Laurens R. Krol, Juliane Pawlitzki, Mahta Mousavi, Lena M. Andreessen,
Thorsten O. Zander

Saliency versus valence in implicit cursor control: First indications of separate cortical processes

Conference paper | Accepted manuscript (Postprint)

This version is available at <https://doi.org/10.14279/depositonce-11334>

Krol, L. R., Pawlitzki, J., Mousavi, M., Andreessen, L. M. & Zander, T. O. (2019). Saliency versus Valence in Implicit Cursor Control: First Indications of Separate Cortical Processes. 2019 IEEE International Conference on Systems, Man and Cybernetics (SMC), 3913–3918.
<https://doi.org/10.1109/smc.2019.8913936>

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Terms of Use

Copyright applies. A non-exclusive, non-transferable and limited right to use is granted. This document is intended solely for personal, non-commercial use.

**WISSEN IM ZENTRUM
UNIVERSITÄTSBIBLIOTHEK**

Technische
Universität
Berlin

Salience versus Valence in Implicit Cursor Control: First Indications of Separate Cortical Processes

Laurens R. Krol^{1,*}, Juliane Pawlitzki¹, Mahta Mousavi², Lena M. Andreessen¹ and Thorsten O. Zander^{1,3}

Abstract—It has previously been shown that passive brain-computer interfacing enables implicit control over a cursor: participants were able to guide a cursor towards a target without being aware of doing so. The control was based on their brain’s automatic responses to the cursor’s autonomous movements. This raises questions with respect to informed consent and privacy of thought. The extent to which these ethical issues are truly relevant depends on the type of cognitive processes targeted by implicit control paradigms: did the cursor indeed have access to the participants’ subjective preferences? It has been argued that the relevant neurophysiological processes may be based on visual salience rather than on the participants’ personal interpretations. We now present a paradigm that aims to differentiate between salience and valence. Based on data from eight initial participants, we present findings that indicate both processes play a role, meaning that valence is not solely responsible, but indeed reflected in the signal.

I. INTRODUCTION

Passive brain-computer interfaces (passive BCIs; pBCIs [1]) use brain activity as input to applications without that input being voluntarily generated or modulated by the user. Instead, the application adapts automatically to naturally occurring user states which are interpreted in real time. A common example of a passive BCI application assesses a user’s mental load while performing a task, in order to support the user (e.g. [2], [3], [4]).

Beyond mental load, passive BCI research demonstrates an increasing ability of pBCI systems to respond intelligently to various types of obtained input [5]. In particular, passive BCI has been shown to enable *implicit control* [6], [7], [8], [9]. Notably, participants have been able to guide a computer cursor towards its desired target without being aware of doing so [7], [8]. This cursor was controlled by the computer and initially moved randomly over the nodes of a grid. Following each cursor movement, the electroencephalogram (EEG) of the observing participant revealed systematic, seemingly involuntary brain activity that could be used to decode the participant’s response to each movement. The computer could thus assess in real time whether or not the cursor’s movements were in line with the goal of the observer, i.e., to reach the target. In the cited experiment, the computer used this information to steer the cursor towards the decoded target. As a result, the participants—who believed themselves to be in a merely observant role—implicitly controlled the cursor.

This highlights potential advantages of pBCI-based implicit interaction, e.g. for effortless control, but also potential dangers. Since the implicit input is not under voluntarily control, this type of neuroadaptive technology poses a danger to informed consent and the privacy of thought [10]. Therefore, it is important to investigate what information exactly can be decoded from EEG during such experiments, and to what extent such information potentially reflects personal opinions and values as opposed to more neutral cognitive processes.

In this paper, we adapt the implicit cursor control experiment [7], [8] in order to further investigate the relevant neurophysiology. The assumption of the original experiment was that the relevant signal isolated by the classifier, i.e. the signal that distinguished ‘good’ and ‘bad’ cursor movements, reflected predictive coding. The framework of predictive coding holds that the human brain performs continuous, automatic prediction of future (neuronal) events, and continuously compares those predictions with their corresponding final perceptions [11], [12]. The assumption was that the participants, having no other cues, predicted that the cursor would go towards their desired target. Deviations from that prediction resulted in a detectable error-related negativity [13], which furthermore scaled linearly with the intensity of the error. This signal allowed the classifier to categorise the cursor’s movements accordingly.

Interpreted as such, predictive coding is a neutral cognitive process that merely produces predictions based on available information. In the original experiment, the desired target was visually highlighted, thus making it a source of information even without further interpretation as to whether reaching this target would be ‘good’ or ‘bad’—other potential targets were not highlighted the same way.

However, a more nuanced perspective has been suggested to interpret the detected event-related potentials (ERPs) [14]. Holroyd and colleagues make a distinction between surprise- or salience-related signals on the one hand, the response of which is modulated by the expectancy of events [13], and the *reward prediction error* signal which responds most strongly to unexpected (i.e. surprising, salient) events but whose sign is inverted for positive versus negative events [15]. This has primarily been investigated in the context of feedback, where a positive event would for example indicate correct performance or monetary reward.

Thus, although the original implicit cursor control experiment may not have contained directly comparable feedback as such, it is worth investigating to what extent the signal that allowed classification may also have “[manifested] as a

¹Team PhyPA, Biological Psychology and Neuroergonomics, Technische Universität Berlin, Berlin, Germany

²Electrical and Computer Engineering, University of California San Diego, La Jolla, CA, USA

³Zander Laboratories B.V., Amsterdam, the Netherlands

*Corresponding author, lrkrol@gmail.com

specific interaction between outcome valence and outcome probability” [16]. However, the original experiment did not allow these two to be differentiated, because salience and valence of the cursor’s movements were coupled. Therefore, we here present a similar but adapted experiment where the visual stimuli and the participant’s task have been designed to allow such a differentiation to be made.

In this design, visual stimuli were left constant between a ‘negative’ and a ‘positive’ condition. This ensured that any salience-related processing, to the extent that this processing is based on the visually presented information, remains constant between conditions. The valence of the visually identical stimuli was manipulated by changing the participant’s task between conditions.

If a separate valence-related signal can indeed be identified in the context of implicit cursor control, and meaningfully exploited by a classifier, this would underscore the earlier-reported potential issues with implicit control.

Here we report initial findings based on the first eight participants of the experiment.

II. METHODS

A. Participants and Equipment

Eight participants (aged 28.3 ± 5.1 years, five female) participated in the experiment. 64-channel EEG was recorded at 5000 Hz using BrainAmp DC amplifiers (Brain Products GmbH, Gilching, Germany), with electrodes arranged according to the international 10-20 system, referenced to FCz. Participants were seated in front of a 27” computer display placed approximately 1 m away from them.

B. Experimental Paradigm

The stimuli closely followed the original implicit cursor control experiment [7], [8]. The main difference was the different grid layout.

Fig. 1. The grid seen by the participants. The centre node is visually highlighted. The cursor is at one of the four possible starting positions.

Participants were shown the grid illustrated in figure 1, with the centre node visually highlighted to be the reference point. At the start of each grid run, the cursor was placed two nodes away from the centre in a horizontal or vertical line. The cursor would then move every three seconds to one of the adjacent nodes, in one of up to eight different directions. The exact moment of each movement could be predicted due to the consistent timing and the same animations as used in the original experiment. Participants were instructed to observe these movements and evaluate with a button press each individual movement as ‘appropriate’ or ‘not appropriate’ with respect to the current goal.

There were two conditions, which we refer to as ‘positive’ and ‘negative’ respectively. The instructions for the positive condition reflected those of the original experiment: Participants were informed that the goal was for the cursor to reach the centre as quickly as possible. Reaching the centre was to be seen as a success; if this had not been achieved after 50 movements, this was to be seen as a failure.

In the negative condition, the goal was for the cursor to stay away from the centre as long as possible. Not reaching the centre after 25 movements was to be seen as a success; reaching the centre was to be seen as a failure. The maximum number of movements per grid was lower in this condition to approximately equalise centre hit ratios.

Thus, keeping visual stimuli identical, these two conditions inverted the meaning of the stimuli. What was ‘appropriate’ in one condition was ‘inappropriate’ in the other.

The otherwise square grid’s edge nodes were removed so as not to create any nodes that may have been perceived to be further away from the centre than others. Such ‘mental opposites’ could have been used by participants to mentally invert the task, for example turning ‘stay away from the centre’ into ‘go to the corners’.

The two conditions were presented within-subjects in two counterbalanced blocks with a break in between. In each block, participants saw and evaluated 800 movements. After every success or failure, a new grid would be started. Short breaks were given between grids.

The cursor moved randomly throughout the experiment. There was no online condition where real-time classifier output influenced the cursor’s behaviour.

C. Feature Extraction and Classification

The recorded EEG data was resampled to 250 Hz and band-pass filtered from 1 to 15 Hz using an FIR filter with a Kaiser window and 250 taps. Features were extracted using the windowed means approach [17], taking the mean amplitudes of eight non-overlapping consecutive time windows of 50 ms each, between 100 and 500 ms following each cursor movement. Regularised linear discriminant analysis (LDA; [18]) was used for classification. The regularisation parameter was calculated to shrink the sample covariance towards a diagonal matrix (i.e. identity times the average of the diagonal elements of the sample covariance) [19]. A 5-fold cross-validation was used to generate estimates of the classifier’s accuracy. When classes had an uneven number of trials, they

were balanced before each calculation by randomly discarding trials from the larger class. The given accuracy estimates are mean values of 10 subsampling repetitions of the larger class.

Class membership of each trial was determined by that trial's cursor movement's angular deviance from a straight path towards the centre. Thus, an angular deviance of 0° indicates a movement leading straight towards the centre, and 180° leads the cursor away. Due to the different layout of the grid, we re-examined the class definitions used in the original experiment. Since no ground truth is available as to what participants would implicitly categorise as 'appropriate' or 'inappropriate', or even 'towards' and 'away', classifier accuracies were calculated for a number of different class definitions. Based on this, a single definition that performed well across conditions was selected for further analyses. (The original experiment has shown that implicit and explicit responses do not necessarily agree. Therefore, an analysis of the participants' explicit responses, the button presses, is not within the scope of the current paper.)

Estimated classification accuracies are given for both the original raw EEG data and ICA-pruned data described next.

D. Independent Component Analysis

The raw EEG data was first resampled to 250 Hz and high-pass filtered at 1 Hz using a Hamming-windowed sinc FIR filter with the -6 dB cutoff point at 0.5 Hz. Bad channels were removed based on visual inspection aided by `clean_rawdata` with a correlation criterion of 0.85, from EEGLAB 14.1.2 [20]. Removed channels were interpolated before re-referencing all channels to the common average. An independent component analysis (ICA; [21]) decomposition was calculated using AMICA 1.5, and dipoles were fitted to this decomposition using DIPFIT 2.x.

Components reflecting cortical and non-cortical activity were identified following the procedures laid out in [22] supported by the ICLabel plug-in [23]. Non-cortical components were discarded from further analyses (see section III).

E. Component Selection, Clustering, and Statistics

For each time window used by the LDA classifier, the corresponding LDA patterns were calculated for each participant and condition using the method published in [24]. These patterns represent the projection of the signal that is isolated by the classifier in order to distinguish between the classes. Just as with scalp EEG data itself, it is possible to 'unmix' these patterns using the previously calculated ICA decomposition's unmixing matrix. The resulting *relevance weights* represent the relative contribution of each independent component (IC) to the pattern, and thus, the relative contribution of each IC to classification in the given time window. This method is described in detail in [25].

It is generally observed that only a small percentage of ICs contribute significantly to classification. We therefore only kept ICs with a weight greater than three standard deviations from the mean in any one time window for each participant.

The remaining ICs thus represented cortical components that contributed significantly to classification. These ICs were

then manually clustered based on their scalp topography, dipole location, and ERPs using EEGLAB.

As the classifier focused on scalp ERPs, the resulting clusters were analysed from an ERP perspective as well. After band-pass filtering the data between 1 and 15 Hz, differences between ERPs were calculated using Student's t-tests for the two independent variables: negative versus positive condition, and towards versus away from the centre, using the class definitions found earlier.

III. RESULTS

The optimal angular deviations were $\leq 27^\circ$ for 'towards the centre' and $> 117^\circ$ for 'away from the centre'. On average across conditions and participants, a range of different definitions of which movements were considered to be going 'towards' and 'away from' the centre produced similar classification accuracies. On average, the selected definition differed less than one percentage point from optimally selected angles on a per-participant, per-condition basis.

The cross-validated classification accuracies of the selected definition of the class boundaries are presented in table I. As classes were balanced, chance level was at 50%. On average, there were approximately 165 trials per class, with significance thus being reached at a classification accuracy of approximately 55% or more [26].

Both before and after cleaning, there was a difference of roughly 10 percentage points between positive and negative conditions. Permutation tests with 10000 iterations showed these differences to be significant ($p < 0.01$). Differences between data sets (before and after cleaning) within the same condition were not significant.

TABLE I
CROSS-VALIDATED CLASSIFICATION ACCURACIES.

Participant	Raw Data		Pruned Data	
	Positive	Negative	Positive	Negative
1	77	66	86	74
2	67	66	68	67
3	75	69	80	70
4	81	74	85	80
5	79	65	82	67
6	90	75	92	80
7	78	62	83	68
8	83	62	84	65
Mean:	79	68	83	71
				%

Because removing all eye, muscle, noise, and undetermined ICs did not decrease the separability of the classes, and because we are currently interested only in neurophysiological effects of the paradigm, we chose to continue only with the ICA-pruned data. Out of the 493 ICs, the selection of cortical ICs left 148 for an average of 18.5 per participant. No participant had less than 13 cortical ICs.

We constructed each of these IC's relevance weight using individual classifiers calibrated on this ICA-pruned data. The sorted distribution of all relevance weights for all participants in both conditions is shown in figure 2. The expected pattern where only a small number of ICs contribute significantly to

classification can clearly be observed. Selecting only those with such a significant contribution left 44 ICs in total, or 5.5 per participant, with each participant having at least 3.

Fig. 2. Cortical independent components for each participant and each condition sorted by their computed relevance weight.

These 44 remaining ICs that contributed significantly to class separability showed clear similarities between participants, most clearly visible with respect to their scalp topographies. The three clusters that were produced are illustrated in figure 3 by their IC scalp topographies.

The largest cluster, consisting of 25 ICs from all participants, contains ICs located primarily in the occipital, and to a lesser extent the parietal lobe. On average, these ICs project most strongly onto occipital sites. The second-largest cluster, with 7 ICs from 6 participants, projects most strongly onto central sites; the individual variability in scalp projections and dipole locations is somewhat greater. The third cluster contains ICs located in the frontal lobe, primarily around the medial prefrontal cortex. This cluster contains 6 ICs from 5 participants. Two remaining ICs do not match any of the other clusters' properties and have thus remained unassigned.

The grand-average ERPs of the 'occipital' and 'frontal' clusters are shown in figure 4. They have been separated by the two independent variables under investigation: condition (positive versus negative) and movement class (towards versus away from the centre).

The statistics show that the occipital cluster consistently exhibits significant differences between movement classes prior to 200 ms, but no consistent significant differences between condition. The frontal cluster consistently exhibits significant differences between condition around 400 ms, as well as between movement classes after 200 ms.

Among others, further significant effects are seen in the occipital cluster between conditions, but only in one class.

The central cluster is not further elaborated here. It showed no significant effects, except for a section around 250 ms following cursor movement limited to the positive condition. It is possible that activity related to the eye muscles are included in this cluster.

IV. DISCUSSION

We adapted the original implicit cursor control experiment [7], [8] to be able to independently manipulate salience and

valence. These two aspects were coupled in the original experiment, although it has been argued that they may contribute independently to the identified signal [14].

By comparing cursor movements that went towards and away from the centre, with the centre thus serving as the one and only salient visual reference point, we could investigate visual salience-related cortical processing across conditions. By keeping these stimuli identical but inverting the valence of the cursor movements between two conditions, we could investigate different affective processing of the same stimuli.

We trained individual classifiers to evaluate the separability of the classes. In both conditions, both classes of movements could be classified with meaningful, above-chance accuracy. However, clear differences were observed between conditions, with the positive condition showing significantly better separability than the negative condition. This rules out a classification signal that is completely driven by visual salience, as this remained constant. We thus assume that the difference between conditions was induced by the different instructions, which inverted the valence of the centre node.

After decomposing the data into independent components using ICA, we focused only on those ICs that contributed significantly to the above-mentioned classification. This was done by translating the classifier's LDA weights into relevance weights for each IC [25]. Those with above-average relevance were clustered into three contributing cortical regions: roughly named an occipital, central, and frontal cluster. The occipital cluster exhibited consistent significant differences between salience-related classes, whereas the frontal cluster exhibited consistent significant differences between valence-related conditions. This supports the hypothesis that these are indeed two separate contributing cortical processes.

However, the effects are not entirely limited to these two results: the frontal cluster also shows differences between classes, albeit at a different time compared to the occipital cluster. This was, however, expected, as the classes of course also represent a difference in valence. The exact timing remains to be investigated. We also do not currently have an explanation for the occipital cluster's significant differences between conditions only for movements away from the centre. The role of the central cluster was inconclusive.

These are the initial results from eight participants. A number of shortcomings must be mentioned.

First of all, the number of participants is relatively low. This is also reflected in the fact that not all eight participants are represented in all clusters. We are currently recording additional participants to increase the results' power.

Additional data should also increase the clustering reliability. EEGLAB's k-means clustering does not result in consistent, unique solutions [27]. Here, we used manual clustering and presented all scalp maps for inspection in figure 3. With additional data, repeated k-means clustering and automatic cluster selection [27] should result in a more

A remaining difficulty concerns the definition of salience. We have here argued with respect to visual salience, i.e., the fact that the centre was the one and only node of the grid that was visually and conceptually more prominent than the others.

It is possible, however, that participants mentally shifted their attention to one of the other nodes [28], thus internally altering its salience. We removed the most obvious candidate nodes for such covert attention shifts by removing the corner nodes. Furthermore, the task is defined with respect to the centre node, thus making this the relevant reference node in either case. Still, we cannot fully control the participants' internal processes. This, too, will be revisited when additional data, including gaze, is available.

As it stands, we have found first indications that separate salience and valence processes play a role in the implicit cursor control paradigm. Although when counted by sheer number the salience ICs form a majority, a consistently valence-related cluster has been identified.

It is thus conceivable that neuroadaptive technology may be able to separately detect and target valence-related brain activity, e.g. using cognitive probes [29]. This may allow personal values to be inferred. While such information can be used productively for e.g. personalisation or productive support [30], it should be handled with due care for data privacy, reinforcing the need for ethical considerations both in research and commercial applications of implicit control and neuroadaptive technology.

ACKNOWLEDGMENT

LRK and JP were supported by the Deutsche Forschungsgemeinschaft (ZA 821/3-1); MM by NIH 5T32MH020002-18 and the Mary Anne Fox Dissertation Year Fellowship.

REFERENCES

- [1] T. O. Zander and C. A. Kothe, "Towards passive brain-computer interfaces: applying brain-computer interface technology to human-machine systems in general," *Journal of Neural Engineering*, vol. 8, no. 2, p. 025005, 2011.
- [2] B. F. Yuksel, K. B. Oleson, L. Harrison, E. M. Peck, D. Afergan, R. Chang, and R. J. K. Jacob, "Learn piano with BACH: An adaptive learning interface that adjusts task difficulty based on brain state," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ser. CHI '16. New York, NY, USA: ACM, 2016, pp. 5372–5384.
- [3] K. C. Ewing, S. H. Fairclough, and K. Gilleade, "Evaluation of an adaptive game that uses EEG measures validated during the design process as inputs to a biocybernetic loop," *Frontiers in Human Neuroscience*, vol. 10, p. 223, 2016.
- [4] T. O. Zander, K. Shetty, R. Lorenz, D. R. Leff, L. R. Krol, A. W. Darzi, K. Gramann, and G.-Z. Yang, "Automated task load detection with electroencephalography: Towards passive brain-computer interfacing in robotic surgery," *Journal of Medical Robotics Research*, vol. 2, no. 1, p. 1750003, 2017.
- [5] L. R. Krol, L. M. Andreessen, and T. O. Zander, "Passive Brain-Computer Interfaces: A Perspective on Increased Interactivity," in *Brain-Computer Interfaces Handbook: Technological and Theoretical Advances*, C. S. Nam, A. Nijholt, and F. Lotte, Eds. Boca Raton, FL, USA: CRC Press, 2018, pp. 69–86.
- [6] M. Rötting, T. O. Zander, S. Trösterer, and J. Dzaack, "Implicit interaction in multimodal human-machine systems," in *Industrial Engineering and Ergonomics*, C. M. Schlick, Ed. Berlin Heidelberg, Germany: Springer, 2009, pp. 523–536.
- [7] T. O. Zander, J. Brönstrup, R. Lorenz, and L. R. Krol, "Towards BCI-based Implicit Control in Human-Computer Interaction," in *Advances in Physiological Computing*, S. H. Fairclough and K. Gilleade, Eds. Berlin, Germany: Springer, 2014, pp. 67–90.
- [8] T. O. Zander, L. R. Krol, N. P. Birbaumer, and K. Gramann, "Neuroadaptive technology enables implicit cursor control based on medial prefrontal cortex activity," *Proceedings of the National Academy of Sciences*, vol. 113, no. 52, pp. 14 898–14 903, 2016.
- [9] I. Iturrate, R. Chavarriaga, L. Montesano, J. Minguez, and J. d. R. Millán, "Teaching brain-machine interfaces as an alternative paradigm to neuroprosthetics control," *Scientific Reports*, vol. 5, p. 13893, 2015.
- [10] G. Mecacci and P. Haselager, "Identifying criteria for the evaluation of the implications of brain reading for mental privacy," *Science and Engineering Ethics*, 2017.
- [11] H. Brown, K. J. Friston, and S. Bestmann, "Active inference, attention and motor preparation," *Frontiers in Psychology*, vol. 2, no. 218, 2011.
- [12] A. Clark, "Whatever next? Predictive brains, situated agents, and the future of cognitive science," *Behavioral and Brain Sciences*, vol. 36, pp. 181–204, 2013.
- [13] C. B. Holroyd and M. G. H. Coles, "The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity," *Psychological Review*, vol. 109, no. 4, pp. 679–708, 2002.
- [14] J. Cockburn and C. B. Holroyd, "Feedback information and the reward positivity," *International Journal of Psychophysiology*, vol. 132, pp. 243–251, 2018.
- [15] T. D. Sambrook and J. Goslin, "A neural reward prediction error revealed by a meta-analysis of ERPs using great grand averages," *Psychological Bulletin*, vol. 141, no. 1, pp. 213–235, 2015.
- [16] S. Heydari and C. B. Holroyd, "Reward positivity: Reward prediction error or salience prediction error?" *Psychophysiology*, vol. 53, no. 8, pp. 1185–1192, 2016.
- [17] B. Blankertz, S. Lemm, M. S. Treder, S. Haufe, and K.-R. Müller, "Single-trial analysis and classification of ERP components – a tutorial," *NeuroImage*, vol. 56, no. 2, pp. 814–825, 2011.
- [18] C. M. Bishop, *Pattern Recognition and Machine Learning*, ser. Information Science and Statistics, M. Jordan, J. Kleinberg, and B. Schölkopf, Eds. New York, NY: Springer, 2006.
- [19] J. Schäfer and K. Strimmer, "A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics," *Statistical Applications in Genetics and Molecular Biology*, vol. 4, no. 1, p. 32, 2008.
- [20] A. Delorme and S. Makeig, "EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis," *Journal of Neuroscience Methods*, vol. 134, no. 1, pp. 9–21, 2004.
- [21] S. Makeig, T.-P. Jung, A. J. Bell, and T. J. Sejnowski, "Independent component analysis of electroencephalographic data," in *Advances in Neural Information Processing Systems*, D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, Eds. Cambridge, MA, USA: MIT Press, 1996, vol. 8, pp. 145–151.
- [22] T.-P. Jung, S. Makeig, C. Humphries, T.-W. Lee, M. J. McKeown, V. Iragui, and T. J. Sejnowski, "Removing electroencephalographic artifacts by blind source separation," *Psychophysiology*, vol. 37, no. 2, pp. 163–178, 2000.
- [23] L. Pion-Tonachini, K. Kreutz-Delgado, and S. Makeig, "ICLabel: An automated electroencephalographic independent component classifier, dataset, and website," *NeuroImage*, 2019.
- [24] S. Haufe, F. Meinecke, K. Görgen, S. Dähne, J.-D. Haynes, B. Blankertz, and F. Bießmann, "On the interpretation of weight vectors of linear models in multivariate neuroimaging," *NeuroImage*, vol. 87, no. 0, pp. 96–110, 2014.
- [25] L. R. Krol, M. Mousavi, V. R. de Sa, and T. O. Zander, "Towards classifier visualisation in 3D source space," in *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2018, pp. 71–76.
- [26] G. R. Müller-Putz, R. Scherer, C. Brunner, R. Leeb, and G. Pfurtscheller, "Better than random? A closer look on BCI results," *International Journal of Bioelectromagnetism*, vol. 10, no. 1, pp. 52–55, 2008.
- [27] M. Klug, L. Gehrke, F. U. Hohlefeld, and K. Gramann, "The BeMoBI Pipeline—Facilitating Mobile Brain/Body Imaging (MoBI) Data Analysis in MATLAB," in *Proceedings of the 3rd International Mobile Brain/Body Imaging Conference*, 2018, pp. 131–132.
- [28] M. I. Posner, "Orienting of attention," *Quarterly Journal of Experimental Psychology*, vol. 32, no. 1, pp. 3–25, 1980.
- [29] L. R. Krol and T. O. Zander, "Towards a conceptual framework for cognitive probing," in *Symbiotic Interaction*, J. Ham, A. Spagnoli, B. Blankertz, L. Gamberini, and G. Jacucci, Eds. Cham, Switzerland: Springer International Publishing, 2018, pp. 74–78.
- [30] M. Mousavi, A. S. Koerner, Q. Zhang, E. Noh, and V. R. de Sa, "Improving motor imagery BCI with user response to feedback," *Brain-Computer Interfaces*, vol. 4, no. 1-2, pp. 74–86, 2017.

Fig. 3. The 44 ICs that were most relevant to classification shown by their scalp topographies, clustered into occipital (left), frontal (upper right) and central (lower right) clusters. Two ICs remained unassigned (bottom centre). The larger topographies are cluster means.

Fig. 4. Blue graphs: Grand-average ERPs of the occipital (left) and frontal (right) clusters separated by condition (negative/positive) and movement class (away/towards). Red graphs: FDR-corrected p-values calculated on the differences between conditions and movement class. Negative: Cursor movements in the 'negative' condition; movements towards the centre were undesirable, away from the centre desirable. Positive: Cursor movements in the 'positive' condition; movements towards the centre were desirable, away from the centre undesirable. Away: Cursor movements that went away from the target. Towards: Cursor movements that went towards the target.