Tech. rep. 00569/13/EN WP 203. 2013. URL: https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf.
- [430] State of California. *California Consumer Privacy Act of 2018*. URL: https://leginfo.ca.gov/faces/codes_displayText.xhtml?lawCode=CIV&division=3.&title=1.81.5.&part=4.&chapter=&article (visited on 11/12/2021).
- [431] Corte, L. D. “Scoping Personal Data: Towards a Nuanced Interpretation of the Material Scope of EU Data Protection Law”. In: *European Journal of Law and Technology* 10.1 (2019). URL: <https://ejlt.org/index.php/ejlt/article/view/672>.
- [432] O’Callaghan, J. “Inferential privacy and artificial intelligence – a new frontier?” In: *Journal of Law & Economic Regulation* 11.2 (2018), pp. 72–89.
- [433] Mittelstadt, B. D. et al. “The ethics of algorithms: Mapping the debate”. In: *Big Data & Society* 3.2 (2016). Publisher: SAGE Publications. DOI: 10.1177/2053951716679679.
- [434] Article 29 Working Party. *Guidelines on the right to data portability*. Tech. rep. 16/EN WP 242 rev.01. 2016. URL: https://ec.europa.eu/newsroom/document.cfm?doc_id=44099.
- [435] Malgieri, G. “Property and (Intellectual) Ownership of Consumers’ Information: A New Taxonomy for Personal Data”. In: *Privacy in Germany* 4 (2016), p. 5. ISSN: 2196-9817. DOI: 10.37307/j.2196-9817.2016.04.05. URL: <https://pingdigital.de/ce/property-and-intellectual-ownership-of-consumers-information-a-new-taxonomy-for-personal-data/detail.html>.
- [436] Kamann, H.-G. and Braun, M. “Art. 16 Recht auf Berichtigung”. German. In: *Datenschutz-Grundverordnung*. Ed. by Eugen Ehmann and Martin Selmayr. Vol. 2. C.H. Beck, 2018, pp. 369–380. ISBN: 978-3-406-72006-2.

REFERENCES

- [437] Miglicco, G. “GDPR is here and it is time to get serious”. In: *Computer Fraud & Security* 2018.9 (2018), pp. 9–12. DOI: 10.1016/S1361-3723(18)30085-X.
- [438] Malgieri, G. “Automated decision-making in the EU Member States: The right to explanation and other “suitable safeguards” in the national legislations”. In: *Computer Law & Security Review* 35.5 (2019). ISSN: 0267-3649. DOI: 10.1016/j.clsr.2019.05.002.
- [439] Molnár-Gábor, F. et al. “Harmonization after the GDPR? Divergences in the rules for genetic and health data sharing in four member states and ways to overcome them by EU measures: Insights from Germany, Greece, Latvia and Sweden”. In: *Seminars in Cancer Biology* (2021). DOI: 10.1016/j.semcancer.2021.12.001.
- [440] Information Commissioner’s Office. *Right to be informed*. URL: <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/right-to-be-informed/> (visited on 02/25/2022).
- [441] Sayaf, R. “Contextual Privacy: The Interplay of Sensitivity and Context”. Doctoral thesis. KU Leuven, 2016. URL: <https://lirias.kuleuven.be/1656619>.
- [442] Wachter, S. “Data protection in the age of big data”. In: *Nature Electronics* 2.1 (2019), pp. 6–7. DOI: 10.1038/s41928-018-0193-y. URL: <https://www.nature.com/articles/s41928-018-0193-y> (visited on 10/14/2021).
- [443] Finck, M. and Pallas, F. “They who must not be identified - distinguishing personal from non-personal data under the GDPR”. In: *International Data Privacy Law* (2020).
- [444] Kasperbauer, T. J. “Protecting health privacy even when privacy is lost”. In: *Journal of Medical Ethics* 46.11 (2020), pp. 768–772. DOI: 10.1136/medethics-2019-105880.
- [445] Mehmood, A. et al. “Protection of big data privacy”. In: *IEEE Access* 4 (2016), pp. 1821–1834. DOI: 10.1109/ACCESS.2016.2558446.
- [446] Ohm, P. “Broken promises of privacy: Responding to the surprising failure of anonymization”. In: *UCLA Law Review* 57 (2010). Publisher: HeinOnline, pp. 1701–1777.
- [447] Narayanan, A. and Felten, E. W. *No silver bullet: De-identification still doesn’t work*. 2014. URL: <http://www.randomwalker.info/publications/no-silver-bullet-de-identification.pdf><http://www.randomwalker.info/publications/no-silver-bullet-de-identification.pdf>.
- [448] Purtova, N. “The law of everything. Broad concept of personal data and future of EU data protection law”. In: *Law, Innovation and Technology* 10.1 (2018), pp. 40–81. DOI: 10.1080/17579961.2018.1452176.
- [449] Shao, Y. et al. “Fast de-anonymization of social networks with structural information”. In: *Data Science and Engineering* 4.1 (2019). Publisher: Springer, pp. 76–92. DOI: 10.1007/s41019-019-0086-8.
- [450] German Data Ethics Commission. *Opinion of the Data Ethics Commission*. 2019. URL: <https://datenethikkommissionhttps://www.bmi.bund.de/SharedDocs/downloads/EN/themen/it-digital-policy/datenethikkommission-abschlussgutachten-lang.html>.
- [451] Hildebrandt, M. and Gutwirth, S. “General Introduction and Overview”. In: *Profiling the European Citizen*. Ed. by M. Hildebrandt and S. Gutwirth. Springer, 2008, pp. 1–13. ISBN: 978-1-4020-6914-7. URL: https://link.springer.com/chapter/10.1007/978-1-4020-6914-7_1.
- [452] GDPR.eu. *GDPR Recital 26*. URL: <https://gdpr.eu/recital-26-not-applicable-to-anonymous-data/> (visited on 02/22/2022).
- [453] Vinocur, N. “We have a huge problem’: European tech regulator despairs over lack of enforcement. 2019. URL: <https://www.politico.com/news/2019/12/27/europe-gdpr-technology-regulation-089605> (visited on 11/16/2021).
- [454] Bodoni, S. “Silicon Valley’s Top Privacy Cop Rejects Claims She’s Too Lax”. In: *Bloomberg.com* (2021). URL: <https://www.bloomberg.com/news/articles/2021-11-18/eu-privacy-enforcement-not-good-enough-top-official-warns> (visited on 02/22/2022).
- [455] Hern, A. “‘Anonymised’ data can never be totally anonymous, says study”. In: *The Guardian* (July 2019). ISSN: 0261-3077. URL: <https://www.theguardian.com/technology/2019/jul/23/anonymised-data-never-be-anonymous-enough-study-finds> (visited on 02/22/2022).

- [456] Medine, D. *Data Protection: Consent is Dead (Long Live Privacy)*. 2021. URL: <https://www.cgdev.org/blog/data-protection-consent-dead-long-live-privacy> (visited on 02/18/2022).
- [457] Kamara, I. and De Hert, P. “Understanding the Balancing Act Behind the Legitimate Interest of the Controller Ground: A Pragmatic Approach”. In: *SSRN Electronic Journal* (2018). DOI: 10.2139/ssrn.3228369. URL: <https://www.ssrn.com/abstract=3228369>.
- [458] Demetzou, K. “Data Protection Impact Assessment: A tool for accountability and the unclarified concept of ‘high risk’ in the General Data Protection Regulation”. In: *Computer Law & Security Review* 35.6 (2019). DOI: 10.1016/j.clsr.2019.105342. URL: <https://www.sciencedirect.com/science/article/pii/S0267364918304357>.
- [459] Jourard, S. M. “Some psychological aspects of privacy”. In: *Law and Contemporary Problems* 31.2 (1966). Publisher: HeinOnline, pp. 307–318. DOI: 10.2307/1190673.
- [460] Goffman, E. “The Presentation of Self in Everyday Life”. In: *Sociology: Exploring the Architecture of Everyday Life Readings*. Ed. by D. M. Newman and J. O’Brien. Pine Forge Press, 1959, pp. 120–129.
- [461] Acquisti, A., Brandimarte, L., and Loewenstein, G. “Privacy and human behavior in the age of information”. In: *Science* 347.6221 (2015), pp. 509–514. DOI: 10.1126/science.aaa1465.
- [462] European Commission. *Proposal for a Regulation of the European Parliament of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*. 2021. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52021PC0206&from=EN> (visited on 03/04/2022).
- [463] AlgorithmWatch. *Draft AI Act: EU needs to live up to its own ambitions in terms of governance and enforcement*. 2021. URL: <https://algorithmwatch.org/en/eu-ai-act-consultation-submission-2021/> (visited on 03/04/2022).
- [464] Lomas, N. *Europe’s AI Act falls far short on protecting fundamental rights, civil society groups warn*. 2021. URL: <https://techcrunch.com/2021/11/30/eu-ai-act-civil-society-recommendations/> (visited on 03/04/2022).
- [465] Stahl, B. C. *EU is cracking down on AI, but leaves a loophole for mass surveillance*. 2021. URL: <http://theconversation.com/eu-is-cracking-down-on-ai-but-leaves-a-loophole-for-mass-surveillance-159421> (visited on 03/04/2022).
- [466] Skelton, S. K. *Ban predictive policing systems in EU AI Act, says civil society*. 2022. URL: <https://www.computerweekly.com/news/252514030/Ban-predictive-policing-systems-in-EU-AI-Act-says-civil-society> (visited on 03/05/2022).
- [467] Fair Trials. *AI Act: EU must ban predictive AI systems in policing and criminal justice*. 2022. URL: <https://www.fairtrials.org/articles/news/ai-act-eu-must-ban-predictive-ai-systems-in-policing-and-criminal-justice/> (visited on 03/05/2022).
- [468] Doctorow, C. *How to Destroy Surveillance Capitalism (Synopsis)*. 2021. URL: <https://www.kobo.com/us/en/ebook/how-to-destroy-surveillance-capitalism> (visited on 11/16/2021).
- [469] Tracking-Free Ads Coalition. *About us*. URL: <https://trackingfreeads.eu/about-us/> (visited on 02/21/2022).
- [470] European Data Protection Supervisor. *Online targeting for political advertising: stricter rules are necessary*. 2022. URL: https://edps.europa.eu/press-publications/press-news/press-releases/2022/online-targeting-political-advertising-stricter_fr (visited on 02/21/2022).
- [471] Jacobs, H. *Is Momentum Shifting Toward a Ban on Behavioral Advertising? – The Markup*. Section: Ask The Markup. 2022. URL: <https://themarkup.org/ask-the-markup/2022/02/03/is-momentum-shifting-toward-a-ban-on-behavioral-advertising> (visited on 03/04/2022).
- [472] Goujard, C. *European Parliament pushes to ban targeted ads based on health, religion or sexual orientation*. Jan. 2022. URL: <https://www.politico.eu/article/european-parliament-bans-use-of-sensitive-personal-data-for-targeted-ads/> (visited on 03/04/2022).

REFERENCES

- [473] Autorité de protection des données. *Decision on the merits 21/2022 of 2 February 2022 - Concerning: Complaint relating to Transparency & Consent Framework (Unofficial translation from Dutch)*. Feb. 2022. URL: <https://www.autoriteprotectiondonnees.be/publications/decision-quant-au-fond-n-21-2022-english.pdf#page111> (visited on 03/04/2022).
- [474] Lomas, N. *IAB fined as TCF found to breach Europe's GDPR*. 2022. URL: <https://social.techcrunch.com/2022/02/02/iab-tcf-gdpr-breaches/> (visited on 03/04/2022).
- [475] Irish Council for Civil Liberties. *GDPR enforcer rules that IAB Europe's consent popups are unlawful*. 2022. URL: <https://www.iccl.ie/news/gdpr-enforcer-rules-that-iab-europes-consent-popups-are-unlawful/> (visited on 03/04/2022).
- [476] Amnesty International. *Facebook Files: How a ban on surveillance advertising can fix Facebook*. 2021. URL: <https://www.amnesty.org/en/latest/campaigns/2021/10/facebook-files-how-a-ban-on-surveillance-advertising-can-fix-facebook/> (visited on 10/26/2021).
- [477] Goodwins, R. "Google's 'Be Evil' business transformation is complete". In: *The Register* (2021). URL: https://www.theregister.com/2021/11/01/google_opinion_column/.
- [478] Witt, A. C. "Excessive Data Collection as a Form of Anticompetitive Conduct: The German Facebook Case". In: *The Antitrust Bulletin* 66.2 (2021). Publisher: SAGE Publications, pp. 276–307. DOI: 10.1177/0003603X21997028.
- [479] "Pokémon GO Caught Millions of Players and Their Data". In: *Information Management* 50.5 (2016). Ed. by V. Wiler, p. 12.
- [480] Nouwens, M. et al. "Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence". In: *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2020). DOI: 10.1145/3313831.3376321.
- [481] Utz, C. et al. "(Un)informed consent: Studying GDPR consent notices in the field". In: *Proceedings of the ACM SIGSAG Conference on Computer and Communications Security*. 2019, pp. 973–990. DOI: 10.1145/3319535.3354212.
- [482] Hagendorff, T. "From privacy to anti-discrimination in times of machine learning". In: *Ethics and Information Technology* 21.4 (Dec. 2019), pp. 331–343. DOI: 10.1007/s10676-019-09510-5.
- [483] AlgorithmWatch. *A paradigm shift in German digital policies? – The newly presented German coalition agreement shows good approaches, but there is need for clarification*. 2021. URL: <https://algorithmwatch.org/en/german-coalition-agreement-2021/> (visited on 02/17/2022).
- [484] European Data Protection Supervisor. *EDPS Opinions on the Digital Services Act and the Digital Markets Act*. 2021. URL: https://edps.europa.eu/press-publications/press-news/press-releases/2021/edps-opinions-digital-services-act-and-digital_en (visited on 02/17/2022).
- [485] Mittelstadt, B. "From Individual to Group Privacy in Big Data Analytics". In: *Philosophy & Technology* 30.4 (2017). Publisher: Springer, pp. 475–494. DOI: 10.1007/s13347-017-0253-7.
- [486] Taylor, L., Floridi, L., and Van der Sloot, B. *Group privacy: New challenges of data technologies*. Springer, 2016. ISBN: 978-3-319-46608-8.
- [487] Kayser-Bril, N. *Automated discrimination: Facebook uses gross stereotypes to optimize ad delivery*. 2020. URL: <https://algorithmwatch.org/en/automated-discrimination-facebook-google/> (visited on 02/22/2022).
- [488] AlgorithmWatch and Open Knowledge Foundation. *OpenSCHUFA*. 2018. URL: <https://openschufa.de/english/> (visited on 02/22/2022).
- [489] Jungblut, P. *Böhmernann über Wahl-Werbung: "Facebook fackelt Demokratie ab"*. German. 2021. URL: <https://www.br.de/nachrichten/kultur/boehmermann-ueber-wahl-werbung-facebook-fackelt-demokratie-ab,SjvZKiF> (visited on 02/22/2022).
- [490] Weizenbaum Institute for the Networked Society. *Jahresbericht 2019/2020*. German. Tech. rep. Berlin, 2021. URL: https://www.weizenbaum-institut.de/media/Publikationen/Jahresberichte/Jahresbericht_2019_20_DE.pdf.

-
- [491] Clarke, L., Williams, O., and Swindells, K. “How Google quietly funds Europe’s leading tech policy institutes”. In: *New Statesman* (July 2021). URL: <https://www.newstatesman.com/science-tech/big-tech/2021/07/how-google-quietly-funds-europe-s-leading-tech-policy-institutes> (visited on 11/16/2021).
 - [492] Kayser-Bril, N. *AlgorithmWatch forced to shut down Instagram monitoring project after threats from Facebook*. 2021. URL: <https://algorithmwatch.org/en/instagram-research-shut-down-by-facebook/> (visited on 11/16/2021).
 - [493] Gilbert, D. “Facebook Just Suspended the Accounts of Some of Its Biggest Critics”. In: (2021). URL: <https://www.vice.com/en/article/n7bkg8/facebook-just-suspended-the-accounts-of-some-of-its-biggest-critics>.
 - [494] AlgorithmWatch. *Under Facebook’s thumb: Platforms must stop suppressing public interest research*. 2021. URL: <https://algorithmwatch.org/en/defend-public-interest-research-on-platforms/> (visited on 11/16/2021).
 - [495] Hert, P. de et al., eds. *The Unaccountable State of Surveillance: Exercising Access Rights in Europe*. Issues in Privacy and Data Protection. Cham: Springer, 2017. ISBN: 978-3-319-47573-8. DOI: 10.1007/978-3-319-47573-8.
 - [496] Mishra, C. and Sharma, A. M. “A review paper on voice analytics”. In: *International Journal of Science Technology and Management* 5.12 (2016), pp. 247–257.
 - [497] Narayanan, A. et al. “Application of Raw Accelerometer Data and Machine-Learning Techniques to Characterize Human Movement Behavior: A Systematic Scoping Review”. In: *Journal of Physical Activity and Health* 17.3 (2020), pp. 360–383. DOI: 10.1123/jpah.2019-0088.
 - [498] Khalil, R. A. et al. “Speech Emotion Recognition Using Deep Learning Techniques: A Review”. In: *IEEE Access* 7 (2019), pp. 117327–117345. DOI: 10.1109/ACCESS.2019.2936124. URL: <https://ieeexplore.ieee.org/document/8805181/>.
 - [499] Castell-Uroz, I., Sole-Pareta, J., and Barlet-Ros, P. “TrackSign: Guided Web Tracking Discovery”. In: 2021. DOI: 10.1109/INFOCOM42981.2021.9488842.
 - [500] Guarino, A. et al. “On Analyzing Third-party Tracking via Machine Learning.” in: *Proceedings of the International Conference on Information Systems Security and Privacy*. Scitepress, 2020, pp. 532–539. DOI: 10.5220/0008972005320539.
 - [501] Chen, Q. et al. “Cookie Swap Party: Abusing First-Party Cookies for Web Tracking”. In: *Proceedings of the Web Conference*. ACM, 2021, pp. 2117–2129. DOI: 10.1145/3442381.3449837.