
Deformation and Transport of Image Data

vorgelegt von

Sebastian Jonas Neumayer (M.Sc.)

an der Fakultät II - Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
- Dr.rer.nat. -

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr. John M. Sullivan (TU Berlin)

Gutachterin: Prof. Dr. Gabriele Steidl (TU Berlin)

Gutachter: Prof. Dr. Alain Trouné (ENS Cachan)

Gutachter: Prof. Dr. Benedikt Wirth (WWU Münster)

Tag der wissenschaftlichen Aussprache: 17.12.2020

Berlin 2021

Acknowledgments

I would like to thank ...

- ... my supervisor Gabriele Steidl for her great guidance and constant support during my time as a PhD student and research assistant both in Kaiserslautern and Berlin,
- ... Alain Trouvé and Benedikt Wirth for serving as a referee for this thesis,
- ... my current colleagues Robert Beinert, Manuel Gräf, Marzieh Hasannasab, Johannes Hertrich and Johannes von Lindheim for a great time during my PhD and many fruitful discussions,
- ... my former colleagues Ronny Bergmann, Jan Henrik Fitschen, Friederike Laus, Max Nimmer, Johannes Persch for enlightening discussions and making my time in Kaiserslautern very memorable,
- ... my coauthors Miroslav Bačák, Alexander Effland, Martin Ehler, Péter Koltai, Lukas Lang, Ozan Öktem, Gerlind Plonka, Martin Rumpf, Carola-Bibiane Schönlieb and Simon Setzer for their collaboration and many fruitful discussions on various topics,
- ... for proofreading: Robert Beinert, Marcel Brusius, Johannes Hertrich and Johannes von Lindheim,
- ... my family, friends and my dear Natalie for all the support over the last years.

Moreover, I would like to acknowledge funding by the “German Research Foundation” (DFG) within the research training group 1932 “Stochastic Models for Innovations in the Engineering Sciences”, project “Stochastic and Analytical Methods from Image Processing for the Characterization of Local Properties of Multiphase Materials” and further financial support by the “Deutsch-Französische Hochschule” for various workshops and meetings.

Abstract

Image and data processing have become increasingly important over recent years. In particular, robust distance measures and transformation methods on large and perturbed data sets are of growing interest. These tasks become even more challenging if the measurement locations are distributed on some manifold or if the data itself has a manifold structure.

In this thesis, we start by examining a time discrete manifold-valued morphing model together with a time continuous counterpart which includes a variational inequality and is inspired by the classical metamorphosis approach. For both models, existence of minimizers is shown under the assumption that the data lies in a finite-dimensional Hadamard manifold. Further, we prove Mosco convergence of the time discrete model to the time continuous one.

Then, we propose to use the real-valued version of this model as variational regularizer for inverse problems if some template image is available. As this incorporates additional information into the reconstruction process, such an approach is particularly useful if the data is sparse. In cases where an even stronger regularization is necessary, we propose to use a more restrictive flow of diffeomorphism based model instead. Additionally, increased robustness to intensity differences between template and ground-truth is achieved by using a distance measure based on normalized cross-correlation. These modifications enable us to obtain promising reconstructions for computed tomography with a very sparse number of angles. For both models existence, stability and convergence for vanishing noise are proven.

Besides metamorphosis and flow of diffeomorphism, we are interested in optimal transport between images. More precisely, we investigate Sinkhorn divergences, which interpolate between optimal transport and discrepancies. Here, we focus on the behavior of the corresponding optimal dual potentials and establish a relation between their limits as the regularization parameter goes to infinity and the optimal potential of the associated discrepancy. Additionally, we compare these different distance measures for a process called dithering.

Finally, we investigate the approximation of probability measures on compact metric spaces and in particular on Riemannian manifolds by measures supported on Lipschitz curves. For this purpose, the approximation quality is assessed in terms of discrepancies as these do not suffer from the so-called curse of dimensionality. While we also investigate the general case via the traveling salesman approach, a strong focus lies on the push-forward of the Lebesgue measure on the unit interval by Lipschitz continuous curves. Here, we utilize a result on the quadrature error achievable by integration with respect to a measure that exactly integrates

all eigenfunctions of the Laplace–Beltrami operator with eigenvalues smaller than a fixed number. In particular, this enables us to provide optimal approximation rates in terms of the curve’s length and Lipschitz constant for measures fulfilling smoothness requirements.

All results established in this thesis are illustrated by numerical proof-of-concept examples and comparisons. These include real-world data experiments in particular for our flow of diffeomorphism based reconstruction model. For the necessary discrepancy evaluations, recent fast Fourier transform techniques on certain manifolds are exploited.

Zusammenfassung

Bild- und Datenverarbeitung haben in den letzten Jahren zunehmend an Bedeutung gewonnen. Dabei haben sich insbesondere robuste Abstandsmaße und Transformationsmethoden für große und gestörte Datensätze als wichtig erwiesen. Diese Aufgaben werden noch komplizierter, wenn die Messpunkte auf einer Mannigfaltigkeit verteilt sind oder wenn die Daten selbst eine Mannigfaltigkeitsstruktur besitzen.

In dieser Arbeit untersuchen wir zunächst ein zeitdiskretes Morphing-Modell für mannigfaltigkeitswertige Bilder zusammen mit einem zeitkontinuierlichen Pendant, welches auf einer Variationsungleichung basiert und vom klassischen Metamorphosis Modell inspiriert ist. Für beide Ansätze wird die Existenz von Minimierern unter der Annahme gezeigt, dass die Daten in einer endlichdimensionalen Hadamard-Mannigfaltigkeit liegen. Außerdem beweisen wir die Mosco-Konvergenz des zeitdiskreten Modells gegen das zeitkontinuierliche Modell.

Falls ein geeignetes Referenzbild verfügbar ist, wollen wir die reellwertige Version dieses Modells als Regularisierer für das Lösen inverser Probleme verwenden. Da hierdurch zusätzliche Information in den Rekonstruktionsprozess eingebunden wird, ist ein solcher Ansatz insbesondere dann nützlich, wenn die Daten unvollständig sind. Ist eine noch stärkere Regularisierung erforderlich, verwenden wir stattdessen ein restriktiveres und auf dem Fluss des Diffeomorphismus basierendes Modell. Zusätzlich kann die Robustheit gegenüber Intensitätsunterschieden zwischen Referenzbild und dem unbekannten Signal durch die Verwendung eines auf der normalisierten Kreuzkorrelation basierenden Abstandsmaßes verbessert werden. Diese Modifikationen ermöglichen es uns, vielversprechende Rekonstruktionen für Computertomografie mit einer sehr geringen Anzahl von Winkeln zu erhalten. Für beide Modelle werden Existenz, Stabilität und Konvergenz für verschwindendes Rauschen gezeigt.

Neben dem Metamorphosis und dem Fluss des Diffeomorphismus Modell sind wir auch an optimalem Transport zwischen Bildern interessiert. Genauer gesagt untersuchen wir Sinkhorn-Divergenzen, die zwischen optimalem Transport und Diskrepanzen interpolieren. In dieser Arbeit konzentrieren wir uns auf das Verhalten der entsprechenden optimalen dualen Potentiale und stellen eine Beziehung zwischen ihren Grenzwerten für wachsende Regularisierungsparameter und dem optimalen Potential der zugehörigen Diskrepanz her. Zusätzlich vergleichen wir die vorgestellten Abstandsmaße für einen Prozess namens “Dithering”.

Schließlich untersuchen wir die Approximation von Wahrscheinlichkeitsmaßen auf kompakten metrischen Räumen und insbesondere auf Riemannschen Mannigfaltigkeiten durch Maße, deren Träger eine lipschitzstetige Kurve ist. Zu diesem Zweck

wird die Approximationsqualität anhand von Diskrepanzen beurteilt, da diese nicht vom sogenannten Fluch der Dimension betroffen sind. Während wir auch den allgemeinen Fall über das Problem des Handlungsreisenden untersuchen, liegt der Fokus auf Bildmaßen vom Lebesgue-Maß auf dem Einheitsintervall bezüglich lipschitzstetiger Kurven. Hierbei verwenden wir ein Ergebnis zum Quadraturfehler bezüglich der Integration mit Maßen, welche Eigenfunktionen des Laplace–Beltrami Operators mit Eigenwerten kleiner einer festen Zahl exakt integrieren. Insbesondere können wir dadurch optimale Approximationsraten in Abhängigkeit der Kurvenlänge und der Lipschitz-Konstante für Maße zeigen, die zusätzliche Glattheitseigenschaften erfüllen.

Alle in dieser Arbeit erzielten Ergebnisse werden durch numerische Konzeptnachweise und Vergleiche sowie Experimente mit realen Daten insbesondere für unser auf dem Fluss des Diffeomorphismus basierenden Rekonstruktionsmodell veranschaulicht. Für die notwendigen Diskrepanzberechnungen werden vor kurzem entwickelte schnelle Fourier-Transformationstechniken auf ausgewählten Mannigfaltigkeiten genutzt.

Contents

List of Publications	xi
1 Introduction	1
2 Morphing of Manifold-Valued Images	11
2.1 Introduction	13
2.2 Preliminaries	15
2.2.1 Hadamard spaces	15
2.2.2 Sobolev spaces and admissible mappings	19
2.3 Minimizers of the space continuous model	20
2.3.1 Space continuous model	20
2.3.2 Model specification	28
2.4 Minimization of the space discrete model	30
2.4.1 Computation of the deformation sequence	30
2.4.2 Computation of the image sequence	35
2.4.3 Multiscale minimization scheme	36
2.5 Numerical examples	36
2.5.1 Images in different color spaces	37
2.5.2 Symmetric positive definite matrices $\mathcal{P}(n)$	38
2.6 Conclusions	41
3 Convergence of the Time Discrete Metamorphosis Model	43
3.1 Introduction	45
3.2 Review and preliminaries	47
3.2.1 Hadamard manifolds	48
3.2.2 Metamorphosis model in Euclidean case	51
3.2.3 Manifold-valued time discrete metamorphosis model	55
3.3 Manifold-valued metamorphosis model	57
3.4 Temporal extension operators	60
3.5 Mosco convergence of time discrete geodesic paths	63
3.6 Conclusion	74
4 Regularization of Inverse Problems via Time Discrete Geodesics	75
4.1 Introduction	77
4.2 Preliminaries	79
4.2.1 Admissible deformations	79

4.2.2	Space of images	80
4.3	Space continuous model	82
4.3.1	Model	82
4.3.2	Existence, stability and convergence	83
4.4	Minimization approaches	90
4.4.1	Spatial discretization	90
4.4.2	Proximal alternating linearized minimization (PALM)	92
4.4.3	Alternating minimization approach	93
4.4.4	Multilevel strategy	95
4.5	Numerical examples	95
4.5.1	Limited angle and sparse CT	97
4.5.2	Superresolution	97
4.6	Conclusions	101
5	Template-Based Image Reconstruction	105
5.1	Introduction	107
5.2	Theoretical results on the transport and continuity equation	111
5.2.1	Continuity equation	111
5.2.2	Transport equation with H^1 regularity	114
5.3	Regularizing properties of template-based image reconstruction	115
5.4	Numerical solution	117
5.5	Numerical examples	122
5.6	Conclusions	128
6	From Optimal Transport to Discrepancy	133
6.1	Introduction	135
6.2	Preliminaries	136
6.3	Discrepancies	139
6.4	Optimal transport and Wasserstein distances	145
6.5	Regularized optimal transport	147
6.6	Sinkhorn divergence	154
6.7	Numerical approach and examples	156
6.8	Conclusions	160
6.A	Basic theorems	161
7	Curve Based Approximation of Measures	163
7.1	Introduction	165
7.2	Probability measures and curves	167
7.3	Discrepancies and RKHS	170
7.4	Approximation of general probability measures	173
7.5	Approximation of probability measures having Sobolev densities	176
7.5.1	Sobolev spaces and lower bounds	177
7.5.2	Upper bounds for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$	178
7.5.3	Upper bounds for $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$ and special manifolds \mathbb{X}	180
7.5.4	Upper bounds for $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ and special manifolds \mathbb{X}	183

7.6	Discretization	188
7.7	Numerical algorithm	189
7.8	Numerical results	192
7.8.1	Parameter choice	192
7.8.2	Quasi-optimal curves on special manifolds	195
7.9	Conclusions	203
7.A	Special manifolds	205
Bibliography		209

List of Publications

Each chapter of this thesis is based on one of the following published articles and preprints. The papers are sorted according to their appearance in this thesis (Chapters 2–7) and the authors are listed in alphabetical order.

- [213] S. Neumayer, J. Persch and G. Steidl. Morphing of Manifold-Valued Images inspired by Discrete Geodesics in Image Spaces. *SIAM Journal on Imaging Sciences*, 11(3):1898–1930, 2018.
Included version: Accepted manuscript
Published version at DOI: 10.1137/17M1150906.
- [98] A. Effland, S. Neumayer and M. Rumpf. Convergence of the Time Discrete Metamorphosis Model on Hadamard Manifolds. *SIAM Journal on Imaging Sciences*, 13(2):557–588, 2020.
Included version: Accepted manuscript
Published version at DOI: 10.1137/19M1247073.
- [214] S. Neumayer, J. Persch and G. Steidl. Regularization of Inverse Problems via Time Discrete Geodesics in Image Spaces. *Inverse Problems*, 35(5):055005, 2019.
Included version: Accepted manuscript
Published version at DOI: 10.1088/1361-6420/ab038a.
- [181] L. Lang, S. Neumayer, O. Öktem and C.-B. Schönlieb. Template-Based Image Reconstruction from Sparse Tomographic Data. *Applied Mathematics & Optimization*, 82:1081–1109, 2020.
Included version: Accepted manuscript
Published version at DOI: 10.1007/s00245-019-09573-2.
- [215] S. Neumayer and G. Steidl. From Optimal Transport to Discrepancy. Accepted for *Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging*, 2020.
Included version: Accepted manuscript
Soon to be published.
- [102] M. Ehler, M. Gräf, S. Neumayer and G. Steidl. Curve Based Approximation of Measures on Manifolds by Discrepancy Minimization. *Foundations of Computational Mathematics*, 2021.
Included version: Accepted manuscript
Published version at DOI: 10.1007/s10208-021-09491-2.

Further publications, which are not contained in this thesis, given in chronological order:

- [144] P. Hagemann, S. Neumayer. Stabilizing Invertible Neural Networks Using Mixture Models. ArXiv Preprint 2009.02994, 2020.
- [178] P. Koltai, J. von Lindheim, S. Neumayer and G. Steidl. Transfer Operators from Optimal Transport Plans for Coherent Set Detection. ArXiv Preprint 2006.16085, 2020.
- [149] C. Hartmann, A. F. Weiss, P. Lechner, W. Volk, S. Neumayer, J. H. Fitschen and G. Steidl. Measurement of Strain, Strain Rate and Crack Evolution in Shear Cutting. *Journal of Materials Processing Technology*, 288:116872, 2021.
- [150] M. Hasannasab, J. Hertrich, S. Neumayer, G. Plonka, S. Setzer and G. Steidl. Parseval Proximal Neural Networks. *Journal of Fourier Analysis and Applications*, 26:59, 2020.
- [14] M. Bačák, J. Hertrich, S. Neumayer and G. Steidl. Minimal Lipschitz and ∞ -Harmonic Extensions of Vector-Valued Functions on Finite Graphs. *Information and Inference: A Journal of the IMA*, 9(4):935–959, 2020.
- [211] S. Neumayer, M. Nimmer, S. Setzer and G. Steidl. On the Rotational Invariant L1-norm PCA. *Linear Algebra and its Applications*, 587:243–270, 2020.
- [210] S. Neumayer, M. Nimmer, S. Setzer and G. Steidl. On the Robust PCA and Weiszfeld’s Algorithm. *Applied Mathematics & Optimization*, 82:1017–1048, 2020.
- [185] J. Lellmann, S. Neumayer, M. Nimmer and G. Steidl. Methods for Finding the Offset in Robust Subspace Fitting. *Proceedings in Applied Mathematics and Mechanics*, 19(1):e201900323, 2019.
- [156] J. Hertrich, M. Bačák, S. Neumayer and G. Steidl. Minimal Lipschitz Extensions for Vector-Valued Functions on Finite Graphs. *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science* 11603, 183–195, Springer, Cham, 2019.
- [212] S. Neumayer, M. Nimmer, G. Steidl and H. Stephani. On a Projected Weiszfeld Algorithm. *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science* 10302, 486–497, Springer, Cham, 2017.

Introduction

The need for processing image data has attracted a lot attention over recent years and the development of fast hardware has lead to many new applications. Typical examples are deformation estimation in image sequences or medical imaging tasks such as object reconstruction from sparse X-ray data. While such applications are easily accessible and understandable for a broad audience, the underneath mathematics is quite involved. Obviously, fast, efficient and reliable approaches for solving such problems are vital and a lively area of research. In this thesis, we focus on imaging models that incorporate explicit knowledge about the particular problem into the solution process, e.g., in form of specifically designed regularization terms. For many problems, such an approach leads to satisfactory solutions with reliable and stable outputs.

Let us formalize our description of images. Assume that some image domain $\Omega \subset \mathbb{R}^n$ is given such as the unit square for example. On this domain, we consider a pair of images $I_A, I_B \in L^2(\Omega)$, i.e., images are interpreted as square integrable functions. In this thesis, we are also interested in two generalizations of this image model. First, as modern data acquisition techniques produce data that is no longer scalar- or vector-valued, we have to deal with images in $L^2(\Omega, \mathcal{H})$, i.e., the image range is some finite-dimensional manifold \mathcal{H} . Applications with manifold-valued data include interferometric synthetic aperture radar (InSAR) [236] with data on the unit circle, electron backscatter diffraction (EBSD) [15] with data in a quotient of the rotation group on \mathbb{R}^3 , or diffusion tensor magnetic resonance imaging (DT-MRI) [20] with symmetric positive definite tensors of size 3×3 as data. An impression of such data is given in Fig. 1.1. All mentioned applications have in common that a treatment with classical Euclidean imaging models is impossible. The second natural generalization is that the image domain is a manifold such as the sphere. This naturally occurs if we take measurements distributed over the complete earth, see Fig. 1.2 for an application in measure approximation by curves as considered in this thesis. Clearly, there are many more applications where measurements are not taken on flat surfaces, e.g., models for aerodynamic properties of cars.

In image processing, two particularly challenging questions are to find “good” distance measures and “nice” transformations between images. Naturally, these

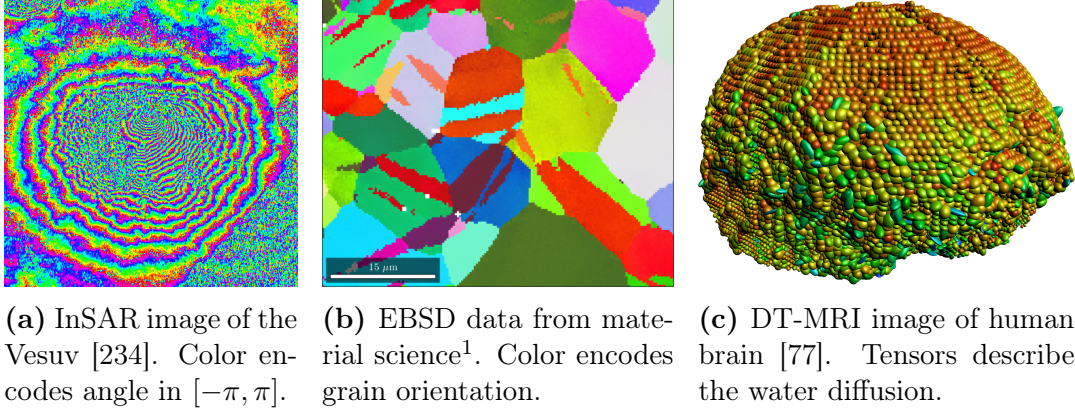


Figure 1.1: Examples of manifold-valued images.

questions are strongly related as the deformation amplitude provides us with a way of measuring distances. Possible applications related to these questions include generative adversarial networks [96], computation of additional frames in movies as it was done for *Willow* [251] or incorporation of template information into sparse computed tomography [116, 208], where the data is limited in order to minimize exposure time of organisms to X-radiation. Regarding the transformation of images, we want to obtain something “natural”, more precisely, the result should look visually appealing for humans. Such a behavior is usually not obtained with a simple pixel-wise linear interpolation between the two images I_A and I_B , which corresponds to geodesics with respect to the standard geometry of the linear space $L^2(\Omega)$. Similarly, the $L^2(\Omega)$ distance does not capture any geometrical information about our images and hence more sophisticated approaches are necessary.


In the following, we aim to give a brief overview over three main topics covered by this thesis. Exact mathematical details are then provided in the corresponding chapters.

Deformation based models Essentially, all such models share the similarity that the image domain is transformed in terms of a diffeomorphism, which usually results in transformations that look more natural to humans. Starting from the flow of diffeomorphism model [71, 95, 261], in which image pixel intensities are transported along trajectories determined by diffeomorphism paths, a whole field of research has emerged. As the space of diffeomorphisms φ has no natural vector space structure, a common choice is to use a subset generated by admissible (smooth) velocity fields v via the ordinary differential equation (ODE)

$$\begin{aligned} \frac{d}{dt}\varphi(t, x) &= v(t, \varphi(t, x)) & \text{for } (t, x) \in [0, 1] \times \Omega, \\ \varphi(0, x) &= x & \text{for } x \in \Omega. \end{aligned} \tag{1.1}$$

In accordance with our informal description, we implicitly define for any $t \in [0, 1]$ a transformed version $I(t, x)$ of the image I_A via $I(t, \varphi(t, x)) = I_A(x)$. This formu-

¹Contained in MTEX 5.6 – A Texture Calculation Toolbox.



(a) Earth's elevation data [134]. (b) Approximation by a curve [102].

Figure 1.2: Approximation of a measure supported on the sphere \mathbb{S}^2 .

lation is usually referred to as Lagrange form of the problem. Equivalently, using the Euler formulation, $I(t, x)$ is given as solution of the transport equation

$$\begin{aligned} \frac{\partial}{\partial t} I(t, x) + v(t, x) \nabla_x I(t, x) &= 0 & \text{for } (t, x) \in [0, 1] \times \Omega, \\ I(0, x) &= I_A(x) & \text{for } x \in \Omega. \end{aligned} \quad (1.2)$$

In the following, let us always denote the set of feasible tuples (I, v) for a partial differential equation (PDE) by \mathcal{A} . Optimal transformation paths between I_A and I_B are then defined as minimizers of

$$\min_{(I, v) \in \mathcal{A}} \lambda \mathcal{E}(v) + \|I(1, \cdot) - I_B\|_{L^2(\Omega)}^2,$$

where the regularizer \mathcal{E} enforces the smoothness of v . Note that this approach can be simplified a lot by using linearized deformations and a Taylor expansion around some initial flow field, leading to optical flow models [7, 41, 164]. We have successfully applied such an optical flow approach for analyzing crack evolution in a shear cutting process, see [149].

However, a transformation purely based on some deformation behavior is often not desirable, e.g., if the images have different mass or topological properties. Therefore, the approach was later extended to the metamorphosis model by Miller, Trounev and Younes [202, 263, 264]. In addition to the transport of pixel intensities, the model also allows variations of the intensities along the trajectories using a so-called source term z , i.e., image paths are solutions of the transport equation

$$\begin{aligned} \frac{\partial}{\partial t} I(t, x) + v(t, x) \nabla_x I(t, x) &= z & \text{for } (t, x) \in [0, 1] \times \Omega, \\ I(0, x) &= I_A(x) & \text{for } x \in \Omega, \\ I(1, x) &= I_B(x) & \text{for } x \in \Omega. \end{aligned} \quad (1.3)$$

In particular, this enables the method to create or erase objects during the transformation process, adding a lot more flexibility. Note that imposing a final condition is reasonable here as the source term z adds enough flexibility for solutions to exist. Equivalently, this problem can be also stated in Lagrange form, where (1.1) is complemented by an (implicit) integral equation for the intensity change

$$\begin{aligned} I(t, \varphi(t, x)) &= I_A(x) + \int_0^t z(s, \varphi(s, x)) \, ds \quad \text{for } (0, x) \in [0, 1] \times \Omega, \\ I(1, x) &= I_B(x) \quad \text{for } x \in \Omega. \end{aligned}$$

For the metamorphosis approach, the optimal transformation path between I_A and I_B is defined as the solution of

$$\min_{(I, v, z) \in \mathcal{A}} \lambda \mathcal{E}_1(v) + \mathcal{E}_2(z),$$

where \mathcal{A} denotes the set of tuples (I, v, z) satisfying (1.3). Again, the regularizers \mathcal{E}_1 and \mathcal{E}_2 enforce sufficient smoothness of v and z , respectively.

Many algorithmic approaches for both flow of diffeomorphism and metamorphosis have been proposed during the last years. In this thesis, we focus on the recent variational time discretization by Berkels, Effland and Rumpf [33] for metamorphosis and a Lagrangian approach for the plain flows of diffeomorphisms setting [193]. Usually, Lagrangian approaches are more efficient for this kind of problems as only an ODE has to be solved. Since the approaches rely on the minimization of energy functionals, they also provide a natural way of defining distances on the space of images. Hence, the metamorphosis model is a suitable choice for equipping the space $L^2(\Omega)$ with a new Riemannian structure. In some sense, this structure encodes how strong the necessary deformation of the domain is and how much mass is blended using the source term. A comprehensive overview is given in the book [285] as well as in the review article [201]. For a historic account see also [199].

Optimal transport Subsequently, we want to discuss optimal transport for measuring distances of images and transforming them. In its classical form, the problem dates back to the 18th century, when the French mathematician Monge published his famous work “Mémoire sur la théorie des déblais et des remblais”. The considered problem can be described as follows: Assume that soil is extracted from the ground in certain places like clay pits and then transported to other places like construction sites. Here, we have to keep in mind that every pit only has a fixed capacity and every construction site has a certain demand. For this particular setup, the locations of the pits and the construction sites are all known and the only unknown is the “optimal” assignment. Clearly, any kind of transportation involves costs that typically scale with the distance between source and target. Summing up, the aim is to minimize the total cost of our transportation process, while not violating capacity and demand. From this description, it should be evident now that optimal transport is a reasonable choice for measuring distances between images as it is able to capture geometric information about the problem. Further, it should

not be surprising anymore that this distance measure is sometimes also called earth mover’s distance.

The problem remained largely open for quite some time, until it was finally solved in a relaxed form by the Soviet mathematician Kantorovich in the 1940s. Over the years, the problem has been greatly generalized and is nowadays a popular choice for tackling various problems in imaging sciences, graphics or machine learning [82]. For good introductions to the topic we refer to the books [6, 243, 274], in which the problem is described in very general settings. If some power of a distance is used as cost function, the resulting transport distance is also known as Wasserstein metric.

In the following, we want to establish a relation to the previous paragraph using the Benamou-Brenier formulation [27] of optimal transport with squared Euclidean cost. This problem is given by

$$\min_{(I,v) \in \mathcal{A}} \int_0^1 \int_{\Omega} v(t,x)^2 I(t,x) \, dx \, dt,$$

where \mathcal{A} consists of all tuples (I, v) satisfying the continuity equation

$$\begin{aligned} \frac{\partial}{\partial t} I(t, x) + \operatorname{div}_x (v(t, x) I(t, x)) &= 0, & \text{for } (t, x) \in [0, 1] \times \Omega, \\ I(0, x) &= I_A, & \text{for } x \in \Omega, \\ I(1, x) &= I_B, & \text{for } x \in \Omega. \end{aligned} \tag{1.4}$$

Naturally, this approach provides us with a way of computing a transformation between I_A and I_B . Note that the continuity equation (1.4) is closely related to the transport equation (1.2) as both problems lead to the same formal Lagrange formulation for the characteristic lines. However, the approaches differ as the regularity requirements for v in the previous section are usually much stronger, i.e., they ensure that solutions in the Lagrange formulation (1.1) exist. In contrast, the unrestricted model based on (1.4) allows very general transformations for which not even continuity is enforced. This increases flexibility but often also leads to less pleasant transformation paths.

Unfortunately, the original optimal transport problem is computationally challenging, since a large linear program has to be solved. To deal with this issue, regularized versions of optimal transport, namely Sinkhorn distances [81], are used as a computationally efficient replacement. Compared to the original problem, they can be computed with matrix scaling algorithms, which are efficiently implementable on GPUs. Note that we have applied such an approach for coherent set detection in [178]. As these distances introduce a bias, Sinkhorn divergences [108] were recently introduced to overcome this problem. For appropriately related transport cost functions and discrepancy kernels, these divergences interpolate between optimal transport if the regularization parameter goes to zero and discrepancies if it goes to infinity.

Discrepancies Now, we want to discuss a way of quantifying the distance between probability densities on Ω . Naturally, this setting includes images interpreted as

densities. In order to simplify matters, assume for now that we only have access to a finite amount of samples drawn according to the densities I_A and I_B . Note that such a sampling approach is usually also used in applications. Then, informally speaking, we want to judge the distance of the corresponding empirical distributions by finding a test function that is large on the points from I_A and small on the points from I_B . Typically, the space of test functions is chosen as the unit ball in a reproducing kernel Hilbert space (RKHS). More precisely, the discrepancy is then obtained by maximizing the difference between the mean function values on the two sample sets over all test functions, i.e., a large value implies that the two input images are different. Clearly, the chosen kernel of the RKHS has a strong effect on the resulting discrepancy. In particular, an appropriately chosen kernel allows to incorporate some geometrical information into the discrepancy, making it a more reasonable choice than the total variation distance of the empirical measures. This can be seen as follows: If the sets of samples are disjoint, the total variation distance is 2. However, as the samples get closer and closer together, it will be increasingly difficult to find a “smooth” test function which is large for points from I_A and small for points from I_B , i.e., the discrepancy gets smaller. Coming back to our original problem, which corresponds to the limit of infinitely many samples, the discrepancy is obtained by maximizing the difference of the expectations $\mathbb{E}_{I_A}(f)$ and $\mathbb{E}_{I_B}(f)$ over all test functions f . Note that this approach is closely related to the 1-Wasserstein distance, where the set of test functions consists of all 1-Lipschitz functions.

The rates for approximating probability measures by atomic or empirical ones with respect to Wasserstein distances depend on the dimension of the underlying spaces, see [68, 177], which is often referred to as the curse of dimensionality. In sharp contrast, approximation rates based on discrepancies can be given independently of the dimension [222]. Discrepancies are linked to a wide range of applications, e.g., the derivation of quadrature rules [222], image dithering and representation [102, 134, 245, 259], generative adversarial networks [96] and multivariate statistical testing [107, 137, 138]. In the last two applications, they are better known as kernel based maximum mean discrepancies (MMDs). Finally, let us remark that discrepancies admit a simple description in the Fourier domain such that the fast Fourier transform can be applied for numerical purposes, leading to excellent scalability of approaches relying on discrepancies.

Organization and contribution of the thesis

In **Chapter 2**, we address morphing (transformation) of manifold-valued images based on the time discrete geodesic path model proposed by Berkels, Effland and Rumpf [33]. For our adapted model, we show existence of a minimizing sequence in $L^2(\Omega, \mathcal{H})$, provided that \mathcal{H} is a Hadamard manifold, together with a corresponding minimizing sequence of admissible diffeomorphisms. To this end, we also show that the continuous manifold-valued functions are dense in $L^2(\Omega, \mathcal{H})$. For our numerical experiments, the problem is spatially discretized using a finite difference approach on staggered grids together with a multilevel strategy, where we focus on the linearized elastic potential as regularizing term. The resulting problem is

solved by alternating the computation of a deformation sequence as solution of manifold-valued registration problems, and the computation of an image sequence based on a fixed sequence of deformations via the solution of a system of equations arising from the corresponding Euler-Lagrange equation. Numerical examples give a proof-of-concept of our ideas.

In **Chapter 3**, we continue analyzing the model introduced in Chapter 2. More precisely, we introduce a novel metamorphosis model for manifold-valued images. Compared to the original Euclidean model proposed by Trounev, Younes and co-workers [202, 264], we replace the PDE constraint with a variational inequality. The chosen approach is justified by establishing an equivalence of both models in the Euclidean setting. Based on this novel model, we are able to show convergence of time discrete geodesic paths to a time continuous geodesic path. In particular, this result establishes existence of geodesic paths for our new model. For our proofs, the joint convexity of the distance function, which characterizes Hadamard manifolds, is a crucial ingredient. A numerical example illustrates the proven convergence behavior.

As we demonstrate in **Chapter 4**, the Euclidean version of the model can be used as variational regularizer for image reconstruction from a template image and indirect, noisy measurements as given, for instance, in X-ray computed tomography. To this end, we combine time discrete metamorphosis with the classical L^2 -TV reconstruction model for inverse problems. For the proposed model, existence of a minimizer, stability with respect to the data, and convergence for vanishing noise are shown. On the numerical side, we propose two minimization procedures, which alternate over the involved sequences of deformations and images in different ways. The updates with respect to the image sequence exploit recent algorithms from convex analysis to minimize the L^2 -TV functional. Similarly as before, the problem is discretized using staggered grids together with a multilevel strategy. Numerical results for sparse and limited angle computed tomography as well as for superresolution demonstrate the usefulness of our method.

A different approach for solving inverse problems based on a template image is presented in **Chapter 5**. Here, images are reconstructed from measurements by solely deforming the given template image, i.e., no source term is incorporated. The required registration step is directly incorporated into the model in form of a PDE that models the registration as either mass- or intensity-preserving transport from the template to the unknown reconstruction. As in Chapter 4, we provide theoretical results for the proposed regularization. Numerically, we solve the problem by extending existing Lagrangian methods [193] and propose a multilevel approach, which is applicable whenever a suitable downsampling procedure for the operator and the measured data is available. To conclude this chapter, we demonstrate the performance of our method for image reconstruction based on some template image and highly undersampled, noisy Radon transform data. Our results indicate that reasonable reconstructions can be obtained when only few measurements are available and demonstrate that the use of the proposed normalized cross-correlation based distance is advantageous if the image intensities between template and unknown image differ substantially.

In **Chapter 6**, we want to investigate distances for images and measures. Coming from optimal transport, we are looking for computationally efficient alternatives. Here, Sinkhorn divergences turned out to be feasible regularizations, i.e., they converge to optimal transport as the regularization parameter ε approaches zero. In the opposite direction, as ε tends to infinity, Sinkhorn divergences with appropriate cost functions converge to discrepancies, for which a transport based interpretation is not possible any more. Here, we are particularly interested in the behavior of the corresponding optimal potentials appearing in the dual formulation of Sinkhorn divergences and discrepancies, respectively. While parts of the contained results are already known, we provide rigorous proofs for relations that cannot be found in the known literature in this generality. At the end of this chapter, we illustrate the limiting process by numerical examples and study the behavior of the distances when used for the approximation of measures by point measures in a process called dithering. Note that the computational cost for computing the distances decreases as we increase the regularization parameter.

In **Chapter 7**, we solely focus on discrepancies as distance measure. In contrast to Wasserstein distances, they do not suffer from the curse of dimensionality. Based on this choice, we study the approximation of probability measures on compact metric spaces and in particular on Riemannian manifolds by measures supported on Lipschitz curves. Potential applications include 3D printing, MRI, wire sculpture construction or Grand Tour computation on the Grassmannian $\mathcal{G}_{2,4}$. While some results cover the general case, special attention is paid to push-forward measures of Lebesgue measures on the unit interval by such curves. Based on a traveling salesman approach, we establish optimal approximation rates in terms of the curve's Lipschitz constant. In case the measures satisfy additional smoothness requirements, improved rates are given. For showing the upper bounds, we incorporate a result on the quadrature error achievable with measures that exactly integrate all eigenfunctions of the Laplace–Beltrami operator with eigenvalues smaller than a fixed number. Having established the theoretical results, we are interested in the numerical solution of the proposed model based on the Fourier formulation of discrepancies. To this end, we present numerical examples for measures on the 2- and 3-dimensional torus, the 2-sphere, the rotation group on \mathbb{R}^3 and the Grassmannian of all 2-dimensional linear subspaces of \mathbb{R}^4 . Our algorithm of choice is a conjugate gradient method on these manifolds, which incorporates second-order information. For efficient gradient and Hessian evaluations within the algorithm, we approximate the given measures by truncated Fourier series and use fast Fourier transform techniques on these manifolds. Our numerical experiments confirm the derived theoretical approximation rates.

Outlook

Working on these topics raised more questions than answers. Consequently, specific tasks for future research are given at the end of each section. Let us just sketch three directions that we intend to pursue in the future.

In cases where a lot of data and compute power is available, machine learn-

ing approaches have become increasingly popular, leading to outstanding results for instance in image classification. Actually quite a few successful learning approaches are inspired by classical imaging models and hence further research into these directions seems promising. In particular, Lipschitz neural networks play an important role when looking for stable methods, which are robust to adversarial attacks. Here, we have recently introduced proximal neural networks [150], which are averaged operators and hence also 1-Lipschitz. In this spirit, it could also be interesting to further explore the usage of Wasserstein metrics and MMDs as distance measure in neural networks, in particular for so-called invertible neural networks [26]. Potentially, a combination of our flow of diffeomorphism based approach from Chapter 5 with neural networks could allow to automatically choose between more than one template image for the reconstruction process.

A second set of problems is related to dimensionality reduction in data, see our recent line of works on principal component analysis [210, 211, 212]. As data becomes increasingly complex and large, such approaches are crucial for efficient workflows. In cases where the inherent data structure is more complex, approximation with curves (or more generally manifolds) instead of straight lines leads to improved results. Our approach for approximating measures in Chapter 7 can be interpreted in such a way, although deeper investigations are clearly necessary.

Finally, many of the transformation and motion models presented throughout this thesis do not have counterparts for manifold-valued images so far. As such data is becoming increasingly important, further research into this direction seems to be necessary. For example, it is not obvious anymore how images should be transformed by diffeomorphisms without an equivalent of the gray-value constancy assumption. In dependence on the application, the manifold data itself may change, as it is the case for EBSD data. Consequently, new models for the action of diffeomorphisms on manifold-valued images are necessary. In the Euclidean setting, such questions do not arise as scalars are rotationally invariant.

Morphing of Manifold-Valued Images inspired by Discrete Geodesics in Image Spaces

Contents

2.1	Introduction	13
2.2	Preliminaries	15
2.2.1	Hadamard spaces	15
2.2.2	Sobolev spaces and admissible mappings	19
2.3	Minimizers of the space continuous model	20
2.3.1	Space continuous model	20
2.3.2	Model specification	28
2.4	Minimization of the space discrete model	30
2.4.1	Computation of the deformation sequence	30
2.4.2	Computation of the image sequence	35
2.4.3	Multiscale minimization scheme	36
2.5	Numerical examples	36
2.5.1	Images in different color spaces	37
2.5.2	Symmetric positive definite matrices $\mathcal{P}(n)$	38
2.6	Conclusions	41

Abstract

The main goal of this chapter¹, published in [213], is to address morphing of manifold-valued images based on the time discrete geodesic paths model of Berkels, Effland and Rumpf [33]. Although for our manifold-valued setting such an interpretation of the energy functional is not available so far, the model is interesting on its own. We prove the existence of a minimizing sequence within the set of $L^2(\Omega, \mathcal{H})$ images having values in a finite-dimensional Hadamard manifold \mathcal{H} together with a minimizing sequence of admissible diffeomorphisms. To this end, we show that the continuous manifold-valued functions are dense in $L^2(\Omega, \mathcal{H})$. Then, we propose a space discrete model based on a finite difference approach on staggered grids, where we focus on the linearized elastic potential in the regularizing term. The numerical minimization alternates between the computation of a deformation sequence between given images via the parallel solution of certain registration problems for manifold-valued images, and the computation of an image sequence with fixed first (template) and last (reference) frame based on a given sequence of deformations via the solution of a system of equations arising from the corresponding Euler-Lagrange equation. Numerical examples give a proof-of-concept for our ideas.

¹First published in [213] in 2018, published by the Society for Industrial and Applied Mathematics (SIAM). Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

2.1 Introduction

Smooth image transition, also known as image morphing, is a frequently addressed task in image processing and computer vision, and there are various approaches to tackle the problem. For example, in feature based morphing only specific features are mapped to each other and the whole deformation is then calculated by interpolation. This was successfully applied, e.g., in the production of the movie *Willow* [251]. We refer to [281, 282] for an overview of similar techniques. This chapter is related to a special kind of image morphing, the so-called metamorphosis introduced by Miller, Trounev and Younes [202, 263, 264]. The metamorphosis model can be considered as an extension of the flow of diffeomorphism model and its large deformation diffeomorphic metric mapping framework [25, 71, 95, 261, 262], in which *each* image pixel is transported along a trajectory determined by a diffeomorphism path. As an extension, the metamorphosis model allows variations of image intensities along pixel trajectories. Solutions via shooting methods were developed, e.g., in [162, 232], where the first reference considers a metamorphosis regression model. A comprehensive overview over the topic is given in the book [285] as well as in the review article [201]. For a historic account see also [199].

This chapter builds up on a time discrete geodesic paths model by Berkels, Effland and Rumpf [33]. We mention that such a variational time discretization in shape spaces was already used in [240, 241], see also [119]. Let $\Omega \subset \mathbb{R}^n$, $n \geq 2$, be an open, bounded, connected domain with Lipschitz boundary. The authors of [33] define a (time) *discrete geodesic* connecting a template image $I_0 := T \in L^2(\Omega, \mathbb{R})$ and a reference image $I_K := R \in L^2(\Omega, \mathbb{R})$, $K \geq 2$, as a minimizing sequence $\mathbf{I} = (I_1, \dots, I_{K-1}) \in L^2(\Omega, \mathbb{R})^{K-1}$ of the *discrete path energy*

$$\mathcal{J}_{\text{BER}}(\mathbf{I}) := \sum_{k=1}^K \inf_{\varphi_k \in \mathcal{A}} \int_{\Omega} W(D\varphi_k) + \gamma |D^m \varphi_k|^2 + \frac{1}{\delta} (I_k \circ \varphi_k - I_{k-1})^2 dx, \quad (2.1)$$

subject to $I_0 = T, I_K = R,$

where $\delta, \gamma > 0$, the function W has to satisfy certain properties, \mathcal{A} is an admissible set of deformations, and the higher order derivatives $D^m \varphi_k$, $m > 1 + \frac{n}{2}$, guarantee a certain smoothness of the deformation. Berkels, Effland, and Rumpf showed that under certain constraints on W minimizers of \mathcal{J}_{BER} converge for $K \rightarrow \infty$ to a minimizer of the continuous geodesic path model of Trounev and Younes [263, 264], where the deformation is regularized by the dissipation density of a Newtonian fluid, which is the linearized elastic potential applied to the time derivative of the diffeomorphism path.

In this chapter, (2.1) is generalized to manifold-valued images in $L^2(\Omega, \mathcal{H})$ and proven that it is well-defined at least for finite-dimensional Hadamard manifolds \mathcal{H} . These are simply connected, complete Riemannian manifolds with non-positive sectional curvature. Typical examples of such Hadamard manifolds are hyperbolic spaces and symmetric positive definite matrices with the affine invariant metric. As an important fact, we use that the distance in Hadamard spaces is jointly convex, which implies weak lower semi-continuity of certain functionals involving the dis-

tance function. Since we use another admissible set than in [33], all proofs are new also for real-valued images. So far we have not established a relation of our model to some kind of time continuous path energy model in the image space $L^2(\Omega, \mathcal{H})$.

Dealing with digital images, we have to introduce a space discrete model. In contrast to the finite element approach in [33], we prefer a finite difference model on a staggered grid. We have used this discretization for gray-value images in [226]. Minimizers are computed using an alternating algorithm fixing either the deformation or the image sequence:

- i) For a fixed image sequence, we have to solve certain registration problems for *manifold-valued images* in parallel to get a sequence $(\varphi_1, \dots, \varphi_K)$ of deformations. Necessary interpolations are performed via Karcher mean computations. There exists a rich literature on registration problems for images with values in the Euclidean space, see [70, 111, 143, 147] and, for an overview, the books of Modersitzki [204, 205].
- ii) For a fixed deformation sequence, we need to find a minimizing image sequence (I_1, \dots, I_{K-1}) of

$$\sum_{k=1}^K d_2^2(I_k \circ \varphi_k, I_{k-1}) \quad \text{subject to} \quad I_0 = T, I_K = R,$$

where d_2 denotes the distance in $L^2(\Omega, \mathcal{H})$. In our manifold-valued setting, this requires to evaluate geodesics between manifold-valued image pixels at several well-defined time steps.

Outline of the chapter: We start with preliminaries on Hadamard spaces in Subsection 2.2.1, where the focus lies on the proof that the uniformly continuous functions mapping into locally compact Hadamard spaces \mathcal{H} are dense in $L^p(\Omega, \mathcal{H})$, $p \in [1, \infty)$. We have not found this result in the literature. Then, we introduce the necessary notation in Sobolev and Hölder spaces in Subsection 2.2.2 and an admissible set of deformations which differs from that in [33]. In particular, our definition guarantees that the concatenation of an image $I \in L^2(\Omega, \mathcal{H})$ with a deformation from our admissible set $I \circ \varphi$ is again an image in $L^2(\Omega, \mathcal{H})$. In Section 2.3, we introduce our space continuous model for manifold-valued images and prove existence of minimizers. Although we could roughly follow the lines in [33], all proofs are new also for the Euclidean setting $\mathcal{H} := \mathbb{R}^d$ due to the different admissible set. Moreover, the manifold-valued image setting requires a nontrivial update of the nice regridding idea from [33, 240, 241] in Theorem 2.7. This also influences the proof of our main result in Theorem 2.9. In Section 2.4, we present the computational approach for the space discrete model, where we propose a finite difference scheme on a staggered grid together with a multiscale technique. Numerical examples are given in Section 2.5. Note that the numerical algorithms are not restricted to Hadamard manifolds. We finish with conclusions in Section 2.6.

2.2 Preliminaries

Throughout this chapter, let $\Omega \subset \mathbb{R}^n$, $n \geq 2$, be an open, bounded, connected domain with Lipschitz boundary.

2.2.1 Hadamard spaces

For an overview on Hadamard spaces we refer to the books [13, 48, 169]. First, recall that a metric space (X, d) is *geodesic* if every two points $x, y \in X$ are connected by a curve $\gamma_{x,y}: [0, 1] \rightarrow X$, called geodesic, such that

$$d(\gamma_{x,y}(s), \gamma_{x,y}(t)) = |s - t|d(\gamma_{x,y}(0), \gamma_{x,y}(1)), \quad \text{for every } s, t \in [0, 1], \quad (2.2)$$

and $\gamma_{x,y}(0) = x$ and $\gamma_{x,y}(1) = y$. A complete metric space (\mathcal{H}, d) is called a *Hadamard space* if it is geodesic and if for every geodesic triangle $\triangle p, q, r \in \mathcal{H}$ and $x \in \gamma_{p,r}$, $y \in \gamma_{q,r}$ we have $d(x, y) \leq |\bar{x} - \bar{y}|$, where \bar{x}, \bar{y} are corresponding points in some comparison triangle $\triangle \bar{p}, \bar{q}, \bar{r} \in \mathbb{R}^2$ having the same side lengths as the geodesic one. By [13, Thm. 1.1.3], this is equivalent to (\mathcal{H}, d) being a complete metric geodesic space with

$$d^2(x, v) + d^2(y, w) \leq d^2(x, w) + d^2(y, v) + 2d(x, y)d(v, w), \quad (2.3)$$

for every $x, y, v, w \in \mathcal{H}$. Inequality (2.3) implies that geodesics are uniquely determined by their endpoints. Later, we restrict our attention to finite-dimensional Hadamard manifolds, which are Hadamard spaces having additionally a Riemannian manifold structure.

A function $f: \mathcal{H} \rightarrow \mathbb{R}$ is called *convex* if for every $x, y \in \mathcal{H}$ the function $f \circ \gamma_{x,y}$ is convex, i.e., if

$$f \circ \gamma_{x,y}(t) \leq (1 - t)f \circ \gamma_{x,y}(0) + tf \circ \gamma_{x,y}(1),$$

for each $t \in [0, 1]$. An important property of Hadamard spaces, which is also fulfilled in more general Busemann spaces, is that the distance is jointly convex, see [13, Prop. 1.1.5], i.e., for two geodesics $\gamma_{x_1, x_2}, \gamma_{y_1, y_2}$ and $t \in [0, 1]$ it holds

$$d(\gamma_{x_1, x_2}(t), \gamma_{y_1, y_2}(t)) \leq (1 - t)d(x_1, y_1) + td(x_2, y_2). \quad (2.4)$$

For a bounded sequence $\{x_n\}_{n \in \mathbb{N}}$ of points $x_n \in \mathcal{H}$, the function $w: \mathcal{H} \rightarrow [0, +\infty)$ defined by

$$w(x; \{x_n\}_{n \in \mathbb{N}}) := \limsup_{n \rightarrow \infty} d^2(x, x_n)$$

has a unique minimizer, which is called the *asymptotic center* of $\{x_n\}_{n \in \mathbb{N}}$, see [13, p. 58]. A sequence $\{x_n\}_{n \in \mathbb{N}}$ *converges weakly* to a point $x \in \mathcal{H}$ if it is bounded and x is the asymptotic center of each subsequence of $\{x_n\}_{n \in \mathbb{N}}$, see [13, p. 103]. The definition of proper and (weakly) lower semi-continuous (lsc) functions carries over from the Hilbert space setting.

2. Morphing of Manifold-Valued Images

The space \mathcal{H} is equipped with the Borel σ -algebra \mathcal{B} . A function $f: \Omega \rightarrow \mathcal{H}$ is (Lebesgue) measurable if $\{\omega \in \Omega : f(\omega) \in B\}$ is a (Lebesgue) measurable set for all $B \in \mathcal{B}$. In the following, we only consider the Lebesgue measure μ on $\Omega \subset \mathbb{R}^n$. A measurable map $f: \Omega \rightarrow \mathcal{H}$ belongs to $\mathcal{L}^p(\Omega, \mathcal{H})$, $p \in [1, \infty]$, if

$$d_p(f, a) < \infty,$$

for any constant mapping $\omega \mapsto a$ to a fixed $a \in \mathcal{H}$, where d_p is defined for two measurable maps f and g by

$$d_p(f, g) := \begin{cases} \left(\int_{\Omega} d^p(f(\omega), g(\omega)) d\omega \right)^{\frac{1}{p}} & p \in [1, \infty), \\ \text{ess sup}_{\omega \in \Omega} d(f(\omega), g(\omega)) & p = \infty. \end{cases}$$

Using the equivalence relation $f \sim g$ if $d_p(f, g) = 0$, the resulting quotient space $L^p(\Omega, \mathcal{H}) := \mathcal{L}^p(\Omega, \mathcal{H}) / \sim$ equipped with d_p becomes a complete metric space. For $p = 2$ it is again a Hadamard space, see [13, Prop. 1.2.18].

By $C(\Omega, \mathcal{H})$ we denote the space of continuous maps from Ω to \mathcal{H} . Next, we show that $C(\Omega, \mathcal{H})$ is dense in $L^p(\Omega, \mathcal{H})$. More precisely, also the uniformly continuous functions are dense. We start by defining simple functions (step functions). A function $g \in L^p(\Omega, \mathcal{H})$ is called a *simple function* if there exists a finite partition $\Omega = \bigcup_{i \in \mathcal{I}} A_i$ into disjoint measurable sets A_i such that $g|_{A_i} = a_i$ for all $i \in \mathcal{I}$. There exists a Hopf-Rinow-like theorem for Hadamard spaces, which states that in locally compact Hadamard spaces closed and bounded sets are compact, see [48, p. 35].

Lemma 2.1. *Let (\mathcal{H}, d) be a locally compact Hadamard space. Then the simple functions are dense in $L^p(\Omega, \mathcal{H})$, $p \in [1, \infty)$.*

Proof. Let $f \in L^p(\Omega, \mathcal{H})$. Then we have for a fixed reference point $a \in \mathcal{H}$ that

$$I_N := \int_{\{d(f(\omega), a) > N\}} d^p(f(\omega), a) d\omega \rightarrow 0$$


as $N \rightarrow \infty$. For an arbitrary $\varepsilon > 0$, we choose $N = N(\varepsilon)$ such that $I_N < \frac{\varepsilon}{2}$ and set

$$\mathcal{A}_0 := \{x \in \mathcal{H} : d(x, a) > N\} \quad \text{and} \quad A_0 := \{\omega \in \Omega : f(\omega) \in \mathcal{A}_0\}.$$

Next, we cover the compact set

$$\mathcal{A} := \{x \in \mathcal{H} : d(x, a) \leq N\}.$$

with open balls of radius $r_\varepsilon := \frac{\varepsilon}{2^{p+1}\mu(\Omega)}$. Since \mathcal{A} is compact, this covering contains a finite subcovering which can be restricted to \mathcal{A} and made disjoint such that $\mathcal{A} = \dot{\bigcup}_{i=1}^M \mathcal{A}_i$ for some $M \in \mathbb{N}$. Fixing any $a_i \in \mathcal{A}_i$, we have $d(x, a_i) \leq 2r_\varepsilon$ for all $x \in \mathcal{A}_i$. Since $\mathcal{A}_i \in \mathcal{B}$ and f is measurable, the sets $A_i := \{\omega \in \Omega : f(\omega) \in \mathcal{A}_i\}$ are measurable. Thus, we obtain a finite disjoint partition $\Omega = \dot{\bigcup}_{i=0}^M A_i$ into measurable



(a) Illustration for the construction of the B_i and K_i . (b) Illustration for the construction of the uniformly continuous approximation h_k .

Figure 2.1: Illustrations to the proof of Theorem 2.2.

sets. Defining the simple function $g: \Omega \rightarrow \mathcal{H}$ by $g|_{A_i} := a_i$, $i = 0, \dots, M$, where $a_0 := a$, we conclude

$$\begin{aligned} \int_{\Omega} d^p(f, g) d\omega &= \int_{A_0} d^p(f(\omega), a_0) d\omega + \sum_{i=1}^M \int_{A_i} d^p(f(\omega), a_i) d\omega \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2\mu(\Omega)} \sum_{i=1}^M \mu(A_i) \leq \varepsilon. \end{aligned}$$

□

Theorem 2.2. *Let (\mathcal{H}, d) be a locally compact Hadamard space. Then the set of uniformly continuous functions mapping from Ω to \mathcal{H} is dense in $L^p(\Omega, \mathcal{H})$, $p \in [1, \infty)$.*

Proof. By Lemma 2.1, it suffices to show that simple functions can be well approximated by uniformly continuous functions.

Let $g: \Omega \rightarrow \mathcal{H}$ be a simple function determined by A_i and corresponding $a_i \in \mathcal{H}$, $i \in \mathcal{I}$ from a finite index set $\mathcal{I} \subset \mathbb{N}$. Set $C := \max_{i,j \in \mathcal{I}} d^p(a_i, a_j)$ and let $\varepsilon > 0$ be arbitrary small. Since the Lebesgue measure is regular, there exist compact sets $K_i \subseteq A_i$ such that

$$\mu(A_i \setminus K_i) \leq \left(\frac{\varepsilon}{2}\right)^p \frac{1}{C |\mathcal{I}|}.$$

Since the K_i are disjoint and compact, there exists $\delta > 0$ such that $\text{dist}(K_i, K_j) > \delta$ and $\text{dist}(K_i, \partial\Omega) > \delta$ for all $i, j \in \mathcal{I}$, where dist is the distance function with respect to the Euclidean norm in \mathbb{R}^n . Then, the open sets

$$B_i = B_{\delta/4}(K_i) := \{\omega \in \Omega : \text{dist}(\omega, K_i) < \delta/4\}, \quad i \in \mathcal{I},$$

are disjoint. Let $B_0 := \Omega \setminus \cup_{i \in \mathcal{I}} B_i$ and $a_0 := a_1$. We define a simple function $\tilde{g}: \Omega \rightarrow \mathcal{H}$ by $\tilde{g}|_{B_i} := a_i$, $i \in \{0\} \cup \mathcal{I}$. For an illustration see Fig. 2.1a. It is related to the original step function g by

$$d_p^p(g, \tilde{g}) = \int_{\Omega \setminus \cup_{i \in \mathcal{I}} K_i} d^p(g, \tilde{g}) d\omega \leq C \sum_{i \in \mathcal{I}} \mu(A_i \setminus K_i) \leq \left(\frac{\varepsilon}{2}\right)^p. \quad (2.5)$$

2. Morphing of Manifold-Valued Images

Next, we approximate \tilde{g} by a sequence of uniformly continuous functions. Let $\gamma_{a,b}$ denote the unique geodesic joining $a, b \in \mathcal{H}$, where $\gamma_{a,b}(0) = a$ and $\gamma_{a,b}(1) = b$. For $k \in \mathbb{N}$, define $h_k: \Omega \rightarrow \mathcal{H}$ by $h_k|_{B_0} := a_0$ and for $i \in \mathcal{I}$,

$$h_k(\omega) := \gamma_{a_0, a_i}(\min\{1, k \text{dist}(B_i^c, \omega)\}), \quad \omega \in B_i,$$

see Fig. 2.1b. Since γ is a geodesic, h_k is by construction continuous on every B_i , $i \in \mathcal{I}$. Further, $h_k|_{\partial B_i} = a_0$ and for any sequence $\{\omega_j\}_{j \in \mathbb{N}}$ converging to some $\hat{\omega} \in \partial B_i$, we have that $h_k(\omega_j)$ converges to $a_0 = h_k(\hat{\omega})$. Hence, h_k is continuous on Ω and by construction constant outside of the compact set $\cup_{i \in \mathcal{I}} \overline{B_i}$. This implies that h_k is even uniformly continuous on Ω . Let

$$\max_{i, j \in \mathcal{I}} \sup_{x \in \gamma_{a_0, a_i}} d^p(x, a_j) \leq D.$$

Then, we get $d^p(\tilde{g}(\omega), h_k(\omega)) \leq D$ for all $k \in \mathbb{N}, \omega \in \Omega$, so that D is an integrable bound of $d^p(\tilde{g}, h_k)$ and $\lim_{k \rightarrow \infty} d(\tilde{g}(\omega), h_k(\omega))^p \rightarrow 0$ pointwise as $k \rightarrow \infty$. By Lebesgue's convergence theorem this implies that

$$\lim_{k \rightarrow \infty} d_p(\tilde{g}, h_k) = 0.$$

Finally, choosing $k \in \mathbb{N}$ such that $d_p(\tilde{g}, h_k) < \frac{\varepsilon}{2}$ we obtain with (2.5) that

$$d_p(g, h_k) \leq d_p(g, \tilde{g}) + d_p(\tilde{g}, h_k) < \varepsilon.$$

This concludes the proof. \square

In Hadamard spaces we have no zero element so that a notion of “compact support” is not available. However, the previous proof shows that we can approximate a function in $L^p(\Omega, \mathcal{H})$ arbitrarily well by a function which is constant outside of a compact set, which can be seen as an equivalent. The previous theorem can be applied to prove the next corollary.

Corollary 2.3. *Let (\mathcal{H}, d) be a locally compact Hadamard space, $f \in L^p(\Omega, \mathcal{H})$, $p \in [1, \infty)$, and $\{\varphi_j\}_{j \in \mathbb{N}}$ be diffeomorphisms on Ω such that $\lim_{j \rightarrow \infty} \|\varphi_j - \hat{\varphi}\|_{L^\infty(\Omega)^n} = 0$ and $|\det D\varphi_j|^{-1} \leq C$ for all $j \in \mathbb{N}$. Then $\limsup_{j \rightarrow \infty} d_p(f \circ \varphi_j, f \circ \hat{\varphi}) = 0$.*

Proof. For $f \in L^p(\Omega, \mathcal{H})$, we can construct by Theorem 2.2 a sequence $\{f_k\}_{k \in \mathbb{N}}$ of uniformly continuous functions with $d_p(f, f_k) < \frac{1}{k}$. Then, we conclude

$$\begin{aligned} d_p(f \circ \varphi_j, f \circ \hat{\varphi}) &\leq d_p(f \circ \varphi_j, f_k \circ \varphi_j) + d_p(f_k \circ \varphi_j, f_k \circ \hat{\varphi}) + d_p(f_k \circ \hat{\varphi}, f \circ \hat{\varphi}) \\ &\leq C d_p(f, f_k) + d_p(f_k \circ \varphi_j, f_k \circ \hat{\varphi}) + C d_p(f_k, f). \end{aligned}$$

By construction, the first and the last term can be made arbitrary small as $k \rightarrow \infty$. Now, let $k \in \mathbb{N}$ be fixed. Since φ_j converges uniformly to $\hat{\varphi}$, we can use the uniform continuity of f_k in order to conclude that $f_k \circ \varphi_j$ converges uniformly to $f_k \circ \hat{\varphi}$. Boundedness of Ω implies that the middle term converges to zero as $j \rightarrow \infty$. \square

In the rest of this chapter, we restrict our considerations to finite-dimensional Hadamard manifolds (\mathcal{H}, d) , i.e., \mathcal{H} has an additional Riemannian manifold structure. Clearly, finite-dimensional Hadamard manifolds are locally compact. By $T_x\mathcal{H}$ we denote the tangential space of \mathcal{H} at $x \in \mathcal{H}$. Then, the geodesics $\gamma_{x,v}$ are determined by their starting point $x \in \mathcal{H}$ and their tangential $v \in T_x\mathcal{H}$ at this point. Further, we need the exponential map $\exp_x: T_x\mathcal{H} \rightarrow \mathcal{H}$ defined by $\exp_x v := \gamma_{x,v}(1)$, and the inverse of the exponential map $\log_x := \exp_x^{-1}: \mathcal{H} \rightarrow T_x\mathcal{H}$.

2.2.2 Sobolev spaces and admissible mappings

Let $C^{k,\alpha}(\overline{\Omega})$, $k \in \mathbb{N}_0$, denote the Hölder space of functions $f \in C^k(\overline{\Omega})$ for which

$$\|f\|_{C^{k,\alpha}(\overline{\Omega})} := \sum_{|\beta| \leq k} \|D^\beta f\|_{C(\overline{\Omega})} + \sum_{|\beta|=k} \sup_{\substack{x,y \in \Omega \\ x \neq y}} \frac{|D^\beta f(x) - D^\beta f(y)|}{|x - y|^\alpha}$$

is finite. With this norm $C^{k,\alpha}(\overline{\Omega})$ is a Banach space.

By $W^{m,p}(\Omega)$, $m \in \mathbb{N}$, $1 \leq p < \infty$, we denote the Sobolev space of functions having weak derivatives up to order m in $L^p(\Omega)$ with norm

$$\|f\|_{W^{m,p}(\Omega)}^p := \int_{\Omega} \sum_{|\alpha| \leq m} |D^\alpha f|^p dx.$$

We apply the usual abbreviation $|D^m f|^p := \sum_{|\alpha|=m} |D^\alpha f|^p$. For $F = (f_\nu)_{\nu=1}^n$, we set $|D^m F|^p = \sum_{\nu=1}^n |D^m f_\nu|^p$. In particular, we are interested in $W^{m,2}(\Omega)$ with $m > 1 + \frac{n}{2}$. In this case, $W^{m,2}(\Omega)$ is compactly embedded into $C^{1,\alpha}(\overline{\Omega})$ for all $\alpha \in (0, m - 1 - \frac{n}{2})$ [4, p. 350, Th. 8.13] and consequently $W^{m,2}(\Omega) \hookrightarrow W^{1,p}(\Omega)$ for all $p \geq 1$.

For $m > 1 + \frac{n}{2}$, we consider the set

$$\mathcal{A} := \{\varphi \in W^{m,2}(\Omega)^n : \det D\varphi > 0 \text{ a.e. in } \Omega, \varphi(x) = x \text{ for } x \in \partial\Omega\},$$

which was used as admissible set of deformations in [33]. By the results of Ball [19], we know that $\varphi(\overline{\Omega}) = \overline{\Omega}$ and φ is injective a.e. By the last property and since $\overline{\Omega}$ is bounded, it follows immediately for all $\varphi \in \mathcal{A}$ that

$$\|\varphi\|_{L^\infty(\Omega)^n} \leq C, \quad \|\varphi\|_{L^2(\Omega)^n} \leq C, \quad (2.6)$$

with constants depending only on Ω . We have $\varphi \in C^{1,\alpha}(\overline{\Omega})^n$ and by the inverse mapping theorem φ^{-1} exists locally around a.e. $x \in \Omega$ and is continuously differentiable on the corresponding neighborhood. However, to guarantee that φ^{-1} is continuous (or, even more, continuously differentiable) further assumptions are required, see [19, Thm. 2]. Take for example the function $\varphi(x) := x^3$ on $\Omega := (-1, 1)$ which is in \mathcal{A} , but $\varphi^{-1}(x) = \text{sgn}(x)|x|^{\frac{1}{3}}$ is not continuously differentiable. Furthermore, $I \in L^2(\Omega, \mathcal{H})$ and $\varphi \in \mathcal{A}$ do not guarantee that $I \circ \varphi \in L^2(\Omega, \mathcal{H})$, as the example $I(x) := x^{-\frac{1}{4}}$ in $L^2([0, 1])$ and $\varphi(x) := x^2$ shows.

Therefore, we introduce, for *small* fixed $\epsilon > 0$, the *admissible set*

$$\mathcal{A}_\epsilon := \{\varphi \in W^{m,2}(\Omega)^n : \det D\varphi \geq \epsilon, \varphi(x) = x \text{ for } x \in \partial\Omega\} \subset \mathcal{A}.$$

Later, in Theorem 2.7, we have to solve a system of equations with entries depending on $\det D\varphi$. For stability reasons, we do not want that the determinant becomes arbitrary small. This can be avoided by introducing ϵ . Moreover, by the inverse mapping theorem, $\varphi \in \mathcal{A}_\epsilon$ is a diffeomorphism, although in general $\varphi^{-1} \notin \mathcal{A}_\epsilon$. Further, $I \in L^2(\Omega, \mathcal{H})$ and $\varphi \in \mathcal{A}_\epsilon$ imply $I \circ \varphi \in L^2(\Omega, \mathcal{H})$. We mention that the space of images $L^\infty(\Omega, \mathbb{R})$ was discussed in the thesis [97]. However, working in the Hadamard space $L^2(\Omega, \mathcal{H})$ simplifies the proofs. In particular, we can use the concept of weak convergence in these spaces.

2.3 Minimizers of the space continuous model

Let \mathcal{H} be a finite-dimensional Hadamard manifold. Due to our application, we call the functions from $L^2(\Omega, \mathcal{H})$ images. Mappings $\varphi \in \mathcal{A}_\epsilon$ can act on images $I \in L^2(\Omega, \mathcal{H})$ via

$$\varphi \cdot I = I \circ \varphi, \quad \varphi \in \mathcal{A}_\epsilon.$$

Inspired by the time discrete geodesic paths model for images with values in \mathbb{R}^n , see [33], we introduce a general morphing model for manifold-valued images and prove existence of minimizers for this model in the next subsection. Although we can basically follow the lines in [33], all proofs are new even for the Euclidean setting due to the different admissible set. Moreover, the manifold-valued setting requires some care when minimizing over the image sequence. Then, in Subsection 2.3.2, we specify the model by choosing the linearized elastic potential as regularizer.

2.3.1 Space continuous model

Let $W: \mathbb{R}^{n,n} \rightarrow \mathbb{R}_{\geq 0}$ be a lsc mapping, $\gamma > 0$, $m > 1 + \frac{n}{2}$ and $K \geq 2$ be an integer. Given a template image and a reference image

$$I_0 = T \in L^2(\Omega, \mathcal{H}), \quad I_K = R \in L^2(\Omega, \mathcal{H}),$$

respectively, we are searching for an image sequence

$$\mathbf{I} := (I_1, \dots, I_{K-1}) \in L^2(\Omega, \mathcal{H})^{K-1},$$

see Fig. 2.2, which minimizes the energy

$$\begin{aligned} \mathcal{J}(\mathbf{I}) &:= \sum_{k=1}^K \inf_{\varphi_k \in \mathcal{A}_\epsilon} \int_{\Omega} W(D\varphi_k) + \gamma |D^m \varphi_k|^2 dx + d_2^2(I_{k-1} \circ \varphi_k, I_k) \\ &= \inf_{\varphi \in \mathcal{A}_\epsilon^K} \sum_{k=1}^K \int_{\Omega} W(D\varphi_k) + \gamma |D^m \varphi_k|^2 dx + d_2^2(I_{k-1} \circ \varphi_k, I_k) \\ &= \inf_{\varphi \in \mathcal{A}_\epsilon^K} \mathcal{J}(\mathbf{I}, \varphi), \end{aligned} \tag{2.7}$$




Figure 2.2: Illustration of the image and the diffeomorphism path

where $\varphi := (\varphi_1, \dots, \varphi_K)$ and

$$\mathcal{J}(\mathbf{I}, \varphi) := \sum_{k=1}^K \int_{\Omega} W(D\varphi_k) + \gamma |D^m \varphi_k|^2 dx + d_2^2(I_{k-1} \circ \varphi_k, I_k).$$

Note that for simplicity of notation we moved the parameter $\frac{1}{\delta}$ of the squared distance term to W , see (2.18), and used a shift in the φ_k .

A typical strategy for minimizing such a functional is alternating minimization over \mathbf{I} and φ . In the following, we show that the corresponding subproblems have a minimizer. This can then be used to show that the whole functional has a minimizer.

First, we fix the image sequence $\mathbf{I} \in L^2(\Omega, \mathcal{H})^{K-1}$ and show that $\mathcal{J}(\mathbf{I}, \cdot)$ has a minimizer $\varphi \in \mathcal{A}^K$. It suffices to prove that each of the registration problems

$$\mathcal{R}(\varphi_k; I_{k-1}, I_k) := \int_{\Omega} W(D\varphi_k) + \gamma |D^m \varphi_k|^2 dx + d_2^2(I_{k-1} \circ \varphi_k, I_k)$$

has a minimizer in \mathcal{A}_ϵ , $k = 1, \dots, K$. Note that we can show the more general result for $\varphi \in \mathcal{A}$ for this functional. If we restrict ourselves to $\varphi \in \mathcal{A}_\epsilon$, the proof can be simplified, see Corollary 2.5.

Theorem 2.4. *Let $W: \mathbb{R}^{n,n} \rightarrow \mathbb{R}_{\leq 0}$ be a lsc mapping with*

$$W(A) = \infty \quad \text{if} \quad \det A \leq 0. \quad (2.8)$$

Further, let $T, R \in L^2(\Omega, \mathcal{H})$ be given. Then there exists $\hat{\varphi} \in \mathcal{A}$ minimizing

$$\mathcal{R}(\varphi; T, R) = \int_{\Omega} W(D\varphi) + \gamma |D^m \varphi|^2 + d^2(T \circ \varphi, R) dx$$

over all $\varphi \in \mathcal{A}$.

Proof. 1. Let $\{\varphi_j\}_{j \in \mathbb{N}}$, $\varphi_j \in \mathcal{A}$, be a minimizing sequence of \mathcal{R} . Then, we have $\mathcal{R}(\varphi_j; T, R) \leq C$ for all $j \in \mathbb{N}$. This implies that $\{\varphi_j\}_{j \in \mathbb{N}}$ has uniformly bounded $W^{m,2}(\Omega)^n$ -seminorm, and by (2.6) the sequence is also uniformly bounded in $L^2(\Omega)^n$. Now, we apply the Gagliardo–Nirenberg inequality, see Theorem 2.12 and Remark 2.13, which states that for all $0 \leq i < m$ it holds

$$\|D^i \varphi_j\|_{L^2(\Omega)} \leq C_1 \|D^m \varphi_j\|_{L^2(\Omega)} + C_2 \|\varphi_j\|_{L^2(\Omega)}.$$

All terms on the right hand side are uniformly bounded. Hence, the $W^{m,2}(\Omega)^n$ norm of $\{\varphi_j\}_{j \in \mathbb{N}}$ is uniformly bounded. Since $W^{m,2}(\Omega)$ is reflexive, there exists a subsequence $\{\varphi_{j_l}\}_{l \in \mathbb{N}}$ that converges weakly to some function $\hat{\varphi}$ in $W^{m,2}(\Omega)^n$. By

2. Morphing of Manifold-Valued Images

the compact embedding $W^{m,2}(\Omega) \hookrightarrow C^{1,\alpha}(\overline{\Omega})$, $\alpha \in (0, m - 1 - \frac{n}{2})$, this subsequence converges strongly to $\hat{\varphi}$ in $C^{1,\alpha}(\overline{\Omega})^n$. We again denote this subsequence by $\{\varphi_j\}_{j \in \mathbb{N}}$.

2. Next, we show that $\hat{\varphi}$ is in the set \mathcal{A} . By the first part, $D\varphi_j$ converges uniformly to $D\hat{\varphi}$. Since W is lsc, this implies

$$W(D\hat{\varphi})(x) \leq \liminf_{j \rightarrow \infty} W(D\varphi_j)(x)$$

for all $x \in \Omega$ and since W is non-negative, we obtain by Fatou's lemma

$$\int_{\Omega} W(D\hat{\varphi}) \, dx \leq \liminf_{j \rightarrow \infty} \int_{\Omega} W(D\varphi_j) \, dx \leq C.$$

By (2.8) this implies $\det D\hat{\varphi} > 0$ a.e. Further, the boundary condition is fulfilled so that $\hat{\varphi} \in \mathcal{A}$. It remains to show that $\lim_{j \rightarrow \infty} \mathcal{R}(\varphi_j; T, R) = \mathcal{R}(\hat{\varphi}; T, R)$.

3. We prove that

$$d_2^2(T \circ \hat{\varphi}, R) \leq \liminf_{j \rightarrow \infty} d_2^2(T \circ \varphi_j, R).$$

Assume $d_2^2(T \circ \hat{\varphi}, R) > \liminf_{j \rightarrow \infty} d_2^2(T \circ \varphi_j, R)$ and for the moment also that $d_2^2(T \circ \hat{\varphi}, R)$ is finite. Then, we can find an $\delta > 0$ such that

$$d_2^2(T \circ \hat{\varphi}, R) - \delta = \liminf_{j \rightarrow \infty} d_2^2(T \circ \varphi_j, R).$$

The sets $\{x \in \Omega : \det D\varphi_j = 0, j \in \mathbb{N}\}$ and $\{x \in \Omega : \det D\hat{\varphi} = 0\}$ have both measure zero. Let Ω_+ be their complementary set, which is open since the φ_j are convergent. By the inverse mapping theorem, we know that φ_j^{-1} and $\hat{\varphi}^{-1}$ are continuously differentiable on Ω_+ . Since Ω_+ is open, we can use monotone convergence to find a compact set \mathcal{K} such that $\mathcal{K} \subset \Omega_+$ and

$$d_2^2(T \circ \hat{\varphi}, R) \leq \int_{\mathcal{K}} d^2(T \circ \hat{\varphi}, R) \, dx + \frac{\delta}{2}.$$

Using that the integrands are non-negative, we get

$$\liminf_{j \rightarrow \infty} \int_{\mathcal{K}} d^2(T \circ \varphi_j, R) \, dx < \int_{\mathcal{K}} d^2(T \circ \hat{\varphi}, R) \, dx. \quad (2.9)$$

For the case $d_2^2(T \circ \hat{\varphi}, R) = \infty$ we obtain the same inequality, since the left-hand side is less than some finite constant C independent of \mathcal{K} . However, \mathcal{K} can be chosen such that the right-hand side gets arbitrary large using monotone convergence. Note that it will be still finite due to the change of variables formula [239, Thm. 7.26].

Now, for the set \mathcal{K} , we obtain

$$\begin{aligned} & \left| \int_{\mathcal{K}} d^2(T \circ \varphi_j, R) - d^2(T \circ \hat{\varphi}, R) \, dx \right| \\ & \leq \int_{\mathcal{K}} (d(T \circ \varphi_j, R) + d(T \circ \hat{\varphi}, R)) |d(T \circ \varphi_j, R) - d(T \circ \hat{\varphi}, R)| \, dx \\ & \leq \int_{\mathcal{K}} (d(T \circ \varphi_j, R) + d(T \circ \hat{\varphi}, R)) d(T \circ \varphi_j, T \circ \hat{\varphi}) \, dx \\ & \leq C_1 \left(\int_{\mathcal{K}} d^2(T \circ \varphi_j, T \circ \hat{\varphi}) \, dx \right)^{\frac{1}{2}} \end{aligned} \quad (2.10)$$

with

$$C_1 := \left(\int_{\mathcal{K}} d^2(T \circ \varphi_j, R) \, dx \right)^{\frac{1}{2}} + \left(\int_{\mathcal{K}} d^2(T \circ \hat{\varphi}, R) \, dx \right)^{\frac{1}{2}} < \infty.$$

This constant is finite as the first term is uniformly bounded by construction and the second term is bounded by construction of \mathcal{K} . Now, Corollary 2.3 implies that the term

$$\int_{\mathcal{K}} d^2(T \circ \varphi_j, T \circ \hat{\varphi}) \, dx$$

in (2.10) converges to zero. This yields a contradiction to (2.9).

4. By the previous steps, we have that the three summands in \mathcal{R} are (weakly) lower semi-continuous. Hence, we get

$$\mathcal{R}(\hat{\varphi}; T, R) \leq \liminf_{j \rightarrow \infty} \int_{\Omega} W(D\varphi_j) + \gamma |D^m \varphi_j|^2 + d^2(T \circ \varphi_j, R) \, dx = \inf_{\varphi \in \mathcal{A}} \mathcal{R}(\varphi),$$

which proves the claim. \square

Following the lines of the previous theorem, we can prove its analogue for \mathcal{A}_ϵ .

Corollary 2.5. *Let $W: \mathbb{R}^{n,n} \rightarrow \mathbb{R}_+$ be a lsc mapping. Further, let $T, R \in L^2(\Omega, \mathcal{H})$ be given. Then there exists $\hat{\varphi} \in \mathcal{A}_\epsilon$ minimizing*

$$\mathcal{R}(\varphi; T, R) := \int_{\Omega} W(D\varphi) + \gamma |D^m \varphi|^2 + d^2(T \circ \varphi, R) \, dx$$

over all $\varphi \in \mathcal{A}_\epsilon$.

Proof. The proof is similar to the previous one, but simplifies as all integrals are defined over Ω and the approximation with compact sets is not necessary. \square

Next, we fix the sequence of admissible mappings $\varphi \in \mathcal{A}_\epsilon^K$ and look for a minimizer of

$$J_\varphi(\mathbf{I}) := \sum_{k=1}^K d_2^2(I_{k-1} \circ \varphi_k, I_k) \quad \text{subject to} \quad I_0 = T, \quad I_K = R.$$

Note that we consider admissible functions in \mathcal{A}_ϵ now, since the composition $I \circ \varphi$ of a deformation φ and an image $I \in L^2(\Omega, \mathcal{H})$ should be again in $L^2(\Omega, \mathcal{H})$.

Lemma 2.6. *For fixed $\varphi \in \mathcal{A}_\epsilon^K$, there exists a sequence $\hat{\mathbf{I}} \in L^2(\Omega, \mathcal{H})^{K-1}$ minimizing J_φ .*

Proof. Let $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$ be a minimizing sequence of J_φ , i.e., $\lim_{j \rightarrow \infty} J_\varphi(\mathbf{I}^{(j)}) = \hat{J}$, where \hat{J} is the infimum of J_φ . Clearly, there exists $C \geq 0$ such that $J_\varphi(\mathbf{I}^{(j)}) \leq C$ for all $j \in \mathbb{N}$. By the triangle inequality, we obtain for some fixed $a \in \mathcal{H}$ that

$$d_2(I_{k+1}^{(j)}, a) \leq d_2(I_k^{(j)} \circ \varphi_{k+1}, I_{k+1}^{(j)}) + d_2(I_k^{(j)} \circ \varphi_{k+1}, a) \leq C + d_2(I_k^{(j)} \circ \varphi_{k+1}, a).$$

Thus, since $\det D\varphi_k \geq \epsilon$ a.e. on Ω for $k = 1, \dots, K$,

$$\begin{aligned} d_2(I_1^{(j)}, a) &\leq C + d_2(T \circ \varphi_1, a) \\ d_2(I_2^{(j)}, a) &\leq C + d_2(I_1^{(j)} \circ \varphi_2, a) = C + \left(\int_{\Omega} d^2(I_1^{(j)} \circ \varphi_2, a) \, dx \right)^{\frac{1}{2}} \\ &= C + \left(\int_{\Omega} d^2(I_1^{(j)}, a) |\det(D\varphi_2 \circ \varphi_2^{-1})|^{-1} \, dx \right)^{\frac{1}{2}} \\ &\leq C + \epsilon^{-\frac{1}{2}} d_2(I_1^{(j)}, a) \leq C + \epsilon^{-\frac{1}{2}} (C + d_2(T \circ \varphi_1, a)). \end{aligned}$$

Continuing this successively, we see that $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$ is bounded in $L^2(\Omega, \mathcal{H})^{K-1}$. From [13, Prop. 3.1.2] we know that a bounded sequence in a Hadamard space has a weakly convergent subsequence $\{\mathbf{I}^{(j_k)}\}_{k \in \mathbb{N}}$. Let $\hat{\mathbf{I}} \in L^2(\Omega, \mathcal{H})^{K-1}$ be its weak limit point. Now $d_2(\cdot, \cdot)$ is a continuous convex function and the same holds true for J_{φ} . Then, by [13, Lem. 3.2.3], the function J_{φ} is weakly lsc, which means that

$$\hat{J} = \lim_{j \rightarrow \infty} J_{\varphi}(\mathbf{I}^{(j)}) = \lim_{k \rightarrow \infty} J_{\varphi}(\mathbf{I}^{(j_k)}) \geq J_{\varphi}(\hat{\mathbf{I}}),$$

so that $\hat{\mathbf{I}}$ is a minimizer of the functional. \square

Theorem 2.7. *For fixed $\varphi \in \mathcal{A}_{\epsilon}^K$, there exists a unique sequence of intermediate images $\mathbf{I} \in L^2(\Omega, \mathcal{H})^{K-1}$ minimizing J_{φ} .*

Proof. By Lemma 2.6, there exists a minimizer of J_{φ} . Setting

$$\begin{aligned} \psi_K &:= \text{Id}, \\ \psi_k &:= \varphi_{k+1} \circ \psi_{k+1} = \varphi_{k+1} \circ \dots \circ \varphi_K, \quad k = K-1, \dots, 0, \end{aligned} \tag{2.11}$$

and substituting $x := \psi_k(y)$ in the k -th summand of J_{φ} , we obtain

$$J_{\varphi}(\mathbf{I}) = \sum_{k=1}^K \int_{\Omega} d^2(I_k \circ \psi_k, I_{k-1} \circ \psi_{k-1}) |\det D\psi_k| \, dy.$$

Using $F_K := I_K = R$, $F_k := I_k \circ \psi_k$, and $w_k := |\det D\psi_k| > 0$, we are concerned with the minimization of

$$\sum_{k=1}^K \int_{\Omega} w_k d^2(F_k, F_{k-1}) \, dx \quad \text{subject to } F_0 = T \circ \psi_0, \, F_K = R.$$

Each image I_k , resp. F_k , appears only in two summands

$$\int_{\Omega} w_k d^2(F_k, F_{k-1}) + w_{k+1} d^2(F_{k+1}, F_k) \, dx.$$

By the Euler-Lagrange equation and since the Riemannian gradient of the squared distance function is $\nabla d^2(\cdot, a)(b) = -2 \log_b a$, we obtain

$$w_k \log_{F_k} F_{k-1} + w_{k+1} \log_{F_k} F_{k+1} = 0, \quad k = 1, \dots, K-1. \tag{2.12}$$




Figure 2.3: Illustration of relation (2.14) for the geodesic (2.13) with $K = 3$, where $\delta, \tilde{\delta}$ are constants canceling out in the fraction.

If $K = 2$, we have only one equation

$$w_1 \log_{F_1} F_0 + w_2 \log_{F_1} F_2 = 0,$$

which implies that F_1 has to be on the geodesic $\gamma_{F_0, F_2}(t), t \in [0, 1]$. Since

$$w_1 |\log_{F_1} F_0| = w_2 |\log_{F_1} F_2|,$$

we obtain

$$F_1 = \gamma_{F_0, F_2}\left(\frac{w_2}{w_1 + w_2}\right).$$

For general $K \geq 3$, the system of equations (2.12) can only be fulfilled if three consecutive points always lie on a geodesic. This is only possible if all points are on the same geodesic $\gamma_{F_0, F_K}(t), t \in [0, 1]$. More precisely,

$$F_k = \gamma_{F_0, F_K}(t_k), \quad k = 1, \dots, K - 1, \quad (2.13)$$

where by (2.12), the t_k are related via

$$\frac{s_k}{s_{k+1}} = \frac{w_{k+1}}{w_k}, \quad s_k := t_k - t_{k-1}, \quad k = 1, \dots, K - 1, \quad (2.14)$$

see Fig. 2.3 for an illustration. It is easy to check that these conditions are fulfilled by

$$s_k := \frac{\alpha_k}{\sum_{i=1}^{K-1} \alpha_i}, \quad \alpha_k := \prod_{\substack{i=1 \\ i \neq k}}^K w_i, \quad k = 1, \dots, K - 1,$$

so that

$$t_k = \sum_{i=1}^k s_i = \frac{\sum_{i=1}^k w_i^{-1}}{\sum_{i=1}^K w_i^{-1}}.$$

The geodesics between F_0 and F_K are unique, so that the resulting points F_k , $k = 1, \dots, K - 1$, are unique as well. As we know that φ_k are diffeomorphisms, the functions ψ_k , $k = 1, \dots, K - 1$, are diffeomorphisms as well. By (2.13), convexity of d^2 , and $F_0, R \in L^2(\Omega, \mathcal{H})$, we see that $F_k \in L^2(\Omega, \mathcal{H})$ and thus $I_k \in L^2(\Omega, \mathcal{H})$, $k = 1, \dots, K - 1$. \square

For proving the next theorem, we need the following corollary.

Corollary 2.8. *Let $\{\varphi^{(j)}\}_{j \in \mathbb{N}}$ with $\varphi^{(j)} \in \mathcal{A}_\epsilon^K$ be a sequence that converges in $C^{1,\alpha}(\overline{\Omega})^{nK}$. For each $j \in \mathbb{N}$, let $\mathbf{I}^{(j)} \in L^2(\Omega, \mathcal{H})^{K-1}$ be the minimizer of $J_{\varphi^{(j)}}$. Then the sequence $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$ converges in $L^2(\Omega, \mathcal{H})^{K-1}$.*

Proof. The proof follows the path of the previous theorem. Similarly as in (2.11), we define

$$\psi_K := \text{Id}, \quad \psi_k^{(j)} := \varphi_{k+1}^{(j)} \circ \psi_{k+1}^{(j)}, \quad F_k^{(j)} := I_k^{(j)} \circ \psi_k^{(j)}, \quad k = 0, \dots, K-1,$$

where $I_0^{(j)} = T$ and $F_K^{(j)} = R$. Clearly, the convergence of $\{\varphi^{(j)}\}_{j \in \mathbb{N}}$ in $C^{1,\alpha}(\overline{\Omega})^n$ implies the convergence of $\{\psi_k^{(j)}\}_{j \in \mathbb{N}}$ to $\hat{\psi}_k$ in $C^{1,\alpha}(\overline{\Omega})^n$ for all $k = 0, \dots, K-1$. Hence, for all $k = 0, \dots, K-1$,

$$w_k^{(j)} := |\det D\psi_k^{(j)}|$$

is uniformly convergent on $\overline{\Omega}$. By construction, $w_k^{(j)} \geq \tilde{\epsilon} > 0$ such that

$$t_k^{(j)} := \frac{\sum_{i=1}^k (w_i^{(j)})^{-1}}{\sum_{i=1}^K (w_i^{(j)})^{-1}}$$

converges pointwise on $\overline{\Omega}$ as $j \rightarrow \infty$. We denote the limit by \hat{t}_k and set

$$\begin{aligned} \hat{F}_0 &:= T \circ \hat{\psi}_0, \\ \hat{F}_k(x) &:= \gamma_{\hat{F}_0(x), R(x)} \circ \hat{t}_k(x), \quad k = 1, \dots, K-1. \end{aligned}$$

Recall that it holds

$$F_k^{(j)}(x) = \gamma_{F_0^{(j)}(x), R(x)} \circ t_k^{(j)}(x).$$

Using Corollary 2.3, we see that $F_0^{(j)} = T \circ \psi_0^{(j)}$ converges in $L^2(\Omega, \mathcal{H})$ to $\hat{F}_0 = T \circ \hat{\psi}_0$. For $k = 1, \dots, K-1$, we obtain

$$\begin{aligned} d_2(F_k^{(j)}, \hat{F}_k) &\leq d_2\left(F_k^{(j)}, \gamma_{\hat{F}_0(x), R(x)} \circ t_k^{(j)}\right) + d_2\left(\gamma_{\hat{F}_0(x), R(x)} \circ t_k^{(j)}, \hat{F}_k\right) \\ &= \left(\int_{\Omega} d^2\left(\gamma_{F_0^{(j)}(x), R(x)} \circ t_k^{(j)}(x), \gamma_{\hat{F}_0(x), R(x)} \circ t_k^{(j)}(x)\right) dx \right)^{\frac{1}{2}} \\ &\quad + \left(\int_{\Omega} d^2\left(\gamma_{\hat{F}_0(x), R(x)} \circ t_k^{(j)}(x), \gamma_{\hat{F}_0(x), R(x)} \circ \hat{t}_k(x)\right) dx \right)^{\frac{1}{2}}. \end{aligned}$$

Since the distance d is jointly convex, we estimate by (2.4) that

$$d^2\left(\gamma_{F_0^{(j)}(x), R(x)} \circ t_k^{(j)}(x), \gamma_{\hat{F}_0(x), R(x)} \circ t_k^{(j)}(x)\right) \leq (1 - t_k^{(j)}(x))^2 d^2(F_0^{(j)}(x), \hat{F}_0(x)).$$

Using this inequality, property (2.2) of geodesics and that $0 < t_k^{(j)} < 1$, we conclude

$$\begin{aligned} d_2(F_k^{(j)}, \hat{F}_k) &\leq \left(\int_{\Omega} d^2(F_0^{(j)}, \hat{F}_0) (1 - t_k^{(j)})^2 dx \right)^{\frac{1}{2}} + \left(\int_{\Omega} d^2(\hat{F}_0, R) |t_k^{(j)} - \hat{t}_k|^2 dx \right)^{\frac{1}{2}} \\ &\leq \left(\int_{\Omega} d^2(F_0^{(j)}, \hat{F}_0) dx \right)^{\frac{1}{2}} + \left(\int_{\Omega} d^2(\hat{F}_0, R) |t_k^{(j)} - \hat{t}_k|^2 dx \right)^{\frac{1}{2}}. \end{aligned}$$

Since the second factor in the last integral converges pointwise to zero as $t_k^{(j)} \rightarrow \hat{t}_k$ in $C^{0,\alpha}(\bar{\Omega})$ and $d^2(\hat{F}_0, R)|t_k^{(j)} - \hat{t}_k|^2$ has the integrable upper bound $4d^2(\hat{F}_0, R)$, the dominated convergence theorem implies that $F_k^{(j)}$ converges in $L^2(\Omega, \mathcal{H})$ to \hat{F}_k , $k = 1, \dots, K-1$.

Finally, setting $\hat{I}_k := \hat{F}_k \circ \hat{\psi}_k^{-1}$, $k = 1, \dots, K-1$, we obtain

$$d_2(I_k^{(j)}, \hat{I}_k) \leq d_2(I_k^{(j)}, \hat{F}_k \circ (\psi_k^{(j)})^{-1}) + d_2(\hat{F}_k \circ (\psi_k^{(j)})^{-1}, \hat{F}_k \circ \hat{\psi}_k^{-1})$$

and since $\{\psi_k^{(j)}\}_{j \in \mathbb{N}}$ converges in $C^{1,\alpha}(\bar{\Omega})^n$ further

$$d_2(I_k^{(j)}, \hat{I}_k) \leq C(\epsilon) \left(d_2(F_k^{(j)}, \hat{F}_k) + d_2(\hat{F}_k, \hat{F}_k \circ \hat{\psi}_k^{-1} \circ \psi_k^{(j)}) \right).$$

As $F_k^{(j)}$ converges in $L^2(\Omega, \mathcal{H})$ to \hat{F}_k and $\psi_k^{(j)}$ converges uniformly to $\hat{\psi}_k$ as $j \rightarrow \infty$, we obtain together with Corollary 2.3 that the right-hand side becomes arbitrary small for j large enough. This concludes the proof. \square

Up to now, we have shown that for a given image sequence $I_0 = T, I_1, \dots, I_{K-1}$, $I_K = R \in L^2(\Omega, \mathcal{H})$ the problem

$$\min_{\varphi \in \mathcal{A}_\epsilon^K} \mathcal{J}(\mathbf{I}, \varphi)$$

has a minimizer and that for given $\varphi \in \mathcal{A}_\epsilon^K$ the problem

$$\min_{\mathbf{I} \in L^2(\Omega, \mathcal{H})^{K-1}} \mathcal{J}(\mathbf{I}, \varphi) \quad \text{subject to} \quad I_0 = T, I_K = R$$

has a unique solution. Using these two results, we can prove that a minimizer of \mathcal{J} in (2.7) exists.

Theorem 2.9. *Let $T, R \in L^2(\Omega, \mathcal{H})$ and $K \geq 2$. Then there exists a sequence $\hat{\mathbf{I}} \in L^2(\Omega, \mathcal{H})^{K-1}$ minimizing \mathcal{J} .*

Proof. Let $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$ with $\mathbf{I}^{(j)} \in L^2(\Omega, \mathcal{H})^{K-1}$ be a minimizing sequence of \mathcal{J} . Then $\mathcal{J}(\mathbf{I}^{(j)}) \leq C$ for all $j \in \mathbb{N}$. By Corollary 2.5, we find for each $\mathbf{I}^{(j)}$ a sequence of diffeomorphisms $\varphi^{(j)}$ such that

$$\mathcal{J}(\mathbf{I}^{(j)}, \varphi^{(j)}) \leq \mathcal{J}(\mathbf{I}^{(j)}, \varphi)$$

for all $\varphi \in \mathcal{A}_\epsilon^K$. We know that $\|D^m \varphi_k^{(j)}\|_{L^2(\Omega)}^2 < C/\gamma$ for all $j \in \mathbb{N}$ and $k = 1, \dots, K$. As in the first part of the proof of Theorem 2.4, we conclude that $\{\varphi_k^{(j)}\}_{j \in \mathbb{N}}$ is bounded in $W^{m,2}(\Omega)^n$ and hence there exists a subsequence $\{\varphi_k^{(j_l)}\}_{l \in \mathbb{N}}$ converging weakly in $W^{m,2}(\Omega)^n$ and strongly in $C^{1,\alpha}(\bar{\Omega})^n$, $0 < \alpha < m - 1 - \frac{n}{2}$, to $\hat{\varphi}_k$. Let us denote the whole subsequence again by $\{\varphi^{(j)}\}_{j \in \mathbb{N}}$.

Using Lemma 2.6, we can replace $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$ by the image sequence $\{\hat{\mathbf{I}}^{(j)}\}_{j \in \mathbb{N}}$, $\hat{\mathbf{I}}^{(j)} \in L^2(\Omega, \mathcal{H})^{K-1}$, minimizing $J_{\varphi^{(j)}}$ so that the energy \mathcal{J} does not increase. From

Corollary 2.8, we know that the sequence $\{\hat{\mathbf{I}}^{(j)}\}_{j \in \mathbb{N}}$ converges in $L^2(\Omega, \mathcal{H})^{K-1}$ to $\hat{\mathbf{I}}$. Thus,

$$\lim_{j \rightarrow \infty} \sum_{k=1}^K d_2^2(\hat{I}_{k-1}^{(j)} \circ \varphi_k^{(j)}, \hat{I}_k^{(j)}) = \sum_{k=1}^K d_2^2(\hat{I}_{k-1} \circ \hat{\varphi}_k, \hat{I}_k). \quad (2.15)$$

By definition of \mathcal{J} we have

$$\mathcal{J}(\hat{\mathbf{I}}) \leq \mathcal{J}(\hat{\mathbf{I}}, \hat{\varphi}).$$

Further, by (2.15), as W is lsc (see Part 4 of the proof of Theorem 2.4), and by construction of $\hat{\mathbf{I}}^{(j)}$, we obtain

$$\mathcal{J}(\hat{\mathbf{I}}) \leq \liminf_{j \rightarrow \infty} \mathcal{J}(\hat{\mathbf{I}}^{(j)}, \varphi^{(j)}) \leq \liminf_{j \rightarrow \infty} \mathcal{J}(\mathbf{I}^{(j)}, \varphi^{(j)}) = \liminf_{j \rightarrow \infty} \mathcal{J}(\mathbf{I}^{(j)}).$$

Thus, $\hat{\mathbf{I}}$ is a minimizer of \mathcal{J} . □

2.3.2 Model specification

To find an image sequence $\mathbf{I} = (I_1, \dots, I_{K-1})$ minimizing \mathcal{J} , we use an alternating minimization scheme. Starting with $\mathbf{I}^{(0)}$, the deformation sequence $\varphi^{(j)}$ is computed by

$$\begin{aligned} \varphi^{(j)} &:= \arg \min_{\varphi \in \mathcal{A}_\epsilon^K} \mathcal{J}(\mathbf{I}^{(j-1)}, \varphi) \\ &= \left(\arg \min_{\varphi_k \in \mathcal{A}_\epsilon} \mathcal{R}(\varphi_k; I_{k-1}^{(j-1)}, I_k^{(j-1)}) \right)_{k=1}^K. \end{aligned} \quad (2.16)$$

To obtain the image sequence $\mathbf{I}^{(j)}$, we compute

$$\begin{aligned} \mathbf{I}^{(j)} &:= \arg \min_{\mathbf{I} \in L^2(\Omega, \mathcal{H})^{K-1}} \mathcal{J}(\mathbf{I}, \varphi^{(j)}) \\ &= \arg \min_{\mathbf{I} \in L^2(\Omega, \mathcal{H})^{K-1}} \sum_{k=1}^K d_2^2(I_{k-1} \circ \varphi_k, I_k) \quad \text{subject to} \quad I_0 = T, I_K = R. \end{aligned} \quad (2.17)$$

While in (2.16) each of the K minimization problems can be tackled separately, (2.17) is a coupled system that can be solved using the approach in the proof of Theorem 2.7.

For computational purposes, we have to specify W in $\mathcal{R}(\varphi) = \mathcal{R}(\varphi; T, R)$. Here, we propose to use the linearized elastic potential, see [204, p. 99],

$$W(D\varphi) := \mu \operatorname{trace}((D\varphi_{\text{sym}} - \mathbb{1})^2) + \frac{\lambda}{2} \operatorname{trace}(D\varphi_{\text{sym}} - \mathbb{1})^2, \quad (2.18)$$

where $\mathbb{1}$ denotes the identity matrix. For this choice, we can reformulate the problem as finding a displacement vector field $v = (v_1, v_2)^T: \Omega \mapsto \mathbb{R}^2$ minimizing

$$\mathcal{R}(v) := \mathcal{S}(v) + \int_{\Omega} \gamma |D^m v(x)|^2 + d^2(T(x - v(x)), R(x)) \, dx, \quad (2.19)$$

where

$$\mathcal{S}(v) := \int_{\Omega} \mu \operatorname{trace} (Dv_{\text{sym}}^T Dv_{\text{sym}}) + \frac{\lambda}{2} \operatorname{trace} (Dv_{\text{sym}})^2 + \eta \|v\|_2^2 dx. \quad (2.20)$$

For $\Omega \subset \mathbb{R}^2$, we have for example

$$Dv_{\text{sym}} = \begin{pmatrix} \partial_x v_1 & \frac{1}{2}(\partial_y v_1 + \partial_x v_2) \\ \frac{1}{2}(\partial_y v_1 + \partial_x v_2) & \partial_y v_2 \end{pmatrix},$$

which is also known as the (Cauchy) strain tensor of the displacement v . As the meaning is clear from the context, we use the same notation \mathcal{R} when addressing φ and v . To get the deformation, we set $\varphi(x) := x - v(x)$.

The term (2.19) is a usual regularizer in the context of registration, see [143, 204, 226]. However, after solving K registration problems separately, the resulting deformations are coupled in the subsequent step to find an optimal image sequence for these deformations. Using $\eta > 0$ in (2.20) together with some conditions on the distance of subsequent images I_k , we can handle the condition $\det D\varphi \geq \epsilon$ as the following remark shows.

Remark 2.10. 1. Assume $\gamma, \eta > 0$. Clearly, it holds

$$\min_{\varphi \in \mathcal{A}_{\epsilon}} \mathcal{R}(\varphi; T, R) \leq \mathcal{R}(\text{id}; T, R) = d_2^2(T, R). \quad (2.21)$$

In a similar way as in the proof of Theorem 2.4, we can use the Gagliardo–Nirenberg inequality together with the compact embedding of $W^{m,2}(\Omega)$ into $C^{1,\alpha}(\bar{\Omega})$ to conclude

$$\|\varphi - \text{id}\|_{C^{1,\alpha}(\bar{\Omega})} \leq C \|\varphi - \text{id}\|_{W^{m,2}(\Omega)} \leq C(\|\varphi - \text{id}\|_{L^2(\Omega)} + \|D^m(\varphi - \text{id})\|_{L^2(\Omega)}).$$

Thus, with the regularization (2.20), the optimal solution $\hat{\varphi}$ of (2.21) must fulfill

$$\|\hat{\varphi} - \text{id}\|_{C^{1,\alpha}(\bar{\Omega})} \leq C d_2(T, R),$$

where the constant C is independent of T and R . Consequently, we obtain that $\|D\hat{\varphi} - \mathbb{1}\|_{L^\infty(\Omega)^{n,n}} \leq C d_2(T, R)$. If $d_2(T, R) \leq C_{\det}$ is sufficiently small, this implies together with the continuity of the determinant that $|\det D\hat{\varphi}| \geq \epsilon$.

2. Now, the argument needs to be extended to the whole problem

$$\min_{\varphi \in \mathcal{A}_{\epsilon}^K, \mathbf{I} \in L^2(\Omega, \mathcal{H})^{K-1}} \sum_{k=1}^K \mathcal{R}(\varphi_k; I_{k-1}, I_k)$$

with given template image $I_0 = T$ and reference image $I_K = R$. For this problem, we use the initialization $\tilde{I}_k = \gamma_{T,R}(\frac{k}{K})$ with the geodesic $\gamma_{T,R}$ and conclude

$$\min_{\varphi \in \mathcal{A}_{\epsilon}^K, \mathbf{I} \in L^2(\Omega, \mathcal{H})^{K-1}} \sum_{k=1}^K \mathcal{R}(\varphi_k; I_{k-1}, I_k) \leq \sum_{k=1}^K \mathcal{R}(\text{id}; \tilde{I}_{k-1}, \tilde{I}_k) = \frac{1}{K} d_2^2(T, R).$$

For every summand it holds $\mathcal{R}(\varphi_k; I_{k-1}, I_k) \leq \frac{1}{K} d_2^2(T, R)$ which is smaller than C_{\det} if $d_2(T, R) \leq C_{\det} \sqrt{K}$. Hence, the optimal deformations satisfy $|\det D\hat{\varphi}_k| \geq \epsilon$ if we use enough images in between.




Figure 2.4: Illustration of the staggered grid, where empty boxes mean zero movement

2.4 Minimization of the space discrete model

In practice, we have to work in a spatially discrete setting. Dealing with digital images, we propose a finite difference model. We have already used such a model as basis of a face colorization method in [226]. In the rest of this chapter, we restrict our attention to two-dimensional images $T, R: \mathcal{G} \rightarrow \mathcal{H}$ defined on the (primal) image grid $\mathcal{G} := \{1, \dots, n_1\} \times \{1, \dots, n_2\}$. The integrals appearing in our space continuous functionals on the integration domain $\overline{\Omega} := [\frac{1}{2}, n_1 + \frac{1}{2}] \times [\frac{1}{2}, n_2 + \frac{1}{2}]$ are discretized by the midpoint quadrature rule, i.e., with pixel values defined on \mathcal{G} .

2.4.1 Computation of the deformation sequence

For discretizing the operators in (2.20), we work as usual on staggered grids. For an application of mimetic grid techniques in optical flow computation see also [286]. Let

$$\mathcal{G}_d := \left\{\frac{3}{2}, \dots, n_1 - \frac{1}{2}\right\} \times \left\{\frac{3}{2}, \dots, n_2 - \frac{1}{2}\right\}$$

be the (inner) dual grid, i.e., \mathcal{G} shifted by $\frac{1}{2}$ in each direction, and

$$\mathcal{G}_1 := \left\{\frac{3}{2}, \dots, n_1 - \frac{1}{2}\right\} \times \{1, \dots, n_2\}, \quad \mathcal{G}_2 := \{1, \dots, n_1\} \times \left\{\frac{3}{2}, \dots, n_2 - \frac{1}{2}\right\}.$$

Further, we consider $v = (v_1, v_2)^T$ with $v_1: \mathcal{G}_1 \rightarrow \mathbb{R}$ and $v_2: \mathcal{G}_2 \rightarrow \mathbb{R}$. In contrast to $\varphi = \text{Id}$ on $\partial\Omega$, we allow flow along the boundary, i.e., $\langle v(x), n(x) \rangle_2 = 0$ for the outer normal $n(x)$ at $x \in \partial\mathcal{G}$. In many applications features intersect with the

boundary, see for example the green stripe in Fig. 2.5 or the top part of Fig. 2.8. Here, movement along the boundary is reasonable even though the theorems from Section 2.3 do not necessarily hold true. To circumvent the gap to the theory, we could embed the images into a larger domain by extending them with a constant value. After the computation the extension could be removed. However, this leads to a higher computational effort as larger images have to be processed and the artificial boundary might influence the original image parts.

Henceforth, we use $\langle v(x), n(x) \rangle_2 = 0$. Regarding the staggered grid, the boundary conditions read

$$\begin{aligned} v_1\left(\frac{3}{2}, x_2\right) &= v_1\left(n_1 - \frac{1}{2}, x_2\right) = 0, & x_2 \in \{1, \dots, n_2\}, \\ v_2\left(x_1, \frac{3}{2}\right) &= v_2\left(x_1, n_2 - \frac{1}{2}\right) = 0, & x_1 \in \{1, \dots, n_1\}. \end{aligned}$$

See Fig. 2.4 for an illustration. We approximate $\partial_x v_1$ for $x = (x_1, x_2)^T \in \mathcal{G}$ by

$$D_{1,x_1} v_1(x) := \begin{cases} 0 & x_1 = 1, \\ v_1\left(x_1 + \frac{1}{2}, x_2\right) - v_1\left(x_1 - \frac{1}{2}, x_2\right) & x_1 = 2, \dots, n_1 - 1, \\ 0 & x_1 = n_1, \end{cases}$$

and $\partial_y v_1$ for $x_1 \in \{1, \dots, n_1 - 1\}$, $x_2 \in \{1, \dots, n_2 - 1\}$ by

$$D_{1,x_2} v_1\left(x + \frac{1}{2}\right) = v_1\left(x_1 + \frac{1}{2}, x_2 + 1\right) - v_1\left(x_1 + \frac{1}{2}, x_2\right),$$

and similarly the derivatives of v_2 . Finally, we obtain

$$\begin{aligned} \mathcal{S}(v) &= \sum_{x \in \mathcal{G}} \mu \left((D_{1,x_1} v_1)^2(x) + (D_{2,x_1} v_2)^2(x) \right) + \frac{\lambda}{2} (D_{1,x_1} v_1(x) + D_{2,x_2} v_2(x))^2 \\ &\quad + \frac{\mu}{2} \sum_{x \in \mathcal{G}_d} (D_{1,x_2} v_1(x) + D_{2,x_1} v_2(x))^2. \end{aligned}$$

As the deformations v are living on grids, it is convenient to use matrix-vector notation. Let us first rewrite $\mathcal{S}(v)$. Using $v_1 := (v_1(x_1 + \frac{1}{2}, x_2))_{x_1, x_2=1}^{n_1-1, n_2} \in \mathbb{R}^{n_1-1, n_2}$, $v_2 := (v_2(x_1, x_2 + \frac{1}{2}))_{x_1, x_2=1}^{n_1, n_2-1} \in \mathbb{R}^{n_1, n_2-1}$ together with the matrices

$$\begin{aligned} D_{1,x_1} &:= \begin{pmatrix} 0 & & & 0 \\ -1 & 1 & & \\ & & \ddots & \\ 0 & & 0 & -1 & 1 \\ 0 & & & 0 & 0 \end{pmatrix} \in \mathbb{R}^{n_1, n_1-1}, \\ D_{1,x_2} &:= \begin{pmatrix} -1 & 1 & & & 0 \\ 0 & -1 & 1 & & \\ & & & \ddots & \\ & & & & -1 & 1 & 0 \\ 0 & & & & 0 & -1 & 1 \end{pmatrix} \in \mathbb{R}^{n_2-1, n_2}, \end{aligned}$$

2. Morphing of Manifold-Valued Images

and similarly $D_{2,x_1} \in \mathbb{R}^{n_1-1,n_1}$ and $D_{2,x_2} \in \mathbb{R}^{n_2,n_2-1}$, we obtain

$$\begin{aligned} \mathcal{S}(v) = & \mu (\|D_{1,x_1} v_1\|_F^2 + \|v_2 D_{2,x_2}^T\|_F^2 + \frac{1}{2} \|v_1 D_{1,x_2}^T + D_{2,x_1} v_2\|_F^2) \\ & + \frac{\lambda}{2} \|D_{1,x_1} v_1 + v_2 D_{2,x_2}^T\|_F^2 + \eta \|v_1\|_F^2 + \eta \|v_2\|_F^2, \end{aligned}$$

where $\|\cdot\|_F$ denotes the Frobenius norm of matrices. Reshaping v_1, v_2 columnwise into vectors $\mathbf{v}_1 \in \mathbb{R}^{(n_1-1)n_2}$ and $\mathbf{v}_2 \in \mathbb{R}^{n_1(n_2-1)}$ (where we keep the notation) and using corresponding Kronecker products \otimes of the matrices above

$$\begin{aligned} \mathbf{D}_{1,x_1} &:= \mathbf{1}_{n_2} \otimes D_{1,x_1}, & \mathbf{D}_{2,x_1} &:= \mathbf{1}_{n_2-1} \otimes D_{2,x_1}, \\ \mathbf{D}_{1,x_2} &:= D_{1,x_2} \otimes \mathbf{1}_{n_1-1}, & \mathbf{D}_{2,x_2} &:= D_{2,x_2} \otimes \mathbf{1}_{n_1}, \end{aligned} \quad (2.22)$$

the regularizing term can be rewritten as

$$\mathcal{S}(\mathbf{v}) = \|\mathbf{S}\mathbf{v}\|_2^2, \quad \mathbf{S} := \begin{pmatrix} \sqrt{\mu} \mathbf{D}_{1,x_1} & 0 \\ 0 & \sqrt{\mu} \mathbf{D}_{2,x_2} \\ \sqrt{\frac{\mu}{2}} \mathbf{D}_{1,x_2} & \sqrt{\frac{\mu}{2}} \mathbf{D}_{2,x_1} \\ \sqrt{\frac{\lambda}{2}} \mathbf{D}_{1,x_1} & \sqrt{\frac{\lambda}{2}} \mathbf{D}_{2,x_2} \\ \sqrt{\eta} \mathbf{1}_{(n_1-1)n_2} & 0 \\ 0 & \sqrt{\eta} \mathbf{1}_{n_1(n_2-1)} \end{pmatrix}, \quad \mathbf{v} = \begin{pmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{pmatrix}. \quad (2.23)$$

For the discretization of the higher order term

$$|D^m \mathbf{v}|^2 = \sum_{|\alpha|=m} \|\mathbf{D}_1^\alpha \mathbf{v}_1\|_2^2 + \|\mathbf{D}_2^\alpha \mathbf{v}_2\|_2^2,$$

we use finite difference matrices $\mathbf{D}_1^\alpha, \mathbf{D}_2^\alpha$ similar to (2.22). Setting $\mathbf{D}_1^m := (\mathbf{D}_1^\alpha)_{|\alpha|=m}$ and $\mathbf{D}_2^m := (\mathbf{D}_2^\alpha)_{|\alpha|=m}$, we extend \mathbf{S} from (2.23) such that all regularization terms are included

$$\mathbf{S} := \begin{pmatrix} \sqrt{\mu} \mathbf{D}_{1,x_1} & 0 \\ 0 & \sqrt{\mu} \mathbf{D}_{2,x_2} \\ \sqrt{\frac{\mu}{2}} \mathbf{D}_{1,x_2} & \sqrt{\frac{\mu}{2}} \mathbf{D}_{2,x_1} \\ \sqrt{\frac{\lambda}{2}} \mathbf{D}_{1,x_1} & \sqrt{\frac{\lambda}{2}} \mathbf{D}_{2,x_2} \\ \sqrt{\eta} \mathbf{1}_{(n_1-1)n_2} & 0 \\ 0 & \sqrt{\eta} \mathbf{1}_{n_1(n_2-1)} \\ \sqrt{\gamma} \mathbf{D}_1^m & 0 \\ 0 & \sqrt{\gamma} \mathbf{D}_2^m \end{pmatrix}.$$

To discretize the data term in (2.19), we need to approximate $T(x - v(x))$. Since v is not defined on \mathcal{G} , we use bilinear interpolation of v at the grid points instead, i.e., the averaged version $Pv = (P_1 v_1, P_2 v_2) := \mathcal{G} \rightarrow \mathbb{R}^2$ given by

$$P_1 v_1(x_1, x_2) := \begin{cases} 0, & x_1 = 1, \\ \frac{1}{2} (v_1(x_1 - \frac{1}{2}, x_2) + v_1(x_1 + \frac{1}{2}, x_2)), & x_1 = 2, \dots, n_1 - 1, \\ 0, & x_1 = n_1, \end{cases}$$

and similarly for $P_2 v_2$ in x_2 -direction. In matrix-vector notation the averaging operator P can be written as $Pv = (P_{n_1}^T v_1, v_2 P_{n_2})$ using

$$P_n = \frac{1}{2} \begin{pmatrix} 1 & 1 & & 0 \\ & & \ddots & \\ 0 & & 0 & 1 & 1 \end{pmatrix} \in \mathbb{R}^{n-1, n}, \quad n \in \mathbb{N}.$$

For the vectorized displacement \mathbf{v} this becomes

$$\mathbf{P}\mathbf{v} := (\mathbf{P}_1 \mathbf{v}_1, \mathbf{P}_2 \mathbf{v}_2) := ((\mathbf{1}_{n_2} \otimes P_{n_1}^T) \mathbf{v}_1, (P_{n_2}^T \otimes \mathbf{1}_{n_1}) \mathbf{v}_2).$$

In general, $x - Pv(x) \notin \mathcal{G}$ and consequently the discrete image T has to be interpolated. To this end, we use a counterpart of bilinear interpolation on manifolds that is based on a reinterpretation of the real valued case as follows. Let $f_{00}, f_{01}, f_{10}, f_{11} \in \mathbb{R}$ be the values at vertices of the unit cell. Then, the bilinear interpolation $f(x)$ at $x = (x_1, x_2)^T \in [0, 1]^2$ is given by

$$\begin{aligned} f(x) &= (1 - x_1)(1 - x_2)f_{00} + (1 - x_1)x_2f_{01} + x_1(1 - x_2)f_{10} + x_1x_2f_{11} \\ &= \arg \min_{f \in \mathbb{R}} \{ (1 - x_1)(1 - x_2)(f - f_{00})^2 + (1 - x_1)x_2(f - f_{01})^2 \\ &\quad + x_1(1 - x_2)(f - f_{10})^2 + x_1x_2(f - f_{11})^2 \}. \end{aligned}$$

The latter formulation, which expresses bilinear interpolation as mean, was applied in [225] to generalize bi- and trilinear interpolation of data $f_{00}, f_{01}, f_{10}, f_{11} \in \mathcal{H}$ using the Karcher mean with appropriate weights

$$\begin{aligned} f(x) &= \arg \min_{f \in \mathcal{H}} \{ (1 - x_1)(1 - x_2)d^2(f, f_{00}) + (1 - x_1)x_2d^2(f, f_{01}) \\ &\quad + x_1(1 - x_2)d^2(f, f_{10}) + x_1x_2d^2(f, f_{11}) \}. \end{aligned} \quad (2.24)$$

In [1], it was shown that this bilinear interpolation leads to $C^0(\Omega, \mathcal{H})$ images, which are $C^1(\Omega, \mathcal{H})$ at points $x = (x_1, x_2)^T$ with neither $x_1 \in \mathbb{Z}$ nor $x_2 \in \mathbb{Z}$. In the following, we write $T(x) \in \mathcal{H}$ for the interpolated values of $T: \mathcal{G} \rightarrow \mathcal{H}$.

In summary, the spatially discrete registration functional is given by

$$\mathcal{R}(\mathbf{v}) := \|\mathbf{S}\mathbf{v}\|_2^2 + \sum_{x \in \mathcal{G}} d^2(T(x - Pv(x)), R(x)).$$

For minimization of this functional, we want to apply a quasi-Newton method. To this end, we need the chain rule for the differential of two concatenated functions $F: \mathcal{H}_1 \rightarrow \mathcal{H}_2$ and $G: \mathcal{H}_0 \rightarrow \mathcal{H}_1$, see [184, Chap. 4], which states

$$D(F \circ G)(x)[\eta] := DF(G(x)) [DG(x)[\eta]],$$

and further the Riemannian gradient of a function $F: \mathcal{H} \rightarrow \mathbb{R}$, which is given by

$$\langle \nabla F(x), \eta \rangle_x := DF(x)[\eta] \quad \text{for all } \eta \in T_x \mathcal{H},$$

where $\langle \cdot, \cdot \rangle_x$ denotes the Riemannian metric on the tangent space $T_x \mathcal{H}$. We start with the gradient of T . Let $\lceil z \rceil$ be the smallest integer larger than z , and $\lfloor z \rfloor$ the largest integer smaller or equal than z . The derivative $\nabla_{x_1} T(x): \mathbb{R}^2 \rightarrow T\mathcal{H}$ is approximated by

$$\nabla_{x_1} T(x) = \frac{1}{\lceil x_1 \rceil - x_1} \log_{T(x)} T(\lceil x_1 \rceil, x_2). \quad (2.25)$$

Note that for $T(\lceil x_1 \rceil, x_2)$ the minimization (2.24) reduces to solving

$$0 = (x_2 - \lfloor x_2 \rfloor) \log_f T(\lceil x_1 \rceil, \lfloor x_2 \rfloor) + (x_2 - \lfloor x_2 \rfloor) \log_f T(\lceil x_1 \rceil, \lfloor x_2 \rfloor)$$

The solution $f = T(\lceil x_1 \rceil, x_2)$ is given analogously to (2.13) by

$$T(\lceil x_1 \rceil, x_2) = \exp_{T(\lceil x_1 \rceil, \lfloor x_2 \rfloor)}((x_2 - \lfloor x_2 \rfloor) \log_{T(\lceil x_1 \rceil, \lfloor x_2 \rfloor)} T(\lceil x_1 \rceil, \lfloor x_2 \rfloor)).$$

Similarly, we obtain the derivative in x_2 -direction.

Assuming $T \in C^1(\Omega, \mathcal{H})$, we can calculate the gradient of \mathcal{R} as

$$\nabla_{\mathbf{v}} \mathcal{R}(\mathbf{v}) = 2\mathbf{S}^T \mathbf{S} \mathbf{v} + \begin{pmatrix} G_1(\mathbf{v}) \\ G_2(\mathbf{v}) \end{pmatrix}, \quad (2.26)$$

where

$$G_i(\mathbf{v}) = \nabla_{\mathbf{v}_i} \left(\sum_{x \in \mathcal{G}} d^2(T(x - Pv(x)), R(x)) \right) (\mathbf{v}).$$

Next, applying the chain rule to $G_i(\mathbf{v})$ results in

$$G_i(\mathbf{v}) = -\mathbf{P}_i^T \left(\nabla_{v_i} d^2(T(x + v), R(x)) (-Pv(x)) \right)_{x \in \mathcal{G}}, \quad (2.27)$$

where $(\cdot)_{x \in \mathcal{G}}$ denotes the vectorized version of the data. Using again the chain rule, the “inner” derivatives are given for $x \in \mathcal{G}$, $v \in \mathbb{R}^2$, and $\eta \in \mathbb{R}$ by

$$\begin{aligned} & D_{v_i} \left(d^2(T(x + \cdot), R(x)) \right) (v) [\eta] \\ &= \left\langle \nabla d^2(\cdot, R(x)) (T(x + v)), D_{v_i} T(x + \cdot) (v) [\eta] \right\rangle_{T(x+v)} \\ &= \left\langle \nabla d^2(\cdot, R(x)) (T(x + v)), \eta \nabla_{x_i} T(x + v) \right\rangle_{T(x+v)} \\ &= -2\eta \left\langle \log_{T(x+v)} R(x), \nabla_{x_i} T(x + v) \right\rangle_{T(x+v)} \\ &= \eta \nabla_{v_i} d^2(T(x + \cdot), R(x)) (v). \end{aligned}$$

Plugging this into (2.27), we obtain

$$G_i(\mathbf{v}) = 2\mathbf{P}_i^T \left(\left\langle \log_{T(x - Pv(x))} R(x), \nabla_{x_i} T(x - Pv(x)) \right\rangle_{T(x - Pv(x))} \right)_{x \in \mathcal{G}}.$$

Using the gradient (2.26), we could solve the registration problems by a gradient descent algorithm with step size τ chosen by line search

$$\mathbf{v}^{(l+1)} = \mathbf{v}^{(l)} - \tau \nabla_{\mathbf{v}} \mathcal{R}(\mathbf{v}^{(l)}).$$

As \mathcal{R} is not convex, we only get convergence of a subsequence to a critical point. The result depends on the starting value, whose choice is described in Subsection 2.4.3. Since Gradient descent algorithms are known to have bad convergence rates, we employ a quasi-Newton method. In such methods, the decent direction is given by

$$\mathbf{D}(\mathbf{v}) = -\mathbf{H}^{-1}(\mathbf{v}^{(l)}) \nabla_{\mathbf{v}} \mathcal{R}(\mathbf{v}^{(l)}),$$

where $\mathbf{H}(\mathbf{v})$ is an approximation of the Hessian of \mathcal{R} . For the actual minimization of \mathcal{R} , we iterate

$$\mathbf{v}^{(l+1)} = \mathbf{v}^{(l)} + \tau \mathbf{D}(\mathbf{v}^{(l)}),$$

with τ chosen by line search. The Hessian is approximated by

$$\mathbf{H}(\mathbf{v}) = 2\mathbf{S}^T \mathbf{S} + \mathbf{J}(\mathbf{v})^T \mathbf{J}(\mathbf{v}),$$

where $\mathbf{J}(\mathbf{v}) \in \mathbb{R}^{(n_1-1)n_2+n_1(n_2-1)}$ is the Jacobian of the data term, i.e.,

$$\mathbf{J}(\mathbf{v}) = \begin{pmatrix} G_1(\mathbf{v})^T & G_2(\mathbf{v})^T \end{pmatrix}.$$

2.4.2 Computation of the image sequence

Given a displacement sequence $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_K)^T$, we can minimize $J_{\varphi} = J_{\mathbf{1}-\mathbf{v}}$ using the construction given in the proof of Theorem 2.7, in particular formula (2.13). For this purpose, we need to evaluate geodesics between the initial image T and final image R to obtain the minimizing image sequence. In particular, for each $x \in \mathcal{G}$ the geodesic $\gamma_{T \circ \psi_0(x), R(x)} \subset \mathcal{H}$ has to be evaluated at the points $t_k(x)$, $k = 1, \dots, K-1$. Here, the grids $\psi_k: \mathcal{G} \rightarrow \Omega$, $k = 0, \dots, K$, are given by

$$\psi_K(x) = x, \quad \psi_k(x) = \psi_{k+1}(x) - Pv_{k+1}(\psi_{k+1}(x)),$$

where $Pv_{k+1}(\psi_{k+1}(x))$ is obtained via bilinear interpolation on \mathbb{R}^2 . The Jacobian of $\varphi_k = \mathbf{1} - v_k$ required for the computation is calculated using forward differences and evaluated at the grid points $\psi_k(x)$ via bilinear interpolation. For each $x \in \mathcal{G}$ and $k = 1, \dots, K-1$, we compute $I_k \circ \psi_k(x)$ by evaluating the geodesics. Note that the intermediate images are calculated at scattered points in Ω , which are in general not on the grid. Finally, the desired values $I_k(x) \in \mathcal{H}$, $x \in \mathcal{G}$, are obtained by linear scattered data interpolation of manifold-valued data.

In the following remark, we detail convergence properties of the alternating minimization scheme (2.16) and (2.17) in the discrete setting.

Remark 2.11 (Convergence of the alternating minimization).

Deformation sequence: *The computation of the deformations is performed by a quasi-Newton method with appropriate step size. Hence, the deformations converge to a critical point of the functional and guarantee that the value of the functional does not increase.*

Image sequence: *The computation of the image sequence is based on the proof of Theorem 2.7, which relies on continuous integrals. For the discrete setting, we have no proof that the functional value decreases, which would imply a decrease of the whole functional \mathcal{J} and the existence of a weakly convergent subsequence of images $\mathbf{I}^{(j_i)}$ in $(\mathcal{H}^{n_1, n_2})^{K+1}$. However, in our numerical examples, we observed a decrease for non-degenerated deformation fields.*

2.4.3 Multiscale minimization scheme

Neither the energy $\mathcal{J}(\mathbf{I}, \varphi)$ nor the registration functional $\mathcal{R}(\mathbf{v})$ is convex. Hence, an initialization close to the optimal solution is desirable. As usual in optical flow and image registration, we apply a coarse-to-fine strategy. First, we iteratively smooth our given images by convolution with a truncated Gaussian and downsample using bilinear interpolation. On the coarsest level we perform a single registration to obtain a deformation. Then, we apply bilinear interpolation to construct an initial deformation on the finer level. Next, we construct $\tilde{K} - 1$, $\tilde{K} < K$, corresponding intermediate images by

$$I_k(x) = \exp_{T(x)} \left(\frac{k}{\tilde{K}} \log_{T(x)} R \circ \varphi^{-1} \left(x - \frac{k}{\tilde{K}} P v(x) \right) \right), \quad \varphi(x) = x - P v(x), \quad (2.28)$$

where T, R are the start and end images at the finer level. The inversion of $\varphi: \mathcal{G} \rightarrow \Omega$ is realized by linear scattered interpolation of the identity

$$\varphi^{-1}(\varphi_1(x_1, x_2), \varphi_1(x_1, x_2)) = (x_1, x_2).$$

Starting with this initial image sequence, we perform alternating minimization to obtain better deformations and intermediate images. After going to the next finer level by bilinear interpolation of the deformations and images, we construct more intermediate images by interpolating between neighboring ones similar to (2.28).

This procedure is repeated until we reach the maximal number K of images and the finest level. Note that going to a finer level may increase the distance between subsequent images. To keep the determinants of the optimal solution $\hat{\varphi}$ bounded away from zero, we can adjust the number of intermediate images according to Remark 2.10. The described multilevel strategy is sketched in Algorithm 2.1.

2.5 Numerical examples

In this section, we present various proof-of-concept examples. While the minimization of (2.16) and (2.17) as well as the multigrid scheme are implemented in MATLAB 2017a, the manifold-valued image processing functions, like filtering, bilinear interpolation, and interpolation of scattered data, are implemented as part of the „Manifold-valued Image Restoration Toolbox“(MVIRT)² by Ronny Bergmann and Johannes Persch. The toolbox uses C++ implementations of the basic manifold

²open source, available at <https://ronnybergmann.net/mvirt/>

Algorithm 2.1 Morphing Algorithm (informal)

```

1:  $T_0 := T, R_0 := R, \mathcal{G}_0 := \mathcal{G}$ 
2: create image stack  $(T_l)_{l=0}^{\text{lev}}, (R_l)_{l=0}^{\text{lev}}$  on  $(\mathcal{G}_l)_{l=0}^{\text{lev}}$  by smoothing and downsampling
3: solve (2.4.1) for  $T_{\text{lev}}, R_{\text{lev}}$  to get  $\tilde{v}$ 
4:  $l \rightarrow \text{lev} - 1$ 
5: use bilinear interpolation to get  $v$  on  $\mathcal{G}_l$  from  $\tilde{v}$ 
6: obtain  $\tilde{K}_l$  images  $\mathbf{I}_l^{(0)}$  from  $T_l, R_l, v$  by (2.28)
7: while  $l \geq 0$  do
8:   find images  $\tilde{\mathbf{I}}_l$  and deformations  $\tilde{\mathbf{v}}_l$  minimizing (2.9) with initialization  $\mathbf{I}_l^{(0)}$ 
9:    $l \rightarrow l - 1$ 
10:  if  $l > 0$  then
11:    use bilinear interpolation to get  $\mathbf{I}_l$  and  $\mathbf{v}_l$  on  $\mathcal{G}_l$ 
12:    for  $k = 1, \dots, \tilde{K}_l$  do
13:      compute  $\tilde{K}_l$  images connecting  $I_{l,k-1}, I_{l,k}$  using (2.28) with  $v_{l,k}$ 
14:  $\mathbf{I} := \mathbf{I}_0$ 

```

functions, like logarithmic and exponential maps, as well as the Karcher means, which are imported into MATLAB using Mex-interfaces.

In all examples, we set $m := 3$ and $\mu = \lambda = \gamma =: \alpha$. The determinants of $D\varphi_k$ in our numerical experiments stayed positive even if using $\eta = 0$.

2.5.1 Images in different color spaces

First, we are interested in the morphing path of color images in different color spaces having a nonlinear structure that is *not* a Hadamard one. We compare

i) the linear RGB color space:

the image path is calculated using $\text{lev} = 7$, where the image size is decreased to 75% per level, $\alpha = 0.00025$, and $\tilde{K}_6 = 3, \tilde{K}_5 = 2, \tilde{K}_4 = 1, \tilde{K}_l = 0, l = 1, 2, 3$, i.e., we decrease the number of new intermediate images while going to finer levels;

ii) the hue-saturation-value (HSV) color space, where the hue is phase valued, i.e., in \mathbb{S}^1 :

the image path is calculated using $\text{lev} = 5$, where the image size is decreased to 60% per level, $\alpha = 0.1$, and $\tilde{K}_5 = 3, \tilde{K}_4 = 2, \tilde{K}_3 = 1, \tilde{K}_l = 0, l = 1, 2$;

iii) the chromaticity-brightness (CB) color space, with \mathbb{S}^2 -valued chromaticity:

the image path is calculated using $\text{lev} = 8$, where the image size is decreased to 75% per level, $\alpha = 0.00025$, and $\tilde{K}_7 = 3, \tilde{K}_6 = 2, \tilde{K}_5 = 1, \tilde{K}_l = 0, l = 1, \dots, 4$.

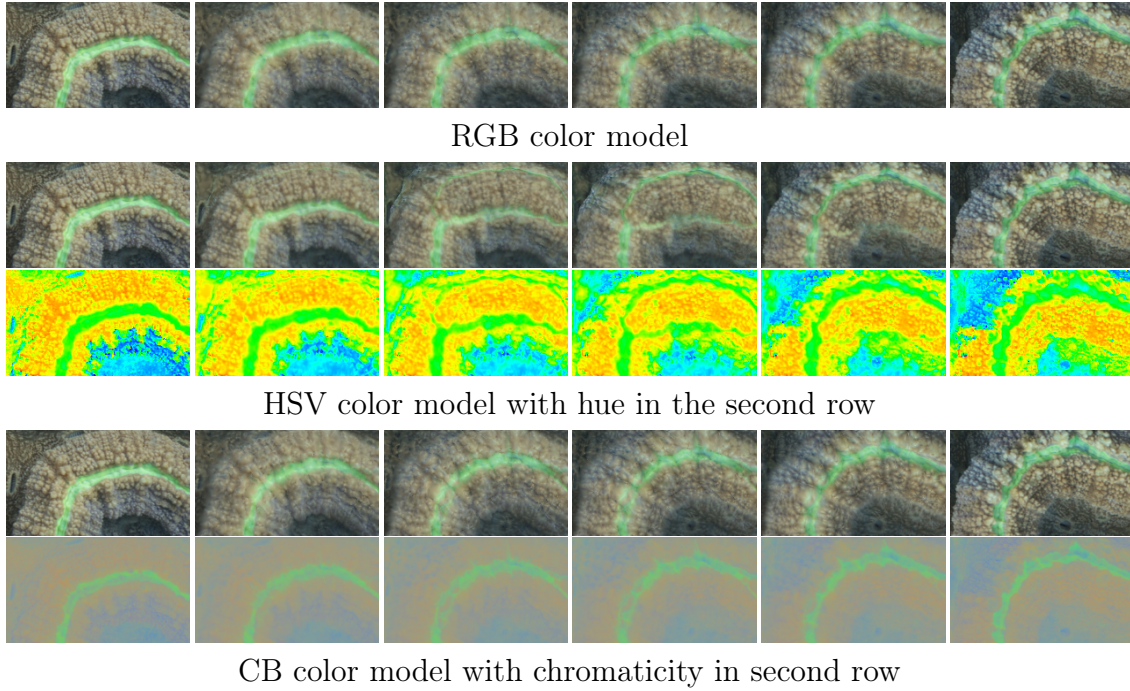



Figure 2.5: Image path between two images of a sponge, using different color models.

In Fig. 2.5 we see the image paths between two images of sponges. We calculated the image paths with 25 images, but only $I_k, k = 0, 4, 9, 14, 19, 24$, are shown. The intermediate images are blurred due to the bilinear interpolation, which could be improved for real-valued images, but the computational cost would be very high for manifold-valued images.


The morphing for the HSV model looks strange when looking at the color image, but is reasonable when considering the hue. In the hue channel the large yellow area is moving, while the green stripe is merging into and emerging out of the boundary. Here, we work on the manifold $\mathbb{S}^1 \times [0, 1]^2 \subset \mathbb{S}^1 \times \mathbb{R}^2$ with the usual product metric. Since the distances in \mathbb{S}^1 are larger than in the interval $[0, 1]$, the hue dominates the morphing. Changing the metric, i.e., the weights for the product metric, could lead to different results with more pleasant color images. Here, we stick to the usual choice to emphasize the importance of the metric. The image path of the CB model is very similar to the RGB path for this image. Looking at the chromaticity, we see that on the right part of the image a small portion of the green color vanishes and appears again close by. This effect could be reduced by lowering λ and μ , but then the deformations become close to irregular. However, on the left side the movement of the green stripe looks smooth, while the background changes as expected.

2.5.2 Symmetric positive definite matrices $\mathcal{P}(n)$

Next, we consider images with values in the manifold of symmetric positive definite $n \times n$ matrices $\mathcal{P}(n)$ with the affine invariant metric [225].



(a) Morphing path between two $\mathcal{P}(3)$ images with rectangular structures.



(b) Morphing path between two $\mathcal{P}(3)$ images with smooth rectangular structures.

Figure 2.6: Comparison of the morphing of sharp and smooth edges using the same set of parameters.

Moving $\mathcal{P}(3)$ rectangle

We start by computing a minimizing discrete path between simple synthetic images to see how edges are preserved. The template and reference images in Fig. 2.6a consist of $3I_3$ matrices in the background and a rectangular part consisting of either the matrix

$$A_T = \begin{pmatrix} 3 & 2 & 1 \\ 2 & 4 & -1 \\ 1 & -1 & 2 \end{pmatrix}, \quad \text{or } A_R = \exp_{3I_3}(2 \log_{3I_3} A_T),$$

which is moved downwards. The matrices are depicted as ellipsoids defined by their eigenvalues and eigenvectors. For this image of size 21×33 , we used $\text{lev} = 2$, where the image size is decreased to 50% per level, and $\alpha = 1$. We calculated 5 intermediate images between existing images on the new level to obtain 7 images in total. The image path looks reasonable except for the smoothing of the rectangle in vertical morphing direction and at its right boundary. The smoothing in the movement direction originates from the bilinear image interpolation model used to obtain the intermediate images and the “smoothness” of the deformations. It is possible to incorporate more sophisticated interpolation methods on manifolds, but this involves higher computational cost. The smoothing on the right side of the rectangle is an effect of the forward differences used in the calculation of the discrete deformations (2.25). This effect could also be reduced by a different discretization of the derivatives. However, for the images with slightly smoothed edges in Fig. 2.6b, our model performs well and does not produce visible artifacts.




Figure 2.7: Morphing path between two artificial $\mathcal{P}(2)$ images with whirl structures.

Whirl $\mathcal{P}(2)$ image

In Fig. 2.7 we compute the discrete path between two $\mathcal{P}(2)$ images, where the final one is constructed by deforming the start image and pushing its values further away from the identity. The artificial deformation is more complicated as in the previous example. For this image of size 64×64 , we used $\text{lev} = 4$, where the image size is decreased to 75% per level, $\alpha = 0.005$, $\tilde{K}_3 = 3$, $\tilde{K}_2 = 2$, $\tilde{K}_1 = 1$. Even though the deformation is more complicated than before, the path shows a reasonable transition from the starting to the final image.

DT-MRI

The morphing path between two slices of the Camino³ [77] is shown in Fig. 2.8. As a preprocessing we inpainted the holes and slightly smoothed the slices using a ℓ_2 – TV regularization with the Douglas-Rachford algorithm [31]. The image path is calculated using $\text{lev} = 5$, where the image size is decreased to 80% per level, $\alpha = 0.025$, and $\tilde{K}_4 = 2$, $\tilde{K}_3 = \tilde{K}_2 = 1$, $\tilde{K}_1 = 0$, i.e., we decrease the number of new intermediate images while going to finer levels. Several interesting effects occur on the path, e.g., the spot of large tensors on the right blends into the background as there is no similar structure in the target. Further, the stripe in the center moves a bit to the bottom left. While big structures on the top merge together, a small

³<http://camino.cs.ucl.ac.uk>

structure at the left boundary separates. As our implementation involves many means, the images on the morphing path are rather smooth. Hence, a comparison with the original Camino slice between template and target image is not reasonable.

2.6 Conclusions

In this chapter, we have shown how the time discrete geodesic path model from [33] can be generalized to manifold-valued images. Indeed, we have used a modified setting and have shown that at least for finite-dimensional Hadamard manifolds minimizers of the space continuous model exist.

We outlined computational details for the smooth transition of discrete images with values on a manifold, where the computations work for more general manifolds than Hadamard ones. It is left to clarify if the solution of the system of equations from the space continuous setting still leads to a decrease of the corresponding discrete functional. In all our numerical examples this seemed true. For more examples of manifold-valued images, see [29]. A generalization of Theorem 2.7 to manifolds with arbitrary curvature seems possible if the images live in compact and convex subsets of the manifold. Compactness ensures existence of a minimizer in $L^2(\Omega, \mathcal{H})$. Convexity is used for the uniqueness of shortest geodesics. Unfortunately, it seems questionable if further theoretical parts can be generalized to other manifolds.

Finally, the most interesting question is if the model converges for $K \rightarrow \infty$ to some meaningful functional such that it can be interpreted as time discretization of some geodesic path in the space of manifold-valued images. Note that we have dealt with the Mosco convergence in Hadamard spaces in [22].

Gagliardo–Nirenberg inequality

Theorem 2.12 (Gagliardo–Nirenberg [220]). *Let $\Omega \subset \mathbb{R}^n$ be a bounded domain satisfying the cone property. For $1 \leq q, r \leq \infty$, suppose that f belongs to $L^q(\Omega)$ and its derivatives of order m to $L^r(\Omega)$. Then for the derivatives $D^j f$, $0 \leq j < m$, the following inequalities hold true with constants C_1, C_2 independent of f :*

$$\|D^j f\|_{L^p(\Omega)} \leq C_1 \|D^m f\|_{L^r(\Omega)}^a \|f\|_{L^q(\Omega)}^{1-a} + C_2 \|f\|_{L^q(\Omega)},$$

where $\frac{1}{p} = \frac{j}{n} + a(\frac{1}{r} - \frac{m}{n}) + (1-a)\frac{1}{q}$ for all $a \in [\frac{j}{m}, 1]$, except for the case $1 < r < \infty$ and $m - j - \frac{n}{r}$ is a non-negative integer, in which the inequality only holds true for $a \in [\frac{j}{m}, 1)$.

Remark 2.13. For $p = q = r = 2$ the inequality simplifies to

$$\begin{aligned} \|D^j f\|_{L^2(\Omega)} &\leq C_1 \|D^m f\|_{L^2(\Omega)}^{\frac{j}{m}} \|f\|_{L^2(\Omega)}^{1-\frac{j}{m}} + C_2 \|f\|_{L^2(\Omega)} \\ &\leq C_1 \|D^m f\|_{L^2(\Omega)} + (C_1 + C_2) \|f\|_{L^2(\Omega)}, \end{aligned}$$

where the second inequality follows by estimating the product with the maximum of both factors.




Figure 2.8: Morphing path between a part of the YZ-slices 49 and 51 of the Camino dataset.

Convergence of the Time Discrete Metamorphosis Model on Hadamard Manifolds

Contents

3.1	Introduction	45
3.2	Review and preliminaries	47
3.2.1	Hadamard manifolds	48
3.2.2	Metamorphosis model in Euclidean case	51
3.2.3	Manifold-valued time discrete metamorphosis model . .	55
3.3	Manifold-valued metamorphosis model	57
3.4	Temporal extension operators	60
3.5	Mosco convergence of time discrete geodesic paths	63
3.6	Conclusion	74

Abstract

Note that this chapter¹ is published in [98]. Continuous image morphing is a classical task in image processing. The metamorphosis model proposed by Trounev, Younes and coworkers [202, 264] casts this problem in the frame of Riemannian geometry and geodesic paths between images. The associated metric in the space of images incorporates dissipation caused by a viscous flow transporting image intensities and its variations along motion paths. In many applications, images are maps from the image domain into a manifold as, e.g., for diffusion tensor imaging (DTI) into the manifold of symmetric positive definite matrices with a suitable Riemannian metric. In this chapter, we propose a generalized metamorphosis model for manifold-valued images, where the image range is a finite-dimensional Hadamard manifold. A corresponding time discrete version was presented in [213] based on the general variational time discretization proposed in [33]. Here, we prove the Mosco convergence of the time discrete metamorphosis functional to the proposed manifold-valued metamorphosis model, which implies the convergence of time discrete geodesic paths to a geodesic path in the (time continuous) metamorphosis model. In particular, the existence of geodesic paths is established. In fact, images as maps into Hadamard manifold are not only relevant in applications, but it is also shown that the joint convexity of the distance function – which characterizes Hadamard manifolds – is a crucial ingredient to establish existence of the metamorphosis model.

¹First published in [98] in 2020, published by the Society for Industrial and Applied Mathematics (SIAM). Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

3.1 Introduction

Image morphing amounts to computing a visually appealing transition of two images such that image features in the reference image are mapped to corresponding image features in the target image whenever possible.

A particular model for image morphing known as image metamorphosis was proposed by Miller, Trounevé, and Younes [202, 264, 263]. It is based on the flow of diffeomorphism model and the large deformation diffeomorphic metric mapping (LDDMM), which dates back to the work of Arnold, Dupuis, Grenander and co-workers [9, 10, 95, 25, 168, 199, 273, 272]. From the perspective of the flow of diffeomorphism model, each point of the reference image is transported to the target image in an energetically optimal way such that the image intensity is preserved along the trajectories of the pixels. The metamorphosis model additionally allows for image intensity modulations along the trajectories by incorporating the magnitude of these modulations, which is reflected by the integrated squared material derivative of the image trajectories as a penalization term in the energy functional. Recently, metamorphosis has been extended to images in reproducing kernel Hilbert spaces [232], to functional shapes [59], and to discrete measures [231]. For a more detailed exposition of these models we refer to [285, 201] and the references therein.

A variational time discretization of the metamorphosis model for images in $L^2(\Omega, \mathbb{R}^m)$ was proposed in [33]. Furthermore, existence of discrete geodesic paths and the Mosco convergence of the time discrete to the time continuous metamorphosis model was proven. The time discrete metamorphosis model has successfully been applied to a variety of imaging applications like image extrapolation [99], Bézier interpolation [100], color transfer [226] or image interpolation in a medical context [32].

Throughout the past years, manifold-valued images have received increased attention, see [21, 72, 186, 277, 30]. Some prominent applications are linked to Hadamard manifold-valued images:

- Diffusion tensor magnetic resonance imaging is an image acquisition method that incorporates in vivo magnetic resonance images of biological tissues driven by local molecular diffusion. The range space of the resulting images is frequently the space of symmetric and positive definite matrices [20, 63, 112, 267].
- Retina data is commonly modeled as images with values in the manifold of univariate non-degenerate Gaussian probability distributions endowed with the Fisher metric [8, 31]. This space is isometric to a hyperbolic space, which can be exploited numerically.

This motivates a generalization of the metamorphosis model as a Riemannian model for spaces of images. In [213], the time discrete metamorphosis model was extended to the set of image $L^2(\Omega, \mathcal{H})$, where \mathcal{H} denotes a finite-dimensional Hadamard manifold. Recall that Hadamard manifolds are Hadamard spaces with a special Riemannian structure having non-positive sectional curvature, details are given below.

In [13], it is revealed that many concepts of Banach spaces can be generalized to Hadamard spaces, which are therefore a proper choice for the analytical treatment of algorithms for manifold-valued images. In particular, the distance in Hadamard spaces is jointly convex, which implies weak lower semi-continuity of certain functionals involving the distance function. Moreover, several analytic properties of Hadamard manifolds presented in Section 3.2, which are crucial for the Mosco convergence, cease to be valid for general manifolds.

In this chapter, we prove the Mosco convergence of the manifold-valued time discrete metamorphosis energy functional originally proposed in [213] to a novel (time continuous) metamorphosis energy functional on Hadamard manifolds. Moreover, we establish the convergence of manifold-valued time discrete geodesic paths to geodesic paths in the proposed manifold-valued metamorphosis model, which coincides with the original metamorphosis energy functional in the Euclidean space. The proof of Mosco convergence in [33] incorporates as an essential ingredient a representation formula for images via integration of the weak material derivative along motion paths for the time continuous metamorphosis model in the Euclidean setting. Here, we no longer make use of such a representation formula. Indeed, our Mosco convergence result can thus be considered as a stronger result even in the case of images as pointwise maps into a Euclidean space.

Outline The chapter is organized as follows. We start with a collection of required notation and symbols in the next paragraph, including the definition of Mosco convergence. In Section 3.2, we discuss the concept of Hadamard spaces and manifolds with an emphasis on important properties of the distance map. Furthermore, we review the classical flow of diffeomorphism and metamorphosis model. Here, we already prove some continuity results on the Lagrange maps associated with a motion field. Finally, we pick up the time discrete metamorphosis model presented in [213]. Section 3.3 is devoted to the presentation of the manifold-valued metamorphosis model. Here, the key point is the suitable definition of a material derivative quantity, which is finally obtained using a variational inequality. We show that the new model for manifold-valued images coincides with the previous model in the Euclidean case. Section 3.4 introduces a method to extend time discrete image paths to time continuous ones as natural prerequisite to prove convergence of the energy functionals on discrete paths to a limit energy functional on continuous paths. Then, in Section 3.5, the main result of this chapter on Mosco convergence is stated and proved. In detail, we show the required liminf-inequality in Theorem 3.12 and the existence of recovery sequences in Theorem 3.14. This finally implies the convergence of discrete geodesic paths in Theorem 3.15 and the existence of a geodesic path for the time continuous metamorphosis model. The proofs generally follow the guideline from [33] for the classical metamorphosis model with conceptual and technical modifications in order to deal with the setup of manifold-valued images.

Notation Throughout this chapter, we assume that the image domain $\Omega \subset \mathbb{R}^n$ is bounded with Lipschitz boundary. Henceforth, we denote time continuous operators

by calligraphic letters and time discrete operators by normal letters. We denote the space of continuous functions and k -times continuously differentiable functions on the image domain Ω by $C^0(\overline{\Omega})$ and $C^k(\overline{\Omega})$, respectively. Hölder spaces of order k with exponent α are denoted by $C^{k,\alpha}(\overline{\Omega})$.

Furthermore, we use the standard notation for Lebesgue and Sobolev spaces, i.e., $L^p(\Omega)$ and $H^m(\Omega) = W^{m,2}(\Omega)$. The associated norms are denoted by $\|\cdot\|_{L^p(\Omega)}$ and $\|\cdot\|_{H^m(\Omega)}$, respectively, and the seminorm in $H^m(\Omega)$ is given by $|\cdot|_{H^m(\Omega)}$. The Sobolev (semi-)norm is defined as

$$|f|_{H^m(\Omega)} = \|D^m f\|_{L^2(\Omega)}, \quad \|f\|_{H^m(\Omega)} = \left(\sum_{j=0}^m |f|_{H^j(\Omega)}^2 \right)^{\frac{1}{2}}.$$

The space $H_0^m(\Omega)$ is the closure of $C_c^\infty(\Omega)$ with respect to $\|\cdot\|_{H^m(\Omega)}$. Derivatives are always in the strong sense, if they exist, or in the weak sense otherwise. The symmetric part of a matrix $A \in \mathbb{R}^{l,l}$ is denoted by $A^{\text{sym}} := \frac{1}{2}(A + A^\top)$. We denote by $GL^+(n)$ the elements of $GL(n)$ with positive determinant, by $\mathbb{1}$ the identity matrix, and by Id the identity map.

Mosco convergence We conclude this section with a brief review of Mosco convergence, which can be seen as a generalization of Γ -convergence. For further details we refer the reader to [84, 206].

Definition 3.1 (Mosco convergence). *Let (X, d) be a metric space and let $\{J_k\}_{k \in \mathbb{N}}$ and J be functionals mapping from X to $\overline{\mathbb{R}}$. Then the sequence J_k is said to converge to J in the sense of Mosco w.r.t. the topology induced by d if*

1. *For every sequence $\{x_k\}_{k \in \mathbb{N}} \subset X$ with $x_k \rightharpoonup x \in X$ it holds that*

$$J(x) \leq \liminf_{k \rightarrow \infty} J_k(x_k).$$

2. *For every $x \in X$ there exists a recovery sequence $\{x_k\}_{k \in \mathbb{N}} \subset X$ such that $x_k \rightarrow x \in X$ and*

$$J(x) \geq \limsup_{k \rightarrow \infty} J_k(x_k).$$

If in 1 the strong convergence of x_k to x in the topology induced by d is required, then J_k is said to Γ -converge to J w.r.t. the topology induced by d .

3.2 Review and preliminaries

In this section, we briefly present some preliminaries of Hadamard manifolds, a short introduction to the metamorphosis model in the Euclidean setting [33], and the manifold-valued time discrete metamorphosis model [213].




Figure 3.1: Comparison triangle in the Euclidean space \mathbb{R}^2 and geodesic triangle on a Hadamard manifold, in which $d(x, y) \leq \|\bar{x} - \bar{y}\|$ is satisfied (Figure adapted from [13, Fig. 1.1]).

3.2.1 Hadamard manifolds

In what follows, a short introduction of Hadamard manifolds is provided and the space of Hölder continuous functions on Hadamard manifolds is analyzed. For further details, we refer the reader to the books [13, 48, 169].

Hadamard manifolds A metric space (X, d) is *geodesic* if every two points $x, y \in X$ are connected by a shortest geodesic curve $\gamma_{x,y}: [0, 1] \rightarrow X$ that is arclength parametrized, i.e., for every $s, t \in [0, 1]$ we have

$$d(\gamma_{x,y}(s), \gamma_{x,y}(t)) = |s - t|d(\gamma_{x,y}(0), \gamma_{x,y}(1)) \quad (3.1)$$

with endpoints $\gamma_{x,y}(0) = x$ and $\gamma_{x,y}(1) = y$. A *geodesic triangle* $\triangle(p, q, r)$ in a geodesic space (X, d) is composed of the vertices $p, q, r \in X$ and three geodesics joining these points. The corresponding comparison triangle $\triangle(\bar{p}, \bar{q}, \bar{r})$ – which is unique up to isometries – is a triangle in the Euclidean space \mathbb{R}^2 with vertices $\bar{p}, \bar{q}, \bar{r} \in \mathbb{R}^2$ such that the three line segments have the same side lengths as the corresponding geodesics of $\triangle(p, q, r)$, i.e.,

$$d(p, q) = \|\bar{p} - \bar{q}\|, \quad d(p, r) = \|\bar{p} - \bar{r}\|, \quad d(r, q) = \|\bar{r} - \bar{q}\|.$$

A complete geodesic space (\mathcal{H}, d) is called a *Hadamard space* if for every geodesic triangle $\triangle(p, q, r) \in \mathcal{H}$ and $x \in \gamma_{p,r}$, $y \in \gamma_{q,r}$ we have $d(x, y) \leq \|\bar{x} - \bar{y}\|$, where \bar{x} and \bar{y} are the corresponding points in the comparison triangle $\triangle(\bar{p}, \bar{q}, \bar{r}) \in \mathbb{R}^2$, see Fig. 3.1. Geodesic spaces satisfying the latter property are also called CAT(0) spaces. By [13, Prop. 1.1.3, Cor. 1.2.5], the geometric CAT(0) condition is equivalent to (\mathcal{H}, d) being a complete geodesic space with

$$d^2(x, v) + d^2(y, w) \leq d^2(x, w) + d^2(y, v) + 2d(x, y)d(v, w)$$

for every $x, y, v, w \in \mathcal{H}$. The most prominent examples of Hadamard spaces are Hilbert spaces and *Hadamard manifolds*, which are defined as complete simply connected Riemannian manifolds with non-positive sectional curvature. Hyperbolic spaces and the manifold of positive definite matrices with the affine invariant metric are examples of Hadamard manifolds. Throughout this chapter, we exclusively consider finite-dimensional Hadamard manifolds, which ensure the existence of unique geodesic curves joining two arbitrary points. Recall that the Hopf–Rinow Theorem ceases to be true for general infinite-dimensional manifolds [176].

A function $f: \mathcal{H} \rightarrow \mathbb{R}$ is *convex* if for every $x, y \in \mathcal{H}$ the function $f \circ \gamma_{x,y}$ is convex, i.e.,

$$f \circ \gamma_{x,y}(t) \leq (1-t)f \circ \gamma_{x,y}(0) + tf \circ \gamma_{x,y}(1)$$

for all $t \in [0, 1]$. In Hadamard spaces the distance is *jointly convex* [13, Prop. 1.1.5], i.e., for two geodesics $\gamma_{x_1, x_2}, \gamma_{y_1, y_2}$ and $t \in [0, 1]$ the relation

$$d(\gamma_{x_1, x_2}(t), \gamma_{y_1, y_2}(t)) \leq (1-t)d(x_1, y_1) + td(x_2, y_2)$$

holds true. Thus, geodesics are uniquely determined by their endpoints. For a bounded sequence $\{x_n\}_{n \in \mathbb{N}} \subset \mathcal{H}$, the function $w: \mathcal{H} \rightarrow [0, +\infty)$ defined by

$$w(x; \{x_n\}_{n \in \mathbb{N}}) := \limsup_{n \rightarrow \infty} d^2(x, x_n) \quad (3.2)$$

has a unique minimizer, which is called the *asymptotic center* of $\{x_n\}_{n \in \mathbb{N}}$, cf. [13, p. 58]. A sequence $\{x_n\}_{n \in \mathbb{N}}$ is said to *converge weakly to a point* $x \in \mathcal{H}$ if it is bounded and x is the asymptotic center of each subsequence of $\{x_n\}_{n \in \mathbb{N}}$, cf. [13, p. 103]. Then, the notion of proper and (weakly) lower semi-continuous functions is analogous to Hilbert spaces.

Next, we consider the Borel σ -algebra \mathcal{B} on \mathcal{H} and on the open and bounded set $\Omega \subset \mathbb{R}^n$. A measurable map $f: \Omega \rightarrow \mathcal{H}$ belongs to $\mathcal{L}^p(\Omega, \mathcal{H})$, $p \in [1, \infty]$, if

$$d_p(f, f_a) < \infty$$

for any constant mapping $f_a(\omega) = a$ with $a \in \mathcal{H}$, where d_p is defined for two measurable maps f and g by

$$d_p(f, g) := \begin{cases} \left(\int_{\Omega} d^p(f(\omega), g(\omega)) \, d\omega \right)^{\frac{1}{p}}, & p \in [1, \infty), \\ \text{ess sup}_{\omega \in \Omega} d(f(\omega), g(\omega)), & p = \infty. \end{cases}$$

Using the equivalence relation $f \sim g$ if $d_p(f, g) = 0$, the resulting quotient space $L^p(\Omega, \mathcal{H}) := \mathcal{L}^p(\Omega, \mathcal{H}) / \sim$ equipped with d_p is a complete metric space and again a Hadamard space if $p = 2$, cf. [13, Prop. 1.2.18]. Finally, for f, g in the *weighted Bochner space* $L^2((0, 1), L^2(\Omega, \mathcal{H}), w)$ with some weight $w \in C^0([0, 1] \times \Omega, [c_1, c_2])$, $0 < c_1 < c_2$, the metric is given by

$$d_2(f, g) = \left(\int_0^1 \int_{\Omega} d(f(t, x), g(t, x))^2 w(t, x) \, dx \, dt \right)^{\frac{1}{2}}.$$

3. Convergence of the Time Discrete Metamorphosis Model

In our proposed model, we observe Hölder continuity of paths in time, which enables pointwise evaluations in time, in particular for $t = 0$ and $t = 1$. Another classical property of Lebesgue spaces also transfers to the Hadamard setting.

Lemma 3.2. *Let $f_k \in L^2((0, 1), L^2(\Omega, \mathcal{H}), w)$ be a convergent sequence with limit f . Then there exists a subsequence which converges a.e. in time as $k \rightarrow \infty$.*

Proof. Since the Chebyshev inequality implies the convergence in measure, we can apply [192, Thm. 5.2.7(i)]. \square

Next, we define subsets of Hölder continuous functions with fixed parameters $\alpha \in (0, 1]$ and $L > 0$ by

$$A_{\alpha, L, w} := \left\{ f \in L^2((0, 1), L^2(\Omega, \mathcal{H}), w) : d_2(f(s), f(t)) \leq L|t - s|^\alpha \ \forall t, s \in [0, 1] \right\}.$$

Theorem 3.3. *The set $A_{\alpha, L, w}$ is closed and convex. In particular, $A_{\alpha, L, w}$ is weakly closed.*

Proof. Closedness: Let $\{f_k\}_{k \in \mathbb{N}} \subset A_{\alpha, L, w}$ be a convergent sequence with limit f . By Lemma 3.2, we get an a.e. convergent subsequence denoted with the same indices. Assume there exists $t \in [0, 1]$ such that this sequence does not converge. Then, we can choose $s \in [0, 1]$ arbitrarily close to t with $d_2(f_k(s), f(s)) \rightarrow 0$ as $k \rightarrow \infty$, which implies

$$d_2(f_k(t), f_l(t)) \leq 2L|t - s|^\alpha + d_2(f_k(s), f_l(s))$$

for all $k, l \in \mathbb{N}$ sufficiently large. Hence, the sequence converges pointwise for every $t \in [0, 1]$. Now, the required Hölder continuity of f follows from

$$d_2(f(s), f(t)) = \lim_{k \rightarrow \infty} d_2(f_k(s), f_k(t)) \leq L|t - s|^\alpha.$$

Convexity: Given $f_1, f_2 \in A_{\alpha, L, w}$ we define a family of geodesics $r \mapsto \gamma_{f_1(s), f_2(s)}(r)$ for $s \in [0, 1]$. Then, we obtain by the joint convexity of the Hadamard metric

$$\begin{aligned} d_2(\gamma_{f_1(s), f_2(s)}(r), \gamma_{f_1(t), f_2(t)}(r)) &\leq (1 - r)d_2(f_1(s), f_1(t)) + rd_2(f_2(s), f_2(t)) \\ &\leq L|t - s|^\alpha, \end{aligned}$$

where we used that geodesics $\gamma_{f_1(s), f_2(s)}$ can be computed pointwise for any $s \in [0, 1]$. Finally, weak closedness in the Bochner space follows by [13, Lem. 3.2.1]. \square

The following lemma is exploited in the proof of Mosco convergence.

Lemma 3.4. *Let (\mathcal{H}, d) be a locally compact Hadamard space. For fixed $p \in [1, \infty)$ let $f \in L^p(\Omega, \mathcal{H})$ and $\{Y_j\}_{j \in \mathbb{N}} \subset C^1(\overline{\Omega}, \overline{\Omega})$ be a sequence of diffeomorphisms such that $|\det(DY_j)|^{-1} \leq C$ for all $j \in \mathbb{N}$, which converges to a diffeomorphism Y in $L^\infty(\Omega)^n$. Then,*

$$\limsup_{j \rightarrow \infty} d_p(f \circ Y_j, f \circ Y) = 0.$$

If in addition $Y_j \rightarrow Y$ in $C^{1, \alpha}(\overline{\Omega})^n$, then $\limsup_{j \rightarrow \infty} d_p(f \circ (Y_j)^{-1}, f \circ Y^{-1}) = 0$.

Proof. See [213, Cor. 3] and [214, Lem. 2.2.2]. \square

The generalization of this result to the space $L^2((0, 1), L^2(\Omega, \mathcal{H}))$ is straightforward.

Corollary 3.5. *Let $\{f_j\}_{j \in \mathbb{N}} \subset L^p(\Omega, \mathcal{H})$, $p \in [1, \infty)$, be a sequence that converges to f in $L^p(\Omega, \mathcal{H})$. Under the assumptions from Lemma 3.4, it holds*

$$\limsup_{j \rightarrow \infty} d_p(f_j \circ Y_j, f \circ Y) = 0 \quad \text{and} \quad \limsup_{j \rightarrow \infty} d_p(f_j \circ (Y_j)^{-1}, f \circ Y^{-1}) = 0.$$

Proof. We prove the first equation only. Using the triangle inequality, it holds that

$$\limsup_{j \rightarrow \infty} d_p(f_j \circ Y_j, f \circ Y) \leq \limsup_{j \rightarrow \infty} C d_p(f_j, f) + \limsup_{j \rightarrow \infty} d_p(f \circ Y_j, f \circ Y) = 0.$$

\square

Again, the result directly generalizes to $L^2((0, 1), L^2(\Omega, \mathcal{H}))$. For a more detailed review on Bochner spaces we refer the reader to [165].

3.2.2 Metamorphosis model in Euclidean case

In this subsection, we briefly introduce the space of images $I: \Omega \rightarrow \mathbb{R}$ equipped with a Riemannian structure from the perspective of the flow of diffeomorphism model and the metamorphosis model. For further details, we refer the reader to the literature mentioned in Section 3.1.

Flow of diffeomorphism In the flow of diffeomorphism model, the temporal evolution of each pixel of the reference image along a trajectory is determined by a *family of diffeomorphisms* $(Y(t))_{t \in [0, 1]}: \bar{\Omega} \rightarrow \mathbb{R}^n$ such that the brightness is preserved. The *brightness constancy assumption*, which is equivalent to the assertion that $t \mapsto I(t, Y(t, x))$ is constant for a.e $x \in \Omega$, is mathematically reflected by a vanishing material derivative $\frac{D}{Dt} I = \dot{I} + v \cdot DI$ along a motion path $(I(t))_{t \in [0, 1]}$ in the space of images, where $v(t) = \dot{Y}(t) \circ Y^{-1}(t)$ denotes the time dependent *Eulerian velocity*. Then, we define for a specific operator L given below the metric and the path energy associated with this family of diffeomorphisms as follows

$$g_{Y(t)}(\dot{Y}(t), \dot{Y}(t)) = \int_{\Omega} L[v(t), v(t)] dx, \quad \mathcal{E}((Y(t))_{t \in [0, 1]}) = \int_0^1 g_{Y(t)}(\dot{Y}(t), \dot{Y}(t)) dt.$$

Throughout this chapter, we consider the *higher order operator*

$$L[v(t), v(t)] = \frac{\lambda}{2} (\text{tr } \varepsilon[v])^2 + \mu \text{tr}(\varepsilon[v]^2) + \gamma |D^m v|^2, \quad (3.3)$$

where $\varepsilon[v] = (Dv)^{\text{sym}}$ refers to the symmetrized part of the Jacobian and $m > 1 + \frac{n}{2}$ as well as $\lambda, \mu, \gamma > 0$ are fixed constants. This particular choice of the operator L originates from fluid mechanics, where the metric $g_{Y(t)}$ refers to a *viscous dissipation* in a multipolar fluid model as described in [216, 136, 135].

3. Convergence of the Time Discrete Metamorphosis Model

If Y_A and Y_B are diffeomorphisms and the energy \mathcal{E} is finite for a general path $(Y(t))_{t \in [0,1]}$ with $Y(0) = Y_A$ and $Y(1) = Y_B$, then using the $H^m(\Omega)$ -coerciveness of the metric $g_{Y(t)}$ (discussed in [28, 95]) the path is already a family of diffeomorphisms. In addition, following [95], an energy minimizing velocity field v exists such that $\frac{d}{dt}Y(t, \cdot) = v(t, Y(t, \cdot))$ for every $t \in [0, 1]$. Furthermore, the corresponding path I for two input images $I_A, I_B \in L^2(\Omega)$ has the particular form $I(t, \cdot) = I_A \circ Y^{-1}(t, \cdot)$.

In what follows, we investigate diffeomorphisms induced by velocity fields in the space

$$\mathcal{V} := H^m(\Omega, \mathbb{R}^n) \cap H_0^1(\Omega, \mathbb{R}^n).$$

The following theorem relates the norm of the induced flow to the integrated norm of the associated velocity field.

Theorem 3.6. *Let $v \in L^2((0, 1), \mathcal{V})$ be a velocity field. Then, there exists a global flow $Y \in C^0([0, 1], H^m(\Omega)^n)$ such that*

$$\begin{aligned} \dot{Y}(t, x) &= v(t, Y(t, x)), \\ Y(0, x) &= x, \end{aligned} \tag{3.4}$$

for all $x \in \Omega$ and a.e. $t \in [0, 1]$. In particular, $Y(t, \cdot)$ is a diffeomorphism for all $t \in [0, 1]$. Further, for $\alpha \in [0, m - 1 - \frac{n}{2})$ the following estimate holds

$$\|Y\|_{C^0([0,1], C^{1,\alpha}(\overline{\Omega}))} + \|Y^{-1}\|_{C^0([0,1], C^{1,\alpha}(\overline{\Omega}))} \leq G \left(\int_0^1 \|v(s, \cdot)\|_{C^{1,\alpha}(\overline{\Omega})} ds \right) \tag{3.5}$$

for a continuous function $G(x) := C(x + 1) \exp(Cx)$. The solution operator from $L^2((0, 1), \mathcal{V})$ to $C^0([0, 1], H^m(\Omega)^n)$ assigning a flow Y to every velocity field v is continuous w.r.t. the weak topology in $L^2((0, 1), \mathcal{V})$ and the $C^0([0, 1] \times \overline{\Omega})$ -topology for Y .

Proof. Existence follows from [49, Thm. 4.4] and weak continuity by [263, Thm. 9]. Although the first result is stated only for \mathbb{R}^n , it is still valid in our setting due to the existence of a linear and continuous extension operator from $H^m(\Omega)$ to $H^m(\mathbb{R}^n)$, which is implied by Stein's extension theorem [254].

The estimate for the first term in (3.5) follows from [263, Lem. 7] and relies on Grönwall's inequality. Let $i \in \{1, \dots, n\}$, $t \in [0, 1]$ and $x, y \in \Omega$. Taking into account [263, Lem. 7], we obtain

$$\|Y\|_{C^0([0,1], C^1(\overline{\Omega}))} \leq C \exp\left(C \int_0^1 \|v(s, \cdot)\|_{C^{1,\alpha}(\overline{\Omega})} ds\right).$$

Applying the triangle inequality and estimating the result using the Hölder conti-

nuity of Dv and Y yields

$$\begin{aligned}
 & |\partial_i Y(t, x) - \partial_i Y(t, y)| \\
 & \leq \int_0^t |Dv(s, Y(s, x)) \cdot \partial_i Y(s, x) - Dv(s, Y(s, y)) \cdot \partial_i Y(s, y)| \, ds \\
 & \leq \int_0^t |Dv(s, Y(s, x)) - Dv(s, Y(s, y))| |\partial_i Y(s, x)| \\
 & \quad + |Dv(s, Y(s, y))| |\partial_i Y(s, x) - \partial_i Y(s, y)| \, ds \\
 & \leq \int_0^t \|v(s, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} \|Y(s, \cdot)\|_{C^1(\bar{\Omega})}^{1+\alpha} |x - y|^\alpha + \|v(s, \cdot)\|_{C^1(\bar{\Omega})} |\partial_i Y(s, x) - \partial_i Y(s, y)| \, ds \\
 & \leq G \left(\int_0^1 \|v(s, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} \, ds \right) |x - y|^\alpha + \int_0^t \|v(s, \cdot)\|_{C^1(\bar{\Omega})} |\partial_i Y(s, x) - \partial_i Y(s, y)| \, ds.
 \end{aligned}$$

By adapting the constant C in the function G , Grönwall's inequality implies

$$|\partial_i Y(t, x) - \partial_i Y(t, y)| \leq G \left(\int_0^1 \|v(s, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} \, ds \right) |x - y|^\alpha,$$

and hence G bounds the first term in (3.5). The second term is estimated similarly by noting that $Y^{-1}(t, \cdot)$ is the flow associated with the (backward) motion field $-v(1 - t, \cdot)$. This proof can be further generalized to $C^0([0, 1], C^{k,\alpha}(\bar{\Omega}))$ -norms provided that m is sufficiently large. \square

Remark 3.7. *Existence results and bounds analogous to those in Theorem 3.6 hold when replacing \mathcal{V} by $C^{1,\alpha}(\bar{\Omega})$ with zero boundary condition [285, Chap. 8]. Furthermore, the mapping $v \rightarrow Y^v$ is Lipschitz continuous in v , i.e.,*

$$\|Y^v(t, \cdot) - Y^{\tilde{v}}(t, \cdot)\|_{C^0(\bar{\Omega})} \leq (1 + C \exp(C)) \int_0^t \|v(s, \cdot) - \tilde{v}(s, \cdot)\|_{C^0(\bar{\Omega})} \, ds,$$

where $C = \int_0^t \|v(s, \cdot)\|_{C^1(\bar{\Omega})} \, ds$, see [285, Eq. (8.16)].

Metamorphosis The metamorphosis model can be regarded as a generalization of the flow of diffeomorphism model, in which the brightness constancy assumption is replaced by a quadratic penalization of the material derivative, which in particular allows for intensity modulations along the trajectories. Thus, as a first attempt, the metric and the path energy in the metamorphosis model associated with the family of images $(I(t))_{t \in [0,1]} : \bar{\Omega} \rightarrow \mathbb{R}^n$ and a penalization parameter $\delta > 0$ are defined as follows

$$g(\dot{I}, \dot{I}) = \min_{v: \bar{\Omega} \rightarrow \mathbb{R}^n} \int_{\Omega} L[v, v] + \frac{1}{\delta} \left(\frac{\partial}{\partial t} I \right)^2 \, dx, \quad \mathcal{E}(I) = \int_0^1 g(\dot{I}, \dot{I}) \, dt. \quad (3.6)$$

Consequently, the flow of diffeomorphism model is formally the limiting case of the metamorphosis model for $\delta \rightarrow 0$.

However, there are two major problems related to (3.6). In general, paths in the space of images do not exhibit any smoothness properties (neither in space nor in time). Thus, the evaluation of the material derivative $(\frac{D}{dt}I)^2$ is not well-defined. Moreover, since different pairs of velocity fields v and material derivatives $\frac{D}{dt}I$ can imply the same time derivative of the image path \dot{I} , the restriction to equivalence classes of pairs $(v, \frac{D}{dt}I)$ is required, where two pairs are equivalent if and only if they induce the same temporal change of the image path \dot{I} .

To tackle both problems, Trouvé and Younes [263] proposed a nonlinear geometric structure in the space of images $L^2(\Omega) := L^2(\Omega, \mathbb{R})$. In detail, for a given velocity field $v \in L^2((0, 1), \mathcal{V})$ and an image path $I \in L^2((0, 1), L^2(\Omega))$ the material derivative is replaced by the function $Z \in L^2((0, 1), L^2(\Omega))$ known as the *weak material derivative*, which is uniquely determined by

$$\int_0^1 \int_{\Omega} \eta Z \, dx \, dt = - \int_0^1 \int_{\Omega} (\partial_t \eta + \operatorname{div}(v\eta)) I \, dx \, dt \quad (3.7)$$

for $\eta \in C_c^\infty((0, 1) \times \Omega)$. Moreover, for all $I \in L^2(\Omega)$ the associated *tangent space* $T_I L^2(\Omega)$ is defined as $T_I L^2(\Omega) = \{I\} \times W/N_I$, where $W = \mathcal{V} \times L^2(\Omega)$ and

$$N_I = \left\{ w = (v, Z) \in W : \int_{\Omega} Z \eta + I \operatorname{div}(\eta v) \, dx = 0 \, \forall \eta \in C_c^\infty(\Omega) \right\}.$$

As usual, the associated *tangent bundle* is given by $TL^2(\Omega) = \bigcup_{I \in L^2(\Omega)} T_I L^2(\Omega)$.

Then, following Trouvé and Younes, a *regular path in the space of images*, denoted by $I \in H^1([0, 1], L^2(\Omega))$, is a curve $I \in C^0([0, 1], L^2(\Omega))$ such that there exists a measurable path $\gamma: [0, 1] \rightarrow TL^2(\Omega)$ with bounded L^2 -norm in space and time and $\pi(\gamma) = \overline{I}$, where $\pi(I, (v, Z)) = I$ refers to the projection onto the image manifold and $(I, (v, Z))$ denotes the equivalence class such that (3.7) holds for all $\eta \in C_c^\infty((0, 1) \times \Omega)$. In this chapter, we use the Lagrange formulation of this equation. Let Y be the coordinate transform given by (3.4). Then, according to [263], the weak material derivative is equivalently determined by

$$I(t, Y(t, \cdot)) - I(s, Y(s, \cdot)) = \int_s^t Z(r, Y(r, \cdot)) \, dr$$

for all $s, t \in [0, 1]$, which can be considered as a Lagrangian version of the classical material derivative. Finally, if we assume the \mathcal{V} -coercivity of the operator L , then the *metamorphosis path energy* for a regular path $I \in H^1([0, 1], L^2(\Omega))$ is defined as

$$\mathcal{E}(I) = \int_0^1 \inf_{(v, Z) \in T_{I(t)} L^2(\Omega)} \int_{\Omega} L[v, v] + \frac{1}{\delta} Z^2 \, dx \, dt. \quad (3.8)$$

The existence of energy minimizing paths in the space of images (known as *geodesic curves*), i.e., solutions of the boundary value problem

$$\min \{ \mathcal{E}(\tilde{I}) : \tilde{I} \in H^1([0, 1], L^2(\Omega)), \tilde{I}(0) = I_A, \tilde{I}(1) = I_B \}$$

for fixed images $I_A, I_B \in L^2(\Omega)$, is proven in [263]. In addition, one can prove the existence of minimizing $(v, Z) \in T_{I(t)} L^2(\Omega)$.

Note that all results in this section can be easily generalized to the space of multichannel or color images $L^2(\Omega, \mathbb{R}^C)$ for $C \geq 2$ color channels with minor modifications.

3.2.3 Manifold-valued time discrete metamorphosis model

Now, we pick up the time discrete metamorphosis model for manifold-valued images, for which the Mosco convergence is studied in this chapter. The model itself was thoroughly analyzed in [213] and extends the variational time discretization of the classical metamorphosis model proposed in [33].

Fix $\gamma, \delta, \varepsilon > 0$ and $m > 1 + \frac{n}{2}$, and let \mathcal{H} be any finite-dimensional Hadamard manifold. For two manifold-valued images $I, \tilde{I} \in L^2(\Omega, \mathcal{H})$ and an *admissible deformation*

$$\varphi \in \mathcal{A}_\varepsilon = \{\varphi \in H^m(\Omega, \Omega) : \det D\varphi > \varepsilon \text{ in } \Omega, \varphi = \text{Id on } \partial\Omega\},$$

the *time discrete energy* for pairs of images is defined as

$$\mathbf{R}(I, \tilde{I}) = \inf_{\varphi \in \mathcal{A}_\varepsilon} \mathbf{R}(I, \tilde{I}, \varphi),$$

where

$$\mathbf{R}(I, \tilde{I}, \varphi) = \int_{\Omega} W(D\varphi) + \gamma \|D^m \varphi\|^2 dx + \frac{1}{\delta} d_2^2(I, \tilde{I} \circ \varphi) \quad (3.9)$$

for an *elastic energy density* W . Here, $d_2^2(\cdot, \cdot)$ replaces the squared L^2 -norm in the time discrete metamorphosis model. The energy \mathbf{R} can be considered as a numerically feasible approximation of the squared Riemannian distance in the underlying image space [241]. Throughout this chapter, we assume that W satisfies the following conditions:

(W1) $W \in C^4(\text{GL}^+(n), \mathbb{R}_0^+)$ is polyconvex.

(W2) There exist constants $C_{W,1}, C_{W,2}, r_W > 0$ such that for all $A \in \text{GL}^+(n)$ the following growth estimates hold true:

$$W(A) \geq C_{W,1} \|A^{\text{sym}} - \mathbf{1}\|^2, \quad \text{if } \|A - \mathbf{1}\| < r_W, \quad (3.10)$$

$$W(A) \geq C_{W,2}, \quad \text{if } \|A - \mathbf{1}\| \geq r_W. \quad (3.11)$$

(W3) The energy density admits the following representation at $\mathbf{1}$:

$$W(\mathbf{1}) = 0, \quad DW(\mathbf{1}) = 0, \quad (3.12)$$

$$\frac{1}{2} D^2 W(\mathbf{1})(A, A) = \frac{\lambda}{2} (\text{tr } A)^2 + \mu \text{tr}((A^{\text{sym}})^2). \quad (3.13)$$

The assumption (W1) is required for the lower semi-continuity of the energy functional. Furthermore, (W2) enforces the convergence of the optimal deformations to the identity in the limit $K \rightarrow \infty$, where K denotes the number of time steps of our

3. Convergence of the Time Discrete Metamorphosis Model

time discrete model to be defined in (3.14). Finally, (W3) ensures the compatibility of W with the elliptic operator L , cf. (3.3). Note that (W1) and (W3) are identical to [33, (W1) and (W3)]. We recall that in [33, (W2)] a growth estimate of the form

$$W(A) \geq C(\det A)^{-s} - C$$

for $s > n - 1$ and a positive constant C instead of (W2) is assumed. This modification additionally requires essentially bounded images in order to ensure that the deformations are homeomorphic. However, in order to use the Hadamard space of square-integrable images, we have to use (W2) instead, which in particular results in diffeomorphic deformations.

The *time discrete path energy* for $K + 1$ images $\mathbf{I} = (I_0, \dots, I_K) \in L^2(\Omega, \mathcal{H})^{K+1}$, $K \geq 2$, is defined as the weighted sum of the discrete energies \mathbf{R} evaluated at consecutive images, i.e.,

$$\mathbf{J}_K(\mathbf{I}) := \inf_{\varphi := (\varphi_1, \dots, \varphi_K) \in \mathcal{A}_\varepsilon^K} \left\{ \mathbf{J}_K(\mathbf{I}, \varphi) := K \sum_{k=1}^K \mathbf{R}(I_{k-1}, I_k, \varphi_k) \right\}. \quad (3.14)$$

The scaling factor K in (3.14) is a natural choice in time discrete geodesic calculus. Indeed, if we sample a continuous path $y: [0, 1] \rightarrow \mathcal{M}$ on a Riemannian manifold (\mathcal{M}, g) at $t_{K,k} = \frac{k}{K}$ for $k = 0, \dots, K$, we obtain from Jensen's inequality

$$\begin{aligned} \sum_{k=1}^K d(y(t_{K,k-1}), y(t_{K,k}))^2 &\leq \sum_{k=1}^K \frac{1}{K} \int_{t_{K,k-1}}^{t_{K,k}} g_{y(t)}(\dot{y}(t), \dot{y}(t)) dt \\ &= \frac{1}{K} \int_0^1 g_{y(t)}(\dot{y}(t), \dot{y}(t)) dt. \end{aligned}$$

A more rigorous justification is given in [241].

For two fixed images $I_A = I_0$ and $I_B = I_K$ in $L^2(\Omega, \mathcal{H})$, any $(K + 1)$ -tuple $\mathbf{I} = (I_0, \dots, I_K) \in L^2(\Omega, \mathcal{H})^{K+1}$ is called a *discrete geodesic curve* if

$$\mathbf{J}_K(\mathbf{I}) \leq \mathbf{J}_K((I_0, \tilde{I}_1, \dots, \tilde{I}_{K-1}, I_K))$$

for all $(\tilde{I}_1, \dots, \tilde{I}_{K-1}) \in L^2(\Omega, \mathcal{H})^{K-1}$. The existence of discrete geodesic curves has been shown in [213, Sec. 3] using (W1) and the properties of the deformation set $\mathcal{A}_\varepsilon^K$. Note that in general neither the discrete geodesic curve nor the associated set of deformations is uniquely determined. The Mosco convergence of a temporal extension of \mathbf{J}_K to \mathcal{E} in the Euclidean case was proven in [33].

Fig. 3.3 shows different discrete geodesic paths for $K = 4, 8, 16$ connecting two synthesized input images of symmetric and positive definite matrices in \mathbb{R}^2 visualized in Fig. 3.2. Here, the colors quantify the geodesic anisotropy index and the eigenvectors of the matrices correspond to the principle axes of the ellipses. Further details concerning the visualization can be found in [203]. For all computations, a finite difference discretization on staggered grids as proposed in [213] was used. In particular, we experimentally observe an indication of convergence for increasing K .




Figure 3.2: The synthesized input images used in all computations, where the diffusion tensors are visualized as ellipsoids color-coded with respect to the geometric anisotropy.

3.3 Manifold-valued metamorphosis model

In this section, we propose a (time continuous) metamorphosis energy functional \mathcal{J} for manifold-valued images in $L^2(\Omega, \mathcal{H})$, where \mathcal{H} is a finite-dimensional Hadamard manifold. This functional substantially differs from the straightforward generalization

$$\inf_{(v, Z) \in \mathcal{C}(I)} \int_0^1 \int_{\Omega} L[v, v] + \frac{1}{\delta} g_I^{\mathcal{H}}(Z, Z) \, dx \, dt$$

of the classical metamorphosis functional in (3.6), where $g_I^{\mathcal{H}}$ is the Hadamard metric at position I on \mathcal{H} . Indeed, a generalization of the weak notion of the material derivative as a tangent vector $Z(x) \in T_{I(x)}\mathcal{H}$ on the Hadamard manifold via a defining equation in the context of a corresponding weak formulation is technically involved. For a given image curve $t \mapsto I(t, Y(t, \cdot))$, the associated tangential vectors at different times are in general contained in different tangent spaces and compactness of the metric $g_I^{\mathcal{H}}$ in the base point on the Hadamard manifold is not to be expected for sequences of paths in $L^2(\Omega, \mathcal{H})$. Hence, we propose a relaxation via an inequality relating distances between images along the motion path and an associated scalar material derivative z , where $z = \|Z\|$ in the Euclidean case of images in $L^2(\Omega, \mathbb{R}^C)$. Most important, this relaxed definition of the material derivative via the variational inequality (3.17) avoids the described technical difficulty in the definition. Furthermore, this formulation will turn out to be suitable for lower semi-continuity considerations that are needed to identify this energy in Section 3.5 as the Mosco limit of the above time discrete path energy and to establish existence of geodesic paths for the novel metamorphosis model. Finally, we prove equivalence of this novel energy functional with the classical metamorphosis model for \mathbb{R}^C -valued images.

The *manifold-valued metamorphosis energy* $\mathcal{J}: L^2((0, 1) \times \Omega, \mathcal{H}) \rightarrow [0, \infty]$ is




Figure 3.3: Time discrete geodesic paths for $K = 4, 8, 16$ (the input images for $K = 16$ are not depicted). Note that the images $I_{4,i}$, $I_{8,2i}$ and $I_{16,4i}$ reflect the increasing similarity expected for larger K in correspondence to the convergence result stated in this chapter.

defined as follows

$$\mathcal{J}(I) := \inf_{(v,z) \in \mathcal{C}(I)} \int_0^1 \int_{\Omega} L[v, v] + \frac{1}{\delta} z^2 \, dx \, dt. \quad (3.15)$$

Here, $\mathcal{C}(I)$ is the set of pairs $(v, z) \in L^2((0, 1), \mathcal{V}) \times L^2((0, 1), L^2(\Omega))$ such that the flow Y defined by

$$\begin{aligned} \dot{Y}(t, x) &= v(t, Y(t, x)) & \text{for } (t, x) \in [0, 1] \times \Omega, \\ Y(0, x) &= x & \text{for } x \in \Omega \end{aligned} \quad (3.16)$$

satisfies for all $t < s \in [0, 1]$ the inequality

$$d(I(t, Y(t, \cdot)), I(s, Y(s, \cdot))) \leq \int_t^s z(r, Y(r, \cdot)) dr. \quad (3.17)$$

Let us verify the equivalence of this new relaxed model with the classical metamorphosis model for \mathbb{R}^C -valued images. In the classical model, the (C -dimensional) material derivative Z is defined via the equation

$$I(t, Y(t, \cdot)) - I(s, Y(s, \cdot)) = \int_t^s Z(r, Y(r, \cdot)) dr \quad (3.18)$$

for all $t < s \in [0, 1]$, whereas the scalar material derivative z obeys the inequality

$$\|I(t, Y(t, \cdot)) - I(s, Y(s, \cdot))\| \leq \int_t^s z(r, Y(r, \cdot)) dr. \quad (3.19)$$

In fact, the equivalence is already implied by the following proposition, which in particular proves that the manifold-valued metamorphosis energy (3.15) coincides with the metamorphosis energy functional (3.8) in the case of \mathbb{R}^C -valued images.

Proposition 3.8. *For every z fulfilling (3.19), there exists a Z fulfilling (3.18) with $z \geq \|Z\|$. Vice versa, for every Z fulfilling (3.18), there exists a z fulfilling (3.19) with $z = \|Z\|$.*

Proof. For given Z , the result follows using the triangle inequality by choosing $z = \|Z\|$. To prove the converse, let z solve (3.19). Taking the L^2 -norm on both sides implies

$$\|I(t, Y(t, \cdot)) - I(s, Y(s, \cdot))\|_{L^2(\Omega)} \leq \int_s^t \|z(r, Y(r, \cdot))\|_{L^2(\Omega)} dr,$$

i.e., the function $t \mapsto I(t, Y(t, x))$ is $AC^2([0, 1], L^2(\Omega))$ in the sense of [6, Def. 1.1.1]. From [6, Rem. 1.1.3], we can additionally infer the a.e. differentiability with derivative $\widehat{Z} \in L^2((0, 1), L^2(\Omega))$ such that

$$I(t, Y(t, x)) - I(0, Y(0, x)) = \int_0^t \widehat{Z}(r, x) dr = \int_0^t Z(r, Y(r, x)) dr$$

with $Z(r, x) := \widehat{Z}(r, X(r, x))$. Here, $X(r, \cdot)$ is the spatial inverse of $Y(r, \cdot)$, which exists due to Theorem 3.6. Now, set

$$B = \{(r, x) \in [0, 1] \times \Omega : z(r, Y(r, x)) < \|Z(r, Y(r, x))\|\}$$

and assume that the Lebesgue measure of B is strictly positive. Note that B can be approximated with finite unions of disjoint semi-open cuboids [255, Thm. 1.4]. Taking into account [6, Thm. 1.1.2, Rem. 1.1.3], we get for every such cuboid $[t_1, t_2] \times D \subset [0, 1] \times \Omega$ that

$$\int_{t_1}^{t_2} \int_D \|Z(t, Y(t, x))\|^2 dx dt \leq \int_{t_1}^{t_2} \int_D z(t, Y(t, x))^2 dx dt.$$

Combining this estimate with the dominated convergence theorem, we conclude

$$\int_B \|Z(t, Y(t, x))\|^2 dx dt \leq \int_B z(t, Y(t, x))^2 dx dt.$$

This yields a contradiction to the definition of the set B . Hence, $z \geq \|Z\|$ a.e. \square

3.4 Temporal extension operators

In this section, temporal extensions of all relevant quantities required for the convergence proof of time discrete metamorphosis are proposed, which in particular allows an explicit solution to the optimality conditions (3.16) and (3.17). We remark that the subsequent construction is similar to [33] with two major modifications, namely the definitions of the interpolated image sequence (3.21) and the weak material derivative (3.24), which are both related to the manifold structure.

For fixed $K \in \mathbb{N}$, let a discrete image path $\mathbf{I}_K = (I_{K,0}, \dots, I_{K,K}) \in L^2(\Omega, \mathcal{H})^{K+1}$ be given. In [213, Sec. 3], existence of corresponding optimal deformations $\boldsymbol{\varphi}_K = (\varphi_{K,1}, \dots, \varphi_{K,K}) \in \mathcal{A}_\varepsilon^K$ satisfying (3.14) is shown. We refer to $\tau = K^{-1}$ as the *time step size* and the image $I_{K,k}$ is associated with the *time step* $t_{K,k} = k\tau$, $k = 0, \dots, K$. For $k = 1, \dots, K$, we define the *discrete transport map* $y_{K,k}: [t_{K,k-1}, t_{K,k}] \times \overline{\Omega} \rightarrow \overline{\Omega}$ as

$$y_{K,k}(t, x) := x + (t - t_{K,k-1})K(\varphi_{K,k}(x) - x).$$

If

$$\max_{k=1, \dots, K} \|\varphi_{K,k} - \text{Id}\|_{C^{1,\alpha}(\overline{\Omega})} < 1, \quad (3.20)$$

we can use [73, Thm. 5.5-1, Thm. 5.5-2] to infer that $\det(Dy_{K,k}(t, \cdot)) > 0$ holds and that $y_{K,k}(t, \cdot)$ is invertible with inverse $x_{K,k}(t, \cdot)$. The validity of this assumption is proven below and is tacitly assumed for all further considerations.

Next, the *extension operator* $I_K^{\text{ext}}: L^2(\Omega, \mathcal{H})^{K+1} \times \mathcal{A}_\varepsilon^K \rightarrow L^2([0, 1], L^2(\Omega, \mathcal{H}))$ is defined for $t \in [t_{K,k-1}, t_{K,k})$ and a.e. $x \in \Omega$ by

$$I_K^{\text{ext}}(\mathbf{I}_K, \boldsymbol{\varphi}_K)(t, x) := \gamma_{I_{K,k-1} \circ x_{K,k}(t, x), I_{K,k} \circ \varphi_{K,k} \circ x_{K,k}(t, x)}(K(t - t_{K,k-1})), \quad (3.21)$$

i.e., by a point on the geodesic connecting $I_{K,k-1} \circ x_{K,k}(t, x)$ and $I_{K,k} \circ \varphi_{K,k} \circ x_{K,k}(t, x)$ on the manifold \mathcal{H} . Thus, I_K^{ext} describes for given \mathbf{I}_K and $\boldsymbol{\varphi}_K$ a blending in geodesic sense along the transport path governed by $y_{K,k}$.

In what follows, we set $w_{K,k} = K(\varphi_{K,k} - \text{Id})$ and define the *piecewise constant (in time) velocity* $w_K = w_K(\boldsymbol{\varphi}_K) \in L^2((0, 1), \mathcal{V})$ as

$$w_K(\boldsymbol{\varphi}_K)|_{[t_{K,k-1}, t_{K,k})} := w_{K,k}.$$

Furthermore, we define the *discrete velocity field* $v_K: \mathcal{V}^K \rightarrow L^2((0, 1), C^{1,\alpha}(\overline{\Omega}))$,

$$v_K(\boldsymbol{\varphi}_K)(t, x) := K(\varphi_{K,k} \circ x_{K,k}(t, x) - x_{K,k}(t, x))$$

for $t \in [t_{K,k-1}, t_{K,k})$ and a.e. $x \in \Omega$, which is constant along time discrete paths.

Note that the extension operator v_K merely admits a $C^{1,\alpha}$ -regularity. To see this, we note that the composition of $f \in C^{1,\alpha}(\bar{\Omega})$ and $g \in C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$ is in $C^{1,\alpha}(\bar{\Omega})$ and the estimate

$$\|f \circ g\|_{C^{1,\alpha}(\bar{\Omega})} \leq \|f\|_{C^1(\bar{\Omega})} (1 + \|g\|_{C^1(\bar{\Omega})}) + [Df \circ g Dg]_\alpha \leq C \|f\|_{C^{1,\alpha}(\bar{\Omega})} (1 + \|g\|_{C^{1,\alpha}(\bar{\Omega})})^2,$$

follows from [110, Props. 1.2.4 and 1.2.7], where $[\cdot]_\alpha$ denotes the Hölder constant. Taking into account [35, Thm. 2.1], we infer that $x_{K,k}(t, \cdot) \in C^{1,\alpha}(\bar{\Omega})$ and

$$Dx_{K,k}(t, \cdot) = K^{-1} \text{Inv}(K^{-1} \mathbf{1} + (t - t_{K,k-1})(D\varphi_{K,k}(x_{K,k}(t, \cdot)) - \mathbf{1})),$$

where $\text{Inv}: GL(n) \rightarrow GL(n)$ denotes the smooth inversion operator. Since Ω is bounded and $x_{K,k}(t, \cdot)$ is a diffeomorphism, we get

$$\|x_{K,k}(t, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} \leq C + \|Dx_{K,k}(t, \cdot)\|_{C^{0,\alpha}(\bar{\Omega})} \leq C(1 + K^{-1} \max_{k=1, \dots, K} \|\varphi_{K,k} - \text{Id}\|_{C^{1,\alpha}(\bar{\Omega})}), \quad (3.22)$$

where the mean value theorem is applied to $x_{K,k}$. This implies $v_K(t, \cdot) \in C^{1,\alpha}(\bar{\Omega})$ and

$$\|v_K(t, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} \leq C \|w_K(t, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} (1 + K^{-1} \|w_K(t, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})})^2. \quad (3.23)$$

As a last preparatory step, we define the *discrete path* $Y_K: [0, 1] \times \bar{\Omega} \rightarrow \bar{\Omega}$ as the concatenation of all small diffeomorphisms $y_{K,k}$ along the motion path. In detail, the mapping is defined for $t \in [0, t_{K,1}]$ by $Y_K(t, x) := y_{K,1}(t, x)$ and then recursively for $k = 2, \dots, K$ and $t \in (t_{K,k-1}, t_{K,k}]$ by

$$Y_K(t, x) := y_{K,k}(t, Y_K(t_{K,k-1}, x))$$

for all $x \in \Omega$. The spatial inverse of Y_K is denoted by X_K . Finally, we define the *material derivative* $z_K \in L^2((0, 1), L^2(\Omega))$ for $t \in [t_{K,k-1}, t_{K,k}]$ as

$$z_K(t, x) := Kd(I_{K,k-1} \circ x_{K,k}(t, x), I_{K,k} \circ \varphi_{K,k} \circ x_{K,k}(t, x)). \quad (3.24)$$

In the following proposition, we prove that the temporal extensions of the images, the velocities, the material derivatives and the discrete paths are indeed an admissible point for the problem, i.e., they satisfy (3.16) and (3.17).

Proposition 3.9 (Admissible extension). *For $\mathbf{I}_K \in L^2(\Omega, \mathcal{H})^{K+1}$ and deformations $\varphi_K \in \mathcal{A}_\varepsilon^K$ satisfying (3.20), the tuple $(I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K), v_K(\varphi_K), Y_K, z_K)$ is a solution to (3.16) and (3.17).*

Proof. By definition, we obtain $Y_K(0, x) = x$ for all $x \in \Omega$. For $t \in [t_{K,k-1}, t_{K,k}]$ and $x \in \Omega$ we get

$$\begin{aligned} \dot{Y}_K(t, x) &= \partial_t y_{K,k}(t, Y_K(t_{K,k-1}, x)) = K(\varphi_{K,k} \circ Y_K(t_{K,k-1}, x) - Y_K(t_{K,k-1}, x)) \\ &= v_K(\varphi_K)(t, Y_K(t, x)). \end{aligned}$$

3. Convergence of the Time Discrete Metamorphosis Model

Therefore, Y_K is a solution of (3.16) in weak sense according to Remark 3.7. A short computation shows for $s \leq t \in [t_{K,k-1}, t_{K,k}]$ that

$$\begin{aligned}
& d(I_K^{\text{ext}}(\mathbf{I}_K, \boldsymbol{\varphi}_K)(t, Y_K(t, x)), I_K^{\text{ext}}(\mathbf{I}_K, \boldsymbol{\varphi}_K)(s, Y_K(s, x))) \\
&= d(\gamma_{I_{K,k-1}, I_{K,k} \circ \varphi_{K,k}}(K(t - t_{K,k-1})) \circ Y_K(t_{K,k-1}, x), \\
&\quad \gamma_{I_{K,k-1}, I_{K,k} \circ \varphi_{K,k}}(K(s - t_{K,k-1})) \circ Y_K(t_{K,k-1}, x)) \\
&= K(t - s) d(I_{K,k-1} \circ Y_K(t_{K,k-1}, x), I_{K,k} \circ \varphi_{K,k} \circ Y_K(t_{K,k-1}, x)) \\
&\leq \int_s^t z_K(r, Y_K(r, x)) dr.
\end{aligned}$$

The first equation follows from the definition of the extension operator (3.21) and for the second equation we exploit the geodesic property (3.1). Finally, the last inequality is implied by the definition of the weak material derivative (3.24). If s and t are not in the same interval, we can use the triangle inequality multiple times, which concludes the proof. \square

The next lemma allows us to bound the $H^m(\Omega)$ -norm of the displacements by a function solely depending on the energy \mathbf{R} .

Lemma 3.10. *Under the assumptions (W1) and (W2), there exists a continuous and monotonically increasing function $\theta: \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ with $\theta(0) = 0$ such that*

$$\|\varphi - \text{Id}\|_{H^m(\Omega)} \leq \theta(\mathbf{R}(I, \tilde{I}, \varphi))$$

for all $I, \tilde{I} \in L^2(\Omega, \mathcal{H})$ and all $\varphi \in \mathcal{A}_\varepsilon$. Furthermore, $\theta(x) \leq C(x + x^2)^{\frac{1}{2}}$ for some constant $C > 0$.

Proof. Set $\overline{\mathbf{R}} = \mathbf{R}(I, \tilde{I}, \varphi)$, which is defined in (3.9). The Gagliardo–Nirenberg inequality [220] implies

$$\|\varphi - \text{Id}\|_{H^m(\Omega)} \leq C(\|\varphi - \text{Id}\|_{L^2(\Omega)} + |\varphi - \text{Id}|_{H^m(\Omega)}). \quad (3.25)$$

By definition of $\overline{\mathbf{R}}$, the $H^m(\Omega)$ -seminorm of the displacement can be controlled as follows

$$|\varphi - \text{Id}|_{H^m(\Omega)} = |\varphi|_{H^m(\Omega)} \leq \sqrt{\frac{\overline{\mathbf{R}}}{\gamma}}. \quad (3.26)$$

Since $\varphi \in H^m(\Omega, \Omega)$ implies $\|\varphi - \text{Id}\|_{L^2(\Omega)} \leq 2\text{diam}(\Omega)$, this already shows for $\alpha \in (0, m - 1 - \frac{n}{2})$ that

$$\|\varphi - \text{Id}\|_{C^{1,\alpha}(\overline{\Omega})} \leq C\|\varphi - \text{Id}\|_{H^m(\Omega)} \leq C + C\sqrt{\overline{\mathbf{R}}}. \quad (3.27)$$

To control the lower order term appearing on the right-hand side of (3.25), we first define the set $\Omega' = \{x \in \Omega : \|D\varphi(x) - \text{Id}\| < r_W\}$. Then, by using (3.10) and (3.11), we obtain

$$|\Omega \setminus \Omega'| C_{W,2} \leq \int_{\Omega} W(D\varphi) dx \leq \overline{\mathbf{R}},$$

which implies $|\Omega \setminus \Omega'| \leq \frac{\bar{R}}{C_{W,2}}$. Hence, by taking into account (3.27), we deduce

$$\begin{aligned} \int_{\Omega} \|(D\varphi)^{\text{sym}} - \mathbf{1}\|^2 dx &= \int_{\Omega'} \|(D\varphi)^{\text{sym}} - \mathbf{1}\|^2 dx + \int_{\Omega \setminus \Omega'} \|(D\varphi)^{\text{sym}} - \mathbf{1}\|^2 dx \\ &\leq \int_{\Omega} \frac{W(D\varphi)}{C_{W,1}} dx + |\Omega \setminus \Omega'| (C + C\sqrt{\bar{R}})^2 \\ &\leq \frac{\bar{R}}{C_{W,1}} + \frac{\bar{R}}{C_{W,2}} (C + C\bar{R}). \end{aligned} \quad (3.28)$$

Thus, the lemma follows from (3.25), where the first term is estimated by combining Korn's inequality with (3.28), and the second term is estimated using (3.26). \square

3.5 Mosco convergence of time discrete geodesic paths

In this section, we prove the Mosco convergence of \mathbf{J}_K to \mathcal{J} as defined in (3.15) and the convergence of time discrete geodesic paths to a time continuous minimizer of \mathcal{J} . The general procedure follows the Mosco convergence proof in the Euclidean setting [33]. Nevertheless we give a comprehensive proof of the convergence result and work out the substantial differences due to the manifold setting. These differences are highlighted throughout the proof. In what follows, we pass to subsequences several times and to increase readability, we frequently avoid relabeling subsequences if obvious. As a first step, the discrete functional $\mathbf{J}_K: L^2(\Omega, \mathcal{H})^{K+1} \times \mathcal{A}_\varepsilon^K \rightarrow [0, \infty]$ is extended to a functional $\mathcal{J}_K: L^2([0, 1], L^2(\Omega, \mathcal{H})) \rightarrow [0, \infty]$ by

$$\mathcal{J}_K(I) = \begin{cases} \inf_{\bar{\varphi}_K \in \mathcal{A}_\varepsilon^K} \{ \mathbf{J}_K(\mathbf{I}_K, \bar{\varphi}_K) : I_K^{\text{ext}}(\mathbf{I}_K, \bar{\varphi}_K) = I \}, & \text{if there exist } (\mathbf{I}_K, \varphi_K) \\ & \text{with } I = I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K), \\ +\infty, & \text{else.} \end{cases} \quad (3.29)$$

The condition $I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K) = I$ has to hold pointwise for every $t \in [0, 1]$ as the involved expressions are continuous in time. In fact, this is only finite for image paths I such that a discrete image path $\mathbf{I}_K \in L^2(\Omega, \mathcal{H})^{K+1}$ and a vector of deformations $\varphi_K \in \mathcal{A}_\varepsilon^K$ with $I = I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K)$ exist. Then, the extended energy coincides with the infimum with respect to the deformation vector for fixed \mathbf{I}_K . The following lemma guarantees that the infimum is actually attained.

Lemma 3.11. *If for a given path $I \in L^2([0, 1], L^2(\Omega, \mathcal{H}))$ a discrete image path $\mathbf{I}_K \in L^2(\Omega, \mathcal{H})^{K+1}$ and a vector of deformations $\varphi_K \in \mathcal{A}_\varepsilon^K$ with $I = I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K)$ exist, then the infimum with respect to the vector of deformations in (3.29) is attained for some $\varphi_K \in \mathcal{A}_\varepsilon^K$.*

Proof. Let $\{\varphi_K^j\}_{j \in \mathbb{N}} \subset \mathcal{A}_\varepsilon^K$ be a minimizing sequence for $\varphi_K \mapsto \mathbf{J}_K(\mathbf{I}_K, \varphi_K)$ that satisfies the equality constraint $I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K^j) = I$ for every $j \in \mathbb{N}$. Due to reflexivity of $H^m(\Omega, \Omega)^K$, a subsequence (not relabeled) exists with $\varphi_K^j \rightharpoonup \bar{\varphi}_K$ in $H^m(\Omega, \Omega)^K$.

3. Convergence of the Time Discrete Metamorphosis Model

The weak lower semi-continuity and the coercivity of $\varphi_K \mapsto \mathbf{J}_K(\mathbf{I}_K, \varphi_K)$ are shown in [213, Thm. 4]. Hence, it remains to prove weak closedness of the equality constraint $I = I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K)$. Since $H^m(\Omega, \Omega)^K \hookrightarrow C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})^K$, we can infer strong convergence of $\varphi_K^j \rightarrow \varphi_K$ in $C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})^K$. Using Corollary 3.5, we conclude for every $t \in [0, 1]$ and a.e. $x \in \Omega$ that

$$I_K^{\text{ext}}(\mathbf{I}_K, \overline{\varphi}_K)(t, x) = \lim_{j \rightarrow \infty} I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K^j)(t, x) = I(t, x)$$

holds true. \square

In what follows, we always use the symbol $\overline{\varphi}_K$ for the minimizing set of deformations for given \mathbf{I}_K . The requirements for Mosco convergence (Definition 3.1) are the liminf-inequality (Theorem 3.12) and the limsup-inequality (Theorem 3.14).

Theorem 3.12 (liminf-inequality). *Under the assumptions (W1), (W2) and (W3), the time discrete path energy \mathcal{J}_K satisfies the liminf-inequality for \mathcal{J} with respect to the $L^2([0, 1], L^2(\Omega, \mathcal{H}))$ -topology.*

Proof. First, let us give a brief outline of the structure of this proof to facilitate reading. Indeed, the different steps of the proof are as follows:

1. *Identification of the image and deformation families.* In the first step, we retrieve \mathbf{I}_K and $\overline{\varphi}_K$ from the path I_K .
2. *Lower semi-continuity of the weak material derivative.* The convergence of the discrete material derivative z_K to a limiting weak material derivative z is shown and lower semi-continuity

$$\int_0^1 \int_{\Omega} z^2 \, dx \, dt \leq \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} d(I_{K,k-1}, I_{K,k} \circ \overline{\varphi}_{K,k})^2 \, dx$$

is verified.

3. *Lower semi-continuity of the viscous dissipation.* Uniform boundedness of the velocity field $w_K = w_K(\overline{\varphi}_K)$ in K is proven, which readily implies $w_K \rightharpoonup v$ in $L^2((0, 1), \mathcal{V})$. Then, the relation

$$\int_0^1 \int_{\Omega} L[v, v] \, dx \, dt \leq \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} W(D\overline{\varphi}_{K,k}) + \gamma \|D^m \overline{\varphi}_{K,k}\|^2 \, dx$$

is shown.

4. *Verifying admissibility of the limit.* In the final step, we prove that (I, v, Y, z) is a solution of (3.16) and (3.17), where Y is the flow associated with v and I is the limit image path.

1. Identification of the image and deformation vectors. Choose a sequence $I_K \in L^2([0, 1], L^2(\Omega, \mathcal{H}))$ that converges weakly to a path $I \in L^2([0, 1], L^2(\Omega, \mathcal{H}))$. If we exclude the trivial case $\liminf_{K \rightarrow \infty} \mathcal{J}_K(I_K) = \infty$ and eventually pass to a subsequence (without relabeling), we may assume

$$\mathcal{J}_K(I_K) \leq \overline{\mathcal{J}} < \infty$$

for all $K \in \mathbb{N}$. By definition of \mathcal{J}_K , this directly implies $I_K = I_K^{\text{ext}}(\mathbf{I}_K, \overline{\varphi}_K)$ with $\mathbf{I}_K = (I_{K,0}, \dots, I_{K,K}) \in L^2(\Omega, \mathcal{H})^{K+1}$ and corresponding optimal deformations $\overline{\varphi}_K = (\overline{\varphi}_{K,1}, \dots, \overline{\varphi}_{K,K}) \in \mathcal{A}_\varepsilon^K$ in (3.29), see Lemma 3.11. In particular, by incorporating Lemma 3.10, we deduce

$$\max_{k=1, \dots, K} \|\overline{\varphi}_{K,k} - \text{Id}\|_{C^{1,\alpha}(\overline{\Omega})} \leq C \max_{k=1, \dots, K} \|\overline{\varphi}_{K,k} - \text{Id}\|_{H^m(\Omega)} \leq C\theta(\overline{\mathcal{J}}K^{-1}) \leq CK^{-\frac{1}{2}}. \quad (3.30)$$

We denote by Y_K , X_K , v_K and z_K the discrete quantities associated with $\overline{\varphi}_K$ defined in Section 3.4, which exist for K sufficiently large.

2. Lower semi-continuity of the weak material derivative. Let us remark that this step resembles the first step of the proof in the Euclidean setting replacing the squared L^2 -norm by the squared distance in the Hadamard manifold.

A straightforward computation shows

$$\begin{aligned} \int_0^1 \int_\Omega z_K^2 dx dt &= \sum_{k=1}^K \int_{t_{K,k-1}}^{t_{K,k}} \int_\Omega K^2 d(I_{K,k-1} \circ x_{K,k}, I_{K,k} \circ \overline{\varphi}_{K,k} \circ x_{K,k})^2 dx dt \\ &= \sum_{k=1}^K \int_{t_{K,k-1}}^{t_{K,k}} \int_\Omega K^2 d(I_{K,k-1}, I_{K,k} \circ \overline{\varphi}_{K,k})^2 \det Dy_{K,k} dx dt. \end{aligned} \quad (3.31)$$

Next, we want to bound the difference of $\det Dy_{K,k}$ and 1 in the L^∞ -norm. Thus, we have

$$Dy_{K,k}(t, x) = \mathbb{1} + K(t - t_{K,k-1})(D\overline{\varphi}_{K,k}(x) - \mathbb{1}).$$

Then, the Lipschitz continuity of the determinant on the ball $B_r(\mathbb{1})$ with associated radius $r = \sup_K \max_{k=1, \dots, K} \|\overline{\varphi}_{K,k} - \text{Id}\|_{H^m(\Omega)} < \infty$ implies

$$\|\det Dy_{K,k} - 1\|_{L^\infty([t_{K,k-1}, t_{K,k}) \times \Omega)} \leq C \|\overline{\varphi}_{K,k} - \text{Id}\|_{C^{1,\alpha}(\overline{\Omega})}.$$

Hence, we can deduce from (3.30) and $t_{K,k} - t_{K,k-1} = K^{-1}$ that

$$\begin{aligned} &\left| \sum_{k=1}^K K^2 \int_{t_{K,k-1}}^{t_{K,k}} \int_\Omega d(I_{K,k-1}, I_{K,k} \circ \overline{\varphi}_{K,k})^2 (\det Dy_{K,k} - 1) dx dt \right| \\ &\leq \delta \overline{\mathcal{J}} C \max_{k=1, \dots, K} \|\overline{\varphi}_{K,k} - \text{Id}\|_{C^{1,\alpha}(\overline{\Omega})} \leq \delta \overline{\mathcal{J}} CK^{-\frac{1}{2}}. \end{aligned}$$

Taking into account the definition of z_K in (3.24), this ultimately leads to

$$\lim_{K \rightarrow \infty} \int_0^1 \int_\Omega z_K^2 dx dt = \lim_{K \rightarrow \infty} K \sum_{k=1}^K \int_\Omega d(I_{K,k-1}, I_{K,k} \circ \overline{\varphi}_{K,k})^2 dx.$$

This also shows the uniform boundedness of $z_K \in L^2((0, 1), L^2(\Omega))$, which implies the existence of a weakly convergent subsequence with limit $z \in L^2((0, 1), L^2(\Omega))$. Hence, using the weak lower semi-continuity of the norm, we get

$$\int_0^1 \int_{\Omega} z^2 dx dt \leq \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} z_K^2 dx dt = \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} d(I_{K,k-1}, I_{K,k} \circ \bar{\varphi}_{K,k})^2 dx.$$

3. Lower semi-continuity of the viscous dissipation. We highlight that this step differs from the corresponding step appearing in [33] due to the modification of the assumption (W2), where the overall structure persists.

Note that the velocity fields $v_K = v_K(\varphi_K)$ are not necessarily in $L^2((0, 1), \mathcal{V})$. First, we show that the sequence $w_K = w_K(\varphi_K) \in L^2((0, 1), \mathcal{V})$ is uniformly bounded in $L^2((0, 1), \mathcal{V})$. To see this, we assume that K is large enough such that $\max_{k=1, \dots, K} \|D\bar{\varphi}_{K,k} - \mathbb{1}\|_{C^0(\bar{\Omega})} < r_W$, see (W2), which is possible due to (3.30). Then, using Korn's inequality, the Poincaré inequality as well as (W2), we obtain

$$\begin{aligned} \int_0^1 \int_{\Omega} \|w_K\|^2 dx dt &\leq C \sum_{k=1}^K \int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} K^2 \|(D\bar{\varphi}_{K,k})^{\text{sym}} - \mathbb{1}\|^2 dx dt \\ &\leq CK \sum_{k=1}^K \int_{\Omega} \frac{W(D\bar{\varphi}_{K,k})}{C_{W,1}} dx \leq \frac{C\bar{\mathcal{J}}}{C_{W,1}}, \\ \int_0^1 \int_{\Omega} \|D^m w_K\|^2 dx dt &= \sum_{k=1}^K \int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} K^2 \|D^m(\bar{\varphi}_{K,k} - \text{Id})\|^2 dx dt \\ &= \sum_{k=1}^K K \int_{\Omega} \|D^m \bar{\varphi}_{K,k}\|^2 dx \leq \frac{\bar{\mathcal{J}}}{\gamma}. \end{aligned}$$

The Gagliardo–Nirenberg inequality implies that the sequence w_K is uniformly bounded in $L^2((0, 1), \mathcal{V})$. By passing to a subsequence (again labeled in the same way), we can deduce $w_K \rightharpoonup v \in L^2((0, 1), \mathcal{V})$ for $K \rightarrow \infty$.

It remains to verify the lower semi-continuity of the sum of the approximate Riemannian distances in (3.14), i.e.,

$$\int_0^1 \int_{\Omega} L[v, v] dx dt \leq \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} W(D\bar{\varphi}_{K,k}) + \gamma \|D^m \bar{\varphi}_{K,k}\|^2 dx.$$

The second order Taylor expansion around $t_{K,k-1}$ of $t \mapsto W(\mathbb{1} + (t - t_{K,k-1})Dw_{K,k})$ evaluated at $t = t_{K,k}$ yields

$$\begin{aligned} W(D\bar{\varphi}_{K,k}) &= W(\mathbb{1}) + K^{-1}DW(\mathbb{1})(Dw_{K,k}) + \frac{1}{2K^2}D^2W(\mathbb{1})(Dw_{K,k}, Dw_{K,k}) + r_{K,k} \\ &= K^{-2} \left(\frac{\lambda}{2} \text{tr}(\varepsilon[w_{K,k}])^2 + \mu \text{tr}(\varepsilon[w_{K,k}]^2) \right) + r_{K,k}, \end{aligned} \quad (3.32)$$

where $r_{K,k}$ denotes a remainder. Here, the lower order terms vanish due to (3.12) and the last equality follows from (3.13). By Taylor's theorem, the definition

of $w_{K,k} = K(\bar{\varphi}_{K,k} - \text{Id})$ and the growth estimate given in (3.30), it follows that the remainder satisfies $\|r_{K,k}\| \leq CK^{-3}\|Dw_{K,k}\|^3$. Then,

$$\begin{aligned} & K \sum_{k=1}^K \int_{\Omega} W(D\bar{\varphi}_{K,k}) + \gamma \|D^m \bar{\varphi}_{K,k}\|^2 dx \\ &= K^{-1} \sum_{k=1}^K \int_{\Omega} \frac{\lambda}{2} \text{tr}(\varepsilon[w_{K,k}])^2 + \mu \text{tr}(\varepsilon[w_{K,k}]^2) + \gamma \|D^m w_{K,k}\|^2 dx + K \sum_{k=1}^K \int_{\Omega} r_{K,k} dx, \end{aligned}$$

and the remainder is of order $K^{-\frac{1}{2}}$. To see this, we apply (3.30), Lemma 3.10 and the uniform bound on the energy to deduce

$$\begin{aligned} & K \sum_{k=1}^K \int_{\Omega} \|r_{K,k}\| dx \leq CK \sum_{k=1}^K \int_{\Omega} K^{-3} \|Dw_{K,k}\|^3 dx \\ & \leq CK \max_{k=1,\dots,K} \|\bar{\varphi}_{K,k} - \text{Id}\|_{C^1(\bar{\Omega})} \sum_{k=1}^K \|\bar{\varphi}_{K,k} - \text{Id}\|_{H^m(\Omega)}^2 \\ & \leq CK \theta(\bar{\mathcal{J}} K^{-1}) \sum_{k=1}^K \theta(\mathbf{R}(I_{K,k-1}, I_{K,k}, \bar{\varphi}_{K,k}))^2 \\ & \leq CK^{\frac{1}{2}} \sum_{k=1}^K \mathbf{R}(I_{K,k-1}, I_{K,k}, \bar{\varphi}_{K,k}) \leq C\bar{\mathcal{J}} K^{-\frac{1}{2}}. \end{aligned}$$

Finally, a standard weak lower semi-continuity argument [83, Thm. 3.20] shows

$$\begin{aligned} & \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} W(D\bar{\varphi}_{K,k}) + \gamma \|D^m \bar{\varphi}_{K,k}\|^2 dx \\ &= \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} \frac{\lambda}{2} \text{tr}(\varepsilon[w_K])^2 + \mu \text{tr}(\varepsilon[w_K]^2) + \gamma \|D^m w_K\|^2 dx dt \\ &\geq \int_0^1 \int_{\Omega} \frac{\lambda}{2} \text{tr}(\varepsilon[v])^2 + \mu \text{tr}(\varepsilon[v]^2) + \gamma \|D^m v\|^2 dx dt, \end{aligned}$$

which implies weak lower semi-continuity of the path energy for $\{I_K\}_{K \in \mathbb{N}}$.

4. Verifying admissibility of the limit. Finally, it remains to verify that (I, v, Y, z) for a suitable Y is a solution of (3.16) and (3.17). We have already pointed out that the manifold-valued metamorphosis energy functional necessitates a variational inequality, which results in significant modifications of this step compared to [33].

Let \tilde{Y} denote the solution of

$$\begin{aligned} \frac{d}{dt} \tilde{Y}(t, x) &= v(t, \tilde{Y}(t, x)) & \text{for } (t, x) \in [0, 1] \times \Omega, \\ \tilde{Y}(0, x) &= x & \text{for } x \in \Omega, \end{aligned} \tag{3.33}$$

3. Convergence of the Time Discrete Metamorphosis Model

which exists due to Theorem 3.6. Furthermore, (3.23) and the uniform boundedness of $w_K \in L^2((0,1), \mathcal{V})$ imply that the sequence v_K is uniformly bounded in $L^2((0,1), C^{1,\alpha}(\overline{\Omega}))$. Incorporating Remark 3.7, we infer that Y_K is uniformly bounded in $C^0([0,1], C^{1,\alpha}(\overline{\Omega}))$, and by exploiting Hölder's inequality we can even show that the sequence is uniformly bounded in $C^{0,1/2}([0,1], C^{1,\alpha}(\overline{\Omega}))$. Hence, by using the compact embedding of Hölder spaces, the sequence Y_K converges strongly to some Y in $C^{0,\beta}([0,1], C^{1,\beta}(\overline{\Omega}))$ for $\beta = \frac{1}{2} \min(\frac{1}{2}, \alpha)$.

It remains to verify that $\tilde{Y} = Y$. To this end, the solutions of (3.33) corresponding to w_K are denoted by \tilde{Y}_K . Then,

$$\|Y - \tilde{Y}\|_{C^0([0,1] \times \overline{\Omega})} \leq \|Y - Y_K\|_{C^0([0,1] \times \overline{\Omega})} + \|Y_K - \tilde{Y}_K\|_{C^0([0,1] \times \overline{\Omega})} + \|\tilde{Y}_K - \tilde{Y}\|_{C^0([0,1] \times \overline{\Omega})}.$$

Here, the first term converges to zero as shown above and the last term converges to zero by the continuous dependence of \tilde{Y}_K on w_K discussed in Theorem 3.6. Then, we can estimate as follows

$$\begin{aligned} \|Y_K - \tilde{Y}_K\|_{C^0([0,1] \times \overline{\Omega})} &\leq C \sum_{k=1}^K \int_{t_{K,k-1}}^{t_{K,k}} \|w_{K,k}(s, x_{K,k}(s, \cdot)) - w_{K,k}(s, \cdot)\|_{C^0(\overline{\Omega})} ds \\ &\leq C \sum_{k=1}^K \int_{t_{K,k-1}}^{t_{K,k}} \|w_{K,k}(s, \cdot)\|_{H^m(\Omega)} \|y_{K,k}(s, \cdot) - \text{Id}\|_{C^0(\overline{\Omega})} ds \\ &\leq C \|w_K\|_{L^2((0,1), H^m(\Omega))} \max_{k=1, \dots, K} \|\bar{\varphi}_{K,k} - \text{Id}\|_{C^0(\overline{\Omega})}. \end{aligned}$$

Here, the first inequality is deduced from Remark 3.7. Further, to derive the second inequality, we exploit Lipschitz continuity of $x \mapsto w_{K,k}(s, x_{K,k}(s, x)) - w_{K,k}(s, x)$, where the Lipschitz constant is bounded by $C\|w_{K,k}(s, \cdot)\|_{H^m(\Omega)}$, and apply the coordinate transform $y_{K,k}(s, \cdot)$. The uniform control of w_K and (3.30) imply $Y = \tilde{Y}$ and by Hölder's inequality $Y \in C^{0,1/2}([0,1], C^{1,\alpha}(\overline{\Omega}))$. Finally, X_K is uniformly bounded in $C^{0,1/2}([0,1], C^{1,\alpha}(\overline{\Omega}))$ due to Remark 3.7. Thus, (3.16) is fulfilled.

Next, note that for $s, t \in [0, 1]$ we obtain

$$\begin{aligned} \int_{\Omega} d(I_K(t, Y_K(t, x)), I_K(s, Y_K(s, x)))^2 dx &\leq \int_{\Omega} \left(\int_t^s z_K(r, Y_K(r, x)) dr \right)^2 dx \\ &\leq |s - t| \left| \int_{\Omega} \int_t^s z_K(r, Y_K(r, x))^2 dr dx \right|. \end{aligned}$$

By uniform boundedness of z_K in $L^2((0,1), L^2(\Omega))$, we get $I_K \circ Y_K \in A_{\frac{1}{2}, L, |\det DY|}$ for some appropriate L . Next, we show weak convergence of a subsequence of $I_K \circ Y_K$ to $I \circ Y \in A_{\frac{1}{2}, L, |\det DY|}$. To this end, we observe

$$\begin{aligned} \limsup_{K \rightarrow \infty} d_2(I_K, I)^2 &= \limsup_{K \rightarrow \infty} \int_0^1 \int_{\Omega} d(I_K(t, Y_K(t, x)), I(t, Y_K(t, x)))^2 |\det DY_K| dx dt \\ &= \limsup_{K \rightarrow \infty} \int_0^1 \int_{\Omega} d(I_K(t, Y_K(t, x)), I(t, Y(t, x)))^2 |\det DY| dx dt. \end{aligned}$$

For the first equality, we incorporate the transformation formula, the second equality follows from the uniform convergence of DY_K , the metric triangle inequality and the convergence of $I(t, Y_K(t, x))$ to $I(t, Y(t, x))$, see Lemma 3.4. To sum up, this proves weak convergence of $I_K \circ Y_K$ according to (3.2), and by Theorem 3.3, the limit is also contained in $A_{\frac{1}{2}, L, |\det DY|}$.

Finally, it remains to verify (3.17). Assume there exist $s < t \in [0, 1]$ such that the set

$$B := \left\{ x \in \Omega : d(I(s, Y(s, x)), I(t, Y(t, x))) > \int_s^t z(r, Y(r, x)) dr \right\}$$

has positive Lebesgue measure. From the joint convexity of the metric $d(\cdot, \cdot)$ and the continuity of point evaluation in time, we infer that $I \mapsto \int_B d(I(s, x), I(t, x)) dx$ is continuous and convex on $A_{\frac{1}{2}, L, |\det DY|}$. Now, this implies weak lower semi-continuity of the mapping, see [13, Lem. 3.2.3], and we obtain

$$\begin{aligned} & \int_B d(I(s, Y(s, x)), I(t, Y(t, x))) dx \\ & \leq \liminf_{K \rightarrow \infty} \int_B d(I_K(s, Y_K(s, x)), I_K(t, Y_K(t, x))) dx \\ & \leq \liminf_{K \rightarrow \infty} \int_B \int_s^t z_K(r, Y_K(r, x)) dr dx = \int_B \int_s^t z(r, Y(r, x)) dr dx, \end{aligned}$$

where the last equality follows from the weak convergence of z_K combined with the strong convergence of Y_K , which also implies the weak convergence of $z_K \circ Y_K$. This yields a contradiction and concludes the proof of the liminf-inequality. \square

In what follows, we prove existence of a recovery sequence and thus establish Mosco convergence. As a preparation, we prove that the infimum in (3.15) is actually attained, where we exploit some results of the proof of Theorem 3.12.

Proposition 3.13. *For $I \in L^2([0, 1], L^2(\Omega, \mathcal{H}))$ with $\mathcal{J}(I) < \infty$, the infimum in (3.15) is attained, i.e., there exists a tuple $(v, z) \in \mathcal{C}(I)$ satisfying (3.16) and (3.17).*

Proof. We first observe that the functional $(v, z) \mapsto \int_0^1 \int_\Omega L[v, v] + \frac{1}{\delta} z^2 dx dt$ is weakly lower semi-continuous and coercive on $\mathcal{C}(I)$, cf. [28]. Since $\mathcal{C}(I)$ is a subset of a reflexive Banach space, it suffices to prove weak closedness of $\mathcal{C}(I)$ to obtain existence of an optimal tuple $(v, z) \in \mathcal{C}(I)$.

Let $\{(v_k, z_k)\}_{k \in \mathbb{N}} \in \mathcal{C}(I)$ be a weakly convergent sequence with limit (v, z) . Due to Theorem 3.6, the corresponding flows Y_k and Y given by (3.16) exist and $Y_k \rightarrow Y$ in $C^0([0, 1] \times \bar{\Omega})$. Further, weak convergence of v_k implies uniform boundedness of $\{v_k\}_{k \in \mathbb{N}}$ in $L^2((0, 1), C^{1, \alpha}(\bar{\Omega}))$. Thus, the reasoning in the paragraph following (3.33) implies that a subsequence of $\{Y_k\}_{k \in \mathbb{N}}$ converges strongly to Y in $C^{0, \beta}([0, 1], C^{1, \beta}(\bar{\Omega}))$ for $\beta = \frac{1}{2} \min(\frac{1}{2}, \alpha)$.

Finally, Lemma 3.4 implies $I(t, Y_k(t, x)) \rightarrow I(t, Y(t, x))$ in $L^2([0, 1], L^2(\Omega, \mathcal{H}))$ and the last part of the proof of Theorem 3.12 shows that (I, Y, z) is a solution of (3.17). \square

Theorem 3.14 (Recovery sequence). *Let $I_A, I_B \in L^2(\Omega, \mathcal{H})$ be fixed input images and let $I \in L^2([0, 1], L^2(\Omega, \mathcal{H}))$ be an image path with $I(0) = I_A$ and $I(1) = I_B$. Then there exists a recovery sequence $\{I_K\}_{K \in \mathbb{N}}$ with $I_K(0) = I_A$ and $I_K(1) = I_B$ for all $K \in \mathbb{N}$ such that the limsup-inequality in Definition 3.1 with respect to the $L^2([0, 1], L^2(\Omega, \mathcal{H}))$ -topology is valid.*

Proof. We proceed in three steps, which follow the usual general guideline to show existence of recovery sequences in the context of Γ -convergence:

1. *Construction of the recovery sequence.*
2. *Verification of the limsup-inequality.*
3. *Identification of the recovery sequence limit.*

1. Construction of the recovery sequence. Compared to [33], our construction avoids the approximation of v and defines the deformations directly. Due to Proposition 3.13, there exist optimal (v, Y, z) corresponding to I satisfying (3.16) and (3.17). Incorporating the flow Y , we define for given $K \in \mathbb{N}$ a vector of diffeomorphisms $\varphi_K = (\varphi_{K,1}, \dots, \varphi_{K,K}) \in H^m(\Omega, \mathbb{R}^n)^K$ by

$$\varphi_{K,k} = Y_{t_{K,k-1}}(t_{K,k}, \cdot),$$

where $Y_a(b, \cdot) := Y(b, Y^{-1}(a, \cdot)) \in H^m(\Omega)^K$ with $a, b \in [0, 1]$. This expression coincides with the evaluation at $t = 1$ of the flow corresponding to the velocity field $v_{a,b}(t, x) := (b - a)v(a + (b - a)t, x)$, i.e., the solution of

$$\begin{aligned} \dot{Y}_{a,b}(t, x) &= v_{a,b}(t, Y_{a,b}(t, x)) & \text{for } (t, x) \in [0, 1] \times \Omega, \\ Y_{a,b}(0, x) &= x & \text{for } x \in \Omega. \end{aligned}$$

Here, v is the velocity field whose existence is postulated in Proposition 3.13. Next, we bound the $C^1(\bar{\Omega})$ -norm of the displacements as follows

$$\begin{aligned} & \max_{k \in \{1, \dots, K\}} \|\varphi_{K,k} - \text{Id}\|_{C^1(\bar{\Omega})} \\ & \leq \sup_{\substack{s, t \in [0, 1] \\ |t-s| \leq K^{-1}}} \|Y_s(t, \cdot) - \text{Id}\|_{C^1(\bar{\Omega})} \leq \sup_{\substack{s, t \in [0, 1] \\ |t-s| \leq K^{-1}}} C \int_0^1 \|v_{s,t}(r, Y_{s,t}(r, \cdot))\|_{H^m(\Omega)} dr \\ & \leq \sup_{\substack{s, t \in [0, 1] \\ |t-s| \leq K^{-1}}} C \left| \int_s^t \|v(r, \cdot)\|_{H^m(\Omega)} dr \right| \leq CK^{-\frac{1}{2}} \sup_{\substack{s, t \in [0, 1] \\ |t-s| \leq K^{-1}}} \left| \int_s^t \|v(r, \cdot)\|_{H^m(\Omega)}^2 dr \right|^{\frac{1}{2}}. \end{aligned} \tag{3.34}$$

For the third inequality, we exploit the estimate

$$\|v(t, Y(t, \cdot))\|_{H^m(\Omega)} \leq C \|v(t, \cdot)\|_{H^m(\Omega)}, \tag{3.35}$$

which follows from [49, Lem. 3.5] and an extension argument as shown in Theorem 3.6, and use the transformation formula. The last inequality is implied by the Cauchy–Schwarz inequality.

Choosing K sufficiently large ensures $\varphi_K \in \mathcal{A}_\varepsilon^K$ and consequently we can apply the temporal extension from Section 3.4. Finally, the recovery sequence is defined as $I_K = I_K^{\text{ext}}(\mathbf{I}_K, \varphi_K)$, where

$$\mathbf{I}_K = (I_{K,0}, I_{K,1}, \dots, I_{K,K}) = (I(t_{K,0}, \cdot), \dots, I(t_{K,K}, \cdot)).$$

2. Verification of the limsup-inequality. Note that this step shares some similarities with the corresponding step in [33] with modifications necessitated by the manifold structure and the different construction. In the following, all terms in the discrete energy $\mathbf{J}_K(\mathbf{I}_K, \varphi_K)$ are estimated separately. For any $k = 1, \dots, K$, we infer using (3.17), Jensen’s inequality and (3.34) that

$$\begin{aligned} & \int_{\Omega} d(I_{K,k-1}, I_{K,k} \circ \varphi_{K,k})^2 dx \\ &= \int_{\Omega} d(I_{K,k-1} \circ Y(t_{K,k-1}, x), I_{K,k} \circ Y(t_{K,k}, x))^2 \det(DY(t_{K,k-1}, x)) dx \\ &\leq \int_{\Omega} \left(\int_{t_{K,k-1}}^{t_{K,k}} z(s, Y(s, x)) ds \right)^2 \det(DY(t_{K,k-1}, x)) dx \\ &\leq K^{-1} \int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} z^2(s, x) \det(DY_s(t_{K,k-1}, x)) dx ds \\ &\leq K^{-1} (1 + CK^{-\frac{1}{2}}) \int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} z^2 dx ds. \end{aligned} \tag{3.36}$$

Recall that $w_{K,k} = K(\varphi_{K,k} - \text{Id})$. Now, the same Taylor argument as in (3.32) implies

$$\int_{\Omega} W(D\varphi_{K,k}) + \gamma \|D^m \varphi_{K,k}\|^2 dx \leq K^{-2} \int_{\Omega} L[w_{K,k}, w_{K,k}] dx + CK^{-3} \int_{\Omega} \|Dw_{K,k}\|^3 dx. \tag{3.37}$$

Summing over the second term on the right hand side and taking into account (3.34), we obtain

$$\sum_{k=1}^K \int_{\Omega} \|Dw_{K,k}\|^3 dx \leq CK^3 \sum_{k=1}^K \|\varphi_{K,k} - \text{Id}\|_{C^1(\overline{\Omega})}^3 \leq CK^{\frac{3}{2}}.$$

A direct application of Jensen’s inequality shows that the lower order term satisfies

$$\begin{aligned} & \int_{\Omega} L[w_{K,k}, w_{K,k}] dx \\ &= \int_{\Omega} L \left[K \int_{t_{K,k-1}}^{t_{K,k}} v(t, Y_{t_{K,k-1}}(t, x)) dt, K \int_{t_{K,k-1}}^{t_{K,k}} v(t, Y_{t_{K,k-1}}(t, x)) dt \right] dx \\ &\leq \int_{\Omega} K \int_{t_{K,k-1}}^{t_{K,k}} L[v(t, Y_{t_{K,k-1}}(t, x)), v(t, Y_{t_{K,k-1}}(t, x))] dt dx. \end{aligned}$$

3. Convergence of the Time Discrete Metamorphosis Model

By using (3.34) and $|\operatorname{tr}(AB)| \leq |\operatorname{tr}(A)| + |\operatorname{tr}(A(B - \mathbb{1}))|$ for $A, B \in \mathbb{R}^{n \times n}$ multiple times, we can estimate the part corresponding to the first summand of L , see (3.3), as follows

$$\begin{aligned}
& \int_{\Omega} \int_{t_{K,k-1}}^{t_{K,k}} \operatorname{tr} \left(D(v(t, Y_{t_{K,k-1}}(t, x))) \right)^2 dt dx \\
&= \int_{\Omega} \int_{t_{K,k-1}}^{t_{K,k}} \operatorname{tr} (Dv(t, Y_{t_{K,k-1}}(t, x)) DY_{t_{K,k-1}}(t, x))^2 dt dx \\
&\leq \int_{\Omega} \int_{t_{K,k-1}}^{t_{K,k}} \operatorname{tr} (Dv(t, Y_{t_{K,k-1}}(t, x)))^2 + \operatorname{tr} (Dv(t, Y_{t_{K,k-1}}(t, x))(\mathbb{1} - DY_{t_{K,k-1}}(t, x)))^2 \\
&\quad + 2 \left| \operatorname{tr} (Dv(t, Y_{t_{K,k-1}}(t, x))) \operatorname{tr} (Dv(t, Y_{t_{K,k-1}}(t, x))(\mathbb{1} - DY_{t_{K,k-1}}(t, x))) \right| dt dx \\
&\leq \int_{\Omega} \int_{t_{K,k-1}}^{t_{K,k}} \operatorname{tr} (Dv(t, Y_{t_{K,k-1}}(t, x)))^2 + C(1 + \|v(t, \cdot)\|_{H^m(\Omega)}^3) K^{-\frac{1}{2}} dt dx \\
&\leq \int_{\Omega} \int_{t_{K,k-1}}^{t_{K,k}} \operatorname{tr} (\varepsilon[v])^2 + C(1 + \|v(t, \cdot)\|_{H^m(\Omega)}^3) K^{-\frac{1}{2}} dt dx.
\end{aligned}$$

For the last inequality, we additionally used the transformation formula and (3.34). The second term in L is estimated analogously

$$\begin{aligned}
& \int_{\Omega} \int_{t_{K,k-1}}^{t_{K,k}} \operatorname{tr} (\varepsilon[v(t, Y_{t_{K,k-1}}(t, x))])^2 dt dx \\
&\leq \int_{\Omega} \int_{t_{K,k-1}}^{t_{K,k}} \operatorname{tr} (\varepsilon[v]^2) + C(1 + \|v(t, \cdot)\|_{H^m(\Omega)}^3) K^{-\frac{1}{2}} dt dx.
\end{aligned}$$

It remains to bound the higher order term appearing in the definition of L . To this end, we use (3.35) and the bound $\|fg\|_{H^{\tilde{m}}} \leq C\|f\|_{H^m}\|g\|_{H^{\tilde{m}}}$ for $f \in H^m(\Omega)$, $g \in H^{\tilde{m}}(\Omega)$ and any $0 \leq \tilde{m} \leq m$, see [166, Lem. 2.3], which results in the following estimates

$$\begin{aligned}
& |v(t, Y_{t_{K,k-1}}(t, \cdot))|_{H^m(\Omega)} \\
&\leq |Dv(t, Y_{t_{K,k-1}}(t, \cdot))|_{H^{m-1}(\Omega)} + \|Dv(t, Y_{t_{K,k-1}}(t, \cdot))D(Y_{t_{K,k-1}}(t, \cdot) - \operatorname{Id})\|_{H^{m-1}(\Omega)} \\
&\leq |Dv(t, Y_{t_{K,k-1}}(t, \cdot))|_{H^{m-1}(\Omega)} + C\|v(t, \cdot)\|_{H^m(\Omega)}\|Y_{t_{K,k-1}}(t, \cdot) - \operatorname{Id}\|_{H^m(\Omega)} \\
&\leq |Dv(t, Y_{t_{K,k-1}}(t, \cdot))|_{H^{m-1}(\Omega)} + C\|v(t, \cdot)\|_{H^m(\Omega)}K^{-\frac{1}{2}}.
\end{aligned}$$

Iterating this argument and applying a change of variables, we obtain for the last term of L

$$\int_{t_{K,k-1}}^{t_{K,k}} |v(t, Y_{t_{K,k-1}}(t, \cdot))|_{H^m(\Omega)}^2 dt \leq \int_{t_{K,k-1}}^{t_{K,k}} |v(t, \cdot)|_{H^m(\Omega)}^2 + C\|v(t, \cdot)\|_{H^m(\Omega)}^2 K^{-\frac{1}{2}} dt. \quad (3.38)$$

By combining the estimate (3.36) with (3.37)–(3.38) for the second inequality below, we get

$$\begin{aligned}
 \mathcal{J}_K(I_K) &\leq K \sum_{k=1}^K \int_{\Omega} W(D\varphi_{K,k}) + \gamma |D^m \varphi_{K,k}|^2 + \frac{1}{\delta} d(I_{K,k-1}, I_{K,k} \circ \varphi_{K,k})^2 dx \\
 &\leq \sum_{k=1}^K \left(\int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} L[v, v] + CK^{-1} |Dw_{K,k}|^3 + \frac{1}{\delta} (1 + CK^{-\frac{1}{2}}) z^2 dx dt \right) \\
 &\leq \int_0^1 \int_{\Omega} L[v, v] + \frac{1}{\delta} z^2 dx dt + CK^{-\frac{1}{2}} + C\frac{1}{\delta} K^{-\frac{1}{2}} = \mathcal{J}(I) + \mathcal{O}(K^{-\frac{1}{2}}),
 \end{aligned}$$

which readily implies the limsup-inequality.

3. Identification of the recovery sequence limit. It remains to verify the convergence $I_K \rightarrow I$ in $L^2([0, 1], L^2(\Omega, \mathcal{H}))$ as $K \rightarrow \infty$. To see this we estimate

$$\begin{aligned}
 &\int_0^1 \int_{\Omega} d(I(s, Y(s, x)), I_K^{\text{ext}}(\mathbf{I}_K, \boldsymbol{\varphi}_K)(s, Y_K(s, x)))^2 dx ds \\
 &= \sum_{k=1}^K \int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} d(I(s, Y(s, x)), I_K^{\text{ext}}(\mathbf{I}_K, \boldsymbol{\varphi}_K)(s, Y_K(s, x)))^2 dx ds \\
 &\leq C \sum_{k=1}^K \left(\int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} K^{-2} z^2(s, Y(s, x)) dx ds \right. \\
 &\quad \left. + \int_{t_{K,k-1}}^{t_{K,k}} \int_{\Omega} d(I_{K,k-1}(Y(t_{K,k-1}, x)), I_K^{\text{ext}}(\mathbf{I}_K, \boldsymbol{\varphi}_K)(s, Y_K(s, x)))^2 dx ds \right) \\
 &\leq CK^{-2} \|z(t, Y(t, x))\|_{L^2((0,1) \times \Omega)}^2.
 \end{aligned}$$

Here, we combined (3.17) with the Cauchy–Schwarz inequality to obtain an estimate for the term $d(I(s, Y(s, x)), I_{K,k}(Y(t_{K,k-1}, x)))$ in the first inequality and used the definition of I_K^{ext} , see (3.21), together with (3.1), (3.17) and the Cauchy–Schwarz inequality in the second inequality. Due to the convergence of Y_K to Y and Corollary 3.5, this readily implies the claimed convergence $I_K^{\text{ext}}(\mathbf{I}_K, \boldsymbol{\varphi}_K) \rightarrow I$. \square

We conclude this section with the desired convergence statement for discrete geodesic paths.

Theorem 3.15 (Convergence of discrete geodesic paths). *Let $I_A, I_B \in L^2(\Omega, \mathcal{H})$ and suppose that the assumptions (W1), (W2) and (W3) hold true. For every $K \in \mathbb{N}$ let I_K be a minimizer of \mathcal{J}_K subject to $I_K(0) = I_A$ and $I_K(1) = I_B$. Then, a subsequence of $\{I_K\}_{K \in \mathbb{N}}$ converges weakly in $L^2([0, 1], L^2(\Omega, \mathcal{H}))$ to a minimizer of the continuous path energy \mathcal{J} as $K \rightarrow \infty$, and the associated sequence of discrete energies converges to the minimal continuous path energy.*

Proof. Using a comparison argument with $v \equiv 0$ and $z(t, x) = d_2(I_A(x), I_B(x))$, we deduce that the path energy \mathcal{J}_K is bounded by $\overline{\mathcal{J}} = \frac{1}{\delta} d_2(I_A, I_B)^2$. For optimal

vectors of images \mathbf{I}_K and deformations $\boldsymbol{\varphi}_K$ in the definition of \mathcal{J}_K , see (3.29), we apply the temporal extension construction from Section 3.4. In particular, $\mathcal{J}_K(\mathbf{I}_K, \boldsymbol{\varphi}_K) \leq \overline{\mathcal{J}}$ for all $K \in \mathbb{N}$. Using (3.22) and (3.31), we conclude that z_K is uniformly bounded in $L^2((0, 1) \times \Omega)$. Next, Remark 3.7 together with (3.23) and (3.30) imply uniform boundedness of Y_K, X_K in $C^0([0, 1], C^{1,\alpha}(\overline{\Omega}))$. Incorporating (3.17), we obtain for $f_a(x) = a$ with $a \in \mathcal{H}$ that

$$d_2(I_K(t, \cdot), f_a) \leq C(d_2(I_K(t, Y_K(t, \cdot)), I_A) + d_2(I_A, f_a)) \leq C(\|z_K\|_{L^2((0,1) \times \Omega)} + 1).$$

Therefore, $\{I_K\}_{K \in \mathbb{N}}$ is uniformly bounded in $L^\infty([0, 1], L^2(\Omega, \mathcal{H}))$ and a subsequence converges weakly to some $I \in L^2([0, 1], L^2(\Omega, \mathcal{H}))$ in $L^2([0, 1], L^2(\Omega, \mathcal{H}))$.

Now, we follow the usual argument and assume that there exists an image path $\tilde{I} \in L^2([0, 1], L^2(\Omega, \mathcal{H}))$ with corresponding optimal tuple $(\tilde{I}, \tilde{v}, \tilde{Y}, z)$, which exists due to Proposition 3.13, satisfying (3.16) and (3.17) such that

$$\mathcal{J}[\tilde{I}] < \mathcal{J}[I]. \quad (3.39)$$

By Theorem 3.14, we get existence of a sequence $\{\tilde{I}_K\}_{K \in \mathbb{N}} \subset L^2((0, 1), L^2(\Omega, \mathcal{H}))$ satisfying $\limsup_{K \rightarrow \infty} \mathcal{J}_K[\tilde{I}_K] \leq \mathcal{J}[\tilde{I}]$. Thus, we obtain applying Theorem 3.12 that

$$\mathcal{J}[I] \leq \liminf_{K \rightarrow \infty} \mathcal{J}_K[I_K] \leq \limsup_{K \rightarrow \infty} \mathcal{J}_K[\tilde{I}_K] \leq \mathcal{J}[\tilde{I}], \quad (3.40)$$

which contradicts (3.39). Hence, I minimizes the continuous path energy over all admissible image paths. Finally, the discrete path energies converge to the limiting path energy along a subsequence, i.e., $\lim_{K \rightarrow \infty} \mathcal{J}_K[I_K] = \mathcal{J}[I]$, which again follows from (3.40) by using $\tilde{I} = I$. \square

3.6 Conclusion

In this chapter, we have introduced a novel metamorphosis functional for manifold-valued images. We specifically considered the case of images as maps into Hadamard manifolds. This choice is at first motivated by applications like DT-MRI images, which are also depicted as examples here. On the other hand, Hadamard manifolds come with the joint convexity of the distance functional. An important aspect of the generalized metamorphosis model for manifold-valued images is the inequality (3.17), which replaces the defining equation for the material derivative in the standard metamorphosis model. As it is shown here, it is in particular the joint convexity of the distance function that allows us to establish this inequality as the limiting constraint for our discrete approximation. Thus, Hadamard manifolds naturally arise in applications and appear to be the proper setup for which the existence and convergence analysis is still possible. Indeed, we picked up a natural time discretization for this model and proved the Mosco convergence to this novel time continuous metamorphosis model. This in particular establishes the existence of solutions for this model, not following or using the approach by Trouné and Younes in [263]. Also numerically, the joint convexity of the distance on Hadamard manifolds is of importance for the convergence of the alternating descent scheme presented in [213].

Regularization of Inverse Problems via Time Discrete Geodesics in Image Spaces

Contents

4.1	Introduction	77
4.2	Preliminaries	79
4.2.1	Admissible deformations	79
4.2.2	Space of images	80
4.3	Space continuous model	82
4.3.1	Model	82
4.3.2	Existence, stability and convergence	83
4.4	Minimization approaches	90
4.4.1	Spatial discretization	90
4.4.2	Proximal alternating linearized minimization (PALM) .	92
4.4.3	Alternating minimization approach	93
4.4.4	Multilevel strategy	95
4.5	Numerical examples	95
4.5.1	Limited angle and sparse CT	97
4.5.2	Superresolution	97
4.6	Conclusions	101

Abstract

This chapter¹, which is published in [214], addresses the solution of inverse problems in imaging given an additional reference image. We combine a modification of the discrete geodesic path model for image metamorphosis with a variational model, actually the L^2 - TV model, for image reconstruction. We prove that the space continuous model has a minimizer that depends in a stable way from the input data. Two minimization procedures, which alternate over the involved sequences of deformations and images in different ways, are proposed. The updates with respect to the image sequence exploit recent algorithms from convex analysis to minimize the L^2 - TV functional. For the numerical computation, we apply a finite difference approach on staggered grids together with a multilevel strategy. We present proof-of-concept numerical results for sparse and limited angle computed tomography as well as for superresolution demonstrating the power of the method.

¹This is an author-created, un-copyedited version of an article accepted for publication/published in Inverse Problems. IOP Publishing Ltd is not responsible for any errors or omissions in this version of the manuscript or any version derived from it. The Version of Record is available online at DOI: 10.1088/1361-6420/ab038a.

4.1 Introduction

In certain applications it makes sense to account for qualitative prior image information to improve the image reconstruction. Typical examples are image super-resolution and computed tomography (CT) with sparsely or limited angle sampled sinogram data. Earlier approaches to incorporate prior knowledge about the image into CT include phase field methods [159, 217], the application of level set techniques, in particular when combining registration with segmentation [258], as well as the utilization of local (shape) descriptors [221, 283]. Recently, a mathematical classification of artifacts from arbitrary incomplete X -ray tomography data using classical filtered backprojection was given in [39]. For earlier papers on this topic the reader may also consult [117, 118, 172, 218].

In this chapter, we incorporate a whole reference image into the reconstruction process and take its deformation towards the image of interest, which is only indirectly given by measurements, into account. Recent work in this direction shows promising results. Schumacher, Modersitzki and Fischer [247] have dealt with combined reconstruction and motion correction in SPECT imaging. Karlsson and Ringh [171] coupled the optimal transport model with inverse problems. Chen and Öktem [64] tackled hard inverse problems with shape priors under the name *indirect image registration* within the large deformation diffeomorphic metric mapping (LDDMM) framework and in an earlier work [223] via linearized deformations. The authors use ODE constrained problem formulations, where the regularization of the deformations exploits reproducing kernel Hilbert spaces. As a drawback, the LDDMM [25, 71, 95, 261, 262] based methods can only deal with images having the same intensities. The metamorphosis model of Miller, Trounev and Younes [202, 263, 264] is an extension of the LDDMM approach allowing variation of the image intensities along trajectories of the pixels. A comprehensive overview over the topic is given in the book [285] as well as in the review article [201]. For a historic account see also [199]. Recently, Gris, Chen and Öktem [139] have enlarged the ideas in [64, 223] to the metamorphosis setting.

Here, we also follow the metamorphosis idea, but in a completely different way than in [139]. We built up on the time discrete geodesic calculus proposed for shape spaces by Rumpf and Wirth [240, 241] and for images by Berkels, Effland and Rumpf [33]. For convergence of the time discrete path model to the metamorphosis one, we refer to these papers. Here, deformations are modeled via a smoothness term and the linearized elastic potential, which is also a usual choice in registration problems. We combine this model with a “usual” variational image reconstruction model, actually the L^2 -TV model, which originated from [238]. Inspired by compressive sensing [55, 93], such variational image reconstruction techniques with sparsity-exploiting priors have achieved impressive reductions in sampling requirements. Besides TV priors, wavelet, shearlet- and curvelet representations [75, 116] were exploited in CT reconstructions with incomplete data.

Let \mathcal{X}, \mathcal{Y} be Hilbert spaces and $A \in L(\mathcal{X}, \mathcal{Y})$ a linear, continuous operator. A typical choice for \mathcal{X} is the space of square integrable function $L^2(\Omega)$ defined over some image domain $\Omega \subset \mathbb{R}^2$. We want to reconstruct an unknown image $I_{\text{orig}} \in \mathcal{X}$




Figure 4.1: Illustration of the image and diffeomorphism path, where $I_{k+1}(x) \approx I_k(\varphi_k^{-1}(x))$, $k = 0, \dots, K - 1$.

having the following information available:

- I1) an image $B = AI_{\text{orig}} + \eta \in \mathcal{Y}$, where η denotes some small error, e.g., due to noise;
- I2) a reference image R that is similar to the original image.

A usual variational model to approximate I_{orig} from B using only I1) is given by

$$\arg \min_{I \in \mathcal{X}} \mathcal{E}(I; B) := \mathcal{D}(I; B) + \alpha \mathcal{P}(I), \quad \alpha \geq 0, \quad (4.1)$$

where \mathcal{D} is a data term and \mathcal{P} a prior or regularizer. A model for edge-preserving image reconstruction is the L^2 -TV model, which will be our model of choice.

To incorporate the reference image R , we want to combine model (4.1) with a modified version of the time discrete geodesic model for image metamorphosis [33]. Given a template image $I_0 = T$ and a reference image $I_K = R$, this model aims to find a chain of smooth deformations $(\varphi_0, \dots, \varphi_{K-1})$ from an appropriately defined admissible set \mathcal{A} together with a sequence of images (I_1, \dots, I_{K-1}) such that the sum of the quadratic distances

$$\sum_{k=0}^{K-1} \|I_k \circ \varphi_k^{-1} - I_{k+1}\|_{\mathcal{X}}^2 \quad (4.2)$$

together with a prior

$$\sum_{k=0}^{K-1} \int_{\Omega} \nu |D^m \varphi_k|^2 + W(D\varphi_k) dx, \quad \nu > 0,$$

on the deformations becomes small, see Fig. 4.1. The first part of the deformation regularization enforces the smoothness of the mappings, while the second term, circumscribed by W , will be chosen as linearized elastic potential. By (4.2), the image sequence (T, I_1, \dots, R) possibly differs from the deformed image sequence $(T, T \circ \varphi_0^{-1}, \dots, T \circ \varphi_0^{-1} \circ \dots \circ \varphi_{K-1}^{-1})$, which makes the model flexible for intensity changes.

For the numerical solution of our model, we propose two different procedures, namely proximal alternating linearized minimization (PALM) [36] and an alternating minimization approach related to [33, 213]. For the later one, recent primal-dual minimization algorithms from convex analysis are merged with a Quasi-Newton approach from image registration.

Outline of the chapter: In Section 4.2, the necessary preliminaries concerning the spaces of deformations and images are introduced. In particular, we highlight properties of the concatenations of admissible deformations and square integrable images. This motivates the modification of the time discrete path model [33] and also of our generalized model for manifold-valued images in [213]. In Section 4.3, our space continuous reconstruction model is established. Since it combines **time discrete** morphing with **inverse** problems we call it TDM-INV. We prove that the functional has a minimizer and that minimizers depend stably on the input data. Further, a convergence result for decreasing noise is provided. Section 4.4 deals with two minimization procedures. For minimizing the image sequence, we incorporate primal-dual algorithms from convex analysis. Further, we explain computational issues in the space discrete setting. The numerical examples in Section 4.5 demonstrate the performance of our approach. We finish with conclusions in Section 4.6.

4.2 Preliminaries

In the rest of this chapter, let $\Omega \subset \mathbb{R}^n$ be a nonempty, open, connected, and bounded set with Lipschitz boundary. In this section, we introduce admissible sets \mathcal{A} of deformations and consider the concatenation of deformations $\varphi \in \mathcal{A}$ with images $I \in L^2(\Omega)$. Note that $I \circ \varphi$ considered in [33] is in general not in $L^2(\Omega)$, while we will see that $I \circ \varphi^{-1} \in L^2(\Omega)$. Therefore, we prefer to modify the time discrete geodesic path model by using the later concatenation. Moreover, this fits better to the original metamorphosis setting of Trounev and Younes. In [97], the image space $L^\infty(\Omega)$ is proposed instead and in [213] the computations are considerably simplified by using a set \mathcal{A}_ε with deformations fulfilling $\det D\varphi \geq \varepsilon$ for some fixed $\varepsilon > 0$.

4.2.1 Admissible deformations

First, we introduce the smoothness spaces of our deformation mappings. Let $C^{k,\alpha}(\overline{\Omega})$, $k \in \mathbb{N}_0$, denote the Hölder space of functions $f \in C^k(\overline{\Omega})$ for which

$$\|f\|_{C^{k,\alpha}(\overline{\Omega})} := \sum_{|\beta| \leq k} \|D^\beta f\|_{C(\overline{\Omega})} + \sum_{|\beta|=k} \sup_{\substack{x,y \in \Omega \\ x \neq y}} \frac{|D^\beta f(x) - D^\beta f(y)|}{|x - y|^\alpha}$$

is finite. Equipped with this norm $C^{k,\alpha}(\overline{\Omega})$ is a Banach space.

By $W^{m,p}(\Omega)$, $m \in \mathbb{N}$, $1 \leq p < \infty$, we denote the Sobolev space of functions having weak derivatives up to order m in $L^p(\Omega)$ with norm

$$\|f\|_{W^{m,p}(\Omega)}^p := \sum_{|\alpha| \leq m} \int_{\Omega} |D^\alpha f|^p dx$$

and seminorm $|D^m f|^p := \sum_{|\alpha|=m} |D^\alpha f|^p$. For vector valued $F = (f_\nu)_{\nu=1}^n$, the component wise norm $|D^m F|^p := \sum_{\nu=1}^n |D^m f_\nu|^p$ is used. The space $W^{m,2}(\Omega)$ with

$m > 1 + \frac{n}{2}$ is of particular interest as it is compactly embedded in $C^{1,\alpha}(\overline{\Omega})$ for all $\alpha \in (0, m - 1 - \frac{n}{2})$, see [4, Thm. 8.13], and consequently also $W^{m,2}(\Omega) \hookrightarrow W^{1,p}(\Omega)$ for all $p \geq 1$.

It is assumed that the deformations φ are elements of the following *admissible set*

$$\mathcal{A} := \{\varphi \in W^{m,2}(\Omega)^n : \det D\varphi > 0 \text{ a.e. in } \Omega, \varphi(x) = x \text{ for } x \in \partial\Omega\},$$

where $m > 1 + \frac{n}{2}$. Then, by a result of Ball [19], φ has the following useful properties

i) $\varphi(\overline{\Omega}) = \overline{\Omega}$.

ii) φ maps measurable sets in $\overline{\Omega}$ to measurable sets in $\overline{\Omega}$ and the change of variables formula

$$\int_B I \circ \varphi \det D\varphi \, dx = \int_{\varphi(B)} I \, dy$$

holds for any measurable set $B \subset \overline{\Omega}$ and any measurable function $I: \overline{\Omega} \rightarrow \mathbb{R}$ provided that one of the above integrals exists.

iii) φ is injective a.e., i.e., the set

$$S := \{x \in \overline{\Omega} : \varphi^{-1}(x) \text{ has more than one element}\}$$

has Lebesgue measure zero.

By property i) and since $\overline{\Omega}$ is bounded, it follows immediately for all $\varphi \in \mathcal{A}$ that

$$\|\varphi\|_{L^\infty(\Omega)^n} \leq C, \quad \|\varphi\|_{L^2(\Omega)^n} \leq C, \quad (4.3)$$

with constants depending only on Ω . By the embedding properties of Sobolev spaces, it holds $\varphi \in C^{1,\alpha}(\overline{\Omega})^n$. Further, by the inverse mapping theorem, φ^{-1} exists locally around a.e. $x \in \Omega$ and is continuously differentiable on the corresponding neighborhood. However, to guarantee that φ^{-1} is continuous (or, even more, continuously differentiable) on Ω , further assumptions are required, see [19, Thm. 2]. A possible counterexample is the function $\varphi(x) := x^3$ on $\Omega := (-1, 1)$, which is in \mathcal{A} but $\varphi^{-1} = \text{sgn}(x)|x|^{\frac{1}{3}}$ is not continuously differentiable.

4.2.2 Space of images

In this chapter, we consider images as functions in $\mathcal{X} = L^2(\Omega)$. Unfortunately, the concatenation of $I \in L^2(\Omega)$ with $\varphi \in \mathcal{A}$ can result in a function

$$I \circ \varphi \notin L^2(\Omega),$$

as the example $I(x) := x^{-\frac{1}{4}}$ in $L^2((0, 1))$ and $\varphi(x) := x^2$ shows. However, this can be avoided by using

$$\varphi \cdot I := I \circ \varphi^{-1} \in L^2(\Omega), \quad (4.4)$$

where the function needs to be defined properly. To this end, let \mathcal{N} be a Borel null set containing S from iii). Then $\mathcal{B} := \Omega \setminus \mathcal{N}$ is a Borel set with $\mu(\mathcal{B}) = \mu(\Omega)$. Note that $\varphi^{-1}(\mathcal{B})$ is itself a Borel set since $\varphi \in W^{m,2}(\Omega)^n$ is measurable. Consider $\varphi^{-1}: \mathcal{B} \rightarrow \varphi^{-1}(\mathcal{B})$ and let $B \subseteq \varphi^{-1}(\mathcal{B})$ be a Borel set. Then, by ii), we see that $(\varphi^{-1})^{-1}(B) = \varphi(B)$ is a Borel set, so that φ^{-1} is a measurable function on \mathcal{B} . For $I \in L^2(\Omega)$ and φ^{-1} as above, the concatenation $I \circ \varphi^{-1}: \mathcal{B} \rightarrow \mathbb{R}$ is measurable if defined as follows

$$I \circ \varphi^{-1}(x) := \begin{cases} I \circ \varphi^{-1}(x) & x \in \mathcal{B}, \\ 0 & \text{otherwise.} \end{cases}$$

Then, (4.4) can be verified by

$$\int_{\Omega} |I \circ \varphi^{-1}|^2 dx = \int_{\mathcal{B}} |I \circ \varphi^{-1}|^2 dx = \int_{\varphi^{-1}(\mathcal{B})} |I|^2 \det D\varphi dy,$$

which is finite since $D\varphi$ has components in $C^{0,\alpha}(\overline{\Omega})$. The same argument can be used to show that $I \circ \varphi^{-1} \in L^p(\Omega)$, $p \in [1, \infty)$, if $I \in L^p(\Omega)$. Further, the following lemma on the image of null sets under deformations φ and φ^{-1} is useful.

Lemma 4.1. *For $\varphi \in \mathcal{A}$, both φ and its pre-image deformation φ^{-1} map null sets to null sets.*

Proof. As φ is Lipschitz continuous, it maps null sets to null sets [279, Thm. 3.33 and it's proof]. Now, assume that there exists a Borel null set \mathcal{N} with $\mu(\varphi^{-1}(\mathcal{N})) > 0$. Using the characteristic function $1_{\mathcal{N}}$ on \mathcal{N} , we get the contradiction

$$0 = \int_{\mathcal{N}} 1_{\mathcal{N}} dx = \int_{\varphi^{-1}(\mathcal{N})} 1_{\mathcal{N}} \circ \varphi \det D\varphi dy = \int_{\varphi^{-1}(\mathcal{N})} \det D\varphi dy > 0.$$

□

Finally, we prove a continuity result for the $L^p(\Omega)$ norm with respect to mappings $\varphi \in \mathcal{A}$.

Lemma 4.2. *Let $I \in L^p(\Omega)$, $p \in [1, \infty)$, and $\{\varphi_j\}_{j \in \mathbb{N}}$ be a sequence of deformations $\varphi_j \in \mathcal{A}$ with $\lim_{j \rightarrow \infty} \|\varphi_j - \hat{\varphi}\|_{C^{1,\alpha}(\Omega)^n} = 0$ for some $\hat{\varphi} \in \mathcal{A}$. Then it holds*

$$\lim_{j \rightarrow \infty} \|I \circ \varphi_j^{-1} - I \circ \hat{\varphi}^{-1}\|_{L^p(\Omega)} = 0.$$

Proof. Since $I \circ \hat{\varphi}^{-1} \in L^p(\Omega)$, there exists a sequence $\{I_k\}_{k \in \mathbb{N}}$ of uniformly continuous functions with $\|I \circ \hat{\varphi}^{-1} - I_k\|_{L^p(\Omega)} \leq \frac{1}{k}$. Using the fact that φ^{-1} maps null sets on null sets, we conclude

$$\begin{aligned} & \|I \circ \varphi_j^{-1} - I \circ \hat{\varphi}^{-1}\|_{L^p(\Omega)} \\ &= \|I \circ \varphi_j^{-1} - I_k \circ \hat{\varphi} \circ \varphi_j^{-1} + I_k \circ \hat{\varphi} \circ \varphi_j^{-1} - I_k + I_k - I \circ \hat{\varphi}^{-1}\|_{L^p(\Omega)} \\ &\leq \left(\int_{\Omega} |I - I_k \circ \hat{\varphi}|^p \det D\varphi_j dx \right)^{\frac{1}{p}} + \left(\int_{\Omega} |I_k \circ \hat{\varphi} - I_k \circ \varphi_j|^p \det D\varphi_j dx \right)^{\frac{1}{p}} + \frac{1}{k}. \end{aligned}$$

Due to the convergence of φ_j , there exists a constant C such that $\det D\varphi_j \leq C$ for all $j \in \mathbb{N}$. Thus,

$$\begin{aligned} & \|I \circ \varphi_j^{-1} - I \circ \hat{\varphi}^{-1}\|_{L^p(\Omega)} \\ & \leq \left(\int_{\Omega} |I - I_k \circ \hat{\varphi}|^p \det D\varphi_j \, dx \right)^{\frac{1}{p}} + C \left(\int_{\Omega} |I_k \circ \hat{\varphi} - I_k \circ \varphi_j|^p \, dx \right)^{\frac{1}{p}} + \frac{1}{k}. \end{aligned}$$

The last term converges to zero as $k \rightarrow \infty$. Now, fix $k \in \mathbb{N}$. Since φ_j converges uniformly to $\hat{\varphi}$, the uniform continuity of I_k can be used to conclude that $I_k \circ \varphi_j$ converges uniformly to $I_k \circ \hat{\varphi}$. Then, boundedness of Ω implies that the second term converges to zero as $j \rightarrow \infty$. For the first term, the uniform continuity of I_k implies that for every $\epsilon > 0$ there exists $j \in \mathbb{N}$ large enough such that

$$\begin{aligned} \left(\int_{\Omega} |I - I_k \circ \hat{\varphi}|^p \det D\varphi_j \, dx \right)^{\frac{1}{p}} & \leq \left(\int_{\Omega} |I - I_k \circ \hat{\varphi}|^p \det D\hat{\varphi} \, dx \right)^{\frac{1}{p}} + \epsilon \\ & = \left(\int_{\Omega} |I \circ \hat{\varphi}^{-1} - I_k|^p \, dx \right)^{\frac{1}{p}} + \epsilon \leq \frac{1}{k} + \epsilon. \end{aligned}$$

This concludes the proof. \square

4.3 Space continuous model

In this section, we establish our space continuous model, which takes the information I1) and I2) into account and prove existence of minimizers, stability and convergence for vanishing noise. These three properties are necessary for a well-defined regularization method.

4.3.1 Model

Starting with the information I1), we are interested in the reconstruction of a two-dimensional image from its measurements based on the variational approach (4.1). In this chapter, the main focus lies on the total variation seminorm as regularizer \mathcal{P} . More precisely, recall that the space of functions of bounded variation $BV(\Omega)$ consists of those functions $I \in L^1_{\text{loc}}(\Omega)$ having weak first order derivatives that are finite Radon measures. For $I \in L^1(\Omega)$, it holds that $I \in BV(\Omega)$ if and only if

$$\text{TV}(I) := \sup \left\{ \int_{\Omega} I \operatorname{div}(\eta) \, dx : \eta \in C_0^\infty(\Omega)^n, |\eta| \leq 1 \right\} < +\infty.$$

Equipped with the norm $\|I\|_{BV} := \|I\|_{L^1(\Omega)} + \text{TV}(I)$, the space $BV(\Omega)$ becomes a Banach space. For $\Omega \subset \mathbb{R}^2$, i.e., two-dimensional domains, $BV(\Omega)$ can be continuously embedded into $L^2(\Omega)$, see [5, Thm. 3.47]. Therefore, we can define

$$\mathcal{P}(I) := \begin{cases} \text{TV}(I) & \text{for } I \in BV(\Omega), \\ +\infty & \text{for } I \in L^2(\Omega) \setminus BV(\Omega). \end{cases} \quad (4.5)$$

It is well-known that \mathcal{P} in (4.5) is a proper, convex and lower semi-continuous (lsc) functional on $L^2(\Omega)$, see [244, Prop. 10.8].

Let $A: L^2(\Omega) \rightarrow \mathcal{Y}$ be a continuous linear operator into a Hilbert space \mathcal{Y} that does not vanish on constant functions and $B \in \mathcal{Y}$. In case of the Radon transform, it holds $\mathcal{Y} = L^2(\mathbb{S}^1, (-1, 1))$. Then, we define the variational reconstruction model

$$\mathcal{E}(I; B) := \frac{1}{2} \|AI - B\|_{\mathcal{Y}}^2 + \alpha \text{TV}(I), \quad \alpha > 0.$$

Note that $\mathcal{E}(I; B)$ is jointly weakly lsc in I and B .

Having a reference image $R \in L^2(\Omega)$ available, we want to add information I2) to the model. To this end, let $W: \mathbb{R}^{2,2} \rightarrow \mathbb{R}_{\geq 0}$ be a lsc mapping and $\nu > 0$, $m > 2$. Throughout the chapter, it is assumed that $K \geq 1$ is an integer. For a sequence $\mathbf{I} := (I_0, \dots, I_{K-1})$ of images in $L^2(\Omega)$ and a sequence of admissible deformations $\varphi := (\varphi_0, \dots, \varphi_{K-1})$, we consider the time discrete geodesic path model

$$\mathcal{F}(\mathbf{I}, \varphi) := \sum_{k=0}^{K-1} \int_{\Omega} W(D\varphi_k) + \nu |D^m \varphi_k|^2 + |I_k \circ \varphi_k^{-1} - I_{k+1}|^2 dx, \quad (4.6)$$

where $I_K := R \in L^2(\Omega)$ is a given reference image. Then, our whole model reads

$$\mathcal{J}(\mathbf{I}, \varphi) := \mathcal{E}(I_0; B) + \beta \mathcal{F}(\mathbf{I}, \varphi) \quad \text{subject to} \quad I_K = R, \quad (4.7)$$

where $\beta > 0$. We call this model TDM-INV model referring to 'time discrete morphing - inverse' problems.

Remark 4.3. *The linearized elastic potential is our choice for W in (4.6). More precisely, rewriting the deformation as $\varphi(x) = x + v(x)$ and introducing the notation of the (Cauchy) strain tensor of the displacement vector field $v = (v_1, v_2)^T: \Omega \mapsto \mathbb{R}^2$ as*

$$Dv_{\text{sym}} := \begin{pmatrix} \partial_x v_1 & \frac{1}{2}(\partial_y v_1 + \partial_x v_2) \\ \frac{1}{2}(\partial_y v_1 + \partial_x v_2) & \partial_y v_2 \end{pmatrix},$$

we apply

$$\mathcal{S}(v) := \int_{\Omega} \mu \text{trace} (Dv_{\text{sym}}^T Dv_{\text{sym}}) + \frac{\lambda}{2} \text{trace} (Dv_{\text{sym}})^2 dx, \quad \nu > 0.$$

Note that the linearized elastic potential is a usual regularizer in the context of registration, see [143, 204, 226].

4.3.2 Existence, stability and convergence

In this section, we prove that there exists a minimizer of \mathcal{J} in (4.7). Based on this, we show its stability with respect to the input data B and the convergence of an image sequence $\{I_0^{(j)}\}_{j \in \mathbb{N}}$ obtained from minimizing the functionals with input data B_j fulfilling $\|AI_0 - B_j\|_{\mathcal{Y}}^2 \leq \delta_j$ for a zero sequence $\{\delta_j\}_{j \in \mathbb{N}}$ and corresponding parameters α_j, β_j decaying faster than δ_j to the limit I_0 .

The existence proof is the hardest part. As usual for functionals in two variables, it is based on three pillars: First, it is shown that a minimizer exists if one of the variables is fixed. In a second step the results are merged to get the overall existence.

Fixing the image sequence \mathbf{I} leads to single registration problems. The proof of Lemma 4.4 follows similar ideas as in [33, 213]. However, since the setting in those papers is different, we prefer to carefully follow the lines and make the necessary modifications to make the chapter self-contained. Fixing φ , it is necessary to deal with the additional term \mathcal{E} and the proof of Lemma 4.5 is different from those in [33, 213], in particular it relies on nested weighted L^2 spaces. Except for the first step, the existence proof of Theorem 4.7 requires completely new estimates compared to [33, 213].

To begin with, we fix an image sequence $\mathbf{I} \in L^2(\Omega)^K$ and show that $J(\mathbf{I}, \cdot)$ has a minimizer $\varphi \in \mathcal{A}^K$. Then, the consideration can be restricted to $\mathcal{F}(\mathbf{I}, \cdot)$ and it suffices to prove that each of the summands

$$\mathcal{R}(\varphi_k; I_k, I_{k+1}) := \int_{\Omega} W(D\varphi_k) + \nu |D^m \varphi_k|^2 \, dx, + |I_k \circ \varphi_k^{-1} - I_{k+1}|^2 \, dx,$$

$k = 0, \dots, K-1$, has a minimizer in \mathcal{A} .

Lemma 4.4. *Let $W: \mathbb{R}^{n,n} \rightarrow \mathbb{R}_{\geq 0}$ be a lsc mapping with the property*

$$W(M) = \infty \quad \text{if} \quad \det M \leq 0. \quad (4.8)$$

Further, let $T, R \in L^2(\Omega)$ be given. Then there exists a minimizer $\hat{\varphi} \in \mathcal{A}$ of

$$\mathcal{R}(\varphi; T, R) := \int_{\Omega} W(D\varphi) + \nu |D^m \varphi|^2 + |T \circ \varphi^{-1} - R|^2 \, dx$$

over all $\varphi \in \mathcal{A}$.

Proof. 1. Let $\{\varphi_j\}_{j \in \mathbb{N}}$, $\varphi_j \in \mathcal{A}$, be a minimizing sequence of \mathcal{R} . Then it holds that $\mathcal{R}(\varphi_j; T, R) \leq C$ for all $j \in \mathbb{N}$. This implies that $\{\varphi_j\}_{j \in \mathbb{N}}$ has uniformly bounded $W^{m,2}(\Omega)^n$ -seminorm, and by (4.3) the sequence is also uniformly bounded in $L^2(\Omega)^n$. Now, we apply the Gagliardo–Nirenberg inequality, see Remark 4.13, which states that for all $0 \leq i < m$ it holds

$$\|D^i \varphi_j\|_{L^2(\Omega)} \leq C_1 \|D^m \varphi_j\|_{L^2(\Omega)} + C_2 \|\varphi_j\|_{L^2(\Omega)}.$$

All terms on the right-hand side are uniformly bounded. Hence, the $W^{m,2}(\Omega)^n$ norm of $\{\varphi_j\}_{j \in \mathbb{N}}$ is uniformly bounded. Since $W^{m,2}(\Omega)$ is reflexive, there exists a subsequence that converges weakly to some function $\hat{\varphi}$ in $W^{m,2}(\Omega)^n$. By the compact embedding $W^{m,2}(\Omega) \hookrightarrow C^{1,\alpha}(\overline{\Omega})$, $\alpha \in (0, m - 1 - \frac{n}{2})$, this subsequence, which is again denoted by $\{\varphi_j\}_{j \in \mathbb{N}}$, converges strongly to $\hat{\varphi}$ in $C^{1,\alpha}(\overline{\Omega})^n$ and hence $D\varphi_j$ converges uniformly to $D\hat{\varphi}$.

2. Next, we show that $\hat{\varphi}$ is in the set \mathcal{A} . Since W is lsc, we conclude

$$\liminf_{j \rightarrow \infty} W(D\varphi_j)(x) \geq W(D\hat{\varphi})(x)$$

for all $x \in \Omega$ and as W is non-negative Fatou's lemma implies

$$\int_{\Omega} W(D\hat{\varphi}) \, dx \leq \liminf_{j \rightarrow \infty} \int_{\Omega} W(D\varphi_j) \, dx \leq C.$$

By incorporating (4.8), this implies $\det D\hat{\varphi} > 0$ a.e. Further, the boundary condition is fulfilled so that $\hat{\varphi} \in \mathcal{A}$.

3. It remains to show that $\hat{\varphi}$ is a minimizer of $\mathcal{R}(\varphi; T, R)$. By Lemma 4.2, it holds $\|T \circ \varphi_j^{-1} - T \circ \hat{\varphi}^{-1}\|_{L^2(\Omega)} \rightarrow 0$ as $j \rightarrow \infty$ so that by the continuity of the norm

$$\|T \circ \hat{\varphi}^{-1} - R\|_{L^2(\Omega)} = \lim_{j \rightarrow \infty} \|T \circ \varphi_j^{-1} - R\|_{L^2(\Omega)}.$$

This together with the previous steps of the proof implies that the three summands in \mathcal{R} are (weakly) lsc. Hence, we obtain

$$\begin{aligned} \mathcal{R}(\hat{\varphi}; T, R) &\leq \liminf_{j \rightarrow \infty} \int_{\Omega} W(D\varphi_j) + \nu |D^m \varphi_j|^2 + |T \circ \varphi_j^{-1} - R|^2 \, dx \\ &= \inf_{\varphi \in \mathcal{A}} \mathcal{R}(\varphi; T, R), \end{aligned}$$

which proves the claim. \square

Next, we fix a sequence of mappings $\varphi \in \mathcal{A}^K$ and ask for a minimizer of $\mathcal{J}(\cdot, \varphi)$.

Lemma 4.5. *Let $A: L^2(\Omega) \rightarrow \mathcal{Y}$ be a continuous linear operator into a Hilbert space \mathcal{Y} that does not vanish on constant functions, $B \in \mathcal{Y}$ and $R \in L^2(\Omega)$. For fixed $\varphi \in \mathcal{A}^K$, there exists a unique image sequence $\mathbf{I} \in L^2(\Omega)^K$ minimizing $\mathcal{J}(\cdot, \varphi)$.*

Proof. We prove lower semi-continuity, coercivity and strict convexity of the functional. Neglecting the constant terms and by changing the indexing of the sum, it suffices to consider

$$J(\mathbf{I}) := \beta \sum_{k=1}^K \int_{\Omega} |I_{k-1} \circ \varphi_{k-1}^{-1} - I_k|^2 \, dx + \mathcal{E}(I_0; B) \quad \text{subject to} \quad I_K = R. \quad (4.9)$$

Setting

$$\begin{aligned} \psi_0 &:= \text{Id}, \\ \psi_k &:= \varphi_{k-1} \circ \psi_{k-1} = \varphi_{k-1} \circ \dots \circ \varphi_0, \quad k = 1, \dots, K, \end{aligned}$$

and substituting $x := \psi_k(y)$ in the k -th summand of (4.9), the functional transforms to

$$J(\mathbf{I}) = \beta \sum_{k=1}^K \int_{\Omega} |I_k \circ \psi_k - I_{k-1} \circ \psi_{k-1}|^2 \det D\psi_k \, dy + \mathcal{E}(I_0; B).$$

Using $\mathbf{F} := (F_0, \dots, F_{K-1})$, where $F_0 := I_0$, $F_k := I_k \circ \psi_k$, and $w_k := \det D\psi_k$, we are concerned with the minimization of

$$\begin{aligned} \tilde{J}(\mathbf{F}) &:= \beta \sum_{k=1}^K \int_{\Omega} |F_k - F_{k-1}|^2 w_k \, dx + \mathcal{E}(F_0; B) \\ &\text{subject to } F_K = R \circ \psi_K. \end{aligned} \quad (4.10)$$

Note that by $0 < w_k \leq C$ a.e. and $w_k = w_{k-1} \det(D\varphi_{k-1} \circ \psi_{k-1})$, the weighted L^2 spaces are nested

$$L^2(\Omega) = L_{w_0}^2(\Omega) \subseteq L_{w_1}^2(\Omega) \subseteq \dots \subseteq L_{w_K}^2(\Omega), \quad (4.11)$$

in particular $F_K \in L_{w_K}^2(\Omega)$ if $R \in L^2(\Omega)$. Further, any minimizer must satisfy $F_0 \in BV(\Omega) \subset L^2(\Omega)$, and by successively considering the integrals in (4.10) we get $F_k \in L_{w_k}^2(\Omega)$. In the following, we set $\beta := 1$ and $\|\cdot\|_{w_k} := \|\cdot\|_{L_{w_k}^2(\Omega)}$ to simplify the notation. Since $g: \mathbb{R}^K \rightarrow \mathbb{R}$ defined by

$$g(f_0, \dots, f_{K-1}) := \sum_{k=1}^K (f_k - f_{k-1})^2 w_k$$

with $w_k > 0$ and f_K fixed is strictly convex, the same holds for the sum of integrals in \tilde{J} . Clearly, this sum can be rewritten as $\sum_{k=1}^{K-1} \|F_k - F_{k-1}\|_{w_k}^2 + \|R \circ \psi_K - F_{K-1}\|_{w_K}^2$, which is continuous. As $\mathcal{E}(F_0; B)$ is proper, convex and lsc, the same holds true for \tilde{J} over $L_{w_0}^2(\Omega) \times \dots \times L_{w_{K-1}}^2(\Omega)$. Thus, \tilde{J} is also weakly lsc [244, Lem. 10.4].

Next, we show that \tilde{J} is coercive. Conversely, assume $\sum_{k=0}^{K-1} \|F_k^{(j)}\|_{w_k} \rightarrow \infty$ with $\tilde{J}(\mathbf{F}^{(j)})$ being bounded. Using the assumptions on A , it holds that $\mathcal{E}(F_0; B)$ is coercive, see [46, Thm. 6.115]. Consequently, $\|F_0^{(j)}\|_1$ is bounded and by (4.11) also $\|F_0^{(j)}\|_{w_1}$ is bounded. Considering successively the integrals in (4.10), this implies that $\|F_k^{(j)}\|_{w_k}$, $k = 1, \dots, K-1$, is bounded, which contradicts our assumption.

Thus, \tilde{J} is coercive and as it is weakly lsc and strictly convex, the functional has a unique minimizer \mathbf{F} . By definition of \mathbf{F} , the unique minimizer of J is given by \mathbf{I} with $I_k = F_k \circ \psi_k^{-1} \in L^2(\Omega)$. \square

For our computations, the following corollary on the minimizer of \tilde{J} in (4.10) with fixed F_0 is useful.

Corollary 4.6. *Let $K \geq 2$ be an integer. Further, let $w_k \in C^{0,\alpha}(\overline{\Omega})$, $k = 1, \dots, K$, fulfill $w_k > 0$ a.e. on Ω and $w_{k+1}/w_k \leq C$, $k = 1, \dots, K-1$. For given $F_0 \in L^2(\Omega)$ and $F_K \in L_{w_K}^2(\Omega)$, the solution of*

$$\arg \min_{F_k \in L_{w_k}^2(\Omega)} \sum_{k=1}^K \int_{\Omega} |F_k - F_{k-1}|^2 w_k \, dx$$

is given by

$$F_k = t_k F_0 + (1 - t_k) F_K, \quad t_k := \frac{\sum_{i=1}^k w_i^{-1}}{\sum_{i=1}^K w_i^{-1}}. \quad (4.12)$$

Proof. Setting the first derivative of the functional to zero, we obtain a.e. on Ω that

$$w_k(F_k - F_{k-1}) + w_{k+1}(F_k - F_{k+1}) = 0, \quad k = 1, \dots, K-1.$$

This can be rewritten as linear system of equations

$$\text{tridiag}(-w_k, w_k + w_{k+1}, -w_{k+1})_{k=1}^{K-1} (F_1, \dots, F_{K-1})^T = (w_1 F_0, 0, \dots, 0, w_K F_K)^T.$$

Since the tridiagonal matrix is irreducible diagonal dominant, the system has a unique solution. Straightforward computation shows that the solution is given by (4.12). \square

Now, we are able to prove the three main results of this section, beginning with existence of minimizers.

Theorem 4.7 (Existence). *Let $R \in L^2(\Omega)$ and $B \in \mathcal{Y}$. Then there exists a minimizer $(\hat{\mathbf{I}}, \hat{\varphi}) \in L^2(\Omega)^K \times \mathcal{A}^K$ of \mathcal{J} .*

Proof. The outline of the proof is as follows. First, we take a minimizing sequence of \mathcal{J} and show that the deformations and the intermediate images have a weakly convergent subsequence. Then, we prove that their concatenation is also weakly convergent and use this to get the weak lower semi-continuity of the functional.

1. Let $\{(\mathbf{I}^{(j)}, \phi^{(j)})\}_{j \in \mathbb{N}}$ be a minimizing sequence of \mathcal{J} . Then $\mathcal{J}(\mathbf{I}^{(j)}, \phi^{(j)}) \leq C$ for all $j \in \mathbb{N}$. By Lemma 4.4, we get for each $\mathbf{I}^{(j)}$ a sequence of diffeomorphisms $\varphi^{(j)}$ such that

$$\mathcal{J}(\mathbf{I}^{(j)}, \varphi^{(j)}) \leq \mathcal{J}(\mathbf{I}^{(j)}, \varphi)$$

for all $\varphi \in \mathcal{A}^K$. Hence, we conclude $\|D^m \varphi_k^{(j)}\|_{L^2(\Omega)}^2 < C/\nu$ for all $j \in \mathbb{N}$ and $k = 0, \dots, K-1$. As in the first part of the proof of Lemma 4.4, we conclude that $\{\varphi_k^{(j)}\}_{j \in \mathbb{N}}$ is bounded in $W^{m,2}(\Omega)^n$. Consequently, there exists a subsequence converging weakly in $W^{m,2}(\Omega)^n$ and strongly in $C^{1,\alpha}(\bar{\Omega})^n$ to $\hat{\varphi}_k$. Set $\hat{\varphi} := (\hat{\varphi}_k)_{k=0}^{K-1}$, denote the subsequence again by $\{\varphi_k^{(j)}\}_{j \in \mathbb{N}}$ and define $\varphi^{(j)} := (\varphi_k^{(j)})_{k=0}^{K-1}$.

2. Since $\mathcal{J}(\mathbf{I}^{(j)}, \varphi^{(j)}) \leq C$ for all $j \in \mathbb{N}$, coercivity of \mathcal{E} implies that $\|I_0^{(j)}\|_{L^2(\Omega)}$ is bounded. Additionally, we conclude for $k = 0, \dots, K-2$ that

$$\begin{aligned} \|I_{k+1}^{(j)}\|_{L^2(\Omega)} &\leq \|I_k^{(j)} \circ (\varphi_k^{(j)})^{-1} - I_{k+1}^{(j)}\|_{L^2(\Omega)} + \|I_k^{(j)} \circ (\varphi_k^{(j)})^{-1}\|_{L^2(\Omega)} \\ &\leq C^{\frac{1}{2}} + \|I_k^{(j)} \circ (\varphi_k^{(j)})^{-1}\|_{L^2(\Omega)}. \end{aligned}$$

Further, $\varphi_k^{(j)}$ is convergent in $C^{1,\alpha}(\bar{\Omega})^n$ and consequently $\det D\varphi_k^{(j)} \leq \tilde{C}$ on Ω for $k = 0, \dots, K-1$. Then, it holds

$$\begin{aligned} \|I_1^{(j)}\|_{L^2(\Omega)} &\leq C^{\frac{1}{2}} + \|I_0^{(j)} \circ (\varphi_0^{(j)})^{-1}\|_{L^2(\Omega)} \\ \|I_2^{(j)}\|_{L^2(\Omega)} &\leq C^{\frac{1}{2}} + \|I_1^{(j)} \circ (\varphi_1^{(j)})^{-1}\|_{L^2(\Omega)} \\ &= C^{\frac{1}{2}} + \left(\int_{\Omega} |I_1^{(j)}|^2 \det D\varphi_1^{(j)} dx \right)^{\frac{1}{2}} \\ &\leq C^{\frac{1}{2}} + \tilde{C}^{\frac{1}{2}} \|I_1^{(j)}\|_{L^2(\Omega)}. \end{aligned}$$

Successive continuation shows that the sequence $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$ is bounded in $L^2(\Omega)^K$. Hence, there exists a weakly convergent subsequence, also denoted by $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$, that converges to $\hat{\mathbf{I}} \in L^2(\Omega)^K$.

3. Next, we show weak convergence of $I_k^{(j)} \circ (\varphi_k^{(j)})^{-1}$ to $\hat{I}_k \circ \hat{\varphi}_k^{-1}$. As the sequence is bounded, it suffices to test with $g \in C_c^\infty(\Omega)$. It holds

$$\int_{\Omega} \left(I_k^{(j)} \circ (\varphi_k^{(j)})^{-1} - \hat{I}_k \circ \hat{\varphi}_k^{-1} \right) g \, dx = \mathcal{I}_1^{(j)} + \mathcal{I}_2^{(j)}$$

with

$$\begin{aligned} \mathcal{I}_1^{(j)} &:= \int_{\Omega} \left(I_k^{(j)} \circ (\varphi_k^{(j)})^{-1} - I_k^{(j)} \circ \hat{\varphi}_k^{-1} \right) g \, dx, \\ \mathcal{I}_2^{(j)} &:= \int_{\Omega} \left(I_k^{(j)} \circ \hat{\varphi}_k^{-1} - \hat{I}_k \circ \hat{\varphi}_k^{-1} \right) g \, dx. \end{aligned}$$

Using the change of variables formula, we obtain

$$\mathcal{I}_2^{(j)} = \int_{\Omega} (I_k^{(j)} - \hat{I}_k) g \circ \hat{\varphi}_k \det D\hat{\varphi}_k \, dx.$$

Since $g \circ \hat{\varphi}_k \det D\hat{\varphi}_k \in L^2(\Omega)$, the weak convergence of $I_k^{(j)}$ to \hat{I}_k implies that $\mathcal{I}_2^{(j)}$ converges to zero as $j \rightarrow \infty$. Using the change of variables formula again, $\mathcal{I}_1^{(j)}$ can be estimated by

$$\begin{aligned} \mathcal{I}_1^{(j)} &= \int_{\Omega} I_k^{(j)} \left(g \circ \varphi_k^{(j)} \det D\varphi_k^{(j)} - g \circ \hat{\varphi}_k \det D\hat{\varphi}_k \right) dx \\ &\leq \|I_k^{(j)}\|_{L^2(\Omega)} \|g \circ \varphi_k^{(j)} \det D\varphi_k^{(j)} - g \circ \hat{\varphi}_k \det D\hat{\varphi}_k\|_{L^2(\Omega)}. \end{aligned}$$

As $\{I_k^{(j)}\}_{j \in \mathbb{N}}$ is bounded, it suffices to show convergence of the second factor. With $g_k^{(j)} := g \circ \varphi_k^{(j)}$ and $\hat{g}_k := g \circ \hat{\varphi}_k$ it follows that

$$\begin{aligned} &\|g_k^{(j)} \det D\varphi_k^{(j)} - \hat{g}_k \det D\hat{\varphi}_k\|_{L^2(\Omega)} \\ &\leq \|g_k^{(j)} \det D\varphi_k^{(j)} - g_k^{(j)} \det D\hat{\varphi}_k\|_{L^2(\Omega)} + \|g_k^{(j)} \det D\hat{\varphi}_k - \hat{g}_k \det D\hat{\varphi}_k\|_{L^2(\Omega)} \\ &\leq C \|\det D\varphi_k^{(j)} - \det D\hat{\varphi}_k\|_{C^0(\bar{\Omega})} + C \|g_k^{(j)} - \hat{g}_k\|_{L^2(\Omega)}. \end{aligned}$$

The first term converges to zero since $D\varphi_k^{(j)}$ is convergent. Uniform convergence of $\varphi_k^{(j)}$ together with uniform continuity of g implies that $g_k^{(j)}$ converges uniformly to \hat{g}_k . Then, boundedness of Ω implies that the second term converges to zero.

4. It remains to show that $(\hat{\mathbf{I}}, \hat{\varphi})$ is a minimizer of $\mathcal{J}(\mathbf{I}, \phi) = \mathcal{E}(I_0; B) + \beta \mathcal{F}(\mathbf{I}, \varphi)$. It holds

$$\begin{aligned} \liminf_{j \rightarrow \infty} \mathcal{F}(\mathbf{I}^{(j)}, \varphi^{(j)}) &\geq \sum_{k=0}^{K-1} \liminf_{j \rightarrow \infty} \int_{\Omega} W(D\varphi_k^{(j)}) \, dx + \nu \liminf_{j \rightarrow \infty} \int_{\Omega} |D^m \varphi_k^{(j)}|^2 \, dx \\ &\quad + \liminf_{j \rightarrow \infty} \|I_k^{(j)} \circ (\varphi_k^{(j)})^{-1} - I_{k+1}^{(j)}\|_{L^2(\Omega)}^2. \end{aligned}$$

The components of $\varphi_k^{(j)}$ converge weakly in $W^{m,2}(\Omega)$, those of $D\varphi_k^{(j)}$ converge in $C^0(\overline{\Omega})$, and $I_k^{(j)}, I_k^{(j)} \circ (\varphi_k^{(j)})^{-1}$ converges weakly in $L^2(\Omega)$. We use this together with the facts that the first summand is lsc, the second one weakly lsc and $\|f - g\|_{L^2(\Omega)}^2$ is weakly lsc (convex and lsc) in both arguments to conclude

$$\liminf_{j \rightarrow \infty} \mathcal{F}(\mathbf{I}^{(j)}, \varphi^{(j)}) \geq \mathcal{F}(\hat{\mathbf{I}}, \hat{\varphi}).$$

As $\mathcal{E}(I_0; B)$ is weakly lsc in I_0 , we obtain

$$\inf_{\mathbf{I}, \phi} \mathcal{J}(\mathbf{I}, \phi) = \liminf_{j \rightarrow \infty} \mathcal{J}(\mathbf{I}^{(j)}, \phi^{(j)}) \geq \liminf_{j \rightarrow \infty} \mathcal{J}(\mathbf{I}^{(j)}, \varphi^{(j)}) \geq \mathcal{J}(\hat{\mathbf{I}}, \hat{\varphi}).$$

□

Next, we prove that the minimizers of \mathcal{J} depend stably on the input data B . To emphasize the dependence of \mathcal{J} on B , we use the notation \mathcal{J}_B instead of \mathcal{J} .

Theorem 4.8 (Stability). *Let $R \in L^2(\Omega)$. Further, let $\{B_j\}_{j \in \mathbb{N}}$ be a sequence in \mathcal{Y} converging to $B \in \mathcal{Y}$. For each $j \in \mathbb{N}$, we choose a minimizer $(\mathbf{I}^{(j)}, \varphi^{(j)})$ of \mathcal{J}_{B_j} . Then, there exists a subsequence of $\{(\mathbf{I}^{(j)}, \varphi^{(j)})\}_{j \in \mathbb{N}}$ that converges weakly to a minimizer $(\hat{\mathbf{I}}, \hat{\varphi})$ of \mathcal{J}_B .*

Proof. 1. Due to the convergence of $\{B_j\}_{j \in \mathbb{N}}$, it holds for every $I_0 \in L^2(\Omega)$ that

$$\mathcal{E}(I_0; B_j) = \frac{1}{2} \|AI_0 - B_j\|_{\mathcal{Y}}^2 + \alpha \text{TV}(I_0) \rightarrow \mathcal{E}(I_0, B).$$

Hence, there exists $C > 0$ with $\mathcal{J}_{B_j}(\mathbf{I}^{(j)}, \varphi^{(j)}) \leq \mathcal{J}_{B_j}(\mathbf{I}^{(1)}, \varphi^{(1)}) \leq C$ for all $j \in \mathbb{N}$. By definition of \mathcal{J}_{B_j} , we get for $j \in \mathbb{N}$ and $k = 0, \dots, K-1$ that $\|D^m \varphi_k^{(j)}\|_{L^2(\Omega)}^2 < C/\nu$. As in the first part of the proof of Lemma 4.4, we conclude that there exists a subsequence converging weakly in $W^{m,2}(\Omega)^n$ and strongly in $C^{1,\alpha}(\overline{\Omega})^n$ to $\hat{\varphi}_k$. Set $\hat{\varphi} := (\hat{\varphi}_k)_{k=0}^{K-1}$. Then, let us denote the subsequence again by $\{\varphi_k^{(j)}\}_{j \in \mathbb{N}}$ and further define $\varphi^{(j)} := (\varphi_k^{(j)})_{k=0}^{K-1}$.

2. Next, we estimate

$$C \geq \mathcal{E}(I_0^{(j)}; B_j) \geq \frac{1}{2} (\|AI_0^{(j)} - B\|_{\mathcal{Y}} - \|B - B_j\|_{\mathcal{Y}})^2 + \alpha \text{TV}(I_0^{(j)}),$$

so that the coercivity of $\frac{1}{2} \|AI_0 - B\|_{\mathcal{Y}}^2 + \alpha \text{TV}(I_0)$ in I_0 implies the boundedness of $\{I_0^{(j)}\}_{j \in \mathbb{N}}$. Now, we can reproduce Steps 2 and 3 from Theorem 4.7 to see that there exists a weakly convergent subsequence, also denoted by $\{\mathbf{I}^{(j)}\}_{j \in \mathbb{N}}$, that converges to $\hat{\mathbf{I}} \in L^2(\Omega)^K$. Additionally, the sequence $I_k^{(j)} \circ (\varphi_k^{(j)})^{-1}$ converges weakly to $\hat{I}_k \circ \hat{\varphi}_k^{-1}$.

3. It remains to show that $(\hat{\mathbf{I}}, \hat{\varphi})$ minimizes \mathcal{J}_B . We can use the lower semi-continuity argument for \mathcal{F} from Theorem 4.7 together with the fact that $\mathcal{E}(I; B)$ is jointly lsc to obtain for any (\mathbf{I}, φ) that

$$\mathcal{J}_B(\hat{\mathbf{I}}, \hat{\varphi}) \leq \liminf_{j \rightarrow \infty} \mathcal{J}_{B_j}(\mathbf{I}^{(j)}, \varphi^{(j)}) \leq \liminf_{j \rightarrow \infty} \mathcal{J}_{B_j}(\mathbf{I}, \varphi) = \mathcal{J}_B(\mathbf{I}, \varphi).$$

The last equality follows from the convergence of B_j together with continuity of \mathcal{E} in B . Hence, $(\hat{\mathbf{I}}, \hat{\varphi})$ is a minimizer of \mathcal{J}_B . □

This section concludes with a convergence result for vanishing noise. Here, we additionally need the dependence of \mathcal{J} on the parameter α (for simplicity we choose $\alpha = \beta$) and hence we write $\mathcal{J}_{\alpha,B}$.

Theorem 4.9 (Convergence). *Let $R \in L^2(\Omega)$ and $B \in \mathcal{Y}$, and suppose that there exists $(\tilde{\mathbf{I}}, \tilde{\varphi})$ such that $A\tilde{I}_0 = B$ and $\mathcal{J}_{1,B}(\tilde{\mathbf{I}}, \tilde{\varphi}) < \infty$. Further, assume that $\alpha: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ satisfies $\alpha(\delta) \rightarrow 0$ and $\delta/\alpha(\delta) \rightarrow 0$ as $\delta \rightarrow 0$. Assume that $\{\delta_j\}_{j \in \mathbb{N}}$ is a sequence of positive numbers converging to 0 and $\{B_j\}_{j \in \mathbb{N}}$ is a sequence in \mathcal{Y} satisfying $\|B - B_j\|_{\mathcal{Y}}^2 \leq \delta_j$ for each j . Let $(\mathbf{I}^{(j)}, \varphi^{(j)})$ be a minimizer of $\mathcal{J}_{\alpha_j, B_j}$, where $\alpha_j := \alpha(\delta_j)$. Then there exists a subsequence of $\{I_0^{(j)}\}_{j \in \mathbb{N}}$ that converges weakly to an image \hat{I}_0 satisfying $A\hat{I}_0 = B$.*

Proof. For every $j \in \mathbb{N}$, it holds

$$\begin{aligned} \|I_0^{(j)}\|_{L^2(\Omega)}^2 &\leq C \operatorname{TV}(I_0^{(j)}) \leq \frac{C}{\alpha_j} \mathcal{J}_{\alpha_j, B_j}(\mathbf{I}^{(j)}, \varphi^{(j)}) \leq \frac{C}{\alpha_j} \mathcal{J}_{\alpha_j, B_j}(\tilde{\mathbf{I}}, \tilde{\varphi}) \\ &= \frac{C}{2\alpha_j} \|B - B_j\|_{\mathcal{Y}}^2 + C \operatorname{TV}(\tilde{\mathbf{I}}) + C\mathcal{F}(\tilde{\mathbf{I}}, \tilde{\varphi}) \\ &\leq \frac{C\delta_j}{2\alpha_j} + C\mathcal{J}_{1,B}(\tilde{\mathbf{I}}, \tilde{\varphi}). \end{aligned}$$

From the assumptions on α and δ we deduce that $\|I_0^{(j)}\|_{L^2(\Omega)}^2$ is bounded. Thus, there is a weakly convergent subsequence with limit \hat{I}_0 . Additionally, it holds that $\|A\hat{I}_0 - B\|_{\mathcal{Y}}^2 \leq \liminf_{j \rightarrow \infty} \|AI_0^{(j)} - B_j\|_{\mathcal{Y}}^2$. Now, we can estimate

$$\|AI_0^{(j)} - B_j\|_{\mathcal{Y}}^2 \leq \mathcal{J}_{\alpha_j, B_j}(\mathbf{I}^{(j)}, \varphi^{(j)}) \leq \mathcal{J}_{\alpha_j, B_j}(\tilde{\mathbf{I}}, \tilde{\varphi}) = \alpha_j \mathcal{J}_{1,B}(\tilde{\mathbf{I}}, \tilde{\varphi}) + \frac{1}{2} \|B - B_j\|_{\mathcal{Y}}^2.$$

Since the two rightmost terms converge to zero, this implies $A\hat{I}_0 = B$. \square

4.4 Minimization approaches

In this section, we propose two different alternating minimization schemes. The first one is known as PALM and updates in each step the deformations and images via proximal computations. Convergence of the whole iteration sequence to a critical point is ensured. The second one just alternates the minimization of the deformations and the images.

Note that solving the *coupled* problem in (\mathbf{I}, φ) , e.g., with a gradient scheme is very time and memory consuming, since all I_k and φ_k are treated at the same time. Moreover, the coupling $I_k^{(j)} \circ \varphi_k^{-1}$ is non-convex and hence it is difficult to provide convergence results for general schemes.

We start with the spatial discretization of \mathcal{J} in (4.7).

4.4.1 Spatial discretization

Dealing with rectangular digital images, we propose a finite difference approach, where we work on staggered grids, see Fig. 4.2. In the following, the spatial discretization is briefly sketched. The domain of the images I is the (primal) grid




Figure 4.2: Illustration of the staggered grid, where empty boxes mean zero movement.

$\mathcal{G} := \{1, \dots, n_1\} \times \{1, \dots, n_2\}$. All integrals are approximated on the integration domain $\overline{\Omega} := [\frac{1}{2}, n_1 + \frac{1}{2}] \times [\frac{1}{2}, n_2 + \frac{1}{2}]$ by the midpoint quadrature rule, i.e., with pixel values defined on \mathcal{G} . Further, it is assumed that for the operator A a discrete version $A: \mathcal{G} \rightarrow Y$ is known, where Y is some finite-dimensional Hilbert space.

First, we discuss the discretization of \mathcal{F} . As regularizer $W(D\varphi)$ we propose the linearized elastic potential $S(v)$ from Remark 4.3 with the variable transformation $v = (v_1, v_2)^T = \varphi - \text{id}$. Using the $\frac{1}{2}$ -shifted grids

$$\mathcal{G}_1 := \{\frac{3}{2}, \dots, n_1 - \frac{1}{2}\} \times \{1, \dots, n_2\}, \quad \mathcal{G}_2 := \{1, \dots, n_1\} \times \{\frac{3}{2}, \dots, n_2 - \frac{1}{2}\},$$

we consider $v = (v_1, v_2)^T$ with $v_1: \mathcal{G}_1 \rightarrow \mathbb{R}$ and $v_2: \mathcal{G}_2 \rightarrow \mathbb{R}$. Then, the spatially discrete version of S reads

$$\begin{aligned} \mathcal{S}(v) = & \mu \left(\|D_{1,x_1} v_1\|_F^2 + \|v_2 D_{2,x_2}^T\|_F^2 + \frac{1}{2} \|v_1 D_{1,x_2}^T + D_{2,x_1} v_2\|_F^2 \right) \\ & + \frac{\lambda}{2} \|D_{1,x_1} v_1 + v_2 D_{2,x_2}^T\|_F^2, \end{aligned}$$

where $\|\cdot\|_F$ is the Frobenius norm of matrices and D_{i,x_j} denotes the forward difference operator (matrix) for v_i in x_j -direction. The higher order term $\int_{\Omega} |D^m \varphi|^2 dx$ with $m = 3$ is discretized by

$$D_3(v) := \sum_{i=0}^3 \left(\|D_{1,x_1}^i v_1 D_{1,x_2}^{3-i}\|_F^2 + \|D_{2,x_1}^i v_2 D_{2,x_2}^{3-i}\|_F^2 + \eta (\|v_1\|_F^2 + \|v_2\|_F^2) \right),$$

where central differences operators are used for the partial derivatives of order two and three. Note that we added the squared Frobenius norm of the v_i , $i = 1, 2$, for a

better control of the displacement value. To cope with the remaining deformation term in (4.6), we approximate $\varphi^{-1} \approx \text{id} - v$ such that the data term simplifies to

$$\int_{\Omega} |I_k^{(j)}(x - v(x)) - I_{k+1}^{(j)}(x)|^2 dx.$$

This integral is evaluated using the midpoint quadrature rule. Since v_i is only defined on \mathcal{G}_i , $i = 1, 2$, the averaged version $Pv = (P_1v_1, P_2v_2)^T: \mathcal{G} \rightarrow \mathbb{R}^2$ is used. In general $x - Pv(x) \notin \mathcal{G}$, so that the image $I_k(x - Pv(x))$ has to be interpolated from its values on \mathcal{G} . For this purpose linear interpolation with an interpolation matrix P_I is used. Note that also interpolation matrices with higher space regularity or splines can be used. Summarizing, the discrete version of (4.6) reads

$$\mathcal{F}(v; \mathbf{I}) := \sum_{k=1}^K \mathcal{S}(v_k) + \nu D_3(v_k) + \sum_{x \in \mathcal{G}} |P_I(x - Pv_k)I_k - P_I(x)I_{k+1}|^2.$$

It remains to discretize \mathcal{E} , which is done by using the midpoint rule for the data term. For the TV-term, forward differences D_{x_i} in x_i , $i = 1, 2$, direction are used

$$\text{TV}(I) := \left\| \sqrt{(D_{x_1}I)^2 + (ID_{x_2}^T)^2} \right\|_1,$$

where the square and the square root are meant componentwise, and $\|\cdot\|_1$ is the sum of the entries of the matrix. Then, the discrete functional reads

$$\mathcal{E}(I; B) = \frac{1}{2} \|AI - B\|_F^2 + \alpha \left\| \sqrt{(D_{x_1}I)^2 + (ID_{x_2}^T)^2} \right\|_1.$$

4.4.2 Proximal alternating linearized minimization (PALM)

Our first approach for the minimization of \mathcal{J} is based on PALM [36, 230]. This algorithm aims to minimize a functional

$$\arg \min_{x_1 \in E_1, x_2 \in E_2} \{H(x_1, x_2) + G_1(x_1) + G_2(x_2)\} \quad (4.13)$$

by iterating

$$\begin{aligned} x_1^{(j+1)} &= \text{prox}_{\tau G_1} \left(x_1^{(j)} - \frac{1}{\tau} \nabla_{x_1} H(x_1^{(j)}, x_2^{(j)}) \right), \\ x_2^{(j+1)} &= \text{prox}_{\sigma G_2} \left(x_2^{(j)} - \frac{1}{\sigma} \nabla_{x_2} H(x_1^{(j+1)}, x_2^{(j)}) \right), \end{aligned} \quad (4.14)$$

where $\tau, \sigma > 0$ and $\text{prox}_{\tau f}(x) := \arg \min_y \frac{1}{2} \|x - y\|_2^2 + \tau f(y)$ denotes the proximal mapping of f , which is uniquely determined for proper, convex and lsc functions f . The convergence result is stated in the following theorem from [36, Thm. 1]. Note that in the theorem the proximal map is also defined for non-convex functions. However, the involved functions in our application are convex, so that no further details on this topic are provided.

Theorem 4.10. *Let E_1, E_2 be Euclidean spaces and $H: E_1 \times E_2 \rightarrow \mathbb{R} \cup \{+\infty\}$, $G_i: E_i \rightarrow \mathbb{R} \cup \{+\infty\}$, $i = 1, 2$, be proper, lsc functions. Assume that H is continuously differentiable with locally Lipschitz continuous gradient and that both $x_i \mapsto \nabla_{x_i} H(x_1, x_2)$ are globally Lipschitz, where the constants $L_1(x_2), L_2(x_1)$ possibly depend on the fixed variable. Let $G_1 + G_2 + H$ in (4.13) fulfill the Kurdyka–Łojasiewicz (KL) property. Further, assume $\tau > L_1(x_2^{(j)})$ and $\sigma > L_2(x_1^{(j)})$ for all $j \in \mathbb{N}$. If the sequence generated by (4.14) is bounded, then it converges to a critical point.*

For our problem we choose the splitting

$$G_1(\mathbf{I}) = \alpha \text{TV}(I_0) + \frac{1}{2} \|AI_0 - B\|_F^2,$$

$$G_2(\mathbf{v}) = 0,$$

$$H(\mathbf{I}, \mathbf{v}) = \beta \left(\sum_{k=0}^{K-1} \sum_{x \in \mathcal{G}} |P_I(x - P_v v_k(x)) I_k - P_I(x) I_{k+1}|^2 + \alpha \mathcal{S}(v_k) + \nu D_3(v_k) \right).$$

Then, the iteration (4.14) reads

$$\begin{aligned} \mathbf{I}^{(j+1)} &= \text{prox}_{\tau G_1} \left(\mathbf{I}^{(j)} - \frac{1}{\tau} \nabla_{\mathbf{I}} H(\mathbf{I}^{(j)}, \mathbf{v}^{(j)}) \right), \\ \mathbf{v}^{(j+1)} &= \mathbf{v}^{(j)} - \frac{1}{\sigma} \nabla_{\mathbf{v}} H(\mathbf{I}^{(j+1)}, \mathbf{v}^{(j)}). \end{aligned}$$

From the structure of H we deduce that the v_k , $k = 0, \dots, K-1$, can be computed separately. Note that the second term in G_1 can also be added to H , but this only makes sense if $\|A^T A\|$ is small as otherwise the Lipschitz constant gets too large. The I_k , $k = 1, \dots, K$, can also be updated separately [36, Sec. 3.6], which possibly improves the Lipschitz constants. If the interpolation matrix P_I originates from smooth piecewise polynomial basis functions and τ, σ are chosen accordingly, all conditions of Theorem 4.10 are satisfied. The proximal map $\text{prox}_{\tau G_1}$ can be computed efficiently by primal-dual algorithms from convex analysis as for example by the Chambolle–Pock algorithm [57, 229]. Indeed there is a vast literature how to solve problems of this kind, see [53, 58] for an overview.

4.4.3 Alternating minimization approach

The computation of $\text{prox}_{\tau G_1}$ with a primal-dual algorithm requires an inner iteration for every step of PALM. If the evaluation of A is computationally expensive, this can result in high computational effort due to many operator evaluations. Our numerical experiments indicated that PALM needs relatively many outer iterations and hence also many evaluations of the operator. Therefore, we want to present a second alternating scheme to minimize $\mathcal{J}(\mathbf{I}, \varphi)$, which needed fewer outer iterations in our experiments.

Starting with $\mathbf{I}^{(0)}, \varphi^{(0)}$ we iterate for $j = 0, \dots$:

1. For $k = 0, 1, \dots, K-1$, we compute

$$\varphi_k^{(j)} = \arg \min_{\varphi_k \in \mathcal{A}} \left\{ \int_{\Omega} W(D\varphi_k) + \nu |D^m \varphi_k|^2 + |I_k^{(j)} \circ \varphi_k^{-1} - I_{k+1}^{(j)}|^2 dx \right\}.$$

2. For given $A \in L(L^2(\Omega), \mathcal{Y})$, $B \in \mathcal{Y}$ and $R \in L^2(\Omega)$, we solve

$$\mathbf{I}^{(j)} = \arg \min_{\mathbf{I} \in (L^2(\Omega))^K} \left\{ \beta \sum_{k=0}^{K-1} \|I_k \circ (\varphi_k^{(j)})^{-1} - I_{k+1}\|_{L^2(\Omega)}^2 + \frac{1}{2} \|AI_0 - B\|_{\mathcal{Y}}^2 + \alpha \text{TV}(I_0) \right\}. \quad (4.15)$$

For the first step the discretization from Section 4.4.1 is applied, resulting in the minimization of

$$\mathcal{R}(v_k; I_k^{(j)}, I_{k+1}^{(j)}) := \mathcal{S}(v_k) + \nu D_3(v_k) + \sum_{x \in \mathcal{G}} |P_I(x - Pv_k)I_k - P_I(x)I_{k+1}|^2, \quad (4.16)$$

for $k = 0, \dots, K-1$. This problem can be solved by a Quasi-Newton method, details can be found in [213, 226].

For the computation of the image sequence in the second step of the algorithm, we use the substitution from the proof of Lemma 4.5. Setting $\psi_k := \varphi_{k-1} \circ \dots \circ \varphi_0$, $w_k := \det D\psi_k$ and $F_0 := I_0$, $F_k := I_k \circ \psi_k$, we can transform (4.15) to

$$\arg \min_{\mathbf{F}} \left\{ \beta \sum_{k=0}^{K-1} \|(F_k - F_{k+1})\sqrt{w_{k+1}}\|_{L^2(\Omega)}^2 + \frac{1}{2} \|AF_0 - B\|_{\mathcal{Y}}^2 + \alpha \text{TV}(F_0) \right\}. \quad (4.17)$$

The functional is discretized on \mathcal{G} , using the approach from Section 4.4.1. We propose to solve the discrete version of (4.17) with a block-coordinate descent that fixes alternately F_0 and $\bar{\mathbf{F}} := (F_1, \dots, F_{K-1})$. For block-coordinate descent the following convergence result was proven in [23, Thm. 14.9, Thm. 14.15], see also [24].

Theorem 4.11. *Let E_1, E_2 be Euclidean spaces and $G: E_1 \times E_2 \rightarrow \mathbb{R} \cup \{+\infty\}$, $G_i: E_i \rightarrow \mathbb{R} \cup \{+\infty\}$, $i = 1, 2$, be proper, convex lsc functions. Assume further that G is continuously differentiable and that the level sets of $G + G_1 + G_2$ are bounded. Then the minimization problem*

$$\arg \min_{x_1 \in E_1, x_2 \in E_2} \{G(x_1, x_2) + G_1(x_1) + G_2(x_2)\}$$

can be solved by alternating minimization in x_1 and x_2 , i.e., every accumulation point of the generated iteration sequence is a minimizer. The convergence rate for the functional values is $\mathcal{O}(\frac{1}{k})$.

For our specific discretized problem (4.17) with

$$\begin{aligned} G(F_0, \bar{\mathbf{F}}) &:= \beta \sum_{x \in \mathcal{G}} |F_0(x) - F_1(x)|^2 w_1(x), \\ G_1(\bar{\mathbf{F}}) &:= \beta \sum_{k=0}^{K-1} \sum_{x \in \mathcal{G}} |F_k(x) - F_{k+1}(x)|^2 w_{k+1}(x), \\ G_2(F_0) &:= \frac{1}{2} \|AF_0 - B\|_{\mathcal{Y}}^2 + \alpha \text{TV}(F_0), \end{aligned}$$

the conditions of the theorem are obviously fulfilled. If F_0 is fixed, Corollary 4.6 implies that the minimizer of $G(F_0, \bar{\mathbf{F}}) + G_1(\bar{\mathbf{F}})$ is given analytically. In the second step of the algorithm we have to minimize, for fixed $\bar{\mathbf{F}}$, the functional

$$G(F_0, \bar{\mathbf{F}}) + G_2(F_0) = \beta \sum_{x \in \mathcal{G}} |F_0(x) - F_1(x)|^2 w_1(x) + \frac{1}{2} \|AF_0 - B\|^2 + \alpha \text{TV}(F_0). \quad (4.18)$$

This can be done efficiently by primal-dual algorithms from convex analysis, see Section 4.4.2 for a discussion. Finally, we use scattered interpolation to obtain the images \mathbf{I} at grid points from \mathbf{F} .

4.4.4 Multilevel strategy

As usual in optical flow and image registration, we apply a coarse-to-fine strategy with $\text{lev} \in \mathbb{N}$ levels if a downsampling procedure for the data and the operator is known. This is the case for our numerical experiments, but clearly it is also possible to use only a single level if no downsampling procedure is known. First, we iteratively smooth our given template image by convolution with a truncated Gaussian and downsampling using bilinear interpolation. Here, special care is necessary for the operator A , as well as for the downsampling procedure of the data B , which is dependent on the operator choice. Both procedures are described in the respective numerical examples.

In order to obtain a deformation on the coarsest level, a single registration is performed with the solution of the L^2 -TV problem, i.e.,

$$I_{0,\text{lev}} = \arg \min_{I: \mathcal{G}_{\text{lev}} \rightarrow \mathbb{R}} \left\{ \frac{1}{2} \|A_{\text{lev}} I - B_{\text{lev}}\|^2 + \alpha \text{TV}(I) \right\}, \quad \alpha > 0, \quad (4.19)$$

where $\text{lev} \in \mathbb{N}$ is number of levels. For better results, the regularization parameters for v are decreased successively as recommended by Modersitzki [205].

After computing a solution on every level, bilinear interpolation is applied to construct an initial deformation on the next finer level. The sequence of $\tilde{K} - 1$, $\tilde{K} < K$, intermediate finer level images is initialized from the end

$$I_k(x) = R\left(x + \frac{k}{K} P v(x)\right), \quad (4.20)$$

where R is the template image at the current level. Using this, we obtain an initial image sequence on this level. The complete multilevel strategy is sketched in Algorithm 4.1 for the alternating minimization scheme presented in Section 4.4.3.

4.5 Numerical examples

In this section, numerical examples demonstrating the potential of the method are presented. The proposed Algorithm 4.1 is implemented using MATLAB. We also implemented the minimization of the TDM-INV model using PALM, but observed higher computation times due to many operator evaluations. For the Radon transform, the computation roughly needed two times as long (about 5-10 minutes). As

Algorithm 4.1 TDM-INV Algorithm (informal)

```

1:  $R_0 := R, B_0 := B, \mathcal{G}_0 := \mathcal{G}$ 
2: create image stack  $(R_l)_{l=0}^{\text{lev}}, (B_l)_{l=0}^{\text{lev}}$  on  $(\mathcal{G}_l)_{l=0}^{\text{lev}}$  by downsampling
3: solve (4.19) for  $B_{\text{lev}}$ 
4: solve (4.16) for  $R_{\text{lev}}, I_{0,\text{lev}}$  to get  $\tilde{v}$ 
5:  $l \rightarrow \text{lev} - 1$ 
6: use bilinear interpolation to get  $v$  on  $\mathcal{G}_l$  from  $\tilde{v}$ 
7: obtain  $\tilde{K}_l$  images  $\mathbf{I}_l^{(0)}$  from  $R_l, v$  by (4.20)
8: while  $l \geq 0$  do
9:     repeat(Alternating outer iteration)
10:         find deformations  $\tilde{\mathbf{v}}_l^{(i+1)}$  minimizing (4.16) for every pair from  $\mathbf{I}_l^{(i)}$ 
11:         initialize  $\mathbf{F}^{(0)} = \mathbf{I}_l^{(i)}$ 
12:         repeat(Alternating inner iteration)
13:             given  $F_0^{(j)}$ , compute  $F_1^{(j+1)}, \dots, F_{K-1}^{(j+1)}$  according to Corollary 4.6
14:             compute  $F_0^{(j+1)}$  as solution of (4.18) using a PD-method using  $F_1^{(j+1)}$ 
15:              $j \rightarrow j + 1$ 
16:         until convergence criterion is reached
17:         compute  $\mathbf{I}_l^{(i)}$  from  $F^{(j)}$  using scattered interpolation
18:          $i \rightarrow i + 1$ 
19:     until convergence criterion is reached
20:      $l \rightarrow l - 1$ 
21:     if  $l > 0$  then
22:         use bilinear interpolation to get  $\mathbf{I}_l$  and  $\mathbf{v}_l$  on  $\mathcal{G}_l$ 
23:         for  $k = 1, \dots, \tilde{K}_l$  do
24:             compute  $\tilde{K}_l$  images connecting  $I_{l,k-1}, I_{l,k}$  using (4.20) with  $v_{l,k}$ 
25:  $\mathbf{I} := \mathbf{I}_0$ 
    
```

comparison a result using PALM is added in the first example. The qualitative differences between the two results are very small and therefore only the results of Algorithm 4.1 are shown in the remaining experiments. Note that PALM might be more favorable if the operator A is simple to evaluate, e.g., if it is sparse.

For representing our images on a grid during the registration step, we applied the Mex interface of the spline library by E. Bertolazzi² with Akima splines. In order to reduce the number of involved parameters in (4.16), we use $\lambda = \mu = \nu = 100\eta$ in all our experiments. Typical choices for the increments \tilde{K} are $\tilde{K}_{\text{lev}-1} = 2$, $\tilde{K}_{\text{lev}-2} = 1$ and $\tilde{K}_i = 0$ for the remaining levels. The other parameters α , β and λ are optimized with respect to SSIM via a gridsearch. For the comparisons, the algorithm parameters are SSIM optimized, too. A GPU implementation is applied for solving the appearing linear systems of equations in the Quasi-Newton method.

In the first part of our experiments, the Radon transform is considered as operator. Among the vast literature on the topic, we refer to the books [154, 179, 208] for a general introduction to CT including some reconstruction methods from in-

²<https://github.com/ebertolazzi/Splines>

complete data and for limited angle tomography to [86, 155, 190]. The second part deals with superresolution, which does not have a continuous counterpart.

4.5.1 Limited angle and sparse CT

We are given a reference image $R \in [0, 1]^{256, 256}$ and sinogram data of a target image $I_{\text{orig}} \in [0, 1]^{256, 256}$, which we want to reconstruct. For the numerical implementation of the (discrete) Radon transform the Astra toolbox [224, 268, 269] is used, which allows more flexibility compared to the built-in MATLAB function.

In our **first example**, the reference image consists of 6 triangular shaped objects, which are deformed to stars in the target image, see Fig. 4.3³. The sinogram is obtained by the Radon transform using 10 measurement directions equally distributed (with steps of 9 degrees) from 0 to 81 degrees, i.e., the measurement angle is limited to less than the half domain. The sinogram is additionally corrupted with 5 percent Gaussian noise. Our goal is to reconstruct the target from the given sinogram data. In the proposed multi grid approach a down-sampling by a factor of 0.5 is used. For the down-sampling of the sinogram, two neighboring rays are averaged and rescaled to the correct intensity. Note that this is easily possible if the number of rays is chosen for example to be 1.5 times the number of pixels per direction. The result of our TDM-INV algorithm is shown in Fig. 4.3c, where the parameters $\text{lev} = 4$, $\lambda = 0.07$, $\alpha = 0.05$ and $\beta = 0.1$ are used. Compared to the reconstruction using the L^2 -TV model (with $\lambda_{\text{TV}} = 0.05$) in Fig. 4.3f, our method is able to better deal with the missing data from 81 to 180 degrees. Visually, the result is almost perfect and also the SSIM value is very good. In Fig. 4.3d the numerical result using PALM is shown. The SSIM and PSNR values are similar to Fig. 4.3c and almost no difference is visible. The difference of both results is depicted in Fig. 4.3e and lies within the color range $[-0.07, 0.06]$.

In the **second example** a more structured image is treated. The given reference image depicts an artificial brain image, and the target can be considered as a deformed version, see Fig. 4.4⁴. The sinogram of the target is created using the Radon transform with 20 measurements equally distributed from 0 to 180 degrees and by adding 5 percent Gaussian noise. For the multi grid approach the procedure from the previous example is used. The result of TDM-INV is shown in Fig. 4.4c and was calculated with the parameters $\text{lev} = 5$, $\lambda = 0.08$, $\alpha = 0.025$ and $\beta = 0.5$. Since our model incorporates the reference information as compensation for the sparse data, the reconstruction is better than the one with the L^2 -TV model (with $\lambda_{\text{TV}} = 0.1$) in Fig. 4.4d.

4.5.2 Superresolution

Here, we are given a reference image $R \in [0, 1]^{256, 256}$ and a low resolution image $B \in [0, 1]^{64, 64}$ obtained by down-sampling of a target image $I_{\text{orig}} \in [0, 1]^{256, 256}$ with

³The images in Fig. 4.3(a) and (b) are taken from the paper [64] and were provided by Barbara Gris and Ozan Öktem.

⁴Available at <http://bigwww.epfl.ch/algorithms/mriphantom/>, see also [141].

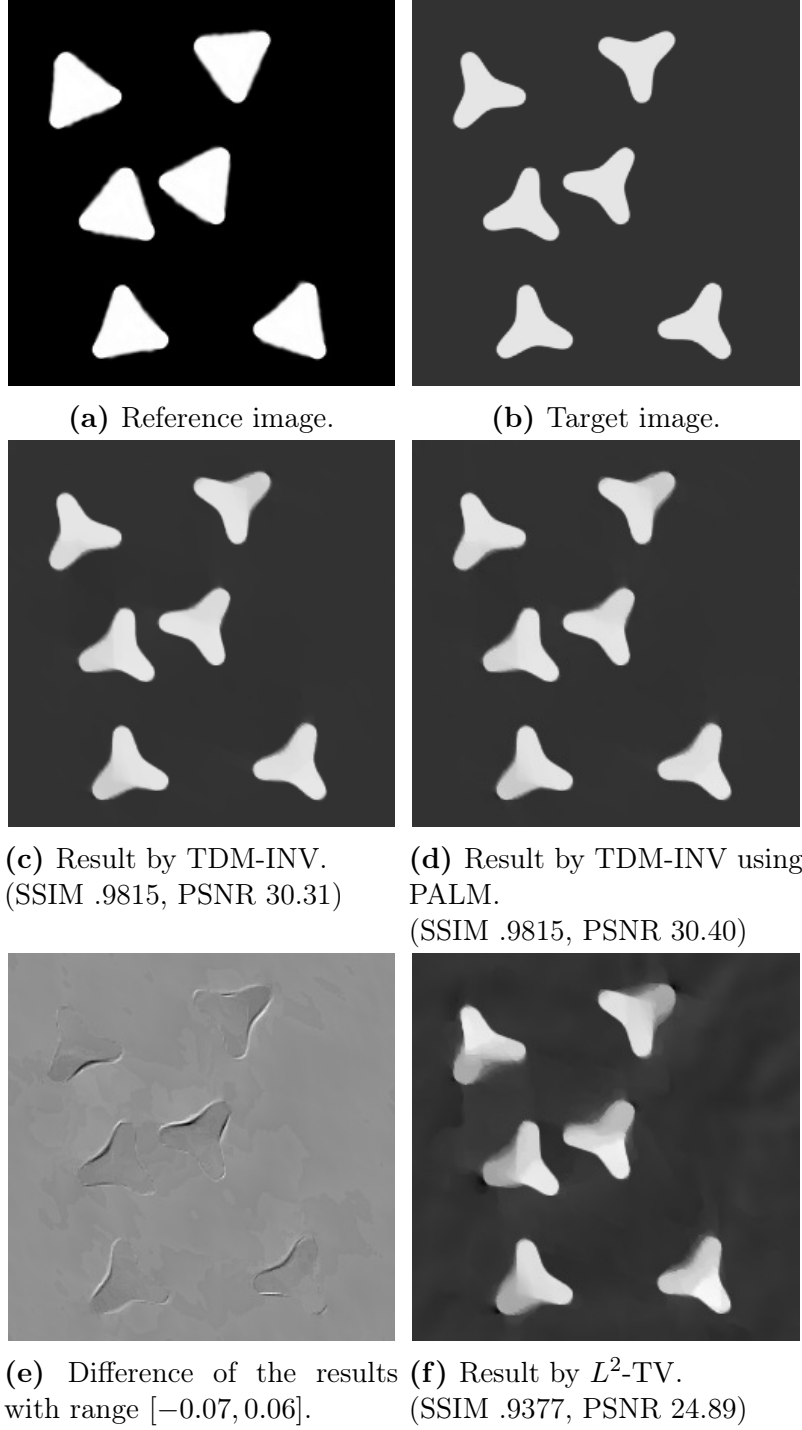


Figure 4.3: Image reconstruction from sparse, limited angle CT measurements from 0 to 81 degrees with 10 angles.




Figure 4.4: Image reconstruction from sparse CT measurements using 20 angles from 0 to 180 degrees.

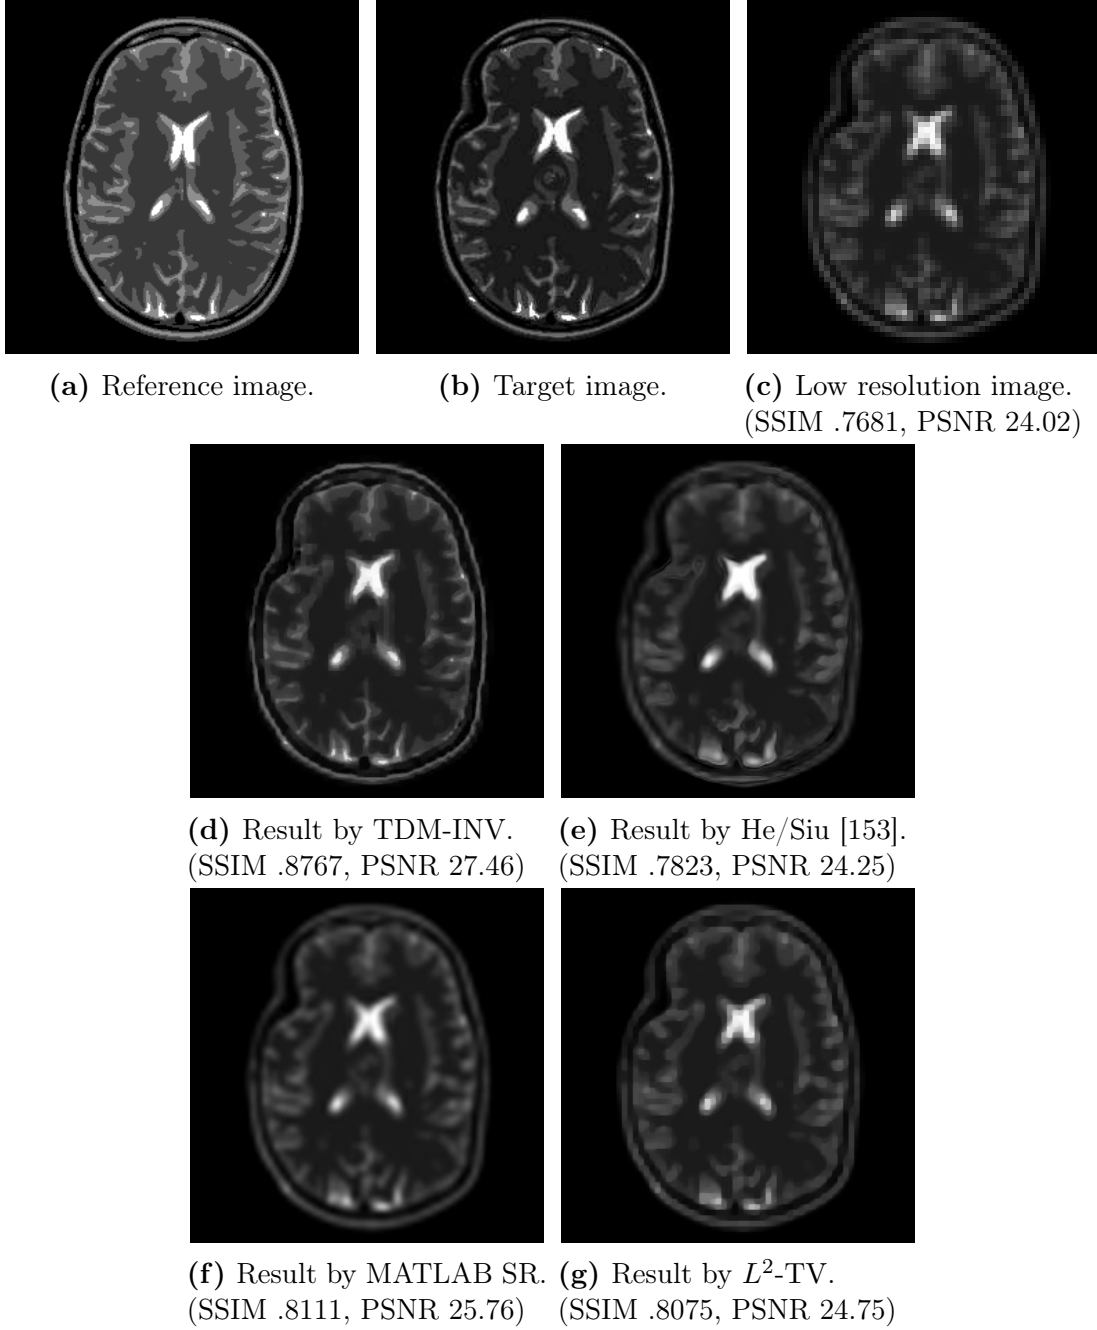


Figure 4.5: Superresolution from 64×64 pixels to 256×256 for brain image.

the down-sampling operator $P_4 \in \mathbb{R}^{256,64}$ given by

$$P_4 = \frac{1}{4} \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & & & & & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ & & & \ddots & & & \ddots & & & \\ 0 & & & & & & & 0 & 1 & 1 & 1 & 1 \end{pmatrix} \in \mathbb{R}^{256,64}.$$

In other words, $B = P_4 I_{\text{orig}} P_4^T$. For the multi grid approach a downscaling with factor 0.5 is applied such that the given image B can be used for the first three levels, i.e., $B_0 = B_1 = B_2 \in \mathbb{R}^{64,64}$. The matrix $P_4 \in \mathbb{R}^{256,64}$ is adapted to $P_2 \in \mathbb{R}^{128,64}$ for the second level and the identity matrix of corresponding size is used for all higher levels.

In our **third example** the same reference and target images as in the second example are used, see Fig. 4.5. The result of TDM-INV is shown in Fig. 4.5d, where the parameters are $\text{lev} = 4$, $\lambda = 0.01$, $\alpha = 0.001$ and $\beta = 2$. First, our method is compared with the single image superresolution method of He and Siu [153], which is based on a self-similarity assumption of the high and low resolution image together with a Gaussian process regression. In contrast to the result obtained by this method in Fig. 4.5e, our result does not have artifacts around the bright features. Using the MATLAB function `imresize`, the best reconstruction is obtained with the “lanczos3” kernel, see Fig. 4.5f, which is still affected by a strong blur. For this example, the L^2 -TV (parameter $\lambda_{\text{TV}} = 0.001$) reconstruction yields the result shown in Fig. 4.5g. Comparing all methods, we see that our method is best at recovering the fine details as well as the overall structure.

In our **last example**, the template image is not only deformed and scaled, but also a new detail is included in the image. As mass can be created on the image path, our method is able to reconstruct also the small detail, cf. Fig. 4.6⁵, where the parameters are chosen as $\text{lev} = 4$, $\lambda = 0.01$, $\alpha = 0.001$ and $\beta = 2$. For this simpler image, our method leads to the best result in SSIM and PSNR. The result produced by [153] in Fig. 4.6e yields almost the same SSIM, but visually the method recovers a lot of background noise. Here, the best result of MATLAB’s `imresize` is given by “bilinear” interpolation, which is still affected by a strong blur. The L^2 -TV approach (parameter $\lambda_{\text{TV}} = 0.001$) works better for this simple image than for the previous example, but is still not able to match our result. Especially the overlapping part in the center of the phantom is only recovered by TDM-INV.

4.6 Conclusions

This chapter merges the edge-preserving L^2 -TV variational model for solving inverse image reconstruction problems with a metamorphosis-inspired approach to utilize information from a reference image. The approach, called TDM-INV, can handle intensity changes between the reference image and the target image that we want to reconstruct. The method gives very good results for artificial images so that

⁵The images used in Fig. 4.6 are based on the ones in [139].

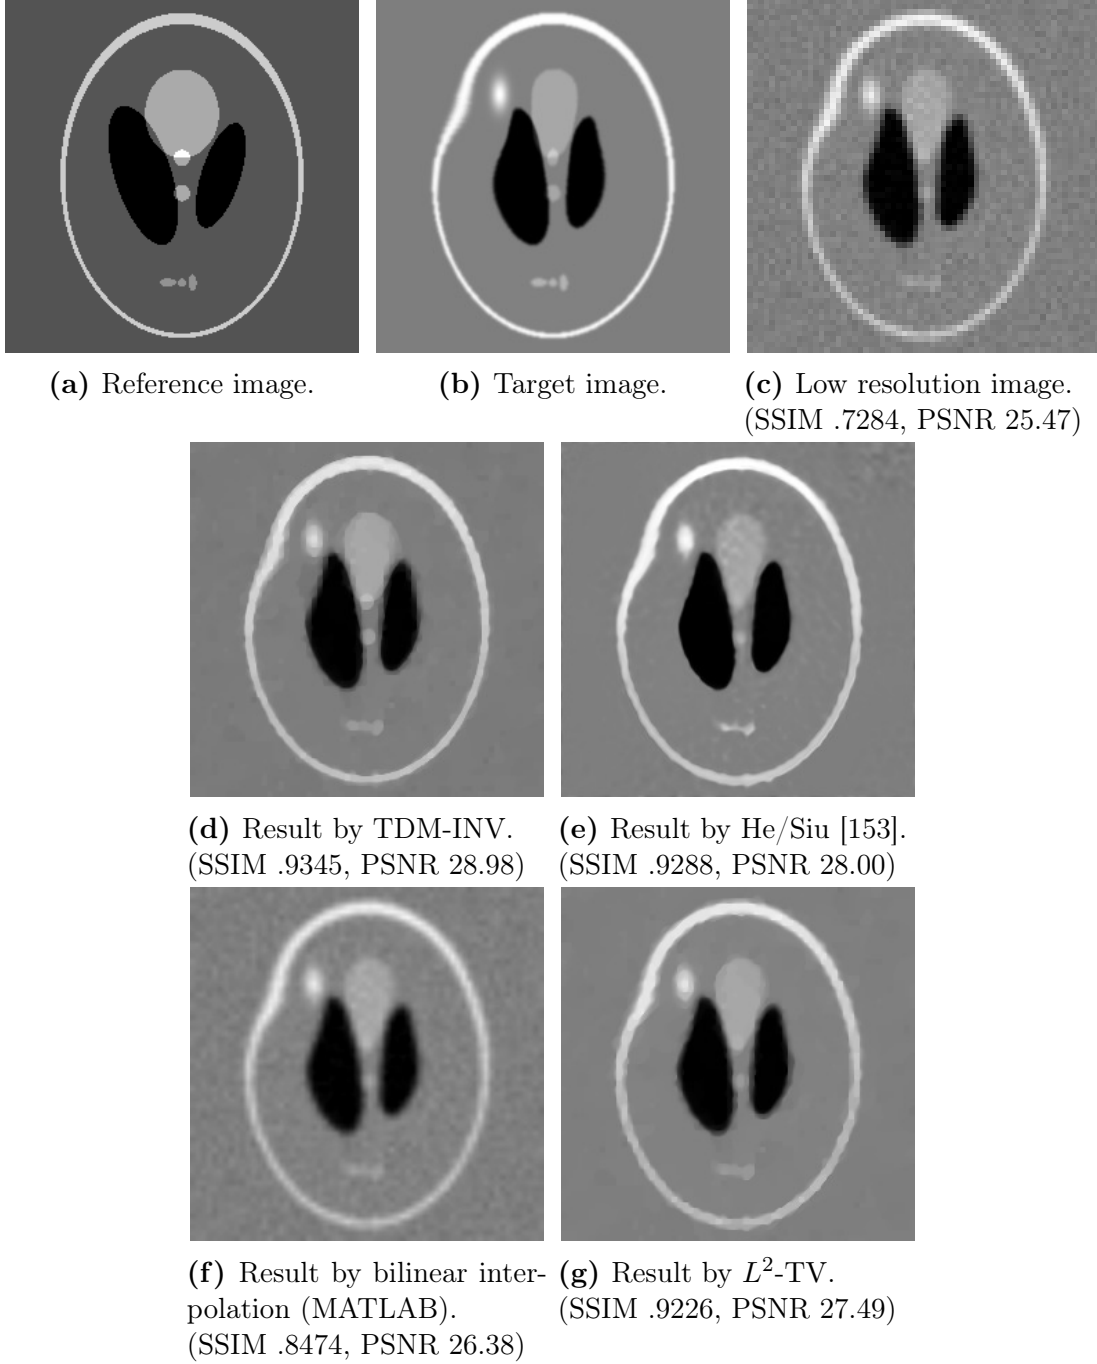


Figure 4.6: Superresolution from 64×64 pixels to 256×256 for Shepp-Logan phantom.

we are looking forward to real-world applications in material sciences or medical imaging, e.g., motion models for organs [103, 126]. Several extensions of the model are possible. Due to the finite difference approach and the design of the method, more sophisticated regularizers than the TV-term can be simply incorporated. Another possible modification would be to apply different transport models, see [191]. Further, the usage of multiple reference images can be taken into account.

Gagliardo–Nirenberg inequality

Theorem 4.12 (Gagliardo–Nirenberg [220]). *Let $\Omega \subset \mathbb{R}^n$ be a bounded domain satisfying the cone property. For $1 \leq q, r \leq \infty$, suppose that f belongs to $L^q(\Omega)$ and its derivatives of order m to $L^r(\Omega)$. Then for the derivatives $D^j f$, $0 \leq j < m$, the following inequalities hold true with constants C_1, C_2 independent of f :*

$$\|D^j f\|_{L^p(\Omega)} \leq C_1 \|D^m f\|_{L^r(\Omega)}^a \|f\|_{L^q(\Omega)}^{1-a} + C_2 \|f\|_{L^q(\Omega)},$$

where $\frac{1}{p} = \frac{j}{n} + a\left(\frac{1}{r} - \frac{m}{n}\right) + (1-a)\frac{1}{q}$ for all $a \in [\frac{j}{m}, 1]$, except for the case $1 < r < \infty$ and $m - j - \frac{n}{r}$ is a non-negative integer, in which the inequality only holds true for $a \in [\frac{j}{m}, 1)$.

Remark 4.13. For $p = q = r = 2$ the inequality simplifies to

$$\begin{aligned} \|D^j f\|_{L^2(\Omega)} &\leq C_1 \|D^m f\|_{L^2(\Omega)}^{\frac{j}{m}} \|f\|_{L^2(\Omega)}^{1-\frac{j}{m}} + C_2 \|f\|_{L^2(\Omega)} \\ &\leq C_1 \|D^m f\|_{L^2(\Omega)} + (C_1 + C_2) \|f\|_{L^2(\Omega)}, \end{aligned}$$

where the second inequality follows by estimating the product with the maximum of both factors.

Template-Based Image Reconstruction from Sparse Tomographic Data

Contents

5.1	Introduction	107
5.2	Theoretical results on the transport and continuity equation .	111
5.2.1	Continuity equation	111
5.2.2	Transport equation with H^1 regularity	114
5.3	Regularizing properties of template-based image reconstruction	115
5.4	Numerical solution	117
5.5	Numerical examples	122
5.6	Conclusions	128

Abstract

In this chapter¹, which is published in [181], we propose a variational regularization approach for the problem of template-based image reconstruction from indirect, noisy measurements as given, for instance, in X-ray computed tomography. An image is reconstructed from such measurements by deforming a given template image. The image registration is directly incorporated into the variational regularization approach in form of a partial differential equation that models the registration as either mass- or intensity-preserving transport from the template to the unknown reconstruction. We provide theoretical results for the proposed variational regularization in both cases. In particular, we prove existence of a minimizer, stability with respect to the data, and convergence for vanishing noise when either of the abovementioned equations is imposed and more general distance functions are used. Numerically, we solve the problem by extending existing Lagrangian methods and propose a multilevel approach that is applicable whenever a suitable down-sampling procedure for the operator and the measured data can be provided. Finally, we demonstrate the performance of our method for template-based image reconstruction from highly undersampled and noisy Radon transform data. We compare results for mass- and intensity-preserving image registration, various regularization functionals, and different distance functions. Our results show that very reasonable reconstructions can be obtained when only few measurements are available and demonstrate that the use of a normalized cross-correlation based distance is advantageous in case the image intensities between the template and the unknown image differ substantially.

¹This is a post-peer-review, pre-copyedit version of an article published in Applied Mathematics and Optimization. The final authenticated version is available online at DOI: 10.1007/s00245-019-09573-2. The article is distributed under the terms of the Creative Commons Attribution 4.0 International License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original authors and the source, provide a link to the Creative Commons license, and indicate if changes were made.

5.1 Introduction

In medical imaging, an image can typically not be observed directly but only through indirect and potentially noisy measurements, as it is the case, for example, in computed tomography (CT) [208]. Due to the severe ill-posedness of the problem, reconstructing an image from measurements is rendered particularly challenging when only few or partial measurements are available. This is, for instance, the case in limited-angle computed tomography [116, 208], where limited-angle data is acquired in order to minimize exposure time of organisms to X-radiation. Therefore, it can be beneficial to impose a priori information on the reconstruction, for instance, in the form of a template image. However, typically neither its exact position nor its exact shape is known.

In *image registration*, the goal is to find a reasonable deformation of a given template image so that it matches a given target image as closely as possible according to a predefined similarity measure, see [205, 204] for an introduction. When the target image is unknown and only given through indirect measurements, it is referred to as *indirect image registration* and has been explored only recently [64, 139, 160, 223]. As a result, a deformation together with a transformed template can be computed from tomographic data. The prescribed template acts as a prior for the reconstruction and, when chosen reasonably close in a deformation sense, gives outstanding reconstructions in situations where only few measurements are available and competing methods such as filtered backprojection [208] or total variation regularization [238] fail, see [64, Sec. 10].

In our setting, deformations are maps from the image domain $\Omega \subset \mathbb{R}^n$, $n \in \mathbb{N}$, to itself together with an action that specifies exactly how such a map deforms elements in the shape space, which in this work is the space $L^2(\Omega, \mathbb{R})$ of gray-scale images supported in the image domain. Natural problems are to characterize admissible deformations and to compute these numerically in an efficient manner.

One possible approach is diffeomorphic image registration, where the set of admissible deformations is restricted to diffeomorphisms in order to preserve the topology of structures within an image [285]. One can, for instance, consider the group of diffeomorphisms together with the composition as group operation. Elements in this group act on gray-scale images by means of the group action and thereby allow for a rich set of non-rigid deformations, as required in many applications. For instance, the geometric group action transforms gray-scale images in a way such that its intensity values are preserved, whereas the mass-preserving group action ensures that, when the image is regarded as a density, the integral over the density is preserved.

A computational challenge in using the above group formalism is that it lacks a natural vector space structure, which is typically desired for the numerical realization of the scheme. Hence, it is convenient to further restrict the set of admissible deformations. One way to obtain diffeomorphic deformations is to perturb the identity map with a displacement vector field. Provided that the vector field is reasonably small and sufficiently regular, the resulting map is invertible [285, Prop. 8.6]. For indirect image registration this idea was pursued in [223].

The basic idea of the *large deformation diffeomorphic metric mapping* (LDDMM) [25, 95, 200, 201, 262, 265, 285] framework is to generate large deformations by considering flows of diffeomorphisms that arise as the solution of an ordinary differential equation (ODE), the so-called flow equation, with velocity fields that stem from a reproducing kernel Hilbert space. In order to ensure that the flow equation admits a unique solution, one typically chooses this vector space so that it can be continuously embedded into $C^1(\Omega, \mathbb{R}^n)$, allowing the application of existence and uniqueness results from Cauchy–Lipschitz theory for ODEs, see [79, Chap. 1] for a brief introduction. In [64], the LDDMM framework is adapted for indirect image registration and the authors prove existence, stability, and convergence of solutions for their variational formulation. Numerically, the problem is solved by gradient descent.

The variational problem associated with LDDMM is typically formulated as an ODE-constrained optimization problem. As the flow equation can be directly related to hyperbolic partial differential equations (PDE) via the method of characteristics [106, Chap. 3.2], the problem can equivalently be rephrased as a PDE-constrained optimization problem [162]. The resulting PDE is determined by the chosen group action, see [64, Sec. 6.1.1] for a brief discussion. For instance, the geometric group action is associated with the transport (or advection) equation, while the mass-preserving group action is associated with the continuity equation. It is important to highlight that the PDE constraint implements both the flow equation and the chosen diffeomorphic group action.

Such an optimal control approach was also pursued for motion estimation and image interpolation [7, 40, 41, 52, 65, 148, 219]. In the terminology of optimal control, the PDE represents the state equation, the velocity field the control, and the transformed image the resulting state. We refer to the books [42, 88, 142, 161] and to the article [157] for a general introduction to PDE-constrained optimization and suitable numerical methods. Let us mention that other methods, such as geodesic shooting [11, 200, 248, 272], exist and constitute particularly efficient numerical approaches. In particular, this direction has recently been combined with machine learning methods [284].

A particularly challenging scenario for diffeomorphic image registration occurs when the target image is not contained in the orbit of the template image under abovementioned group action of diffeomorphisms. For instance, this could happen in the case of the geometric group action due to the appearance of new structures in the target image or due to a discrepancy between the image intensities of the template and the target image. A possible solution is provided by the *metamorphosis* framework [202, 232, 263, 264], which is an extension to LDDMM that allows for modulations of the image intensities along characteristics of the flow. The image intensities change according to an additional flow equation with an unknown source. See [285, Chap. 13] for a general introduction and, for instance, [162] for an application to magnetic resonance imaging. Let us also mention [214], which adopts a discrete geodesic path model for the purpose of image reconstruction, and [191], in which the metamorphosis model is combined with optimal transport.

In [139], the metamorphosis framework is adapted for indirect image registration.

The authors prove that their formulation constitutes a well-defined regularization method by showing existence, stability, and convergence of solutions. However, in the setting where only few measurements, e.g., a few directions in CT, are available, reconstruction of appearing or disappearing structures seems very challenging.

Therefore, in order to obtain robustness with respect to differences in the intensities between the transformed template and the sought target image, we follow a different approach. We consider not only the standard *sum-of-squared differences* (*SSD*) but also a distance that is based on the *normalized cross-correlation* (*NCC*) [205, Chap. 7.2], as it is invariant with respect to a scaling of the image intensities.

While image registration itself is already an ill-posed inverse problem that requires regularization [105], the indirect setting as described above is intrinsically more challenging. It can be phrased as an inverse problem, where measurements (or observations) $g \in Y$ are related to an unknown quantity $f \in X$ via the operator equation

$$K \circ f = g + n^\delta. \quad (5.1)$$

Here, $K: X \rightarrow Y$ is a (not necessarily linear) operator that models the data acquisition, often by means of a physical process, n^δ are measurement errors such as noise, and X and Y are Banach spaces. When f constitutes an image and g are tomographic measurements, solving (5.1) is often referred to as *image reconstruction*.

We use a variational scheme [244] to solve the inverse problem of indirect image registration, which can be formulated as a PDE-constrained optimization problem [64, Sec. 6.1.1]. It is given by

$$\begin{aligned} \min_{v \in V} J_{\gamma,g}(v), \\ \text{s.t. } C(v), \end{aligned} \quad (5.2)$$

where $J_{\gamma,g}: V \rightarrow [0, +\infty]$ is the functional

$$v \mapsto D(K \circ f_v(T, \cdot), g) + \gamma \|v\|_V^2. \quad (5.3)$$

Here, V is an admissible vector space with norm $\|\cdot\|_V$, $D: Y \times Y \rightarrow \mathbb{R}_{\geq 0}$ is a data fidelity term that quantifies the misfit of the solution against the measurements, and $\gamma > 0$ is a regularization parameter. Moreover, $f_v(T, \cdot): \Omega \rightarrow \mathbb{R}$ denotes the evaluation at time $T > 0$ of the (weak) solution of $C(v)$, which is either the Cauchy problem

$$C(v) = \begin{cases} \frac{\partial}{\partial t} f(t, x) + v(t, x) \nabla_x f(t, x) = 0, & \text{for } (t, x) \in [0, T] \times \Omega, \\ f(0, x) = f_0(x), & \text{for } x \in \Omega, \end{cases}$$

governed by the transport equation, or

$$C(v) = \begin{cases} \frac{\partial}{\partial t} f(t, x) + \operatorname{div}_x(v(t, x) f(t, x)) = 0, & \text{for } (t, x) \in [0, T] \times \Omega, \\ f(0, x) = f_0(x), & \text{for } x \in \Omega, \end{cases}$$

involving the continuity equation. Here, $f_0 \in L^2(\Omega, \mathbb{R})$ denotes an initial condition, which in our case is the template image.

The main goals of this article are the following. First, to study variational and regularizing properties of problem (5.2), and to develop efficient numerical methods for solving it. Second, to investigate alternative choices of distance functions D , such as the abovementioned NCC-based distance. Third, to demonstrate experimentally that excellent reconstructions can be computed from highly undersampled and noisy Radon transform data.

Our numerical approach is based on the Lagrangian methods developed in [193], called LagLDDMM. In contrast to most existing approaches, which are mainly first-order methods (see [193] for a brief classification and discussion), LagLDDMM uses a Gauss–Newton–Krylov method paired with Lagrangian solvers for the hyperbolic PDEs listed above. The characteristics associated with these PDEs are computed with an explicit Runge–Kutta method. One of the main advantages of this approach is that Lagrangian methods are unconditionally stable with regard to the admissible step size. Furthermore, the approach limits numerical diffusion and, in order to evaluate the gradient or the Hessian required for optimization, does not require the storage of multiple space-time vector fields or images at intermediate time instants. The scheme can also be implemented matrix-free.

In comparison to abovementioned existing methods for indirect image registration, such as [64, 139, 160, 223], our method is conceptually different in several ways. The first difference concerns the discretization. While [64, 139, 223] are mainly based on small deformations and use reproducing kernel Hilbert spaces, our method relies on nonparametric registration. The main advantages are that it directly allows for a multilevel approach and no kernel parameters need to be chosen. Moreover, due to the flexibility of the underlying framework it is straightforward to extend our method to parametric registration. Second, our approach relies on second-order methods for optimization by using a Gauss–Newton method paired with line search, while the other methods mainly rely on gradient descent. This allows for a fast decrease of the objective within only few iterations. Third, our method allows to easily exchange the underlying PDE solver. Essentially, any solver can be used as long as it can be differentiated efficiently. The used explicit Runge–Kutta method has the advantage that it does not require the storage of multiple images or repeated interpolation of the template, which can potentially lead to a blurred solution. Finally, let us mention that [160] is conceptually different since both a deformation and a template image are computed. Our main focus, however, are applications where only very few and noisy measurements are available and the problem of estimating an additional template seems highly underdetermined in such situations.

Contributions The contributions of this article are as follows. First, we provide the necessary theoretical background on (weak) solutions of the continuity and the transport equation, and recapitulate existence and uniqueness theory for characteristic curves for the associated ODE. In contrast to the results derived in [64], where

the template image is assumed to be contained in the space $SBV(\Omega, \mathbb{R}) \cap L^\infty(\Omega, \mathbb{R})$ of essentially bounded functions with special bounded variation, our results only require $L^2(\Omega, \mathbb{R})$ regularity. In addition, by using results from [92], we are able to consider the transport equation in the setting with H^1 regularity of vector fields in space as well as in time and with bounded divergence. Moreover, we show the existence of a minimizer of the problem (5.2), stability with respect to the data, and convergence for vanishing noise.

Second, in order to solve the problem numerically, we follow a *discretize-then-optimize* approach and extend the LagLDDMM framework [193] to the indirect setting. The library itself is an extension of FAIR [205] and, as a result, our implementation provides great flexibility regarding the selected PDE, and can easily be extended to other distances as well as to other regularization functionals. The source code of our MATLAB implementation is available online.²

Finally, we present numerical results for the abovementioned distances and PDEs. To the best of our knowledge, the results obtained for indirect image reconstruction based on the continuity equation are entirely novel. Moreover, we propose to use the NCC-based distance instead of SSD whenever the image intensities of the template and the unknown target are far apart, and show its numerical feasibility.

5.2 Theoretical results on the transport and continuity equation

In this section, we review the necessary theoretical background, and state results on the existence and stability of weak solutions of the transport and the continuity equation. Compared to [64], our results are stronger since we do not require space regularity of the template image.

5.2.1 Continuity equation

In what follows, we consider well-posedness of the continuity equation, which arises in the LDDMM framework using the mass-preserving group action via the method of characteristics. The regularity assumptions on v are such that we can apply the theory from [263].

Let $\Omega \subset \mathbb{R}^n$ be a bounded, open, convex domain with Lipschitz boundary and let $T > 0$. In the following, we examine the continuity equation

$$\begin{cases} \frac{\partial}{\partial t} f(t, x) + \operatorname{div}_x(v(t, x)f(t, x)) = 0 & \text{for } (t, x) \in [0, T] \times \Omega, \\ f(0, x) = f_0(x) & \text{for } x \in \Omega, \end{cases} \quad (5.4)$$

with coefficients $v \in L^2([0, T], \mathcal{V})$ and initial condition $f_0 \in L^2(\Omega, \mathbb{R})$, where \mathcal{V} is a Banach space that is continuously embedded into $C_0^{1,\alpha}(\Omega, \mathbb{R}^n)$ for some $\alpha > 0$. Here, $C_0^{1,\alpha}(\Omega, \mathbb{R}^n)$ denotes the closure of $C_c^\infty(\Omega, \mathbb{R}^n)$ under the $C^{1,\alpha}$ norm. Note that such

²<https://doi.org/10.5281/zenodo.2598138>

velocity fields can be continuously extended to the boundary. Clearly, (5.4) has to be understood in a weak sense, i.e., a function $f \in C^0([0, T], L^2(\Omega, \mathbb{R}))$ is said to be a weak solution of (5.4) if

$$\int_0^T \int_{\Omega} f(t, x) (v(t, x) \nabla_x \eta(t, x) + \frac{\partial}{\partial t} \eta(t, x)) \, dx \, dt + \int_{\Omega} f_0(x) \eta(0, x) \, dx = 0 \quad (5.5)$$

holds for all $\eta \in C_c^\infty([0, T] \times \Omega)$. The corresponding characteristic ODE is

$$\begin{cases} \frac{d}{dt} X(t, s, x) = v(t, X(t, s, x)) & \text{for } (t, s, x) \in [0, T]^2 \times \Omega, \\ X(s, s, x) = x & \text{for } x \in \Omega. \end{cases} \quad (5.6)$$

In this notation, the first argument of X is the time dependence, the second the initial time, and the third the initial space coordinate. The following theorem is a reformulation of [263, Thms. 1 and 9] and characterizes solutions of (5.6).

Theorem 5.1. *Let $v \in L^2([0, T], \mathcal{V})$ and $s \in [0, T]$ be given. There exists a unique global solution $X(\cdot, s, \cdot) \in C^0([0, T], C^1(\overline{\Omega}, \mathbb{R}^n))$ such that $X(s, s, x) = x$ for all $x \in \Omega$ and*

$$\frac{d}{dt} X(t, s, x) = v(t, X(t, s, x))$$

in weak sense. The solution operator $X_v: L^2([0, T], \mathcal{V}) \rightarrow C^0([0, T] \times \overline{\Omega}, \mathbb{R}^n)$ assigning a flow X_v to every velocity field v is continuous with respect to the weak topology in $L^2([0, T], \mathcal{V})$.

As $X(0, t, X(t, 0, x)) = x$, we can directly conclude that $X(t, 0, \cdot)$ is a diffeomorphism for every $t \in [0, T]$. Now, the diffeomorphism $X(0, t, x)$ can be used to characterize solutions of (5.4) as follows.

Proposition 5.2. *If $v \in L^2([0, T], \mathcal{V})$, then the unique weak solution of (5.4), as defined in (5.5), is given by $f(t, x) = \det(\mathcal{D}_x X(0, t, x)) f_0(X(0, t, x))$, where $\mathcal{D}_x X$ denotes the Jacobian of X .*

Proof. The proof is divided in three steps. First, we show that f satisfies the regularity conditions of weak solutions. For this purpose, the first step is to show $X(0, \cdot, \cdot) \in C^0([0, T], C^0(\overline{\Omega}, \mathbb{R}^n))$, i.e., that the flow is continuous in the initial values. Clearly, $X(0, t, \cdot) \in C^0(\overline{\Omega}, \mathbb{R}^n)$ for every $t \in [0, T]$. For an arbitrary sequence $t_i \rightarrow t$ we get

$$\|X(0, t_i, \cdot) - X(0, t, \cdot)\|_{C^0(\overline{\Omega})} \leq \|\mathcal{D}_x X(0, t_i, \cdot)\|_{C^0(\overline{\Omega})} \|\text{Id} - X(t_i, t, \cdot)\|_{C^0(\overline{\Omega})} \rightarrow 0,$$

where the first factor is bounded due to [263, Lem. 9]. Next, using the sequence $X_i(\cdot) = X(0, t_i, \cdot)$, it follows $f_0(X(0, \cdot, \cdot)) \in C^0([0, T], L^2(\Omega, \mathbb{R}))$, where the continuity in time follows from [213, Cor. 3]. Then, by differentiating $X(0, t, X(t, 0, x)) = x$ and rearranging the terms we obtain

$$\det(\mathcal{D}_x X(0, \cdot, \cdot)) = \det(\mathcal{D}_x X(\cdot, 0, \cdot))^{-1} (X(0, \cdot, \cdot)) \in C^0([0, T] \times \overline{\Omega}),$$

since all expressions are continuous. Finally, we conclude $f \in C^0([0, T], L^2(\Omega, \mathbb{R}))$, which follows from

$$\begin{aligned} & \|f(t, \cdot) - f(t_i, \cdot)\|_{L^2(\Omega)} \\ & \leq \|\det(\mathcal{D}_x X(0, t, x)) - \det(\mathcal{D}_x X(0, t_i, x))\|_{C^0(\overline{\Omega})} \|f_0 \circ X(0, t, x)\|_{L^2(\Omega)} \\ & \quad + \|\det(\mathcal{D}_x X(0, t_i, x))\|_{C^0(\overline{\Omega})} \|f_0 \circ X(0, t, x) - f_0 \circ X(0, t_i, x)\|_{L^2(\Omega)}, \end{aligned}$$

since both summands converge to zero.

The second step is to show that (5.5) is satisfied. Note that $X(\cdot, 0, x)$ is differentiable in t for a.e. $t \in [0, T]$, since it is absolutely continuous by definition. By inserting f into (5.5) and using the transformation formula, we get

$$\begin{aligned} & \int_0^T \int_{\Omega} f(t, x) (v(t, x) \nabla_x \eta(t, x) + \frac{\partial}{\partial t} \eta(t, x)) \, dx \, dt + \int_{\Omega} f_0(x) \eta(0, x) \, dx \\ & = \int_0^T \int_{\Omega} \det(\mathcal{D}_x X(t, 0, x)) f(t, X(t, 0, x)) \frac{d}{dt} \eta(t, X(t, 0, x)) \, dx \, dt + \int_{\Omega} f_0(x) \eta(0, x) \, dx \\ & = \int_0^T \int_{\Omega} f_0(x) \frac{d}{dt} \eta(t, X(t, 0, x)) \, dx \, dt + \int_{\Omega} f_0(x) \eta(0, x) \, dx = 0. \end{aligned}$$

For the last equality we used that $\eta(t, X(t, 0, x))$ is absolutely continuous.

The last step is to prove uniqueness of weak solutions, i.e., that every solution has the given form. Let f_1, f_2 be two different solutions, then we can find a t such that $\|f_1(t, \cdot) - f_2(t, \cdot)\|_{L^2(\Omega)} > 0$. By continuity in time we can find an interval I of length $\delta > 0$ that contains t , and a constant $c > 0$ such that

$$\|f_1(s, \cdot) - f_2(s, \cdot)\|_{L^2(\Omega)} \geq c$$

for all $s \in I$. However, weak solutions are unique in $L^\infty([0, T], L^2(\Omega, \mathbb{R}))$, see [92, Cor. II.1], where we used the embedding of \mathcal{V} into $C_0^1(\Omega, \mathbb{R}^n)$. This yields a contradiction. \square

Additionally, we can state and prove the following stability result for solutions of (5.4).

Proposition 5.3 (Stability). *Let $v_i \rightharpoonup v$ in $L^2([0, T], \mathcal{V})$ and f_i denote the weak solution of (5.4) corresponding to v_i . Then for every $t \in [0, T]$, there exists a subsequence, also denoted with f_i , such that $f_i(t, \cdot) \rightarrow f(t, \cdot)$ in $L^2(\Omega, \mathbb{R})$.*

Proof. The solution of (5.6) corresponding to v_i is denoted by X_i . Fix an arbitrary $t \in [0, T]$. From Theorem 5.1 we conclude $\|X_i(0, t, \cdot) - X(0, t, \cdot)\|_{C^0(\overline{\Omega})} \rightarrow 0$. Further, [97, Thm. 3.1.10] implies that $X_i(0, t, \cdot)$ is uniformly bounded for all $i \in \mathbb{N}$ in $C^{1,\alpha}(\overline{\Omega})$, implying $f_0 \circ X_i(0, t, \cdot) \rightarrow f_0 \circ X(0, t, \cdot)$ in $L^2(\Omega, \mathbb{R})$ by [213, Cor. 3].

It is left to show that a subsequence, also denoted by X_i , exists such that $X_i(0, t, \cdot) \rightarrow X(0, t, \cdot)$ in $C^1(\overline{\Omega}, \mathbb{R}^n)$. This concludes the proof since it also implies the convergence of $\det(\mathcal{D}_x X_i(0, t, \cdot)) \rightarrow \det(\mathcal{D}_x X(0, t, \cdot))$ in $C^0(\overline{\Omega})$. However, $X_i(0, t, \cdot)$ is uniformly bounded in $C^{1,\alpha}(\overline{\Omega}, \mathbb{R}^n)$ and it follows that $\mathcal{D}_x X_i(0, t, \cdot)$ is uniformly bounded in $C^{0,\alpha}(\overline{\Omega}, \mathbb{R}^{n \times n})$. By using the compact embedding of $C^{0,\alpha}(\overline{\Omega}, \mathbb{R}^{n \times n})$ into $C^0(\overline{\Omega}, \mathbb{R}^{n \times n})$, see [127, Lem. 6.33], there exists a subsequence of $X_i(0, t, \cdot)$ that converges to $X(0, t, \cdot)$ in $C^1(\overline{\Omega}, \mathbb{R}^n)$. \square

5.2.2 Transport equation with H^1 regularity

Here, we prove well-posedness of the transport equation, which arises in the LD-DMM framework using the geometric group action. Compared to the previous section, the space regularity assumptions on v are weaker and fit the setting in [92].

The transport equation reads as

$$\begin{cases} \frac{\partial}{\partial t} f(t, x) + v(t, x) \nabla_x f(t, x) = 0 & \text{for } (t, x) \in [0, T] \times \Omega, \\ f(0, x) = f_0(x) & \text{for } x \in \Omega, \end{cases} \quad (5.7)$$

with coefficients

$$v \in A := \{v \in H^1([0, T] \times \Omega)^n \cap L^2([0, T], H_0^1(\Omega)^n) : \|\operatorname{div}_x v\|_{L^\infty([0, T] \times \Omega)} \leq C\}$$

for some fixed constant C and initial value $f_0 \in L^2(\Omega, \mathbb{R})$. The admissible set A consists of all H^1 functions that are zero on the boundary of the spatial domain and have bounded divergence in L^∞ norm.

Note that the set A is a subset of $H^1([0, T] \times \Omega)^n$, which is closed and convex so that it is a weakly closed subset of a reflexive Banach space. In the following, we only check that A is closed. Let v_i be a convergent sequence in A with limit v . Since the two involved spaces are Banach spaces, we only have to check that v satisfies the constraint. Assume that $\|\operatorname{div}_x v\|_{L^\infty([0, T] \times \Omega)} > C$, then there exists a set B with positive measure and an $\epsilon > 0$ such that for all $x \in B$ we have $|\operatorname{div}_x v(x)| \geq C + \epsilon$. Hence, we get $\|\operatorname{div}_x v_i - \operatorname{div}_x v\|_{L^2([0, T] \times \Omega)} \geq \sqrt{\mu(B)}\epsilon$, which contradicts the convergence in H^1 .

Again, (5.7) has to be understood in weak sense so that $f \in C^0([0, T], L^2(\Omega, \mathbb{R}))$ is said to be a solution of (5.7) if it satisfies

$$\int_0^T \int_\Omega f(t, x) (\operatorname{div}_x (v(t, x) \eta(t, x)) + \frac{\partial}{\partial t} \eta(t, x)) \, dx \, dt + \int_\Omega f_0(x) \eta(0, x) \, dx = 0$$

for all $\eta \in C_c^\infty([0, T] \times \Omega)$. The next theorem is an existence and stability result, see [92, Cors. II.1 and II.2, Thm. II.5].

Theorem 5.4 (Existence & Stability). *For every $v \in A$ there exists a unique weak solution $f \in C^0([0, T], L^2(\Omega, \mathbb{R}))$ of (5.7). If $v_i \in A$ converges to $v \in A$ in the norm of $L^2([0, T] \times \Omega, \mathbb{R}^n)$, then the corresponding sequence of weak solutions $f_i \in C^0([0, T], L^2(\Omega, \mathbb{R}))$ converges to f in $C^0([0, T], L^2(\Omega, \mathbb{R}))$.*

Proof. The existence and uniqueness of weak solutions follows from [92, Corrs. II.1 and II.2]. Note that these solutions are also renormalized due to [92, Thm. II.3].

We recast the second part of the theorem such that it has the exact form of [92, Thm. II.5]. First, note that both the velocity fields and the initial condition can be extended to \mathbb{R}^n by zero outside of Ω due to boundary condition of A . Due to the conditions on v , the weak formulation is equivalent to the one for the extension in the \mathbb{R}^n setting. The uniform boundedness condition on f_i is satisfied since Ω is bounded. \square

Corollary 5.5. *Let $v_i \rightharpoonup v \in A$ with the inner product of $H^1([0, T] \times \Omega)^n$. Then f_i converges to f in $C^0([0, T], L^2(\Omega, \mathbb{R}))$.*

Proof. Combine Theorem 5.4 with the compact embedding of $H^1([0, T] \times \Omega)^n$ into $L^2([0, T] \times \Omega)^n$, i.e., the Rellich embedding theorem [4, A6.4]. \square

Remark 5.6. *Note that the same arguments can be used if we use higher spatial regularity, such as H^2 , in this section. From a numerical point of view, the bound on the divergence is always satisfied for C large enough if we use linear interpolation for the velocities on a fixed grid. Here we use that all norms are equivalent on finite-dimensional spaces.*

5.3 Regularizing properties of template-based image reconstruction

In this section, we prove regularizing properties of template-based reconstruction as defined in (5.2). Recall that the problem reads

$$\begin{aligned} \min_{v \in V} D(K \circ f_v(T, \cdot), g) + \gamma \|v\|_V^2, \\ \text{s.t. } C(v), \end{aligned}$$

where $C(v)$ is the Cauchy problem with either the transport or the continuity equation. The admissible set \mathcal{V} is chosen such that the regularity requirements stated in the previous section are satisfied. For the following considerations we require these assumptions on K and D :

1. The operator K is continuous, $D(\cdot, g)$ is lower semi-continuous for each $g \in Y$, and $D(g, \cdot)$ is continuous for each $g \in Y$.
2. If f_n, g_n are two convergent sequences with limits f and g , respectively, then D must satisfy $\liminf_{n \rightarrow \infty} D(f_n, g) \leq \liminf_{n \rightarrow \infty} D(f_n, g_n)$.
3. If $D(f, g) = 0$, then $f = g$.

In particular, the requirements on D are satisfied if D is a metric. The obtained results are along the lines of [64] but are adapted to our setting and notation. For simplicity we stick to the notation of the continuity equation, but want to mention that the same derivations hold for the transport equation with coefficients in the set A . First, we prove that a minimizer of the problem exists.

Proposition 5.7 (Existence). *For every $f_0 \in L^2(\Omega, \mathbb{R})$, the functional $J_{\gamma, g}$ defined in (5.3) has a minimizer.*

Proof. The idea of the proof is to construct a minimizing sequence that is weakly convergent and then use that the functional is weakly lower semi-continuous. Let us consider a sequence v_n such that $J_{\gamma, g}(v_n)$ converges to $\inf_v J_{\gamma, g}(v)$. By construction of the functional, v_n is bounded in $L^2([0, T], \mathcal{V})$ and hence there exists

a subsequence, also denoted with v_n , such that $v_n \rightharpoonup v_\infty$. By Proposition 5.3, there exists a subsequence, also denoted with v_n , such that $f_{v_n}(T, \cdot) \rightarrow f_{v_\infty}(T, \cdot)$ in $L^2(\Omega, \mathbb{R})$. With this at hand, we are able to prove weak lower semi-continuity of the data term. Indeed, as K is continuous, from $f_{v_n}(T, \cdot) \rightarrow f_{v_\infty}(T, \cdot)$ we get $K \circ f_{v_n}(T, \cdot) \rightarrow K \circ f_{v_\infty}(T, \cdot)$. Since $D(\cdot, g)$ is lower semi-continuous, we obtain that $D(K \circ f_{v_\infty}(T, \cdot), g) \leq \liminf_{n \rightarrow \infty} D(K \circ f_{v_n}(T, \cdot), g)$. This concludes the proof as $J_{\gamma, g}$ is weakly lower semi-continuous, and hence $J_{\gamma, g}(v_\infty) \leq \inf_v J_{\gamma, g}(v)$. \square

Next, we state a stability result.

Proposition 5.8 (Stability). *Let $f_0 \in L^2(\Omega, \mathbb{R})$ and $\gamma > 0$. Let g_n be a sequence in Y converging to $g \in Y$. For each n , we choose v_n as minimizer of J_{γ, g_n} . Then, there exists a subsequence of v_n that converges weakly towards a minimizer v of $J_{\gamma, g}$.*

Proof. By the properties of D it holds, for every n , that

$$\|v_n\|_V^2 \leq \frac{1}{\gamma} J_{\gamma, g_n}(v_n) \leq \frac{1}{\gamma} J_{\gamma, g_n}(0) = \frac{1}{\gamma} D(K \circ f_0, g_n) \rightarrow \frac{1}{\gamma} D(K \circ f_0, g) < \infty.$$

Hence, v_n is bounded in $L^2([0, T], \mathcal{V})$ and there exists a subsequence, also denoted with v_n , such that $v_n \rightharpoonup v$. Further, it holds $\gamma \|v\|_V^2 \leq \gamma \liminf_{n \rightarrow \infty} \|v_n\|_V^2$.

By passing to a subsequence and by using Proposition 5.3, we deduce that $f_{v_n}(T, \cdot) \rightarrow f_v(T, \cdot)$. Together with the convergence of g_n and the convergence property of D this implies

$$D(K \circ f_v(T, \cdot), g) \leq \liminf_{n \rightarrow \infty} D(K \circ f_{v_n}(T, \cdot), g) \leq \liminf_{n \rightarrow \infty} D(K \circ f_{v_n}(T, \cdot), g_n).$$

Thus, for any \tilde{v} , it holds that

$$J_{\gamma, g}(v) \leq \liminf_{n \rightarrow \infty} \gamma \|v_n\|_V^2 + D(K \circ f_{v_n}(T, \cdot), g_n) = \liminf_{n \rightarrow \infty} J_{\gamma, g_n}(v_n) \leq \liminf_{n \rightarrow \infty} J_{\gamma, g_n}(\tilde{v}),$$

since v_n minimizes J_{γ, g_n} . Then, as $J_{\gamma, g_n}(\tilde{v})$ converges to $J_{\gamma, g}(\tilde{v})$ by the assumptions on D , we deduce $J_{\gamma, g}(v) \leq J_{\gamma, g}(\tilde{v})$ and hence that v minimizes $J_{\gamma, g}$. \square

Finally, we state a convergence result for the method.

Proposition 5.9 (Convergence). *Let $f_0 \in L^2(\Omega, \mathbb{R})$ and $g \in Y$, and suppose that there exists $\hat{v} \in L^2([0, T], \mathcal{V})$ such that $K \circ f_{\hat{v}}(T, \cdot) = g$. Further, assume that $\gamma: \mathbb{R}_{>0} \mapsto \mathbb{R}_{>0}$ satisfies $\gamma(\delta) \rightarrow 0$ and $\delta/\gamma(\delta) \rightarrow 0$ as $\delta \rightarrow 0$. Now, let δ_n be a sequence of positive numbers converging to 0 and assume that g_n is a data sequence satisfying $D(g, g_n) \leq \delta_n$ for each $n \in \mathbb{N}$. Let v_n be a minimizer of J_{γ_n, g_n} , where $\gamma_n = \gamma(\delta_n)$. Then there exists a subsequence of v_n that converges weakly towards an element v such that $K \circ f_v(T, \cdot) = g$.*

Proof. For every n , it holds that

$$\|v_n\|_V^2 \leq \frac{1}{\gamma_n} J_{\gamma_n, g_n}(v_n) \leq \frac{1}{\gamma_n} J_{\gamma_n, g_n}(\hat{v}) = \frac{1}{\gamma_n} (D(g, g_n) + \gamma_n \|\hat{v}\|_V^2) \leq \frac{\delta_n}{\gamma_n} + \|\hat{v}\|_V^2.$$

From the requirements on γ and δ we deduce that v_n is bounded in $L^2([0, T], \mathcal{V})$ and then that up to an extraction, v_n converges weakly to some v in $L^2([0, T], \mathcal{V})$.

Further, it holds $D(K \circ f_v(T, \cdot), g) \leq \liminf_{n \rightarrow \infty} D(K \circ f_{v_n}(T, \cdot), g_n)$ with the same arguments as in the previous proposition. Finally, for every n , it holds that

$$D(K \circ f_{v_n}(T, \cdot), g_n) \leq J_{\gamma_n, g_n}(v_n) \leq J_{\gamma_n, g_n}(\hat{v}) = D(g, g_n) + \gamma_n \|\hat{v}\|_V^2,$$

where the two rightmost terms both converge to zero. Thus, $K \circ f_v(T, \cdot) = g$ by the assumptions on D . \square

We conclude with a remark on data discrepancy functionals that satisfy the conditions and will be used in our numerical experiments in Section 5.5.

Remark 5.10. *We now assume that the data space Y is a real Hilbert space. Clearly, the conditions are satisfied if $D_{\text{SSD}}(f, g) = \|f - g\|_Y^2$. We will only check the convergence condition. It holds*

$$\liminf_{n \rightarrow \infty} \|f_n - g\|_Y^2 = \liminf_{n \rightarrow \infty} \|f_n - g_n\|_Y^2 + 2\langle f_n - g_n, g_n - g \rangle + \|g - g_n\|_Y^2,$$

where the last two terms converge to zero since convergent sequences are bounded.

Another function that satisfies the conditions is $D_{\text{NCC}}: Y \setminus \{0\} \times Y \setminus \{0\} \rightarrow [0, 1]$ with

$$D_{\text{NCC}}(f, g) = 1 - \frac{\langle f, g \rangle^2}{\|f\|_Y^2 \|g\|_Y^2},$$

which is based on the NCC. First, note that $\tilde{D}(\cdot, g) = \langle \cdot, g \rangle^2 / \|g\|_Y^2$ and the function $\|\cdot\|_Y^{-2}$ are continuous. Thus, we get that $D_{\text{NCC}}(\cdot, g)$ is continuous. By symmetry, this also holds for $D_{\text{NCC}}(g, \cdot)$. It remains to check the convergence property

$$\begin{aligned} \lim_{n \rightarrow \infty} 1 - D_{\text{NCC}}(f_n, g) &= \lim_{n \rightarrow \infty} \frac{(\langle f_n, g - g_n \rangle + \langle f_n, g_n \rangle)^2}{\|f_n\|_Y^2 \|g\|_Y^2} = \lim_{n \rightarrow \infty} \frac{\langle f_n, g_n \rangle^2}{\|f_n\|_Y^2 \|g\|_Y^2} \\ &= \lim_{n \rightarrow \infty} \frac{\langle f_n, g_n \rangle^2}{\|f_n\|_Y^2 \|g_n\|_Y^2} = \lim_{n \rightarrow \infty} 1 - D_{\text{NCC}}(f_n, g_n). \end{aligned}$$

From this we conclude $\liminf_{n \rightarrow \infty} D_{\text{NCC}}(f_n, g) = \liminf_{n \rightarrow \infty} D_{\text{NCC}}(f_n, g_n)$. Unfortunately, $D_{\text{NCC}}(f, g) = 0$ only implies $f = cg$, with $c \in \mathbb{R}$.

5.4 Numerical solution

The focus of this section is to approximately solve the problem (5.2). Our approach is based on the Lagrangian methods developed in [193] and the inexact multilevel Gauss–Newton method used in [205]. Both methods and their necessary modifications are briefly outlined here.

As customary in PDE-constrained optimization [88, Chap. 3], we eliminate the state equation by defining a *control-to-state* operator, which parameterizes the final state $f_v(T, \cdot)$ in terms of the unknown velocities v . With a slight abuse of notation, we define this solution map as

$$\begin{aligned} S: V &\rightarrow L^2(\Omega, \mathbb{R}), \\ v &\mapsto f_v(T, \cdot) =: f(v). \end{aligned} \tag{5.8}$$

Here, f_v denotes the unique solution to either the transport or the continuity equation, as defined in Section 5.2. As a result, we obtain the reduced form of (5.2):

$$\min_{v \in V} D(K \circ f(v), g) + \gamma R(v). \quad (5.9)$$

Here, $R: V \rightarrow \mathbb{R}_{\geq 0}$ is a regularization functional that can be written as

$$R(v) = \frac{1}{2} \int_0^T \int_{\Omega} \|Bv(t, x)\|^2 dx dt \quad (5.10)$$

with B denoting a linear (vectorial) differential operator.

In this work, we consider the operators $B = \nabla_x$ and $B = \Delta_x$, which are also used in [193]. We refer to the resulting functionals R as *diffusion* and *curvature regularization* functionals, respectively. Note that B can as well be chosen to incorporate derivatives with respect to time.

Among the operators above, we also consider a regularization functional that resembles the norm of the space $V = L^2([0, T], H_0^3(\Omega, \mathbb{R}^n))$. This particular choice is motivated by the fact that, for $n = \{2, 3\}$, the space $H_0^3(\Omega, \mathbb{R}^n)$ can be continuously embedded in $C_0^{1,\alpha}(\Omega, \mathbb{R}^n)$, for some $\alpha > 0$, so that the results in Section 5.2 hold. The norm of V is given by

$$\|v\|_V^2 = \frac{1}{2} \int_0^T \|v(t, \cdot)\|_{L^2(\Omega, \mathbb{R}^n)}^2 dt + \frac{1}{2} \int_0^T |v(t, \cdot)|_{H^3(\Omega, \mathbb{R}^n)}^2 dt. \quad (5.11)$$

Here, $|\cdot|_{H^k(\Omega, \mathbb{R}^n)}$ denotes the usual H^k -seminorm including only the highest-order partial derivatives. By the Gagliardo–Nirenberg inequality, (5.11) is equivalent to the usual norm of $L^2([0, T], H_0^3(\Omega, \mathbb{R}^n))$. To simplify numerical optimization, we omit the requirement that v is compactly supported in Ω and minimize over $L^2([0, T], H^3(\Omega, \mathbb{R}^n))$.

In order to solve problem (5.9), we follow a discretize-then-optimize strategy. Without loss of generality, we assume that the domain is $\Omega = (0, 1)^n$. We partition it into a regular grid consisting of m^n equally sized cells of edge length $h_X = 1/m$ in every coordinate direction.

The template image $f_0 \in L^2(\Omega, \mathbb{R})$ is assumed to be sampled at cell-centered locations $\mathbf{x}_c \in \mathbb{R}^{m^n}$, giving rise to its discrete version $\mathbf{f}_0(\mathbf{x}_c) \in \mathbb{R}^{m^n}$. The template image is interpolated on the cell-centered grid by means of cubic B-spline interpolation as outlined in [205, Chap. 3.4].

Similarly, the time domain is assumed to be $[0, 1]$ and is partitioned into m_t equally sized cells of length $h_t = 1/m_t$. We assume that the unknown velocities $v: [0, 1] \times \Omega \rightarrow \mathbb{R}^n$ are sampled at cell-centered locations in space as well as at cell-centered locations in time, leading to a vector of unknowns $\mathbf{v} \in \mathbb{R}^N$, where $N = (m_t + 1) \cdot n \cdot m^n$ is the total number of unknowns of the finite-dimensional minimization problem.

Lagrangian solver In order to compute the solution map $f(v)$ numerically, i.e., to solve the hyperbolic PDEs (5.4) and (5.7), the Lagrangian solver in [193] follows

a two-step approach. First, given a vector $\mathbf{v} \in \mathbb{R}^N$ of velocities, the ODE (5.6) is solved approximately using a fourth-order Runge–Kutta (RK4) method with N_t equally spaced time steps of size Δt . For simplicity, we follow the presentation in [193] based on an explicit first-order Euler method and refer to [193, Sec. 3.1] for the full details.

Given initial points $\mathbf{x} \in \mathbb{R}^{m^n}$ and velocities $\mathbf{v} \in \mathbb{R}^N$, an approximate solution $\mathbf{X}_{\mathbf{v}}: [0, 1]^2 \times \mathbb{R}^{m^n} \rightarrow \mathbb{R}^{m^n}$ is given by

$$\mathbf{X}_{\mathbf{v}}(0, t_{k+1}, \mathbf{x}) = \mathbf{X}_{\mathbf{v}}(0, t_k, \mathbf{x}) + \Delta t \mathbf{I}(\mathbf{v}, t_k, \mathbf{X}_{\mathbf{v}}(0, t_k, \mathbf{x})), \quad (5.12)$$

for all $k = 0, 1, \dots, N_t - 1$, with initial condition $\mathbf{X}_{\mathbf{v}}(0, 0, \mathbf{x}) = \mathbf{x}$. Here, we need the componentwise interpolation $\mathbf{I}(\mathbf{v}, t_k, \mathbf{X}_{\mathbf{v}}(0, t_k, \mathbf{x}))$ of \mathbf{v} at time $t_k = k\Delta t$ and at the points $\mathbf{X}_{\mathbf{v}}(0, t_k, \mathbf{x})$. Note that, since the characteristic curves for both PDEs coincide, this step is identical regardless of which PDE we impose.

The second step computes approximate intensities of the final state $f_v(1, \cdot)$. This step depends on the particular PDE. For the transport equation, in order to compute the intensities at the grid points \mathbf{x}_c , we follow characteristic curves backwards in time, which is achieved by setting $\Delta t = -1/N_t$ in (5.12). The deformed template is then given by

$$\mathbf{f}(\mathbf{v}) = \mathbf{f}_0 \circ \mathbf{X}_{\mathbf{v}}(1, 0, \mathbf{x}_c), \quad (5.13)$$

where $\mathbf{f}_0 \in \mathbb{R}^{m^n}$ is the interpolation of the discrete template image.

For the continuity equation, a *particle-in-cell* (PIC) method is proposed by [193], see [67] for details. The density of particles that are initially located at grid points \mathbf{x}_c is represented by a linear combination of basis functions, which are then shifted by following the characteristics computed in the first step. To determine the final density at grid points, exact integration over the grid cells is performed. By setting $\Delta t = 1/N_t$ in (5.12), the transformed template can be computed as

$$\mathbf{f}(\mathbf{v}) = \mathbf{F} \circ \mathbf{X}_{\mathbf{v}}(0, 1, \mathbf{x}_c) \mathbf{f}_0(\mathbf{x}_c), \quad (5.14)$$

where $\mathbf{F} \in \mathbb{R}^{N \times N}$ is the push-forward matrix that computes the integrals over the shifted basis functions. See [193, Sec. 3.1] for its detailed specification using linear, compactly supported basis functions. By design, the method is mass-preserving at the discrete level.

Numerical optimization Let us denote by $\mathbf{K}: \mathbb{R}^N \rightarrow \mathbb{R}^M$, $M \in \mathbb{N}$, a finite-dimensional, Fréchet differentiable approximation of the (not necessarily linear) operator $K: L^2(\Omega, \mathbb{R}) \rightarrow Y$. With the application to CT in mind, we outline a discretization of (5.9) suitable for the n -dimensional Radon transform, which maps a function on \mathbb{R}^n into the set of its integrals over the hyperplanes in \mathbb{R}^n [208, Chap. 2].

An element $K \circ f(v) \in Y$ is a function on the unit cylinder $\mathbb{S}^{n-1} \times \mathbb{R}$ of \mathbb{R}^{n+1} , where \mathbb{S}^{n-1} is the $(n-1)$ -dimensional unit sphere. We discretize this unit cylinder as follows. First, we sample $p \in \mathbb{N}$ directions from \mathbb{S}^{n-1} . When $n = 2$, as it is the case in our experiments in Section 5.5, directions are parameterized by angles from

the interval $[0, 180]$ degrees. For simplicity, we say (slightly imprecise) that we take one measurement in each direction. Second, similarly to the sampling of Ω , we use an interval $(0, 1)$ instead of \mathbb{R} and partition it into q equally sized cells of length $h_Y = 1/q$. Depending on n and the diameter of Ω , the interval length may require adjustment. Each measurement i is then sampled at cell-centered points $\mathbf{y}_c \in \mathbb{R}^q$ and denoted by $\mathbf{g}_i(\mathbf{y}_c) \in \mathbb{R}^q$. All measurements are then concatenated into a vector $\mathbf{g} := \mathbf{g}(\mathbf{y}_c) \in \mathbb{R}^M$, where $M = p \cdot q$.

The finite-dimensional optimization problem in abstract form is then given by

$$\min_{\mathbf{v} \in \mathbb{R}^N} \{J_{\gamma, \mathbf{g}}(\mathbf{v}) := D(\mathbf{K} \circ \mathbf{f}(\mathbf{v}), \mathbf{g}) + \gamma R(\mathbf{v})\}, \quad (5.15)$$

where D and R are chosen to be discretizations of a distance and of (5.10), respectively.

In further consequence, we approximate integrals using a midpoint quadrature rule. As we are mainly interested in the setting where only few directions are given, we disregard integration over the unit sphere. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^M$, the corresponding approximations of the distance based on sum-of-squared-differences and the NCC-based distance are then

$$D_{\text{SSD}}(\mathbf{x}, \mathbf{y}) \approx \frac{h_Y}{2} (\mathbf{x} - \mathbf{y})^\top (\mathbf{x} - \mathbf{y}) \quad \text{and} \quad D_{\text{NCC}}(\mathbf{x}, \mathbf{y}) \approx 1 - \frac{(\mathbf{x}^\top \mathbf{y})^2}{\|\mathbf{x}\|^2 \|\mathbf{y}\|^2}, \quad (5.16)$$

respectively. See [205, Chaps. 6.2 and 7.2] for details. Note that, due to cancellation, no (spatial) discretization parameter occurs in the approximation of the NCC above.

Moreover, we approximate the regularization functional in (5.10) with

$$R(\mathbf{v}) \approx \frac{h_t h_x^2}{2} \mathbf{v}^\top \mathbf{B}^\top \mathbf{B} \mathbf{v}, \quad (5.17)$$

where $\mathbf{B} \in \mathbb{R}^{N \times N}$ is a finite-difference discretization of the differential operator in (5.10), analogous to [204, Chap. 8.5]. In our implementation, we use zero Neumann boundary conditions and pad the spatial domain to mitigate boundary effects arising from the discretization of the operator.

In order to apply (inexact) Gauss–Newton optimization to problem (5.15), we require first- and (approximate) second-order derivatives of $J_{\gamma, \mathbf{g}}(\mathbf{v})$. By application of the chain rule, we obtain

$$\frac{\partial}{\partial \mathbf{v}} J_{\gamma, \mathbf{g}}(\mathbf{v}) = \frac{\partial}{\partial \mathbf{v}} \mathbf{f}(\mathbf{v})^\top \frac{\partial}{\partial \mathbf{f}} \mathbf{K} \circ \mathbf{f}(\mathbf{v})^\top \frac{\partial}{\partial \mathbf{x}} D(\mathbf{K} \circ \mathbf{f}(\mathbf{v}), \mathbf{g}) + \gamma \frac{\partial}{\partial \mathbf{v}} R(\mathbf{v}),$$

where $\partial \mathbf{K} / \partial \mathbf{f}$ is the Fréchet derivative of \mathbf{K} and $\partial \mathbf{f}(\mathbf{v}) / \partial \mathbf{v}$ is the derivative of the solution map (5.8) with respect to the velocities, which is given below.

The partial derivatives of the distance functions (5.16) with respect to its first argument are given by

$$\frac{\partial}{\partial \mathbf{x}} D_{\text{SSD}}(\mathbf{x}, \mathbf{y}) = h_Y (\mathbf{x} - \mathbf{y}) \quad \text{and} \quad \frac{\partial^2}{\partial \mathbf{x}^2} D_{\text{SSD}}(\mathbf{x}, \mathbf{y}) = h_Y \mathbf{I}_N, \quad (5.18)$$

where $\mathbf{I}_N \in \mathbb{R}^{N \times N}$ is the identity matrix of size N , and

$$\frac{\partial}{\partial \mathbf{x}} D_{\text{NCC}}(\mathbf{x}, \mathbf{y}) = -2 \frac{(\mathbf{x}^\top \mathbf{y}) \mathbf{y}}{\|\mathbf{x}\|^2 \|\mathbf{y}\|^2} + 2 \frac{(\mathbf{x}^\top \mathbf{y})^2 \mathbf{x}}{\|\mathbf{x}\|^4 \|\mathbf{y}\|^2},$$

respectively. Moreover, the derivatives of (5.17) are given by

$$\frac{\partial}{\partial \mathbf{v}} R(\mathbf{v}) = h_t h_X^n \mathbf{B}^\top \mathbf{B} \mathbf{v} \quad \text{and} \quad \frac{\partial^2}{\partial \mathbf{v}^2} R(\mathbf{v}) = h_t h_X^n \mathbf{B}^\top \mathbf{B}.$$

In order to obtain an efficient iterative second-order method for solving (5.15), we require an approximation of the Hessian $\mathbf{H} \in \mathbb{R}^{N \times N}$ that balances the following tradeoff. Ideally, it is reasonably efficient to compute, consumes limited memory (sparsity is desired), and has sufficient structure so that preconditioning can be used. However, each iteration of the Gauss–Newton method should also provide a suitable descent direction. For these reasons, we approximate the Hessian by

$$\begin{aligned} \mathbf{H}(\mathbf{v}) = \frac{\partial^2}{\partial \mathbf{v}^2} J_{\gamma, \mathbf{g}}(\mathbf{v}) &\approx \frac{\partial}{\partial \mathbf{v}} \mathbf{f}(\mathbf{v})^\top \frac{\partial}{\partial \mathbf{f}} \mathbf{K} \circ \mathbf{f}(\mathbf{v})^\top \frac{\partial^2}{\partial \mathbf{x}^2} D(\mathbf{K} \circ \mathbf{f}(\mathbf{v}), \mathbf{g}) \frac{\partial}{\partial \mathbf{f}} \mathbf{K} \circ \mathbf{f}(\mathbf{v}) \frac{\partial}{\partial \mathbf{v}} \mathbf{f}(\mathbf{v}) \\ &\quad + \gamma h_t h_X^n \mathbf{B}^\top \mathbf{B} + \epsilon \mathbf{I}_N, \end{aligned}$$

where $\epsilon > 0$ ensures positive semi-definiteness. For simplicity, the term involving $\partial^2 \mathbf{f}(\mathbf{v}) / \partial \mathbf{v}^2$ is omitted and, regardless of the chosen distance, we use the second derivative in (5.18) as an approximation of $\partial^2 D(\mathbf{x}, \mathbf{y}) / \partial \mathbf{x}^2$. In our numerical experiments we found that this choice works well for the problem considered in Section 5.5.

It remains to discuss the derivative of the solution map. For the transport equation, the application of the chain rule to (5.13) yields

$$\frac{\partial}{\partial \mathbf{v}} \mathbf{f}(\mathbf{v}) = \nabla_x \mathbf{f}_0(\mathbf{X}_\mathbf{v}(1, 0, \mathbf{x}_c)) \frac{\partial}{\partial \mathbf{v}} \mathbf{X}_\mathbf{v}(1, 0, \mathbf{x}_c),$$

where $\nabla_x \mathbf{f}_0$ denotes the gradient of the interpolation of the template image and $\partial \mathbf{X}_\mathbf{v} / \partial \mathbf{v}$ is the derivative of the endpoints of the characteristic curves with respect to the velocities, see below. Similarly, for the solution map (5.14) that corresponds to the continuity equation, we obtain

$$\frac{\partial}{\partial \mathbf{v}} \mathbf{f}(\mathbf{v}) = \frac{\partial}{\partial \mathbf{X}_\mathbf{v}} (\mathbf{F} \circ \mathbf{X}_\mathbf{v}(0, 1, \mathbf{x}_c) \mathbf{f}_0(\mathbf{x}_c)) \frac{\partial}{\partial \mathbf{v}} \mathbf{X}_\mathbf{v}(0, 1, \mathbf{x}_c).$$

Here, $\partial \mathbf{F} / \partial \mathbf{X}_\mathbf{v}$ is the derivative of the push-forward matrix with respect to the endpoints of the characteristics, again see [193, Sec. 3.1].

If explicit time stepping methods are used to solve the ODE (5.6), the partial derivative $\partial \mathbf{X}_\mathbf{v} / \partial \mathbf{v}$ can be computed recursively. For example, for the forward Euler approach in (5.12) it is given by

$$\begin{aligned} \frac{\partial}{\partial \mathbf{v}} \mathbf{X}_\mathbf{v}(0, t_{k+1}, \mathbf{x}_c) &= \frac{\partial}{\partial \mathbf{v}} \mathbf{X}_\mathbf{v}(0, t_k, \mathbf{x}_c) + \Delta t \frac{\partial}{\partial \mathbf{v}} \mathbf{I}(\mathbf{v}, t_k, \mathbf{X}_\mathbf{v}(0, t_k, \mathbf{x}_c)) \\ &\quad + \Delta t \frac{\partial}{\partial \mathbf{X}_\mathbf{v}} \mathbf{I}(\mathbf{v}, t_k, \mathbf{X}_\mathbf{v}(0, t_k, \mathbf{x}_c)) \frac{\partial}{\partial \mathbf{v}} \mathbf{X}_\mathbf{v}(0, t_k, \mathbf{x}_c), \end{aligned}$$

for all $k = 0, 1, \dots, N_t - 1$, with $\partial \mathbf{I} / \partial \mathbf{v}$ and $\partial \mathbf{I} / \partial \mathbf{X}_\mathbf{v}$ being the derivatives of the interpolation schemes with respect to the velocities and with respect to the endpoints of the characteristics, respectively. We refer to [205, Chap. 3.5] for details. The case where characteristics are computed backwards in time can be handled similarly.

In order to solve the finite-dimensional minimization problem (5.15), we apply a inexact Gauss–Newton–Krylov method, which proceeds as follows. Given an initial guess $\mathbf{v}^{(0)} = \mathbf{0}$, we update the velocities in each iteration $i = 0, 1, \dots$ by

$\mathbf{v}^{(i+1)} = \mathbf{v}^{(i)} + \mu \delta \mathbf{v}$ until a termination criterion is satisfied. Here, $\mu \in \mathbb{R}$ denotes a step size that is determined via Armijo line search and $\delta \mathbf{v} \in \mathbb{R}^N$ is the solution to the linear system

$$\mathbf{H}(\mathbf{v}^{(i)}) \delta \mathbf{v} = -\frac{\partial}{\partial \mathbf{v}} J_{\gamma, \mathbf{g}}(\mathbf{v}^{(i)}). \quad (5.19)$$

For details on the stopping criteria and the line search we refer to [205, Chap. 6.3.3]. We solve the system (5.19) approximately by means of a preconditioned conjugate gradient (PCG) method, which can be implemented matrix-free whenever the derivative of \mathbf{K} and its adjoint can be computed matrix-free. See [193, Sec. 3.2] for further details on the preconditioning.

Due to the non-convexity of (5.9) and to speed up computation, we use a multi-level strategy in order to reduce the risk of ending up in a local minimum, see [143]. On each level, we use a subsampled version of the velocities that were computed on the previous, coarser discretization as initial guess.

While standard image registration typically uses resampling of the template and the target image [205, Chap. 3.7], the approach described here requires multilevel versions of the operator \mathbf{K} together with a suitable method for resampling the measurements \mathbf{g} . We stress that, if these are not available, optimization can as well just be performed on the finest discretization level.

In the following, we assume that \mathbf{K} is a discretization of the Radon transform [208], which is a linear operator, and outline a suitable procedure for creating multi-level versions of the operator and the measured data. The former is easily achieved with a computational backend such as Astra [268, 269], which allows to explicitly specify the number of grid cells used to discretize the measurement geometry. For the sake of simplicity, we restrict the presentation here to the case where $n = 2$, i.e., $\Omega \subset \mathbb{R}^2$, and \mathbf{K} is linear.

Let us assume that the number of grid cells used to discretize Ω at the finest level is $m = 2^\ell$, $\ell \in \mathbb{N}$. In our experiments, we set the number of grid cells of the one-dimensional measurement domain $(0, 1)$ at the current level $k \leq \ell$ to $q^{(k)} = 1.5 \cdot 2^{(k)}$ and set the length of each cell to $h_Y^{(k)} = 1/q^{(k)}$. Then, a multilevel representation of each measurement \mathbf{g}_i , $i \leq p$, at cell-centered grid points $\mathbf{y}_j = (j - 1/2)h_Y^{(k-1)}$ is given by

$$\mathbf{g}_i^{(k-1)}(\mathbf{y}_j) := \left(\mathbf{g}_i^{(k)}(\mathbf{y}_j) + \mathbf{g}_i^{(k)}(\mathbf{y}_j + h_Y^{(k)}) \right) / 4,$$

where the denominator arises from averaging over two neighboring grid points and dividing the edge length of the imaging domain Ω in each coordinate direction in half. The approach can be extended to higher dimensions.

5.5 Numerical examples

In our numerical experiments we use the Radon transform [208] as operator. Other choices are possible and, assuming that one has access to a suitable resampling procedure for the measured data, the multilevel strategy can be applied as well. The aim here is to investigate the reconstruction quality with different regularization functionals, distances, and noise levels for both PDE constraints. We show synthetic




Figure 5.1: Synthetic example based on an artificial brain image [141] that has been deformed manually. We generated six Radon transform measurements that correspond to six equally spaced angles from the interval $[0, 60]$ degrees.

examples for the settings $n = 2$ and $n = 3$, and nonsynthetic examples for $n = 2$ using real X-ray tomography data. In the synthetic case, all shown reconstructions are computed from measurements taken from at most 10 directions (angles) sampled from intervals within $[0, 180]$ degrees.

All computations are performed using an Intel Xeon E5-2630 v4 2.2 GHz server equipped with 128 GB RAM and an NVIDIA Quadro P6000 GPU featuring 24 GB of memory. The GPU is only used for computing the Radon transform of 3D volumes.

Before we proceed, we give a brief idea of suitable parameter choices. For the multilevel approach we use in each synthetic example 32×32 pixels at the coarsest level and 128×128 pixels at the finest level, i.e., $\ell = 7$. The size of the reconstructed images in the nonsynthetic examples is 128×128 . Again, three levels are used. In the synthetic 3D example, the reconstructed volume is $32 \times 32 \times 32$ and the coarsest level is $8 \times 8 \times 8$.

We use time dependent velocity fields with only one time step, i.e., $n_t = 1$, since this keeps the computational cost reasonable and suffices for our examples. The characteristics are computed using five Runge–Kutta steps, i.e., $N_t = 5$.

The spatial regularization parameter depends on the chosen regularization functional and the noise level, and is chosen in the order of 10^{-3} , 10^0 , and 10^3 for third-order, curvature, and diffusion regularization, respectively, in the noise-free case and using the NCC-based distance. The temporal regularization parameter is less sensitive and is chosen in the order of 10^2 . Furthermore, the parameter corresponding to the norm of $L^2(\Omega, \mathbb{R}^n)$ in (5.11) is set to 10^{-6} .

In our first example, we investigate different regularization functionals with different noise levels together with the transport equation. The target is 2D Radon transform data based on a digital brain image and the template is a deformed version thereof, see Fig. 5.1. Since we want to focus on the behavior of the regularization functionals, we do not treat the continuity equation here. The data is generated

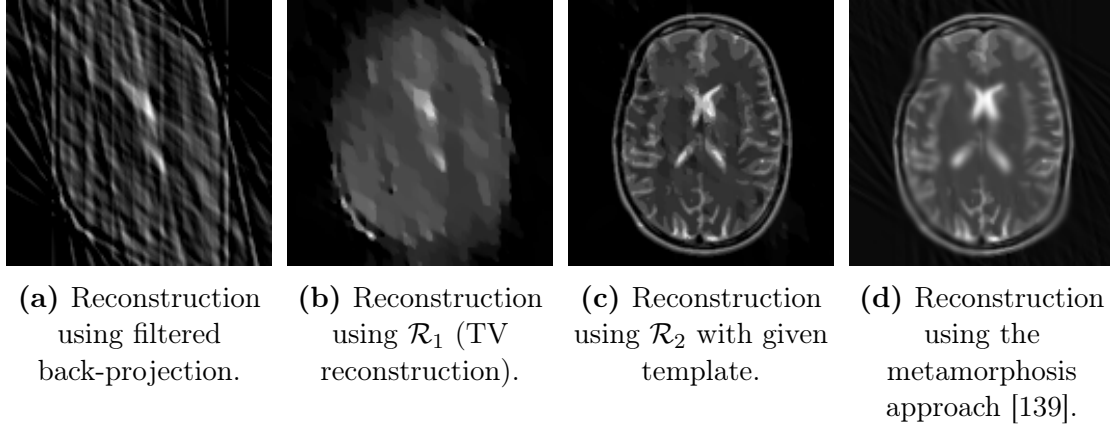


Figure 5.2: Comparison of different reconstruction models applied to an artificial brain image [205] that has been deformed manually. We generated six measurements that correspond to six equally spaced angles from the interval $[0, 60]$ degrees.

using parallel beam tomography with only six equally distributed angles from the interval $[0, 60]$ degrees and corrupted with Gaussian white noise of different levels.

Fig. 5.2 shows results obtained from the generated noise-free measurements using four existing methods. In Fig. 5.2a filtered backprojection is used. In Fig. 5.2b and 5.2c, the following two total variation regularization-based models, see [55],

$$\min_{\mathbf{u}} \|\mathbf{K}\mathbf{u} - \mathbf{g}\|^2 + \gamma \mathcal{R}_i(\mathbf{u}),$$

with $\mathcal{R}_1(\mathbf{u}) := \text{TV}(\mathbf{u})$, $\mathcal{R}_2(\mathbf{u}) := \text{TV}(\mathbf{u} - \mathbf{f}_0)$, and $\gamma > 0$ are used. Here, $\mathcal{R}_2(\mathbf{u})$ incorporates template information. Approximate minimizers of both functionals are computed using the primal-dual hybrid gradient method [57]. For the case of filtered backprojection, the standard MATLAB implementation is used. The results in Fig. 5.2a and 5.2c highlight why more sophisticated methods, such as the proposed template-based approach, are necessary to obtain satisfying reconstructions in this setting, and illustrate the challenges when dealing with very sparse data.

As outlined in Section 5.1, one possibility is the metamorphosis approach [139]. In Fig. 5.2d we show a result obtained with this method using the recommended parameters, i.e., 200 gradient descent steps are performed and the regularization parameters are set to $\gamma = 10^{-5}$ and $\tau = 1$. Observe the change in image intensities compared to Fig. 5.1a and the blur in the heavily deformed regions.

In Fig. 5.3, we show results for the different noise levels and different regularization functionals computed with our approach. All results are obtained using the NCC-based distance. As expected, the quality of the reconstruction gets worse for higher noise levels and, consequentially, larger regularization parameters are necessary. Since data is acquired from only six directions, the influence of the noise is very strong. Especially for the diffusive regularization we need to choose large regularization parameters for higher noise levels, see Fig. 5.3c. Since diffusion corresponds to first-order regularization, it is much easier to reconstruct the noise with “rough” deformations. Overall, we found that second- and third-order regular-

ization performed similar when appropriate regularization parameters are chosen. Even though some theoretical results only hold for higher-order regularity, second-order regularization seems sufficient for our use case. The computation time for the results in Fig. 5.3 is between 200 and 700 seconds.

In the second example, see Fig. 5.4, we compare the behavior of the SSD and the NCC-based distance. The example consists of two different hands that, in addition, are rotated relative to each other. Here, the deformation is much larger than in the previous example, but still fairly regular. The data is generated similarly to the previous example, but with only five angles from the interval $[0, 75]$ degrees. Note also that the intensities of the template and target image are different (roughly by a factor of two). First, we discuss the transport equation. The intensity difference is a serious issue if we use the SSD distance, as we can see in Fig. 5.4e. The hand is deformed into a smaller version in order to compensate the differences. If we use the NCC-based distance instead, which can deal with such discrepancies, the result is much better from a visual point of view. The shapes are well-aligned. The resulting SSIM value is still low, which is not surprising as SSIM is not invariant with respect to intensity differences between perfectly aligned images. However, neither of the two approaches is able to remove or create any of the additional (noise) structures in the images. For the combination SSD with continuity equation, no satisfactory results could be obtained. Since no change of intensity is possible by changing the size of the hand, part of it is moved outside of the image. This behavior could potentially be corrected if other boundary conditions are used in the implementation. Therefore, we do not provide an example image for this case. Using the NCC-based distance, the results look similar as for the transport equation with slightly worse SSIM value. These results suggest that the NCC-based distance is a more robust choice that avoids unnatural deformations, which would be required in the case of SSD to compensate intensity differences. In this example, the computation times are between 50 and 325 seconds.

In the next example, see Fig. 5.5, we compare the continuity equation with the transport equation as constraint together with the NCC-based distance. The continuity equation allows for limited change of mass along the deformation path. Since the intensity change scales with the determinant of the Jacobian, bigger changes are only possible if areas are compressed or extended a lot. In the presented example this occurs only to a mild extent. For this example, the continuity equation and the transport equation yield visually similar results with minor differences in the SSIM value. As in the previous examples, higher-order regularization is beneficial and artifacts occur for the diffusion regularization. The computation time amounts to roughly 64 to 360 seconds in this example.

In Fig. 5.6, we created an artificial pair of images consisting of a disk to show the possibilities of intensity changes when using the continuity equation as a constraint. Both template and unknown image are constructed so that their total mass is equal. The measurements are generated as before using only five angles uniformly distributed in the interval $[0, 90]$ degrees. Furthermore, we use curvature regularization. For the transport equation we observe that the shape is matched, but the intensity is not correct, see Fig. 5.6d. If we use the continuity equation instead,

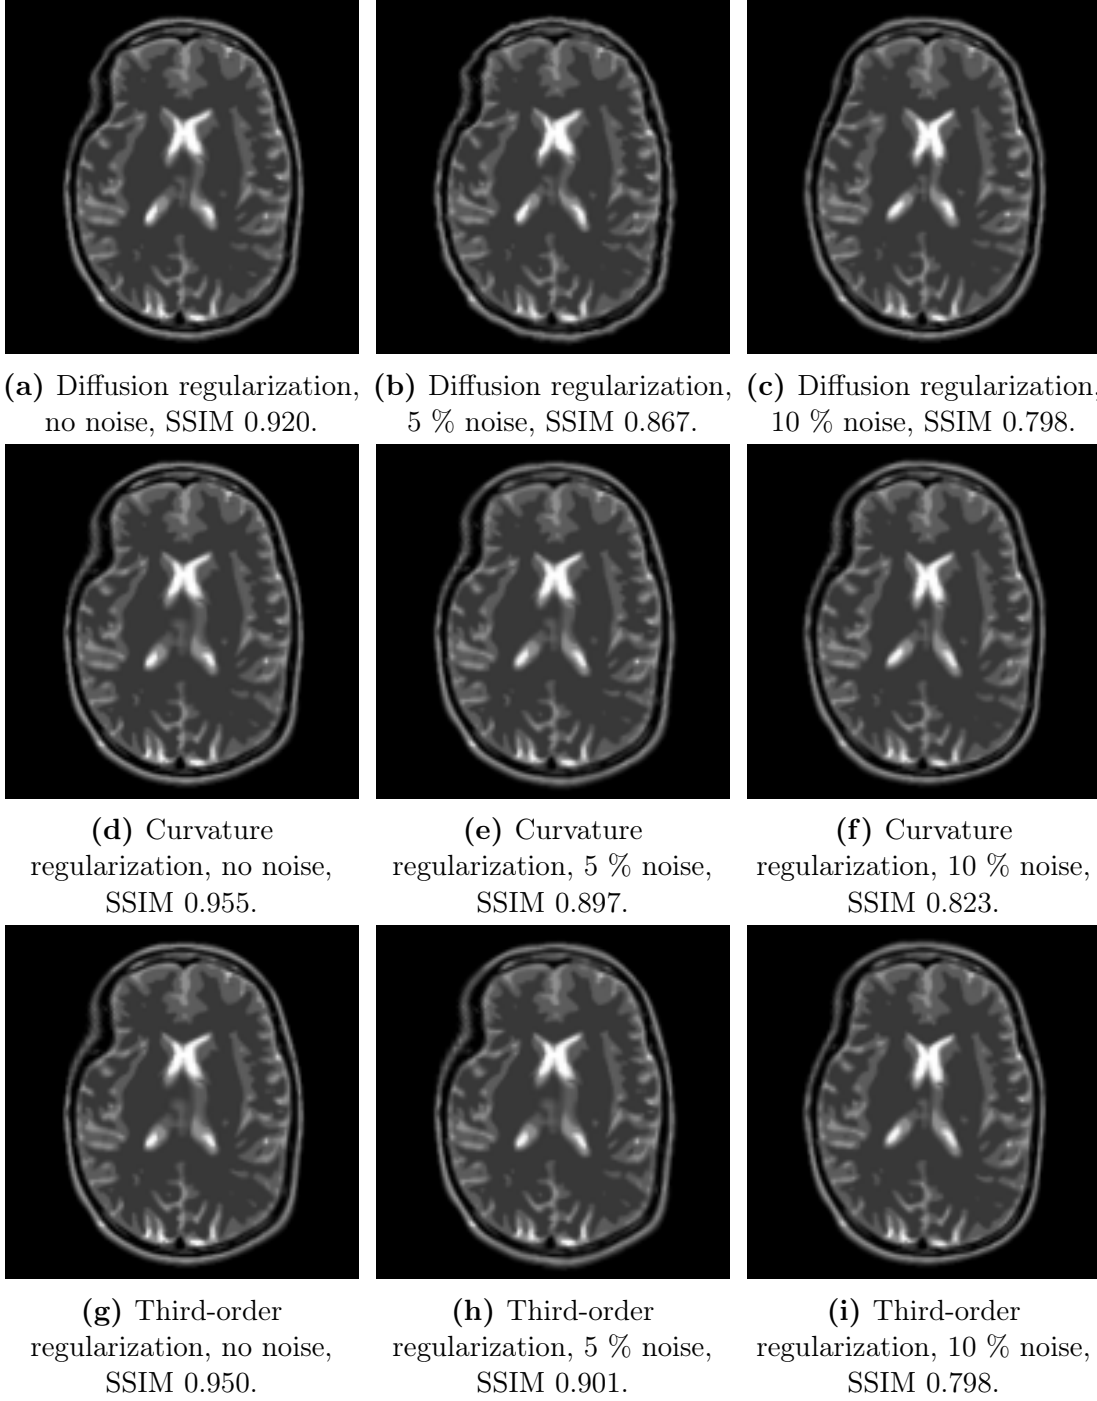


Figure 5.3: Reconstructions for the artificial brain image in Fig. 5.1 using our method and different regularization functionals. Note that only six measurements were used. The measured data was corrupted with noise of different levels.




Figure 5.4: Reconstructions of manually deformed Hand [205] images with different image intensity levels using our method. We generated five measurements that correspond to five equally spaced angles from the interval $[0, 75]$ degrees and added five percent noise.

intensity changes are possible, which can be observed in Fig. 5.6e. The computation time for the two results are 90 and 500 seconds, respectively.

In order to demonstrate the practicality of our method, we compute results from nonsynthetic X-ray tomography data [50, 146], which are available online.^{3,4} See Fig. 5.7 for these two examples ('lotus' and 'walnut'). The template is generated by applying filtered backprojection to the full measurements and by subsequently deforming it. Then, this deformed template is used in our method to compute a reconstruction from only few measurement directions. The computation times are roughly 80 and 600 seconds in these examples. In both nonsynthetic examples, the use of the NCC-based distance proved crucial and no satisfactory result could be obtained using SSD.

In Fig. 5.8, we demonstrate that our framework is also capable of reconstructing

³<https://doi.org/10.5281/zenodo.1254204>

⁴<https://doi.org/10.5281/zenodo.1254206>

3D volumes. Here, we use the SSD distance together with curvature regularization and the transport equation. We apply the 3D Radon transform to obtain ten measurements from angles within $[0, 180]$. The total computation time is roughly 800 seconds.

All in all, our results demonstrate that, given a suitable template image, very reasonable reconstructions can efficiently be obtained from only a few measurements, even in the presence of noise. Moreover, our examples show that the NCC-based distance adds robustness to the approach with regard to discrepancies in the image intensities.

5.6 Conclusions

Overall, our numerical examples show that our implementation yields good results, as long as the deformation between template and target is fairly regular. By using the NCC-based distance, robustness with respect to intensity differences between the template and the target image can be achieved. As already mentioned in the introduction, we do not follow the metamorphosis approach, since there is too much flexibility in the model and the source term likely reproduces noise and artifacts if the data is too limited. It is left for further research to investigate possible adaptations of the model that allow for the appearance of new objects or structures in the reconstruction without reproducing noise or artifacts. Possibly, the results of our method can be used as better template for other algorithms that require template information. Finally, note that due to the great flexibility of the FAIR library, it is also possible to use a great variety of regularization functionals for the velocities and other distances, see [205, Chaps. 7 and 8]. Additionally, our implementation is not necessarily restricted to the Radon transform and essentially every (continuous) operator can be used. The multilevel approach can be applied as long as a meaningful resampling procedure for the operator and the measured data can be provided.

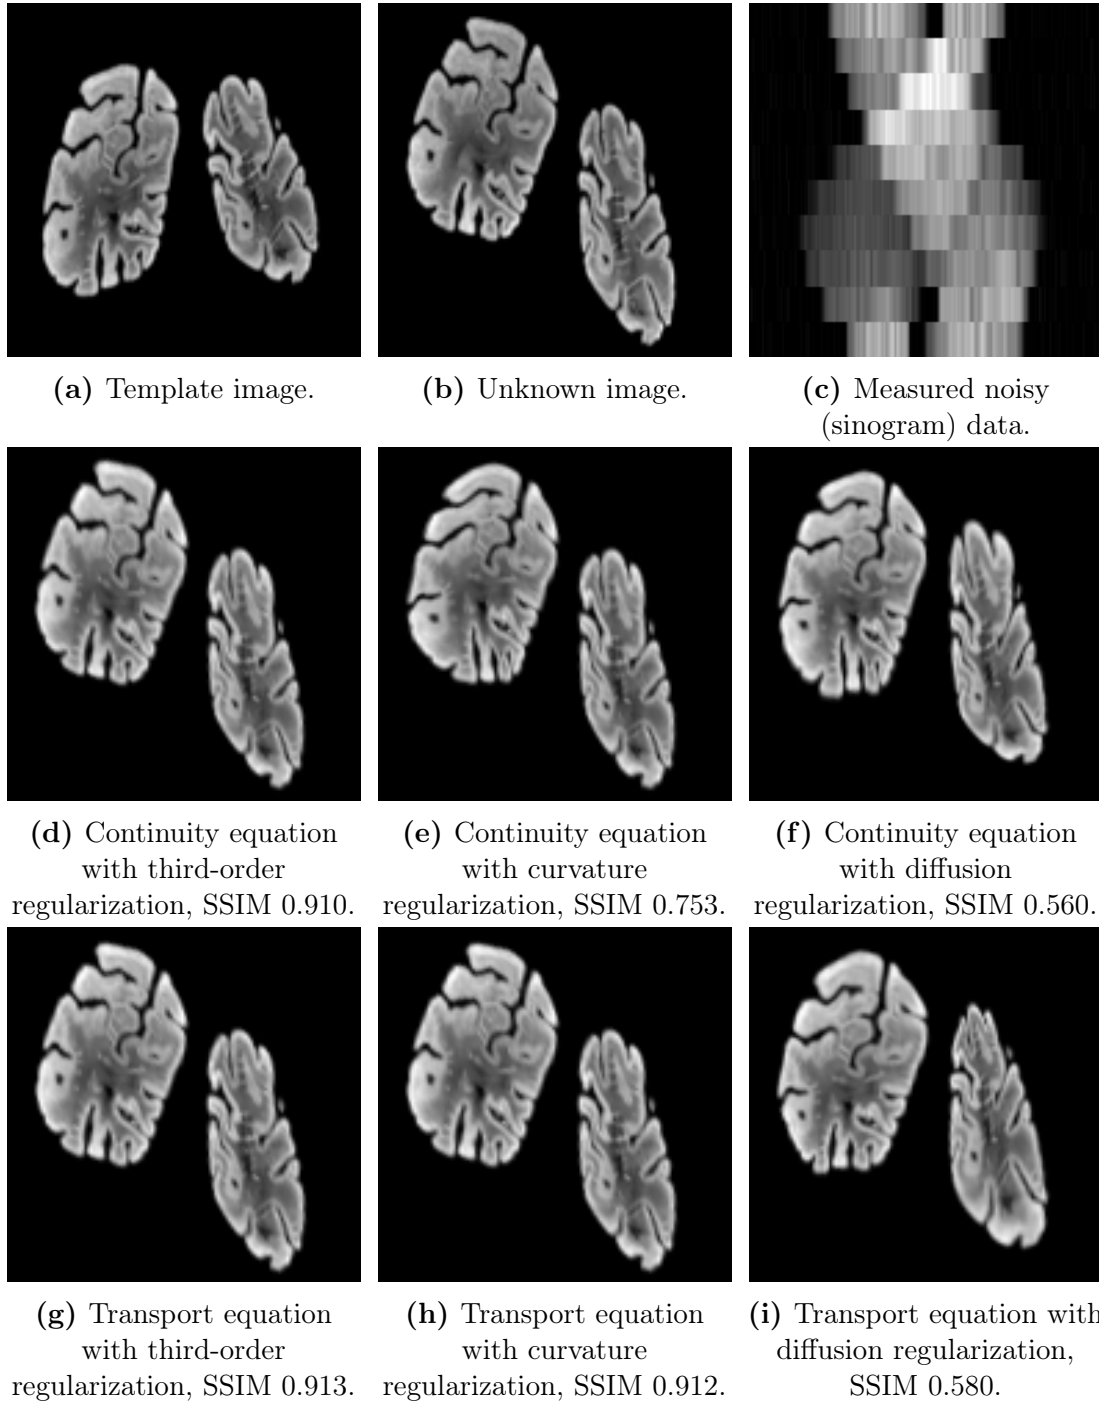


Figure 5.5: Reconstructions for the HNSP [205] image using our approach, different regularization functionals, and different PDE constraints. Here, ten measurements corresponding to ten angles equally distributed in the interval $[0, 180]$ degrees were taken. The measured data was corrupted with five percent noise.

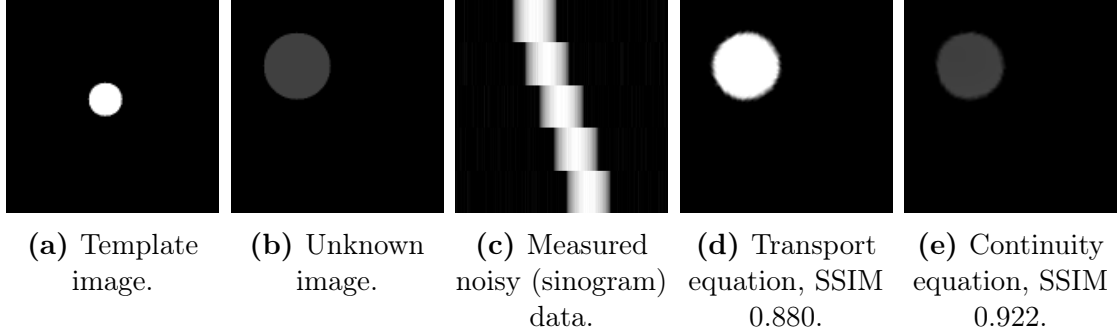


Figure 5.6: Reconstructions of an image showing a disk obtained with our method. Five measurements were taken at directions corresponding to five angles equally distributed in $[0, 90]$ degrees. As before, five percent noise was added.

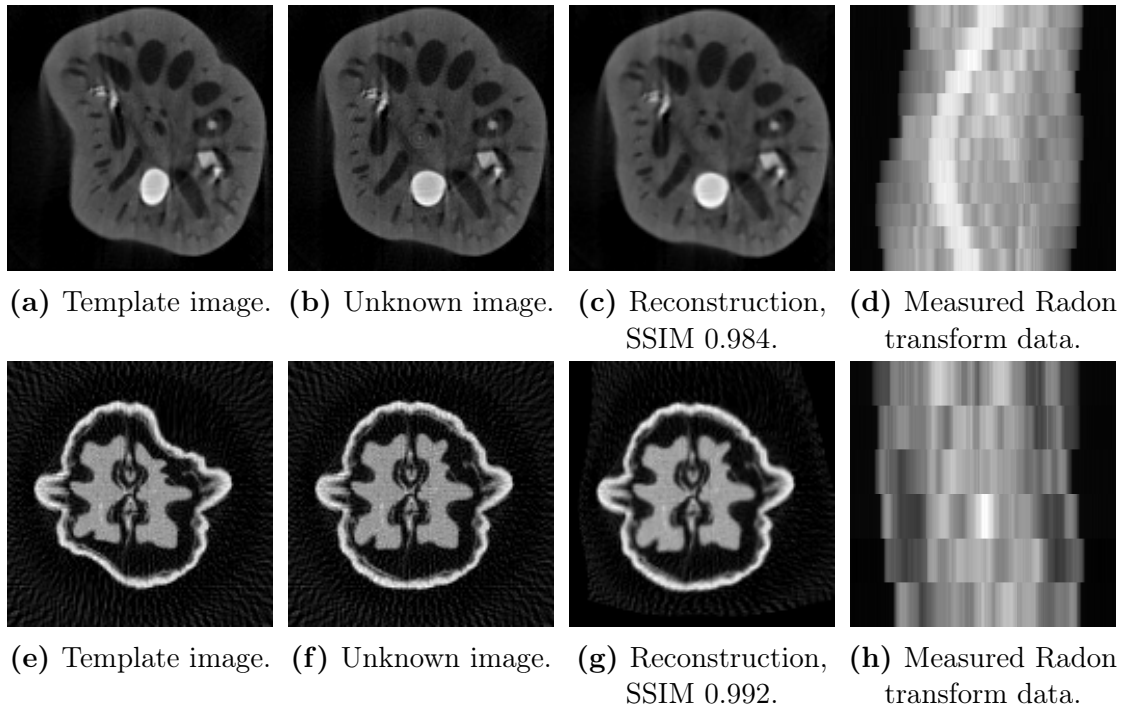


Figure 5.7: Reconstructions based on nonsynthetic X-ray tomographic measurements [50, 146] computed with our method using the transport equation together with curvature regularization. Measurements from twelve and six directions with angles in $[0, 180]$ degrees were used.




Figure 5.8: Reconstruction of a 3D volume (‘mice3D’, see [205]) using our method. In Fig. 5.8a, Fig. 5.8b, and Fig. 5.8d, slices (left to right, top to bottom) of each volume along the third coordinate direction are shown. In Fig. 5.8c, slices of the 3D Radon transform measurements are shown. Each slice corresponds to one measurements direction. In total, only ten measurements were taken at angles equally distributed in $[0, 180]$ degrees. As before, five percent noise was added.

From Optimal Transport to Discrepancy

Contents

6.1	Introduction	135
6.2	Preliminaries	136
6.3	Discrepancies	139
6.4	Optimal transport and Wasserstein distances	145
6.5	Regularized optimal transport	147
6.6	Sinkhorn divergence	154
6.7	Numerical approach and examples	156
6.8	Conclusions	160
6.A	Basic theorems	161

Abstract

This chapter¹ is published in [215]. A common way to quantify the “distance” between measures is via their discrepancy, also known as maximum mean discrepancy (MMD). Discrepancies are related to Sinkhorn divergences S_ε with appropriate cost functions as $\varepsilon \rightarrow \infty$. In the opposite direction, if $\varepsilon \rightarrow 0$, Sinkhorn divergences approach another important distance between measures, namely the Wasserstein distance or more generally optimal transport “distance”. In this chapter, we investigate the limiting process for arbitrary measures on compact sets and Lipschitz continuous cost functions. In particular, we are interested in the behavior of the corresponding optimal potentials $\hat{\varphi}_\varepsilon$, $\hat{\psi}_\varepsilon$ and $\hat{\varphi}_K$ appearing in the dual formulation of the Sinkhorn divergences and discrepancies, respectively. While part of the results are known, we provide rigorous proofs for some relations which we have not found in this generality in the literature. Finally, we demonstrate the limiting process by numerical examples and show the behavior of the distances when used for the approximation of measures by point measures in a process called dithering.

¹Accepted and soon to be published in Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging by Springer Nature.




Figure 6.1: Approximation of a measure on \mathbb{S}^2 by an empirical measure [134] (left) and a measure supported on a curve [102] (right) using discrepancies as objective function to minimize.

6.1 Introduction

The approximation of probability measures based on their discrepancies is a well examined problem in approximation and complexity theory [180, 194, 222]. Discrepancies appear in a wide range of applications, e.g., in the derivation of quadrature rules [222], the construction of designs [89], image dithering and representation [102, 134, 245, 259], see also Fig. 6.1, generative adversarial networks [96] and multivariate statistical testing [107, 137, 138]. In the last two applications, they are also called kernel based maximum mean discrepancies (MMDs).

On the other hand, optimal transport (OT) “distances” and in particular Wasserstein distances became very popular for tackling various problems in imaging sciences, graphics or machine learning [82]. There exists a large amount of papers both on the theory and applications of OT, for image dithering with Wasserstein distances see, e.g., [61, 129, 183].

Recently, regularized versions of OT for an efficient numerical treatment, known as Sinkhorn divergences [81], were used as replacement of OT in data science. Note that such regularization ideas are also investigated in the earlier works [242, 249, 280, 287]. For appropriately related transport cost functions and discrepancy kernels, the Sinkhorn divergences interpolate between the OT distance if the parameter goes to zero and the discrepancy if it goes to infinity [108]. In this chapter, the convergence behavior is examined for general measures on compact sets. Since cost functions applied in practice are mainly Lipschitz, we restrict our attention to such costs. This simplifies some proofs, since the theorem of Arzelà–Ascoli can be utilized. To make the paper self-contained, we provide most of the proofs although some

of them are not novel and the corresponding papers are cited in the context. For estimating approximation rates when approximating measures by those of certain subsets, see, e.g., [68, 102, 122, 222], the dual form of the discrepancy, respectively of the (regularized) Wasserstein distance, plays an important role. Therefore, we are interested in the properties of the optimal dual potentials for varying regularization parameters. In Proposition 6.13 we prove that the optimal dual potentials converge uniformly to certain functions as $\varepsilon \rightarrow \infty$. Then, in Corollary 6.15, we see that the normalized difference of these limiting functions coincides with the optimal potential in the dual form of the discrepancy if the cost function and the kernel are appropriately related. This behavior is underlined by a numerical example.

This chapter is organized as follows: Section 6.2 recalls basic results on measures, the Kullback-Leibler (KL) divergence and from convex analysis. In Section 6.3, we introduce discrepancies, in particular their dual formulation. Since these rely on positive definite kernels, we have a closer look at positive definite and conditionally positive definite kernels. Optimal transport and in particular Wasserstein distances are considered in Section 6.4. In Section 6.5, we investigate the limiting processes for the KL regularized OT distances, when the regularization parameter goes to zero or infinity. Some results in Proposition 6.8 are novel in this generality; Proposition 6.13 seems to be new as well. Remark 6.7 highlights why the KL divergence should be preferred as regularizer instead of the (neg)-entropy when dealing with non-discrete measures. KL regularized OT does not fulfill $\text{OT}_\varepsilon(\mu, \mu) = 0$, which motivates the definition of the Sinkhorn divergence S_ε in Section 6.6. Further, we prove Γ -convergence to the discrepancy as $\varepsilon \rightarrow \infty$ if the cost function of the Sinkhorn divergence is adapted to the kernel defining the discrepancy. Section 6.7 underlines the results on the limiting process by numerical examples. Further, we provide an example on the dithering of the standard Gaussian when Sinkhorn divergences with respect to different regularization parameters ε are involved. Finally, conclusions and directions of future research are given in Section 6.8.

6.2 Preliminaries

Measures Let \mathbb{X} be a compact Polish space (separable, complete metric space) with metric $d_{\mathbb{X}}$. By $\mathcal{B}(\mathbb{X})$ we denote the Borel σ -algebra on \mathbb{X} and by $\mathcal{M}(\mathbb{X})$ the linear space of all finite signed Borel measures on \mathbb{X} , i.e., all $\mu: \mathcal{B}(\mathbb{X}) \rightarrow \mathbb{R}$ satisfying $\mu(\mathbb{X}) < \infty$ and for any sequence $\{B_k\}_{k \in \mathbb{N}} \subset \mathcal{B}(\mathbb{X})$ of pairwise disjoint sets the relation $\mu(\cup_{k=1}^{\infty} B_k) = \sum_{k=1}^{\infty} \mu(B_k)$. In the following, the subset of non-negative measures is denoted by $\mathcal{M}^+(\mathbb{X})$. The *support of a measure* μ is defined as the closed set

$$\text{supp}(\mu) := \{x \in \mathbb{X} : B \subset \mathbb{X} \text{ open, } x \in B \implies \mu(B) > 0\}.$$

The *total variation* measure of $\mu \in \mathcal{M}(\mathbb{X})$ is defined by

$$|\mu|(B) := \sup \left\{ \sum_{k=1}^{\infty} |\mu(B_k)| : \bigcup_{k=1}^{\infty} B_k = B, B_k \text{ pairwise disjoint} \right\}.$$

With the norm $\|\mu\|_{\mathcal{M}} = |\mu|(\mathbb{X})$ the space $\mathcal{M}(\mathbb{X})$ becomes a Banach space. By $C(\mathbb{X})$ we denote the Banach space of continuous real-valued functions on \mathbb{X} equipped with the norm $\|\varphi\|_{C(\mathbb{X})} := \max_{x \in \mathbb{X}} |\varphi(x)|$. The space $\mathcal{M}(\mathbb{X})$ can be identified via Riesz' representation theorem with the dual space of $C(\mathbb{X})$ and the weak-* topology on $\mathcal{M}(\mathbb{X})$ gives rise to the *weak convergence of measures*. More precisely, a sequence $\{\mu_k\}_{k \in \mathbb{N}} \subset \mathcal{M}(\mathbb{X})$ converges *weakly* to μ and we write $\mu_k \rightharpoonup \mu$, if

$$\lim_{k \rightarrow \infty} \int_{\mathbb{X}} \varphi \, d\mu_k = \int_{\mathbb{X}} \varphi \, d\mu \quad \text{for all } \varphi \in C(\mathbb{X}).$$

For a non-negative, finite measure μ and $p \in [1, \infty)$, let $L^p(\mathbb{X}, \mu)$ be the Banach space (of equivalence classes) of complex-valued functions with norm

$$\|f\|_{L^p(\mathbb{X}, \mu)} = \left(\int_{\mathbb{X}} |f|^p \, d\mu \right)^{\frac{1}{p}} < \infty.$$

A measure $\nu \in \mathcal{M}(\mathbb{X})$ is *absolutely continuous* with respect to μ and we write $\nu \ll \mu$ if for every $A \in \mathcal{B}(\mathbb{X})$ with $\mu(A) = 0$ we have $\nu(A) = 0$. If $\mu, \nu \in \mathcal{M}^+(\mathbb{X})$ satisfy $\nu \ll \mu$, then the *Radon-Nikodym derivative* $\sigma_\nu \in L^1(\mathbb{X}, \mu)$ (also denoted by $\frac{d\nu}{d\mu}$) exists and $\nu = \sigma_\nu \mu$. Further, $\mu, \nu \in \mathcal{M}(\mathbb{X})$ are *mutually singular* and we write $\mu \perp \nu$ if two disjoint sets $X_\mu, X_\nu \in \mathcal{B}(\mathbb{X})$ exist such that $\mathbb{X} = X_\mu \cup X_\nu$ and for every $A \in \mathcal{B}(\mathbb{X})$ we have $\mu(A) = \mu(A \cap X_\mu)$ and $\nu(A) = \nu(A \cap X_\nu)$. For any $\mu, \nu \in \mathcal{M}^+(\mathbb{X})$, there exists a unique *Lebesgue decomposition* of μ with respect to ν given by $\mu = \sigma_\mu \nu + \mu^\perp$, where $\sigma_\mu \in L^1(\mathbb{X}, \nu)$ and $\mu^\perp \perp \nu$.

By $\mathcal{P}(\mathbb{X})$ we denote the *set of Borel probability measures* on \mathbb{X} , i.e., non-negative Borel measures with $\mu(\mathbb{X}) = 1$. This set is *weakly compact*, i.e., compact with respect to the weak-* topology. Note that there is an ambiguity in the notation as the above usual weak-* convergence is called weak convergence in stochastics. In Section 6.4, we introduce a metric on $\mathcal{P}(\mathbb{X})$ such that it becomes a Polish space.

Convex analysis The following can be found, e.g., in [46]. Let V be a real Banach space with dual V^* , i.e., the space of real-valued continuous linear functionals on V . We use the notation $\langle v, x \rangle = v(x)$, $v \in V^*, x \in V$. For $F: V \rightarrow (-\infty, +\infty]$, the *domain* of F is given by $\text{dom} F := \{x \in V : F(x) \in \mathbb{R}\}$. If $\text{dom} F \neq \emptyset$, then F is called *proper*. The *subdifferential* of $F: V \rightarrow (-\infty, +\infty]$ at a point $x_0 \in \text{dom} F$ is defined as

$$\partial F(x_0) := \{v \in V^* : F(x) \geq F(x_0) + \langle v, x - x_0 \rangle\},$$

and $\partial F(x_0) = \emptyset$ if $x_0 \notin \text{dom} F$. The *Fenchel conjugate* $F^*: V^* \rightarrow (-\infty, +\infty]$ is given by

$$F^*(v) = \sup_{x \in V} \{\langle v, x \rangle - F(x)\}.$$

If $F: V \rightarrow (-\infty, +\infty]$ is convex and lower semi-continuous (lsc) at $x \in \text{dom} F$, then

$$v \in \partial F(x) \quad \Leftrightarrow \quad x \in \partial F^*(v). \quad (6.1)$$

By $\Gamma_0(V)$ we denote the set of proper, convex, lsc functions mapping from V to $(-\infty, +\infty]$. Let W be another real Banach space. Then, for $F \in \Gamma_0(V)$, $G \in \Gamma_0(W)$ and a linear, bounded operator $A: V \rightarrow W$ with the property that there exists $x \in \text{dom} F$ such that G is continuous at Ax , the following Fenchel–Rockafellar duality relation is fulfilled

$$\sup_{x \in V} \{-F(-x) - G(Ax)\} = \inf_{w \in W^*} \{F^*(A^*w) + G^*(w)\}, \quad (6.2)$$

see [104, Thm. 4.1, p. 61], where we consider

$$\sup_{x \in V} \{-F(-x) - G(Ax)\} = - \inf_{x \in V} \{F(-x) + G(Ax)\}$$

as primal problem with respect to the notation in [104]. If the optimal (primal) solution \hat{x} exists, it is related to any optimal (dual) solution \hat{w} by

$$A\hat{x} \in \partial G^*(\hat{w}), \quad (6.3)$$

see [104, Prop. 4.1].

Kullback–Leibler divergence A function $f: [0, +\infty) \rightarrow [0, +\infty]$ is called *entropy function*, if it is convex, lsc and $\text{dom} f \cap (0, +\infty) \neq \emptyset$. The corresponding recession constant is given by $f'_\infty = \lim_{x \rightarrow \infty} \frac{f(x)}{x}$. For every $\mu, \nu \in \mathcal{M}^+(\mathbb{X})$ with Lebesgue decomposition $\mu = \sigma_\mu \nu + \mu^\perp$, the *f-divergence* is defined as

$$D_f(\mu, \nu) = \int_{\mathbb{X}} f \circ \sigma_\mu \, d\nu + f'_\infty \mu^\perp(\mathbb{X}). \quad (6.4)$$

In case that $f'_\infty = \infty$ and $\mu^\perp(\mathbb{X}) = 0$, we make the usual convention $\infty \cdot 0 = 0$. The *f-divergence* fulfills $D_f(\mu, \nu) \geq 0$ for all $\mu, \nu \in \mathcal{M}^+(\mathbb{X})$ with equality if and only if $\mu = \nu$, and is in general neither symmetric nor satisfies a triangle inequality. The associated mapping $D_f: \mathcal{M}^+(\mathbb{X}) \times \mathcal{M}^+(\mathbb{X}) \rightarrow [0, +\infty]$ is jointly convex and weakly lsc, see [188, Cor. 2.9]. The *f-divergence* can be written in the dual form

$$D_f(\mu, \nu) = \sup_{\varphi \in C(\mathbb{X})} \int_{\mathbb{X}} \varphi \, d\mu - \int_{\mathbb{X}} f^* \circ \varphi \, d\nu,$$

see [188, Rem. 2.10]. Hence, $D_f(\cdot, \nu)$ is the Fenchel conjugate of $H: C(\mathbb{X}) \rightarrow \mathbb{R}$ given by $H(\varphi) := \int_{\mathbb{X}} f^* \circ \varphi \, d\nu$. If f^* is differentiable, we directly deduce from (6.1) that

$$\varphi \in \partial_\mu D_f(\mu, \nu) \Leftrightarrow \mu = \nabla H(\varphi) \Leftrightarrow \mu = \nabla f^* \circ \varphi \, \nu. \quad (6.5)$$

In the following, we focus on the *Shannon–Boltzmann entropy* function and its Fenchel conjugate given by

$$f(x) = x \log(x) - x + 1 \quad \text{and} \quad f^*(x) = \exp(x) - 1$$

with the agreement $0 \log 0 = 0$. The corresponding *f-divergence* is the *Kullback–Leibler divergence* $\text{KL}: \mathcal{M}^+(\mathbb{X}) \times \mathcal{M}^+(\mathbb{X}) \rightarrow [0, +\infty]$. For $\mu, \nu \in \mathcal{M}^+(\mathbb{X})$ with

existing Radon-Nikodym derivative $\sigma_\mu = \frac{d\mu}{d\nu}$ of μ with respect to ν , formula (6.4) can be written as

$$\text{KL}(\mu, \nu) := \int_{\mathbb{X}} \log(\sigma_\mu) d\mu + \nu(\mathbb{X}) - \mu(\mathbb{X}). \quad (6.6)$$

In case that the above Radon-Nikodym derivative does not exist, (6.4) implies $\text{KL}(\mu, \nu) = +\infty$. For $\mu, \nu \in \mathcal{P}(\mathbb{X})$ the last two summands in (6.6) cancel each other. Hence, we have for discrete measures $\mu = \sum_{j=1}^n \mu_j \delta_{x_j}$ and $\nu = \sum_{j=1}^n \nu_j \delta_{x_j}$ with $\mu_j, \nu_j \geq 0$ and $\sum_{j=1}^n \mu_j = \sum_{j=1}^n \nu_j = 1$ that

$$\text{KL}(\mu, \nu) = \sum_{j=1}^n \log\left(\frac{\mu_j}{\nu_j}\right) \mu_j.$$

Further, the KL divergence is strictly convex with respect to the first variable. Due to the Fenchel conjugate pairing

$$H(\varphi) = \int_{\mathbb{X}} \exp(\varphi) - 1 d\nu \quad \text{and} \quad H^*(\mu) = \text{KL}(\mu, \nu), \quad (6.7)$$

the derivative relation (6.5) simplifies to

$$\varphi \in \partial_\mu \text{KL}(\mu, \nu) \quad \Leftrightarrow \quad \mu = e^\varphi \nu \quad \Leftrightarrow \quad \varphi = \log\left(\frac{d\mu}{d\nu}\right). \quad (6.8)$$

Finally, note that the KL divergence and the total variation norm $\|\cdot\|_{\mathcal{M}}$ are related by the *Pinsker inequality* $\|\mu - \nu\|_{\mathcal{M}}^2 \leq \text{KL}(\mu, \nu)$.

6.3 Discrepancies

In this section, we introduce the notation of discrepancies and have a closer look at (conditionally) positive definite kernels. In particular, we emphasize how conditionally positive definite kernels can be modified to positive definite ones.

Let $\sigma_{\mathbb{X}} \in \mathcal{M}(\mathbb{X})$ be non-negative with $\text{supp}(\sigma_{\mathbb{X}}) = \mathbb{X}$. The given definition of discrepancies is based on symmetric, positive definite, continuous kernels. There is a close relation to general discrepancies related to measures on $\mathcal{B}(\mathbb{X})$, see [222]. Recall that a symmetric function $K: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ is *positive definite* if for any finite number $n \in \mathbb{N}$ of points $x_j \in \mathbb{X}$, $j = 1, \dots, n$, the relation

$$\sum_{i,j=1}^n a_i a_j K(x_i, x_j) \geq 0$$

is satisfied for all $(a_j)_{j=1}^n \in \mathbb{R}^n$ and *strictly positive definite* if strict inequality holds for all $(a_j)_{j=1}^n \neq 0$. Assuming that $K \in C(\mathbb{X} \times \mathbb{X})$ is symmetric, positive definite, we know by Mercer's theorem [80, 195, 257] that there exists an orthonormal basis

$\{\phi_k : k \in \mathbb{N}\}$ of $L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ and non-negative coefficients $\{\alpha_k\}_{k \in \mathbb{N}} \in \ell_1$ such that K has the Fourier expansion

$$K(x, y) = \sum_{k=0}^{\infty} \alpha_k \phi_k(x) \overline{\phi_k(y)} \quad (6.9)$$

with absolute and uniform convergence of the right-hand side. If $\alpha_k > 0$ for some $k \in \mathbb{N}_0$, the corresponding function ϕ_k is continuous. Every function $f \in L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ has a Fourier expansion

$$f = \sum_{k=0}^{\infty} \hat{f}_k \phi_k, \quad \hat{f}_k := \int_{\mathbb{X}} f \overline{\phi_k} d\sigma_{\mathbb{X}}.$$

Moreover, for $k \in \mathbb{N}_0$ with $\alpha_k > 0$, the *Fourier coefficients* of $\mu \in \mathcal{P}(\mathbb{X})$ are well-defined by

$$\hat{\mu}_k := \int_{\mathbb{X}} \overline{\phi_k} d\mu.$$

The kernel K gives rise to a *reproducing kernel Hilbert space* (RKHS). More precisely, the function space

$$H_K(\mathbb{X}) := \left\{ f \in L^2(\mathbb{X}, \sigma_{\mathbb{X}}) : \sum_{k=0}^{\infty} \alpha_k^{-1} |\hat{f}_k|^2 < \infty \right\}$$

equipped with the inner product and the corresponding norm

$$\langle f, g \rangle_{H_K(\mathbb{X})} = \sum_{k=0}^{\infty} \alpha_k^{-1} \hat{f}_k \overline{\hat{g}_k}, \quad \|f\|_{H_K(\mathbb{X})} = \langle f, f \rangle_{H_K(\mathbb{X})}^{\frac{1}{2}} \quad (6.10)$$

forms a Hilbert space with reproducing kernel, i.e.,

$$\begin{aligned} K(x, \cdot) &\in H_K(\mathbb{X}) && \text{for all } x \in \mathbb{X}, \\ f(x) &= \langle f, K(x, \cdot) \rangle_{H_K(\mathbb{X})} && \text{for all } f \in H_K(\mathbb{X}), x \in \mathbb{X}. \end{aligned} \quad (6.11)$$

Note that $f \in H_K(\mathbb{X})$ implies $\hat{f}_k = 0$ if $\alpha_k = 0$, in which case we make the convention $\alpha_k^{-1} \hat{f}_k = 0$ in (6.10). Indeed, $H_K(\mathbb{X})$ is the closure of the linear span of $\{K(x_j, \cdot) : x_j \in \mathbb{X}\}$ with respect to the norm (6.10). The space $H_K(\mathbb{X})$ is continuously embedded in $C(\mathbb{X})$ and hence point evaluations in $H_K(\mathbb{X})$ are continuous. Since the series in (6.9) converges uniformly and the functions ϕ_k are continuous, the function

$$\|K(x, \cdot)\|_{H_K(\mathbb{X})} = \left\| \sum_{k=0}^{\infty} \alpha_k \phi_k(x) \overline{\phi_k(\cdot)} \right\|_{H_K(\mathbb{X})} = \left(\sum_{k=0}^{\infty} \alpha_k |\phi_k(x)|^2 \right)^{\frac{1}{2}}$$

is also continuous so that we have $\int_{\mathbb{X}} \|K(x, \cdot)\|_{H_K(\mathbb{X})} d\mu(x) < \infty$. By the definition of Bochner integrals, see [165, Prop. 1.3.1], we have for any $\mu \in \mathcal{P}(\mathbb{X})$ that

$$\int_{\mathbb{X}} K(x, \cdot) d\mu(x) \in H_K(\mathbb{X}). \quad (6.12)$$

For $\mu, \nu \in \mathcal{M}(\mathbb{X})$, the *discrepancy* $\mathcal{D}_K(\mu, \nu)$ is defined as norm of the linear operator $T: H_K \rightarrow \mathbb{R}$ with $\varphi \mapsto \int_{\mathbb{X}} \varphi d\xi$,

$$\mathcal{D}_K(\mu, \nu) = \max_{\|\varphi\|_{H_K(\mathbb{X})} \leq 1} \int_{\mathbb{X}} \varphi d\xi, \quad (6.13)$$

where $\xi := \mu - \nu$, see [128, 222]. If $\mu_n \rightharpoonup \mu$ and $\nu_n \rightharpoonup \nu$ as $n \rightarrow \infty$, then also $\mu_n \otimes \nu_n \rightharpoonup \mu \otimes \nu$. Thus, continuity of K implies that $\lim_{n \rightarrow \infty} \mathcal{D}_K(\mu_n, \nu_n) = \mathcal{D}_K(\mu, \nu)$. Since

$$\int_{\mathbb{X}} \varphi d\xi = \int_{\mathbb{X}} \langle \varphi, K(x, \cdot) \rangle_{H_K(\mathbb{X})} d\xi(x) = \left\langle \varphi, \int_{\mathbb{X}} K(x, \cdot) d\xi(x) \right\rangle_{H_K(\mathbb{X})},$$

we obtain by Schwarz' inequality that the optimal dual potential (up to the sign) is given by

$$\hat{\varphi}_K = \frac{\int_{\mathbb{X}} K(x, \cdot) d\xi(x)}{\left\| \int_{\mathbb{X}} K(x, \cdot) d\xi(x) \right\|_{H_K(\mathbb{X})}} = \frac{\int_{\mathbb{X}} K(x, \cdot) d\mu(x) - \int_{\mathbb{X}} K(x, \cdot) d\nu(x)}{\left\| \int_{\mathbb{X}} K(x, \cdot) d\mu(x) - \int_{\mathbb{X}} K(x, \cdot) d\nu(x) \right\|_{H_K(\mathbb{X})}}. \quad (6.14)$$

In the following, it is always clear from the context if the Fourier transform of the function or the optimal dual potential is meant. Further, Riesz' representation theorem implies

$$\mathcal{D}_K(\mu, \nu) = \max_{\|\varphi\|_{H_K(\mathbb{X})} \leq 1} \int_{\mathbb{X}} \varphi d\xi = \left\| \int_{\mathbb{X}} K(x, \cdot) d\xi(x) \right\|_{H_K(\mathbb{X})},$$

so that we conclude by Fubini's theorem and (6.11) that

$$\begin{aligned} \mathcal{D}_K^2(\mu, \nu) &= \left\| \int_{\mathbb{X}} K(x, \cdot) d\xi(x) \right\|_{H_K(\mathbb{X})}^2 = \int_{\mathbb{X}^2} K d(\xi \otimes \xi) \\ &= \int_{\mathbb{X}^2} K d(\mu \otimes \mu) + \int_{\mathbb{X}^2} K d(\nu \otimes \nu) - 2 \int_{\mathbb{X}^2} K d(\mu \otimes \nu). \end{aligned} \quad (6.15)$$

By (6.9), we finally get

$$\mathcal{D}_K^2(\mu, \nu) = \sum_{k=0}^{\infty} \alpha_k |\hat{\mu}_k - \hat{\nu}_k|^2, \quad (6.16)$$

where the summation runs over all $k \in \mathbb{N}_0$ with $\alpha_k > 0$.

Remark 6.1. (Relation to attraction-repulsion functionals) We briefly consider the relation to attraction-repulsion functionals motivated from electrostatic halftoning, see [245, 259]. Let $\nu = w dx$ be fixed, for example a continuous (normalized) image with gray values in $[0, 1]$ represented by $w: \mathbb{X} \rightarrow [0, 1]$, where pure black is the largest value of w and white the smallest one. Then, looking for a discrete measure $\mu = \frac{1}{M} \sum_{j=1}^M \delta(\cdot - p_j)$ that approximates ν by minimizing the squared discrepancy is equivalent to solving the minimization problem

$$\arg \min_{p \in \mathbb{R}^M} \left\{ \underbrace{\frac{1}{2M} \sum_{i,j=1}^M K(p_i, p_j)}_{\text{repulsion}} - \underbrace{\sum_{i=1}^M \int_{\mathbb{X}} w(x) K(x, p_i)}_{\text{attraction}} \right\}.$$

For $K(x, y) = h(\|x - y\|)$ and an decreasing function $h: [0, +\infty) \rightarrow \mathbb{R}$, it becomes clear that

- the first term is minimal if the points are far away from each other, implying a *repulsion*;
- the second (negative) term becomes maximal if for large $w(x)$, there are many points positioned in this area; so it can be considered as an *attraction* steered by w .

Kernels. In this paragraph, we want to have a closer look at appropriate kernels. Recall that for symmetric, positive definite kernels $K_i \in C(\mathbb{X} \times \mathbb{X})$, $i = 1, 2$, and $\alpha > 0$, the kernels αK_1 , $K_1 + K_2$, $K_1 \cdot K_2$ and $\exp(K_1)$ are again positive definite, see [256, Lems. 4.5 and 4.6].

Of special interest are so-called radial kernels of the form

$$K(x, y) := h(d_{\mathbb{X}}(x, y)),$$

where $h: [0, +\infty) \rightarrow \mathbb{R}$. In the following, the discussion is restricted to compact sets \mathbb{X} in \mathbb{R}^d and the Euclidean distance $d_{\mathbb{X}}(x, y) = \|x - y\|$. Many results on positive definite functions on \mathbb{R}^d go back to Schoenberg [246] and Micchelli [198]. For a good overview, we refer to [278], where some of the following statements can be found. Clearly, restricting positive definite kernels on \mathbb{R}^d to compact subsets \mathbb{X} results in positive definite kernels on \mathbb{X} . The radial kernels related to the Gaussian, which are quite popular in MMDs, and the inverse multiquadric given by

$$h(r) = e^{-r^2/c^2} \quad \text{and} \quad h(r) = (c^2 + r^2)^{-p}, \quad c, p > 0,$$

are known to be strictly positive definite on \mathbb{R}^d for every $d \in \mathbb{N}$. Further, the following compactly supported functions h give rise to positive definite kernels in \mathbb{R}^d :

$$h(r) = (1 - r)_+^p, \quad p \geq \left\lfloor \frac{d}{2} \right\rfloor + 1, \quad (6.17)$$

where $\lfloor a \rfloor$ denotes the largest integer less or equal than $a \in \mathbb{R}$ and $a_+ := \max(a, 0)$.

In connection with Wasserstein distances, we are interested in (negative) powers of distances $K(x, y) = \|x - y\|^p$, $p > 0$, related to the functions $h(r) = r^p$. Unfortunately, all these functions are not positive definite! By (6.17), we know that $\tilde{K}(x, y) = 1 - \|x - y\|$ is positive definite in one dimension $d = 1$. A more general result for the Euclidean distance is given in the following proposition.

Proposition 6.2. *Let $K(x, y) = -\|x - y\|$. For every compact set $\mathbb{X} \subset \mathbb{R}^d$, there exists a constant $C > 0$ such that the function*

$$\tilde{K}(x, y) := C - \|x - y\|$$

is positive definite on \mathbb{X} . Further, for $\mu, \nu \in \mathcal{P}(\mathbb{X})$, it holds

$$\mathcal{D}_{\tilde{K}}^2(\mu, \nu) = \mathcal{D}_K^2(\mu, \nu) \quad \text{and} \quad \hat{\varphi}_{\tilde{K}} = \hat{\varphi}_K.$$

Proof. In [131, Cor. 2.15] it was shown that \tilde{K} is positive definite. The rest follows in a straightforward way from (6.15) and (6.14) regarding that μ and ν are probability measures. \square

Some interesting functions such as negative powers of Euclidean distances or the smoothed distance function $\sqrt{c^2 + \|x - y\|^2}$, $0 < c \ll 1$, are conditionally positive definite. Let $\Pi_{m-1}(\mathbb{R}^d)$ denote the $\binom{d+m-1}{d}$ -dimensional space of polynomials on \mathbb{R}^d of absolute degree (sum of exponents) $\leq m - 1$. A function $K: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ is *conditionally positive definite of order m* if for all points $x_1, \dots, x_n \in \mathbb{R}^d$, $n \in \mathbb{N}$, the relation

$$\sum_{i,j=1}^n a_i a_j K(x_i, x_j) \geq 0 \quad (6.18)$$

holds true for all $a_1, \dots, a_n \in \mathbb{R}$ satisfying

$$\sum_{i=1}^n a_i P(x_i) = 0 \quad \text{for all } P \in \Pi_{m-1}(\mathbb{R}^d).$$

If strong inequality holds in (6.18) except for $a_i = 0$ for all $i = 1, \dots, n$, then K is called *strictly conditionally positive definite of order m* . In particular, for $m = 1$, the condition (6.18) relaxes to $\sum_{i=1}^n a_i = 0$.

The radial kernels related to the following functions are strictly conditionally positive definite of order m on \mathbb{R}^d :

$$\begin{aligned} h(r) &= (-1)^{\lceil p \rceil} (c^2 + r^2)^p, & p > 0, p \notin \mathbb{N}, m = \lceil p \rceil, \\ h(r) &= (-1)^{\lceil p/2 \rceil} r^p, & p > 0, p \notin 2\mathbb{N}, m = \lceil p/2 \rceil, \\ h(r) &= (-1)^{k+1} r^{2k} \log(r), & k \in \mathbb{N}, m = k + 1, \end{aligned}$$

where $\lceil a \rceil$, denotes the smallest integer larger or equal than $a \in \mathbb{R}$. The first group of functions are called multiquadric and the last group is known as thin plate splines. In connection with Wasserstein distances, the second group of functions is of interest.

By the following lemma, it is easy to turn conditionally positive definite functions into positive definite ones. However, only for conditionally positive definite functions of order $m = 1$, the discrepancy remains the same.

Lemma 6.3. *Let $\Xi := \{u_k : k = 1, \dots, N\}$ with $N := \binom{d+m-1}{m-1}$ be a set of points such that $P(u_k) = 0$ for all $k = 1, \dots, N$, $P \in \Pi_{m-1}(\mathbb{R}^d)$, is only fulfilled for the zero polynomial. Denote by $\{P_k : k = 1, \dots, N\}$ the set of Lagrangian basis polynomials with respect to Ξ , i.e., $P_k(u_j) = \delta_{jk}$. Let $K \in C(\mathbb{X} \times \mathbb{X})$ be a symmetric conditionally positive definite kernel of order m .*

i) Then

$$\begin{aligned}\tilde{K}(x, y) := & K(x, y) - \sum_{j=1}^N P_j(x)K(u_j, y) - \sum_{k=1}^N P_k(y)K(x, u_k) \\ & + \sum_{j,k=1}^N P_j(x)P_k(y)K(u_j, u_k)\end{aligned}$$

is a positive definite kernel.

ii) If μ and ν have the same moments up to order $m - 1$, then they satisfy $\mathcal{D}_K^2(\mu, \nu) = \mathcal{D}_{\tilde{K}}^2(\mu, \nu)$.

iii) In particular, we have for $m = 1$, $\mu, \nu \in \mathcal{P}(\mathbb{X})$ and any fixed $u \in \mathbb{X}$ that

$$\tilde{K}(x, y) = K(x, y) - K(u, y) - K(x, u) + K(u, u) \quad (6.19)$$

and

$$\begin{aligned}\mathcal{D}_{\tilde{K}}^2(\mu, \nu) &= \mathcal{D}_K^2(\mu, \nu), \\ \hat{\varphi}_{\tilde{K}} &= \frac{\int_{\mathbb{X}} K(x, \cdot) d\mu(x) - \int_{\mathbb{X}} K(x, \cdot) d\nu(x) + c_\nu - c_\mu}{\left\| \int_{\mathbb{X}} K(x, \cdot) d\mu(x) - \int_{\mathbb{X}} K(x, \cdot) d\nu(x) + c_\nu - c_\mu \right\|_{H_K(\mathbb{X})}},\end{aligned}$$

where

$$c_\mu := \int_{\mathbb{X}} K(x, u) d\mu(x) \quad \text{and} \quad c_\nu := \int_{\mathbb{X}} K(x, u) d\nu(x). \quad (6.20)$$

Proof. i) This part follows by straightforward computation, see [278, Thm. 10.18].

ii) Assuming that μ and ν have the same moments up to order $m - 1$, i.e.,

$$p_j = \int_{\mathbb{X}} P_j(x) d\mu(x) = \int_{\mathbb{X}} P_j(x) d\nu(x), \quad j = 1, \dots, N,$$

and abbreviating for the symmetric kernels

$$c_{\mu,j} := \int_{\mathbb{X}} K(u_j, y) d\mu(y), \quad c_{\nu,j} := \int_{\mathbb{X}} K(u_j, y) d\nu(y),$$

we obtain by definition of \tilde{K} that

$$\begin{aligned}& \mathcal{D}_{\tilde{K}}^2(\mu, \nu) \\ &= \int_{\mathbb{X}^2} \tilde{K} d(\mu \otimes \mu) + \int_{\mathbb{X}^2} \tilde{K} d(\nu \otimes \nu) - 2 \int_{\mathbb{X}^2} \tilde{K} d(\mu \otimes \nu) \\ &= \mathcal{D}_K^2(\mu, \nu) - \sum_{j=1}^N p_j(c_{\mu,j} + c_{\nu,j}) - \sum_{k=1}^N p_j(c_{\mu,k} + c_{\nu,k}) + 2 \sum_{j,k=1}^N p_j p_k K(u_j, u_k) \\ &\quad + \sum_{j=1}^N p_j(c_{\mu,j} + c_{\nu,j}) + \sum_{k=1}^N p_j(c_{\mu,k} + c_{\nu,k}) - 2 \sum_{j,k=1}^N p_j p_k K(u_j, u_k) \\ &= \mathcal{D}_K^2(\mu, \nu).\end{aligned}$$

iii) Let $m = 1$. Then we have for the optimal dual potential in (6.14) related to $\mathcal{D}_{\tilde{K}}$ that

$$\begin{aligned}\hat{\varphi}_{\tilde{K}} &= \frac{\int_{\mathbb{X}} \tilde{K}(x, \cdot) d\mu(x) - \int_{\mathbb{X}} \tilde{K}(x, \cdot) d\nu(x)}{\left\| \int_{\mathbb{X}} \tilde{K}(x, \cdot) d\mu(x) - \int_{\mathbb{X}} \tilde{K}(x, \cdot) d\nu(x) \right\|_{H_K(\mathbb{X})}} \\ &= \frac{\int_{\mathbb{X}} K(x, \cdot) d\mu(x) - \int_{\mathbb{X}} K(x, \cdot) d\nu(x) + c_\nu - c_\mu}{\left\| \int_{\mathbb{X}} K(x, \cdot) d\mu(x) - \int_{\mathbb{X}} K(x, \cdot) d\nu(x) + c_\nu - c_\mu \right\|_{H_K(\mathbb{X})}}.\end{aligned}$$

□

6.4 Optimal transport and Wasserstein distances

The following discussion about optimal transport is based on [6, 82, 243], where many aspects simplify due to the compactness of \mathbb{X} and the assumption that the cost c is Lipschitz continuous. Let $\mu, \nu \in \mathcal{P}(\mathbb{X})$ and $c \in C(\mathbb{X} \times \mathbb{X})$ be a non-negative, symmetric and Lipschitz continuous function. Then, the *Kantorovich problem of optimal transport* (OT) reads

$$\text{OT}(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{X}^2} c d\pi, \quad (6.21)$$

where $\Pi(\mu, \nu)$ denotes the set of joint probability measures π on \mathbb{X}^2 with marginals μ and ν . In our setting, the OT functional $\pi \mapsto \int_{\mathbb{X}^2} c d\pi$ is weakly continuous, (6.21) has a solution and every such minimizer $\hat{\pi}$ is called optimal transport plan. In general, we can not expect the optimal transport plan to be unique. However, if \mathbb{X} is a compact subset of a separable Hilbert space, $c(x, y) = \|x - y\|_{\mathbb{X}}^p$, $p \in (1, \infty)$, and either μ or ν is *regular*, see [6, Def. 6.2.2] for the technical definition, then (6.21) has a unique solution. Instead of giving the exact definition, we want to remark that for $\mathbb{X} = \mathbb{R}^d$ the regular measures are precisely the ones which have a density with respect to the Lebesgue measure.

The c -transform $\varphi^c \in C(\mathbb{X})$ of $\varphi \in C(\mathbb{X})$ is defined as

$$\varphi^c(y) = \min_{x \in \mathbb{X}} \{c(x, y) - \varphi(x)\}.$$

Note that φ^c has the same Lipschitz constant as c . A function $\varphi^c \in C(\mathbb{X})$ is called c -concave if it is the c -transform of some function $\varphi \in C(\mathbb{X})$.

The dual formulation of the OT problem (6.21) reads

$$\text{OT}(\mu, \nu) = \max_{\substack{(\varphi, \psi) \in C(\mathbb{X})^2 \\ \varphi(x) + \psi(y) \leq c(x, y)}} \int_{\mathbb{X}} \varphi d\mu + \int_{\mathbb{X}} \psi d\nu. \quad (6.22)$$

Maximizing pairs are essentially of the form $(\varphi, \psi) = (\hat{\varphi}, \hat{\varphi}^c)$ for some c -concave function $\hat{\varphi}$ and fulfill $\hat{\varphi}(x) + \hat{\varphi}^c(y) = c(x, y)$ in $\text{supp}(\hat{\pi})$, where $\hat{\pi}$ is any optimal transport plan. The function $\hat{\varphi}$ is called (Kantorovich) potential for the couple (μ, ν) . If $(\hat{\varphi}, \hat{\psi})$ is an optimal pair, clearly also $(\hat{\varphi} - C, \hat{\psi} + C)$ with $C \in \mathbb{R}$ is optimal

and manipulations outside of $\text{supp}(\mu)$ and $\text{supp}(\nu)$ do not change the functional value. But even if we exclude such manipulations, the optimal dual potentials are in general not unique as Example 6.4 shows.

Example 6.4. We choose $\mathbb{X} = [0, 1]$, $c(x, y) = |x - y|$, $\mu = \delta_0/2 + \delta_1/2$ and $\nu = \delta_{0.1}/2 + \delta_{0.9}/2$. Then, $\text{OT}(\mu, \nu) = 0.1$ with the unique optimal transport plan $\hat{\pi} = \frac{1}{2}\delta_{0,0.1} + \frac{1}{2}\delta_{1,0.9}$. Optimal dual potentials are given by

$$\hat{\varphi}_1(x) = \begin{cases} 0.1 - x & \text{for } x \in [0, 0.1], \\ x - 0.9 & \text{for } x \in [0.9, 1], \\ 0 & \text{else,} \end{cases} \quad \text{and} \quad \hat{\varphi}_2(x) = \begin{cases} 0.2 - x & \text{for } x \in [0, 0.2], \\ x - 0.9 & \text{for } x \in [0.9, 1], \\ 0 & \text{else.} \end{cases}$$

Clearly, these potentials do not differ only by a constant.

Remark 6.5. Note that the space $C(\mathbb{X})^2$ in the dual problem could also be replaced with $C(\text{supp}(\mu)) \times C(\text{supp}(\nu))$. Using the Tietze extension theorem, any feasible point of the restricted problem can be extended to a feasible point of the original problem and hence the problems coincide. If the problem is restricted, all other concepts have to be adapted accordingly.

For $p \in [1, \infty)$, the p -Wasserstein distance W_p between $\mu, \nu \in \mathcal{P}(\mathbb{X})$ is defined by

$$W_p(\mu, \nu) := \left(\min_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{X}^2} d(x, y)^p d\pi(x, y) \right)^{\frac{1}{p}}.$$

It is a metric on $\mathcal{P}(\mathbb{X})$, which metrizes the weak topology. Indeed, due to compactness of \mathbb{X} , we have that $\mu_k \rightharpoonup \mu$ if and only if $\lim_{k \rightarrow \infty} W_p(\mu_k, \mu) = 0$.

For $1 \leq p \leq q < \infty$ it holds $W_p \leq W_q$. The distance W_1 is also called *Kantorovich-Rubinstein distance* or *Earth's mover distance*. Here, it holds $\varphi^c = -\varphi$ and the dual problem reads

$$W_1(\mu, \nu) = \max_{|\varphi|_{\text{Lip}(\mathbb{X})} \leq 1} \int_{\mathbb{X}} \varphi d\xi, \quad \xi := \mu - \nu,$$

where the maximum is taken over all Lipschitz continuous functions with Lipschitz constant bounded by 1. This looks similar to the discrepancy (6.13), but the space of test functions is larger for W_1 .

The distance W_1 is related to W_p by

$$W_1(\mu, \nu) \leq W_p(\mu, \nu) \leq CW_1(\mu, \nu)^{\frac{1}{p}}$$

with a constant $0 \leq C < \infty$ depending on $\text{diam}(\mathbb{X})$ and p .

6.5 Regularized optimal transport

In this section, we give a self-contained introduction to continuous *regularized optimal transport*. For $\mu, \nu \in \mathcal{P}(\mathbb{X})$ and $\varepsilon > 0$, regularized OT is defined as

$$\text{OT}_\varepsilon(\mu, \nu) := \min_{\pi \in \Pi(\mu, \nu)} \left\{ \int_{\mathbb{X}^2} c \, d\pi + \varepsilon \text{KL}(\pi, \mu \otimes \nu) \right\}. \quad (6.23)$$

Compared to the original OT problem, we will see in the numerical part that OT_ε can be efficiently solved numerically, see also [82]. Moreover, OT_ε has the following properties.

Lemma 6.6.

- i) *There is a unique minimizer $\hat{\pi}_\varepsilon \in \mathcal{P}(\mathbb{X}^2)$ of (6.23) with finite value.*
- ii) *The function OT_ε is weakly continuous and Fréchet differentiable.*
- iii) *For any $\mu, \nu \in \mathcal{P}(\mathbb{X})$ and $\varepsilon_1, \varepsilon_2 \in [0, \infty]$ with $\varepsilon_1 \leq \varepsilon_2$ it holds*

$$\text{OT}_{\varepsilon_1}(\mu, \nu) \leq \text{OT}_{\varepsilon_2}(\mu, \nu).$$

Proof. i) First, note that $\mu \otimes \nu$ is a feasible point and hence the infimum is finite. Existence of minimizers follows as the functional is weakly lsc and $\Pi(\mu, \nu) \subset \mathcal{P}(\mathbb{X}^2)$ is weakly compact. Uniqueness follows since $\text{KL}(\cdot, \mu \otimes \nu)$ is strictly convex.

ii) The proof uses the dual formulation in Proposition 6.9, see [108, Prop. 2].

iii) Let $\hat{\pi}_{\varepsilon_2}$ be the minimizer for $\text{OT}_{\varepsilon_2}(\mu, \nu)$. Then, it holds

$$\begin{aligned} \text{OT}_{\varepsilon_2}(\mu, \nu) &= \int_{\mathbb{X}^2} c \, d\hat{\pi}_{\varepsilon_2} + \varepsilon_2 \text{KL}(\hat{\pi}_{\varepsilon_2}, \mu \otimes \nu) \\ &\geq \int_{\mathbb{X}^2} c \, d\hat{\pi}_{\varepsilon_2} + \varepsilon_1 \text{KL}(\hat{\pi}_{\varepsilon_2}, \mu \otimes \nu) \geq \text{OT}_{\varepsilon_1}(\mu, \nu). \end{aligned}$$

□

Note that in special cases, e.g., for absolutely continuous measures, see [56, 187], it is possible to show convergence of the optimal solutions $\hat{\pi}_\varepsilon$ to an optimal solution of $\text{OT}(\mu, \nu)$ as $\varepsilon \rightarrow 0$. However, we are not aware of a fully general result. An extension of entropy regularization to unbalanced OT is discussed in [69].

Originally, entropic regularization was proposed in [81] for *discrete* probability measures with the negative entropy E , see also [227],

$$\widetilde{\text{OT}}_\varepsilon(\mu, \nu) := \min_{\pi \in \Pi(\mu, \nu)} \left\{ \int_{\mathbb{X}^2} c \, d\pi + \varepsilon E(\pi) \right\}, \quad E(\pi) := \sum_{i,j=1}^n \log(p_{ij}) p_{ij} = \text{KL}(\pi, \lambda \otimes \lambda),$$

where λ denotes the counting measure. For $\pi \in \Pi(\mu, \nu)$ it is easy to check that

$$E(\pi) = \text{KL}(\pi, \mu \otimes \nu) + \sum_{i,j=1}^n \log(\mu_i \nu_j) \mu_i \nu_j = \text{KL}(\pi, \mu \otimes \nu) + \text{KL}(\mu \otimes \nu, \lambda \otimes \lambda),$$

i.e., the minimizers are independent of the chosen regularization. For non-discrete measures, special care is necessary as the following remark shows.

Remark 6.7. (KL($\pi, \mu \otimes \nu$) versus $E(\pi)$ regularization) Since the entropy is only defined for measures with densities, we consider compact sets $\mathbb{X} \subset \mathbb{R}^d$ equipped with the normalized Lebesgue measure λ and $\mu, \nu \ll \lambda$ with densities $\sigma_\mu, \sigma_\nu \in L^1(\mathbb{X})$. For $\pi \ll \lambda \otimes \lambda$ with density σ_π the entropy is defined by

$$E(\pi) = \int_{\mathbb{X}^2} \log(\sigma_\pi) \sigma_\pi d(\lambda \otimes \lambda) = \text{KL}(\pi, \lambda \otimes \lambda).$$

Note that for any $\pi \in \Pi(\mu, \nu)$ we have

$$\pi \ll \mu \otimes \nu \iff \pi \ll \lambda \otimes \lambda,$$

where the right implication follows directly and the left one can be seen as follows: If $\pi \ll \lambda \otimes \lambda$ with density $\sigma_\pi \in L^1(\mathbb{X} \times \mathbb{X})$, then

$$0 = \int_{\{z \in \mathbb{X} : \sigma_\mu(z) = 0\}} \int_{\mathbb{X}} \sigma_\pi(x, y) dy dx.$$

Consequently, we get $\sigma_\pi(x, y) = 0$ a.e. on $\{z \in \mathbb{X} : \sigma_\mu(z) = 0\} \times \mathbb{X}$ (for any representative of σ_μ). The same reasoning is applicable to $\mathbb{X} \times \{z \in \mathbb{X} : \sigma_\nu(z) = 0\}$. Thus,

$$\pi = \sigma_\pi (\lambda \otimes \lambda) = \frac{\sigma_\pi(x, y)}{\sigma_\mu(x)\sigma_\nu(y)} (\mu \otimes \nu),$$

where the quotient is defined as zero if σ_μ or σ_ν vanish. Hence, the left implication also holds true.

If $\text{KL}(\mu \otimes \nu, \lambda \otimes \lambda) < \infty$, we conclude for any $\pi \ll \lambda \otimes \lambda$ with $\pi \in \Pi(\mu, \nu)$ that the following expressions are well-defined

$$\begin{aligned} & \text{KL}(\pi, \lambda \otimes \lambda) - \text{KL}(\mu \otimes \nu, \lambda \otimes \lambda) \\ &= \int_{\mathbb{X}^2} \log(\sigma_\pi) d\pi - \int_{\mathbb{X}^2} \log\left(\frac{d(\mu \otimes \nu)}{d(\lambda \otimes \lambda)}\right) d(\mu \otimes \nu) \\ &= \text{KL}(\pi, \mu \otimes \nu) + \int_{\mathbb{X}^2} \log(\sigma_\mu(x)\sigma_\nu(y)) d\pi(x, y) - \int_{\mathbb{X}^2} \log(\sigma_\mu(x)\sigma_\nu(y)) d\mu(x) d\nu(y) \\ &= \text{KL}(\pi, \mu \otimes \nu). \end{aligned}$$

Consequently, in this case we also have $\widetilde{\text{OT}}_\varepsilon(\mu, \nu) = \text{OT}_\varepsilon(\mu, \nu) + \varepsilon \text{KL}(\mu \otimes \nu, \lambda \otimes \lambda)$. The crux is the condition $\text{KL}(\mu \otimes \nu, \lambda \otimes \lambda) < \infty$, which is equivalent to μ, ν having finite entropy, i.e., σ_μ, σ_ν are in a so-called Orlicz space $L \log L$ [209]. The authors in [74] considered the entropy as regularization (with continuous cost function) and pointed out that $\widetilde{\text{OT}}_\varepsilon(\mu, \nu)$ admits a (finite) minimizer exactly in this case. However, we have seen that we can avoid this existence trouble if we regularize with $\text{KL}(\pi, \mu \otimes \nu)$ instead, which therefore seems to be a more natural choice. A comparison of the settings and a more general existence discussion based on merely continuous cost functions can be also found in [90].

Another possibility is to use quadratic regularization instead, see [189] for more details. In connection with discrepancies, we are especially interested in the limiting case $\varepsilon \rightarrow \infty$. The next proposition is basically known, see [82, 108]. However, we have not found it in this generality in the literature.

Proposition 6.8.

i) It holds $\lim_{\varepsilon \rightarrow \infty} \text{OT}_\varepsilon(\mu, \nu) = \text{OT}_\infty(\mu, \nu)$, where

$$\text{OT}_\infty(\mu, \nu) := \int_{\mathbb{X}^2} c \, d(\mu \otimes \nu).$$

ii) It holds $\lim_{\varepsilon \rightarrow 0} \text{OT}_\varepsilon(\mu, \nu) = \text{OT}(\mu, \nu)$.

Proof. i) For $\pi = \mu \otimes \nu$, we have

$$\int_{\mathbb{X}^2} c \, d\pi + \varepsilon \text{KL}(\pi, \mu \otimes \nu) = \text{OT}_\infty(\mu, \nu)$$

and consequently $\limsup_{\varepsilon \rightarrow \infty} \text{OT}_\varepsilon(\mu, \nu) \leq \text{OT}_\infty(\mu, \nu)$. In particular, the optimal transport plan $\hat{\pi}_\varepsilon$ satisfies $\limsup_{\varepsilon \rightarrow \infty} \varepsilon \text{KL}(\hat{\pi}_\varepsilon, \mu \otimes \nu) \leq \text{OT}_\infty(\mu, \nu)$. Since KL is weakly lsc, we conclude that the sequence of minimizers $\hat{\pi}_\varepsilon$ satisfies $\hat{\pi}_\varepsilon \rightharpoonup \mu \otimes \nu$ as $\varepsilon \rightarrow \infty$. Hence, we obtain the desired result from

$$\begin{aligned} \liminf_{\varepsilon \rightarrow \infty} \text{OT}_\varepsilon(\mu, \nu) &= \liminf_{\varepsilon \rightarrow \infty} \int_{\mathbb{X}^2} c \, d\hat{\pi}_\varepsilon + \varepsilon \text{KL}(\hat{\pi}_\varepsilon, \mu \otimes \nu) \\ &\geq \liminf_{\varepsilon \rightarrow \infty} \int_{\mathbb{X}^2} c \, d\hat{\pi}_\varepsilon = \text{OT}_\infty(\mu, \nu). \end{aligned}$$

ii) This part is more involved and follows from Proposition 6.13 ii). \square

Similar as OT in (6.22), its regularized version OT_ε can be written in dual form, see [69, 74].

Proposition 6.9. *The (pre-)dual problem of OT_ε is given by*

$$\begin{aligned} \text{OT}_\varepsilon(\mu, \nu) &= \sup_{(\varphi, \psi) \in C(\mathbb{X})^2} \left\{ \int_{\mathbb{X}} \varphi \, d\mu + \int_{\mathbb{X}} \psi \, d\nu \right. \\ &\quad \left. - \varepsilon \int_{\mathbb{X}^2} \exp\left(\frac{\varphi(x) + \psi(y) - c(x, y)}{\varepsilon}\right) - 1 \, d(\mu \otimes \nu) \right\}. \end{aligned} \quad (6.24)$$

If optimal dual solutions $\hat{\varphi}_\varepsilon$ and $\hat{\psi}_\varepsilon$ exist, they are related to the optimal transport plan $\hat{\pi}_\varepsilon$ by

$$\hat{\pi}_\varepsilon = \exp\left(\frac{\hat{\varphi}_\varepsilon(x) + \hat{\psi}_\varepsilon(y) - c(x, y)}{\varepsilon}\right) \mu \otimes \nu. \quad (6.25)$$

Proof. Let us consider $F \in \Gamma_0(C(\mathbb{X})^2)$, $G \in \Gamma_0(C(\mathbb{X}^2))$ with Fenchel conjugates $F^* \in \Gamma_0(\mathcal{M}(\mathbb{X})^2)$, $G^* \in \Gamma_0(\mathcal{M}(\mathbb{X}^2))$ together with a linear bounded operator $A: C(\mathbb{X})^2 \rightarrow C(\mathbb{X}^2)$ with adjoint operator $A^*: \mathcal{M}(\mathbb{X}^2) \rightarrow \mathcal{M}(\mathbb{X})^2$ defined by

$$\begin{aligned} F(\varphi, \psi) &= \int_{\mathbb{X}} \varphi \, d\mu + \int_{\mathbb{X}} \psi \, d\nu, \\ G(\varphi) &= \varepsilon \int_{\mathbb{X}^2} \exp\left(\frac{\varphi - c}{\varepsilon}\right) - 1 \, d(\mu \otimes \nu), \\ A(\varphi, \psi)(x, y) &= \varphi(x) + \psi(y). \end{aligned}$$

Then, (6.24) has the form of the left-hand side in (6.2). Incorporating (6.7), we get

$$G^*(\pi) = \int_{\mathbb{X}} c \, d\pi + \varepsilon \text{KL}(\pi, \mu \otimes \nu).$$

Using the indicator function ι_C defined by $\iota_C(x) := 0$ for $x \in C$ and $\iota_C(x) := +\infty$ otherwise, we have

$$\begin{aligned} F^*(A^*\pi) &= \sup_{(\varphi, \psi) \in C(\mathbb{X})^2} \langle A^*\pi, (\varphi, \psi) \rangle - \int_{\mathbb{X}} \varphi \, d\mu - \int_{\mathbb{X}} \psi \, d\nu \\ &= \sup_{(\varphi, \psi) \in C(\mathbb{X})^2} \langle \pi, \varphi(x) + \psi(y) \rangle - \int_{\mathbb{X}} \varphi \, d\mu - \int_{\mathbb{X}} \psi \, d\nu \\ &= \iota_{\Pi(\mu, \nu)}(\pi). \end{aligned}$$

Now, the duality relation follows from (6.2).

If the optimal solution $(\hat{\varphi}_\varepsilon, \hat{\psi}_\varepsilon)$ exists, we can apply (6.3) and (6.8) to obtain

$$\hat{\varphi}_\varepsilon(x) + \hat{\psi}_\varepsilon(y) = c + \log \left(\frac{d\hat{\pi}_\varepsilon}{d(\mu \otimes \nu)} \right),$$

which yields (6.25). \square

Remark 6.10. Using the Tietze extension theorem, we could also replace the space $C(\mathbb{X})^2$ by $C(\text{supp}(\mu)) \times C(\text{supp}(\nu))$.

Note that the last term in (6.24) is a smoothed version of the associated constraint $\varphi(x) + \psi(y) \leq c(x, y)$ appearing in (6.22). Clearly, the values of φ and ψ are only relevant on $\text{supp}(\mu)$ and $\text{supp}(\nu)$, respectively. Further, for any $\varphi, \psi \in C(\mathbb{X})$ and $C \in \mathbb{R}$, the potentials $\varphi + C, \psi - C$ realize the same value in (6.24).

For fixed φ or ψ , the corresponding maximizing potentials in (6.24) are given by

$$\hat{\psi}_{\varphi, \varepsilon} = T_{\mu, \varepsilon}(\varphi) \text{ on } \text{supp}(\nu) \quad \text{and} \quad \hat{\varphi}_{\psi, \varepsilon} = T_{\nu, \varepsilon}(\psi) \text{ on } \text{supp}(\mu),$$

respectively. Here, $T_{\mu, \varepsilon}: C(\mathbb{X}) \rightarrow C(\mathbb{X})$ is defined as

$$T_{\mu, \varepsilon}(\varphi)(x) := -\varepsilon \log \left(\int_{\mathbb{X}} \exp \left(\frac{\varphi(y) - c(x, y)}{\varepsilon} \right) d\mu(y) \right). \quad (6.26)$$

Therefore, any pair of optimal potentials $\hat{\varphi}_\varepsilon$ and $\hat{\psi}_\varepsilon$ must satisfy

$$\hat{\psi}_\varepsilon = T_{\mu, \varepsilon}(\hat{\varphi}_\varepsilon) \text{ on } \text{supp}(\nu), \quad \hat{\varphi}_\varepsilon = T_{\nu, \varepsilon}(\hat{\psi}_\varepsilon) \text{ on } \text{supp}(\mu).$$

For every $\varphi \in C(\mathbb{X})$ and $C \in \mathbb{R}$, it holds $T_{\mu, \varepsilon}(\varphi + C) = T_{\mu, \varepsilon}(\varphi) + C$. Hence, $T_{\mu, \varepsilon}$ can be interpreted as an operator on the quotient space $C(\mathbb{X})/\mathbb{R}$, where $f_1, f_2 \in C(\mathbb{X})$ are equivalent if they differ by a real constant. This space can be equipped with the *oscillation norm*

$$\|f\|_{\circ, \infty} := \frac{1}{2}(\max f - \min f)$$

and for $f \in C(\mathbb{X})/\mathbb{R}$ there is a representative $\bar{f} \in C(\mathbb{X})$ with $\|f\|_{\circ, \infty} = \|\bar{f}\|_\infty$. Finally, it is possible to restrict the domain of $T_{\mu, \varepsilon}$ to $C(\text{supp}(\mu))$ and $C(\text{supp}(\mu))/\mathbb{R}$, respectively. This interpretation is useful for showing convergence of the Sinkhorn algorithm. In the next lemma, we collect a few properties of $T_{\mu, \varepsilon}$, see also [122, 271].

Lemma 6.11.

i) For any measure $\mu \in P(\mathbb{X})$, $\varepsilon > 0$ and $\varphi \in C(\mathbb{X})$, the function $T_{\mu,\varepsilon}(\varphi) \in C(\mathbb{X})$ has the same Lipschitz constant as c and satisfies

$$T_{\mu,\varepsilon}(\varphi)(x) \in \left[\min_{y \in \text{supp}(\mu)} c(x, y) - \varphi(y), \max_{y \in \text{supp}(\mu)} c(x, y) - \varphi(y) \right]. \quad (6.27)$$

ii) For fixed $\mu \in \mathcal{P}(\mathbb{X})$, the operator $T_{\mu,\varepsilon}: C(\text{supp}(\mu)) \rightarrow C(\mathbb{X})$ is 1-Lipschitz. Additionally, the operator $T_{\mu,\varepsilon}: C(\text{supp}(\mu))/\mathbb{R} \rightarrow C(\mathbb{X})/\mathbb{R}$ is κ -Lipschitz with $\kappa < 1$.

Proof. i) For $x_1, x_2 \in \mathbb{X}$ (possibly changing the naming of the variables) we obtain

$$\begin{aligned} & |T_{\mu,\varepsilon}(\varphi)(x_1) - T_{\mu,\varepsilon}(\varphi)(x_2)| \\ &= \varepsilon \left| \log \int_{\mathbb{X}} \exp\left(\frac{\varphi(y) - c(x_2, y)}{\varepsilon}\right) d\mu(y) - \log \int_{\mathbb{X}} \exp\left(\frac{\varphi(y) - c(x_1, y)}{\varepsilon}\right) d\mu(y) \right| \\ &= \varepsilon \log \left(\int_{\mathbb{X}} \exp\left(\frac{\varphi(y) - c(x_2, y)}{\varepsilon}\right) d\mu(y) / \int_{\mathbb{X}} \exp\left(\frac{\varphi(y) - c(x_1, y)}{\varepsilon}\right) d\mu(y) \right). \end{aligned}$$

Incorporating the L -Lipschitz continuity of c , we get

$$\exp\left(\frac{c(x_1, y) - c(x_2, y)}{\varepsilon}\right) \leq \exp\left(\frac{|c(x_1, y) - c(x_2, y)|}{\varepsilon}\right) \leq \exp\left(\frac{L}{\varepsilon}|x_1 - x_2|\right),$$

so that

$$\int_{\mathbb{X}} \exp\left(\frac{\varphi(y) - c(x_2, y)}{\varepsilon}\right) d\mu(y) \leq \exp\left(\frac{L}{\varepsilon}|x_1 - x_2|\right) \int_{\mathbb{X}} \exp\left(\frac{\varphi(y) - c(x_1, y)}{\varepsilon}\right) d\mu(y).$$

Thus, $T_{\mu,\varepsilon}(\varphi)$ is Lipschitz continuous

$$|T_{\mu,\varepsilon}(\varphi)(x_1) - T_{\mu,\varepsilon}(\varphi)(x_2)| \leq \varepsilon \log\left(\exp\left(\frac{L}{\varepsilon}|x_1 - x_2|\right)\right) = L|x_1 - x_2|.$$

Finally, (6.27) follows directly from (6.26) since μ is a probability measure.

ii) For any $x \in \mathbb{X}$ and $\varphi_1, \varphi_2 \in C(\text{supp}(\mu))$ it holds

$$\begin{aligned} T_{\mu,\varepsilon}(\varphi_1)(x) - T_{\mu,\varepsilon}(\varphi_2)(x) &= \int_0^1 \frac{d}{dt} T_{\mu,\varepsilon}(\varphi_1 + t(\varphi_2 - \varphi_1))(x) dt \\ &= \int_0^1 \int_{\mathbb{X}} (\varphi_1(z) - \varphi_2(z)) \rho_{t,x}(z) d\mu(z) dt \end{aligned} \quad (6.28)$$

with

$$\rho_{t,x} := \frac{\exp((t\varphi_2 + (1-t)\varphi_1 - c(x, \cdot))/\varepsilon)}{\int_{\mathbb{X}} \exp((t\varphi_2(z) + (1-t)\varphi_1(z) - c(x, z))/\varepsilon) d\mu(z)}.$$

This directly implies

$$\|T_{\mu,\varepsilon}(\varphi_1) - T_{\mu,\varepsilon}(\varphi_2)\|_{\infty} \leq \sup_{x \in \text{supp}(\mu)} \int_0^1 \int_{\mathbb{X}} |\varphi_1(z) - \varphi_2(z)| \rho_{t,x}(z) d\mu(z) dt \leq \|\varphi_1 - \varphi_2\|_{\infty}.$$

In order to show the second claim, we choose representatives φ_1 and φ_2 such that $\|\varphi_1 - \varphi_2\|_\infty = \|\varphi_1 - \varphi_2\|_{\circ, \infty}$. Given $x, y \in \mathbb{X}$, we conclude using (6.28) that

$$\begin{aligned} & \frac{1}{2} (T_{\mu, \varepsilon}(\varphi_1)(x) - T_{\mu, \varepsilon}(\varphi_2)(x) - T_{\mu, \varepsilon}(\varphi_1)(y) + T_{\mu, \varepsilon}(\varphi_2)(y)) \\ &= \frac{1}{2} \int_0^1 \int_{\mathbb{X}} (\varphi_1(z) - \varphi_2(z)) (\rho_{t,x}(z) - \rho_{t,y}(z)) \, d\mu(z) \, dt \\ &\leq \|\varphi_1 - \varphi_2\|_{\circ, \infty} \frac{1}{2} \int_0^1 \|\rho_{t,x} - \rho_{t,y}\|_{L^1(\mu)} \, dt. \end{aligned} \tag{6.29}$$

For all $z \in \mathbb{X}$ with $p_{t,x}(z) \geq p_{t,y}(z)$, we can estimate

$$p_{t,x}(z) - p_{t,y}(z) \leq p_{t,x}(z) (1 - \exp(-2L \operatorname{diam}(\mathbb{X})/\varepsilon))$$

and similarly for $z \in \mathbb{X}$ with $p_{t,y}(z) \geq p_{t,x}(z)$. Hence, we obtain

$$\begin{aligned} \|\rho_{t,x} - \rho_{t,y}\|_{L^1(\mu)} &\leq \int_{\mathbb{X}} (1_{\{p_{t,x} \geq p_{t,y}\}} p_{t,x} + 1_{\{p_{t,y} > p_{t,x}\}} p_{t,y}) (1 - \exp(-2L \operatorname{diam}(\mathbb{X})/\varepsilon)) \, d\mu \\ &\leq 2(1 - \exp(-2L \operatorname{diam}(\mathbb{X})/\varepsilon)). \end{aligned}$$

Finally, inserting this into (6.29) implies

$$\|T_{\mu, \varepsilon}(\varphi_1) - T_{\mu, \varepsilon}(\varphi_2)\|_{\circ, \infty} \leq (1 - \exp(-2L \operatorname{diam}(\mathbb{X})/\varepsilon)) \|\varphi_1 - \varphi_2\|_{\circ, \infty}.$$

□

Now, we are able to prove existence of an optimal solution $(\hat{\varphi}_\varepsilon, \hat{\psi}_\varepsilon)$.

Proposition 6.12. *The optimal potentials $\hat{\varphi}_\varepsilon, \hat{\psi}_\varepsilon \in C(\mathbb{X})$ exist and are unique on $\operatorname{supp}(\mu)$ and $\operatorname{supp}(\nu)$, respectively (up to the additive constant).*

Proof. Let $\varphi_n, \psi_n \in C(\mathbb{X})$ be maximizing sequences of (6.24). Using the operator $T_{\mu, \varepsilon}$, these can be replaced by

$$\tilde{\psi}_n = T_{\mu, \varepsilon}(\varphi_n) \quad \text{and} \quad \tilde{\varphi}_n = T_{\nu, \varepsilon} \circ T_{\mu, \varepsilon}(\varphi_n),$$

which are Lipschitz continuous with the same constant as c by Lemma 6.11 i) and therefore uniformly equi-continuous. Next, we can choose some $x_0 \in \operatorname{supp}(\mu)$ and w.l.o.g. assume $\tilde{\psi}_n(x_0) = 0$. Due to the uniform Lipschitz continuity, the potentials $\tilde{\psi}_n$ are uniformly bounded and by (6.27) the same holds true for $\tilde{\varphi}_n$. Now, the theorem of Arzelà–Ascoli implies that both sequences contain convergent subsequences. Since the functional in (6.24) is continuous, we can readily infer the existence of optimal potentials $\hat{\varphi}_\varepsilon, \hat{\psi}_\varepsilon \in C(\mathbb{X})$. Due to the uniqueness of $\hat{\pi}_\varepsilon$, (6.25) implies that $\hat{\varphi}_\varepsilon|_{\operatorname{supp}(\mu)}$ and $\hat{\psi}_\varepsilon|_{\operatorname{supp}(\nu)}$ are uniquely determined up to an additive constant. □

Combining the optimality condition (6.26) and (6.24), we directly obtain for any pair of optimal solutions

$$\text{OT}_\varepsilon(\mu, \nu) = \int_{\mathbb{X}} \hat{\varphi}_\varepsilon \, d\mu + \int_{\mathbb{X}} \hat{\psi}_\varepsilon \, d\nu. \quad (6.30)$$

Adding, e.g., the additional constraint

$$\int_{\mathbb{X}} \varphi \, d\mu = \frac{1}{2} \text{OT}_\infty(\mu, \nu), \quad (6.31)$$

the restricted optimal potentials $\hat{\varphi}_\varepsilon|_{\text{supp}(\mu)}$ and $\hat{\psi}_\varepsilon|_{\text{supp}(\nu)}$ are unique. The next proposition investigates the limits of the potentials as $\varepsilon \rightarrow 0$ and $\varepsilon \rightarrow \infty$.

Proposition 6.13.

i) If (6.31) is satisfied, the restricted potentials $\hat{\varphi}_\varepsilon|_{\text{supp}(\mu)}$ and $\hat{\psi}_\varepsilon|_{\text{supp}(\nu)}$ converge uniformly for $\varepsilon \rightarrow \infty$ to

$$\begin{aligned} \hat{\varphi}_\infty(x) &= \int_{\mathbb{X}} c(x, y) \, d\nu(y) - \frac{1}{2} \text{OT}_\infty(\mu, \nu), \\ \hat{\psi}_\infty(y) &= \int_{\mathbb{X}} c(x, y) \, d\mu(x) - \frac{1}{2} \text{OT}_\infty(\mu, \nu), \end{aligned}$$

respectively.

ii) For $\varepsilon \rightarrow 0$ every accumulation point of $(\hat{\varphi}_\varepsilon|_{\text{supp}(\mu)}, \hat{\psi}_\varepsilon|_{\text{supp}(\nu)})$ can be extended to an optimal dual pair for $\text{OT}(\mu, \nu)$ satisfying (6.31). In particular, $\lim_{\varepsilon \rightarrow 0} \text{OT}_\varepsilon(\mu, \nu) = \text{OT}(\mu, \nu)$.

Proof. i) Since \mathbb{X} is bounded, the Lipschitz continuity of the potentials together with (6.31) implies that all $\hat{\varphi}_\varepsilon$ are uniformly bounded on $\text{supp}(\mu)$. Then, we conclude for $y \in \text{supp}(\nu)$ using l'Hôpital's rule, dominated convergence and (6.31) that

$$\begin{aligned} & \lim_{\varepsilon \rightarrow \infty} \hat{\psi}_\varepsilon(y) \\ &= \lim_{\varepsilon \rightarrow \infty} - \frac{\int_{\mathbb{X}} (\hat{\varphi}_\varepsilon(x) - c(x, y)) \exp((\hat{\varphi}_\varepsilon(x) - c(x, y))/\varepsilon) \, d\mu(x)}{\int_{\mathbb{X}} \exp((\hat{\varphi}_\varepsilon(x) - c(x, y))/\varepsilon) \, d\mu(x)} \\ &= \lim_{\varepsilon \rightarrow \infty} \int_{\mathbb{X}} c(x, y) \exp((\hat{\varphi}_\varepsilon(x) - c(x, y))/\varepsilon) - \hat{\varphi}_\varepsilon(x) \exp((\hat{\varphi}_\varepsilon(x) - c(x, y))/\varepsilon) \, d\mu(x) \\ &= \int_{\mathbb{X}} c(x, y) \, d\mu(x) - \lim_{\varepsilon \rightarrow \infty} \int_{\mathbb{X}} \hat{\varphi}_\varepsilon(x) \left(\exp((\hat{\varphi}_\varepsilon(x) - c(x, y))/\varepsilon) - 1 \right) + \hat{\varphi}_\varepsilon(x) \, d\mu(x) \\ &= \int_{\mathbb{X}} c(x, y) \, d\mu(x) - \frac{1}{2} \text{OT}_\infty(\mu, \nu). \end{aligned}$$

Again, a similar reasoning, incorporating (6.27), can be applied for $\hat{\varphi}_\varepsilon$. Finally, note that pointwise convergence of uniformly Lipschitz continuous functions on compact sets implies uniform convergence.

ii) By continuity of the integral, we can directly infer that (6.31) is satisfied for any accumulation point. Note that for any fixed $\varphi \in C(\mathbb{X})$, $x \in \mathbb{X}$ and $\varepsilon \rightarrow 0$ it holds

$$T_{\mu,\varepsilon}(\varphi)(x) \rightarrow \min_{y \in \text{supp}(\mu)} c(x, y) - \varphi(y),$$

see [108, Prop. 9], which by uniform Lipschitz continuity of $T_{\mu,\varepsilon}(\varphi)$ directly implies the convergence in $C(\mathbb{X})$. Let $\{(\hat{\varphi}_{\varepsilon_j}, \hat{\psi}_{\varepsilon_j})\}_j$ be a subsequence converging to $(\hat{\varphi}_0, \hat{\psi}_0) \in C(\text{supp}(\mu)) \times C(\text{supp}(\nu))$. Then, we have

$$\begin{aligned} \hat{\psi}_0 &= \lim_{j \rightarrow \infty} \hat{\psi}_{\varepsilon_j} = \lim_{j \rightarrow \infty} T_{\mu,\varepsilon_j}(\hat{\varphi}_{\varepsilon_j}) \\ &= \lim_{j \rightarrow \infty} \left(T_{\mu,\varepsilon_j}(\hat{\varphi}_{\varepsilon_j}) - T_{\mu,\varepsilon_j}(\hat{\varphi}_0) + T_{\mu,\varepsilon_j}(\hat{\varphi}_0) \right). \end{aligned}$$

By Lemma 6.11 ii), it holds

$$\|T_{\mu,\varepsilon_j}(\hat{\varphi}_{\varepsilon_j}) - T_{\mu,\varepsilon_j}(\hat{\varphi}_0)\|_\infty \leq \|\hat{\varphi}_{\varepsilon_j} - \hat{\varphi}_0\|_\infty$$

and we conclude

$$\hat{\psi}_0 = \lim_{j \rightarrow \infty} T_{\mu,\varepsilon_j}(\hat{\varphi}_0) = \min_{y \in \text{supp}(\mu)} c(\cdot, y) - \hat{\varphi}_0(y).$$

Similarly, we get

$$\hat{\varphi}_0 = \min_{y \in \text{supp}(\nu)} c(\cdot, y) - \hat{\psi}_0(y).$$

Thus, $(\hat{\varphi}_0, \hat{\psi}_0)$ can be extended to a feasible point in $C(\mathbb{X})^2$ of (6.22) by Remark 6.5.

Due to continuity of (6.30) and since OT_ε is monotone in ε , this implies

$$\lim_{j \rightarrow \infty} \text{OT}_{\varepsilon_j}(\mu, \nu) = \int_{\mathbb{X}} \hat{\varphi}_0 d\mu + \int_{\mathbb{X}} \hat{\psi}_0 d\nu \leq \text{OT}(\mu, \nu) \leq \lim_{j \rightarrow \infty} \text{OT}_{\varepsilon_j}(\mu, \nu).$$

Hence, the extended potentials are optimal for (6.22). Since the subsequence choice was arbitrary, this also shows Proposition 6.8 ii). \square

So far we cannot show the convergence of the potentials for $\varepsilon \rightarrow 0$ for the fully general case. Essentially, our approach would require that all $T_{\mu,\varepsilon}$ are contractive with a uniform constant $\beta < 1$, which is not the case. Note that if we assume that the unregularized potentials satisfying (6.31) are unique, then ii) directly implies convergence of the restricted dual potentials, see also [34, Thm. 3.3] and [76]. Nevertheless, we always observed convergence in our numerical examples.

6.6 Sinkhorn divergence

The functional OT_ε is biased, i.e., in general $\min_\nu \text{OT}_\varepsilon(\nu, \mu) \neq \text{OT}_\varepsilon(\mu, \mu)$. Hence, the usage as distance measure is meaningless, which motivates the introduction of the *Sinkhorn divergence*

$$S_\varepsilon(\mu, \nu) = \text{OT}_\varepsilon(\mu, \nu) - \frac{1}{2} \text{OT}_\varepsilon(\mu, \mu) - \frac{1}{2} \text{OT}_\varepsilon(\nu, \nu).$$

Indeed, it was shown that S_ε is non-negative, bi-convex and metrizes the convergence in law under mild assumptions [108]. Clearly, we have $S_0 = \text{OT}$. By (6.14) and Proposition 6.13, we obtain the following corollary.

Corollary 6.14. *Assume that $K \in C(\mathbb{X} \times \mathbb{X})$ is symmetric and positive definite. Set $c(x, y) := -K(x, y)$. Then, it holds $S_\infty(\mu, \nu) = \frac{1}{2} \mathcal{D}_K^2(\mu, \nu)$ and the optimal dual potential $\hat{\varphi}_K$ realizing $\mathcal{D}_K(\mu, \nu)$ is related to the uniform limits $\hat{\varphi}_\infty, \hat{\psi}_\infty$ of $\hat{\varphi}_\varepsilon, \hat{\psi}_\varepsilon$ in $\text{OT}_\varepsilon(\mu, \nu)$ with constraint (6.31) by*

$$\hat{\varphi}_K = \frac{\hat{\varphi}_\infty - \hat{\psi}_\infty}{\|\hat{\varphi}_\infty - \hat{\psi}_\infty\|_{H_K(\mathbb{X})}}.$$

Note that (6.12) already implies that for the chosen c it holds $\hat{\varphi}_\infty, \hat{\psi}_\infty \in H_K(\mathbb{X})$. By Corollary 6.14, we have for $c(x, y) := -K(x, y)$ that $S_\infty(\mu, \nu) = \frac{1}{2} \mathcal{D}_K^2(\mu, \nu)$ if $K \in C(\mathbb{X} \times \mathbb{X})$ is symmetric, positive definite. For the cost $c(x, y) = \|x - y\|^p$ of the classical p -Wasserstein distance, we have already seen in Section 6.3 that $K(x, y) = -c(x, y)$ is not positive definite. However, at least for $p = 1$ the kernel is conditionally positive definite of order 1 and can be tuned by Proposition 6.2 to a positive definite kernel by adding a constant, which neither changes the value of the discrepancy nor of the optimal dual potential. More generally, we have the following corollary.

Corollary 6.15. *Let $K \in C(\mathbb{X} \times \mathbb{X})$ be symmetric, conditionally positive definite of order 1, and let \tilde{K} be the corresponding positive definite kernel in (6.19). Then we have for $c = -\tilde{K}$ that*

$$S_\infty(\mu, \nu) = \frac{1}{2} \mathcal{D}_K^2(\mu, \nu)$$

and for the optimal dual potentials

$$\begin{aligned} \hat{\varphi}_\infty(x) &= \int_{\mathbb{X}} -K(x, y) \, d\nu(y) + \frac{1}{2} \int_{\mathbb{X}^2} K \, d(\mu \otimes \nu) + K(x, \xi) + \frac{1}{2}(c_\nu - c_\mu - K(\xi, \xi)), \\ \hat{\psi}_\infty(y) &= \int_{\mathbb{X}} -K(x, y) \, d\mu(x) + \frac{1}{2} \int_{\mathbb{X}^2} K \, d(\mu \otimes \nu) + K(\xi, y) + \frac{1}{2}(c_\mu - c_\nu - K(\xi, \xi)), \end{aligned}$$

with some fixed $\xi \in \mathbb{X}$ and c_μ, c_ν defined as in (6.20).

Proof. By Corollary 6.14 and Lemma 6.3, we obtain

$$S_\infty(\mu, \nu) = \frac{1}{2} \mathcal{D}_{\tilde{K}}(\mu, \nu)^2 = \frac{1}{2} \mathcal{D}_K(\mu, \nu)^2.$$

The second claim follows by Proposition 6.13. \square

In the following, we want to characterize the convergence of the functional $S_\varepsilon(\cdot, \nu)$ in the limiting cases $\varepsilon \rightarrow 0$ and $\varepsilon \rightarrow \infty$ for fixed $\nu \in \mathcal{P}(\mathbb{X})$. Recall that a sequence $\{F_n\}_{n \in \mathbb{N}}$ of functionals $F_n: \mathcal{P}(\mathbb{X}) \rightarrow (-\infty, +\infty]$ is said to Γ -converge to $F: \mathcal{P}(\mathbb{X}) \rightarrow (-\infty, +\infty]$ if the following two conditions are fulfilled for every $\mu \in \mathcal{P}(\mathbb{X})$, see [44]:

- i) $F(\mu) \leq \liminf_{n \rightarrow \infty} F_n(\mu_n)$ whenever $\mu_n \rightarrow \mu$,
- ii) there is a sequence $\{\mu_n\}_{n \in \mathbb{N}}$ with $\mu_n \rightarrow \mu$ and $\limsup_{n \rightarrow \infty} F_n(\mu_n) \leq F(\mu)$.

The importance of Γ -convergence relies in the fact that every cluster point of minimizers of $\{F_n\}_{n \in \mathbb{N}}$ is a minimizer of F .

Proposition 6.16. *It holds $S_\varepsilon(\cdot, \nu) \xrightarrow{\Gamma} S_\infty(\cdot, \nu)$ as $\varepsilon \rightarrow \infty$ and $S_\varepsilon(\cdot, \nu) \xrightarrow{\Gamma} \text{OT}(\cdot, \nu)$ as $\varepsilon \rightarrow 0$.*

Proof. In both cases the lim sup-inequality follows from Proposition 6.8 by choosing for some fixed $\mu \in \mathcal{P}(\mathbb{X})$ the constant sequence $\mu_n = \mu$, $n \in \mathbb{N}$.

Concerning the lim inf-inequality, we first treat the case $\varepsilon \rightarrow \infty$. Let $\mu_n \rightarrow \mu$ and $\varepsilon_n \rightarrow \infty$. Since $\text{OT}_\varepsilon(\mu, \nu)$ is increasing with ε , it holds for every fixed $m \in \mathbb{N}$ that

$$\begin{aligned} \liminf_{n \rightarrow \infty} S_{\varepsilon_n}(\mu_n, \nu) &= \liminf_{n \rightarrow \infty} (\text{OT}_{\varepsilon_n}(\mu_n, \nu) - \tfrac{1}{2} \text{OT}_{\varepsilon_n}(\mu_n, \mu_n) - \tfrac{1}{2} \text{OT}_{\varepsilon_n}(\nu, \nu)) \\ &\geq \liminf_{n \rightarrow \infty} (\text{OT}_m(\mu_n, \nu) - \tfrac{1}{2} \text{OT}_\infty(\mu_n, \mu_n)) - \tfrac{1}{2} \text{OT}_\infty(\nu, \nu). \end{aligned}$$

Due to the weak continuity of OT_m and OT_∞ , we obtain

$$\liminf_{n \rightarrow \infty} S_{\varepsilon_n}(\mu_n, \nu) \geq \text{OT}_m(\mu, \nu) - \tfrac{1}{2} \text{OT}_\infty(\mu, \mu) - \tfrac{1}{2} \text{OT}_\infty(\nu, \nu).$$

Letting $m \rightarrow \infty$, Proposition 6.8 implies the lim inf-inequality.

Next, we consider $\varepsilon \rightarrow 0$. Let $\mu_n \rightarrow \mu$ and $\varepsilon_n \rightarrow 0$. With similar arguments as above we obtain for any fixed $m \in \mathbb{N}$ that

$$\liminf_{n \rightarrow \infty} S_{\varepsilon_n}(\mu_n, \nu) \geq \liminf_{n \rightarrow \infty} (\text{OT}(\mu_n, \nu) - \tfrac{1}{2} \text{OT}_m(\mu_n, \mu_n)) - \tfrac{1}{2} \text{OT}_m(\nu, \nu)$$

and weak continuity of OT_m and OT implies

$$\liminf_{n \rightarrow \infty} S_{\varepsilon_n}(\mu_n, \nu) \geq \text{OT}(\nu, \mu) - \tfrac{1}{2} \text{OT}_m(\mu, \mu) - \tfrac{1}{2} \text{OT}_m(\nu, \nu).$$

Using again Proposition 6.8, we verify the lim inf-inequality. \square

6.7 Numerical approach and examples

In this section, we discuss the Sinkhorn algorithm for computing OT_ε based on the (pre)-dual form (6.24) and show some numerical examples. As pointed out in Remark 6.10, we can restrict the potentials and the update operator (6.26) to $\text{supp}(\mu)$ and $\text{supp}(\nu)$, respectively. In particular, this restriction results in a discrete problem if both input measures are atomic. For a fixed starting iterate $\psi^{(0)}$, the Sinkhorn algorithm iterates are defined as

$$\begin{aligned} \varphi^{(i+1)} &= T_{\nu, \varepsilon}(\psi^{(i)}), \\ \psi^{(i+1)} &= T_{\mu, \varepsilon}(\varphi^{(i+1)}). \end{aligned}$$




Figure 6.2: Energy values between S_0 and S_∞ for two given measures on $[0, 1]$ and cost function $c(x, y) = |x - y|$. Every blue dot corresponds to the position and the weight of a Dirac measure.

Equivalently, we could rewrite the scheme with just one potential and the following update $\psi^{(i+1)} = T_{\mu, \varepsilon} \circ T_{\nu, \varepsilon}(\psi^{(i)})$. According to Lemma 6.11, the operator $T_{\mu, \varepsilon} \circ T_{\nu, \varepsilon}$ is contractive and hence the Banach fixed point theorem implies that the algorithm converges linearly. Note that it suffices to enforce the additional constraint (6.31) after the Sinkhorn scheme by adding an appropriately chosen constant. Then, the value of $\text{OT}_\varepsilon(\mu, \nu)$ can be computed from the optimal potentials using (6.30). Here, we do not want to go into more detail on implementation issues, since this is not the main scope of this chapter. The numerical examples merely serve as an illustration of the theoretical results. All computations in this section are performed using GEOMLOSS, a publicly available PyTorch implementation for regularized optimal transport. Implementation details can be found in Feydy et al. [108] and in the corresponding GitHub repository.

Demonstration of convergence results. In the following, we present a numerical toy example for illustrating the convergence results from the previous sections. First, we want to verify the interpolation behavior of $S_\varepsilon(\mu, \nu)$ between $\text{OT}(\mu, \nu)$ and $\mathcal{D}_K(\mu, \nu)$. We choose $\mathbb{X} = [0, 1]$, $c(x, y) = |x - y|$ and the probability measures μ and ν depicted in Fig. 6.2. The resulting energies $S_\varepsilon(\mu, \nu)$ in log-scale are plotted in the same figure.

We observe that the values converge as shown in Proposition 6.8 and that the change mainly happens in the interval $[10^{-2}, 10^1]$. Additionally, the numerical results indicate $S_{\varepsilon_1}(\mu, \nu) \leq S_{\varepsilon_2}(\mu, \nu)$ for $\varepsilon_1 > \varepsilon_2$, which is the opposite behavior as for OT_ε where the energies increase, see Lemma 6.6 iii). So far we are not aware of any theoretical result in this direction for $S_\varepsilon(\mu, \nu)$.

Next, we investigate the behavior of the corresponding optimal potentials $\hat{\varphi}_\varepsilon$ and $\hat{\psi}_\varepsilon$ in (6.24). The convergence of the potentials as shown in Proposition 6.13 iii) is numerically verified in Fig. 6.3. Further, the corresponding potentials $\hat{\varphi}_\varepsilon$ are depicted in Fig. 6.4 and the differences $\hat{\varphi}_\varepsilon - \hat{\psi}_\varepsilon$ are depicted in Fig. 6.5. According to Corollary 6.14, this difference is related to the optimal potential $\hat{\varphi}_K$ in the dual formulation of the related discrepancy. The shape of the potentials ranges from something almost linear for small ε to something more quadratic for large ε . Again,

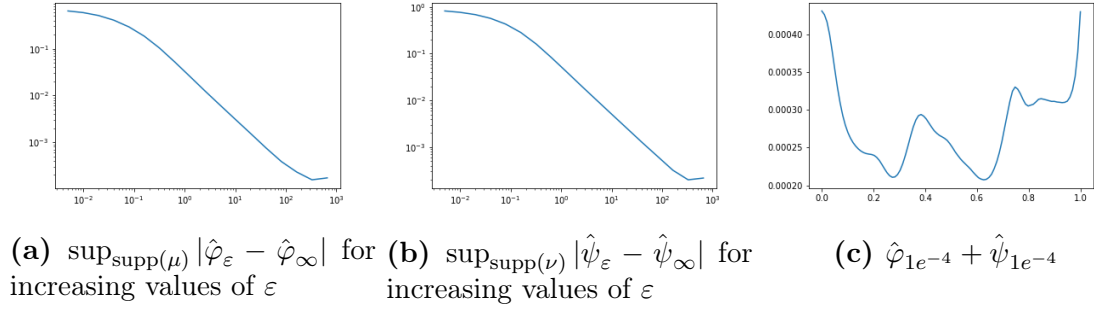


Figure 6.3: Numerical verification of Prop. 6.13 and of $\hat{\psi}_\varepsilon \approx -\hat{\varphi}_\varepsilon$ for small ε .

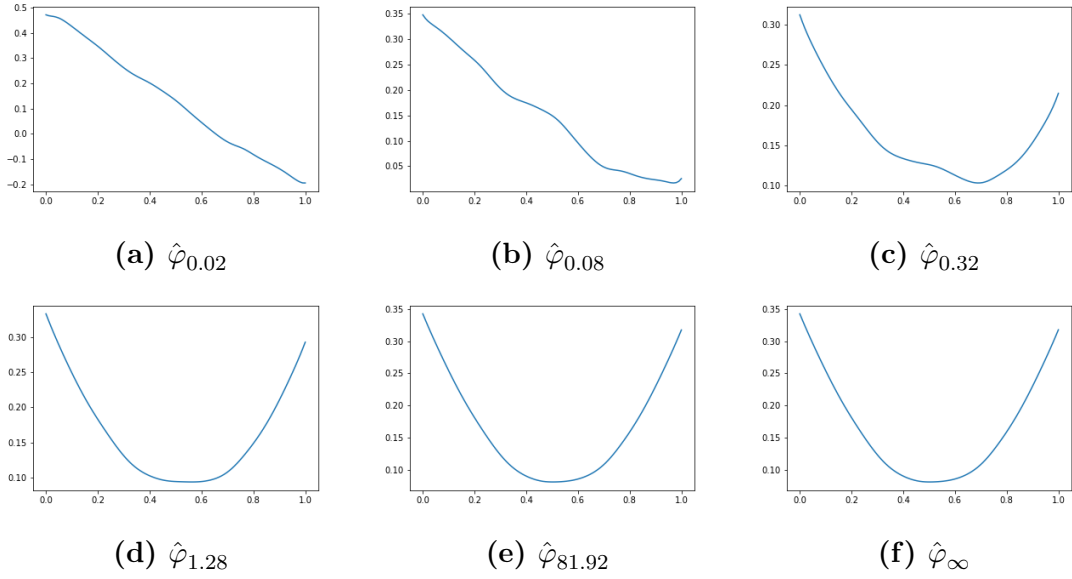


Figure 6.4: Optimal potentials $\hat{\varphi}_\varepsilon$ in $\text{OT}_\varepsilon(\mu, \nu)$ for increasing values of ε .

we observe that the changes mainly happen for ε in the interval $[10^{-2}, 10^1]$ and that numerical instabilities start to occur for $\varepsilon > 10^3$. For small values of ε , we actually observe numerical convergence and that the relation $\hat{\psi}_\varepsilon \approx -\hat{\varphi}_\varepsilon$ holds true, see Fig. 6.3c. This fits the theoretical findings for $W_1(\mu, \nu)$ in Section 6.4.

Dithering results. Now, we want to take a short glimpse at a more involved problem. In the following, we investigate the influence of using S_ε with different values ε as approximation quality measure in dithering. For this purpose, we choose $\mathbb{X} = [-1, 1]^2$, $c(x, y) = |x - y|$ and $\mu = C \exp(-9\|x\|^2/2)(\lambda \otimes \lambda)$, where $C \in \mathbb{R}$ is a normalizing constant. In order to deal with a fully discrete problem, μ is approximated by an atomic measure with 90×90 spikes on a regular grid. Then, we approximate μ with a measure $\nu \in \mathcal{P}_{\text{emp}}^{400}(\mathbb{X})$ (empirical measure with 400 spikes) in terms of the following objective function

$$\min_{\nu \in \mathcal{P}_{\text{emp}}^{400}(\mathbb{X})} S_\varepsilon(\mu, \nu). \quad (6.32)$$

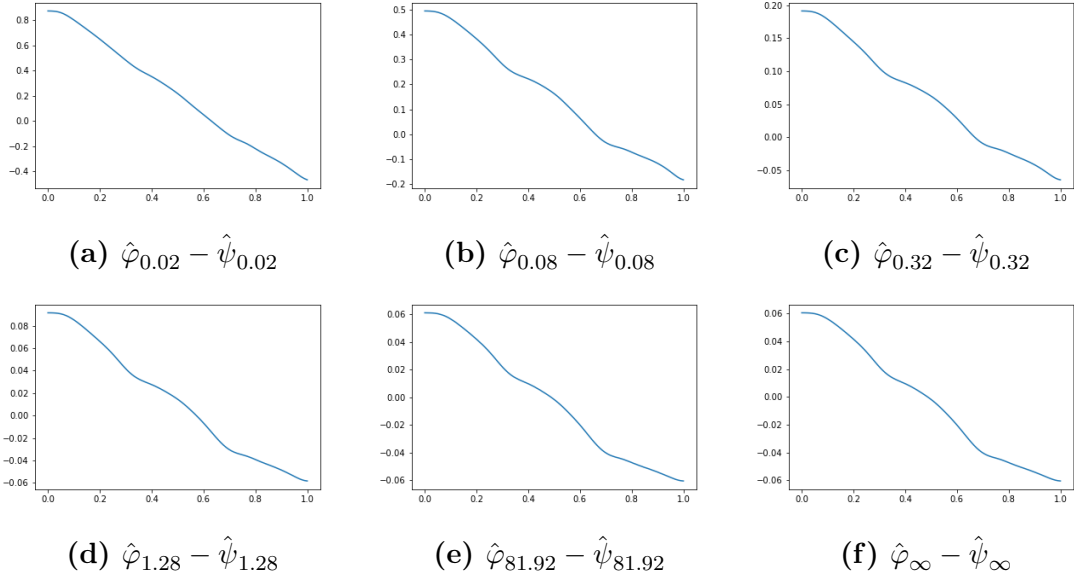


Figure 6.5: Difference $\hat{\varphi}_{\varepsilon} - \hat{\psi}_{\varepsilon}$ of optimal potentials in $\text{OT}_{\varepsilon}(\mu, \nu)$ for increasing ε , where the normalized function $\hat{\varphi}_{\infty} - \hat{\psi}_{\infty}$ coincides with the optimal dual potential $\hat{\varphi}_K$ in the discrepancy by Corollary 6.15.

For solving this problem, we can equivalently minimize over the positions of the equally weighted Dirac spikes in ν . Hence, we need the gradient of S_{ε} with respect to these positions. If $\varepsilon = \infty$, this gradient is given by an analytic expression. Otherwise, we can apply automatic differentiation tools to the Sinkhorn algorithm in order to compute a numerical gradient, see [108] for more details. Here, it is important to ensure high enough numerical precision and to perform enough Sinkhorn iterations. In any case, the gradient serves as input for the L-BFGS-B (Quasi-Newton) method in which the Hessian is approximated in a memory efficient way [54]. The numerical results are depicted in Fig. 6.6, where all examples are iterated to high numerical precision. Numerically, we nicely observe the convergence of $S_{\varepsilon}(\mu, \hat{\nu})$ in the limits $\varepsilon \rightarrow 0$ and $\varepsilon \rightarrow \infty$ as implied from the Γ -convergence result in Proposition 6.16. Visually, the result using Fourier methods is most appealing. Differences could be caused by the different numerical approaches. In particular, the minimization of (6.32) is quite challenging and our applied approach is pretty straight forward without including any special knowledge about the problem. Noteworthy, the Fourier method uses a truncation of $S_{\infty} = \frac{1}{2} \mathcal{D}_K^2$ in the Fourier domain, see (6.16), namely

$$\sum_{k=0}^N \alpha_k |\hat{\mu}_k - \hat{\nu}_k|^2, \quad N := 128$$

as target functional, see [134]. The value of S_{∞} for the Fourier method is slightly larger than the result using optimization of S_{∞} directly. Since the computational cost increases as ε gets smaller, we suggest to choose $\varepsilon \approx 1$ or to directly stick with discrepancies. This also avoids that the approximation rates suffer from the

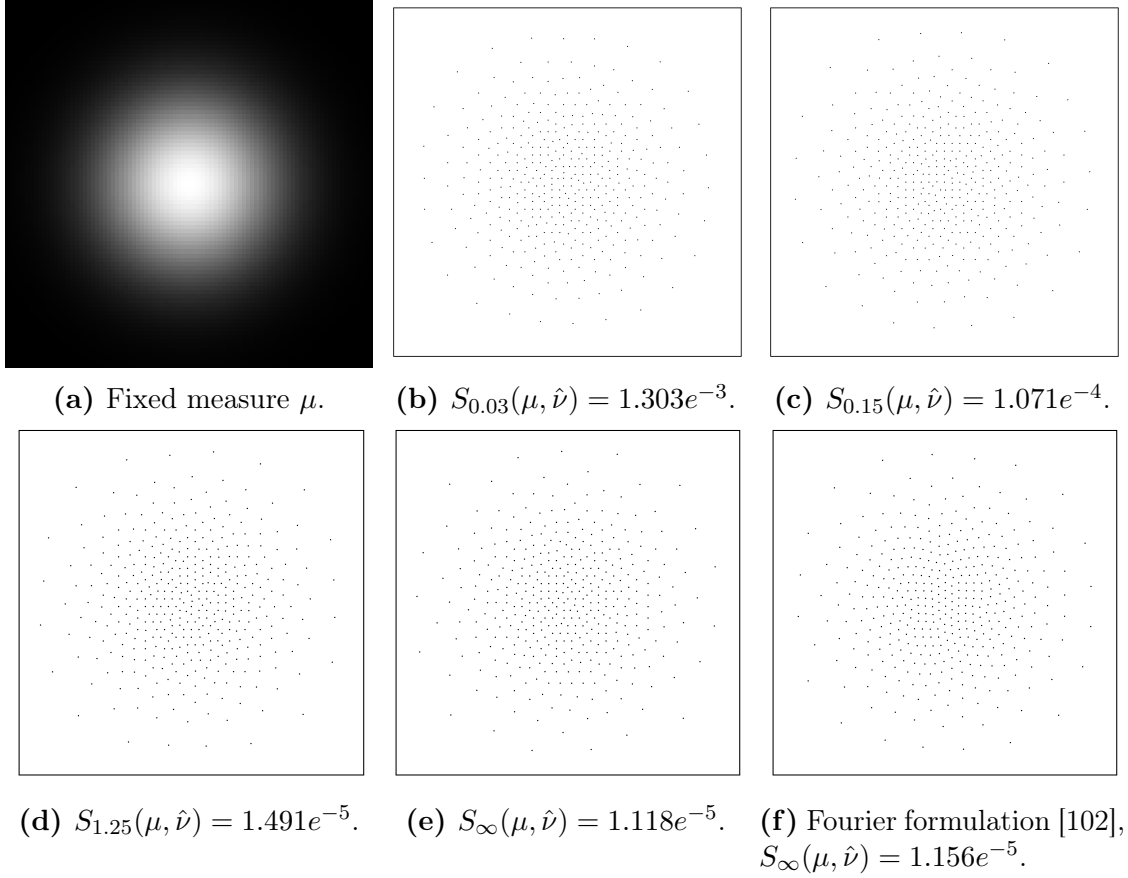


Figure 6.6: Optimal approximations $\hat{\nu}$ and corresponding energies $S_{\varepsilon}(\mu, \hat{\nu})$ for increasing ε .

so-called curse of dimensionality.

Finally, note that we sampled μ with a lot more points than we used for the dithering. If not enough points are used, we would observe clustering of the dithered measure around the positions of μ . One possibility to avoid such a behavior for S_{ε} could be to use the semi-discrete approach described in [123], avoiding any sampling of the measure μ . In the Fourier based approach, this issue was less pronounced.

6.8 Conclusions

In this chapter, we examined the behavior of the Sinkhorn divergences S_{ε} as $\varepsilon \rightarrow \infty$ and $\varepsilon \rightarrow 0$, with focus on the first case, which leads to discrepancies for appropriate cost functions and kernels. We considered a quite general scenario of measures involving, e.g., convex combinations of measures with densities and point measures (spikes). Besides application questions, some open theoretical problem are left. While OT_{ε} is monotone increasing in ε for any cost function c , we observed numerically for $c(x, y) = \|x - y\|$ that S_{ε} is monotone decreasing. Further, in Proposition 6.13 ii), we were not able to show convergence of the whole sequence

of optimal potentials $\{(\hat{\varphi}_\varepsilon, \hat{\psi}_\varepsilon)\}_\varepsilon$ without further assumptions so far.

6.A Basic theorems

We frequently apply the theorem of Arzelà–Ascoli. By definition, a sequence $\{f_n\}_{n \in \mathbb{N}}$ of continuous functions on \mathbb{X} is *uniformly bounded*, if there exists a constant $M \geq 0$ independent of n and x such that for all f_n and all $x \in \mathbb{X}$ it holds $|f_n(x)| \leq M$. The sequence is said to be *uniformly equi-continuous* if, for every $\varepsilon > 0$, there exists a $\delta > 0$ such that for all functions f_n

$$|f_n(x) - f_n(y)| < \varepsilon$$

whenever $d_{\mathbb{X}}(x, y) < \delta$.

Theorem 6.17. (*Arzelà–Ascoli*) *Let $\{f_n\}_{n \in \mathbb{N}}$ be a uniformly bounded, uniformly equi-continuous sequence of continuous functions on \mathbb{X} . Then, the sequence has a uniformly convergent subsequence.*

For the dual problems, we want to extend continuous functions from $A \subset \mathbb{X}$ to the whole space, which is possible by the following theorem. In the standard version, the theorem comes without the bounds, but they can be included directly since min and max of two continuous functions are again continuous functions.

Theorem 6.18. (*Tietze Extension Theorem*) *Let a closed subset $A \subset \mathbb{X}$ and a continuous function $f: A \rightarrow \mathbb{R}$ be given. If $g, h \in C(\mathbb{X})$ are such that $g \leq h$ and $g(x) \leq f(x) \leq h(x)$ for all $x \in A$, then there exists a continuous function $F: \mathbb{X} \rightarrow \mathbb{R}$ such that $F(x) = f(x)$ for all $x \in A$ and $g(x) \leq F(x) \leq h(x)$ for all $x \in \mathbb{X}$.*

Curve Based Approximation of Measures on Manifolds by Discrepancy Minimization

Contents

7.1	Introduction	165
7.2	Probability measures and curves	167
7.3	Discrepancies and RKHS	170
7.4	Approximation of general probability measures	173
7.5	Approximation of probability measures having Sobolev densities	176
7.5.1	Sobolev spaces and lower bounds	177
7.5.2	Upper bounds for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$	178
7.5.3	Upper bounds for $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$ and special manifolds \mathbb{X} .	180
7.5.4	Upper bounds for $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ and special manifolds \mathbb{X} .	183
7.6	Discretization	188
7.7	Numerical algorithm	189
7.8	Numerical results	192
7.8.1	Parameter choice	192
7.8.2	Quasi-optimal curves on special manifolds	195
7.9	Conclusions	203
7.A	Special manifolds	205

Abstract

The approximation of probability measures on compact metric spaces and in particular on Riemannian manifolds by atomic or empirical ones is a classical task in approximation and complexity theory with a wide range of applications. This chapter¹, which is published in [102], is instead concerned with the approximation by measures supported on Lipschitz curves. Special attention is paid to push-forward measures of Lebesgue measures on the unit interval by such curves. Using the discrepancy as distance between measures, we prove optimal approximation rates in terms of the curve's length and Lipschitz constant. Having established the theoretical convergence rates, we are interested in the numerical minimization of the discrepancy between a given probability measure and the set of push-forward measures of Lebesgue measures on the unit interval by Lipschitz curves. We present numerical examples for measures on the 2- and 3-dimensional torus, the 2-sphere, the rotation group on \mathbb{R}^3 and the Grassmannian of all 2-dimensional linear subspaces of \mathbb{R}^4 . Our algorithm of choice is a conjugate gradient method on these manifolds, which incorporates second-order information. For efficient gradient and Hessian evaluations within the algorithm, we approximate the given measures by truncated Fourier series and use fast Fourier transform techniques on these manifolds. At this point I also want to thank my colleague Manuel Gräf who provided the numerical implementation for this chapter.

¹This is a post-peer-review, pre-copyedit version of an article published in Foundations of Computational Mathematics. The final authenticated version is available online at DOI: 10.1007/s10208-021-09491-2. The article is distributed under the terms of the Creative Commons Attribution 4.0 International License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original authors and the source, provide a link to the Creative Commons license, and indicate if changes were made.

7.1 Introduction

The approximation of probability measures by atomic or empirical ones based on their discrepancies is a well examined problem in approximation and complexity theory [180, 194, 222] with a wide range of applications, e.g., in the derivation of quadrature rules and in the construction of designs. Recently, discrepancies were also used in image processing for dithering [134, 245, 259], i.e., for representing a gray-value image by a finite number of black dots, and in generative adversarial networks [96].

Besides discrepancies, Optimal Transport (OT) and in particular Wasserstein distances have emerged as powerful tools to compare probability measures in recent years, see [82, 274] and the references therein. In fact, so-called Sinkhorn divergences, which are computationally much easier to handle than OT, are known to interpolate between OT and discrepancies [108]. For the sample complexity of Sinkhorn divergences we refer to [122]. The rates for approximating probability measures by atomic or empirical ones with respect to Wasserstein distances depend on the dimension of the underlying spaces, see [68, 177]. In contrast, approximation rates based on discrepancies can be given independently of the dimension [222], i.e., they do not suffer from the curse of dimensionality. Additionally, we should keep in mind that the computation of discrepancies does not involve a minimization problem, which is a major drawback of OT and Sinkhorn divergences. Moreover, discrepancies admit a simple description in Fourier domain and hence the use of fast Fourier transforms is possible, leading to better scalability than the aforementioned methods.

Instead of point measures, we are interested in approximations with respect to measures supported on curves. More precisely, we consider push-forward measures of probability measures $\omega \in \mathcal{P}([0, 1])$ by Lipschitz curves of bounded speed, with special focus on absolutely continuous measures $\omega = \rho\lambda$ and the Lebesgue measure $\omega = \lambda$. In this chapter, we focus on approximation with respect to discrepancies. For related results on quadrature and approximation on manifolds, we refer to [109, 140, 196, 197] and the references therein. An approximation model based on the 2-Wasserstein distance was proposed in [183]. That work exploits completely different techniques than ours both in the theoretical and numerical part. Finally, we want to point out a relation to principal curves which are used in computer science and graphics for approximating distributions approximately supported on curves [151, 152, 173, 152, 175]. For the interested reader, we further comment on this direction of research in Remark 7.5 and in the conclusions. Next, we want to motivate our framework by numerous potential applications:

- In MRI sampling [43, 60], it is desirable to construct sampling curves with short sampling times (short curve) and high reconstruction quality. Unfortunately, these requirements usually contradict each other and finding a good trade-off is necessary. Experiments demonstrating the power of this novel approach on a real-world scanner are presented in [182].
- For laser engraving [183] and 3D printing [66], we require nozzle trajectories

based on our (continuous) input densities. Compared to the approach in [66], where points given by Lloyd's method are connected as a solution of the TSP (traveling salesman problem), our method jointly selects the points and the corresponding curve. This avoids the necessity of solving a TSP, which can be quite costly, although efficient approximations exist. Further, it is not obvious that the fixed initial point approximation is a good starting point for constructing a curve.

- The model can be used for wire sculpture creation [3]. In view of this, our numerical experiment presented in Fig. 7.5 can be interpreted as a building plan for a wire sculpture of the Spock head, namely of a 2D surface. Clearly, the approach can be also used to create images similar to TSP Art [170], where images are created from points by solving the corresponding TSP.
- In a more manifold related setting, the approach can be used for grand tour computation on $\mathcal{G}_{2,4}$ [12], see also our numerical experiment in Fig. 7.11. More technical details are provided in the corresponding section.

Our contribution is two-fold. On the theoretical side, we provide estimates of the approximation rates in terms of the maximal speed of the curve. First, we prove approximation rates for general probability measures on compact Ahlfors d -regular length spaces \mathbb{X} . These spaces include many compact sets in the Euclidean space \mathbb{R}^d , e.g., the unit ball or the unit cube as well as d -dimensional compact Riemannian manifolds without boundary. The basic idea consists in combining the known convergence rates for approximation by atomic measures with cost estimates for the traveling salesman problem. As for point measures, the approximation rate $L^{d/(2d-2)} \leq L^{-1/2}$ for general $\omega \in \mathcal{P}([0, 1])$ and $L^{d/(3d-2)} \leq L^{-1/3}$ for $\omega = \lambda$ in terms of the maximal Lipschitz constant (speed) L of the curves does not crucially depend on the dimension of \mathbb{X} . In particular, the second estimate improves a result given in [61] for the torus.

If the measures fulfill additional smoothness properties, these estimates can be improved on compact, connected, d -dimensional Riemannian manifolds without boundary. Our results are formulated for absolutely continuous measures (with respect to the Riemannian measure) having densities in the Sobolev space $H^s(\mathbb{X})$, $s > d/2$. In this setting, the optimal approximation rate becomes roughly speaking $L^{-s/(d-1)}$. Our proofs rely on a general result of Brandolini et al. [45] on the quadrature error achievable by integration with respect to a measure that exactly integrates all eigenfunctions of the Laplace–Beltrami with eigenvalues smaller than a fixed number. Hence, we need to construct measures supported on curves that fulfill the above exactness criterion. More precisely, we construct such curves for the d dimensional torus \mathbb{T}^d , the spheres \mathbb{S}^d , the rotation group $\mathrm{SO}(3)$ and the Grassmannian $\mathcal{G}_{2,4}$.

On the numerical side, we are interested in finding (local) minimizers of discrepancies between a given continuous measure and those from the set of push-forward measures of the Lebesgue measure by bounded Lipschitz curves. This problem is tackled numerically on \mathbb{T}^2 , \mathbb{T}^3 , \mathbb{S}^2 as well as $\mathrm{SO}(3)$ and $\mathcal{G}_{2,4}$ by switching to the

Fourier domain. The minimizers are computed using the method of conjugate gradients (CG) on manifolds, which incorporates second order information in form of a multiplication by the Hessian. Thanks to the approach in the Fourier domain, the required gradients and the calculations involving the Hessian can be performed efficiently by fast Fourier transform techniques at arbitrary nodes on the respective manifolds. Note that in contrast to our approach, semi-continuous OT minimization relies on Laguerre tessellations [87], which are not available in the required form on the 2-sphere, $\text{SO}(3)$ or $\mathcal{G}_{2,4}$.

This chapter is organized as follows: In Section 7.2 we give the necessary preliminaries on probability measures. In particular, we introduce the different sets of measures supported on Lipschitz curves that are used for the approximation. Note that measures supported on continuous curves of finite length can be equivalently characterized by push-forward measures of probability measures by Lipschitz curves. Section 7.3 provides the notation on reproducing kernel Hilbert spaces and discrepancies including their representation in the Fourier domain. Section 7.4 contains our estimates of the approximation rates for general given measures and different approximation spaces of measures supported on curves. Following the usual lines in approximation theory, we are then concerned with the approximation of absolutely continuous measures with density functions lying in Sobolev spaces. Our main results on the approximation rates of smoother measures are contained in Section 7.5, where we distinguish between the approximation with respect to the push-forward of general measures $\omega \in \mathcal{P}[0, 1]$, absolute continuous measures and the Lebesgue measure on $[0, 1]$. In Section 7.6 we formulate our numerical minimization problem. Our numerical algorithms of choice are briefly described in Section 7.7. For a comprehensive description of the algorithms on the different manifolds, we refer to respective papers. Section 7.8 contains numerical results demonstrating the practical feasibility of our findings. Conclusions are drawn in Section 7.9. Finally, Appendix 7.A briefly introduces the different manifolds \mathbb{X} used in our numerical examples together with the Fourier representation of probability measures on \mathbb{X} .

7.2 Probability measures and curves

In this section, the basic notation on measure spaces is provided, see [5, 113], with focus on probability measures supported on curves. At this point, let us assume that

\mathbb{X} is a compact metric space endowed with a bounded non-negative Borel measure $\sigma_{\mathbb{X}} \in \mathcal{M}(\mathbb{X})$ such that $\text{supp}(\sigma_{\mathbb{X}}) = \mathbb{X}$. Further, we denote the metric by $d_{\mathbb{X}}$.

Additional requirements on \mathbb{X} are added along the way and notations are explained below. By $\mathcal{B}(\mathbb{X})$ we denote the Borel σ -algebra on \mathbb{X} and by $\mathcal{M}(\mathbb{X})$ the linear space of all finite signed Borel measures on \mathbb{X} , i.e., the space of all $\mu: \mathcal{B}(\mathbb{X}) \rightarrow \mathbb{R}$ satisfying $\mu(\mathbb{X}) < \infty$ and for any sequence $(B_k)_{k \in \mathbb{N}} \subset \mathcal{B}(\mathbb{X})$ of pairwise disjoint sets

the relation $\mu(\bigcup_{k=1}^{\infty} B_k) = \sum_{k=1}^{\infty} \mu(B_k)$. The *support of a measure* μ is the closed set

$$\text{supp}(\mu) := \{x \in \mathbb{X} : B \subset \mathbb{X} \text{ open, } x \in B \implies \mu(B) > 0\}.$$

For $\mu \in \mathcal{M}(\mathbb{X})$ the total variation measure is defined by

$$|\mu|(B) := \sup \left\{ \sum_{k=1}^{\infty} |\mu(B_k)| : \bigcup_{k=1}^{\infty} B_k = B, B_k \text{ pairwise disjoint} \right\}.$$

With the norm $\|\mu\|_{\mathcal{M}} = |\mu|(\mathbb{X})$ the space $\mathcal{M}(\mathbb{X})$ becomes a Banach space. By $\mathcal{C}(\mathbb{X})$ we denote the Banach space of continuous real-valued functions on \mathbb{X} equipped with the norm $\|\varphi\|_{\mathcal{C}(\mathbb{X})} := \max_{x \in \mathbb{X}} |\varphi(x)|$. The space $\mathcal{M}(\mathbb{X})$ can be identified via Riesz' theorem with the dual space of $\mathcal{C}(\mathbb{X})$ and the weak-* topology on $\mathcal{M}(\mathbb{X})$ gives rise to the *weak convergence of measures*, i.e., a sequence $(\mu_k)_k \subset \mathcal{M}(\mathbb{X})$ converges *weakly* to μ and we write $\mu_k \rightharpoonup \mu$, if

$$\lim_{k \rightarrow \infty} \int_{\mathbb{X}} \varphi d\mu_k = \int_{\mathbb{X}} \varphi d\mu \quad \text{for all } \varphi \in \mathcal{C}(\mathbb{X}).$$

For a non-negative, finite measure μ , let $L^p(\mathbb{X}, \mu)$ be the Banach space (of equivalence classes) of complex-valued functions with norm

$$\|f\|_{L^p(\mathbb{X}, \mu)} = \left(\int_{\mathbb{X}} |f|^p d\mu \right)^{\frac{1}{p}} < \infty.$$

By $\mathcal{P}(\mathbb{X})$ we denote the space of Borel probability measures on \mathbb{X} , i.e., non-negative Borel measures with $\mu(\mathbb{X}) = 1$. This space is *weakly compact*, i.e., compact with respect to the topology of weak convergence. We are interested in the approximation of measures in $\mathcal{P}(\mathbb{X})$ by probability measures supported on points and curves in \mathbb{X} . To this end, we associate with $x \in \mathbb{X}$ a probability measure δ_x with values $\delta_x(B) = 1$ if $x \in B$ and $\delta_x(B) = 0$ otherwise.

The *atomic probability measures* at N points are defined by

$$\mathcal{P}_N^{\text{atom}}(\mathbb{X}) := \left\{ \sum_{k=1}^N w_k \delta_{x_k} : (x_k)_{k=1}^N \in \mathbb{X}^N, (w_k)_{k=1}^N \in [0, 1]^N, \sum_{k=1}^N w_k = 1 \right\}.$$

In other words, $\mathcal{P}_N^{\text{atom}}(\mathbb{X})$ is the collection of probability measures, whose support consists of at most N points. Further restriction to equal mass distribution leads to the *empirical probability measures* at N points denoted by

$$\mathcal{P}_N^{\text{emp}}(\mathbb{X}) := \left\{ \frac{1}{N} \sum_{k=1}^N \delta_{x_k} : (x_k)_{k=1}^N \in \mathbb{X}^N \right\}.$$

In this chapter, we are interested in the approximation by measures having their support on curves. Let $\mathcal{C}([a, b], \mathbb{X})$ denote the set of closed, continuous curves $\gamma: [a, b] \rightarrow \mathbb{X}$. Although our presented experiments involve solely closed curves,

some applications might require open curves. Hence, we want to point out that all of our approximation results still hold without this requirement. Upper bounds would not get worse and we have not used the closedness for the lower bounds on the approximation rates. The *length* of a curve $\gamma \in \mathcal{C}([a, b], \mathbb{X})$ is given by

$$\ell(\gamma) := \sup_{\substack{a \leq t_0 \leq \dots \leq t_n \leq b \\ n \in \mathbb{N}}} \sum_{k=1}^n d_{\mathbb{X}}(\gamma(t_k), \gamma(t_{k-1})).$$

If $\ell(\gamma) < \infty$, then γ is called *rectifiable*. By reparametrization, see [145, Thm. 3.2], the image of any rectifiable curve in $\mathcal{C}([a, b], \mathbb{X})$ can be derived from the set of closed *Lipschitz continuous curves*

$$\text{Lip}(\mathbb{X}) := \{\gamma \in \mathcal{C}([0, 1], \mathbb{X}) : \exists L \in \mathbb{R} \text{ with } d_{\mathbb{X}}(\gamma(s), \gamma(t)) \leq L|s - t| \forall s, t \in [0, 1]\}.$$

The *speed* of a curve $\gamma \in \text{Lip}(\mathbb{X})$ is defined a.e. by the metric derivative

$$|\dot{\gamma}|(t) := \lim_{s \rightarrow t} \frac{d_{\mathbb{X}}(\gamma(s), \gamma(t))}{|s - t|}, \quad t \in [0, 1],$$

cf. [6, Sec. 1.1]. The optimal Lipschitz constant $L = L(\gamma)$ of a curve γ is given by $L(\gamma) = \|\dot{\gamma}\|_{\infty([0, 1])}$. For a constant speed curve it holds $L(\gamma) = \ell(\gamma)$.

We aim to approximate measures in $\mathcal{P}(\mathbb{X})$ from those of the subset

$$\mathcal{P}_L^{\text{curv}}(\mathbb{X}) := \{\nu \in \mathcal{P}(\mathbb{X}) : \exists \gamma \in \mathcal{C}([a, b], \mathbb{X}), \text{supp}(\nu) \subset \gamma([a, b]), \ell(\gamma) \leq L\}. \quad (7.1)$$

This space is quite large and in order to define further meaningful subsets, we derive an equivalent formulation in terms of push-forward measures. For $\gamma \in \mathcal{C}([0, 1], \mathbb{X})$, the *push-forward* $\gamma_*\omega \in \mathcal{P}(\mathbb{X})$ of a probability measure $\omega \in \mathcal{P}([0, 1])$ is defined by $\gamma_*\omega(B) := \omega(\gamma^{-1}(B))$ for $B \in \mathcal{B}(\mathbb{X})$. We directly observe $\text{supp}(\gamma_*\omega) = \gamma(\text{supp}(\omega))$. By the following lemma, $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$ consists of the push-forward of measures in $\mathcal{P}([0, 1])$ by constant speed curves.

Lemma 7.1. *The space $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$ in (7.1) is equivalently given by*

$$\mathcal{P}_L^{\text{curv}}(\mathbb{X}) = \{\gamma_*\omega : \gamma \in \text{Lip}(\mathbb{X}) \text{ has constant speed } L(\gamma) \leq L, \omega \in \mathcal{P}([0, 1])\}. \quad (7.2)$$

Proof. Let $\nu \in \mathcal{P}_L^{\text{curv}}(\mathbb{X})$ as in (7.1). If $\text{supp}(\nu)$ consists of a single point $x \in \mathbb{X}$ only, then the constant curve $\gamma \equiv x$ pushes forward an arbitrary δ_t for $t \in [a, b]$, which shows that ν is contained in (7.2).

Suppose that $\text{supp}(\nu)$ contains at least two distinct points and let $\gamma \in \mathcal{C}([a, b], \mathbb{X})$ with $\text{supp}(\nu) \subset \gamma([a, b])$ and $\ell(\gamma) < \infty$. According to [51, Prop. 2.5.9], there exists a continuous curve $\tilde{\gamma} \in \text{Lip}(\mathbb{X})$ with constant speed $\ell(\gamma)$ and a continuous non-decreasing function $\varphi: [a, b] \rightarrow [0, 1]$ with $\gamma = \tilde{\gamma} \circ \varphi$. Now, define $f: \mathbb{X} \rightarrow [0, 1]$ by $f(x) := \min\{\tilde{\gamma}^{-1}(x)\}$. This function is measurable, since for every $t \in [0, 1]$ it holds that

$$\{x \in \mathbb{X} : f(x) \leq t\} = \{x \in \mathbb{X} : \min\{\tilde{\gamma}^{-1}(x)\} \leq t\} = \tilde{\gamma}([0, t])$$

is compact. Due to $\text{supp}(\nu) \subset \tilde{\gamma}([0, 1])$, we can define $\omega := f_*\nu \in \mathcal{P}([0, 1])$. By construction, ω satisfies $\tilde{\gamma}_*\omega(B) = \omega(\tilde{\gamma}^{-1}(B)) = \nu(f^{-1} \circ \tilde{\gamma}^{-1}(B)) = \nu(B)$ for all $B \in \mathcal{B}(\mathbb{X})$. This concludes the proof. \square

The set $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$ contains $\mathcal{P}_N^{\text{atom}}(\mathbb{X})$ if L is sufficiently large compared to N and \mathbb{X} is sufficiently nice, cf. Section 7.4. It is reasonable to ask for more restrictive sets of approximation measures, e.g., when $\omega \in \mathcal{P}([0, 1])$ is assumed to be absolutely continuous. For the Lebesgue measure λ on $[0, 1]$, we consider

$$\mathcal{P}_L^{\text{a-curv}}(\mathbb{X}) := \{\gamma_*\omega : \gamma \in \text{Lip}(\mathbb{X}), L(\gamma) \leq L, \omega = \rho\lambda \in \mathcal{P}([0, 1]), L(\rho) \leq L\}.$$

In the literature [61, 183], the special case of push-forward of the Lebesgue measure $\omega = \lambda$ on $[0, 1]$ by Lipschitz curves in \mathbb{T}^d was discussed and successfully used in certain applications [43, 60]. Therefore, we also consider approximations from

$$\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X}) := \{\gamma_*\lambda : \gamma \in \text{Lip}(\mathbb{X}), L(\gamma) \leq L\}.$$

It is obvious that our probability spaces related to curves are nested,

$$\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X}) \subset \mathcal{P}_L^{\text{a-curv}}(\mathbb{X}) \subset \mathcal{P}_L^{\text{curv}}(\mathbb{X}).$$

Hence, one may expect that establishing good approximation rates is most difficult for $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ and easier for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$.

7.3 Discrepancies and RKHS

The aim of this section is to introduce the way we quantify the distance (“discrepancy”) between two probability measures. To this end, choose a continuous, symmetric function $K: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ that is positive definite, i.e., for any finite number $n \in \mathbb{N}$ of points $x_j \in \mathbb{X}$, $j = 1, \dots, n$, the relation

$$\sum_{i,j=1}^n a_i a_j K(x_i, x_j) \geq 0$$

is satisfied for all $a_j \in \mathbb{R}$, $j = 1, \dots, n$. We know by Mercer’s theorem [80, 195, 257] that there exists an orthonormal basis $\{\phi_k : k \in \mathbb{N}\}$ of $L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ and non-negative coefficients $(\alpha_k)_{k \in \mathbb{N}} \in \ell_1$ such that K has the Fourier expansion

$$K(x, y) = \sum_{k=0}^{\infty} \alpha_k \phi_k(x) \overline{\phi_k(y)} \quad (7.3)$$

with absolute and uniform convergence of the right-hand side. If $\alpha_k > 0$ for some $k \in \mathbb{N}_0$, the corresponding function ϕ_k is continuous. Every function $f \in L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ has a Fourier expansion

$$f = \sum_{k=0}^{\infty} \hat{f}_k \phi_k, \quad \hat{f}_k := \int_{\mathbb{X}} f \overline{\phi_k} d\sigma_{\mathbb{X}}.$$

The kernel K gives rise to a *reproducing kernel Hilbert space* (RKHS). More precisely, the function space

$$H_K(\mathbb{X}) := \left\{ f \in L^2(\mathbb{X}, \sigma_{\mathbb{X}}) : \sum_{k=0}^{\infty} \alpha_k^{-1} |\hat{f}_k|^2 < \infty \right\}$$

equipped with the inner product and the corresponding norm

$$\langle f, g \rangle_{H_K(\mathbb{X})} = \sum_{k=0}^{\infty} \alpha_k^{-1} \hat{f}_k \overline{\hat{g}_k}, \quad \|f\|_{H_K(\mathbb{X})} = \sqrt{\langle f, f \rangle_{H_K(\mathbb{X})}} \quad (7.4)$$

forms a Hilbert space with reproducing kernel, i.e.,

$$\begin{aligned} K(x, \cdot) &\in H_K(\mathbb{X}) && \text{for all } x \in \mathbb{X}, \\ f(x) &= \langle f, K(x, \cdot) \rangle_{H_K(\mathbb{X})} && \text{for all } f \in H_K(\mathbb{X}), x \in \mathbb{X}. \end{aligned}$$

Note that $f \in H_K(\mathbb{X})$ implies $\hat{f}_k = 0$ if $\alpha_k = 0$, in which case we make the convention $\alpha_k^{-1} \hat{f}_k = 0$ in (7.4). The space $H_K(\mathbb{X})$ is the closure of the linear span of $\{K(x_j, \cdot) : x_j \in \mathbb{X}\}$ with respect to the norm (7.4), and $H_K(\mathbb{X})$ is continuously embedded in $C(\mathbb{X})$. In particular, the point evaluations in $H_K(\mathbb{X})$ are continuous.

The *discrepancy* $\mathcal{D}_K(\mu, \nu)$ is defined as the dual norm on $H_K(\mathbb{X})$ of the linear operator $T: H_K(\mathbb{X}) \rightarrow \mathbb{C}$ with $\varphi \mapsto \int_{\mathbb{X}} \varphi d(\mu - \nu)$:

$$\mathcal{D}_K(\mu, \nu) = \max_{\|\varphi\|_{H_K(\mathbb{X})} \leq 1} \left| \int_{\mathbb{X}} \varphi d(\mu - \nu) \right|, \quad (7.5)$$

see [128, 222]. Note that this looks similar to the 1-Wasserstein distance, where the space of test functions consists of Lipschitz continuous functions and is larger. Since

$$\int_{\mathbb{X}} \varphi d\mu = \int_{\mathbb{X}} \langle \varphi, K(x, \cdot) \rangle_{H_K(\mathbb{X})} d\mu(x) = \left\langle \varphi, \int_{\mathbb{X}} K(x, \cdot) d\mu(x) \right\rangle_{H_K(\mathbb{X})},$$

we obtain by Riesz's representation theorem

$$\max_{\|\varphi\|_{H_K(\mathbb{X})} \leq 1} \int_{\mathbb{X}} \varphi d\mu = \left\| \int_{\mathbb{X}} K(x, \cdot) d\mu(x) \right\|_{H_K(\mathbb{X})},$$

which yields by Fubini's theorem, (7.3), (7.4) and symmetry of K that

$$\mathcal{D}_K^2(\mu, \nu) = \iint_{\mathbb{X} \times \mathbb{X}} K d\mu d\mu - 2 \iint_{\mathbb{X} \times \mathbb{X}} K d\mu d\nu + \iint_{\mathbb{X} \times \mathbb{X}} K d\nu d\nu \quad (7.6)$$

$$= \sum_{k=0}^{\infty} \alpha_k |\hat{\mu}_k - \hat{\nu}_k|^2, \quad (7.7)$$

where the *Fourier coefficients* of $\mu, \nu \in \mathcal{P}(\mathbb{X})$ are well-defined for k with $\alpha_k \neq 0$ by

$$\hat{\mu}_k := \int_{\mathbb{X}} \overline{\phi_k} d\mu, \quad \hat{\nu}_k := \int_{\mathbb{X}} \overline{\phi_k} d\nu.$$

Remark 7.2. The Fourier coefficients $\hat{\mu}_k$ and $\hat{\nu}_k$ depend on both K and $\sigma_{\mathbb{X}}$, but the identity (7.6) shows that $\mathcal{D}_K(\mu, \nu)$ only depends on K . Thus, our approximation rates do not depend on the choice of $\sigma_{\mathbb{X}}$. On the other hand, our numerical algorithms in Section 7.7 depend on ϕ_k and hence on the choice of $\sigma_{\mathbb{X}}$.

If $\mu_n \rightharpoonup \mu$ and $\nu_n \rightharpoonup \nu$ as $n \rightarrow \infty$, then also $\mu_n \otimes \nu_n \rightharpoonup \mu \otimes \nu$. Therefore, the continuity of K implies that $\lim_{n \rightarrow \infty} \mathcal{D}_K(\mu_n, \nu_n) = \mathcal{D}_K(\mu, \nu)$, so that \mathcal{D}_K is continuous with respect to weak convergence in both arguments. Thus, for any weakly compact subset $P \subset \mathcal{P}(\mathbb{X})$, the infimum

$$\inf_{\nu \in P} \mathcal{D}_K(\mu, \nu)$$

is actually a minimum. All of the subsets introduced in the previous section are weakly compact.

Lemma 7.3. *The sets $\mathcal{P}_N^{\text{atom}}(\mathbb{X})$, $\mathcal{P}_N^{\text{emp}}(\mathbb{X})$, $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$, $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$, and $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ are weakly compact.*

Proof. It is well-known that $\mathcal{P}_N^{\text{atom}}(\mathbb{X})$ and $\mathcal{P}_N^{\text{emp}}(\mathbb{X})$ are weakly compact.

We show that $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$ is weakly compact. In view of (7.2), let $(\gamma_k)_{k \in \mathbb{N}}$ be Lipschitz curves with constant speed $L(\gamma_k) \leq L$ and $(\omega_k)_{k \in \mathbb{N}} \subset \mathcal{P}([0, 1])$. Since $\mathcal{P}([0, 1])$ is weakly compact, we can extract a subsequence $(\omega_{k_j})_{j \in \mathbb{N}}$ with weak limit $\hat{\omega} \in \mathcal{P}([0, 1])$. Now, we observe that $d_{\mathbb{X}}(\gamma_{k_j}(s), \gamma_{k_j}(t)) \leq L|s - t|$ for all $j \in \mathbb{N}$. Since \mathbb{X} is compact, the Arzelà–Ascoli theorem implies that there exists a subsequence of $(\gamma_{k_j})_{j \in \mathbb{N}}$ which converges uniformly towards $\hat{\gamma} \in \text{Lip}(\mathbb{X})$ with $L(\hat{\gamma}) \leq L$. Then, $\hat{\nu} := \hat{\gamma}_* \hat{\omega}$ fulfills $\text{supp}(\hat{\nu}) \subset \hat{\gamma}([0, 1])$, so that $\hat{\nu} \in \mathcal{P}_L^{\text{curv}}(\mathbb{X})$ by (7.1). Thus, $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$ is weakly compact.

The proof for $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$ and $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ is analogous and hence omitted. \square

Remark 7.4. (*Discrepancies and Convolution Kernels*) Let $\mathbb{X} = \mathbb{T}^d := \mathbb{R}^d / \mathbb{Z}^d$ be the torus and $h \in \mathcal{C}(\mathbb{T}^d)$ be a function with Fourier series

$$h(x) = \sum_{k \in \mathbb{Z}^d} \hat{h}_k e^{2\pi i \langle k, x \rangle}, \quad \hat{h}_k := \int_{\mathbb{T}^d} h(x) e^{-2\pi i \langle k, x \rangle} d\sigma_{\mathbb{T}^d}(x),$$

which converges in $L^2(\mathbb{T}^d)$ so that $\sum_k |\hat{h}_k|^2 < \infty$. Assume that $\hat{h}_k \neq 0$ for all $k \in \mathbb{Z}^d$. We consider the special Mercer kernel

$$K(x, y) := \sum_{k \in \mathbb{Z}^d} |\hat{h}_k|^2 e^{2\pi i \langle k, x - y \rangle} = \sum_{k \in \mathbb{Z}^d} |\hat{h}_k|^2 \cos(2\pi \langle k, x - y \rangle)$$

with associated discrepancy \mathcal{D}_h via (7.6), i.e., $\phi_k(x) = e^{2\pi i \langle k, x \rangle}$, $\alpha_k = |\hat{h}_k|^2$, $k \in \mathbb{Z}^d$ in (7.3). The convolution of h with $\mu \in \mathcal{M}(\mathbb{T}^d)$ is the function $h * \mu \in \mathcal{C}(\mathbb{T}^d)$ defined by

$$(h * \mu)(x) := \int_{\mathbb{T}^d} h(x - y) d\mu(y).$$

By the convolution theorem for Fourier transforms it holds $(\widehat{h * \mu})_k = \hat{h}_k \hat{\mu}_k$, $k \in \mathbb{Z}^d$, and we obtain by Parseval's identity for $\mu, \nu \in \mathcal{M}(\mathbb{T}^d)$ and (7.7) that

$$\|h * (\mu - \nu)\|_{L^2(\mathbb{T}^d)}^2 = \|(\hat{h}_k (\hat{\mu}_k - \hat{\nu}_k))_{k \in \mathbb{Z}^d}\|_{\ell_2}^2 = \sum_{k \in \mathbb{Z}^d} |\hat{h}_k|^2 |\hat{\mu}_k - \hat{\nu}_k|^2 = \mathcal{D}_h^2(\mu, \nu).$$

In image processing, metrics of this kind were considered in [61, 114, 259].

Remark 7.5. (*Relations to Principal Curves*) A similar concept, sharing the common theme of “a curve which passes through the middle of a distribution” with the intention of our chapter, is that of principle curves. The notion of principal curves has been developed in a statistical framework and was successfully applied in statistics and machine learning, see [124, 173, 175]. The idea is to generalize the concept of principal components with just one direction to so-called self-consistent (principal) curves. In the seminal paper [151], the authors showed that these principal curves γ are critical points of the energy functional

$$E(\gamma, \mu) = \int_{\mathbb{X}} \|x - \text{proj}_{\gamma}(x)\|_2^2 d\mu(x), \quad (7.8)$$

where μ is a given probability measure on \mathbb{X} and $\text{proj}_{\gamma}(x) = \arg\min_{y \in \gamma} \|x - y\|_2$ is a projection of a point $x \in \mathbb{X}$ on γ . This notion has also been generalized to Riemannian manifolds in [152], see also [175] for an application on the sphere. Further investigation of principal curves in the plane, cf. [94], showed that self-consistent curves are not (local) minimizers, but saddle points of (7.8). Moreover, the existence of such curves is established only for certain classes of measures, such as elliptical ones. By additionally constraining the length of curves minimizing (7.8), these unfavorable effects were eliminated, cf. [173]. In comparison to the objective (7.8), the discrepancy (7.6) averages for fixed $x \in \mathbb{X}$ the distance encoded by K to any point on γ , instead of averaging over the squared minimal distances to γ .

7.4 Approximation of general probability measures

Given $\mu \in \mathcal{P}(\mathbb{X})$, the estimates²

$$\min_{\nu \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_N^{\text{emp}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim N^{-\frac{1}{2}}, \quad (7.9)$$

are well-known, cf. [131, Cor. 2.8]. Here, the constant hidden in \lesssim depends on \mathbb{X} and K but is independent of μ and $N \in \mathbb{N}$. In this section, we are interested in approximation rates with respect to measures supported on curves.

Our approximation rates for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$ are based on those for $\mathcal{P}_N^{\text{atom}}(\mathbb{X})$ combined with estimates for the traveling salesman problem (TSP). Let $\text{TSP}_{\mathbb{X}}(N)$ denote the worst case minimal cost tour in a fully connected graph G of N arbitrary nodes represented by $x_1, \dots, x_N \in \mathbb{X}$ and edges with cost $d_{\mathbb{X}}(x_i, x_j)$, $i, j = 1, \dots, N$. Similarly, let $\text{MST}_{\mathbb{X}}(N)$ denote the worst case cost of the minimal spanning tree of G . To derive suitable estimates, we require that \mathbb{X} is Ahlfors d -regular (sometimes also called Ahlfors-David d -regular), i.e., there exists $0 < d < \infty$ such that

$$\sigma_{\mathbb{X}}(B_r(x)) \sim r^d, \quad \text{for all } x \in \mathbb{X}, \quad 0 < r \leq \text{diam}(\mathbb{X}), \quad (7.10)$$

²We use the symbols \lesssim and \gtrsim to indicate that the corresponding inequalities hold up to a positive constant factor on the respective right-hand side. The notation \sim means that both relations \lesssim and \gtrsim hold. The dependence of the constants on other parameters shall either be explicitly stated or clear from the context.

where $B_r(x) = \{y \in \mathbb{X} : d_{\mathbb{X}}(x, y) \leq r\}$ and the constants in \sim do not depend on x or r . Note that d is not required to be an integer and turns out to be the Hausdorff dimension. For \mathbb{X} being the unit cube the following lemma was proved in [253].

Lemma 7.6. *If \mathbb{X} is a compact Ahlfors d -regular metric space, then there is a constant $0 < C_{\text{TSP}} < \infty$ depending on \mathbb{X} such that*

$$\text{TSP}_{\mathbb{X}}(N) \leq C_{\text{TSP}} N^{1-\frac{1}{d}}.$$

Proof. Using (7.10) and the same covering argument as in [252, Lem. 3.1], we see that for every choice $x_1, \dots, x_N \in \mathbb{X}$, there exist $i \neq j$ such that $d_{\mathbb{X}}(x_i, x_j) \lesssim N^{-1/d}$, where the constant depends on \mathbb{X} .

Let $S = \{x_1, \dots, x_N\}$ be an arbitrary selection of N points from \mathbb{X} . First, we choose x_i and x_j with $d_{\mathbb{X}}(x_i, x_j) \leq cN^{-1/d}$. Then, we form a minimal spanning tree T of $S \setminus \{x_i\}$ and augment the tree by adding the edge between x_i and x_j . This construction provides us with a spanning tree and hence we can estimate $\text{MST}_{\mathbb{X}}(N) \leq \text{MST}_{\mathbb{X}}(N-1) + cN^{-1/d}$. Iterating the argument, we deduce

$$\text{MST}_{\mathbb{X}}(N) \lesssim N^{1-\frac{1}{d}},$$

cf. [253]. Finally, the standard relation $\text{TSP}_{\mathbb{X}}(N) \leq 2\text{MST}_{\mathbb{X}}(N)$ for edge costs satisfying the triangular inequality concludes the proof. \square

To derive a curve in \mathbb{X} from a minimal cost tour in the graph, we require the additional assumption that \mathbb{X} is a *length space*, i.e., a metric space with

$$d_{\mathbb{X}}(x, y) = \inf\{\ell(\gamma) : \gamma \text{ a continuous curve that connects } x \text{ and } y\},$$

cf. [48, 51]. Thus, for the rest of this section, we are assuming that

\mathbb{X} is a compact Ahlfors d -regular length space.

In this case, Lemma 7.6 yields the next proposition.

Proposition 7.7. *If \mathbb{X} is a compact Ahlfors d -regular length space, then it holds $\mathcal{P}_N^{\text{atom}}(\mathbb{X}) \subset \mathcal{P}_{C_{\text{TSP}}N^{1-1/d}}^{\text{curv}}(\mathbb{X})$.*

Proof. The Hopf-Rinow Theorem for metric measure spaces, see [48, Chap. I.3] and [51, Thm. 2.5.28], yields that every pair of points $x, y \in \mathbb{X}$ can be connected by a geodesic, i.e., there is $\gamma \in \text{Lip}(\mathbb{X})$ with constant speed and $\ell(\gamma|_{[s,t]}) = d_{\mathbb{X}}(\gamma(s), \gamma(t))$ for all $0 \leq s \leq t \leq 1$. Thus, for any pair $x, y \in \mathbb{X}$, there is a constant speed curve $\gamma_{x,y} \in \text{Lip}(\mathbb{X})$ of length $\ell(\gamma_{x,y}) = d_{\mathbb{X}}(x, y)$ with $\gamma_{x,y}(0) = x$, $\gamma_{x,y}(1) = y$, cf. [51, Rem. 2.5.29]. For $\mu_N \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})$, let $\{x_1, \dots, x_N\} = \text{supp}(\mu_N)$. The minimal cost tour in Lemma 7.6 leads to a curve $\gamma \in \text{Lip}(\mathbb{X})$, so that $\mu_N = \gamma_*\omega \in \mathcal{P}_L^{\text{curv}}(\mathbb{X})$ for an appropriate measure $\omega \in \mathcal{P}_N^{\text{atom}}([0, 1])$. \square

Proposition 7.7 enables us to transfer approximation rates from $\mathcal{P}_N^{\text{atom}}(\mathbb{X})$ to $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$.

Theorem 7.8. For $\mu \in \mathcal{P}(\mathbb{X})$, it holds with a constant depending on \mathbb{X} and K that

$$\min_{\nu \in \mathcal{P}_L^{\text{curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{d}{2d-2}}.$$

Proof. Choose $\alpha = \frac{d-1}{d}$. For L large enough, set $N := \lfloor (L/C_{\text{TSP}})^{\frac{1}{\alpha}} \rfloor \in \mathbb{N}$, so that we observe $\mathcal{P}_N^{\text{atom}}(\mathbb{X}) \subset \mathcal{P}_L^{\text{curv}}(\mathbb{X})$. According to (7.9), we obtain

$$\min_{\nu \in \mathcal{P}_L^{\text{curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim N^{-\frac{1}{2}} \lesssim L^{-\frac{1}{2\alpha}}.$$

□

Next, we derive approximation rates for $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$ and $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$.

Theorem 7.9. For $\mu \in \mathcal{P}(\mathbb{X})$, we have with a constant depending on \mathbb{X} and K that

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{d}{3d-2}}. \quad (7.11)$$

Proof. Let $\alpha = \frac{d-1}{d}$, $d \geq 2$. For L large enough, set $N := \lfloor L^{\frac{2}{2\alpha+1}} / \text{diam}(\mathbb{X}) \rfloor \in \mathbb{N}$. By (7.9), there is a set of points $\{x_1, \dots, x_N\} \subset \mathbb{X}$ such that

$$\mathcal{D}_K(\mu, \nu_N) \lesssim N^{-\frac{1}{2}} \lesssim L^{-\frac{1}{2\alpha+1}}, \quad \nu_N := \frac{1}{N} \sum_{j=1}^N \delta_{x_j}. \quad (7.12)$$

Let these points be ordered as a solution of the corresponding TSP. Set $x_0 := x_N$ and $\tau_i := d_{\mathbb{X}}(x_i, x_{i+1})/L$, $i = 0, \dots, N-1$. Note that

$$N \leq L^{\frac{2}{2\alpha+1}} / \text{diam}(\mathbb{X}) \leq L/d_{\mathbb{X}}(x_i, x_{i+1}),$$

so that $\tau_i \leq N^{-1}$ for all $i = 0, \dots, N-1$. We construct a closed curve $\gamma_L: [0, 1] \rightarrow \mathbb{X}$ that rests in each x_i for a while and then rushes from x_i to x_{i+1} . As in the proof of Proposition 7.7, \mathbb{X} being a compact length space enables us to choose $\gamma_i \in \text{Lip}(\mathbb{X})$ with $\gamma_i(0) = x_i$, $\gamma_i(1) = x_{i+1}$ and $L(\gamma_i) = d_{\mathbb{X}}(x_i, x_{i+1})$. For $i = 0, \dots, N-1$, we define

$$\gamma_L(t) := \begin{cases} x_i & \text{for } t \in [\frac{i}{N}, \frac{i+1}{N} - \tau_i), \\ \gamma_i(\frac{1}{\tau_i}(t - \frac{i+1}{N} + \tau_i)) & \text{for } t \in [\frac{i+1}{N} - \tau_i, \frac{i+1}{N}). \end{cases}$$

By construction, $L(\gamma_L)$ is bounded by $\min_i d(x_i, x_{i+1})\tau_i^{-1} \leq L$. Defining the measure $\nu := (\gamma_L)_* \lambda \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$, the related discrepancy can be estimated by

$$\begin{aligned} \mathcal{D}_K(\mu, \nu) &= \sup_{\|\varphi\|_{H_K(\mathbb{X})} \leq 1} \left| \int_{\mathbb{X}} \varphi d\mu - \int_0^1 \varphi \circ \gamma_L d\lambda \right| \\ &\leq \mathcal{D}_K(\mu, \nu_N) + \sup_{\|\varphi\|_{H_K(\mathbb{X})} \leq 1} \sum_{i=0}^{N-1} \left(\tau_i |\varphi(x_i)| + \left| \int_{\frac{i+1}{N} - \tau_i}^{\frac{i+1}{N}} \varphi \circ \gamma_L d\lambda \right| \right). \end{aligned}$$

The relation (7.12) yields $\mathcal{D}_K(\mu, \nu_N) \leq C L^{-\frac{1}{2\alpha+1}}$ with some constant $C > 0$. Since for $\varphi \in H_K(\mathbb{X})$ it holds $\|\varphi\|_{L^\infty(\mathbb{X})} \leq C_K \|\varphi\|_{H_K(\mathbb{X})}$ with $C_K := \sup_{x \in \mathbb{X}} \sqrt{K(x, x)}$, we finally obtain by Lemma 7.6

$$\begin{aligned} \mathcal{D}_K(\mu, \nu) &\leq C L^{-\frac{1}{2\alpha+1}} + 2 C_K \sum_{i=0}^{N-1} \tau_i \leq C L^{-\frac{1}{2\alpha+1}} + 2 C_K C_{\text{TSP}} \frac{N^\alpha}{L} \\ &\leq (C + 2 C_K C_{\text{TSP}} / \text{diam}(\mathbb{X})) L^{-\frac{1}{2\alpha+1}}. \end{aligned}$$

□

Note that many compact sets in \mathbb{R}^d are compact Ahlfors d -regular length spaces with respect to the Euclidean metric and the normalized Lebesgue measure such as the unit ball or the unit cube. Moreover many compact connected manifolds with or without boundary satisfy these conditions. All assumptions in this section are indeed satisfied for d -dimensional connected, compact Riemannian manifolds without boundary equipped with the Riemannian metric and the normalized Riemannian measure. The latter setting is studied in the subsequent section to refine our investigations on approximation rates.

Remark 7.10. For $\mathbb{X} = \mathbb{T}^d$ with $d \in \mathbb{N}$, the estimate

$$\min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{1}{d}}. \quad (7.13)$$

was derived in [61] provided that K satisfies an additional Lipschitz condition, where the constant in (7.13) depends on d and K . The rate coincides with our rate in (7.11) for $d = 2$ and is worse for higher dimensions as $\frac{d}{3d-2} > \frac{1}{3}$ for all $d \geq 3$.

7.5 Approximation of probability measures having Sobolev densities

To study approximation rates in more detail, we follow the standard strategy in approximation theory and take additional smoothness properties into account. We shall therefore consider μ with a density satisfying smoothness requirements. To define suitable smoothness spaces, we make additional structural assumptions on \mathbb{X} . Throughout the remaining part of the chapter, we suppose that

\mathbb{X} is a d -dimensional connected, compact Riemannian manifold without boundary equipped with the *Riemannian metric* $d_{\mathbb{X}}$ and the *normalized Riemannian measure* $\sigma_{\mathbb{X}}$.

In the first part of this section, we introduce the necessary background on Sobolev spaces and derive general lower bounds for the approximation rates. Then, we focus on upper bounds in the rest of the section. So far, we only have general upper bounds for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$. In case of the smaller spaces $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$ and $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$, we have to restrict to special manifolds \mathbb{X} in order to obtain bounds. For a better overview, all theorems related to approximation rates are named accordingly.

7.5.1 Sobolev spaces and lower bounds

In order to define a smoothness class of functions on \mathbb{X} , let $-\Delta$ denote the (negative) Laplace–Beltrami operator on \mathbb{X} . It is self-adjoint on $L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ and has a sequence of positive, non-decreasing eigenvalues $(\lambda_k)_{k \in \mathbb{N}}$ (with multiplicities) with a corresponding orthonormal complete system of smooth eigenfunctions $\{\phi_k : k \in \mathbb{N}\}$. Every function $f \in L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ has a Fourier expansion

$$f = \sum_{k=0}^{\infty} \hat{f}(k) \phi_k, \quad \hat{f}(k) := \int_{\mathbb{X}} f \overline{\phi_k} d\sigma_{\mathbb{X}}.$$

The *Sobolev space* $H^s(\mathbb{X})$, $s > 0$, is the set of all functions $f \in L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ with distributional derivative $(I - \Delta)^{s/2} f \in L^2(\mathbb{X}, \sigma_{\mathbb{X}})$ and norm

$$\|f\|_{H^s(\mathbb{X})} := \|(I - \Delta)^{s/2} f\|_{L^2(\mathbb{X}, \sigma_{\mathbb{X}})} = \left(\sum_{k=0}^{\infty} (1 + \lambda_k)^s |\hat{f}(k)|^2 \right)^{\frac{1}{2}}.$$

For $s > d/2$, the space $H^s(\mathbb{X})$ is continuously embedded into the space of Hölder continuous functions of degree $s - d/2$, and every function $f \in H^s(\mathbb{X})$ has a uniformly convergent Fourier series, see [235, Thm. 5.7]. Actually, $H^s(\mathbb{X})$, $s > d/2$, is a RKHS with reproducing kernel

$$K(x, y) := \sum_{k=0}^{\infty} (1 + \lambda_k)^{-s} \phi_k(x) \overline{\phi_k(y)}.$$

Hence, the discrepancy $\mathcal{D}_K(\mu, \nu)$ satisfies (7.5) with $H_K(\mathbb{X}) = H^s(\mathbb{X})$. Clearly, each kernel of the above form with coefficients having the same decay as $(1 + \lambda_k)^{-s}$ for $k \rightarrow \infty$ gives rise to a RKHS that coincides with $H^s(\mathbb{X})$ with an equivalent norm. Appendix 7.A contains more details of the above discussion for the torus \mathbb{T}^d , the sphere \mathbb{S}^d , the special orthogonal group $\text{SO}(3)$ and the Grassmannian $\mathcal{G}_{k,d}$.

Now, we are in the position to establish lower bounds on the approximation rates. Again, we want to remark that our results still hold if we drop the requirement that the approximating curves are closed.

Theorem 7.11 (Lower bound). *For $s > d/2$ suppose that $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Assume that μ is absolutely continuous with respect to $\sigma_{\mathbb{X}}$ with a continuous density ρ . Then, there are constants depending on \mathbb{X} , K , and ρ such that*

$$\begin{aligned} N^{-\frac{s}{d}} &\lesssim \min_{\nu \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_N^{\text{emp}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu), \\ L^{-\frac{s}{d-1}} &\lesssim \min_{\nu \in \mathcal{P}_L^{\text{curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu). \end{aligned}$$

Proof. The proof is based on the construction of a suitable fooling function to be used in (7.5) and follows [45, Thm. 2.16]. There exists a ball $B \subset \mathbb{X}$ with $\rho(x) \geq \epsilon = \epsilon(B, \rho)$ for all $x \in B$ and $\sigma_{\mathbb{X}}(B) > 0$, which is chosen as the support

of the constructed fooling functions. We shall verify that for every $\nu \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})$ there exists $\varphi \in H^s(\mathbb{X})$ such that φ vanishes on $\text{supp}(\nu)$ but

$$\int_B \varphi \, d\mu \gtrsim \|\varphi\|_{H^s(\mathbb{X})} N^{-\frac{s}{d}}, \quad (7.14)$$

where the constant depends on \mathbb{X} , K , and ρ . For small enough δ we can choose $2N$ disjoint balls in B with diameters $\delta N^{-1/d}$, see also [125]. For $\nu \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})$, there are N of these balls that do not intersect with $\text{supp}(\nu)$. By putting together bump functions supported on each of the N balls, we obtain a non-negative function φ supported in B that vanishes on $\text{supp}(\nu)$ and satisfies (7.14), with a constant that depends on ϵ , cf. [45, Thm. 2.16]. This yields

$$\left| \int_{\mathbb{X}} \varphi \, d\mu - \int_{\mathbb{X}} \varphi \, d\nu \right| = \int_B \varphi \, d\mu \gtrsim \|\varphi\|_{H^s(\mathbb{X})} N^{-\frac{s}{d}}.$$

The inequality for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$ is derived in a similar way. Given a continuous curve $\gamma: [0, 1] \rightarrow \mathbb{X}$ of length L , choose N such that $L \leq \delta N N^{-1/d}$. By taking half of the radius of the above balls, there are $2N$ pairwise disjoint balls of radius $\frac{\delta}{2} N^{-1/d}$ contained in B with pairwise distances at least $\delta N^{-1/d}$. Any curve of length $\delta N N^{-1/d}$ intersects at most N of those balls. Hence, there are N balls of radius $\frac{\delta}{2} N^{-1/d}$ that do not intersect $\text{supp}(\gamma)$. As above, this yields a fooling function φ satisfying (7.14), which ends the proof. \square

7.5.2 Upper bounds for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$

In this section, we derive upper bounds that match the lower bounds in Theorem 7.11 for $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$. Our analysis makes use of the following theorem, which was already proved for $\mathbb{X} = \mathbb{S}^d$ in [158].

Theorem 7.12. [45, Thm. 2.12] *Assume that $\nu_r \in \mathcal{P}(\mathbb{X})$ provides an exact quadrature for all eigenfunctions φ_k of the Laplace–Beltrami operator with eigenvalues $\lambda_k \leq r^2$, i.e.,*

$$\int_{\mathbb{X}} \varphi_k \, d\sigma_{\mathbb{X}} = \int_{\mathbb{X}} \varphi_k \, d\nu_r. \quad (7.15)$$

Then, it holds for every function $f \in H^s(\mathbb{X})$, $s > d/2$, that there is a constant depending on \mathbb{X} and s with

$$\left| \int_{\mathbb{X}} f \, d\sigma_{\mathbb{X}} - \int_{\mathbb{X}} f \, d\nu_r \right| \lesssim r^{-s} \|f\|_{H^s(\mathbb{X})}.$$

For our estimates it is important that the number of eigenfunctions of the Laplace–Beltrami operator on \mathbb{X} belonging to eigenvalues with $\lambda_k \leq r^2$ is of order r^d , see [62, Chap. 6.4] and [163, Thm. 17.5.3, Cor. 17.5.8]. This is known as Weyl’s estimates on the spectrum of an elliptic operator. For some special manifolds, the eigenfunctions are explicitly given in the appendix. In the following lemma, the result from Theorem 7.12 is rewritten in terms of discrepancies and generalized to absolutely continuous measures with densities $\rho \in H^s(\mathbb{X})$.

Lemma 7.13. *For $s > d/2$ suppose that $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms and that $\nu_r \in \mathcal{P}(\mathbb{X})$ satisfies (7.15). Let $\mu \in \mathcal{P}(\mathbb{X})$ be absolutely continuous with respect to $\sigma_{\mathbb{X}}$ with density $\rho \in H^s(\mathbb{X})$. For sufficiently large r , the measures $\tilde{\nu}_r := \frac{\rho}{\beta_r} \nu_r \in \mathcal{P}(\mathbb{X})$ with $\beta_r := \int_{\mathbb{X}} \rho d\nu_r$ are well defined and there is a constant depending on \mathbb{X} and K with*

$$\mathcal{D}_K(\mu, \tilde{\nu}_r) \lesssim \|\rho\|_{H^s(\mathbb{X})} r^{-s}.$$

Proof. Note that $H^s(\mathbb{X})$ is a Banach algebra with respect to addition and multiplication [78], in particular, for $f, g \in H^s(\mathbb{X})$ we have $fg \in H^s(\mathbb{X})$ with

$$\|fg\|_{H^s(\mathbb{X})} \leq \|f\|_{H^s(\mathbb{X})} \|g\|_{H^s(\mathbb{X})}. \quad (7.16)$$

By Theorem 7.12, we obtain for all $\varphi \in H^s(\mathbb{X})$ that

$$\left| \int_{\mathbb{X}} \varphi \rho d\sigma_{\mathbb{X}} - \int_{\mathbb{X}} \varphi \rho d\nu_r \right| \lesssim r^{-s} \|\varphi \rho\|_{H^s(\mathbb{X})} \lesssim r^{-s} \|\varphi\|_{H^s(\mathbb{X})} \|\rho\|_{H^s(\mathbb{X})}. \quad (7.17)$$

In particular, this implies for $\varphi \equiv 1$ that

$$|1 - \beta_r| \lesssim r^{-s} \|\rho\|_{H^s(\mathbb{X})}. \quad (7.18)$$

Then, application of the triangle inequality results in

$$\left| \int_{\mathbb{X}} \varphi d\mu - \int_{\mathbb{X}} \varphi d\tilde{\nu}_r \right| \leq \left| \int_{\mathbb{X}} \varphi d\mu - \int_{\mathbb{X}} \varphi \rho d\nu_r \right| + \left| \int_{\mathbb{X}} \varphi \rho \frac{\beta_r - 1}{\beta_r} d\nu_r \right|.$$

According to (7.17), the first summand is bounded by $\lesssim r^{-s} \|\varphi\|_{H^s(\mathbb{X})} \|\rho\|_{H^s(\mathbb{X})}$. It remains to derive matching bounds on the second term. Hölder's inequality yields

$$\left| \int_{\mathbb{X}} \varphi \rho \frac{\beta_r - 1}{\beta_r} d\nu_r \right| \lesssim \|\varphi\|_{L^\infty(\mathbb{X})} |\beta_r - 1| \lesssim \|\varphi\|_{H^s(\mathbb{X})} r^{-s} \|\rho\|_{H^s(\mathbb{X})},$$

where the last inequality is due to $H^s(\mathbb{X}) \hookrightarrow L^\infty(\mathbb{X})$ and (7.18). \square

Using the previous lemma, we derive optimal approximation rates for $\mathcal{P}_N^{\text{atom}}(\mathbb{X})$ and $\mathcal{P}_L^{\text{curv}}(\mathbb{X})$.

Theorem 7.14 (Upper bounds). *For $s > d/2$ suppose that $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Assume that μ is absolutely continuous with respect to $\sigma_{\mathbb{X}}$ with density $\rho \in H^s(\mathbb{X})$. Then, there are constants depending on \mathbb{X} and K such that*

$$\min_{\nu \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim \|\rho\|_{H^s(\mathbb{X})} N^{-\frac{s}{d}}, \quad (7.19)$$

$$\min_{\nu \in \mathcal{P}_L^{\text{curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim \|\rho\|_{H^s(\mathbb{X})} L^{-\frac{s}{d-1}}. \quad (7.20)$$

Proof. By [45, Lem. 2.11] and since the Laplace–Beltrami has $N \sim r^d$ eigenfunctions belonging to eigenvectors $\lambda_k < r^2$, there exists a measure $\nu_r \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})$ that satisfies (7.15). Hence, (7.15) is satisfied with $r \sim N^{1/d}$, where the constants depend on \mathbb{X} and K . Thus, Lemma 7.13 with $\tilde{\nu}_r \in \mathcal{P}_N^{\text{atom}}(\mathbb{X})$ leads to (7.19).

The assumptions of Lemma 7.6 are satisfied, so that analogous arguments as in the proof of Theorem 7.8 yield $\mathcal{P}_N^{\text{atom}}(\mathbb{X}) \subset \mathcal{P}_L^{\text{curv}}(\mathbb{X})$ with suitable $N \sim L^{d/(d-1)}$. Hence, (7.19) implies (7.20). \square

7.5.3 Upper bounds for $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$ and special manifolds \mathbb{X}

To establish upper bounds for the smaller space $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$, restriction to special manifolds is necessary. The basic idea consists in the construction of a curve and a related measure ν_r such that all eigenfunctions of the Laplace–Beltrami operator belonging to eigenvalues smaller than a certain value are exactly integrated by this measure and then applying Lemma 7.13 for estimating the minimum of discrepancies. We begin with the torus.

Theorem 7.15 (Torus). *Let $\mathbb{X} = \mathbb{T}^d$ with $d \in \mathbb{N}$, $s > d/2$ and suppose that $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Then, for any absolutely continuous measure $\mu \in \mathcal{P}(\mathbb{X})$ with Lipschitz continuous density $\rho \in H^s(\mathbb{X})$, there exists a constant depending on d , K , and ρ such that*

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s}{d-1}}.$$

Proof. 1. First, we construct a closed curve γ_r such that the trigonometric polynomials from $\Pi_r(\mathbb{T}^d)$, see (7.33) in the appendix, are exactly integrated along this curve. Clearly, the polynomials in $\Pi_r(\mathbb{T}^{d-1})$ are exactly integrated at equispaced nodes $x_{\mathbf{k}} = \frac{\mathbf{k}}{n}$, $\mathbf{k} = (k_1, \dots, k_{d-1}) \in \mathbb{N}_0^{d-1}$, $0 \leq k_i \leq n-1$, with weights $1/n^{d-1}$, where $n := r+1$. Set $z(\mathbf{k}) := k_1 + k_2 n + \dots + k_{d-1} n^{d-2}$ and consider the curves

$$\gamma_{\mathbf{k}}: I_{\mathbf{k}} := \left[\frac{z(\mathbf{k})}{n^{d-1}}, \frac{z(\mathbf{k})+1}{n^{d-1}} \right] \rightarrow \mathbb{T}^d \quad \text{with} \quad \gamma_{\mathbf{k}}(t) := \left(\frac{x_{\mathbf{k}}}{n^{d-1}t} \right).$$

Then, each element in Π_r^d is exactly integrated along the union of these curves, i.e., using $I := \{0, \dots, n-1\}^{d-1}$, we have

$$\int_{\mathbb{T}^d} p \, d\sigma_{\mathbb{T}^d} = \sum_{\mathbf{k} \in I} \int_{I_{\mathbf{k}}} p \circ \gamma_{\mathbf{k}} \, d\lambda, \quad p \in \Pi_r^d.$$

The argument is repeated for every other coordinate direction, so that we end up with dn^{d-1} curves mapping from an interval of length $\frac{1}{dn^{d-1}}$ to \mathbb{T}^d . The intersection points of these curves are considered as vertices of a graph, where each vertex has $2d$ edges. Consequently, there exists an Euler path $\gamma_r: [0, 1] \rightarrow \mathbb{T}^d$ through the vertices build from all curves. It has constant speed dn^{d-1} and the polynomials Π_r^d are exactly integrated along γ_r , i.e.,

$$\int_{\mathbb{T}^d} p \, d\sigma_{\mathbb{T}^d} = \int_{\mathbb{T}^d} p \, d\gamma_{r*} \lambda, \quad p \in \Pi_r^d.$$

2. Next, we apply Lemma 7.13 for $\nu_r = \gamma_{r*} \lambda$. We observe $\tilde{\nu}_r = \gamma_{r*}((\rho \circ \gamma_r)/\beta_r \lambda)$ and deduce $L(\rho \circ \gamma_r/\beta_r) \leq L(\gamma_r)L(\rho)/\beta_r \lesssim r^{d-1} \sim L$ as $\beta_r \sim 1$. Here, constants depend on d , K , and ρ . \square

Now, we provide approximation rates for $\mathbb{X} = \mathbb{S}^d$.

Theorem 7.16 (Sphere). *Let $\mathbb{X} = \mathbb{S}^d$ with $d \geq 2$, $s > d/2$ and suppose that $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Then, we have for any absolutely continuous measure $\mu \in \mathcal{P}(\mathbb{X})$ with Lipschitz continuous density $\rho \in H^s(\mathbb{X})$ that there is a constant depending on d , K , and ρ with*

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s}{d-1}}.$$

Proof. 1. First, we construct a constant speed curve $\gamma_r: [0, 1] \rightarrow \mathbb{S}^d$ and a probability measure $\omega_r = \rho_r \lambda$ with Lipschitz continuous density $\rho_r: [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ such that for all $p \in \Pi_r(\mathbb{S}^d)$, it holds

$$\int_{\mathbb{S}^d} p \, d\sigma_{\mathbb{S}^d} = \int_0^1 p \circ \gamma_r \, d\omega_r. \quad (7.21)$$

Utilizing spherical coordinates

$$x_1 = \cos \theta_1, \quad x_2 = \sin \theta_1 \cos \theta_2, \quad \dots, \quad x_d = \prod_{j=1}^{d-1} \sin \theta_j \cos \phi, \quad x_{d+1} = \prod_{j=1}^{d-1} \sin \theta_j \sin \phi, \quad (7.22)$$

where $\theta_k \in [0, \pi]$, $k = 1, \dots, d-1$, and $\phi \in [0, 2\pi)$, we obtain

$$\int_{\mathbb{S}^d} p \, d\sigma_{\mathbb{S}^d} = \int_0^\pi c_d \sin(\theta_1)^{d-1} \int_{\mathbb{S}^{d-1}} p(\cos(\theta_1), \sin(\theta_1)\tilde{x}) \, d\sigma_{\mathbb{S}^{d-1}}(\tilde{x}) \, d\theta_1, \quad (7.23)$$

where $c_d := (\int_0^\pi \sin(\theta)^{d-1} \, d\theta)^{-1}$. There exist nodes $\tilde{x}_i \in \mathbb{S}^{d-1}$ and positive weights a_i , $i = 1, \dots, n \sim r^{d-1}$, with $\sum_{i=1}^n a_i = 1$, such that for all $p \in \Pi_r(\mathbb{S}^{d-1})$ it holds

$$\int_{\mathbb{S}^{d-1}} p \, d\sigma_{\mathbb{S}^{d-1}} = \sum_{i=1}^n a_i p(\tilde{x}_i).$$

To see this, substitute $u_k = \sin \theta_k$, $k = 2, \dots, d-1$, apply Gaussian quadrature with nodes $[(r+1)/2]$ and corresponding weights to exactly integrate over u_k , and equispaced nodes and weights $1/(2r+1)$ for the integration over ϕ as, e.g., in [275]. Then, we define $\gamma_r: [0, 1] \rightarrow \mathbb{S}^d$ for $t \in [(i-1)/n, i/n]$, $i = 1, \dots, n$, by

$$\gamma_r(t) := \gamma_{r,i}(2\pi nt), \quad \gamma_{r,i}(\alpha) := (\cos(\alpha), \sin(\alpha)\tilde{x}_i), \quad \alpha \in [0, 2\pi].$$

Since $(1, 0, \dots, 0) = \gamma_{r,i}(0) = \gamma_{r,i}(2\pi)$ for all $i = 1, \dots, n$, the curve is closed. Furthermore, $\gamma_r(t)$ has constant speed since for $i = 1, \dots, n$, i.e.,

$$|\dot{\gamma}_r|(t) = |\dot{\gamma}_{r,i}|(2\pi nt) = 2\pi n \sim r^{d-1}.$$

Next, the density $\rho_r: [0, 1] \rightarrow \mathbb{R}$ is defined for $t \in [(i-1)/n, i/n]$, $i = 1, \dots, n$, by

$$\rho_r(t) := \rho_{r,i}(2\pi nt), \quad \rho_{r,i}(\alpha) := a_i c_d \pi n |\sin(\alpha)|^{d-1}, \quad \alpha \in [0, 2\pi].$$

We directly verify that ρ_r is Lipschitz continuous with $L(\rho_r) \lesssim \max_i a_i n^2$. By [115], the quadrature weights fulfill $a_i \lesssim \frac{1}{r^{d-1}}$ so that $L(\rho_r) \lesssim n^2 r^{-(d-1)} \sim r^{d-1}$. By

definition of the constant c_d and weights a_i , we see that ρ_r is indeed a probability density

$$\begin{aligned} \int_0^1 \rho_r d\lambda &= \sum_{i=1}^n \int_{\frac{i-1}{n}}^{\frac{i}{n}} \rho_{r,i}(2\pi nt) dt = \frac{1}{2\pi n} \sum_{i=1}^n \int_0^{2\pi} \rho_{r,i}(\alpha) d\alpha \\ &= \frac{c_d}{2} \sum_{i=1}^n a_i \int_0^{2\pi} |\sin(\theta)|^{d-1} d\theta = 1. \end{aligned}$$

For $p \in \Pi_r(\mathbb{S}^d)$, we obtain

$$\begin{aligned} &\int_0^1 p \circ \gamma_r \rho_r d\lambda \\ &= \sum_{i=1}^n \int_{\frac{i-1}{n}}^{\frac{i}{n}} p(\gamma_{r,i}(2\pi nt)) \rho_{r,i}(2\pi nt) dt = \int_0^{2\pi} \frac{1}{2\pi n} \sum_{i=1}^n p(\gamma_{r,i}(\alpha)) \rho_{r,i}(\alpha) d\alpha \\ &= \frac{c_d}{2} \int_0^{2\pi} |\sin(\alpha)|^{d-1} \sum_{i=1}^n a_i p(\cos(\alpha), \sin(\alpha) \tilde{x}_i) d\alpha \\ &= \frac{c_d}{2} \int_0^\pi |\sin(\alpha)|^{d-1} \sum_{i=1}^n a_i \left(p(\cos(\alpha), \sin(\alpha) \tilde{x}_i) + p(-\cos(\alpha), -\sin(\alpha) \tilde{x}_i) \right) d\alpha. \end{aligned}$$

Without loss of generality, p is chosen as a homogeneous polynomial of degree $k \leq r$, i.e., $p(tx) = t^k p(x)$. Then,

$$\int_0^1 p \circ \gamma_r \rho_r d\lambda = \frac{1 + (-1)^k}{2} \int_0^\pi c_d |\sin(\alpha)|^{d-1} \sum_{i=1}^n a_i p(\cos(\alpha), \sin(\alpha) \tilde{x}_i) d\alpha,$$

and regarding that for fixed $\alpha \in [0, 2\pi]$ the function $\tilde{x} \mapsto p(\cos(\alpha), \sin(\alpha) \tilde{x})$ is a polynomial of degree at most r on \mathbb{S}^{d-1} , we conclude

$$\int_0^1 p \circ \gamma_r \rho_r d\lambda = \frac{1 + (-1)^k}{2} \int_0^\pi c_d |\sin(\alpha)|^{d-1} \int_{\mathbb{S}^{d-1}} p(\cos(\alpha), \sin(\alpha) \tilde{x}) d\sigma_{\mathbb{S}^{d-1}}(\tilde{x}) d\alpha.$$

Now, the assertion (7.21) follows from (7.23) and since $\int_{\mathbb{S}^d} p d\sigma_{\mathbb{S}^d} = 0$ if k is odd.

2. Next, we apply Lemma 7.13 for $\nu_r = \gamma_{r*} \rho_r$, from which we obtain that $\tilde{\nu}_r = \gamma_{r*}((\rho \circ \gamma_r) \rho_r / \beta_r \lambda)$. As all ρ_r are uniformly bounded by construction and ρ is bounded due to continuity, we conclude using $L(\rho_r) \lesssim r^{d-1}$ and $L(\gamma_r) \sim r^{d-1}$ that

$$L(\rho \circ \gamma_r \rho_r / \beta_r) \leq (L(\rho \circ \gamma_r) \|\rho_r\|_\infty + L(\rho_r) \|\rho\|_\infty) / \beta_r \lesssim (L(\rho) + \|\rho\|_\infty) r^{d-1},$$

which concludes the proof. \square

Finally, we derive approximation rates for $\mathbb{X} = \text{SO}(3)$.

Corollary 7.17 (Special orthogonal group). *Let $\mathbb{X} = \text{SO}(3)$, $s > 3/2$ and suppose $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Then, we have for any absolutely continuous measure $\mu \in \mathcal{P}(\mathbb{X})$ with Lipschitz continuous density $\rho \in H^s(\mathbb{X})$ that*

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s}{d-1}},$$

where the constant depends on K and ρ .

Proof. 1. For fixed $L \sim r^2$, we shall construct a curve $\gamma_r: [0, 1] \rightarrow \text{SO}(3)$ with $L(\gamma_r) \lesssim L$ and a probability measure $\omega_r = \rho_r \lambda$ with density $\rho_r: [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ and $L(\rho_r) \lesssim L$, such that

$$\int_{\text{SO}(3)} p \, d\sigma_{\text{SO}(3)} = \int_{\text{SO}(3)} p \, d\gamma_{r*}(\rho_r \lambda).$$

We use the fact that the sphere \mathbb{S}^3 is a double covering of $\text{SO}(3)$. That is, there is a surjective two-to-one mapping $a: \mathbb{S}^3 \rightarrow \text{SO}(3)$ satisfying $a(x) = a(-x)$, $x \in \mathbb{S}^3$. Moreover, we know that $a: \mathbb{S}^3 \rightarrow \text{SO}(3)$ is a local isometry, see [130], i.e., it respects the Riemannian structures, implying the relations $\sigma_{\text{SO}(3)} = a_* \sigma_{\mathbb{S}^3}$ and

$$d_{\text{SO}(3)}(a(x_1), a(x_2)) = \min(d_{\mathbb{S}^3}(x_1, x_2), d_{\mathbb{S}^3}(x_1, -x_2)).$$

It also maps $\Pi_r(\text{SO}(3))$ into $\Pi_{2r}(\mathbb{S}^3)$, i.e., $p \in \Pi_r(\text{SO}(3))$ implies $p \circ a \in \Pi_{2r}(\mathbb{S}^3)$. Now, let $\tilde{\gamma}_r: [0, 1] \rightarrow \mathbb{S}^3$ and $\tilde{\omega}_r$ be given as in the first part of the proof of Theorem 7.16 for $d = 3$, i.e., $\tilde{\gamma}_{r*} \tilde{\omega}_r$ satisfies (7.21) with $L(\tilde{\gamma}_r) \lesssim L$ and $\tilde{\omega}_r = \tilde{\rho}_r \lambda$ with $L(\tilde{\rho}_r) \lesssim L$.

We now define a curve γ_r in $\text{SO}(3)$ by

$$\gamma_r: [0, 1] \rightarrow \text{SO}(3), \quad \gamma_r(t) := a \circ \tilde{\gamma}_{2r}(t),$$

and let $\omega_r := \tilde{\omega}_{2r}$. For $p \in \Pi_r(\text{SO}(3))$, the push-forward measure $\gamma_{r*} \omega_r$ leads to

$$\begin{aligned} \int_{\text{SO}(3)} p \, d\sigma_{\text{SO}(3)} &= \int_{\text{SO}(3)} p \, da_* \sigma_{\mathbb{S}^3} = \int_{\mathbb{S}^3} p \circ a \, d\sigma_{\mathbb{S}^3} \\ &= \int_{\mathbb{S}^3} p \circ a \, d\tilde{\gamma}_{2r*} \tilde{\omega}_{2r} = \int_{\text{SO}(3)} p \, d\gamma_{r*} \omega_r. \end{aligned}$$

Hence, property (7.15) is satisfied for $\gamma_{r*} \omega_r = \gamma_{r*}(\tilde{\rho}_{2r} \lambda)$.

2. The rest follows along the lines of step 2. in the proof of Theorem 7.16. \square

7.5.4 Upper bounds for $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ and special manifolds \mathbb{X}

To derive upper bounds for the smallest space $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$, we need the following specification of Lemma 7.13.

Lemma 7.18. *For $s > d/2$ suppose that $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Let $\mu \in \mathcal{P}(\mathbb{X})$ be absolutely continuous with respect to $\sigma_{\mathbb{X}}$ with positive density $\rho \in H^s(\mathbb{X})$. Suppose that $\nu_r := \gamma_{r*} \lambda$ with $\gamma_r \in \text{Lip}(\mathbb{X})$ satisfies (7.15) and let $\beta_r := \int_{\mathbb{X}} \rho \, d\nu_r$. Then, for sufficiently large r ,*

$$g: [0, 1] \rightarrow [0, 1], \quad g(t) := \frac{1}{\beta_r} \int_0^t \rho \circ \gamma_r \, d\lambda$$

is well-defined and invertible. Moreover, $\tilde{\gamma}_r := \gamma_r \circ g^{-1}$ satisfies $L(\tilde{\gamma}_r) \lesssim L(\gamma_r)$ and

$$\mathcal{D}_K(\mu, \tilde{\gamma}_{r*} \lambda) \lesssim r^{-s}, \tag{7.24}$$

where the constants depend on \mathbb{X} , K , and ρ .

Proof. Since ρ is continuous, there is $\epsilon > 0$ with $\rho \geq \epsilon$. To bound the Lipschitz constant $L(\tilde{\gamma}_r)$, we apply the mean value theorem together with the definition of g and the fact that $(g^{-1})'(s) = 1/g'(g^{-1}(s))$ to obtain

$$|\tilde{\gamma}_r(s) - \tilde{\gamma}_r(t)| \leq L(\gamma_r)|g^{-1}(s) - g^{-1}(t)| \leq L(\gamma_r) \frac{\beta_r}{\epsilon} |s - t|.$$

Using (7.18), this can be further estimated for sufficiently large r as

$$|\tilde{\gamma}_r(s) - \tilde{\gamma}_r(t)| \lesssim L(\gamma_r) \frac{1 + \|\rho\|_{H^s(\mathbb{X})} r^{-s}}{\epsilon} |s - t| \lesssim L(\gamma_r) \frac{2}{\epsilon} |s - t|.$$

To derive (7.24), we aim to apply Lemma 7.13 with $\nu_r = \gamma_{r*}\lambda$. We observe

$$\tilde{\nu}_r = \frac{\rho}{\beta_r} \gamma_{r*}\lambda = \gamma_{r*}\left(\frac{\rho \circ \gamma_r}{\beta_r} \lambda\right) = \gamma_{r*}(g'\lambda) = (\gamma_r \circ g^{-1})_*\lambda = \tilde{\gamma}_{r*}\lambda,$$

so that Lemma 7.13 indeed implies (7.24). \square

In comparison to Theorem 7.15, we now trade the Lipschitz condition on ρ with the positivity requirement, which enables us to cover $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$.

Theorem 7.19 (Torus). *Let $\mathbb{X} = \mathbb{T}^d$ with $d \in \mathbb{N}$, $s > d/2$ and suppose that $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Then, for any absolutely continuous measure $\mu \in \mathcal{P}(\mathbb{X})$ with positive density $\rho \in H^s(\mathbb{X})$, there is a constant depending on d , K , and ρ with*

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s}{d-1}}.$$

Proof. The first part of the proof is identical to the proof of Theorem 7.15. Instead of Lemma 7.13 though, we now apply Lemma 7.18 for γ_r and $\rho_r \equiv 1$. Hence, $\tilde{\gamma}_r = \gamma_r \circ g_r^{-1}$ satisfies $L(\tilde{\gamma}_r) \leq \frac{\beta_r}{\epsilon} d(2r+1)^{d-1} \lesssim r^{d-1}$, so that $\tilde{\gamma}_{r*}\lambda$ satisfies (7.24) and is in $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ with $L \sim r^{d-1}$. \square

The construction on $\mathbb{X} = \mathbb{S}^d$ for $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X})$ in the proof of Theorem 7.16 is not compatible with $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$. Thus, the situation is different from the torus, where we have used the same underlying construction and only switched from Lemma 7.13 to Lemma 7.18. Now, we present a new construction for $\mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$, which is tailored to $\mathbb{X} = \mathbb{S}^2$. In this case, we can transfer the ideas of the torus, but with Gauss-Legendre quadrature points.

Theorem 7.20 (2-sphere). *Let $\mathbb{X} = \mathbb{S}^2$, $s > 1$ and suppose $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Then, we have for any absolutely continuous measure $\mu \in \mathcal{P}(\mathbb{X})$ with positive density $\rho \in H^s(\mathbb{X})$ that there is a constant depending on K and ρ with*

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-s}.$$

Proof. 1. We construct closed curves such that the spherical polynomials from $\Pi_r(\mathbb{S}^2)$, see (7.35) in the appendix, are exactly integrated along this curve. It suffices to show this for the polynomials $p(x) = x^{k_1}x^{k_2}x_3^{k_3} \in \Pi_r(\mathbb{S}^2)$ with $k_1 + k_2 + k_3 \leq r$ restricted to \mathbb{S}^2 . We select $n = \lceil (r+1)/2 \rceil$ Gauss-Legendre quadrature points $u_j = \cos(\theta_j) \in [-1, 1]$ and corresponding weights $2\omega_j$, $j = 1, \dots, n$. Note that $\sum_{j=1}^n \omega_j = 1$. Using spherical coordinates $x_1 = \cos(\theta)$, $x_2 = \sin(\theta) \cos(\phi)$, and $x_3 = \sin(\theta) \sin(\phi)$ with $(\theta, \phi) \in [0, \pi] \times [0, 2\pi]$, we obtain

$$\begin{aligned} \int_{\mathbb{S}^2} p \, d\sigma_{\mathbb{S}^2} &= \frac{1}{4\pi} \int_0^{2\pi} \cos(\phi)^{k_2} \sin(\phi)^{k_3} \int_0^\pi \cos(\theta)^{k_1} \sin(\theta)^{k_2+k_3} \sin(\phi) \, d\theta \, d\phi \\ &= \frac{1}{4\pi} \int_0^{2\pi} \cos(\phi)^{k_2} \sin(\phi)^{k_3} \int_{-1}^1 u^{k_1} (1-u^2)^{\frac{k_2+k_3}{2}} \, du \, d\phi, \end{aligned}$$

see also [276]. If $k_2 + k_3$ is odd, then the integral over ϕ becomes zero. If $k_2 + k_3$ is even, the inner integrand is a polynomial of degree $\leq r$. In both cases we get

$$\int_{\mathbb{S}^2} p \, d\sigma_{\mathbb{S}^2} = \frac{1}{2\pi} \sum_{j=1}^n \omega_j \int_0^{2\pi} p(\cos(\theta_j), \sin(\theta_j) \cos(\phi), \sin(\theta_j) \sin(\phi)) \, d\phi.$$

Substituting in each summand $\phi = 2\pi t/\omega_j$, $j = 1, \dots, n$, yields

$$\int_{\mathbb{S}^2} p \, d\sigma_{\mathbb{S}^2} = \sum_{j=1}^n \int_0^{\omega_j} p \circ \gamma_j \, d\lambda,$$

where $\gamma_j: [0, \omega_j] \rightarrow \mathbb{S}^2$ is defined by

$$\gamma_j(t) := (\cos(\theta_j), \sin(\theta_j) \cos(2\pi t/\omega_j), \sin(\theta_j) \sin(2\pi t/\omega_j)),$$

and has constant speed $L(\gamma_j) = 2\pi \sin(\theta_j)/\omega_j$. The lower bound $\omega_j \gtrsim \frac{1}{n} \sin(\theta_j)$, cf. [115], implies that $L(\gamma_j) \lesssim n$. Defining a curve $\tilde{\gamma}: [0, 1] \rightarrow \mathbb{S}^2$ piecewise via

$$\tilde{\gamma}|_{[0, s_1]} = \gamma_1, \quad \tilde{\gamma}|_{[s_1, s_2]} = \gamma_2(\cdot - s_1), \quad \dots, \quad \tilde{\gamma}|_{[s_{n-1}, 1]} = \gamma_n(\cdot - s_{n-1}),$$

where $s_j := \omega_1 + \dots + \omega_j$, we obtain

$$\int_{\mathbb{S}^2} p \, d\sigma_{\mathbb{S}^2} = \int_0^1 p \, d\tilde{\gamma}_* \lambda, \quad p \in \Pi_r(\mathbb{S}^2).$$

Further, the curve satisfies $L(\tilde{\gamma}) \lesssim r$.

As with the torus, we now “turn” the sphere (or switch the position of ϕ) so that we get circles along orthogonal directions. This large collection of circles is indeed connected. As with the torus, each intersection point has an incoming and outgoing part of a circle, so that all this corresponds to a graph, where again each vertex has an even number of “edges”. Hence, there is an Euler path inducing our final curve $\gamma_r: [0, 1] \rightarrow \mathbb{S}^2$ with piecewise constant speed $L(\gamma_r) \lesssim r$ satisfying

$$\int_{\mathbb{S}^2} p \, d\sigma_{\mathbb{S}^2} = \int_0^1 p \, d(\gamma_{r*} \lambda), \quad p \in \Pi_r(\mathbb{S}^2).$$

2. Let $r \sim L$. Analogous to the end of the proof of Theorem 7.19, Lemma 7.18 now yields the assertion. \square

To get the approximation rate for $\mathbb{X} = \mathcal{G}_{2,4}$, we make use of its double covering $\mathbb{X} = \mathbb{S}^2 \times \mathbb{S}^2$, cf. Remark 7.27.

Theorem 7.21 (Grassmannian). *Let $\mathbb{X} = \mathcal{G}_{2,4}$, $s > 2$ and suppose $H_K(\mathbb{X}) = H^s(\mathbb{X})$ holds with equivalent norms. Then, we have for any absolutely continuous measure $\mu \in \mathcal{P}(\mathbb{X})$ with positive density $\rho \in H^s(\mathbb{X})$ that there exists a constant depending on K and ρ with*

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \leq \min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s}{3}}.$$

Proof. By Remark 7.27 in the appendix, we know that $\mathcal{G}_{2,4} \cong \mathbb{S}^2 \times \mathbb{S}^2 / \{\pm 1\}$ so that it remains to prove the assertion for $\mathbb{X} = \mathbb{S}^2 \times \mathbb{S}^2$.

There exist pairwise distinct points $\{x_1, \dots, x_N\} \subset \mathbb{S}^2$ such that $\frac{1}{N} \sum_{j=1}^N \delta_{x_j}$ satisfies (7.15) on \mathbb{S}^2 with $N \sim r^2$, cf. [37, 38]. On the other hand, let $\tilde{\gamma}$ be the curve on \mathbb{S}^2 constructed in the proof of Theorem 7.20, so that $\tilde{\gamma}_* \lambda$ satisfies (7.15) on \mathbb{S}^2 with $\ell(\tilde{\gamma}) \leq L(\tilde{\gamma}) \sim r$. Let us introduce the virtual point $x_{N+1} := x_1$. The curve $\tilde{\gamma}([0, 1])$ contains a great circle. Thus, for each pair x_j and x_{j+1} there is $O_j \in O(3)$ such that $x_j, x_{j+1} \in \Gamma_j := O_j \tilde{\gamma}([0, 1])$. It turns out that the set on $\mathbb{S}^2 \times \mathbb{S}^2$ given by $\bigcup_{j=1}^N (\{x_j\} \times \Gamma_j) \cup (\Gamma_j \times \{x_{j+1}\})$ is connected. We now choose $\gamma_j := O_j \tilde{\gamma}$ and know that the union of the trajectories of the set of curves

$$t \mapsto (x_j, \gamma_j(t)), \quad t \mapsto (\gamma_j(t), x_{j+1}), \quad j = 1, \dots, N,$$

is connected. Combinatorial arguments involving Euler paths, see Theorems 7.15 and 7.20, lead to a curve γ with $\ell(\gamma) \leq L(\gamma) \sim NL(\tilde{\gamma}) \sim r^3$, so that $\gamma_* \lambda$ satisfies (7.15). The remaining part follows along the lines of the proof of Theorem 7.16. \square

Our approximation results can be extended to diffeomorphic manifolds, e.g., from \mathbb{S}^2 to ellipsoids, see also the 3D-torus example in Section 7.8. To this end, recall that we can describe the Sobolev space $H^s(\mathbb{X})$ using local charts, see [260, Sec. 7.2]. The exponential maps $\exp_x: T_x \mathbb{X} \rightarrow \mathbb{X}$ give rise to local charts $(\tilde{B}_x(r_0), \exp_x^{-1})$, where $\tilde{B}_x(r_0) := \{y \in \mathbb{X} : d_{\mathbb{X}}(x, y) < r_0\}$ denotes the geodesic balls around x with the injectivity radius r_0 . If $\delta < r_0$ is chosen small enough, there exists a uniformly locally finite covering of \mathbb{X} by a sequence of balls $(\tilde{B}_{x_j}(\delta))_j$ with a corresponding smooth resolution of unity $(\psi_j)_j$ with $\text{supp}(\psi_j) \subset \tilde{B}_{x_j}(\delta)$, see [260, Prop. 7.2.1]. Then, an equivalent Sobolev norm is given by

$$\|f\|_{H^s(\mathbb{X})} := \left(\sum_{j=1}^{\infty} \|(\psi_j f) \circ \exp_{x_j}\|_{H^s(\mathbb{R}^d)}^2 \right)^{\frac{1}{2}}, \quad (7.25)$$

where $(\psi_j f) \circ \exp_{x_j}$ is extended to \mathbb{R}^d by zero, see [260, Thm. 7.4.5]. Using Definition (7.25), we are able to pull over results from the Euclidean setting.

Proposition 7.22. *Let $\mathbb{X}_1, \mathbb{X}_2$ be two d -dimensional connected, compact Riemannian manifolds without boundary, which are $s + 1$ diffeomorphic with $s > d/2$.*

Assume that for $H_K(\mathbb{X}_2) = H^s(\mathbb{X}_2)$ and every absolutely continuous measure μ with positive density $\rho \in H^s(\mathbb{X}_2)$ it holds

$$\min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s}{d-1}},$$

where the constant depends on \mathbb{X}_2 , K , and ρ . Then, the same property holds for \mathbb{X}_1 , where the constant additionally depends on the diffeomorphism.

Proof. Let $f: \mathbb{X}_2 \rightarrow \mathbb{X}_1$ denote such a diffeomorphism and $\rho \in H^s(\mathbb{X}_1)$ the density of the measure μ on \mathbb{X}_1 . Any curve $\tilde{\gamma}: [0, 1] \rightarrow \mathbb{X}_2$ gives rise to a curve $\gamma: [0, 1] \rightarrow \mathbb{X}_1$ via $\gamma = f \circ \tilde{\gamma}$, which for every $\varphi \in H^s(\mathbb{X}_1)$ satisfies

$$\left| \int_{\mathbb{X}_1} \varphi \rho d\sigma_{\mathbb{X}_1} - \int_0^1 \varphi \circ \gamma d\lambda \right| = \left| \int_{\mathbb{X}_2} (\varphi \rho) \circ f |\det(J_f)| d\sigma_{\mathbb{X}_2} - \int_0^1 \varphi \circ f \circ \tilde{\gamma} d\lambda \right|,$$

where J_f denotes the Jacobian of f . Now, note that $\varphi \circ f, \rho \circ f |\det(J_f)| \in H^s(\mathbb{X}_2)$, see (7.16) and [260, Thm. 4.3.2], which is lifted to manifolds using (7.25). Hence, we can define a measure $\tilde{\mu}$ on \mathbb{X}_2 through the probability density $\rho \circ f |\det(J_f)|$. Choosing $\tilde{\gamma}_L$ as a realization for some minimizer of $\inf_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}} \mathcal{D}(\tilde{\mu}, \nu)$, we can apply the approximation result for \mathbb{X}_2 and estimate for $\gamma_L = f \circ \tilde{\gamma}_L$ that

$$\left| \int_{\mathbb{X}_1} \varphi \rho d\sigma_{\mathbb{X}_1} - \int_0^1 \varphi \circ \gamma_L d\lambda \right| \lesssim L^{-\frac{s}{d-1}} \|\varphi \circ f\|_{H^s(\mathbb{X}_2)} \lesssim L^{-\frac{s}{d-1}} \|\varphi\|_{H^s(\mathbb{X}_1)},$$

where the second estimate follows from [260, Thm. 4.3.2]. Now, $L(\gamma_L) \leq L(f)L$ implies

$$\inf_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s}{d-1}}.$$

□

Remark 7.23. Consider a probability measure μ on \mathbb{X} such that the dimension d_μ of its support is smaller than the dimension d of \mathbb{X} . Then, μ does not have any density with respect to $\sigma_{\mathbb{X}}$. If $\text{supp}(\mu)$ is itself a d_μ -dimensional connected, compact Riemannian manifold \mathbb{Y} without boundary, we switch from \mathbb{X} to \mathbb{Y} . Sobolev trace theorems and reproducing kernel Hilbert space theory imply that the assumption $H_K(\mathbb{X}) = H^s(\mathbb{X})$ leads to $H_{K'}(\mathbb{Y}) = H^{s'}(\mathbb{Y})$, where $K' := K|_{\mathbb{Y} \times \mathbb{Y}}$ is the restricted kernel and $s' = s - (d - d_\mu)/2$, cf. [121]. If, for instance, \mathbb{Y} is diffeomorphic to \mathbb{T}^{d_μ} (or \mathbb{S}^{d_μ} with $d_\mu = 2$), and μ has a positive density $\rho \in H^{s'}(\mathbb{Y})$ with respect to $\sigma_{\mathbb{Y}}$, then Theorem 7.19 (or 7.20) and Proposition 7.22 eventually yield

$$\min_{\nu \in \mathcal{P}_L^{\lambda\text{-curv}}} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s'}{d_\mu-1}}.$$

If $\text{supp}(\mu)$ is a proper subset of \mathbb{Y} , we are able to analyze approximations with $\mathcal{P}_L^{\text{a-curv}}(\mathbb{Y})$. First, we observe that the analogue of Proposition 7.22 also holds for $\mathcal{P}_L^{\text{a-curv}}(\mathbb{X}_1), \mathcal{P}_L^{\text{a-curv}}(\mathbb{X}_2)$ when the positivity assumption on ρ is replaced with the

Lipschitz requirement as in Theorems 7.15 and 7.16. If, for instance, \mathbb{Y} is diffeomorphic to \mathbb{T}^{d_μ} or \mathbb{S}^{d_μ} and μ has a Lipschitz continuous density $\rho \in H^{s'}(\mathbb{Y})$ with respect to $\sigma_{\mathbb{Y}}$, then Theorems 7.15 and 7.16, and Proposition 7.22 eventually yield

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}} \mathcal{D}_K(\mu, \nu) \lesssim L^{-\frac{s'}{d_\mu-1}}.$$

7.6 Discretization

In our numerical experiments, we are interested in determining minimizers of

$$\min_{\nu \in \mathcal{P}_L^{\text{a-curv}}(\mathbb{X})} \mathcal{D}_K^2(\mu, \nu). \quad (7.26)$$

Defining $A_L := \{\gamma \in \text{Lip}(\mathbb{X}) : L(\gamma) \leq L\}$ and using the indicator function

$$\iota_{A_L}(\gamma) := \begin{cases} 0 & \text{if } \gamma \in A_L, \\ +\infty & \text{otherwise,} \end{cases}$$

we can rephrase problem (7.26) as a minimization problem over curves

$$\min_{\gamma \in \mathcal{C}([0,1], \mathbb{X})} \mathcal{J}_L(\gamma),$$

where $\mathcal{J}_L(\gamma) := \mathcal{D}_K^2(\mu, \gamma_*\lambda) + \iota_{A_L}(\gamma)$. As \mathbb{X} is a connected Riemannian manifold, we can approximate curves in A_L by piecewise shortest geodesics with N parts, i.e., by curves from

$$A_{L,N} := \left\{ \gamma \in A_L : \gamma|_{[(i-1)/N, i/N]} \text{ is a shortest geodesic for } i = 1, \dots, N \right\}.$$

Next, we approximate the Lebesgue measure on $[0, 1]$ by $e_N := \frac{1}{N} \sum_{i=1}^N \delta_{i/N}$ and consider the minimization problems

$$\min_{\gamma \in \mathcal{C}([0,1], \mathbb{X})} \mathcal{J}_{L,N}(\gamma), \quad (7.27)$$

where $\mathcal{J}_{L,N}(\gamma) := \mathcal{D}_K^2(\mu, \gamma_*e_N) + \iota_{A_{L,N}}(\gamma)$. Since $\text{ess sup}_{t \in [0,1]} |\dot{\gamma}|(t) = L(\gamma)$, the constraint $L(\gamma) \leq L$ can be reformulated as $\int_0^1 (|\dot{\gamma}|(t) - L)_+^2 dt = 0$.¹ Hence, using $x_i = \gamma(i/N)$, $i = 1, \dots, N$, $x_0 = x_N$ and regarding that $|\dot{\gamma}|(t) = N d_{\mathbb{X}}(x_{i-1}, x_i)$ for $t \in (\frac{i-1}{N}, \frac{i}{N})$, problem (7.27) is rewritten in the computationally more suitable form

$$\min_{(x_1, \dots, x_N) \in \mathbb{X}^N} \mathcal{D}_K^2\left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{x_i}\right) \quad \text{s.t.} \quad \frac{1}{N} \sum_{i=1}^N (N d_{\mathbb{X}}(x_{i-1}, x_i) - L)_+^2 = 0. \quad (7.28)$$

This discretization is motivated by the next proposition. To this end, recall that a sequence $(f_N)_{N \in \mathbb{N}}$ of functions $f_N : \mathbb{X} \rightarrow (-\infty, +\infty]$ is said to Γ -converge to $f : \mathbb{X} \rightarrow (-\infty, +\infty]$ if the following two conditions are fulfilled for each $x \in \mathbb{X}$, see [44]:

¹For $r \in \mathbb{R}$, we use the notation $r_+ = \begin{cases} r, & r \geq 0, \\ 0, & \text{otherwise.} \end{cases}$

- i) $f(x) \leq \liminf_{N \rightarrow \infty} f_N(x_N)$ whenever $x_N \rightarrow x$,
- ii) there is a sequence $(y_N)_{N \in \mathbb{N}}$ with $y_N \rightarrow x$ and $\limsup_{N \rightarrow \infty} f_N(y_N) \leq f(x)$.

The importance of Γ -convergence relies in the fact that every cluster point of minimizers of $(f_N)_{N \in \mathbb{N}}$ is a minimizer of f . Note that for non-compact manifolds \mathbb{X} an additional equi-coercivity condition would be required.

Proposition 7.24. *The sequence $(\mathcal{J}_{L,N})_{N \in \mathbb{N}}$ is Γ -convergent with limit \mathcal{J}_L .*

Proof. 1. First, we verify the lim inf-inequality. Let $(\gamma_N)_{N \in \mathbb{N}}$ with $\lim_{N \rightarrow \infty} \gamma_N = \gamma$, i.e., the sequence satisfies $\sup_{t \in [0,1]} d_{\mathbb{X}}(\gamma(t), \gamma_N(t)) \rightarrow 0$. By excluding the trivial case $\liminf_{N \rightarrow \infty} \mathcal{J}_{L,N}(\gamma_N) = \infty$ and restricting to a subsequence $(\gamma_{N_k})_{k \in \mathbb{N}}$, we may assume $\gamma_{N_k} \in A_{L,N_k} \subset A_L$. Since A_L is closed, we directly infer $\gamma \in A_L$. It holds $e_N \rightarrow \lambda$, which is equivalent to the convergence of Riemann sums for $f \in C[0,1]$, and hence also $\gamma_{N*} e_N \rightarrow \gamma_* dr$. By the weak continuity of \mathcal{D}_K^2 , we obtain

$$\mathcal{J}_L(\gamma) = \mathcal{D}_K^2(\mu, \gamma_* \lambda) = \lim_{N \rightarrow \infty} \mathcal{D}_K^2(\mu, \gamma_{N*} e_N) = \liminf_{N \rightarrow \infty} \mathcal{J}_{L,N}(\gamma_N). \quad (7.29)$$

2. Next, we prove the lim sup-inequality, i.e., we are searching for a sequence $(\gamma_N)_{N \in \mathbb{N}}$ with $\gamma_N \rightarrow \gamma$ and $\limsup_{N \rightarrow \infty} \mathcal{J}_{L,N}(\gamma_N) \leq \mathcal{J}_L(\gamma)$. First, we may exclude the trivial case $\mathcal{J}_L(\gamma) = \infty$. Then, γ_N is defined on every interval $[(i-1)/N, i/N]$, $i = 1, \dots, N$, as a shortest geodesic from $\gamma((i-1)/N)$ to $\gamma(i/N)$. By construction we have $\gamma_N \in A_{L,N}$. From $\gamma, \gamma_N \in A_L$ we conclude

$$\begin{aligned} \sup_{t \in [0,1]} d_{\mathbb{X}}(\gamma(t), \gamma_N(t)) &= \max_{i=1, \dots, N} \sup_{t \in [(i-1)/N, i/N]} d_{\mathbb{X}}(\gamma(t), \gamma_N(t)) \\ &\leq \max_{i=1, \dots, N} \sup_{t \in [(i-1)/N, i/N]} d_{\mathbb{X}}(\gamma(t), \gamma(i/N)) + d_{\mathbb{X}}(\gamma_N(i/N), \gamma_N(t)) \leq \frac{2L}{N}, \end{aligned}$$

implying $\gamma_N \rightarrow \gamma$. Similarly as in (7.29), we infer $\limsup_{N \rightarrow \infty} \mathcal{J}_{L,N}(\gamma_N) \leq \mathcal{J}_L(\gamma)$. \square

In the numerical part, we use the penalized form of (7.28) and minimize

$$\min_{(x_1, \dots, x_N) \in \mathbb{X}^N} \mathcal{D}_K^2\left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{x_i}\right) + \frac{\lambda}{N} \sum_{i=1}^N (N d_{\mathbb{X}}(x_{i-1}, x_i) - L)_+^2, \quad \lambda > 0. \quad (7.30)$$

7.7 Numerical algorithm

For a detailed overview on Riemannian optimization we refer to [233] and the books [2, 266]. In order to minimize (7.30), we have a closer look at the discrepancy term. By (7.6) and (7.7), the discrepancy can be represented as follows

$$\begin{aligned} \mathcal{D}_K^2\left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{x_i}\right) &= \frac{1}{N^2} \sum_{i,j=1}^N K(x_i, x_j) - 2 \sum_{i=1}^N \int_{\mathbb{X}} K(x_i, x) d\mu(x) + \iint_{\mathbb{X} \times \mathbb{X}} K d\mu d\mu \\ &= \sum_{k=0}^{\infty} \alpha_k \left| \hat{\mu}_k - \frac{1}{N} \sum_{i=1}^N \varphi_k(x_i) \right|^2. \end{aligned}$$

Both formulas have pros and cons: The first formula allows for an exact evaluation only if the expressions $\Phi(x) := \int_{\mathbb{X}} K(x, y) d\mu(y)$ and $\int_{\mathbb{X}} \Phi d\mu$ can be written in closed forms. In this case the complexity scales quadratically in the number of points N . The second formula allows for exact evaluation only if the kernel has a finite expansion (7.3). In that case the complexity scales linearly in N .

Our approach is to use kernels fulfilling $H_K(\mathbb{X}) = H^s(\mathbb{X})$, $s > d/2$, and approximating them by their truncated representation with respect to the eigenfunctions of the Laplace–Beltrami operator

$$K_r(x, y) := \sum_{k \in \mathcal{I}_r} \alpha_k \varphi_k(x) \overline{\varphi_k(y)}, \quad \mathcal{I}_r := \{k : \varphi_k \in \Pi_r(\mathbb{X})\}.$$

Then, we finally aim to minimize

$$\min_{x \in \mathbb{X}^N} F(x) := \sum_{k \in \mathcal{I}_r} \alpha_k \left(\hat{\mu}_k - \frac{1}{N} \sum_{i=1}^N \varphi_k(x_i) \right)^2 + \frac{\lambda}{N} \sum_{i=1}^N (Nd_{\mathbb{X}}(x_{i-1}, x_i) - L)_+^2, \quad (7.31)$$

where $\lambda > 0$. Our algorithm of choice is the nonlinear conjugate gradient (CG) method with Armijo line search as outlined in Algorithm 7.1 with notation and implementation details described in the comments after Remark 7.25, see [85] for Euclidean spaces. Note that the notation is independent of the special choice of \mathbb{X} in our comments. The proposed method is of “exact conjugacy” and uses the second order derivative information provided by the Hessian. For the Armijo line search itself, the sophisticated initialization in Algorithm 7.2 is used, which also incorporates second order information via the Hessian. The main advantage of the CG method is its simplicity together with fast convergence at low computational cost. Indeed, Algorithm 7.1, together with Algorithm 7.2 replaced by an exact line search, converges under suitable assumptions superlinearly, more precisely dN -step quadratically towards a local minimum, cf. [250, Thm. 5.3] and [131, Sec. 3.3.2, Thm. 3.27].

Remark 7.25. *The objective in (7.31) violates the smoothness requirements whenever $x_{k-1} = x_k$ or $d_{\mathbb{X}}(x_{k-1}, x_k) = L/N$. However, we observe numerically that local minimizers of (7.31) do not belong to this set of measure zero. This means in turn, if a local minimizer has a positive definite Hessian, then there is a local neighborhood where the CG method (with exact line search) permits a superlinear convergence rate. We do indeed observe this behavior in our numerical experiments.*

Let us briefly comment on Algorithm 7.1 for $\mathbb{X} \in \{\mathbb{T}^2, \mathbb{T}^3, \mathbb{S}^2, \text{SO}(3), \mathcal{G}_{2,4}\}$ which are considered in our numerical examples. For additional implementation details we refer to [131]. By $\gamma_{x,d}$ we denote the geodesic with $\gamma_{x,d}(0) = x$ and $\dot{\gamma}_{x,d}(0) = d$. Besides evaluating the geodesics $\gamma_{x^{(k)}, d^{(k)}}(\tau^{(k)})$ in the first iteration step, we have to compute the parallel transport of $d^{(k)}$ along the geodesics in the second step. Furthermore, we need to compute the Riemannian gradient $\nabla_{\mathbb{X}^N} F$ and products of the Hessian $H_{\mathbb{X}^N} F$ with vectors d , which are approximated by the finite difference

$$H_{\mathbb{X}^N} F(x)d \approx \frac{\|d\|}{h} \left(\nabla_{\mathbb{X}^N} F(\gamma_{x,hd/\|d\|}) - \nabla_{\mathbb{X}^N} F(x) \right), \quad h := 10^{-8}.$$

Algorithm 7.1 (CG Method with Restarts)**Parameters:** maximal iterations $k_{\max} \in \mathbb{N}$ **Input:** twice differentiable function $F: \mathbb{X}^N \rightarrow [0, \infty)$, initial point $x^{(0)} \in \mathbb{X}^N$ **Initialization:** $g^{(0)} := \nabla_{\mathbb{X}^N} F(x^{(0)})$, $d^{(0)} := -g^{(0)}$, $r := 0$ **for** $k := 0, \dots, k_{\max}$ **do** $x^{(k+1)} := \gamma_{x^{(k)}, d^{(k)}}(\tau^{(k)})$ where $\tau^{(k)}$ is determined by Algorithm 7.2 $\tilde{d}^{(k)} := \dot{\gamma}_{x^{(k)}, d^{(k)}}(\tau^{(k)})$ $g^{(k+1)} := \nabla_{\mathbb{X}^N} F(x^{(k+1)})$

$$\beta^{(k)} := \begin{cases} \frac{\langle \tilde{d}^{(k)}, H_{\mathbb{X}^N} F(x^{(k+1)}) g^{(k+1)} \rangle}{\langle \tilde{d}^{(k)}, H_{\mathbb{X}^N} F(x^{(k+1)}) \tilde{d}^{(k)} \rangle}, & \langle \tilde{d}^{(k)}, H_{\mathbb{X}^N} F(x^{(k+1)}) \tilde{d}^{(k)} \rangle \neq 0, \\ 0, & \text{else} \end{cases}$$

 $d^{(k+1)} := -g^{(k+1)} + \beta^{(k)} \tilde{d}^{(k)}$ **if** $\langle d^{(k+1)}, g^{(k+1)} \rangle > 0$ or $(k+1) \equiv r \pmod{N \dim(\mathbb{X})}$ **then** $d^{(k+1)} = -g^{(k+1)}$ $r := k + 1$ **Output:** iteration sequence $x^{(0)}, x^{(1)}, \dots \in \mathbb{X}^N$ **Algorithm 7.2 (Armijo Line Search)****Parameters:** $0 < a < \frac{1}{2}$, $0 < b < 1$, maximal iterations $k_{\max} \in \mathbb{N}$ **Input:** smooth function $F: \mathbb{X}^N \rightarrow [0, \infty)$, start point $x \in \mathbb{X}^N$, descent direction $d \in T_x \mathbb{X}^N$ **Initialization:** $k := 0$,

$$\tau^{(0)} := \begin{cases} \left| \frac{\langle d, \nabla_{\mathbb{X}^N} F(x) \rangle}{\langle d, H_{\mathbb{X}^N} F(x) d \rangle} \right|, & \langle d, H_{\mathbb{X}^N} F(x) d \rangle \neq 0, \\ 1, & \text{else} \end{cases}$$

while $f \circ \gamma_{x,d}(\tau^{(k)}) - F(x) \geq a\tau^{(k)} \langle \nabla_{\mathbb{X}^N} F(x), d \rangle$ and $k < k_{\max}$ **do** $\tau^{(k+1)} := b\tau^{(k)}$ $k := k + 1$ **Output:** $\tau^{(k)}$ (success if $k \leq k_{\max}$)

The computation of the gradient of the *penalty term* in (30) is done by applying the chain rule and noting that for $x \mapsto d_{\mathbb{X}}(x, y)$, we have $\nabla_{\mathbb{X}} d_{\mathbb{X}}(x, y) = \log_x y / d_{\mathbb{X}}(x, y)$, $x \neq y$ with the logarithmic map \log on \mathbb{X} , while the distance is not differentiable for $x = y$. Concerning the later point, see Remark 5. The evaluation of the gradient of the penalty term at a point in \mathbb{X}^N requires only $\mathcal{O}(N)$ arithmetic operations. The computation of the Riemannian gradient of the *data term* in (30) is done analytically via the gradient of the eigenfunctions φ_k of the Laplace–Beltrami operator. Then, the evaluation of the gradient of the whole data term at given points can be done efficiently by *fast Fourier transform* (FFT) techniques at non-equispaced nodes using the NFFT software package of Potts et al. [174]. The overall complexity of

\mathbb{X}	Reference	Complexity
\mathbb{T}^d	[134], [131, Sec. 5.2.1]	$\mathcal{O}(r^d \log(r) + N)$
\mathbb{S}^2	[133, 134], [131, Sec. 5.2.2]	$\mathcal{O}(r^2 \log^2(r) + N)$
$\text{SO}(3)$	[130, 132], [131, Sec. 5.2.3]	$\mathcal{O}(r^3 \log^2(r) + N)$
$\mathcal{G}_{2,4}$	[91]	$\mathcal{O}(r^4 \log^2(r) + N)$

Table 7.1: References for implementation details of Alg. 7.1 (left) and arithmetic complexity for the evaluations per iteration for the different manifolds (right).

the algorithm and references for the computation details for the above manifolds are given in Table 7.1.

7.8 Numerical results

In this section, we underline our theoretical results by numerical examples. We start by studying the parameter choice in our numerical model. Then, we provide examples for the approximation of absolutely continuous measures with densities in $H^s(\mathbb{X})$, $s > d/2$, by push-forward measures of the Lebesgue measure on $[0, 1]$ by Lipschitz curves for the manifolds $\mathbb{X} \in \{\mathbb{T}^2, \mathbb{T}^3, \mathbb{S}^2, \text{SO}(3), \mathcal{G}_{2,4}\}$. Supplementary material can be found on our webpage.

7.8.1 Parameter choice

We like to emphasize that the optimization problem (7.31) is highly nonlinear and the objective function has a large number of local minimizers, which appear to increase exponentially in N . In order to find for fixed L reasonable (local) solutions of (7.26), we carefully adjust the parameters in problem (7.31), namely the number of points N , the polynomial degree r in the kernel truncation, and the penalty parameter λ . In the following, we suppose that $\dim(\text{supp}(\mu)) = d \geq 2$.

- i) **Number of points N :** Clearly, N should not be too small compared to L . However, from a computational perspective it should also be not too large since the optimization procedure is hampered by the vast number of local minimizers. From the asymptotic of the path lengths of TSP in Lemma 7.6, we conclude that $N \gtrsim \ell(\gamma)^{d/(d-1)}$ is a reasonable choice, where $\ell(\gamma) \leq L$ is the length of the resulting curve γ going through the points.
- ii) **Polynomial degree r :** Based on the proofs of the theorems in Subsection 7.5.4 it is reasonable to choose

$$r \sim L^{\frac{1}{d-1}} \sim N^{\frac{1}{d}}.$$

- iii) **Penalty parameter λ :** If λ is too small, we cannot enforce that the points approximate a regular curve, i.e., $L/N \gtrsim d_{\mathbb{X}}(x_{k-1}, x_k)$. Otherwise, if λ is

too large the optimization procedure is hampered by the rigid constraints. Hence, to find a reasonable choice for λ in dependence on L , we assume that the minimizers of (7.31) treat both terms proportionally, i.e., for $N \rightarrow \infty$ both terms are of the same order. Therefore, our heuristic is to choose the parameter λ such that

$$\min_{x_1, \dots, x_N} \mathcal{D}_K^2\left(\mu, \frac{1}{N} \sum_{k=1}^N \delta_{x_k}\right) \sim N^{-\frac{2s}{d}} \sim \frac{\lambda}{N} \sum_{k=1}^N (Nd_{\mathbb{X}}(x_{k-1}, x_k) - L)_+^2.$$

On the other hand, assuming that for the length $\ell(\gamma) = \sum_{k=1}^N d_{\mathbb{X}}(x_{k-1}, x_k)$ of a minimizer γ we have $\ell(\gamma) \sim L \sim N^{(d-1)/d}$, so that $Nd_{\mathbb{X}}(x_{k-1}, x_k) \sim L$, the value of the penalty term behaves like

$$\frac{\lambda}{N} \sum_{k=1}^N (Nd_{\mathbb{X}}(x_{k-1}, x_k) - L)_+^2 \sim \lambda L^2 \sim \lambda N^{\frac{2d-2}{d}}.$$

Hence, a reasonable choice is

$$\lambda \sim L^{\frac{-2s-2(d-1)}{d-1}} \sim N^{\frac{-2s-2(d-1)}{d}}. \quad (7.32)$$

Remark 7.26. In view of Remark 7.23 the relations in i)-iii) become

$$N \sim L^{\frac{d_\mu}{d_\mu-1}}, \quad r \sim N^{\frac{1}{d_\mu}} \sim L^{\frac{1}{d_\mu-1}}, \quad \lambda \sim L^{\frac{-2s-3d_\mu+d+2}{d_\mu-1}} \sim N^{\frac{-2s-3d_\mu+d+2}{d_\mu}}.$$

In the rest of this subsection, we aim to provide some numerical evidence for the parameter choice above. We restrict our attention to the torus $\mathbb{X} = \mathbb{T}^2$ and the kernel K given in (7.34) with $d = 2$ and $s = 3/2$. Choose μ as the Lebesgue measure on \mathbb{T}^2 . From (7.32), we should keep in mind $\lambda \sim N^{-5/2} \sim L^{-5}$.

Influence of N and λ . We fix $L = 4$ and a large polynomial degree $r = 128$ for truncating the kernel. For any $\lambda_i = 0.1 \cdot 2^{-5i/2}$, $i = 1, \dots, 4$, we compute local minimizers with $N_j = 10 \cdot 2^j$, $j = 1, \dots, 4$. More precisely, keeping λ_i fixed we start with $N_1 = 20$ and refine successively the curves by inserting the midpoints of the line segments connecting consecutive points and applying a local minimization with this initialization. The results are depicted in Fig. 7.1. For fixed λ (fixed row) we can clearly notice that the local minimizers converge towards a smooth curve for increasing N . Moreover, the diagonal images correspond to the choice $\lambda = 0.1(N/10)^{-5/2}$, where we can already observe good approximation of the curves emerging to the right of it. This should provide some evidence that the choice of the penalty parameter λ and the number of points N discussed above is reasonable. Indeed, for $\lambda \rightarrow \infty$ we observe $L(\gamma) \rightarrow \ell(\gamma) \rightarrow L = 4$.

Influence of the polynomial degree r . In Fig. 7.2 we illustrate the local minimizers of (7.31) for fixed Lipschitz parameters $L_i = 2^i$ and corresponding regularization weights $\lambda_i = 0.2 \cdot L_i^{-5}$, $i = 1, \dots, 4$, (rows) in dependence on the polynomial degrees $r_j = 8 \cdot 2^j$, $j = 1, \dots, 5$ (columns). According to the previous experiments, it seems reasonable to choose $N = 20L^2$. Note, that the (numerical)




Figure 7.1: Influence of N and λ on local minimizers of (7.31) for the Lebesgue measure on \mathbb{T}^2 , $L = 4$ and $r = 128$. Results for increasing N (column-wise) and decreasing $\lambda = 0.1 \cdot 2^{-5i/2}$, $i = 1, \dots, 4$, (row-wise). Here, the curve length increases for decreasing λ or increasing N , until stagnation for sufficient small λ or large N . For all minimizer the distance between consecutive points is around $\ell(\gamma)/N$.




Figure 7.2: Influence of r on the local minimizer of (7.31) for the Lebesgue measure on \mathbb{T}^2 . Column-wise we increase $r = 16, 32, 64, 128, 256$ and row-wise we increase $L = 2, 4, 8, 16$, where $\lambda = 0.2L^{-5}$ and $N = 20L^2$. Note that the degree r steers the resolution of the curves. It appears that the spacing of the curves is bounded by r^{-1} .

choice of λ leads to curves with length $\ell(\gamma) \approx 2L$. In Fig. 7.2 we observe that for $r = cL$ the corresponding local minimizers have common features. For instance, if $c = 4$ (i.e., $r \approx \ell(\gamma)$) the minimizers have mostly vertical and horizontal line segments. Furthermore, for fixed r it appears that the length of the curves increases linearly with L until L exceeds $2r$, from where it remains unchanged. This observation can be explained by the fact that there are curves of bounded length cr which provide exact quadratures for degree r .

7.8.2 Quasi-optimal curves on special manifolds

In this subsection, we give numerical examples for $\mathbb{X} \in \{\mathbb{T}^2, \mathbb{T}^3, \mathbb{S}^2, \text{SO}(3), \mathcal{G}_{2,4}\}$. Since the objective function in (7.31) is highly non-convex, the main problem is to find nearly optimal curves $\gamma_L \in \mathcal{P}_L^{\lambda\text{-curv}}(\mathbb{X})$ for increasing L . Our heuristic is as

follows:

- i) We start with a curve $\gamma_{L_0}: [0, 1] \rightarrow \mathbb{X}$ of small length $\ell(\gamma) \approx L_0$ and solve the problem (7.31) for increasing $L_i = cL_{i-1}$, $c > 1$, where we choose the parameters N_i , λ_i and r_i in dependence on L_i as described in the previous subsection. In each step a local minimizer is computed using the CG method with 100 iterations. Then, the obtained minimizer γ_i serves as the initial guess in the next step, which is obtained by inserting the midpoints.
- ii) In case that the resulting curves γ_i have non-constant speed, each is refined by increasing λ_i and N_i . Then, the resulting problem is solved with the CG method and γ_i as initialization. Details on the parameter choice are given in the according examples.

The following examples show that this recipe indeed enables us to compute “quasi-optimal” curves, meaning that the obtained minimizers have optimal decay in the discrepancy.

2d-Torus \mathbb{T}^2 . In this example we illustrate how well a gray-valued image (considered as probability density) may be approximated by an almost constant speed curve. The original image of size 170x170 is depicted in the bottom-right corner of Fig. 7.3. Its Fourier coefficients $\hat{\mu}_{k_1, k_2}$ are computed by a discrete Fourier transform (DFT) using the FFT algorithm and normalized appropriately. The kernel K is given by (7.34) with $d = 2$ and $s = 3/2$.

We start with $N_0 = 96$ points on a circle given by the formula

$$x_{0,k} = \left(\frac{1}{5} \cos(2\pi k/N_0), \frac{1}{5} \sin(2\pi k/N_0) \right), \quad k = 0, \dots, N_0.$$

Then, we apply our procedure for $i = 0, \dots, 11$ with parameters

$$L_i = 0.97 \cdot 2^{\frac{i+5}{2}}, \quad \lambda_i = 100 \cdot L_i^{-5}, \quad N_i = 96 \cdot 2^i \sim L_i^2 \quad r_i = \lfloor 2^{\frac{i+11}{2}} \rfloor \sim L_i,$$

chosen such that the length of the local minimizer γ_i satisfies $\ell(\gamma_i) \approx 2^{(i+5)/2}$ and the maximal speed is close to L_i .

To get nearly constant speed curves γ_i , see ii), we increase λ_i by a factor of 100, N_i by a factor of 2 and set $L_i := 2^{(i+5)/2}$. Then, we apply the CG method with maximal 100 iterations and i restarts. The results are depicted in Fig. 7.3. Note that the complexity for the evaluation of the function in (7.31) scales roughly as $N \sim L^2$. In Fig. 7.4 we observe that the decay-rate of the squared discrepancy $\mathcal{D}_K^2(\mu, \nu)$ in dependence on the Lipschitz constant L matches indeed the theoretical findings of Theorem 7.19.

3D-Torus \mathbb{T}^3 . The aim of this example is two-fold. First, it shows that the algorithm works pretty well in three dimensions. Second, we are able to approximate any compact surface in the three-dimensional space by a curve. We construct a measure μ supported around a two-dimensional surface by taking samples from




Figure 7.3: Local minimizers of (7.31) for the image at bottom right.




Figure 7.4: Squared discrepancy between the measure μ given by the image in Fig. 7.3 and the computed local minimizers (black dots) on \mathbb{T}^2 in log-scale. The blue line corresponds to the optimal decay-rate in Theorem 7.19.

Spock's head³ and placing small Gaussian peaks at the sampling points, i.e., the density is given for $x \in [-\frac{1}{2}, \frac{1}{2}]$ by

$$\rho(x) := c^{-1} \sum_{p \in S} e^{-30000 \|p-x\|_2^2}, \quad c := \int_{[-\frac{1}{2}, \frac{1}{2}]^3} \sum_{p \in S} e^{-30000 \|p-x\|_2^2} dx,$$

where $S \subset [-\frac{1}{2}, \frac{1}{2}]^3$ is the discrete sampling set. From a numerical point of view it holds $\dim(\text{supp}(\mu)) = 2$. The Fourier coefficients are again computed by a DFT and the kernel K is given by (7.34) with $d = 3$ and $s = 2$ so that $H_K = H^2(\mathbb{T}^3)$.

We start with $N_0 = 100$ points on a smooth curve given by the formula

$$x_{0,k} = \left(\frac{3}{10} \cos(2\pi k/N_0), \frac{3}{10} \sin(2\pi k/N_0), \frac{3}{10} \sin(4\pi k/N_0) \right), \quad k = 0, \dots, N_0.$$

Then, we apply our procedure for $i = 0, \dots, 8$ with parameters, cf. Remark 7.26,

$$L_i = 2^{\frac{i+5}{2}}, \quad \lambda_i = 10 \cdot L_i^{-5}, \quad N_i = 100 \cdot 2^i \sim L_i^2, \quad r_i = \lfloor 2^{\frac{i+5}{2}} \rfloor \sim L_i.$$

To get nearly constant speed curves γ_i , we increase λ_i by a factor of 100, N_i by a factor of 2 and set $L_i := 2^{(i+6)/2}$. Then, we apply the CG method with maximal 100 iterations and one restart to the previously found curve γ_i . The results are illustrated in Fig. 7.5. Note that the complexity of the function evaluation in (7.31) scales roughly as $N^{3/2} \sim L^3$. In Fig. 7.6 we depict the squared discrepancy $\mathcal{D}_K^2(\mu, \nu)$ of the computed curves. For small Lipschitz constants, say $L(\gamma) \leq 50$, we observe a decrease of approximately $L(\gamma)^{-3}$, which matches the optimal decay-rate for measures supported on surfaces as discussed in Remark 7.23.

2-Sphere \mathbb{S}^2 . Next, we approximate a gray-valued image on the sphere \mathbb{S}^2 by an almost constant speed curve. The image represents the earth's elevation data provided by MATLAB, given by samples $\rho_{i,j}$, $i = 1, \dots, 180$, $j = 1, \dots, 360$, on the grid

$$x_{i,j} := \left(\sin(i \frac{\pi}{180}) \sin(j \frac{\pi}{180}), \sin(i \frac{\pi}{180}) \cos(j \frac{\pi}{180}), \cos(i \frac{\pi}{180}) \right).$$

The Fourier coefficients are computed by discretizing the Fourier integrals, i.e.,

$$\hat{\mu}_k^m := \begin{cases} \frac{1}{180 \cdot 360} \sum_{i=1}^{180} \sum_{j=1}^{360} \rho_{i,j} \overline{Y_k^m(x_{i,j})} \sin(i \frac{\pi}{180}), & 1 \leq k \leq 2m+1, m \leq 180, \\ 0, & \text{else,} \end{cases}$$

followed by a suitable normalization such that $\hat{\mu}_0^0 = 1$. The corresponding sums are efficiently computed by an adjoint non-equispaced fast spherical Fourier transform (NFSFT), see [228]. The kernel K is given by (7.36). Similar to the previous examples, we apply our procedure for $i = 0, \dots, 12$ with parameters

$$L_i = 9.7 \cdot 2^{\frac{i}{2}}, \quad \lambda_i = 100 \cdot L_i^{-5}, \quad N_i = 100 \cdot 2^i \sim L_i^2, \quad r_i = \lfloor L_i \rfloor \sim L_i.$$

To get nearly constant speed curves, we increase λ_i by a factor of 100, N_i by a factor of 2 and set $L_i := L_0 2^{i/2}$. Then, we apply the CG method with maximal 100

³<http://www.cs.technion.ac.il/~vitus/mingle/>




Figure 7.5: Local minimizers of (7.31) for a measure μ concentrated on a surface (head of Spock) in \mathbb{T}^3 .




Figure 7.6: Squared discrepancy between the measure μ given by the surface in Fig. 7.5 and the computed local minimizers (black dots) on \mathbb{T}^3 in log-scale. The blue line corresponds to the optimal decay-rate in Theorem 7.19.




Figure 7.7: Local minimizers of (7.31) for μ given by the earth's elevation data on the sphere \mathbb{S}^2 .




Figure 7.8: Squared discrepancy between the measure μ and the computed local minimizers (black dots) in log-scale. The blue line corresponds to the optimal decay-rate in Theorem 7.20.

iterations and one restart to the previously constructed curves γ_i . The results for $i = 6, 8, 10, 12$ are depicted in Fig. 7.7. Note that the complexity of the function evaluation in (7.31) scales roughly as $N \sim L^2$. In Fig. 7.8 we observe that the decay-rate of the squared discrepancy $\mathcal{D}_K^2(\mu, \nu)$ in dependence on the Lipschitz constant matches indeed the theoretical findings in Theorem 7.20.

3D-Rotations $\text{SO}(3)$. There are several possibilities to parameterize the rotation group $\text{SO}(3)$. We apply those by Euler angles and an axis-angle representation for visualization. Euler angles $(\varphi_1, \theta, \varphi_2) \in [0, 2\pi) \times [0, \pi] \times [0, 2\pi)$ correspond to rotations $\text{Rot}(\varphi_1, \theta, \varphi_2)$ in $\text{SO}(3)$ that are the successive rotations around the axes e_3, e_2, e_3 by the respective angles. Then, the Haar measure of $\text{SO}(3)$ is determined by

$$d\mu_{\text{SO}(3)}(\varphi_1, \theta, \varphi_2) = \frac{1}{8\pi^2} \sin(\theta) d\varphi_1 d\theta d\varphi_2.$$

We are interested in the full three-dimensional doughnut

$$D = \{\text{Rot}(\varphi_1, \theta, \varphi_2) : 0 \leq \theta \leq \frac{\pi}{2}, 0 \leq \varphi_1, \varphi_2 \leq 2\pi\} \subset \text{SO}(3).$$

Next, we want to approximate the Haar measure $\mu = \mu_D$ restricted to D , i.e., with normalization we consider the measure defined for $f \in C(\text{SO}(3))$ by

$$\int_{\text{SO}(3)} f d\mu_D = \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{\frac{\pi}{2}} \int_0^{2\pi} f(\varphi_1, \theta, \varphi_2) \sin(\theta) d\varphi_1 d\theta d\varphi_2.$$

The Fourier coefficients of μ_D can be explicitly computed by

$$\hat{\mu}_{l,l'}^k = \begin{cases} P_{k-1}(0) - P_{k+1}(0), & l, l' = 0, k \geq 0, \\ 0, & l, l' \neq 0, \end{cases}$$

where P_k are the Legendre polynomials. The kernel K is given by (7.37) with $d = 3$ and $s = 2$. For $i = 0, \dots, 8$ the parameters are chosen as

$$L_i = 0.93 \cdot 2^{\frac{2i+12}{3}}, \quad \lambda_i = 10 \cdot L_i^{-4}, \quad N_i = 64 \cdot 2^i \sim L_i^2, \quad r_i = \lfloor 2^{\frac{i+9}{3}} \rfloor \sim L_i^{\frac{1}{2}}.$$

Here, we use a CG method with 100 iterations and one restart. Step ii) appears to be not necessary. Note that the complexity for the function evaluations in (7.31) scales roughly as $N \sim L^{3/2}$.

The constructed curves are illustrated in Fig. 7.9, where we utilized the following visualization: Every rotation $R(\alpha, r) \in \text{SO}(3)$ is determined by a rotation axis $r = (r_1, r_2, r_3) \in \mathbb{S}^2$ and a rotation angle $\alpha \in [0, \pi]$, i.e.,

$$R(\alpha, r)x = r(r^T x) + \cos(\alpha)((r \times x) \times r) + \sin(\alpha)(r \times x).$$

Setting $q := (\cos(\frac{\alpha}{2}), \sin(\frac{\alpha}{2})r) \in \mathbb{S}^3$ with $r \in \mathbb{S}^2$ and $\alpha \in [0, 2\pi]$, see (7.22), we observe that the same rotation is generated by $-q = (\cos(\frac{2\pi-\alpha}{2}), \sin(\frac{2\pi-\alpha}{2})(-r)) \in \mathbb{S}^3$, in other words $\text{SO}(3) \cong \mathbb{S}^3/\{\pm 1\}$. Then, by applying the stereographic projection $\pi(q) = (q_2, q_3, q_4)/(1+q_1)$, we map the upper hemisphere onto the three dimensional unit ball. Note that the equatorial plane of \mathbb{S}^3 is mapped onto the sphere \mathbb{S}^2 , hence on the surface of the ball antipodal points have to be identified. In other words, the rotation $R(\alpha, r)$ is plotted as the point

$$\pi(q) = \frac{\sin(\frac{\alpha}{2})}{1 + \cos(\frac{\alpha}{2})} r = \tan(\frac{\alpha}{4}) r \in \mathbb{R}^3.$$

In Fig. 7.10 we observe that the decay-rate of $\mathcal{D}_K^2(\mu, \nu)$ in dependence on the Lipschitz constant L matches the theoretical findings in Corollary 7.17.

The 4-dimensional Grassmannian $\mathcal{G}_{2,4}$. Here, we aim to approximate the Haar measure of the Grassmannian $\mathcal{G}_{2,4}$ by a curve of almost constant speed. As this curve samples the space $\mathcal{G}_{2,4}$ quite evenly, it could be used for the grand tour, a technique to analyze high-dimensional data by their projections onto two-dimensional subspaces, cf. [12].

The kernel K of the Haar measure is given by (7.38) and the Fourier coefficients are given by $\hat{\mu}_{m,m'}^{k,k'} = \delta_{m,0} \delta_{m',0} \delta_{k,0} \delta_{k',0}$. For $i = 0, \dots, 8$ the parameters are chosen as

$$L_i = 0.91 \cdot 2^{\frac{3i+16}{4}}, \quad \lambda_i = 100 \cdot L_i^{-\frac{11}{3}}, \quad N_i = 128 \cdot 2^i \sim L_i^2, \quad r_i = \lfloor 2^{\frac{3i+16}{12}} \rfloor + 1 \sim L_i^{\frac{1}{3}}.$$

Here, we use a CG method with 100 iterations and one restart. Our experiments suggest that step ii) is not necessary. Note that the complexity for the function evaluation in (7.31) scales roughly as $N \sim L^{3/2}$.




Figure 7.9: Local minimizers of (7.31) for the Haar measure μ_D of three-dimensional doughnut D in the rotation group $SO(3)$ with a color scheme for better visibility of the 3D structure.




Figure 7.10: Squared discrepancy between the measure μ_D and the computed local minimizers (black dots) in log-scale. The blue line corresponds to the optimal decay-rate in Corollary 7.17.

The computed curves are illustrated in Fig. 7.11, where we use the following visualization. By Remark 7.27, there exists an isometric one-to-one mapping $P: \mathbb{S}^2 \times \mathbb{S}^2 / \{\pm 1\} \rightarrow \mathcal{G}_{2,4}$. Using this relation, we plot the point $P(u, v) \in \mathcal{G}_{2,4}$ by two antipodal points $z_1 = u + v$, $z_2 = -u - v \in \mathbb{R}^3$ together with the RGB color-coded vectors $\pm u$.⁴ More precisely, $R = (1 \mp u_1)/2$, $G = (1 \mp u_2)/2$, $B = (1 \mp u_3)/2$. This means a curve $\gamma(t) \in \mathcal{G}_{2,4}$ only intersects itself if the corresponding curve $z(t) \in \mathbb{R}^3$ intersects and has the same colors at the intersection point. In Fig. 7.12 we observe that the decay-rate of the squared discrepancy $\mathcal{D}_K^2(\mu, \nu)$ in dependence on the Lipschitz constant L matches indeed the theoretical findings in Theorem 7.21.

7.9 Conclusions

In this chapter, we provided approximation results for general probability measures on compact Ahlfors d -regular metric spaces \mathbb{X} by

- i) measures supported on continuous curves of finite length, which are actually push-forward measures of probability measures on $[0, 1]$ by Lipschitz curves;
- ii) push-forward measures of absolutely continuous probability measures on $[0, 1]$ by Lipschitz curves;
- iii) push-forward measures of the Lebesgue measure on $[0, 1]$ by Lipschitz curves.

Our estimates rely on discrepancies between measures. In contrast to Wasserstein distances, these estimates do not reflect the curse of dimensionality.

In approximation theory, a natural question is how the approximation rates improve as the “measures become smoother”. Therefore, we considered absolutely continuous probability measures with densities in Sobolev spaces, where we have to restrict ourselves to compact Riemannian manifolds \mathbb{X} . We proved lower estimates for all three approximation spaces i)-iii). Concerning upper estimates, we gave a result for the approximation space i). Unfortunately, we were not able to show similar results for the smaller approximation spaces ii) and iii). Nevertheless, for these cases, we could provide results for the d -dimensional torus, the d -sphere, the three-dimensional rotation group and the Grassmannian $\mathcal{G}_{2,4}$, which are all of interest on their own. Numerical examples on these manifolds underline our theoretical findings.

Our results can be seen as starting point for future research. Clearly, we want to have more general results also for the approximation spaces ii) and iii). We hope that our research leads to further practical applications. It would be also interesting to consider approximation spaces of measures supported on higher dimensional submanifolds as, e.g., surfaces.

Recently, results on the principal component analysis (PCA) on manifolds were obtained. It may be interesting to see if some of our approximation results can

⁴Note that the decomposition of $z \in \mathbb{R}^3$ with $0 < \|z\| < 2$ into u and v is not unique. There is a one-parameter family of points $u_s, v_s \in \mathbb{S}^2$ such $z = u_s + v_s$. The point $z = 0$ has a two-dimensional ambiguity $v = -u$, $u \in \mathbb{S}^2$ and the point $z \in 2\mathbb{S}^2$ has a unique pre-image $v = u = \frac{1}{2}z$.




Figure 7.11: Local minimizers of (7.31) for the Haar measure of the Grassmannian $\mathcal{G}_{2,4}$.




Figure 7.12: The squared discrepancy between the Haar measure μ and the computed local minimizers (black dots) in log-scale. Here, the blue line corresponds to the optimal decay-rate, cf. Theorem 7.21.

be also modified for the setting of principal curves, cf. Remark 7.5. In contrast to [173, Thm. 1] that bounds the discretization error for fixed length, we were able to provide precise error bounds for the discrepancy in dependence on the Lipschitz constant L of γ and the smoothness of the density $d\mu$.

7.A Special manifolds

Here, we introduce the main examples that are addressed in the numerical part. The measure $\sigma_{\mathbb{X}}$ is always the normalized Riemannian measure on the manifold \mathbb{X} . Note that for simplicity of notation all eigenspaces are complex in this section. We are interested in the following special manifolds.

Example 1: $\mathbb{X} = \mathbb{T}^d$. For $\mathbf{k} \in \mathbb{Z}^d$, set $|\mathbf{k}|^2 := k_1^2 + \dots + k_d^2$ and $|\mathbf{k}|_\infty := \max\{|k_1|, \dots, |k_d|\}$. Then $-\Delta$ has eigenvalues $\{4\pi^2|\mathbf{k}|^2\}_{\mathbf{k} \in \mathbb{Z}^d}$ with eigenfunctions $\{e^{2\pi i \langle \mathbf{k}, \cdot \rangle}\}_{\mathbf{k} \in \mathbb{Z}^d}$. The space of d -variate trigonometric polynomials of degree r ,

$$\Pi_r(\mathbb{T}^d) := \text{span}\{e^{2\pi i \langle \mathbf{k}, x \rangle} : |\mathbf{k}|_\infty \leq r\} \quad (7.33)$$

has dimension $(2r+1)^d$ and contains the eigenspaces belonging to eigenvalues smaller than $4\pi^2 r^2$. As kernel for H^s , $s = (d+1)/2$, we use in our numerical examples

$$K(x, y) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (1 + |\mathbf{k}|_2^2)^{-\frac{d+1}{2}} e^{2\pi i \langle \mathbf{k}, x-y \rangle} = \sum_{\mathbf{k} \in \mathbb{Z}^d} (1 + |\mathbf{k}|_2^2)^{-\frac{d+1}{2}} \cos(2\pi \langle \mathbf{k}, x-y \rangle). \quad (7.34)$$

Example 2: $\mathbb{X} = \mathbb{S}^d \subset \mathbb{R}^{d+1}$, $d \geq 1$. We use distance $d_{\mathbb{S}^d}(x, z) = \arccos(\langle x, z \rangle)$. The Laplace–Beltrami operator $-\Delta$ on \mathbb{S}^d has the eigenvalues $\{k(k+d-1)\}_{k \in \mathbb{N}}$ with the spherical harmonics of degree k ,

$$\{Y_l^k : l = 1, \dots, Z(d, k)\}, \quad Z(d, k) := (2k + d - 1) \frac{\Gamma(k+d-1)}{\Gamma(d)\Gamma(k+1)}$$

as corresponding orthonormal eigenfunctions [207]. The span of eigenfunctions with eigenvalues smaller than $r(r+d-1)$ is given by

$$\Pi_r(\mathbb{S}^d) := \text{span}\{Y_l^k : k = 0, \dots, r, l = 1, \dots, Z(d, k)\}. \quad (7.35)$$

It has dimension $\sum_{k=0}^r Z(d, k) = \frac{(d+2r)\Gamma(d+r)}{\Gamma(d+1)\Gamma(r+1)} \sim r^d$ and coincides with the space of polynomials of total degree r in d variables restricted to the sphere. As kernel for $H^s(\mathbb{S}^2)$, $s = 3/2$, we use

$$\begin{aligned} K(x, y) &= \frac{1}{3} + \sum_{k=1}^{\infty} \frac{2}{(2k-1)(2k+1)(2k+3)} \sum_{l=1}^{2k+1} Y_l^k(x) \overline{Y_l^k(y)} \\ &= \frac{1}{3} + \sum_{k=1}^{\infty} \frac{2}{(2k-1)(2k+3)} P_k(\langle x, y \rangle) = 1 - \frac{1}{2} \|x - y\|_2^2 \end{aligned} \quad (7.36)$$

with the Legendre polynomials P_k . Note that the coefficients have the decay as $(k(k+1))^{-3/2}$.

Example 3: $\mathbb{X} = \text{SO}(3)$. This 3-dimensional manifold is equipped with the distance $d_{\text{SO}(3)}(x, y) = \arccos((\text{trace}(x^\top y) - 1)/2)/2$. The eigenvalues of $-\Delta$ are $\{k(k+1)\}_{k=0}^\infty$ and the (normalized) *Wigner-D functions* $\{\mathcal{D}_{l,l'}^k : l, l' = -k, \dots, k\}$ provide an orthonormal basis for $L^2(\text{SO}(3))$, cf. [270]. The span of eigenspaces belonging to eigenvalues smaller than $r(r+1)$ is

$$\Pi_r(\text{SO}(3)) := \text{span}\{\mathcal{D}_{l,l'}^k : k = 0, \dots, r, l, l' = -k, \dots, k\}$$

and has dimension $(r+1)(2r+1)(2r+3)/3$. In the numerical part, we use the following kernel for $H^s(\text{SO}(3))$, $s = 2$,

$$\begin{aligned} K(x, y) &= \frac{\pi}{8} - \frac{1}{3} + \sum_{k=1}^{\infty} \frac{1}{(2k-1)(2k+1)^2(2k+3)} \sum_{l=-k}^k \sum_{l'=-k}^k \mathcal{D}_{l,l'}^k(x) \overline{\mathcal{D}_{l,l'}^k(y)} \quad (7.37) \\ &= \frac{\pi}{8} - \frac{1}{3} + \sum_{k=1}^{\infty} \frac{1}{(2k-1)(2k+1)(2k+3)} U_{2k}\left(\frac{1}{2}\sqrt{\text{tr}(x^\top y) + 1}\right) \\ &= \frac{\pi}{8} - \pi \frac{\sqrt{2}}{16} \|x - y\|_{\text{F}}, \end{aligned}$$

where U_k are the Chebyshev polynomials of the second kind.

Example 4: $\mathbb{X} = \mathcal{G}_{2,4}$. For integers $1 \leq s < r$, the (s, r) -Grassmannian is the collection of all s -dimensional linear subspaces of \mathbb{R}^r and carries the structure of a closed Riemannian manifold. By identifying a subspace with the orthogonal projector onto this subspace, the Grassmannian becomes

$$\mathcal{G}_{s,r} := \{x \in \mathbb{R}^{r \times r} : x^\top = x, x^2 = x, \text{rank}(x) = s\}.$$

In our context, the cases $\mathcal{G}_{1,2}$, $\mathcal{G}_{1,3}$, and $\mathcal{G}_{2,3}$ can essentially be treated by the spheres \mathbb{S}^1 and \mathbb{S}^2 . The simplest Grassmannian that is algebraically different is $\mathcal{G}_{2,4}$. It is a 4-dimensional manifold and the geodesic distance between $x, y \in \mathcal{G}_{2,4}$ is given by

$$d_{\mathcal{G}_{2,4}}(x, y) = \sqrt{2} \sqrt{\theta_1^2(x, y) + \theta_2^2(x, y)},$$

where $\theta_1(x, y)$ and $\theta_2(x, y)$ are the principal angles between the subspaces associated to x and y , respectively. The terms $\cos(\theta_1(x, y))^2$ and $\cos(\theta_2(x, y))^2$ correspond to the two largest singular values of the product xy . The eigenvalues of $-\Delta$ on $\mathcal{G}_{2,4}$ are $4(\lambda_1^2 + \lambda_2^2 + \lambda_1)$, where λ_1 and λ_2 run through all integers with $\lambda_1 \geq \lambda_2 \geq 0$, cf. [16, 17, 18, 101, 167, 237]. The associated eigenfunctions are denoted by φ_l^λ with $l = 1, \dots, Z(\lambda)$, where $Z(\lambda) = (1 + \lambda_1 + \lambda_2)\eta(\lambda_2)$ and $\eta(\lambda_2) = 1$ if $\lambda_2 = 0$ and 2 if $\lambda_2 > 0$ cf. [120, (24.29) and (24.41)] as well as [17, 18].

The space of polynomials of total degree r on $\mathbb{R}^{16} \cong \mathbb{R}^{4 \times 4}$ restricted to $\mathcal{G}_{2,4}$ is

$$\Pi_r(\mathcal{G}_{2,4}) := \text{span}\{\varphi_l^\lambda : \lambda_1 + \lambda_2 \leq r, l = 1, \dots, Z(\lambda)\}.$$

It contains all eigenfunctions φ_l^λ with $4(\lambda_1^2 + \lambda_2^2 + \lambda_1) < 2(r+1)(r+2)$, cf. [47, Thm. 5].

For $H^s(\mathcal{G}_{2,4})$ with $s = 5/2$, we chose the kernel

$$K(x, y) = \sum_{\lambda_1 \geq \lambda_2 \geq 0} (1 + \lambda_1^2 + \lambda_2^2)^{-\frac{5}{2}} \sum_{l=1}^{Z(\lambda)} \varphi_l^\lambda(x) \overline{\varphi_l^\lambda(y)}. \quad (7.38)$$

Remark 7.27. *It is well-known that $\mathbb{S}^2 \times \mathbb{S}^2$ is a double covering of $\mathcal{G}_{2,4}$. More precisely, there is an isometric one-to-one mapping $P: \mathbb{S}^2 \times \mathbb{S}^2 / \{\pm 1\} \rightarrow \mathcal{G}_{2,4}$ given by*

$$P(u, v) = P(-u, -v) := \frac{1}{2} \begin{pmatrix} 1 + u^\top v & -(u \times v)^\top \\ -u \times v & uv^\top + vu^\top + (1 - u^\top v)I_3 \end{pmatrix},$$

cf. [91]. Moreover, the φ_l^λ are essentially tensor products of spherical harmonics, which enables transferring the non-equispaced fast Fourier transform from $\mathbb{S}^2 \times \mathbb{S}^2$ to $\mathcal{G}_{2,4}$, see [91] for details.

Bibliography

- [1] P.-A. Absil, P.-Y. Gouzenbourger, P. Striowski, and B. Wirth. Differentiable piecewise-Bézier surfaces on Riemannian manifolds. *SIAM J. Imaging Sci.*, 9(4):1788–1828, 2016.
- [2] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, 2008.
- [3] E. Akleman, Q. Xing, P. Garigipati, G. Taubin, J. Chen, and S. Hu. Hamiltonian cycle art: Surface covering wire sculptures and duotone surfaces. *Comput. Graph.*, 37(5):316–332, 2013.
- [4] H. W. Alt. *Lineare Funktionalanalysis: Eine anwendungsorientierte Einführung*. Springer, Berlin, 2002.
- [5] L. Ambrosio, N. Fusco, and D. Pallara. *Functions of Bounded Variation and Free Discontinuity Problems*. Oxford University Press, New York, 2000.
- [6] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Birkhäuser, Basel, 2005.
- [7] R. Andreev, O. Scherzer, and W. Zulehner. Simultaneous optical flow and source estimation: Space–time discretization and preconditioning. *Appl. Numer. Math.*, 96:72–81, 2015.
- [8] J. Angulo and S. Velasco-Forero. Morphological processing of univariate Gaussian distribution-valued images based on Poincaré upper-half plane representation. In *Geometric Theory of Information*. Springer, Cham, 2014.
- [9] V. Arnold. Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l’hydrodynamique des fluides parfaits. *Ann. Inst. Fourier*, 16(1):319–361, 1966.
- [10] V. I. Arnold and B. A. Khesin. *Topological Methods in Hydrodynamics*, volume 125 of *Applied Mathematical Sciences*. Springer, New York, 1998.
- [11] J. Ashburner. A fast diffeomorphic image registration algorithm. *Neuroimage*, 38(1):95–113, 2007.
- [12] D. Asimov. The Grand Tour: A tool for viewing multidimensional data. *SIAM J. Sci. Stat. Comput.*, 6(1):28–143, 1985.

- [13] M. Bačák. *Convex Analysis and Optimization in Hadamard Spaces*, volume 22 of *De Gruyter Series in Nonlinear Analysis and Applications*. De Gruyter, Berlin, 2014.
- [14] M. Bačák, J. Hertrich, S. Neumayer, and G. Steidl. Minimal Lipschitz and ∞ -harmonic extensions of vector-valued functions on finite graphs. *Inf. Inference*, 2020.
- [15] F. Bachmann, R. Hielscher, and H. Schaeben. Grain detection from 2D and 3D EBSD data—specification of the MTEX algorithm. *Ultramicroscopy*, 111(12):1720–1733, 2011.
- [16] C. Bachoc. Linear programming bounds for codes in Grassmannian spaces. *IEEE Trans. Inf. Th.*, 52(5):2111–2125, 2006.
- [17] C. Bachoc, E. Bannai, and R. Coulangenon. Codes and designs in Grassmannian spaces. *Discrete Math.*, 277(1-3):15–28, 2004.
- [18] C. Bachoc, R. Coulangenon, and G. Nebe. Designs in Grassmannian spaces and lattices. *J. Algebr. Comb.*, 16(1):5–19, 2002.
- [19] J. M. Ball. Global invertibility of Sobolev functions and the interpenetration of matter. *Proc. R. Soc. Edinb. A*, 88(3-4):315–328, 1981.
- [20] P. J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophys. J.*, 66(1):259–267, 1994.
- [21] M. Bačák, R. Bergmann, G. Steidl, and A. Weinmann. A second order nonsmooth variational model for restoring manifold-valued images. *SIAM J. Sci. Comput.*, 38(1):567–597, 2016.
- [22] M. Bačák, M. Montag, and G. Steidl. Convergence of functions and their Moreau envelopes on Hadamard spaces. *J. Approx. Theory*, 224:1–12, 2017.
- [23] A. Beck. *First-Order Methods in Optimization*, volume 25 of *MOS-SIAM Series on Optimization*. SIAM, Philadelphia, 2017.
- [24] A. Beck and L. Tetruashvili. On the convergence of block coordinate descent type methods. *SIAM J. Optim.*, 23(4):2037–2060, 2013.
- [25] M. F. Beg, M. I. Miller, A. Trounev, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.*, 61(2):139–157, 2005.
- [26] J. Behrmann, W. Grathwohl, R. T. Chen, D. Duvenaud, and J.-H. Jacobsen. Invertible residual networks. In *Proc. of Machine Learning Research*, volume 97, pages 573–582, 2019.

-
- [27] J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numer. Math.*, 84(3):375–393, 2000.
 - [28] B. Benešová and M. Kružík. Weak lower semicontinuity of integral functionals and applications. *SIAM Rev.*, 59(4):703–766, 2017.
 - [29] R. Bergmann, R. H. Chan, R. Hielscher, J. Persch, and G. Steidl. Restoration of manifold-valued images by half-quadratic minimization. *Inverse Probl. Imaging*, 10(2):281–304, 2016.
 - [30] R. Bergmann, F. Laus, J. Persch, and G. Steidl. Processing manifold-valued images. Technical report, SIAM News October 2017, 2017.
 - [31] R. Bergmann, J. Persch, and G. Steidl. A parallel Douglas–Rachford algorithm for restoring images with values in symmetric Hadamard manifolds. *SIAM J. Imaging Sci.*, 9(3):901–937, 2016.
 - [32] B. Berkels, M. Buchner, A. Effland, M. Rumpf, and S. Schmitz-Valckenberg. GPU-based image geodesics for optical coherence tomography. In *Bildverarbeitung für die Medizin 2017*, pages 68–73, 2017.
 - [33] B. Berkels, A. Effland, and M. Rumpf. Time discrete geodesic paths in the space of images. *SIAM J. Imaging Sci.*, 8(3):1457–1488, 2015.
 - [34] R. J. Berman. The Sinkhorn algorithm, parabolic optimal transport and geometric Monge–Ampère equations. *Numer. Math.*, 145(4):771–836, 2020.
 - [35] B. Bojarski, P. Hajłasz, and P. Strzelecki. Sard’s theorem for mappings in Hölder and Sobolev spaces. *Manuscr. Math.*, 118(3):383–397, 2005.
 - [36] J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Math. Program.*, 146(1-2):459–494, 2014.
 - [37] A. Bondarenko, D. Radchenko, and M. Viazovska. Optimal asymptotic bounds for spherical designs. *Ann. Math.*, 178(2):443–452, 2013.
 - [38] A. Bondarenko, D. Radchenko, and M. Viazovska. Well-separated spherical designs. *Constr. Approx.*, 41(1):93–112, 2015.
 - [39] L. Borg, J. Friel, J. S. Jørgensen, and E. T. Quinto. Analyzing reconstruction artifacts from arbitrary incomplete X-ray CT data. *SIAM J. Imaging Sci.*, 11(4):2786–2814, 2018.
 - [40] A. Borzì, K. Ito, and K. Kunisch. An optimal control approach to optical flow computation. *Int. J. Numer. Methods Fluids*, 40(1-2):231–240, 2002.
 - [41] A. Borzì, K. Ito, and K. Kunisch. Optimal control formulation for determining optical flow. *SIAM J. Sci. Comput.*, 24(3):818–847, 2003.

- [42] A. Borzì and V. Schulz. *Computational Optimization of Systems Governed by Partial Differential Equations*, volume 8 of *Computational Science & Engineering*. SIAM, Philadelphia, 2012.
- [43] C. Boyer, N. Chauffert, P. Ciuciu, J. Kahn, and P. Weiss. On the generation of sampling schemes for magnetic resonance imaging. *SIAM J. Imaging Sci.*, 9(4):2039–2072, 2016.
- [44] A. Braides. *Γ -Convergence for Beginners*. Oxford University Press, Oxford, 2002.
- [45] L. Brandolini, C. Choirat, L. Colzani, G. Gigante, R. Seri, and G. Travaglini. Quadrature rules and distribution of points on manifolds. *Ann. Scuola Norm.-Sci.*, 13(4):889–923, 2014.
- [46] K. Bredies and D. Lorenz. *Mathematische Bildverarbeitung*. Vieweg+Teuber, Wiesbaden, 2011.
- [47] A. Breger, M. Ehler, and M. Gräf. Quasi Monte Carlo integration and kernel-based function approximation on Grassmannians. In *Frames and Other Bases in Abstract and Function Spaces: Novel Methods in Harmonic Analysis*, pages 333–353. Birkhäuser, Basel, 2017.
- [48] M. Bridson and A. Häfliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *A Series of Comprehensive Studies in Mathematics*. Springer, Berlin, 1999.
- [49] M. Bruveris and F.-X. Vialard. On completeness of groups of diffeomorphisms. *J. Eur. Math. Soc.*, 19(5):1507–1544, 2017.
- [50] T. A. Bubba, A. Hauptmann, S. Huotari, J. Rimpeläinen, and S. Silta-nen. Tomographic X-ray data of a lotus root filled with attenuating objects. *arXiv:1609.07299*, 2016.
- [51] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry*, volume 33 of *Graduate Studies in Mathematics*. Amer. Math. Soc., Providence, 2001.
- [52] M. Burger, H. Dirks, and C.-B. Schönlieb. A variational model for joint motion estimation and image reconstruction. *SIAM J. Imaging Sci.*, 11(1):94–128, 2018.
- [53] M. Burger, A. Sawatzky, and G. Steidl. First order algorithms in variational image processing. In *Splitting Methods in Communication, Imaging, Science, and Engineering*, pages 345–407. Springer, Cham, 2016.
- [54] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.*, 16(5):1190–1208, 1995.

-
- [55] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Commun. Pure Appl. Math.*, 59(8):1207–1223, 2006.
 - [56] G. Carlier, V. Duval, G. Peyré, and B. Schmitzer. Convergence of entropic schemes for optimal transport and gradient flows. *SIAM J. Math. Anal.*, 49(2):1385–1418, 2017.
 - [57] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.*, 40(1):120–145, 2011.
 - [58] A. Chambolle and T. Pock. An introduction to continuous optimization for imaging. *Acta Numer.*, 25:161–319, 2016.
 - [59] N. Charon, B. Charlier, and A. Trounev. Metamorphoses of functional shapes in Sobolev spaces. *Found. Comp. Math.*, 18(6):1535–1596, 2018.
 - [60] N. Chauffert, P. Ciuciu, J. Kahn, and P. Weiss. Variable density sampling with continuous trajectories. *SIAM J. Imaging Sci.*, 7(4):1962–1992, 2014.
 - [61] N. Chauffert, P. Ciuciu, J. Kahn, and P. Weiss. A projection method on measures sets. *Constr. Approx.*, 45(1):83–111, 2017.
 - [62] I. Chavel. *Eigenvalues in Riemannian Geometry*. Academic Press, Orlando, 1984.
 - [63] C. Ched'hotel, D. Tschumperlé, R. Deriche, and O. Faugeras. Regularizing flows for constrained matrix-valued images. *J. Math. Imaging Vis.*, 20(1-2):147–162, 2004.
 - [64] C. Chen and O. Öktem. Indirect image registration with large diffeomorphic deformations. *SIAM J. Imaging Sci.*, 11(1):575–617, 2018.
 - [65] K. Chen and D. A. Lorenz. Image sequence interpolation using optimal control. *J. Math. Imaging Vis.*, 41(3):222–238, 2011.
 - [66] Z. Chen, Z. Shen, J. Guo, J. Cao, and X. Zeng. Line drawing for 3D printing. *Comput. Graph.*, 66:85–92, 2017.
 - [67] A. Chertock and A. Kurganov. On a practical implementation of particle methods. *Appl. Numer. Math.*, 56(10-11):1418–1431, 1999.
 - [68] J. Chevallier. Uniform decomposition of probability measures: Quantization, clustering and rate of convergence. *J. Appl. Probab.*, 55(4):1037–1045, 2018.
 - [69] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard. Scaling algorithms for unbalanced optimal transport problems. *Math. Comp.*, 87(314):2563–2609, 2018.

- [70] G. E. Christensen and H. J. Johnson. Consistent image registration. *IEEE Trans. Med. Imaging*, 20(7):568–582, 2001.
- [71] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. Deformable templates using large deformation kinematics. *IEEE Trans. Image Process.*, 5(10):1435–1447, 1996.
- [72] R. Ciak, M. Hirzmann, and O. Scherzer. Regularization with metric double integrals of functions with values in a set of vectors. *J. Math. Imaging Vis.*, 61(6):824–848, 2019.
- [73] P. G. Ciarlet. *Mathematical Elasticity: Vol. I*, volume 20 of *Studies in Mathematics and its Applications*. North-Holland, Amsterdam, 1988.
- [74] C. Clason, D. Lorenz, H. Mahler, and B. Wirth. Entropic regularization of continuous optimal transport problems. *J. Math. Anal. Appl.*, 494(1):124432, 2021.
- [75] F. Colonna, G. Easley, K. Guo, and D. Labate. Radon transform inversion using the shearlet representation. *Appl. Comput. Harmon. Anal.*, 29(2):232–250, 2010.
- [76] R. Cominetti and J. San Martín. Asymptotic analysis of the exponential penalty trajectory in linear programming. *Math. Programming*, 67(1-3):169–187, 1994.
- [77] P. A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, page 2759, Seattle, US, 2006.
- [78] T. Coulhon, E. Russ, and V. Tardivel-Nachef. Sobolev algebras on Lie groups and Riemannian manifolds. *Amer. J. Math.*, 123(2):283–342, 2001.
- [79] G. Crippa. *The Flow Associated to Weakly Differentiable Vector Fields*. PhD thesis, Scuola Normale Superiore di Pisa, 2007.
- [80] F. Cucker and S. Smale. On the mathematical foundations of learning. *Bull. Amer. Math. Soc.*, 39(1):1–49, 2002.
- [81] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems 26*, pages 2292–2300, 2013.
- [82] M. Cuturi and G. Peyré. Computational optimal transport. *Found. Trends Mach. Learn.*, 11(5-6):355–607, 2019.
- [83] B. Dacorogna. *Direct Methods in the Calculus of Variations*, volume 78 of *Applied Mathematical Sciences*. Springer, New York, 2nd edition, 2008.

-
- [84] G. Dal Maso. *An Introduction to Γ -convergence*, volume 8 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser, Boston, 1993.
 - [85] J. W. Daniel. The conjugate gradient method for linear and nonlinear operator equations. *SIAM J. Numer. Anal.*, 4(1):10–26, 1967.
 - [86] M. E. Davison. The ill-conditioned nature of the limited angle tomography problem. *SIAM J. Appl. Math.*, 43(2):428–448, 1983.
 - [87] F. de Gournay, J. Kahn, and L. Lebrat. Differentiation and regularity of semi-discrete optimal transport with respect to the parameters of the discrete measure. *Numer. Math.*, 141(2):429–453, 2019.
 - [88] J. C. De los Reyes. *Numerical PDE-Constrained Optimization*. Springer, Cham, 2015.
 - [89] P. Delsarte, J. M. Goethals, and J. J. Seidel. Spherical codes and designs. *Geom. Dedicata*, 6:363–388, 1977.
 - [90] S. Di Marino and A. Gerolin. An optimal transport approach for the Schrödinger bridge problem and convergence of Sinkhorn algorithm. *J. Sci. Comput.*, 85:27, 2020.
 - [91] J. Dick, M. Ehler, M. Gräf, and C. Krattenthaler. Spectral decomposition of discrepancy kernels on the Euclidean ball, the special orthogonal group, and the Grassmannian manifold. *arXiv:1909.12334*, 2019.
 - [92] R. J. DiPerna and P.-L. Lions. Ordinary differential equations, transport theory and Sobolev spaces. *Invent. Math.*, 98(3):511–547, 1989.
 - [93] D. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, 2006.
 - [94] T. Duchamp and W. Stuetzle. Extremal properties of principal curves in the plane. *Ann. Stat.*, 24(4):1511–1520, 1996.
 - [95] P. Dupuis, U. Grenander, and M. I. Miller. Variational problems on flows of diffeomorphisms for image matching. *Q. Appl. Math.*, 56(3):587–600, 1998.
 - [96] G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. In *Proc. of the 31st Conference on Uncertainty in Artificial Intelligence*, pages 258–267, 2015.
 - [97] A. Effland. *Discrete Riemannian Calculus and A Posteriori Error Control on Shape Spaces*. PhD thesis, University of Bonn, 2017.
 - [98] A. Effland, S. Neumayer, and M. Rumpf. Convergence of the time discrete metamorphosis model on Hadamard manifolds. *SIAM J. Imaging Sci.*, 13(2):557–588, 2020.

- [99] A. Effland, M. Rumpf, and F. Schäfer. Image extrapolation for the time discrete metamorphosis model: Existence and applications. *SIAM J. Imaging Sci.*, 11(1):834–862, 2018.
- [100] A. Effland, M. Rumpf, S. Simon, K. Stahn, and B. Wirth. Bézier curves in the space of images. In *Scale Space and Variational Methods in Computer Vision*, pages 372–384. Springer, Cham, 2015.
- [101] M. Ehler and M. Gräf. Reproducing kernels for the irreducible components of polynomial spaces on unions of Grassmannians. *Constr. Approx.*, 49(1):29–58, 2018.
- [102] M. Ehler, M. Gräf, S. Neumayer, and G. Steidl. Curve based approximation of measures on manifolds by discrepancy minimization. *Foundations of Computational Mathematics*, 2021.
- [103] J. Ehrhardt and C. Lorenz. *4D Modeling and Estimation of Respiratory Motion for Radiation Therapy*. Springer, Berlin, 2013.
- [104] I. Ekeland and R. Témam. *Convex Analysis and Variational Problems*. SIAM, Philadelphia, 1999.
- [105] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*, volume 375 of *Mathematics and its Applications*. Springer Netherlands, Dordrecht, 1996.
- [106] L. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. Amer. Math. Soc., Providence, 1998.
- [107] V. A. Fernández, M. J. Gamero, and J. M. García. A test for the two-sample problem based on empirical characteristic functions. *Comput. Stat. Data Anal.*, 52(7):3730–3748, 2008.
- [108] J. Feydy, T. Séjourné, F.-X. Vialard, S. Amari, A. Trounev, and G. Peyré. Interpolating between optimal transport and MMD using Sinkhorn divergences. In *Proc. of Machine Learning Research*, volume 89, pages 2681–2690. PMLR, 2019.
- [109] F. Filbir and H. N. Mhaskar. Marcinkiewicz–Zygmund measures on manifolds. *J. Complex.*, 27(6):568–596, 2011.
- [110] R. Fiorenza. *Hölder and locally Hölder Continuous Functions, and Open Sets of Class $C^k, C^{k,\lambda}$* . Frontiers in Mathematics. Birkhäuser, Basel, 2016.
- [111] B. Fischer and J. Modersitzki. Curvature based image registration. *J. Math. Imaging Vis.*, 18(1):81–85, 2003.
- [112] T. Fletcher and S. Joshi. Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Process.*, 87(2):250–262, 2007.

-
- [113] I. Fonseca and G. Leoni. *Modern Methods in the Calculus of Variations: L^p Spaces*. Springer, New York, 2007.
 - [114] M. Fornasier, J. Haskovec, and G. Steidl. Consistency of variational continuous-domain quantization via kinetic theory. *Appl. Anal.*, 92(6):1283–1298, 2013.
 - [115] K.-J. Förster and K. Petras. On estimates for the weights in Gaussian quadrature in the ultraspherical case. *Math. Comp.*, 55(191):243–264, 1990.
 - [116] J. Friel. Sparse regularization in limited angle tomography. *Appl. Comput. Harmon. Anal.*, 34(1):117–141, 2013.
 - [117] J. Friel and E. T. Quinto. Characterization and reduction of artifacts in limited angle tomography. *Inverse Prob.*, 29(12):125007, 2013.
 - [118] J. Friel and E. T. Quinto. Limited data problems for the generalized Radon transform in \mathbb{R}^n . *SIAM J. Math. Anal.*, 48(4):2301–2318, 2016.
 - [119] M. Fuchs, B. Jüttler, O. Scherzer, and H. Yang. Shape metrics based on elastic deformations. *J. Math. Imaging Vis.*, 35(1):68–102, 2009.
 - [120] W. Fulton and J. Harris. *Representation Theory: A First Course*. Springer, New York, 1991.
 - [121] E. Fuselier and G. B. Wright. Scattered data interpolation on embedded submanifolds with restricted positive definite kernels: Sobolev error estimates. *SIAM J. Numer. Anal.*, 50(3):1753–1776, 2012.
 - [122] A. Genevay, L. Chizat, F. Bach, M. Cuturi, and G. Peyré. Sample complexity of Sinkhorn divergences. In *Proc. of Machine Learning Research*, volume 89, pages 1574–1583. PMLR, 2019.
 - [123] A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems 29*, pages 3440–3448, 2016.
 - [124] S. Gerber and R. Whitaker. Regularization-free principal curve estimation. *J. Mach. Learn. Res.*, 14(1):1285–1302, 2013.
 - [125] G. Gigante and P. Leopardi. Diameter bounded equal measure partitions of Ahlfors regular metric measure spaces. *Discrete Comput. Geom.*, 57(2):419–430, 2017.
 - [126] F. Gigengack, X. Jiang, M. Dawood, and K. P. Schäfers. *Motion Correction in Thoracic Positron Emission Tomography*. Springer, Cham, 2015.
 - [127] D. Gilbarg and N. S. Trudinger. *Elliptic Partial Differential Equations of Second Order*, volume 224 of *Classics in Mathematics*. Springer, Berlin, 2001.

- [128] M. Gnewuch. Weighted geometric discrepancies and numerical integration on reproducing kernel Hilbert spaces. *J. Complex.*, 28(1):2–17, 2012.
- [129] F. D. Goes, K. Breeden, V. Ostromoukhov, and M. Desbrun. Blue noise through optimal transport. *ACM Trans. Graphics*, 31(6):171–182, 2012.
- [130] M. Gräf. A unified approach to scattered data approximation on \mathbb{S}^3 and $\text{SO}(3)$. *Adv. Comput. Math.*, 37(3):379–392, 2012.
- [131] M. Gräf. *Efficient Algorithms for the Computation of Optimal Quadrature Points on Riemannian Manifolds*. PhD thesis, TU Chemnitz, 2013.
- [132] M. Gräf and D. Potts. Sampling sets and quadrature formulae on the rotation group. *Numer. Funct. Anal. Optim.*, 30(7-8):665–688, 2009.
- [133] M. Gräf and D. Potts. On the computation of spherical designs by a new optimization approach based on fast spherical Fourier transforms. *Numer. Math.*, 119(4):699–724, 2011.
- [134] M. Gräf, M. Potts, and G. Steidl. Quadrature errors, discrepancies and their relations to halftoning on the torus and the sphere. *SIAM J. Sci. Comput.*, 34(5):2760–2791, 2013.
- [135] A. E. Green and R. S. Rivlin. Multipolar continuum mechanics. *Arch. Rational Mech. Anal.*, 17(2):113–147, 1964.
- [136] A. E. Green and R. S. Rivlin. Simple force and stress multipoles. *Arch. Rational Mech. Anal.*, 16(5):325–353, 1964.
- [137] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A kernel method for the two-sample-problem. In *Advances in Neural Information Processing Systems 19*, pages 513–520, 2007.
- [138] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *J. Mach. Learn. Res.*, 13(1):723–773, 2012.
- [139] B. Gris, C. Chen, and O. Öktem. Image reconstruction through metamorphosis. *Inverse Prob.*, 36(2):025001, 2020.
- [140] K. Gröchenig, J. L. Romero, J. Unnikrishnan, and M. Vetterli. On minimal trajectories for mobile sampling of bandlimited fields. *Appl. Comput. Harmon. Anal.*, 39(3):487–510, 2015.
- [141] M. Guerquin-Kern, L. Lejeune, K. P. Pruessmann, and M. Unser. Realistic analytical phantoms for parallel magnetic resonance imaging. *IEEE Trans. Med. Imaging*, 31(3):626–636, 2012.
- [142] M. D. Gunzburger. *Perspectives in Flow Control and Optimization*, volume 5 of *Advances in Design and Control*. SIAM, Philadelphia, 2003.

-
- [143] E. Haber and J. Modersitzki. A multilevel method for image registration. *SIAM J. Sci. Comput.*, 27(5):1594–1607, 2006.
 - [144] P. Hagemann and S. Neumayer. Stabilizing invertible neural networks using mixture models. *arXiv:2009.02994*, 2020.
 - [145] P. Hajlasz. Sobolev spaces on metric-measure spaces. In *Heat Kernels and Analysis on Manifolds, Graphs, and Metric Spaces*, volume 338 of *Contemp. Math.*, pages 173–218. Amer. Math. Soc., Providence, 2003.
 - [146] K. Hämäläinen, L. Harhanen, A. Kallonen, A. Kujanpää, E. Niemi, and S. Siltanen. Tomographic X-ray data of walnut. *arXiv:1502.04064*, 2015.
 - [147] J. Han, B. Berkels, M. Droske, J. Hornegger, M. Rumpf, C. Schaller, J. Scorzin, and H. Urbach. Mumford–Shah model for one-to-one edge matching. *IEEE Trans. Image Process.*, 16(11):2720–2732, 2007.
 - [148] G. L. Hart, C. Zach, and M. Niethammer. An optimal control approach for deformable registration. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 9–16. IEEE, New York, 2009.
 - [149] C. Hartmann, A. F. Weiss, P. Lechner, W. Volk, S. Neumayer, J. H. Fitschen, and G. Steidl. Measurement of strain, strain rate and crack evolution in shear cutting. *J. Mater. Process. Technol.*, 288:116872, 2021.
 - [150] M. Hasannasab, J. Hertrich, S. Neumayer, G. Plonka, S. Setzer, and G. Steidl. Parseval proximal neural networks. *J. Fourier Anal. Appl.*, 26:59, 2020.
 - [151] T. Hastie and W. Stuetzle. Principal curves. *J. Am. Stat. Assoc.*, 84(406):502–516, 1989.
 - [152] S. Hauberg. Principal curves on Riemannian manifolds. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(9):1915–1921, 2015.
 - [153] H. He and W.-C. Siu. Single image super-resolution using Gaussian process regression. In *2011 IEEE Conference on Computer Vision and Pattern Recognition*, pages 449–456. IEEE, New York, 2011.
 - [154] G. T. Herman. *Image Reconstructions from Projections: The Fundamentals of Computerized Tomography*. Springer, London, 2nd edition, 2009.
 - [155] G. T. Herman and R. Davidi. Image reconstruction from a small number of projections. *Inverse Prob.*, 24(4):045011, 2008.
 - [156] J. Hertrich, M. Bačák, S. Neumayer, and G. Steidl. Minimal Lipschitz extensions for vector-valued functions on finite graphs. In *Scale Space and Variational Methods in Computer Vision*, pages 183–195. Springer, Cham, 2019.

- [157] R. Herzog and K. Kunisch. Algorithms for PDE-constrained optimization. *GAMM-Mitteilungen*, 33(2):163–176, 2010.
- [158] K. Hesse, H. N. Mhaskar, and I. H. Sloan. Quadrature in Besov spaces on the Euclidean sphere. *J. Complex.*, 23(4-6):528–552, 2007.
- [159] R. J. Hewett, I. Jermyn, M. Heath, and F. Kamalabadi. A phase field method for tomographic reconstruction from limited data. In *Proc. of the British Machine Vision Conference*, pages 1–11. BMVA Press, 2012.
- [160] J. Hinkle, M. Szegedi, B. Wang, B. Salter, and S. Joshi. 4D CT image reconstruction with diffeomorphic motion model. *Med. Image Anal.*, 16(6):1307–1316, 2012.
- [161] M. Hinze, R. Pinnau, M. Ulbrich, and S. Ulbrich. *Optimization with PDE Constraints*, volume 23 of *Mathematical Modelling: Theory and Applications*. Springer, New York, 2009.
- [162] Y. Hong, S. Joshi, M. Sanchez, M. Styner, and M. Niethammer. Metamorphic geodesic regression. In *Proc. of International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 197–205, 2012.
- [163] L. Hörmander. *The Analysis of Linear Partial Differential Operators I*. Springer, Berlin, 1983.
- [164] B. K. Horn and B. G. Schunck. Determining optical flow. *Artif. Intell.*, 17(1-3):185–203, 1981.
- [165] T. Hytönen, J. van Neerven, M. Veraar, and L. Weis. *Analysis in Banach spaces-Vol. I: Martingales and Littlewood-Paley theory*, volume 63 of *A Series of Modern Surveys in Mathematics*. Springer, Cham, 2016.
- [166] H. Inci, T. Kappeler, and P. Topalov. On the regularity of the composition of diffeomorphisms. *Mem. Am. Math. Soc.*, 226:vi+60, 2013.
- [167] A. T. James and A. G. Constantine. Generalized Jacobi polynomials as spherical functions of the Grassmann manifold. *Proc. London Math. Soc.*, 29(3):174–192, 1974.
- [168] S. C. Joshi and M. I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Trans. Image Process.*, 9(8):1357–1370, 2000.
- [169] J. Jost. *Nonpositive Curvature: Geometric and Analytic Aspects*. Lectures in Mathematics ETH Zürich. Birkhäuser, Basel, 1997.
- [170] C. S. Kaplan and R. Bosch. TSP art. In *Renaissance Banff: Mathematics, Music, Art, Culture*, pages 301–308. Bridges Conference, 2005.

-
- [171] J. Karlsson and A. Ringh. Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport. *SIAM J. Imaging Sci.*, 10(4):1935–1962, 2017.
- [172] A. I. Katsevich. Local tomography for the limited-angle problem. *J. Math. Anal. Appl.*, 213(1):160–182, 1997.
- [173] B. Kégl, A. Krzyzak, T. Linder, and K. Zeger. Learning and design of principal curves. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(3):281–297, 2000.
- [174] J. Keiner, S. Kunis, and D. Potts. Using NFFT3 – a software library for various nonequispaced fast Fourier transforms. *ACM Trans. Math. Software*, 36(4):1–30, 2009.
- [175] J.-H. Kim, J. Lee, and H.-S. Oh. Spherical principal curves. *IEEE Trans. Pattern. Anal. Mach. Intell.*, 2020.
- [176] W. P. A. Klingenberg. *Riemannian Geometry*, volume 1 of *De Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, 2nd edition, 1995.
- [177] B. Kloeckner. Approximation by finitely supported measures. *ESAIM Control Opt. Calc. Var.*, 18(2):343–359, 2012.
- [178] P. Koltai, J. von Lindheim, S. Neumayer, and G. Steidl. Transfer operators from optimal transport plans for coherent set detection. *arXiv:2006.16085*, 2020.
- [179] P. Kuchment. *The Radon Transform and Medical Imaging*. SIAM, Philadelphia, 2014.
- [180] L. Kuipers and H. Niederreiter. *Uniform Distribution of Sequences*. Wiley, New York, 1974.
- [181] L. Lang, S. Neumayer, O. Öktem, and C.-B. Schönlieb. Template-based image reconstruction from sparse tomographic data. *Appl. Math. Optim.*, 82:1081–1109, 2020.
- [182] C. Lazarus, P. Weiss, N. Chauffert, F. Mauconduit, L. El Gueddari, C. Destrieux, I. Zemmoura, A. Vignaud, and P. Ciuciu. SPARKLING: Variable-density k-space filling curves for accelerated T_2^* -weighted MRI. *Magn. Reson. Med.*, 81(6):3643–3661, 2019.
- [183] L. Lebrat, F. de Gournay, J. Kahn, and P. Weiss. Optimal transport approximation of 2-dimensional measures. *SIAM J. Imaging Sci.*, 12(2):762–787, 2019.
- [184] J. M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, 2012.

- [185] J. Lellmann, S. Neumayer, M. Nimmer, and G. Steidl. Methods for finding the offset in robust subspace fitting. *PAMM*, 19(1):e201900323, 2019.
- [186] J. Lellmann, E. Strekalovskiy, S. Koetter, and D. Cremers. Total variation regularization for functions with values in a manifold. In *2013 IEEE International Conference on Computer Vision*, pages 2944–2951, 2013.
- [187] C. Léonard. From the Schrödinger problem to the Monge–Kantorovich problem. *J. Funct. Anal.*, 262(4):1879–1920, 2012.
- [188] M. Liero, A. Mielke, and G. Savaré. Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures. *Invent. Math.*, 211(3):969–1117, 2018.
- [189] D. Lorenz, P. Manns, and C. Meyer. Quadratically regularized optimal transport. *Appl. Math. Optim.*, 2019.
- [190] A. K. Louis. Incomplete data problems in X-ray computerized tomography. *Numer. Math.*, 48(3):251–262, 1986.
- [191] J. Maas, M. Rumpf, C.-B. Schönlieb, and S. Simon. A generalized model for optimal transport of images including dissipation and density modulation. *ESAIM Math. Model. Numer. Anal.*, 49(6):1745–1769, 2015.
- [192] P. Malliavin. *Integration and Probability*, volume 157 of *Graduate Texts in Mathematics*. Springer, New York, 1995.
- [193] A. Mang and L. Ruthotto. A Lagrangian Gauss–Newton–Krylov solver for mass- and intensity-preserving diffeomorphic image registration. *SIAM J. Sci. Comput.*, 39(5):860–885, 2017.
- [194] J. Matousek. *Geometric Discrepancy*, volume 18 of *Algorithms and Combinatorics*. Springer, Berlin, 2010.
- [195] J. Mercer. Functions of positive and negative type and their connection with the theory of integral equations. *Philos. Trans. Roy. Soc. London Ser. A*, 209(441-458):415–446, 1909.
- [196] H. N. Mhaskar. Eignets for function approximation on manifolds. *Appl. Comput. Harmon. Anal.*, 29(1):63–87, 2010.
- [197] H. N. Mhaskar. Approximate quadrature measures on data-defined spaces. In *Contemporary Computational Mathematics - A Celebration of the 80th Birthday of Ian Sloan*. Springer, Cham, 2018.
- [198] C. A. Micchelli. Interpolation of scattered data: Distance matrices and conditionally positive definite functions. *Constr. Approx.*, 2(1):11–22, 1986.

-
- [199] M. I. Miller, A. Trouné, and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annu. Rev. Biomed. Eng.*, 4(1):375–405, 2002.
 - [200] M. I. Miller, A. Trouné, and L. Younes. Geodesic shooting for computational anatomy. *J. Math. Imaging Vis.*, 24(2):209–228, 2006.
 - [201] M. I. Miller, A. Trouné, and L. Younes. Hamiltonian systems and optimal control in computational anatomy: 100 years since D’Arcy Thompson. *Annu. Rev. Biomed. Eng.*, 17:447–509, 2015.
 - [202] M. I. Miller and L. Younes. Group actions, homeomorphisms, and matching: A general framework. *Int. J. Comput. Vis.*, 41(1-2):61–84, 2001.
 - [203] M. Moakher and P. G. Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. In *Visualization and Processing of Tensor Fields*, pages 285–298. Springer, Berlin, 2006.
 - [204] J. Modersitzki. *Numerical Methods for Image Registration*. Oxford University Press, Oxford, 2004.
 - [205] J. Modersitzki. *FAIR: Flexible Algorithms for Image Registration*. SIAM, Philadelphia, 2009.
 - [206] U. Mosco. Convergence of convex sets and of solutions of variational inequalities. *Adv. Math.*, 3(4):510–585, 1969.
 - [207] C. Müller. *Spherical Harmonics*, volume 17 of *Lecture Notes in Mathematics*. Springer, Berlin, 1992.
 - [208] F. Natterer. *The Mathematics of Computerized Tomography*, volume 32 of *Classics in Applied Mathematics*. SIAM, Philadelphia, 2001.
 - [209] I. Navrotskaya and P. J. Rabier. $L \log L$ and finite entropy. *Adv. Nonlinear Anal.*, 2(4):379–387, 2013.
 - [210] S. Neumayer, M. Nimmer, S. Setzer, and G. Steidl. On the robust PCA and Weiszfeld’s algorithm. *Appl. Math. Optim.*, 82:1017–1048, 2020.
 - [211] S. Neumayer, M. Nimmer, S. Setzer, and G. Steidl. On the rotational invariant L_1 -norm PCA. *Linear Algebra Appl.*, 587:243–270, 2020.
 - [212] S. Neumayer, M. Nimmer, G. Steidl, and H. Stephani. On a projected Weiszfeld algorithm. In *Scale Space and Variational Methods in Computer Vision*, pages 486–497. Springer, Cham, 2017.
 - [213] S. Neumayer, J. Persch, and G. Steidl. Morphing of manifold-valued images inspired by discrete geodesics in image spaces. *SIAM J. Imaging Sci.*, 11(3):1898–1930, 2018.

- [214] S. Neumayer, J. Persch, and G. Steidl. Regularization of inverse problems via time discrete geodesics in image spaces. *Inverse Prob.*, 35(5):055005, 2019.
- [215] S. Neumayer and G. Steidl. From optimal transport to discrepancy. *To appear in Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging*, 2020.
- [216] J. Nečas and M. Šilhavý. Multipolar viscous fluids. *Quart. Appl. Math.*, 49(2):247–265, 1991.
- [217] L. H. Nguyen, S. K. F. Stoter, T. Baum, J. S. Kirschke, M. Ruess, Z. Yosibash, and D. Schillinger¹. Phase-field boundary conditions for the voxel finite cell method: Surface-free stress analysis of CT-based bone structures. *Int. J. Numer. Method. Biomed. Eng.*, 33(12):1–34, 2017.
- [218] L. V. Nguyen. How strong are streak artifacts in limited angle computed tomography? *Inverse Prob.*, 31(5):055003, 2015.
- [219] M. Niethammer, G. L. Hart, and C. Zach. An optimal control approach for the registration of image time-series. In *Proc. of the 48th IEEE Conference on Decision and Control*, pages 262–270. IEEE, New York, 2009.
- [220] L. Nirenberg. An extended interpolation inequality. *Ann. Scuola Norm.-Sci.*, 20(4):733–737, 1966.
- [221] M. A. Nogueira, P. H. Abreu, P. Martins, P. Machado¹, H. Duarte, and J. Santos. Image descriptors in radiology images: A systematic review. *Artif. Intell. Rev.*, 47(4):531–559, 2017.
- [222] E. Novak and H. Wozniakowski. *Tractability of Multivariate Problems. Volume II*, volume 12 of *EMS Tracts in Mathematics*. EMS Publishing House, Zürich, 2010.
- [223] O. Öktem, C. Chen, N. O. Domaniç, P. Ravikumar, and C. Bajaj. Shape-based image reconstruction using linearized deformations. *Inverse Prob.*, 33(3):035004, 2017.
- [224] W. Palenstijn, K. Batenburg, and J. Sijbers. Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs). *J. Struct. Biol.*, 176(2):250–253, 2011.
- [225] X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *Int. J. Comput. Vis.*, 66(1):41–66, 2006.
- [226] J. Persch, F. Pierre, and G. Steidl. Exemplar-based face colorization using image morphing. *J. Imaging*, 3(4):48, 2017.
- [227] G. Peyré. Entropic Wasserstein gradient flows. *SIAM J. Imaging Sci.*, 8(4):2323–2351, 2015.

-
- [228] G. Plonka, D. Potts, G. Steidl, and M. Tasche. *Numerical Fourier Analysis*. Birkhäuser, Basel, 2019.
- [229] T. Pock, A. Chambolle, D. Cremers, and H. Bischof. A convex relaxation approach for computing minimal partitions. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 810–817, 2009.
- [230] T. Pock and S. Sabach. Inertial proximal alternating linearized minimization (iPALM) for nonconvex and nonsmooth problems. *SIAM J. Imaging Sci.*, 9(4):1756–1787, 2016.
- [231] C. L. Richardson and L. Younes. Computing metamorphoses between discrete measures. *J. Geom. Mech.*, 5(1):131–150, 2013.
- [232] C. L. Richardson and L. Younes. Metamorphosis of images in reproducing kernel Hilbert spaces. *Adv. Comput. Math.*, 42(3):573–603, 2016.
- [233] W. Ring and B. Wirth. Optimization methods on Riemannian manifolds and their application to shape space. *SIAM J. Optim.*, 22(2):596–627, 2012.
- [234] F. Rocca, C. Prati, and A. M. Guarnieri. Possibilities and limits of SAR interferometry. In *Proc. of the First Latino-American Seminar on Radar Remote Sensing – Image Processing Technique*, pages 15–26, 1997.
- [235] J. Roe. *Elliptic Operators, Topology and Asymptotic Methods*. Longman, Harlow, 2nd edition, 1998.
- [236] P. A. Rosen, S. Hensley, I. R. Joughin, F. K. Li, S. N. Madsen, E. Rodriguez, and R. M. Goldstein. Synthetic aperture radar interferometry. *Proc. IEEE*, 88(3):333–382, 2000.
- [237] A. Roy. Bounds for codes and designs in complex subspaces. *J. Algebr. Comb.*, 31(1):1–32, 2010.
- [238] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1-4):259–268, 1992.
- [239] W. Rudin. *Real and Complex Analysis*. McGraw-Hill, New York, 3rd edition, 1987.
- [240] M. Rumpf and B. Wirth. Discrete geodesic calculus in shape space and applications in the space of viscous fluidic objects. *SIAM J. Imaging Sci.*, 6(4):2581–2602, 2013.
- [241] M. Rumpf and B. Wirth. Variational time discretization of geodesic calculus. *IMA J. Numer. Anal.*, 35(3):1011–1046, 2015.
- [242] L. Rüschendorf. Convergence of the iterative proportional fitting procedure. *Ann. Statist.*, 23(4):1160–1174, 1995.

- [243] F. Santambrogio. *Optimal Transport for Applied Mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser, Basel, 2015.
- [244] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational Methods in Imaging*. Springer, New York, 2009.
- [245] C. Schmaltz, P. Gwosdek, A. Bruhn, and J. Weickert. Electrostatic halftoning. *Comp. Graph. For.*, 29(8):2313–2327, 2010.
- [246] I. J. Schoenberg. Metric spaces and completely monotone functions. *Ann. Math.*, 39(4):811–841, 1938.
- [247] H. Schumacher, J. Modersitzki, and B. Fischer. Combined reconstruction and motion correction in SPECT imaging. *IEEE Trans. Nucl. Sci.*, 56(1):73–80, 2009.
- [248] N. Singh, J. Hinkle, S. Joshi, and P. T. Fletcher. A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In *Proc. of the 10th International Symposium on Biomedical Imaging*, pages 127–142. IEEE, New York, 2013.
- [249] R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Ann. Math. Statist.*, 35(2):876–879, 1964.
- [250] S. T. Smith. Optimization techniques on Riemannian manifolds. In *Hamiltonian and Gradient Flows, Algorithms and Control*, volume 3 of *Fields Inst. Commun.*, pages 113–136. Amer. Math. Soc., Providence, 1994.
- [251] D. B. Smythe. A two-pass mesh warping algorithm for object transformation and image interpolation. Technical report, ILM Technical Memo Department, Lucasfilm Ltd, 1990.
- [252] J. M. Steele. Growth rates of Euclidean minimum spanning trees with power weighted edges. *Ann. Probab.*, 16(4):1767–1787, 1988.
- [253] J. M. Steele and T. L. Snyder. Worst-case growth rates of some classical problems of combinatorial optimization. *SIAM J. Comput.*, 18(2):278–287, 1989.
- [254] E. M. Stein. *Singular Integrals and Differentiability Properties of Functions*, volume 30 of *Princeton Mathematical Series*. Princeton University Press, Princeton, 1970.
- [255] E. M. Stein and R. Shakarchi. *Real Analysis: Measure Theory, Integration, and Hilbert Spaces*, volume 3 of *Princeton Lectures in Analysis*. Princeton University Press, Princeton, 2005.
- [256] I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, New York, 2008.

-
- [257] I. Steinwart and C. Scovel. Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constr. Approx.*, 35(3):363–417, 2011.
 - [258] P. Swierczynski, B. W. Papiez, J. A. Schnabel, and C. Macdonald. A level-set approach to joint image segmentation and registration with application to CT lung imaging. *Comput. Med. Imaging Graph.*, 65:58–68, 2018.
 - [259] T. Teuber, G. Steidl, P. Gwosdek, C. Schmaltz, and J. Weickert. Dithering by differences of convex functions. *SIAM J. Imaging Sci.*, 4(1):79–108, 2011.
 - [260] H. Triebel. *Theory of Function Spaces II*. Birkhäuser, Basel, 1992.
 - [261] A. Trouvé. An infinite dimensional group approach for physics based models in pattern recognition. *Technical report*, 1995.
 - [262] A. Trouvé. Diffeomorphisms groups and pattern matching in image analysis. *Int. J. Comput. Vis.*, 28(3):213–221, 1998.
 - [263] A. Trouvé and L. Younes. Local geometry of deformable templates. *SIAM J. Math. Anal.*, 37(2):17–59, 2005.
 - [264] A. Trouvé and L. Younes. Metamorphoses through Lie group action. *Found. Comput. Math.*, 5(2):173–198, 2005.
 - [265] A. Trouvé and L. Younes. Shape spaces. In *Handbook of Mathematical Methods in Imaging*, pages 1759–1817. Springer, New York, 2015.
 - [266] C. Udriște. *Convex Functions and Optimization Methods on Riemannian Manifolds*, volume 297 of *Mathematics and its Applications*. Springer, Dordrecht, 1994.
 - [267] T. Valkonen, K. Bredies, and F. Knoll. Total generalized variation in diffusion tensor imaging. *SIAM J. Imaging Sci.*, 6(1):487–525, 2013.
 - [268] W. van Aarle, W. J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabravolski, J. De Beenhouwer, K. J. Batenburg, and J. Sijbers. Fast and flexible X-ray tomography using the ASTRA toolbox. *Opt. Express*, 24(22):25129–25147, 2016.
 - [269] W. van Aarle, W. J. Palenstijn, J. De Beenhouwer, T. Altantzis, S. Bals, K. J. Batenburg, and J. Sijbers. The ASTRA toolbox: A platform for advanced algorithm development in electron tomography. *Ultramicroscopy*, 157:35–47, 2015.
 - [270] D. Varshalovich, A. Moskalev, and V. Khersonskii. *Quantum Theory of Angular Momentum*. World Scientific, Singapore, 1988.
 - [271] F.-X. Vialard. An elementary introduction to entropic regularization and proximal methods for numerical optimal transport. Lecture, May 2019.

- [272] F.-X. Vialard, L. Risser, D. Rueckert, and C. J. Cotter. Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. *Int. J. Comput. Vis.*, 97(2):229–241, 2011.
- [273] F.-X. Vialard and F. Santambrogio. Extension to BV functions of the large deformation diffeomorphisms matching approach. *Comptes Rendus Math.*, 347(1-2):27–32, 2009.
- [274] C. Villani. *Topics in Optimal Transportation*. Amer. Math. Soc., Providence, 2003.
- [275] G. Wagner. On means of distances on the surface of a sphere II (upper bounds). *Pacific J. Math.*, 154(2):381–396, 1992.
- [276] G. Wagner and B. Volkmann. On averaging sets. *Monatsh. Math.*, 111(1):69–78, 1991.
- [277] A. Weinmann, L. Demaret, and M. Storath. Total variation regularization for manifold-valued data. *SIAM J. Imaging Sci.*, 7(4):2226–2257, 2014.
- [278] H. Wendland. *Scattered Data Approximation*, volume 17 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2004.
- [279] R. L. Wheeden and A. Zygmund. *Measure and Integral: An Introduction to Real Analysis*. CRC Press, Boca Raton, 2015.
- [280] A. G. Wilson. The use of entropy maximising models in the theory of trip distribution, mode split and route split. *J. Transp. Econ. Policy*, 3(1):108–126, 1969.
- [281] G. Wolberg. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, 1990.
- [282] G. Wolberg. Image morphing: A survey. *Vis. Comput.*, 14(8):360–372, 1998.
- [283] W. Yang, L. Zhong, L. Lin, Y. Chen, Z. Lu, S. Liu, Y. Wu, Q. Feng, and W. Chen. Predicting CT image from MRI data through feature matching with learned nonlinear local descriptors. *IEEE Trans. Med. Imaging*, 37(4):977–987, 2018.
- [284] X. Yang, R. Kwitt, M. Styner, and M. Niethammer. Quicksilver: Fast predictive image registration—a deep learning approach. *Neuroimage*, 158:378–396, 2017.
- [285] L. Younes. *Shapes and Diffeomorphisms*. Springer, Berlin, 2010.
- [286] J. Yuan, C. Schnörr, and G. Steidl. Simultaneous higher order optical flow estimation and decomposition. *SIAM J. Sci. Comput.*, 29(6):2283–2304, 2007.

- [287] G. U. Yule. On the methods of measuring association between two attributes.
J. R. Stat. Soc., 75(6):579–652, 1912.