

Geodesics, Splines, and Embeddings in Spaces of Images

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von
Marko Rajković
aus
Nevesinje, Republika Srpska, Bosnien und Herzegowina

Bonn, Februar 2023

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

1. Gutachter: Prof. Dr. Martin Rumpf
2. Gutachter: Prof. Dr. Benedikt Wirth

Tag der Promotion: 07. 06. 2023
Erscheinungsjahr: 2023

Abstract

The main contribution of this thesis concerns two larger mathematical topics. As a first topic, we study the geodesic and spline interpolation of images in the framework of the metamorphosis model. The second topic is learning low bending and low distortion image manifold embeddings.

The metamorphosis model, originally introduced by Miller, Trounev, and Younes, is a well-known approach for image morphing: the problem of computing a visually appealing transition between two images such that semantically corresponding regions are mapped onto each other. The Riemannian manifold of images based on the metamorphosis model defines the shortest geodesic paths interpolating two images as minimizers of the path energy measuring the viscous dissipation caused by the motion field and dissipation caused by the material derivative of the image intensity change along the motion paths. Starting from the general framework for a variational time discretization in geodesic calculus developed by Rumpf and Wirth, a variational time discretization of the metamorphosis model for square-integrable images was proposed by Berkels, Effland, and Rumpf. They showed the existence of discrete geodesic paths, the Mosco-convergence of the time discrete to the time continuous model, and provided numerical results based on a finite element discretization.

The above-mentioned approach does not consider the deep features that go beyond the obvious ones based on grayscale or color representation, leading to a blending along the geodesic paths instead of flow-induced geometric transformations. Furthermore, strong smoothness, implied by the homogeneous and isotropic variational prior for the deformation fields, prevents depth discontinuities of the scene, preventing the preservation of sharp interfaces, such as object boundaries, along the geodesic paths. To overcome the problems originating from the color-based matching, in this thesis, we advocate a metamorphosis model in a deep feature space, which amounts to replacing the input images with feature vectors that combine image intensities and semantic information generated by a feature extraction operator. The deep feature extraction operator is represented by the deep convolution neural network introduced by Visual Geometry Group in Oxford (VGG). To explicitly allow for discontinuities in the deformation fields, we apply an anisotropic regularization operator depending on the magnitude of image gradients. As our main contribution, following the approach by Berkels, Effland, and Rumpf, we prove the existence of discrete geodesic paths for the deep feature metamorphosis model with an anisotropic regularization and discuss its Mosco-convergence to the appropriate time continuous model. In particular, this implies the convergence of time discrete geodesic paths to their time continuous counterparts and establishes the existence of time continuous geodesics as minimizers of the time continuous model. The fully discrete model is based on the finite difference approximation of deformation fields and cubic Hermite spline approximation of the warping operation. The numerical results, obtained by optimization via the iPALM algorithm, developed by Pock and Sabach, allow efficient and robust computation of morphing sequences that visually outperform the color-based morphing. The described approach is based on a publication [1] as a joint project with Alexander Effland, Erich Kobler, Thomas Pock, and Martin Rumpf.

The above approach allows for a smooth and natural interpolation between two fixed images. The task that naturally follows is finding an as smooth as possible interpolation between multiple fixed key frame images as data points. Both the color-based and the deep feature-based piecewise geodesic interpolation of images in the framework of the metamorphosis model lack smoothness at the interpolation data points, leading to jerking at the corresponding times when animating the resulting family of images. Cubic splines, extensively studied in several different shape spaces, are a widespread and versatile solution for this problem. To this end, picking up the general observation of cubic splines in Euclidean space as minimizers of the integral of squared acceleration along the interpolation path, we propose a spline interpolation approach in the framework of the metamorphosis model. Our approach is based on the minimization of a functional which combines quadratic functionals of the Eulerian motion acceleration and the second material derivative of the image intensity as the proper notion of image intensity acceleration. This physically intuitive separation of flow and image intensity acceleration, however, does not lead to a fully Riemannian model. Indeed, penalization of the squared covariant derivative of the path velocity in the Riemannian metric would lead to an interwoven model of these different types of acceleration. Building upon the above-discussed model, we propose a variational time discretization of the spline model and study the convergence to the time continuous model in the context of Mosco-convergence. As a byproduct, this also allows

for establishing the existence of metamorphosis splines for given key frame images as minimizers of the time continuous spline functional. The time discretization is complemented by a spatial discretization based on finite differences and a stable B-spline approximation of warping operation. A variety of numerical results, based on numerical optimization via the iPALM algorithm, demonstrates the robustness and versatility of the proposed method in applications, and highlights smoothness, stability, and different qualitative properties in comparison with the piecewise geodesic interpolation. The described approach had been initially introduced in [2], while its rigorous mathematical analysis was given in [3]. Both present a joint effort with Jorge Justiniano and Martin Rumpf.

Processing of high-dimensional data in their original form, with underlying very high- or even infinite-dimensional manifold, an example of which is the metamorphosis model, is a challenging task, both in theoretical and numerical computation senses. To this end, it represents a substantial simplification if the data lies on a low-dimensional manifold. This leads to the task of the identification of given data sets consisting of objects embedded in a very high-dimensional space with points in a low-dimensional manifold, as one of the central machine learning tasks. Autoencoders, which consist of an encoder and a decoder, represent a widely used tool for solving this task. The encoder embeds the input data manifold into a lower-dimensional latent space, while the decoder represents the inverse map, providing a parametrization of the data manifold by the manifold in latent space. Further data processing tasks, an example of which is data interpolation, may be substantially simplified by the good regularity and structure of the embedded manifold. To this end, we propose and analyze a regularization for learning the encoder component of an autoencoder, based on a loss functional that prefers isometric and extrinsically flat embeddings. The input data for the training consists of pairs of nearby points on the input manifold, measured in terms of sampling radius, together with their Riemannian distance and Riemannian average. These values can be efficiently evaluated due to the hand-crafted construction of the datasets. The discrete sampling loss functional is computed via Monte Carlo integration and it depends on the employed sampling strategy. We first obtain the sampling-dependent nonlocal continuous loss functionals as the limit as the number of samples tends to infinity. As the main contribution, the Mosco-limit of these functionals, as the sampling radius tends to zero, is identified as a purely local geometric loss functional that again promotes isometry and flatness of the embedding. Numerical tests, using image data that encodes different explicitly given data manifolds, show that smooth manifold embeddings into latent space are obtained. Furthermore, due to the promotion of extrinsic flatness, these embeddings allow reasonable approximation of interpolation between nearby points on the manifold by the linear interpolation in the latent space. The described approach had been initially proposed in [4], while its rigorous mathematical analysis was presented in [5]. Both are the results of a joint project with Juliane Braunsmann, Martin Rumpf, and Benedikt Wirth.

References

- [1] A. Effland, E. Kobler, T. Pock, M. Rajković, and M. Rumpf. Image morphing in deep feature spaces: theory and applications. *J. Math. Imaging Vis.*, 63(2):309–327, 2021. <https://doi.org/10.1007/s10851-020-00974-5>
- [2] J. Justiniano, M. Rajković, and M. Rumpf. Splines for image metamorphosis. In *Proc. of International Conference on Scale Space and Variational Methods in Computer Vision*, volume 12679 of *Lecture Notes in Computer Science*, pages 463–475. Springer, Cham, 2021. https://doi.org/10.1007/978-3-030-75549-2_37
- [3] J. Justiniano, M. Rajković, and M. Rumpf. Consistent approximation of interpolating splines in image metamorphosis. *J. Math. Imaging Vis.*, 65(1):29–52, 2023. <https://doi.org/10.1007/s10851-022-01128-5>
- [4] J. Braunsmann, M. Rajković, M. Rumpf, and B. Wirth. Learning low bending and low distortion manifold embeddings. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 4416–4424, 2021. <https://doi.org/10.1109/CVPRW53098.2021.00498>
- [5] J. Braunsmann, M. Rajković, M. Rumpf, and B. Wirth. Convergent autoencoder approximation of low bending and low distortion manifold embeddings. *arXiv preprint, submitted*, 2022. <https://arxiv.org/abs/2208.10193>


Contents

List of Figures	VII
1 Introduction	3
1.1 Geodesic and Spline Interpolation in Metamorphosis Model	3
1.2 Learning of Low Bending and Low Distortion Manifold Embeddings	7
2 Preliminaries	11
2.1 Basic Setup for Analysis on Euclidean Spaces	11
2.1.1 Domain	11
2.1.2 Function Spaces	12
2.1.3 Inequalities and Embeddings	13
2.2 Differential Geometry	15
2.2.1 Function Spaces on Riemannian Manifolds	19
2.3 Neural Networks	21
2.3.1 Sets of Functions as Realizations of Sets of Neural Networks	24
2.4 Γ - and Mosco-convergence	25
3 Foundations of the Metamorphosis Model	27
3.1 Riemannian Geodesics and Splines	27
3.1.1 Time Continuous Setting	27
3.1.2 Variational Time Discretization	28
3.2 Image Metamorphosis Model	30
3.2.1 Flow of Diffeomorphism	30
3.2.2 Metamorphosis in Image Space	31
3.2.3 The Flow Equation	32
3.3 iPALM Algorithm	33
4 Image Morphing in Deep Feature Space with Anisotropic Regularization	37
4.1 Time Continuous Model	38
4.2 Variational Time Discretization	42
4.3 Temporal Extension Operators	48
4.4 Convergence of Discrete Geodesic Paths	49
4.5 Fully Discrete Model	60
4.6 Numerical Optimization	62
4.7 Numerical Results	63
4.8 Extensions, Conclusion and Outlook	69
5 Spline Interpolation in Image Metamorphosis Model	73
5.1 Time Continuous Model	74
5.2 Variational Time Discretization	76
5.3 Temporal Extension Operators	81
5.4 Convergence of Discrete Splines	83
5.5 Relaxation of the Warping Constraint	91

5.6	Fully Discrete Model	93
5.7	Numerical Optimization	95
5.8	Numerical Results	96
5.9	Conclusion and Outlook	101
6	Learning Low Bending and Low Distortion Manifold Embeddings	103
6.1	A Low Bending and Low Distortion Regularization for Encoders	104
6.1.1	A Low Bending and Low Distortion Loss Functional	104
6.1.2	Monte Carlo Limit for Dense Sampling	106
6.2	Existence of Minimizer to Nonlocal Energy	109
6.3	Convergence for Vanishing Sampling Radius	113
6.4	Autoencoder Architecture and Training Procedure	116
6.5	Numerical Results	117
6.5.1	Input Manifolds	118
6.5.2	Visualization of the Embeddings	119
6.5.3	Linear Interpolation in Latent Space	123
6.6	Extensions, Conclusion and Outlook	125
	Bibliography	129

List of Figures

1.1	Time discrete geodesic sequences of animal photos for the simple RGB model (first row) and deep feature model (third row), along with a zoom of the mouth region with magnification factor 4 (second/fourth row). Note that the novel deep feature-based model has significantly fewer blending artifacts, as indicated by the arrows.	4
1.2	Left: Time discrete spline (first row) and time discrete piecewise geodesic (second row) with framed key frame images: Gaussian distributions with different center and mass. Right: Euclidean splines for the input parameters versus splines for image metamorphosis extracted from the numerical results in postprocessing. The plot is in (x, y, m) -coordinates, with (x, y) denoting the center of mass and m the mass of the distribution.	5
1.3	Results for the sundial dataset. From left to right: Generating the data by casting a shadow of a vertical rod from all possible positions of a sun on a hemisphere. Training data: image pairs with their geodesic average and distance, induced by the geometry of the hemisphere. The latent manifold, which is a subset of \mathbb{R}^{16} , projected into \mathbb{R}^3 via PCA, showing that the embedded manifold clearly reflects the geometric characteristics of the hidden manifold: the hemisphere. Decoder outputs for the orange points in latent space.	8
4.1	An example of input image and the typical look of the corresponding feature channels on different levels (first to the fifth, from left to right, <i>cf.</i> Table 4.1) of multilevel semantic decomposition (plotted with the inverted values for a better visibility).	38
4.2	The visual representation of the anisotropy weight for the same input image as in Figure 4.1 and values of the parameters $\rho = 1000, \sigma_1 = 0.5, \sigma_2 = 2, \xi = 10^{-6}$. We observe that the lowest values (corresponding to the almost black grayscale intensity) are indeed obtained in the proximity of edges/discontinuities.	39
4.3	Time discrete geodesic sequences of self-portraits by van Gogh for the RGB model (first row) and deep feature model (fifth row), along with a zoom of the ear region with magnification factor 4 (second/sixth row), the associated sequences of anisotropy weights (third/seventh row), and color-coded displacement fields $\phi_k - \mathbb{1}$ (fourth/eighth row). Note that the intensity-based approach leads to blending artifacts indicated by the arrows, which are resolved in the deep feature-based model.	65
4.4	Time discrete geodesic sequences of animal photos for the RGB model (first row) and deep feature model (fifth row), along with a zoom of the mouth region with magnification factor 4 (second/sixth row), the associated sequences of anisotropy weights (third/seventh row), and color-coded displacement fields $\phi_k - \mathbb{1}$ (fourth/eighth row). Note that the novel deep feature-based model has significantly fewer blending artifacts as indicated by the arrows.	66
4.5	Pairs of time discrete geodesic paths using the deep feature model and corresponding color-coded displacement fields for paintings of US presidents (first/second row) as well as for paintings of Catherine the Great (third/forth row).	67
4.6	Visualization of the anisotropy in RGB model for a significant smaller value $\rho = 200$ compared to Figure 4.3: anisotropy weight (left) and color-coded displacement field (right) for $k = 12$	67
4.7	The part of the face region for the example with the van Gogh self-portraits and variation of the parameter δ for RGB model (first to third row) and deep feature model (fourth to sixth row).	68
4.8	Representation of the velocity field $\tilde{\mathbf{v}}$ and the corresponding deformation ζ_B of Ω_{MN} [Gua22, Figure 2, Example 4].	70

4.9	Numerical results for input from Figure 4.8 with $K = 8$ and $\delta = 10^{-3}$: geodesic path ζ_k (first row), its zoom version on the central right part of Ω_{MN} (second row), deformations ϕ_k (third row), and forward flow ψ_k (fourth row) [Gua22, Figure 9].	71
4.10	Left: Values of $\ \mathbb{1} - \zeta_k \circ \psi_k\ _{L^2(\Omega_{MN})}^2$ for $k = 1, \dots, 8$ and $\delta = 1, 10^{-1}, 10^{-2}, 10^{-3}$ [Gua22, Figure 12 (d)]. Right: Deformation of Ω_{MN} under $\zeta_B \circ \psi_8$ for $\delta = 1, 10^{-1}, 10^{-2}, 10^{-3}$, from left to right [Gua22, Figure 12, column 4].	71
5.1	Left: Schematic drawing of the Hermite interpolation $y_k^K(x)$ (blue) on the time interval $[(k - \frac{1}{2})/K, (k + \frac{1}{2})/K]$ together with the discrete acceleration $a_k^K(x)$ (red). Right: Image extension $\mathcal{U}^K[\mathbf{u}^K, \Phi^K](\cdot, x)$ along a path $(y_t^K(x))_{t \in [0,1]}$, plotted against time. Dots represent the values u_k^K , $k = 0, \dots, K$, and crosses the “half-way” values $\frac{1}{2}(u_k^K + u_{k-1}^K)$, $k = 1, \dots, K$, along the discrete transport path.	82
5.2	Right: B -spline basis b^i with segments colored in different colors. Left: Interpolation function s for $M = 16$	94
5.3	Left: Time discrete spline with framed key frame images (first row), color-coded discrete displacement/velocity field (second row), discrete second order material derivative (third row) and color-coded discrete acceleration field (fourth row), for the Gaussians example and values of the parameters $\delta = 5 \cdot 10^{-3}$, $\sigma = 1$, $\theta = 5 \cdot 10^{-5}$. The colors and their intensities indicate the direction and the intensity of the field, as indicated by the color wheel on the left. Right: Euclidean spline and piecewise geodesic interpolation for the input parameters versus spline and piecewise geodesic interpolation in metamorphosis model, extracted from the numerical results in postprocessing. The plots are given in (x, y, m) -coordinates, with (x, y) denoting the center of mass and m the mass of the distribution.	97
5.4	Left: Time discrete spline (top row) and piecewise geodesic (middle row) interpolation with framed key frames. The bottom row shows the difference in intensity between the different interpolations, using the color map -0.35  0.35 . Right: Width of the interpolated shape measured at the horizontal axis of symmetry (in number of pixels) for a spline interpolation (orange) and piecewise geodesic interpolation (green), showing the concavities in the spline interpolation ($\delta = 5 \cdot 10^{-3}$, $\sigma = 1$, $\theta = 5 \cdot 10^{-4}$).	97
5.5	Time discrete spline with framed fixed images (first and second row), first order material derivative variable \bar{z} (third and fourth row), second order material derivative with energies comparison (fifth and sixth row), and color-coded acceleration field with energies comparison (seventh and eighth row), for the values of the parameters $\delta = 2 \cdot 10^{-2}$, $\sigma = 2$, $\theta = 8 \cdot 10^{-4}$. The graphics on the right, in rows four and six, show the spline (orange) time plots of the L^2 -norm of the actual second order material derivative $\hat{\mathbf{w}}_k$ and the dissipation energy density $\ \mathbb{W}_A(\nabla_{MN} \mathbf{a}_k)\ _{L^1_{MN}}$, reflecting the motion acceleration, respectively. This is compared to the corresponding piecewise geodesic interpolation (green) (not visualized here, cf. Figure 5.8).	98
5.6	Left: Time discrete spline with framed fixed images (top and bottom row). Right: Energy density norm of acceleration flow $\ \mathbb{W}_A(\nabla_{MN} \mathbf{a}_k)\ _{L^1_{MN}}$ (top), and L^2 -norm of the actual second order material derivative $\hat{\mathbf{w}}_k$ (bottom). Parameter values: $\delta = 10^{-3}$, $\sigma = 2$, $\theta = 2 \cdot 10^{-5}$	99
5.7	Top and bottom rows: Two time discrete splines with key-frames images differing in shape and color. Parameter values: $\delta = 8 \cdot 10^{-3}$, $\sigma = 2.5$, $\theta = 2 \cdot 10^{-4}$. For visualization purposes, only even-numbered frames \mathbf{u}_k are depicted.	99
5.8	Top row: Image \mathbf{u}_{11} for the human face example, second order material derivative $\hat{\mathbf{w}}_k$ and acceleration field \mathbf{a}_k for $k = 8$, for the time discrete piecewise geodesic (left image of each panel pair) and spline (right image of each panel). The pairs of material derivatives and the acceleration fields are jointly scaled to reflect the differences in intensities. Bottom row: Same visualization with image \mathbf{u}_4 from the letter example.	100

5.9	First row: The original frames extracted from the video showing white blood cell (neutrophil) chasing a Staphylococcus aureus bacterium. Second row: Time discrete spline with framed fixed images. Third row: Time discrete piecewise geodesic with framed fixed images. Fourth and fifth row: Fully discrete second order material derivative for the discrete spline and the the discrete piecewise geodesic interpolation, respectively. The values of parameters are $\delta = 4 \cdot 10^{-2}$, $\sigma = 2.5$, $\theta = 1.6 \cdot 10^{-4}$	100
6.1	A sketch of the cone condition (M2). The parametrization $\Pi_x = \exp_x \circ \iota_x$ maps the cone $\mathcal{C}_{r_0, \kappa} \subset U_x \subset \mathbb{R}^m$ into \mathcal{M} and the neighborhoods $U_x^\epsilon \subset U_x$ onto $D_\epsilon^{\mathcal{M}}(x)$	105
6.2	Results for dataset (S) , for $\lambda = 0$ and $\epsilon = \pi/8$ (with maximal distance being π). From left to right: a sketch of the sundial configuration, training data (image pairs with their geodesic average and distance), the latent manifold $\phi(\mathcal{M})$ projected into \mathbb{R}^3 via PCA, decoder outputs for the orange points in latent space. The encoder was trained separately from the decoder and the training was stopped after the value of the functional $E^{S_\epsilon}[\phi]$ evaluated on a test set did not decrease for 10 epochs. Pairs with distance below 0.01π were rejected. The decoder was trained until an accuracy of 10^{-5} was reached on the same test set.	119
6.3	Latent manifold $\phi(\mathcal{M})$ for dataset (S) for $\lambda = 5$ (colored points as in Figure 6.2). The other parameters and the training procedure are the same as for Figure 6.2.	120
6.4	Visualization of the results for dataset (R) . We used $\epsilon = \frac{\pi}{8}$ (with maximal distance being $\frac{\pi}{2}$). Encoder and decoder were trained at the same time with $\kappa = 0.1$, and training was stopped once a reconstruction error R on a test set reached a threshold of $3.5 \cdot 10^{-3}$. Pairs with distance below $\frac{1}{20}$ of the maximal distance were rejected.	121
6.5	Visualization of results for dataset (A) . We used $\epsilon = \frac{1}{4\sqrt{2}}$ (with maximal distance being $\frac{1}{\sqrt{2}}$). Encoder and decoder were trained at the same time with $\kappa = 0.1$, and training was stopped once a reconstruction error R on a test set reached a threshold of $2 \cdot 10^{-3}$. Pairs with distance below $\frac{1}{20}$ of the maximal distance were rejected.	122
6.6	Visualization of the results for dataset (E) for $\lambda = 1$ and $\epsilon = \frac{\pi}{2}$ for different ways of rendering ellipses: as the characteristic function (left) and as its smoothed version (6.28) with $k = 3$ (right). In each case, two input triples and their distances (with maximal distance ≈ 4.3) are shown. In both cases, there are 4 relevant directions. PCA coordinates 1, 3, 4 and 2, 3, 4, respectively, are used to visualize the latent manifolds. The choice is taken for consistency of visualization, where the change of the relevant dimension is due to the symmetry of the center position parameters c in the dataset. On the bottom, example interpolations are visualized. Their corresponding interpolation path in latent space is shown in the PCA plot in the color of the frame. The encoder and decoder were trained at the same time with $\kappa = 0.1$, and training was stopped after 500 epochs.	123
6.7	Comparison of the interpolation error, for different options for the regularization, for datasets (R) (left) and (A) (right). The encoder was trained separately from the decoder for 500 epochs, and the weights yielding the best value for the functional $E^{S_\epsilon}[\phi]$ evaluated on a test set were chosen. Pairs with distance below $\frac{1}{20}$ of the respective maximal distance were rejected. The decoder was trained until an accuracy of $3.5 \cdot 10^{-3}$ for (R) and $2 \cdot 10^{-3}$ for (A) was reached.	124
6.8	Average interpolation error for training with different sampling radii ϵ , computed across a test set consisting of random pairs with arbitrary distance on the upper hemisphere. The encoder was trained separately from the decoder, and the training was stopped after the value of the functional $E^{S_\epsilon}[\phi]$, evaluated on a different test set, did not decrease for 10 epochs. Pairs with the distance below 0.01π were rejected. The decoder was trained until an accuracy of 10^{-5} was reached on the test set. This was repeated for three different initializations of the network weights, obtained as outputs of training the encoder.	125

6.9 The visualization of the results of the embedding of Gaussian ellipses and geodesic interpolation on the obtained latent manifold. We trained the encoder and the decoder jointly on around 10000 images with $\kappa = 0.1$. The bending parameter is $\lambda = 1$, while the standard deviation parameter σ is $\frac{1}{10}$ of the maximal distance of the training points. Visualization is obtained by the projection to the first 3 PCA dimensions (out of 4 relevant). The results clearly show the advantages of the geodesic interpolation on the manifold in comparison with the simple linear interpolation in the latent space. In particular, the shape of the ellipse is preserved for the geodesic interpolation, which is not the case for the linear interpolation. 126

Acknowledgements

FIRST and foremost I would like to express my deep gratitude to my advisor Prof. Dr. Martin Rumpf, who introduced me to many interesting topics in mathematics and provided excellent guidance and constant support throughout my doctoral studies.

I highly appreciate the willingness of Prof. Dr. Benedikt Wirth for co-reviewing this thesis, and Prof. Dr. Juan Jose Lopez Velázquez and Prof. Dr. Heiko Röglin for being part of the doctoral committee.

I would like to thank my coauthors: Juliane Braunsman, Prof. Dr. Alexander Effland, Jorge Justiniano, Dr. Erich Kobler, Prof. Dr. Thomas Pock, Prof. Dr. Martin Rumpf, and Prof. Dr. Benedikt Wirth for interesting collaborations and many useful discussions.

I am also grateful to Prof. Dr. Thomas Pock and his group for hosting me during my research visit to Graz.

Special thanks go to Dr. Stefan Simon, Jorge Justiniano, Josua Sassen, and Christoph Smoch for proofreading parts of my thesis. Besides, I would like to thank other current and former colleagues and personnel at the Institute for Numerical Simulation, University of Bonn: Dr. Martin Lenz, Dr. Behrend Heeren, Florine Hartwig, Mara Guastini, Franziska Henrich, Karsten Rezny, Carole Rossignol, and Gabi Sodoge-Stork for many useful discussions, but also for providing assistance, guidance, and an always pleasant atmosphere.

I am very grateful to all of my teachers and everyone involved in my education to this day.

Last but not least, I want to thank my family for their honest love and everlasting support. They are the main reason for all of my achievements. In particular, I would like to thank my beloved wife Ivana for her patience, encouragement, and motivation.

Funding I gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) via project 211504053 - Collaborative Research Center 1060 and Germany's Excellence Strategy project 390685813 - Hausdorff Center for Mathematics. Furthermore, both during my Master's and Doctoral studies, I was supported by the Bonn International Graduate School of Mathematics. Besides their financial support, I am grateful to Karen Bingel from the BIGS administration, for her willingness to help in any situation, especially in my first days in Bonn. During my Bachelor's studies, I was financially supported by the Scholarship Fund of University of Primorska and "Dr Milan Jelić" Fund, to whom I am also very grateful.

Chapter 1

Introduction

THE contribution of this thesis consists of studying two larger mathematical topics. The first topic is the geodesic and spline interpolation of images in the framework of the metamorphosis model. The second topic concerns learning low bending and low distortion manifold embeddings. In this chapter, we give a basic introduction to both of these topics. In Chapter 2, we introduce basic mathematical notions and notation used throughout the thesis. Chapter 3 introduces the metamorphosis model and tools necessary for the study of geodesic interpolation in deep feature space with anisotropy regularization which is given in Chapter 4 and spline interpolation studied in Chapter 5. Finally, Chapter 6 introduces a theoretical framework and presents numerical results for the learning of low bending and low distortion manifold embeddings.

1.1 Geodesic and Spline Interpolation in Metamorphosis Model

In mathematical and computer vision, morphing is the problem of computing a visually appealing transition of two shapes such that semantically corresponding regions are mapped onto each other. For image morphing, one can think of a transition that gives a smooth video-like animation with fixed initial and final frames. To define this notion in mathematical terms, the energy cost of such transitions is defined such that the desired morphing transition path will have the smallest possible cost. The problems of searching for the shortest path connecting two fixed points as a minimizer of a certain path energy is the basis of *geodesic interpolation* as a generalization of the notion of the straight line to suitably defined Riemannian manifolds. Finding an appropriate manifold structure and the underlying path energy in the context of image morphing is a well-studied problem. Considering that we are interested in a fluid-like transition which properly morphs shapes, it is not a surprise that the motivation for one of the most well-known approaches comes from fluid dynamics and the study of diffeomorphisms. This model, called the *flow of diffeomorphism model* [CRM96, DGM98], is based on Arnold's approach [Arn66] to study flows of ideal fluids by a geodesic formulation in the space of measure-preserving diffeomorphisms. In terms of image morphing, each point of the reference image is optimally transported to the target image with respect to the energy measuring the total dissipation of the underlying diffeomorphic flow. The key assumption of this model is the *brightness constancy*, *i.e.*, the image intensity is preserved along the trajectories of the pixels. The *metamorphosis model*, originally introduced by Miller, Trounev, and Younes [MY01, TY05b, TY05a], builds upon the flow of diffeomorphism model by additionally allowing for image intensity modulations along the trajectories. This is achieved by introducing a metric on the manifold of squared-integrable images that, in addition to the flow of diffeomorphism metric, includes the magnitude of these modulations, reflected by the integrated squared material derivative of image trajectories as a penalization term. Starting from the general framework for a variational time discretization of geodesic calculus [RW15], a variational time discretization of the metamorphosis model for square-integrable images was proposed in [BER15]. Moreover, the existence of time discrete geodesic paths as well as the Mosco-convergence [Mos69] of the time discrete to the time continuous metamorphosis model was proven. For numerical computations, they applied an alternating algorithm to the fully discrete model based on (non)conforming finite element discretizations.

Metamorphosis model in deep feature space with anisotropic regularization

The above-mentioned models study the morphing of images in their grayscale or color-based representation. This representation, however, does not carry any semantic information, which reflects on the quality of the image morphing paths. In particular, this frequently leads to an unnatural morphing along the obtained geodesic paths typified by a blending instead of flow-induced geometric transformations. Furthermore, the standard metamorphosis model assumes the strong smoothness of the underlying flow fields, leading to improper behavior of shape interfaces along the geodesic paths. This resulted in a need to use a fourth color channel representing a manual segmentation of image regions and a color adaptation for some examples in [BER15]. To address these problems, in Chapter 4, based on publications [EKPR19, EKP⁺21] where the author of this thesis contributed to theoretical and experimental studies in the latter, we advocate a metamorphosis model in a deep feature space, which amounts to replacing the input images by feature vectors combining image intensities and semantic information generated by a feature extraction operator represented by a deep convolutional neural network [SZ14]. To explicitly allow for approximate discontinuities in the deformation fields, we further incorporate an anisotropic flow regularization based on anisotropic edge diffusion [PM90]. We follow the analogous variational time discretization as in [BER15] and prove the existence of discrete geodesic paths and their convergence to the continuous counterpart. The spatial discretization is based on a finite difference approximation of derivatives of deformations and Hermite spline approximation of the warping operation. For numerical optimization, an inertial proximal alternating linearized minimization approach of the iPALM algorithm [PS16] is used. Based on the algorithm implementation presented in [EKPR19], we provide numerical experiments comparing the results of the original intensity-based approaches with the results after the incorporation of semantic deep features, showing the superiority of the latter.



Figure 1.1: Time discrete geodesic sequences of animal photos for the simple RGB model (first row) and deep feature model (third row), along with a zoom of the mouth region with magnification factor 4 (second/fourth row). Note that the novel deep feature-based model has significantly fewer blending artifacts, as indicated by the arrows.

Splines for image metamorphosis

The natural question following image morphing between two fixed images is the task of finding a smooth interpolation between several key frame images. The piecewise geodesic interpolation represents an attempt to solve the task, but the resulting interpolant lacks smoothness at the interpolation data points, leading to jerking at the corresponding times when animating the resulting family of images. This is the general observation for piecewise geodesic interpolation of several fixed points also on other manifolds and shape spaces. A widespread and versatile solution for this problem is a (*cubic*) *spline interpolation*. On Euclidean spaces, in contrast to the geodesic interpolation as minimization of the squared velocity, cubic splines are classically defined by de Boor [dB63] as minimizers of the integral of the squared acceleration. In a Riemannian manifold context, Noakes *et al.* [NHP89] introduced Riemannian cubic splines as stationary paths of the integrated squared covariant derivative of the velocity. This motivates the splines for image metamorphosis model defined in Chapter 5 (based on publications [JRR21, JRR23] to which author of the thesis contributed), as minimizers of a functional which is a sum of quadratic functionals of the Eulerian motion acceleration and the second material derivative representing acceleration in the change of intensities along motion paths. This physically intuitive separation of flow and image intensity acceleration, however, does not lead to a fully Riemannian model, since the integral over the squared covariant derivative of the path velocity in the Riemannian metric would lead to an interwoven model of these different types of acceleration. For the proposed model, we define a variational time discretization and a space discretization naturally following already introduced approaches in the geodesic interpolation context. Furthermore, we prove the convergence of time discrete splines to their continuous counterpart and provide numerical examples showing the smoothness and stability of spline interpolation in comparison to the piecewise geodesic interpolation.




Figure 1.2: Left: Time discrete spline (first row) and time discrete piecewise geodesic (second row) with framed key frame images: Gaussian distributions with different center and mass. Right: Euclidean splines for the input parameters versus splines for image metamorphosis extracted from the numerical results in postprocessing. The plot is in (x, y, m) -coordinates, with (x, y) denoting the center of mass and m the mass of the distribution.

Related work

We present, a far from exhaustive, recapitulation of works related to the flow of diffeomorphism and metamorphosis models. For a more detailed exposition of some of these works, we refer the reader to [You10, MTY15] and the references therein. We also mention some works in the direction of spline interpolation.

Flow of diffeomorphism The study of image morphing through the flow of diffeomorphism model is closely connected to *diffeomorphic image registration* where the goal is to find a reasonable diffeomorphic deformation of a given template image so that it matches a given target image as closely as possible. Modeling this deformation as a diffeomorphic flow that arises as the solution of the *flow equation*, with an appropriate choice for the space of velocity fields, leads to the large deformation diffeomorphic metric mapping (LDDMM) framework [CRM96]. These models were initiated by Trouné [Tro95, Tro98] who constructed a distance in the space of deformations by exploiting Lie group methods, where deformations are viewed as actions of an infinitesimal transformation group on images. Since then, they were intensively studied from many different perspectives and applied in numerous

contexts. The variational problem associated with the LDDMM framework can be formulated as a constrained optimization problem, leading to the study of the problem through the optimal control methods. In this context, the flow equation represents the state equation, the velocity field the control, and the transformed image the resulting state [HZN09, VRRC12, VRRH12]. Efficient gradient descent algorithms were proposed for Euler–Lagrange equations for minimizing deformation fields [BMTY05] and for the path energy of the flow of diffeomorphism model by using Jacobi fields [You07]. Geodesic shooting in the LDDMM framework, initially studied by Miller *et al.* [MTY06] from the perspective of the conservation of the initial momentum in Lagrangian coordinates, led to some efficient numerical approaches [VRRC12, SHJF13, ZF19]. Unlike these methods, which first try to understand the optimization of the continuous problem and then build the corresponding discretization approach, there are numerical methods taking the discretize-then-optimize perspective [Mod09, MR17, PNVM20]. More recent approaches taking this perspective are based on (deep) neural network parametrization of the velocity fields or diffeomorphic flows. Since suitable training data of velocity fields/flows is rarely available, the unsupervised methods, which take pairs of template and target images as an input, are more popular. These methods include [YKSN17] based on learning of initial momentum and geodesic shooting, [DBGS19] based on learning stationary velocity fields and approximate geodesic shooting in the sense of [ACPA06], and [AAS22, WJW⁺22] which parametrize time-dependent velocity fields as deep neural networks with residual connections. The flow of diffeomorphism model and the LDDMM framework were applied in numerous science and engineering applications, but most successfully and frequently in medicine. One of the earliest applications of the flow of diffeomorphism model to medical data was proposed in [BMTY02], where the diffeomorphisms represent deformations of anatomic reference structures. Some more recent applications include joint image reconstruction and motion estimation in four-dimensional computed tomography (CT) [HSW⁺12, CGÖ19], indirect image registration in two-dimensional CT [CÖ18], template-based image reconstruction from sparse tomographic data [LNÖS20], molecular digital pathology with tissue magnetic resonance imaging (MRI) [SWTM22], and emotion conversion [SHCV22].

Metamorphosis model The origins of the image metamorphosis model, as a generalization of the flow of diffeomorphism model in which images are allowed to vary in time, can be traced to Miller and Younes [MY01] where a left-invariant distance was defined on the product space of a transformation group and the family of images. Trouvé and Younes [TY05b] rigorously introduced the metamorphosis metric in the space of images and analyzed the geodesic equation for this metric. In [TY05a], they proved the existence of geodesic curves for square-integrable images and the (local) existence and uniqueness of solutions to the initial value problem for the geodesic equation in the case of weakly differentiable images. Holm *et al.* [HTY09] studied the geodesic equation in a Lagrangian formulation of the model and proved the existence of a solution for both the boundary value and the initial value problem in the case of measure-valued images. This approach was used by [RY16] for a robust shooting method for images contained in a reproducing kernel Hilbert space, and [FGG21], where a semi-Lagrangian scheme with geodesic shooting was introduced. In [MFG⁺22], this geodesic equation was solved by a deep residual network with further use of a segmentation mask as a local regularizer, with an application to the alignment of images of brain tumors. A further application of deep learning perspective to the metamorphosis model can be found in [BVCD20], where a variational autoencoder framework was developed for joint shape and appearance representations. This is achieved by observing training images as metamorphic transformations of a prototype image and imposing norm equivalence between the metamorphosis norm and the latent space norm, while the velocity field is taken to be static and representable by a discrete reproducing kernel Hilbert space. Hong *et al.* [HJS⁺12] proposed a metamorphosis regression model and developed a shooting method to reliably recover initial momenta. Richardson and Younes [RY13] extended the metamorphosis framework to discrete measures and provided a rigorous analysis of singularities that can be generated along the geodesic path. Charon *et al.* [CCT18] generalized the metamorphosis model to functional shapes representing scalar-valued signals such as curves or finite-dimensional submanifolds. In [HC22], a metamorphosis model in the space of geometric measures (varifolds) is introduced by allowing dynamical change of measures weights which is penalized by the Fisher–Rao metric. [BC20] studies the Eulerian formulation of the model in sense of a constrained minimization problem leading to a geodesic equation corresponding to a system of advection and continuity equations that can be solved analytically. In [BAC19], a metamorphosis setting was introduced in the space of landmarks in which non-diffeomorphic evolution is locally allowed, but penalized. Neumayer *et al.* [NPS18] extended the time discrete model introduced by [BER15] to Hadamard manifold-valued images, while the proof of convergence to the time continuous model is given in [ENR20]. Based on the time discrete metamorphosis model developed by Berkels *et al.* [BER15], Efland *et al.*

introduced the Bézier curves interpolation in the space of images [ERS⁺15] and discrete extrapolation [ERS18]. Further applications of the image metamorphosis model include image geodesics for optical coherence tomography [BBE⁺17], exemplar-based face colorization [PPS17], sparse and limited angle computerized tomography and super-resolution of images [NPS19], joint (spatio-temporal) tomographic reconstruction and registration [GCÖ20], template based image reconstruction suitable for sparse data [NT21], and brain tumor registration with restriction of image intensity change to certain areas [FMO⁺22].

Spline interpolation As already mentioned, in Euclidean space cubic splines are defined as minimizers of the integral of the squared acceleration [dB63], while in a Riemannian context cubic splines are defined as stationary paths of the integrated squared covariant derivative of the velocity [NHP89]. Today, there is a variety of spline approaches in nonlinear spaces and with applications to shape spaces. Trouvé and Vialard [TV12] studied a second-order shape functional in the space of landmarks based on an optimal control approach. Singh *et al.* [SVN15] introduced an optimal control method involving a functional measuring the motion acceleration in a flow of diffeomorphisms ansatz for image regression. Tahraoui and Vialard [TV19] considered a second-order variational model on the group of diffeomorphisms of the interval $[0, 1]$. They proposed a relaxed model leading to a Fisher–Rao functional, as a convex functional on the space of measures. Vialard [Via20] showed the existence of a minimizer of the Riemannian acceleration energy on the group of diffeomorphisms endowed with a right-invariant Sobolev metric of high order. Benamou *et al.* [BGV19] and Chen *et al.* [CCG18] discussed spline interpolation in the space of probability measures endowed with the Wasserstein metric. Thereby, energy splines are defined as minimizers of the action functional on Wasserstein space which involves the acceleration of measure-valued paths, sharing similarities with the spline functional in the space of images that we present in this thesis. The initial computational intractability of such approaches is tackled by a relaxation based on multi-marginal optimal transport and entropic regularization. The transport problem this approach aims at solving might not have a Monge solution, which was solved by Chewi *et al.* [CCG⁺21] by constructing measure-valued splines, dubbed transport splines. The existence of Riemannian splines defined in the sense of [NHP89] was shown in [HRW18]. This paper also introduces a variational time discretization. The existence of time discrete splines was shown together with their convergence towards the continuous spline which follows from Γ –convergence of time discrete spline energies towards their continuous counterparts. This method was applied for spline interpolation on embedded finite-dimensional manifolds, on a high-dimensional manifold of discrete shells with applications in surface processing (see also [HRS⁺16]), and the infinite-dimensional shape manifold of viscous rods. Following this approach, a time discrete approximation of splines in Wasserstein space is studied in [JRE23]. In this case, the time discretization is based on Wasserstein-2 distances and the generalized Wasserstein barycenter. The existence of time discrete splines is investigated and consistency to the time continuous model is shown for Gaussian distributions.

1.2 Learning of Low Bending and Low Distortion Manifold Embeddings

Processing of high-dimensional data in their original form, with underlying very high- or even infinite-dimensional manifold, example of which are the above-mentioned models, is a challenging task, both in theoretical and numerical computation senses. To this end, in everyday applications, it is often observed that apparently high-dimensional data actually forms a low-dimensional manifold. This observation, called the *manifold hypothesis*, is the basis for one of the central machine learning tasks: identify given data sets consisting of objects embedded in a high-dimensional space with points in a (suitably generated) low-dimensional *latent manifold*. Such a representation is frequently defined via an *encoder map* of an *autoencoder* acting on input objects and mapping into a low-dimensional Euclidean *latent space*. The latent manifold is then the image of this encoder map. The associated *decoder* maps back points from latent space to points in the input space, *i.e.*, to the so-called *hidden manifold* (due to its a priori unknown topology and geometry). Encoder and decoder are trained as deep neural networks (giving rise to the name *deep manifold learning*) via the minimization of a loss functional called *reconstruction loss* that compares the input data with its image under the composition of the encoder and decoder mapping. To make the processing of the data by processing the *latent codes* meaningful, the latent manifold geometry should be as close as possible to the geometry of the hidden manifold. However, measuring only the reconstruction loss is not enough to accurately recover the geometry present in the data. In particular, a major deficit of autoencoders is that they frequently fail to reproduce the statistical distribution of the input data to the latent space. To resolve this

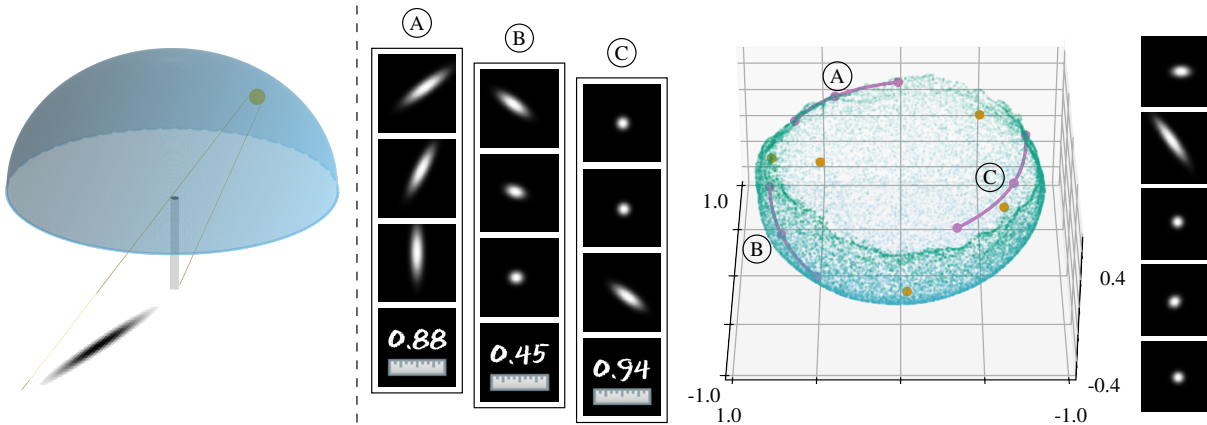


Figure 1.3: Results for the sundial dataset. From left to right: Generating the data by casting a shadow of a vertical rod from all possible positions of a sun on a hemisphere. Training data: image pairs with their geodesic average and distance, induced by the geometry of the hemisphere. The latent manifold, which is a subset of \mathbb{R}^{16} , projected into \mathbb{R}^3 via PCA, showing that the embedded manifold clearly reflects the geometric characteristics of the hidden manifold: the hemisphere. Decoder outputs for the orange points in latent space.

issue, among others, approaches promoting *low distortion* or *isometry i.e.*, length-preserving of encoder maps, are proposed. Furthermore, embeddings which lead to *flatness* of the latent manifold are also favorable, since in that case, the simplicity of the geometry of Euclidean spaces would lead to simplifying the downstream tasks such as interpolation or classification.

In Chapter 6, based on publications [BRRW21, BRRW22] to which the author of the thesis contributed, we study embeddings from high-dimensional input manifolds of images into low-dimensional latent spaces. Taking the previous observations as a motivation, we propose a loss functional consisting of a penalty term measuring the distortion of the distance, thus promoting approximately isometric latent space embeddings under the encoder mapping, and a penalty term penalizing large bending, thus promoting flatness. Our input data consists of Monte Carlo sampled point pairs, together with their Riemannian distance and Riemannian average on the input manifold, where we assume that for each such pair, this additional information is well-defined and easily computed. This data is used for obtaining the first and second order difference quotients which approximate directional components of the Riemannian gradient (Jacobian) and Hessian, respectively. The discrete loss functional depends on a parameter ϵ defined as the maximal distance of each point pair sampled on the input manifold. For an increasing number of samples, we obtain a Monte Carlo limit of the initial sampling loss as a nonlocal loss functional represented by a double integral over the pairs of input manifold points with an additional density depending on the sampling strategy. We prove the existence of a minimizer for the nonlocal loss functional in the class of functions representable by deep neural networks, where we suppose certain bounds on the architecture of the networks as a function of ϵ . The main theorem identifies the Mosco-limit of these nonlocal Monte Carlo limits for sampling distance $\epsilon \rightarrow 0$ as a local geometric loss functional of embedding maps from the input manifold to the latent space. This loss functional consists of a distortion and a bending energy part as proper limits of their discrete counterparts. As a corollary, we prove the convergence of the minimizers of the nonlocal energies to their continuous counterpart. In the numerical experiments, we use digital image data so that the dimension of our input space equals the number of pixels times the number of colors. These images encode different explicitly given low-dimensional data manifolds, corresponding to the hidden manifold. We visualize the obtained latent manifolds (point clouds in latent space) by projecting it onto directions resulting from a principal component analysis (PCA) [Jol02]. Our numerical experiments confirm that the bending loss significantly increases *smoothness* of the resulting latent space manifold compared to a pure isometry loss, which would for instance also allow irregular Nash–Kuiper embeddings [Kui55b, Kui55a, Nas54]. Furthermore, as one possible application of our approach, penalization of the bending results in the fact that the decoder maps of the linear interpolations in the latent space represent reasonable interpolations on the data manifold.

Related work

We first recall some classical approaches to the task of identification of high-dimensional objects with low-dimensional latent manifolds. Linear approaches to this task include *principal component analysis (PCA)* [Jol02] and *singular value decomposition (SVD)* [MF11], where the data is projected onto spaces determined by the eigenvectors and singular vectors, respectively, of the corresponding matrices built from the data. Classical nonlinear approaches are based on neighborhood graphs of high-dimensional input data points. The output consisting of low-dimensional vectors is computed by preserving pairwise geodesic distances in case of *Isomap* [TDSL00] or local linearity neighborhoods in *Local Linear Embedding* [RS00]. Further approaches are so-called *spectral embeddings* where the construction of neighborhood graphs is based on spectral decomposition and eigenvalues of Laplace–Beltrami operator (*Laplacian Eigenmaps* [BN01]), Hessian operator (*Hessian Eigenmaps* [DG03]) and heat kernels (*Diffusion Maps* [CL06]). All of these models give fixed latent coordinates for the input data and thus are not suitable for application on previously unseen data samples.

As mentioned above, different strategies to favor the smoothness of the encoder and decoder mappings and thus the regularity of the latent manifold have been investigated. Some of the well-known and widely used methods are: promoting *sparsity* of latent representation [RPC⁺07], *contractive autoencoders* [RVM⁺11] which use loss functionals penalizing the norm of the Jacobian of the encoder mapping, *denoising autoencoders* [VLBM08] which make the learned latent representations robust to partial noise-corruption of the input pattern, and *variational autoencoders* [KW13] which regularize the latent representation to match a tractable probability distribution.

There are also related approaches that use loss functionals to promote isometric and/or flat embeddings. In [ST19], a loss functional measures the difference between distances in the pushforward manifold metric and distances in latent space to obtain an approximately isometric encoder map. The loss functional from [PTBK19] compares Euclidean distances in latent space with geodesic distances on the input manifold. In [KZSN20], *isometric encoder* maps are used to more accurately pushforward distributions from input to latent space using a loss function based on Shannon-Rate-Distortion. Extending these methods, [NKS21] shows that variational autoencoders can be mapped to an implicit isometric embedding with a scale factor derived from the posterior parameter. The loss functional from [AGL20] promotes isometry of the decoder map, by penalizing deviation from a non-orthogonal Jacobian matrix, and that the encoder is a pseudo-inverse of the decoder, by enforcing the Jacobians of decoder and encoder to be transposes of each other. This structure was used in [CRM⁺22] for learning injective normalizing flows since it allows for splitting of manifold learning and density estimation. In [PLD⁺20], isometric embeddings in latent space are learned to obtain standardized data coordinates from scientific measurements. To this end, the Jacobian is approximated via normally distributed sampling around each data point (so-called bursts) and the deviation of the local covariance of bursts from the identity is used to measure the lack of orthogonality of the Jacobian. A loss functional which is based on measuring the cosine of the angle between feature vectors in different media spaces is proposed in [ZYY⁺18] for cross-modal retrieval. This is then incorporated in a denoising autoencoder and combined with a distance-preserving common space. In [YYSP21], these approaches are extended by studying autoencoders which (approximately) preserve not only distances but also angles and areas. This is achieved by a loss functional depending on eigenvalues of the pullback metric defined in terms of Jacobians of en- and decoder. In [KCM21], the spectral method of Laplacian Eigenmaps was combined with a loss which penalizes, in terms of the Lipschitz constants, the deviation of the embedding map and its inverse from isometry.

One of the reasons for penalizing the bending of embeddings is to enable meaningful interpolation between data points via an affine interpolation in the latent space. A more advanced approach to the geometry of latent manifolds includes the observation that the initial data manifolds often come with a metric encoding the cost of local variations on the manifold. For sufficiently regular data manifolds it is shown in [SKTF18] how to transfer this metric to the latent manifold turning it into a *Riemannian manifold*. This in principle allows computing shortest paths, exponential maps, and parallel transport in latent space, which the decoder can pushforward to the data manifold. Latent spaces of variational autoencoders were first studied as a Riemannian manifold in [CKK⁺18]. In [STT⁺18], a convex latent distribution is promoted by training a *generative adversarial network (GAN)* [GPAM⁺20] on latent space interpolations. In [BRRG18], a regularizer based on GANs is proposed to ensure visually realistic interpolations in latent space. To this end, the adversarial regularizer tries to make the decoding of interpolations in latent space indistinguishable from real data points. To extend this method, [QLC⁺19] treats linear interpolation in each dimension of the latent representation differently, and explicitly considers the distribution of the latent representations matching it with a prior using the maximum mean discrepancy framework. In [CKF⁺20], the flatness of the latent space for variational autoencoders was promoted by penalizing deviation of the decoder induced Cauchy–

Green deformation tensor from the scaled identity. A generalized definition of interpolation via the training of a discriminator was proposed in [OYHO20] which allows checking that the interpolated point belongs to the original dataset. In this approach, an additional smoothness loss is used based on differentiation along interpolation paths in latent space. While the method in [BRRG18] relies solely on an adversarial network to discriminate between real data and interpolations, the approach in [OYHO20] also suggests including ground truth interpolation data. Besides interpolation, isometric and flat embeddings were also used for other downstream processing tasks: classification and clustering [AGL20, BRRG18], clustering for anomaly detection [ZLC⁺21], and novelty detection [AKAD22].

Chapter 2

Preliminaries

IN this chapter, we introduce some basic notation and recall general definitions and properties of objects used throughout this thesis. In Section 2.1, we recall the basic definitions of function classes on Euclidean spaces and their properties. In Section 2.2, we collect notions and results from differential geometry, in particular, Riemannian manifolds and some properties of function spaces on compact Riemannian manifolds. In Section 2.3, we introduce the notation for the study and the application of neural networks and present some results on function spaces consisting of functions that are realizations of neural networks. Finally, Section 2.4 introduces the notions of convergence that will be the subject of study in the central theoretical results of this thesis.

2.1 Basic Setup for Analysis on Euclidean Spaces

In this section, we recall the definitions and basic properties of classical function classes on Euclidean spaces, including some fundamental inequalities, embeddings, and convergence results. For proofs of these results, we refer the reader to some of the classical books on the theory of functional analysis and partial differential equations, including [AF03, Alt06, Bre11].

2.1.1 Domain

Definition 2.1.1 (Geometric properties of domains, [AF03, Section 4.1]). Let $\Omega \subset \mathbb{R}^n$ be a bounded *domain*, i.e., any non-empty, open, connected, and bounded set with boundary $\partial\Omega$.

- (i) The domain Ω has a *Lipschitz boundary* if for each $x \in \partial\Omega$ there exists a neighborhood $U = U(x)$ of x and a Lipschitz function $f : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ (see the definition below) such that

$$\Omega \cap U(x) = \{y = (y_1, \dots, y_n) \in U(x) : y_n > f(y_1, \dots, y_{n-1})\}.$$

- (ii) Let $\sphericalangle(x, y)$ denote the angle between $x, y \in \mathbb{R}^n$. A *finite cone* with *vertex* at the origin associated with the triple $(v, \rho, \kappa) \in \mathbb{R}^n \times \mathbb{R}^+ \times (0, \frac{\pi}{2}]$ (axis direction, height and aperture angle of the cone) is the set

$$\mathcal{C}(v, \rho, \kappa) := \{x \in \mathbb{R}^n : 0 \leq |x| \leq \rho, \sphericalangle(x, v) \leq \kappa\}.$$

The cone with the same triple and vertex $x \in \mathbb{R}^n$ is given by $x + \mathcal{C}(v, \rho, \kappa)$. The domain Ω satisfies the *cone condition* if there exists a finite cone $\mathcal{C}(v, \rho, \kappa)$ such that each $x \in \Omega$ is the vertex of a cone which is contained in Ω and congruent to $\mathcal{C}(v, \rho, \kappa)$.

- (iii) The domain is *strongly Lipschitz* if Ω is a Lipschitz domain that satisfies the cone condition.

2.1.2 Function Spaces

Given a sufficiently smooth function $f : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$, we denote by $\partial_{x_i}^j$ the j -times partial derivative in direction of x_i , and for a multi-index $\beta = (\beta_1, \dots, \beta_n) \in \mathbb{N}^n$ we write

$$D^\beta f := \partial_{x_1}^{\beta_1} \cdots \partial_{x_n}^{\beta_n} f, \quad |\beta| := \sum_{i=1}^n \beta_i, \quad \text{and for } k \in \mathbb{N} : D^k f := \{D^\beta f, |\beta| = k\} \text{ with } D^0 f = f.$$

For $k \in \mathbb{N}$, we use the standard notation for *the Banach space of continuous or continuously differentiable functions* $C^k(\overline{\Omega})$ on a compact set $\overline{\Omega} \subset \mathbb{R}^n$, which is endowed with the norm $\|\cdot\|_{C^k(\overline{\Omega})}$ and the seminorm $|\cdot|_{C^k(\overline{\Omega})}$ given by

$$\|f\|_{C^0(\overline{\Omega})} := \sup_{x \in \overline{\Omega}} |f(x)|, \quad |f|_{C^k(\overline{\Omega})} := \sum_{|\beta|=k} \|D^\beta f\|_{C^0(\overline{\Omega})}, \quad \|f\|_{C^k(\overline{\Omega})} := \sum_{l=0}^k |f|_{C^l(\overline{\Omega})}.$$

Moreover, we set $C^\infty(\overline{\Omega}) := \bigcap_{k \geq 0} C^k(\overline{\Omega})$. In addition, for $k \in \mathbb{N}$ and $\alpha \in (0, 1]$ we observe *the space of Hölder continuously differentiable functions* $C^{k,\alpha}(\overline{\Omega})$. This is a Banach space when equipped with the norm $\|\cdot\|_{C^{k,\alpha}(\overline{\Omega})}$, and we denote by $|\cdot|_{C^{k,\alpha}(\overline{\Omega})}$ the associated seminorm, *i.e.*,

$$|f|_{C^{k,\alpha}(\overline{\Omega})} := \sum_{|\beta|=k} \sup \left\{ \frac{|D^\beta f(x) - D^\beta f(y)|}{|x-y|^\alpha} : x, y \in \overline{\Omega}, x \neq y \right\}, \quad \|f\|_{C^{k,\alpha}(\overline{\Omega})} := \|f\|_{C^k(\overline{\Omega})} + |f|_{C^{k,\alpha}(\overline{\Omega})}.$$

For simplicity, we set $C^{k,0}(\overline{\Omega}) = C^k(\overline{\Omega})$ for any $k \in \mathbb{N}$. In the case $k = 0, \alpha = 1$, we say that function f is a *Lipschitz function* and call $|\cdot|_{C^{0,1}}$ the *Lipschitz constant*, often denoted by L . If $\alpha > 1$ by $C^{k,\alpha}$ we mean $C^{k+\lfloor \alpha \rfloor, \alpha - \lfloor \alpha \rfloor}$. For an extensive study of functions with Hölder continuity, we refer the reader to [Fio16]. This class can be observed as an instance of a wider class of functions allowing a *modulus of continuity*. We say that a monotonically increasing function $\omega : [0, +\infty) \rightarrow [0, +\infty)$ is a modulus of continuity of a function $f : \Omega \rightarrow \mathbb{R}$ if

$$|f(x) - f(y)| \leq \omega(|x - y|), \quad \forall x, y \in \Omega; \quad \omega(0) = \lim_{h \rightarrow 0} \omega(h) = 0.$$

We denote the class of k -times differentiable functions on $\overline{\Omega}$ whose k -th derivative has the modulus of continuity ω by $C_\omega^k(\overline{\Omega})$. Finally, the subscript c always indicates compactly supported continuous or differentiable functions.

We use the standard notation for *Lebesgue spaces* with the norm defined by

$$\|f\|_{L^p(\Omega)} := \left(\int_\Omega |f|^p dx \right)^{\frac{1}{p}} \quad \text{for } 1 \leq p < \infty, \quad \|f\|_{L^\infty(\Omega)} := \inf_{N \subset \Omega: |N|=0} \sup_{x \in \Omega \setminus N} |f(x)|.$$

For *Sobolev spaces* on a domain Ω (see [AF03, Alt06, Bre11] for details), we use the notation $W^{m,p}(\Omega)$ for $m \in \mathbb{N}$ and $p \in [1, \infty]$. In particular, we set $H^m(\Omega) := W^{m,2}(\Omega)$. The associated (semi)norms read as

$$|f|_{W^{m,p}(\Omega)} := \sum_{|\beta|=m} \|D^\beta f\|_{L^p(\Omega)},$$

$$\|f\|_{W^{m,p}(\Omega)} := \left(\sum_{l=0}^m |f|_{W^{l,p}(\Omega)}^p \right)^{\frac{1}{p}}, \quad \text{for } 1 \leq p < \infty, \quad \|f\|_{W^{m,\infty}(\Omega)} := \max_{0 \leq l \leq m} \sum_{|\beta|=l} \|D^\beta f\|_{L^\infty(\Omega)}.$$

Furthermore, $W_0^{m,p}(\Omega)$ is the closure of $C_c^\infty(\Omega)$ w.r.t. the norm $\|\cdot\|_{W^{m,p}(\Omega)}$, and we set $W^{0,p}(\Omega) = L^p(\Omega)$. For clarity of arguments, we often write also the codomain in the notation, *e.g.* $L^p(\Omega, \mathbb{R}^n)$.

The above notation is naturally extended to *Bochner spaces* of the form $X(I, Y)$, where I is some interval (usually a subset of $[0, 1]$) and X, Y are some of the above-defined spaces. We will use \dot{f} to denote the time derivative of function f . For more details on Bochner spaces, we refer the reader to [HvNVW16].

Given a sequence of functions $\{f^j\}_{j \in \mathbb{N}}$, we say it *converges strongly* to a function f if $\lim_{j \rightarrow \infty} \|f - f^j\| = 0$, where $\|\cdot\|$ is one of the above-defined norms.

To define the weak convergence on a function space X , we observe its dual space X' and say that a sequence of functions $\{f^j\}_{j \in \mathbb{N}}$ *weakly converges* to f with notation $f^j \rightharpoonup f$ if and only if $\phi(f^j) \rightarrow \phi(f)$ in \mathbb{R} for every $\phi \in X'$ (see [Alt06, Chapter 8] and [Bre11, Chapter 3] for details).

A classical property of Lebesgue and Sobolev spaces that will be used numerous times in this thesis is *reflexivity*. In short words, spaces $L^p(\Omega)$, $W^{k,p}(\Omega)$ for $1 < p < \infty$ are weakly precompact, meaning that uniformly bounded sequence in those spaces have a weakly convergent subsequence in the same space (cf. [Alt06, Section 8.8] and [Bre11, Section 3.5]). The main application of *compact embedding* results (see Theorem 2.1.8 below) is to transfer this weak convergence to the strong convergence in the larger space.

The other concept that will rely on the definition of weak convergence is (*sequential*) *weak lower semicontinuity*. We say that a functional $F : X \rightarrow \mathbb{R}$ is (sequentially) weak lower semicontinuous if for every sequence $\{f^j\}_{j \in \mathbb{N}}$ with $f^j \rightharpoonup f$ it holds $\liminf_{j \rightarrow \infty} F(f^j) \geq F(f)$. In particular, strongly lower semicontinuous and convex functionals are weakly lower semicontinuous [Bre11, Proposition 3.5 and Corollary 3.9].

In the following section, we recall some basic inequalities which hold for functions in the above-defined function spaces, convergence criteria, and embeddings between function spaces.

2.1.3 Inequalities and Embeddings

The following fundamental results are of the vital importance throughout this thesis. While we use a special notation for some of the constants in the following results, we remark that throughout the thesis the symbol C frequently indicates a generic finite positive constant.

We give the first two results in a more general context of measure spaces since we use them both on Euclidean domains with the standard Lebesgue measure and on Riemannian manifolds with the Riemann–Lebesgue measure (see Subsection 2.2.1).

Theorem 2.1.2 (Fatou's lemma, [Alt06, Theorem A3.20]). *Let (S, Σ, μ) be a measure space and $\{f^j\}_{j \in \mathbb{N}}$ be a sequence of Σ -measurable non-negative functions from S to \mathbb{R} . Then it holds*

$$\int_S \liminf_{j \rightarrow \infty} f^j \, d\mu \leq \liminf_{j \rightarrow \infty} \int_S f^j \, d\mu.$$

Theorem 2.1.3 (Dominated convergence theorem, [Alt06, Theorem A3.21]). *Let (S, Σ, μ) be a measure space and $\{f^j\}_{j \in \mathbb{N}}$ and g be Σ -measurable functions from S to \mathbb{R} . Furthermore,*

- g is integrable, i.e., $\int_S |g| \, d\mu < \infty$,
- $|f^j| \leq g$ μ -almost everywhere for all $j \in \mathbb{N}$,
- $f^j \rightarrow f$ μ -almost everywhere as $j \rightarrow \infty$.

Then $\{f^j\}_{j \in \mathbb{N}}$ and f are integrable and

$$\lim_{j \rightarrow \infty} \int_S f^j \, d\mu = \int_S f \, d\mu.$$

Theorem 2.1.4 (Poincaré's inequality, [Alt06, Theorem 6.7]). *Let $\Omega \subset \mathbb{R}^n$ be a open and bounded set. Then there exists a constant $C_P = C_P(\Omega, n, p)$ such that*

$$\|f\|_{L^p(\Omega)} \leq C_P \|Df\|_{L^p(\Omega)}, \quad \forall f \in W_0^{1,p}(\Omega).$$

Theorem 2.1.5 (Gagliardo–Nirenberg interpolation inequality, [Nir66]). *Let $\Omega \subset \mathbb{R}^n$ be a bounded domain satisfying the cone condition. If $f \in L^2(\Omega)$ and $D^m f \in L^2(\Omega)$, then there exist constants $C_{GN,1} = C_{GN,1}(\Omega, m, n) > 0$ and $C_{GN,2} = C_{GN,2}(\Omega, m, n) > 0$ such that for every $j \in \{1, \dots, m-1\}$*

$$\|f\|_{H^j(\Omega)} \leq C_{GN,1} \|f\|_{H^m(\Omega)}^{\frac{j}{m}} \|f\|_{L^2(\Omega)}^{1-\frac{j}{m}} + C_{GN,2} \|f\|_{L^2(\Omega)} \leq C_{GN,1} \|f\|_{H^m(\Omega)} + (C_{GN,1} + C_{GN,2}) \|f\|_{L^2(\Omega)}.$$

Theorem 2.1.6 (Korn's inequality, [Cia88, Section 6.3]). *Let $\Omega \subset \mathbb{R}^n$ be a bounded domain with a Lipschitz boundary. Then there exists a constant $C_{Korn} = C_{Korn}(\Omega, n) > 0$ such that*

$$C_{Korn}^{-1} \|f\|_{H^1(\Omega)} \leq \|(Df)^{\text{sym}}\|_{L^2(\Omega)} \leq C_{Korn} \|f\|_{H^1(\Omega)}, \quad \forall f \in H_0^1(\Omega),$$

where $(Df)^{\text{sym}} := \frac{1}{2}(Df + (Df)^T)$ is the *symmetric gradient*.

The following classical result allows for an approximation of L^p -functions, for $1 \leq p < \infty$, with smooth functions, and it will be used throughout this thesis.

Theorem 2.1.7 ([Alt06, Theorem 4.15]). *Let $\Omega \subset \mathbb{R}^n$ be a bounded domain. Then the set $C_c^\infty(\Omega)$ is dense in $L^p(\Omega)$ for $1 \leq p < \infty$. In particular, given $f \in L^p(\Omega)$ there exists a sequence of functions $\{f^j\}_{j \in \mathbb{N}} \in C_c^\infty(\Omega)$ and a function $C : \mathbb{R}^+ \times \mathbb{N} \rightarrow [0, \infty)$ such that*

$$\|f - f^j\|_{L^p(\Omega)} \leq 2^{-j}, \quad \|Df^j\|_{L^\infty(\Omega)} \leq C(\|f\|_{L^p(\Omega)}, j), \quad \forall j \in \mathbb{N}.$$

The following theorem comprises the central embedding theorem for Sobolev and Hölder functions that we will frequently refer to.

Theorem 2.1.8 (Embedding theorem for Sobolev and Hölder functions, [Alt06, Theorem 10.13]). *Let $\Omega \subset \mathbb{R}^n$ be a bounded domain with a Lipschitz boundary.*

(i) *If $m_1, m_2 \in \mathbb{N}$ and $p_1, p_2 \in [1, \infty)$ satisfy*

$$m_1 - \frac{n}{p_1} \geq m_2 - \frac{n}{p_2}, \quad m_1 \geq m_2, \quad (2.1)$$

then a continuous embedding $W^{m_1, p_1}(\Omega) \hookrightarrow W^{m_2, p_2}(\Omega)$ exists and for all $f \in W^{m_1, p_1}(\Omega)$ one obtains

$$\|f\|_{W^{m_2, p_2}(\Omega)} \leq C(\Omega, n, m_1, m_2, p_1, p_2) \|f\|_{W^{m_1, p_1}(\Omega)}.$$

If both inequalities in (2.1) are strict, then the embedding is compact.

(ii) *If $m \in \mathbb{N}^+$ and $p \in [1, \infty)$ are given such that*

$$m - \frac{n}{p} \geq k + \alpha \quad \text{for any } \alpha \in (0, 1), k \in \mathbb{N} \quad (2.2)$$

holds true, then a continuous embedding $W^{m, p}(\Omega) \hookrightarrow C^{k, \alpha}(\overline{\Omega})$ exists such that for all $f \in W^{m, p}(\Omega)$ a representative \tilde{f} of f exists with $\tilde{f}(x) = f(x)$ for a.e. $x \in \Omega$ and

$$\|\tilde{f}\|_{C^{k, \alpha}(\overline{\Omega})} \leq C(\Omega, n, m, p, k, \alpha) \|f\|_{W^{m, p}(\Omega)}.$$

If the inequality (2.2) is strict, then the embedding $W^{m, p}(\Omega) \hookrightarrow C^{k, \alpha}(\overline{\Omega})$ is additionally compact.

For the space of continuous functions, we have a classical criterion for convergence and compactness.

Theorem 2.1.9. (i) *(Arzelá–Ascoli’s theorem for compact domains, [Alt06, Theorem 4.12]). Let K be a compact subset of \mathbb{R}^n and let a sequence of functions $\{f^j\}_{j \in \mathbb{N}} \in C^0(K)$ satisfy*

- $\{f^j\}_{j \in \mathbb{N}}$ *is uniformly bounded in $C^0(K)$,*
- $\{f^j\}_{j \in \mathbb{N}}$ *is uniformly equicontinuous, i.e., for every $\epsilon > 0$, there exists a $\delta > 0$ such that $|f^j(x) - f^j(y)| < \epsilon$ for every $|x - y| < \delta$ and for all $j \in \mathbb{N}$.*

Then $\{f^j\}_{j \in \mathbb{N}}$ has a subsequence which converges to some $f \in C^0(K)$ in $C^0(K)$ -norm.

(ii) *(Arzelá–Ascoli’s theorem for Bochner spaces, [Sim87]). Let X be a Banach space and let a sequence of functions $\{f^j\}_{j \in \mathbb{N}} \in C^0([0, 1], X)$ satisfy*

- $f_t^j \rightarrow f_t$ *in X for every $t \in (0, 1)$*
- $\{f^j\}_{j \in \mathbb{N}}$ *is uniformly equicontinuous in time, i.e., for every $\epsilon > 0$, there exists a $\delta > 0$ such that $\|f_{t_1}^j - f_{t_2}^j\|_X < \epsilon$, for every $|t_1 - t_2| < \delta$ and for all $j \in \mathbb{N}$.*

Then $f^j \rightarrow f \in C^0([0, 1], X)$.

The following result is a corollary of the above theorem and the fact that uniformly bounded Hölder continuous families of functions (same applies for other choices of the joint modulus of continuity) satisfy the equicontinuity conditions.

Corollary 2.1.10 ([Alt06, Theorem 10.6]). *Let $k_1, k_2 \in \mathbb{N}$ and $\alpha_1, \alpha_2 \in (0, 1]$ with $k_1 + \alpha_1 \geq k_2 + \alpha_2$. Let $\Omega \subset \mathbb{R}^n$ be open and bounded (with Lipschitz boundary if $k_1 > 0$). Then $C^{k_1, \alpha_1}(\overline{\Omega}) \hookrightarrow C^{k_2, \alpha_2}(\overline{\Omega})$. If the inequality is strict the embedding is in addition compact. The analogous results hold for the Bochner spaces $C^{l, \beta}([0, 1], C^{k, \alpha}(\overline{\Omega}))$.*

2.2 Differential Geometry

In this section, we present a short survey of concepts from differential geometry, which will be frequently used in the rest of this thesis. We focus on finite-dimensional Riemannian manifolds (see e.g. [dC92, Lee03]), while in the later chapters we will revisit some of the introduced notions for Riemannian manifolds modeled on general (possibly infinite-dimensional) Banach or Hilbert spaces, studied in [Lan95, Kli95, Sak96, War83].

Definition 2.2.1 (Topological manifold, chart [Lee03, Definition 1.1]). Let \mathcal{M} be a Hausdorff topological space with a countable basis. We say that \mathcal{M} is an *m-dimensional topological manifold (with boundary)* if for each point $x \in \mathcal{M}$ there exists a neighborhood \mathcal{M}_x and a homeomorphism φ_x from \mathcal{M}_x to an open subset of \mathbb{R}^m or a relatively open subset of $\mathbb{R}_+^m := \{(x_1, \dots, x_m) : x_m \geq 0\}$. We call such a pairing $(\mathcal{M}_x, \varphi_x)$ a *chart*, or to make the distinction a *regular* or a *boundary chart*, respectively. Points lying in regular charts are called *interior points*, while points lying in a boundary chart are called *boundary points*.

This definition allows local identification of manifolds with Euclidean domains. Analogously, if the neighborhoods on \mathcal{M} can be identified in the above manner with a Banach space or Hilbert space then we say that \mathcal{M} is a *Banach or Hilbert manifold*. This setting will be put into use in Section 3.1 and we do not study it further at this point.

To do calculus on manifolds we need further properties of charts.

Definition 2.2.2 (Differentiable manifold, [dC92, Chapter 0, Definition 2.1]). Let \mathcal{M} be an *m-dimensional topological manifold*.

- (i) We say that $(\mathcal{M}_i, \varphi_i)_{i \in I}$ is a *family of charts* if $(\mathcal{M}_i)_{i \in I}$ is an *open covering* of \mathcal{M} and $(\mathcal{M}_i, \varphi_i)$ is a chart for every $i \in I$.
- (ii) A *C^k -differentiable atlas* ($k \geq 1$) is a family of charts $(\mathcal{M}_i, \varphi_i)_{i \in I}$ such that

$$\varphi_j \circ \varphi_i^{-1} : \varphi_i(\mathcal{M}_i \cap \mathcal{M}_j) \rightarrow \varphi_j(\mathcal{M}_i \cap \mathcal{M}_j)$$

is a C^k -diffeomorphism for each pair of indices $i, j \in I$. Two differentiable atlases are equivalent if the union of these atlases is a differentiable atlas.

- (iii) A *C^k -differentiable structure* is an equivalence class of C^k -differentiable atlases. A *C^k -differentiable manifold* is a topological manifold \mathcal{M} endowed with a C^k -differentiable structure.

Whether a point is an interior or a boundary point does not depend on the choice of charts [Lee03, Theorem 1.46]. The set of all interior points is denoted by $\text{Int}(\mathcal{M})$ and called the *interior of \mathcal{M}* and its complement is the *boundary of \mathcal{M}* denoted by $\partial\mathcal{M} = \mathcal{M} \setminus \text{Int}(\mathcal{M})$. In particular, x is a boundary point if the m -th component of $\varphi(x)$ equals 0 for all (boundary) charts φ it lies in. Here and for the rest of this thesis, if not otherwise stated, we assume \mathcal{M} to be *smooth* in the sense that k is sufficiently large or $k = \infty$. If for every family of charts there exists a finite subfamily of charts we say that the \mathcal{M} is *compact*.

Further properties and examples of topological and differentiable manifolds (with boundary) can be found in [Lee03, Chapter 1] and [dC92, Chapter 0].

We now proceed to definitions and properties of some of the key concepts on differentiable manifolds. The above-defined behavior of charts on intersections of their domains allows us to, for convenience, assume the existence of a global chart (\mathcal{M}, φ) .

Definition 2.2.3 (Tangent space, canonical basis, tangent bundle, [dC92, Chapter 0, Definition 2.6]). Let \mathcal{M} be a smooth *m-dimensional manifold*.

- (i) The *tangent space* $T_x\mathcal{M}$ of \mathcal{M} at $x \in \mathcal{M}$ is defined as

$$T_x\mathcal{M} := \{\dot{\gamma}(0) \text{ where } \gamma : (-\epsilon, \epsilon) \rightarrow \mathcal{M} \text{ is a smooth curve with } \gamma(0) = x, \epsilon > 0\}.$$

- (ii) For every $x \in \mathcal{M}$, we have that $T_x\mathcal{M}$ is a *m-dimensional vector space*. Indeed, if $x = \varphi^{-1}(\xi)$ we have $T_x\mathcal{M} = \text{span}\{\partial_{\xi_i} \varphi^{-1}(\xi)\}_{i=1}^m$. We denote this basis by $\{e_i\}_{i=1}^m$ and call it *the canonical basis of $T_x\mathcal{M}$* .

(iii) We call the set $T\mathcal{M} := \{(x, v), x \in \mathcal{M}, v \in T_x\mathcal{M}\}$ the *tangent bundle* of \mathcal{M} .

Notice that for $x \in \partial\mathcal{M}$ the tangent space also contains the outward pointing vectors $\sum_{i=1}^m v_i e_i(\xi)$ with $\{v_i\}_{i=1}^m \in \mathbb{R}$ and $v_m < 0$. Having defined the notion of tangent space, we can now study functions mapping from the manifold to a Euclidean space or a manifold.

Definition 2.2.4 (Scalar and vector fields, [dC92, Chapter 0, Definition 5.1]). Let \mathcal{M} be a smooth m -dimensional manifold.

- (i) A *scalar function* or a *scalar field* $f : \mathcal{M} \rightarrow \mathbb{R}$ is said to be C^k -differentiable (or smooth for $k = \infty$) if $f \circ \varphi^{-1} : U \subset \mathbb{R}^m \rightarrow \mathbb{R}$ is C^k or smooth, respectively. We denote the set of all smooth scalar fields by $\mathcal{F}(\mathcal{M})$.
- (ii) A *vector field* X on \mathcal{M} is a mapping $X : \mathcal{M} \rightarrow T\mathcal{M}$. For $I \subset \mathbb{R}$, a vector field $X : I \rightarrow T\mathcal{M}$ along a curve $\gamma : I \rightarrow \mathcal{M}$ fulfills $X(t) = X(\gamma(t)) \in T_{\gamma(t)}\mathcal{M}$. Thus, for $\xi \in U$ we can write $X(\varphi^{-1}(\xi)) = \sum_{i=1}^m v_i(\xi) e_i(\xi)$, with $v_i : U \subset \mathbb{R}^m \rightarrow \mathbb{R}$. We say X is smooth if and only if the functions v_i are smooth. We denote the set of all smooth vector fields by $\mathfrak{X}(\mathcal{M})$.
- (iii) For $X \in \mathfrak{X}(\mathcal{M})$ and $f \in \mathcal{F}(\mathcal{M})$ we define $X(f)(x) := \sum_{i=1}^m v_i(\xi) \frac{\partial f \circ \varphi^{-1}}{\partial \xi_i}(\xi)$, where $\xi = \varphi(x)$.

Definition 2.2.5 (Differentiable mappings between manifolds, immersions and embeddings, [dC92, Chapter 0, Definitions 2.1 and 3.1]). Let \mathcal{M}^1 and \mathcal{M}^2 be two smooth manifolds with the corresponding smooth atlases $(\mathcal{M}_i^1, \varphi_i^1)_{i \in I}$ and $(\mathcal{M}_j^2, \varphi_j^2)_{j \in J}$.

- (i) A mapping $F : \mathcal{M}^1 \rightarrow \mathcal{M}^2$ is said to be C^k -differentiable if

$$\varphi_j^2 \circ F \circ \left(\varphi_i^1 \Big|_{\mathcal{M}_i^1 \cap F^{-1}(\mathcal{M}_j^2)} \right)^{-1} : \varphi_i^1(\mathcal{M}_i^1 \cap F^{-1}(\mathcal{M}_j^2)) \rightarrow \varphi_j^2(\mathcal{M}_j^2)$$

is C^k -differentiable for any $(i, j) \in I \times J$. This definition is independent of the choice of families of charts.

- (ii) Given a differentiable mapping $F : \mathcal{M}^1 \rightarrow \mathcal{M}^2$, a point $x \in \mathcal{M}^1$ and $v \in T_x\mathcal{M}^1$ we define the *differential* of F as the linear mapping $dF_x : T_x\mathcal{M}^1 \rightarrow T_{F(x)}\mathcal{M}^2$ by $dF_x(v) := (F \circ \gamma)'(0)$ where $\gamma : (-\epsilon, \epsilon) \rightarrow \mathcal{M}^1$ is an arbitrary smooth curve with $\gamma(0) = x$ and $\dot{\gamma}(0) = v$.
- (iii) A differentiable mapping is said to be an *immersion* if dF_x is injective for all $x \in \mathcal{M}^1$.
- (iv) If an immersion is a homeomorphism onto $F(\mathcal{M}^1)$, where $F(\mathcal{M}^1)$ has the topology induced from \mathcal{M}^2 , we say that F is an *embedding*.

A crucial notion in the interaction of scalar and vector fields is the notion of an affine connection. This leads to the definition of covariant derivative which is of the central importance for developing the geodesic calculus.

Definition 2.2.6 (Affine connection, covariant derivative, [dC92, Chapter 2, Definition 2.1]).

- (i) An *affine connection* on a smooth manifold \mathcal{M} is a mapping

$$\nabla : \mathfrak{X}(\mathcal{M}) \times \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{X}(\mathcal{M}), \quad (X, Y) \mapsto \nabla_X Y$$

such that

$$\begin{aligned} \nabla_{f_1 X_1 + f_2 X_2} Y &= f_1 \nabla_{X_1} Y + f_2 \nabla_{X_2} Y, \\ \nabla_X (f_1 Y_1 + f_2 Y_2) &= f_1 \nabla_X Y_1 + X(f_1) Y_1 + f_2 \nabla_X Y_2 + X(f_2) Y_2, \end{aligned}$$

for all $X, X_1, X_2, Y, Y_1, Y_2 \in \mathfrak{X}(\mathcal{M})$ and $f_1, f_2 \in \mathcal{F}(\mathcal{M})$.

- (ii) We refer to $\nabla_X Y$ as the *covariant derivative* of Y w.r.t. X .
- (iii) If $X(t) = Y(\gamma(t))$ for some $Y \in \mathfrak{X}(\mathcal{M})$ then we define $\frac{D}{dt} X := \nabla_{\dot{\gamma}} Y$. The correspondence $X \mapsto \frac{D}{dt} X$ is called *covariant derivative of X along γ* [dC92, Chapter 2, Proposition 2.3].

So far, we worked with (sufficiently) smooth manifolds. To study further notions of interest in this thesis, we now define a Riemannian manifold as a manifold with a metric defined on tangent space(s).

Definition 2.2.7 (Riemannian metric, Riemannian manifold, [dC92, Chapter 1, Definition 2.1]).

- (i) A *Riemannian metric* on \mathcal{M} is a mapping $g : x \mapsto g_x$ such that $g_x : T_x\mathcal{M} \times T_x\mathcal{M} \rightarrow \mathbb{R}$ is a bilinear, symmetric and positive-definite form, and $\xi \mapsto g_{ij}(\xi) := g_{\varphi^{-1}(\xi)}(e_i(\xi), e_j(\xi))$ is a smooth function on $\varphi(\mathcal{M})$, where $\{e_i\}_{i=1}^m$ is the canonical basis.
- (ii) A manifold \mathcal{M} equipped with a Riemannian metric g is referred to as a *Riemannian manifold* with notation (\mathcal{M}, g) .

The Riemannian metric gives rise to a norm on tangent spaces: given $v \in T_x\mathcal{M}$, we will use notation $\|v\|_g := \sqrt{g_x(v, v)}$ for this norm, where point x will be clear from the context.

The following theorem introduces a distinct affine connection, called the *Levi-Civita connection*, with the associated derivative called *Levi-Civita derivative*. This connection is unique on every Riemannian manifold once a Riemannian metric is chosen. For the rest of this thesis, ∇ will always refer to Levi-Civita connection.

Theorem 2.2.8 (Levi-Civita connection, [dC92, Chapter 2, Theorem 3.6]). *On a Riemannian manifold (\mathcal{M}, g) there exists a unique affine connection ∇ such that*

- (i) ∇ is torsion-free or symmetric on \mathcal{M} , i.e.,

$$\nabla_X Y - \nabla_Y X = [X, Y] \quad \forall X, Y \in \mathfrak{X}(\mathcal{M}),$$

where Lie bracket is given by [dC92, Chapter 0, Lemma 5.2]

$$[\cdot, \cdot] : \mathfrak{X}(\mathcal{M}) \times \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{X}(\mathcal{M}), \quad [X, Y](f) = X(Y(f)) - Y(X(f))$$

for two vector fields $X, Y \in \mathfrak{X}(\mathcal{M})$ and all $f \in \mathcal{F}(\mathcal{M})$.

- (ii) The covariant derivative associated with ∇ is Riemannian, i.e.,

$$Xg(Y, Z) = g(\nabla_X Y, Z) + g(Y, \nabla_X Z) \quad \forall X, Y, Z \in \mathfrak{X}(\mathcal{M}).$$

We now introduce some of the crucial concepts as applications of the notion of covariant derivative. The first is the notion of *parallel transport* along a curve on the manifold. This is a concept of transferring tangent vectors along the path and ensuring that they “stay parallel w.r.t. the metric”, which later enables the transfer of geometric data along a specified path.

Proposition 2.2.9 (Parallel vector field, parallel transport, [dC92, Chapter 2, Definition 2.1 and Proposition 2.6]).

- (i) A smooth vector field $X \in \mathfrak{X}(\mathcal{M})$ is called parallel w.r.t. a smooth curve $\gamma : I \rightarrow \mathcal{M}$, if $\frac{DX}{dt}(\gamma(t)) = 0$ for all t in an open interval $I \subset \mathbb{R}$.
- (ii) Let $I \subset \mathbb{R}$ be an open interval, $t_0, t_1 \in I$, and $\gamma : I \rightarrow \mathcal{M}$ be a smooth curve. Then for every $X_0 \in T_{\gamma(t_0)}\mathcal{M}$ there exists a unique parallel vector field $X \in \mathfrak{X}(\mathcal{M})$ along γ such that $X(t_0) = X_0$ and $X(t) \in T_{\gamma(t)}\mathcal{M}$ for $t \in I$, and the mapping

$$\text{PT}_{t_0 \rightarrow t_1}^\gamma : T_{\gamma(t_0)}\mathcal{M} \rightarrow T_{\gamma(t_1)}\mathcal{M}, \quad X(t_0) \mapsto X(t_1)$$

is a linear and isometric isomorphism, which is called parallel transport.

This definition is a starting point for one of the central notions in Riemannian geometry: a geodesic curve.

Definition 2.2.10 (Geodesic curve, [dC92, Chapter 3, Definition 2.1]). A smooth curve $\gamma : I \rightarrow \mathcal{M}$ is called a *geodesic curve* or simply a *geodesic* if it satisfies the *geodesic equation*

$$\frac{D}{dt}\dot{\gamma}(t) = 0, \quad \forall t \in I. \quad (2.3)$$

In particular, the velocity vector of a geodesic curve is transported parallelly along the curve. Furthermore, by [dC92, Chapter 2, Proposition 3.2] having chosen the Levi–Civita connection we have

$$\frac{d}{dt} \left(g_{\gamma(t)}(\dot{\gamma}, \dot{\gamma}) \right) = 2g_{\gamma(t)} \left(\frac{D}{dt} \dot{\gamma}(t), \dot{\gamma}(t) \right) = 0,$$

which implies that $g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))$ is constant along a geodesic curve. This property is known as *arc length parametrization*.

The following result concerns the existence and uniqueness of geodesics. We will restrict ourselves in further discussion to manifolds without boundary (or simply just observe $\text{Int}(\mathcal{M})$) as properties of geodesics on the boundary and near the boundary often ask for more careful observation (due to the presence of the outward pointing tangent vectors).

Theorem 2.2.11 (Local existence and uniqueness of geodesics as solutions of the geodesic initial value problem). *For any point $x \in \mathcal{M}$ there exists a (largest) $\epsilon > 0$ such that for all $v \in T_x \mathcal{M}$ with $\|v\|_g < \epsilon$ there exists a unique geodesic $\gamma : [0, 1] \rightarrow \mathcal{M}$ with $\gamma(0) = x$ and $\dot{\gamma}(0) = v$.*

Proof. The existence and the uniqueness are first shown on an interval $(-\delta, \delta)$ for some $\delta > 0$ (cf. [dC92, Chapter 3, Proposition 2.5]). To this end, one notices that the geodesic equation $\frac{D}{dt} \dot{\gamma}(t)(\gamma(t)) = 0$ is a second-order ordinary differential equation (ODE). Existence and uniqueness then follow from the Picard-Lindelöf theorem for the solutions of ODEs with prescribed initial conditions. Furthermore, we note that γ depends smoothly on both x and v . Finally, we use homogeneity of geodesics [dC92, Chapter 3, Lemma 2.6] which states that given a geodesic $\gamma : (-\delta, \delta) \rightarrow \mathcal{M}$ with $\gamma(0) = x$, $\dot{\gamma}(0) = v$ the curve with the same initial point and velocity αv for some $\alpha > 0$ is a geodesic on the interval $(-\frac{\delta}{\alpha}, \frac{\delta}{\alpha})$. \square

We will denote this unique mapping by $t \mapsto \gamma(t, x, v)$, $t \in [0, 1]$. We have the following results.

Proposition 2.2.12 (Exponential and logarithm map, injectivity radius, normal neighborhood [dC92, Chapter 3, Proposition 2.9]).

(i) *For every $x \in \mathcal{M}$ there exists $\epsilon > 0$, called the injectivity radius at x , such that the exponential map at x*

$$\exp_x : U_x^\epsilon := \{v \in T_x \mathcal{M} : \|v\|_g < \epsilon\} \rightarrow \mathcal{M}, \quad \exp_x(v) := \gamma(1, x, v)$$

is well defined and a diffeomorphism.

(ii) *We call $\exp_x(U_x^\epsilon) \subset \mathcal{M}$ the normal neighborhood of x and define the logarithm map at x by*

$$\log_x : U_x^\epsilon \rightarrow T_x \mathcal{M}, \quad \log_x(y) := (\exp_x^{-1})(y).$$

Geometrically, $\log_x(y)$ is the initial velocity of the unique geodesic connecting x and y .

Geodesic curves are often characterized as *locally shortest curves*. To specify this notion we first introduce the notions of the length of a curve and the distance on a manifold.

Definition 2.2.13. (i) The *length* of a curve $\gamma : [0, 1] \rightarrow \mathcal{M}$ is defined by

$$\mathcal{L}[(\gamma(t))_{t \in [0, 1]}] := \int_0^1 \sqrt{g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt.$$

(ii) Given points $x, y \in \mathcal{M}$ we define their (*Riemannian*) *distance* $d_{\mathcal{M}} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ by

$$d_{\mathcal{M}}(x, y) := \inf \{ \mathcal{L}[\gamma(t)] : \gamma : [0, 1] \rightarrow \mathcal{M} \text{ piecewise differentiable, } \gamma(0) = x, \gamma(1) = y \}.$$

Notice that the definition of the length is invariant under reparametrization. The following proposition illustrates the distance minimizing properties of geodesic curves.

Proposition 2.2.14 ([dC92, Chapter 3, Proposition 3.6 and Corollary 3.9]). *Let $x \in \mathcal{M}$ and y be in the normal neighborhood of x . Let $\gamma, \tilde{\gamma} : [0, 1] \rightarrow \mathcal{M}$ be two piecewise differentiable curves connecting x and y , where γ is geodesic. Then $\mathcal{L}[\gamma] \leq \mathcal{L}[\tilde{\gamma}]$ with equality if and only if $\gamma([0, 1]) = \tilde{\gamma}([0, 1])$, i.e., γ and $\tilde{\gamma}$ are equal up to reparametrization.*

Conversely, let $\gamma : [0, 1] \rightarrow \mathcal{M}$ be a piecewise differentiable curve connecting $x, y \in \mathcal{M}$ that is parametrized proportional to arc length. If $\mathcal{L}[\gamma]$ is less or equal to the length of any other piecewise differentiable curve connecting x and y , then γ is a smooth geodesic.

This implies that if y is in the normal neighborhood of x the shortest path is along the geodesic and we have $d_{\mathcal{M}}(x, y) = \|\log_x(y)\|_g$ due to the arc length parametrization of geodesics. Another relevant notion concerning the curves is the path energy.

Definition 2.2.15. The *path energy* of a curve $\gamma : [0, 1] \rightarrow \mathcal{M}$ is defined by

$$\mathcal{E}[(\gamma(t))_{t \in [0, 1]}] := \int_0^1 g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t)) dt. \quad (2.4)$$

Remark 2.2.16. A direct application of the Cauchy–Schwarz inequality shows that

$$\mathcal{L}[(\gamma(t))_{t \in [0, 1]}] \leq \sqrt{\mathcal{E}[(\gamma(t))_{t \in [0, 1]}]},$$

and equality holds if and only if $g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t)) = \text{const}$, i.e., for curves with constant speed and arc length parametrization, including geodesics. Path energy is, unlike path length, dependent on parametrization, but it is a convex functional which presents an advantage in the study of the existence of minimizers. Therefore, it is important to observe that a minimizer of \mathcal{E} is indeed a minimizer of \mathcal{L} . To this end, let $(\gamma(t))_{t \in [0, 1]}$ be a minimizer of \mathcal{E} . Suppose that there exists a curve $(\tilde{\gamma}(t))_{t \in [0, 1]}$ such that $\mathcal{L}[\gamma] > \mathcal{L}[\tilde{\gamma}]$. Due to the invariance of length under the parametrization, we may assume $g_{\tilde{\gamma}(t)}(\dot{\tilde{\gamma}}(t), \dot{\tilde{\gamma}}(t)) \equiv \mathcal{L}[\tilde{\gamma}]$. Then we have

$$\mathcal{E}[\tilde{\gamma}] := \mathcal{L}[\tilde{\gamma}]^2 < \mathcal{L}[\gamma]^2 \leq \mathcal{E}[\gamma],$$

which is a contradiction to the choice of γ . Thus, for every $(\tilde{\gamma}(t))_{t \in [0, 1]}$ we have $\mathcal{L}[\tilde{\gamma}]^2 \geq \mathcal{L}[\gamma]^2$, as we wanted to show. Furthermore, one can show from Euler–Lagrange equation of path energy that a minimizer satisfies the geodesic equation (2.3) [HRW18, Remark 2.12].

For definitions of all of the above notions, we used the advantage of having a finite-dimensional manifold by using the canonical basis. To be able to define these notions for infinite-dimensional manifolds, we comment on the coordinate-free definition of covariant derivative which allows us to define other notions of interest. Definitions and properties of parallel transport, geodesics, and other notions obtained from covariant derivatives are analogous to the already existing. For more details see [Lan95].

Definition 2.2.17 (Covariant derivative, [Lan95, VIII, Theorem 3.1]). Let $\gamma : I \rightarrow \mathcal{M}$ be a curve and $X : I \rightarrow T\mathcal{M}$ a vector field along γ . We define the covariant derivative $\frac{D}{dt}X$ along γ at $x = \gamma(t)$ for $t \in I$ implicitly by

$$g_x\left(\frac{D}{dt}X(t), Y\right) = g_x\left(\dot{X}(t) + \Gamma_x(X(t), \dot{\gamma}(t)), Y\right), \quad \forall Y \in T_x\mathcal{M},$$

where the *Christoffel operator* $\Gamma_x : T_x\mathcal{M} \times T_x\mathcal{M} \rightarrow T_x\mathcal{M}$ is defined implicitly by [Lan95, VIII, Theorem 4.2]

$$g_x(\Gamma_x(X, Y), Z) = \frac{1}{2}\left(Xg_x(Y, Z) + Yg_x(X, Z) - Zg_x(X, Y)\right), \quad \forall X, Y, Z \in T_x\mathcal{M}.$$

2.2.1 Function Spaces on Riemannian Manifolds

In this section, we introduce some function spaces consisting of functions mapping from a smooth and compact m -dimensional Riemannian manifold (\mathcal{M}, g) to the Euclidean space. We are in particular interested in definitions of appropriate notions of derivatives (of the first and the second order), and Lebesgue and Sobolev spaces, for which we will need to introduce the integration on manifolds via the Riemann–Lebesgue measure.

In Definition 2.2.4, we already defined a scalar function $f : \mathcal{M} \rightarrow \mathbb{R}$ to be of class C^k , $k \geq 1$ if $f \circ \varphi^{-1} : \mathbb{R}^m \rightarrow \mathbb{R}$ is of that class. To define derivatives of such a function at a point on manifold, we observe the chart defined by the exponential map at that point (cf. Definition 2.2.12). This leads to the definition of the *Riemannian gradient (Jacobian)* $\text{grad } f(x) \in T_x \mathcal{M}$ implicitly via the identity [AMS09, Section 3.6]

$$\frac{d}{dt}(f \circ \exp_x)(tv)|_{t=0} = g(\text{grad } f(x), v), \quad \forall v \in T_x \mathcal{M}. \quad (2.5)$$

Analogously, the *Riemannian Hessian* $\text{Hess } f$ at x is the linear operator $\text{Hess } f(x) : T_x \mathcal{M} \rightarrow T_x \mathcal{M}$ given by [AMS09, Proposition 5.5.4]

$$\frac{d^2}{dt^2}(f \circ \exp_x(tv))|_{t=0} = g_x(\text{Hess } f(x)[v], v), \quad \forall v \in T_x \mathcal{M}. \quad (2.6)$$

It can also be defined by $\text{Hess } f(x)[v] = \nabla_v \text{grad } f(x)$ [AMS09, Definition 5.5.1], where ∇ is the Levi-Civita connection (cf. Theorem 2.2.8) and ∇_v the covariant derivative in direction v , showing that $\text{Hess } f(x)$ is a symmetric operator due to the symmetry property of the Levi-Civita connection. It is, thus, a symmetric endomorphism on $T_x \mathcal{M}$, denoted by $\text{Hess } f(x) \in L_{\text{sym}}(T_x \mathcal{M}, T_x \mathcal{M})$.

For a function $f : \mathcal{M} \rightarrow \mathbb{R}^l$ we have $\text{grad } f(x) = (\text{grad } f_1(x), \dots, \text{grad } f_l(x)) \in (T_x \mathcal{M})^l$ and $\text{Hess } f(x) = (\text{Hess } f_1(x), \dots, \text{Hess } f_l(x)) : T_x \mathcal{M} \rightarrow (T_x \mathcal{M})^l$, where the above definitions are applied componentwise. The application of the Riemannian metric g on $(T_x \mathcal{M})^l \times T_x \mathcal{M}$ is to be understood componentwise, i.e., we will use the notation $g(A, v) = (g(A_j, v))_{j=1, \dots, l}$ for an l -tuple $A = (A_1, \dots, A_l)$ of tangent vectors. Similarly, we write $g(A, B) = \sum_{j=1}^l g(A_j, B_j)$ for $A, B \in (T_x \mathcal{M})^l$ and $g(H, K) = \sum_{j=1}^m g(H[v_j], K[v_j])$ for linear operators $H, K : T_x \mathcal{M} \rightarrow T_x \mathcal{M}$, where v_1, \dots, v_m is any orthonormal basis of $T_x \mathcal{M}$.

We will denote the functions with the continuous second order derivatives in the above sense by $C^2(\mathcal{M}, \mathbb{R}^l)$. One could define also the higher order derivatives in an analogous way, but in this thesis, we will be interested only into the first and second order derivatives so we omit further discussion.

The following class of functions we are interested in are Lipschitz continuous functions. A mapping $f : \mathcal{M} \rightarrow \mathbb{R}$ is said to be *Lipschitz continuous* with *Lipschitz constant* $L(f)$ if for all $x \in \mathcal{M}$ and all y in the normal neighbourhood (cf. Definition 2.2.12) of x it holds [Bou23, Definition 10.40]

$$|f(x) - f(y)| \leq L(f)d_{\mathcal{M}}(x, y),$$

where $d_{\mathcal{M}}$ is the distance of the points on the manifold defined in Definition 2.2.13. If f is differentiable its gradient is Lipschitz continuous with Lipschitz constant $L_{\text{grad}}(f)$ if [Bou23, Definition 10.44]

$$g_x(\text{PT}_{0 \leftarrow 1}^\gamma \text{grad } f(y) - \text{grad } f(x), \text{PT}_{0 \leftarrow 1}^\gamma \text{grad } f(y) - \text{grad } f(x)) \leq L_{\text{grad}}(f)d_{\mathcal{M}}(x, y), \quad (2.7)$$

where $\gamma : [0, 1] \rightarrow \mathcal{M}$ is the unique minimizing geodesic connecting x to y (cf. Definition 2.2.10) and $\text{PT}_{0 \leftarrow 1}^\gamma : T_y \mathcal{M} \rightarrow T_x \mathcal{M}$ denotes the parallel transport along γ from $\gamma(1) = y$ to $\gamma(0) = x$ (cf. Proposition 2.2.9). If f is twice differentiable, its Hessian is Lipschitz continuous with Lipschitz constant $L_{\text{Hess}}(f)$ if [Bou23, Definition 10.49]

$$g_x(\text{PT}_{0 \leftarrow 1}^\gamma \text{Hess } f(y) \text{PT}_{1 \leftarrow 0}^\gamma - \text{Hess } f(x), \text{PT}_{0 \leftarrow 1}^\gamma \text{Hess } f(y) \text{PT}_{1 \leftarrow 0}^\gamma - \text{Hess } f(x)) \leq L_{\text{Hess}}(f)d_{\mathcal{M}}(x, y), \quad (2.8)$$

where $\text{PT}_{0 \leftarrow 1}^\gamma$ is the inverse of $\text{PT}_{1 \leftarrow 0}^\gamma$. In all the cases, the estimates should hold for every $x \in \mathcal{M}$ and every y in the normal neighborhood of x . The second and the third definition are special cases of definitions of Lipschitz continuity for general vector fields and linear mappings from $T\mathcal{M}$ to $T\mathcal{M}$. The function space of twice differentiable functions with Lipschitz continuous second derivative is denoted by $C^{2,1}(\mathcal{M}, \mathbb{R})$. In the case of functions mapping to \mathbb{R}^l , the above definitions are meant in the corresponding componentwise sense (as defined above), and we use the notation $C^{2,1}(\mathcal{M}, \mathbb{R}^l)$.

To define the Lebesgue and Sobolev spaces on \mathcal{M} , we first need to introduce the natural measure for integration of functions $f : \mathcal{M} \rightarrow \mathbb{R}$. To this end, let us have a family of charts $(\mathcal{M}_i, \varphi_i)_{i \in I}$ and a partition of unity $(\alpha_i)_{i \in I}$ subordinated to the covering $(\mathcal{M}_i)_{i \in I}$, i.e., for every x there exists a neighborhood such that all but finitely many of functions are zero, $\sum_{i \in I} \alpha_i = 1$ and $\text{supp } \alpha_i \subset \mathcal{M}_i$ for every $i \in I$. Then we define the *Riemann-Lebesgue measure* of continuous function f , i.e., the integration of the function f on \mathcal{M} by [Heb96, Section 1.2]

$$V_g(f) := \int_{\mathcal{M}} f dV_g := \sum_{i \in I} \int_{\varphi_i(\mathcal{M}_i)} (\alpha_i \sqrt{\det g_i} f) \circ \varphi_i^{-1} dx, \quad (2.9)$$

where g_i is the matrix representation of metric g in the canonical basis $\{e_j = \partial_{\xi_j} \varphi_i^{-1}(\xi)\}_{j=1}^m$ (cf. Definition 2.2.7). The standard Lebesgue integration is well-defined due to the continuity of f and compactness of \mathcal{M} and the sum is finite due to the same reason. Finally, one extends the definition to the general integrable functions $f : \mathcal{M} \rightarrow \mathbb{R}$ by approximating the function with continuous functions (cf. [Sak96, Section II.5]).

The *Lebesgue spaces* $L^p(\mathcal{M})$, for $1 \leq p < \infty$, can be defined as the closure of the set of smooth functions under the norm

$$\|f\|_{L^2(\mathcal{M})} := \left(\int_{\mathcal{M}} |f|^2 dV_g \right)^{\frac{1}{2}}.$$

The *Sobolev spaces* $H^1(\mathcal{M})$ and $H^2(\mathcal{M})$ are defined as the closure of $C^\infty(\mathcal{M})$ under the norms

$$\begin{aligned} \|f\|_{H^1(\mathcal{M})}^2 &:= \int_{\mathcal{M}} |f|^2 + g(\text{grad } f, \text{grad } f) dV_g, \\ \|f\|_{H^2(\mathcal{M})}^2 &:= \int_{\mathcal{M}} |f|^2 + g(\text{grad } f, \text{grad } f) + g(\text{Hess } f, \text{Hess } f) dV_g, \end{aligned}$$

where in the case of functions mapping to \mathbb{R}^l we use the above-defined summation over the components. The second and the third term in the above expression are the $L^2(\mathcal{M}, T\mathcal{M})$ - and $L^2(\mathcal{M}, L_{\text{sym}}(T\mathcal{M}, T\mathcal{M}))$ -norms, respectively.

The higher order Sobolev spaces, with possibly different exponents, can be defined analogously (see also definitions of Lebesgue and Sobolev spaces in Subsection 2.1.2), but we omit this discussion here and refer the reader to [Heb96]. We discuss a few properties of the above-defined Sobolev spaces that will be of use. First, by [Heb96, Proposition 2.3], we have that $H^2(\mathcal{M}, \mathbb{R}^l)$ and $H^1(\mathcal{M}, \mathbb{R}^l)$ are reflexive. Furthermore, the following result allows us to estimate the Lebesgue norm of the function by the norm of the gradient.

Theorem 2.2.18 (Poincaré's inequality, [Heb96, Theorem 2.10]). *Let (\mathcal{M}, g) be a smooth compact m -dimensional manifold. Then there exists a constant $C_P(\mathcal{M}, m)$ such that*

$$\|f - \bar{f}\|_{L^2(\mathcal{M})} \leq C_P \|\text{grad } f\|_{L^2(\mathcal{M}, T\mathcal{M})}, \quad \forall f \in H^1(\mathcal{M}),$$

where $\bar{f} := \frac{1}{V_g(\mathcal{M})} \int_{\mathcal{M}} f dV_g$ with $V_g(\mathcal{M}) := \int_{\mathcal{M}} 1 dV_g$.

This result holds also on Euclidean spaces and should be compared to Theorem 2.1.4. In particular, for functions with zero mean, i.e., $\bar{f} = 0$, we have the analogous result as in that case. The spaces of functions with zero mean, as subspaces of previously defined function spaces, will be denoted by $\dot{H}^2(\mathcal{M})$ and $\dot{L}^2(\mathcal{M})$, respectively.

The following result is analogous to Theorem 2.1.8. It holds in the same generality as the mentioned result, but we use only a part of it, as we did not introduce Sobolev spaces on a Riemannian manifold in the full generality.

Theorem 2.2.19 (Rellich embedding theorem for compact manifolds, [Heb96, Theorem 2.9]). *Let (\mathcal{M}, g) be a smooth compact m -dimensional manifold. Then there exist compact embeddings $H^2(\mathcal{M}) \hookrightarrow H^1(\mathcal{M}) \hookrightarrow L^2(\mathcal{M})$.*

2.3 Neural Networks

Neural networks, stemming from the works from 1940s [MP43], represent an algorithmic approach to learning with the aim of building a theory of artificial intelligence inspired by the functionality of the human brain. This is reflected in the structure of neural networks which in a broad context consist of neurons arranged in layers and connected by weighted edges. Today, neural networks, especially those with a large number of layers, are state-of-the-art technology for a wide variety of applications in mathematics, sciences, and engineering: solving PDEs [HJE18, Y⁺18], image classification [KSH12, SZ14], speech recognition [HDY⁺12], natural language processing [YHPC18], object recognition [HZRS16], game intelligence [SHM⁺16] and numerous other applications. This application directed development is recent and it is due to use of the modern software and hardware technologies which allowed faster computation together with better memory management. There is a plethora of books and review papers on neural networks. We mention here and refer in this thesis to the books [GBC16] and [GK22] which represent both an introduction to the field of deep learning (study of neural networks with multiple layers) and a review of modern research in this field, both from the theoretical and application oriented perspectives.

In this section, we introduce the basic notions and notations for the study of neural networks. For more details, we refer the reader to the above-mentioned books and references therein. We later study properties of function spaces consisting of functions representable by neural networks.

Definition 2.3.1 (Neural network). We call a family of tuples $\Phi = ((A_l, b_l, \sigma_l))_{l=1}^L$ with weights $A_l \in \mathbb{R}^{n_l, n_{l-1}, c_l}$, biases $b_l \in \mathbb{R}^{n_l, c_l}$, and activation functions $\sigma_l : \mathbb{R} \rightarrow \mathbb{R}$ a *neural network architecture* or simply a *neural network*. We call n_0 the input dimension and n_L the output dimension. We refer to L as the number of layers, n_l is the number of neurons in layer l , and $N = \sum_{l=1}^L n_l$ as the number of neurons of Φ , while c_l is called the number of channels in layer l . The network is called *deep* if $L \geq 3$, and otherwise it is said to be *shallow*.

We define the *realization* of Φ as the function $R(\Phi) : \mathbb{R}^{n_0, c_0} \rightarrow \mathbb{R}^{n_L, c_L}$, $x \mapsto x_L$ with

$$\begin{aligned} x_0 &= x, \\ x_l &:= \sigma_l(A_l x_{l-1} + b_l), \quad l = 1, \dots, L-1, \\ x_L &:= A_L x_{L-1} + b_L, \end{aligned}$$

where elements of $(x_l)_{l=1}^L$ are called *neurons*. Here, σ_l , $l = 1, \dots, L$ act componentwise, σ_L is the identity function, and

$$(A_l x_{l-1})_c = \sum_{i=1}^{c_{l-1}} (A_l)_c (x_{l-1})_i, \quad c = 1, \dots, c_l. \quad (2.10)$$

Weights, biases, and activation functions are optimized to approximate a task dependent unknown function based on a given set of input-output value pairs. This process is called *supervised learning*. This optimization is usually accomplished through some version of the stochastic gradient descent algorithm, where the gradients are computed via backpropagation through the network [RHW86].

This type of neural network is often called *fully connected feedforward neural network*. The term feedforward refers to the structure where realization is defined in a completely forward manner and there are no feedback connections in which outputs of the model are fed back into itself. Full connectedness means that every neuron in one layer is connected to every neuron in the other layer. Furthermore, we observe that every element of A_l interacts with only one element per channel of x_{l-1} . One can also have, e.g. for images, $x_l \in \mathbb{R}^{n_l, m_l, c_l}$ in which case $b_l \in \mathbb{R}^{n_l, m_{l-1}, c_l}$. If we are working with images, it is often advantageous to have a structure that is still local but allows for the interaction of more than one pair of weights and inputs. This is achieved through the structure of *convolutional neural networks*, where one replaces the matrix vector multiplication in (2.10) by the cross-correlation operation $w \star x$.

Definition 2.3.2 (Convolutional neural network). A *convolutional layer* is defined by a kernel $w \in \mathbb{R}^{w_v, w_h, c_{out}}$, a bias b , and parameters for padding $p = (p_d, p_u, p_l, p_r)$ and stride $s = (s_v, s_h)$. Given $x \in \mathbb{R}^{v_{in}, h_{in}, c_{in}}$, its realization is defined by the function $x \mapsto w \star \bar{x} + b$, where

$$\bar{x} \in \mathbb{R}^{v_{in}+p_d+p_u, h_{in}+p_l+p_r, c_{in}}, \quad \bar{x}(i, j) = \begin{cases} x(i, j), & i \in [p_d+1, p_d+v_{in}] \text{ and } j \in [p_l+1, p_l+h_{in}] \\ 0, & \text{else,} \end{cases}$$

$$(w \star \bar{x})(i, j) := \sum_{c=1}^{c_{in}} \sum_{k=1}^{w_v} \sum_{l=1}^{w_h} w(k, l) \bar{x}(s_v(i-1) + k, s_h(j-1) + l, c). \quad (2.11)$$

The corresponding dimension for bias is $b \in \mathbb{R}^{b_v, b_h, c_{out}}$ where $b_v = \lfloor \frac{v_{in} + p_l + p_r - w_v}{s_v} + 1 \rfloor$ and $b_h = \lfloor \frac{h_{in} + p_d + p_u - w_h}{s_h} + 1 \rfloor$. The parameters w_v, w_h are called the *kernel size* and c_{in}, c_{out} are the *number of input/output channels*.

A *convolutional neural network* is a family of convolutional layers together with activation functions with realization defined by the sequential applications of realizations of the layers.

Let us notice that the fully connected feedforward network architecture also includes convolutional neural networks. Indeed, one can write the cross-correlation operation in terms of matrix vector multiplication by writing

rows of kernels one after another accompanied by the necessary number of zeros. Convolutional neural networks were defined already in the 1980s [L⁺89], but they achieved wide recognition after the tremendous success of the convolutional neural network introduced in [KSH12] for image classification tasks. For more details on convolutional neural networks, we refer the reader to [GBC16, Chapter 9] and references therein.

Besides convolutional layers, popular choices are so-called *pooling layers*. Here, instead of cross-correlation, one computes maximum/minimum/average (max-, min-, average-pooling) on parts of the input which are of the size of the kernel. Observe that average-pooling is a convolutional layer with kernel inputs being reciprocal of kernel size. While the pooling operation decreases the dimensionality of the input, *upsampling* operation increases the dimensionality by using interpolation of the nearest already existing values. The other types of layers/neural networks we mention here are *residual networks* [HZRS16] where *skip connections* or *shortcuts* are used to jump over some layers, and *recurrent networks* [RHW86], [GBC16, Section 10] where output from a neuron is allowed to affect subsequent input to the same neuron, allowing the network to exhibit temporal dynamic behavior.

After discussing the structure of the weights, let us study the activation functions. Although they can also be learned through optimization (see e.g. [CP17]), they are usually fixed throughout the training process. A non-exhaustive table of most commonly used functions is given in [PRV21, Table 1]. One of the most commonly used functions, and in particular used by networks in this thesis, is (*leaky*) *rectification linear unit*

$$\text{ReLU}(x) := \max(0, x), \quad \text{LeakyReLU}_\alpha(x) := \max(x, \alpha x) = x + \text{ReLU}((\alpha - 1)x), \quad 0 \leq \alpha \leq 1. \quad (2.12)$$

The reasons for its success lay in the fact that it is cheap to compute, promotes sparsity in data representation [GBB11], and partially solves the problem of vanishing-gradients [BCV13]. What one immediately notices is that these functions are continuous, but not differentiable at $x = 0$. One possibility to circumvent this issue is to use *rectification power unit* $\text{RePU}(x) := \max(0, x^p)$, for $p > 1$. The other option is to approximate these functions smoothly. To this end, we define, for $\beta > 0$

$$\begin{aligned} \text{ReLU}_\beta(x) &:= \frac{1}{\beta} \log(1 + \exp(\beta x)), \\ \text{LeakyReLU}_{\alpha, \beta}(x) &:= \frac{1}{\beta} \log(\exp(\beta x) + \exp(\alpha \beta x)) = x + \text{ReLU}_\beta((\alpha - 1)x). \end{aligned} \quad (2.13)$$

The function ReLU_β is also known as *softplus* and it represents an approximation of ReLU for large values of β , since one has

$$\max(0, x) \leq \frac{1}{\beta} (\log(1 + \exp(\beta x)) \leq \max(0, x) + \frac{\log(2)}{\beta}, \quad \forall x \in \mathbb{R}.$$

So far, we observed neural networks defined on standard Euclidean space. In the second part of this thesis, we will make use of neural networks defined on a smooth Riemannian manifold (\mathcal{M}, g) of dimension m which is (smoothly) embedded in \mathbb{R}^n (cf. Definition 2.2.5). To this end, notice that we can write every $f : \mathcal{M} \rightarrow \mathbb{R}$ as $f = \sum_{i=1}^I f_i = \sum_{i=1}^I f \alpha_i$, where $\{\alpha_i\}_{i=1}^I$ is a smooth partition of unity with a compact support subordinated to a finite open cover $\{\mathcal{M}_i\}_{i=1}^I$ of \mathcal{M} . We suppose that every \mathcal{M}_i can be parametrized as the graph of a function over a subset of Eulerian coordinates, i.e.,

$$\mathcal{M}_i = \{F_i(x_{d_1}, \dots, x_{d_m}) = (\gamma_1(x_{d_1}, \dots, x_{d_m}), \dots, x_{d_1}, \dots, \gamma_n(x_{d_1}, \dots, x_{d_m})), (x_{d_1}, \dots, x_{d_m}) \in V_i \subset \mathbb{R}^m\},$$

where $\gamma_j : \mathbb{R}^m \rightarrow \mathbb{R}, j \in \{1, \dots, n\} \setminus \{d_1, \dots, d_m\}$ are smooth mappings, and $F_i : V_i \rightarrow \mathcal{M}_i, i = 1, \dots, I$ are smooth diffeomorphisms [Lee03, Theorem 4.5] and linear as they are inverse of orthogonal projection onto the coordinates $(x_{d_1}, \dots, x_{d_m})$ [Lee03, Corollary 4.13]. Then, a neural network Φ on \mathcal{M} is defined in the form

$$\Phi := \sum_{i=1}^I \Phi \alpha_i := \sum_{i=1}^I \Phi_i := \sum_{i=1}^I \tilde{\Phi}_i \circ F_i^{-1}, \quad (2.14)$$

where $\tilde{\Phi}_i$ is a neural network defined on V_i with mn neurons more than the network Φ_i (since F_i^{-1} is linear). One analogously defines $\tilde{f}_i := f \circ F_i : V_i \rightarrow \mathbb{R}$ for every $f : \mathcal{M} \rightarrow \mathbb{R}$. For more details see [BGKP19, Section 7] and references therein.

2.3.1 Sets of Functions as Realizations of Sets of Neural Networks

In this section, we study sets of functions that are defined as realizations of certain sets of neural networks. These sets are usually defined in terms of networks with upper bounds on the number of neurons and layers. Furthermore, besides the choice of activation functions, bounds on the norm of weights and biases are also seen to be of influence to properties of the resulting sets of realizations. To this end, one introduces the notion of the norm of the network by

$$\|\Phi\|_{NN} := \max_{l=1,\dots,L} \max_{\substack{i=1,\dots,n_l \\ j=1,\dots,n_{l-1} \\ c=1,\dots,c_l}} \max(|(A_l)_{i,j,c}|, |(b_l)_{i,c}|). \quad (2.15)$$

In [PRV21, Theorem 3.1], it is shown that the set of realizations of a set of neural networks defined on a closed cube in \mathbb{R}^d with an upper bound on the number of neurons and layers (at least 2) is not closed with respect to the L^p -norm for any $p \in [1, \infty)$ and most of the widely used activation functions including ReLU, LeakyReLU, and softplus (see [PRV21, Table 1] for other examples). However, if in addition to this upper bound on the architecture size, there is an upper bound on the norm $\|\cdot\|_{NN}$, one can show that for any continuous activation function the set of realizations of networks defined on any compact set is closed with respect to the L^p -, $p \in [1, \infty)$, and C^0 -norms [PRV21, Proposition 3.5]. This is due to the compactness of the set of networks and the continuity of the realization as a mapping from the set of neural networks with the $\|\cdot\|_{NN}$ -norm to the set of continuous functions with the standard norm [PRV21, Proposition 4.1], and can be further extended to the compactness in $W^{k,p}$ - and C^k -norms, if activation functions are in $C^k(\mathbb{R})$. As a part of this process, one can obtain bounds on derivatives of realizations in terms of the architecture of the underlying networks and their norms.

One of the crucial reasons for the success of neural networks is the fact that realizations of neural networks can approximate almost any arbitrary function. This is a question of *expressivity of neural networks*:

Given a function f and $\epsilon > 0$, is there a neural network Φ such that $\|R(\Phi) - f\| < \epsilon$?

This question needs additional details, in particular on the class of functions to which f belongs and the norm in which we measure the approximation. Furthermore, the naturally following question is the one on the size and the norm of the approximation network.

Some of the pioneering contributions to answer this question were given by Hornik *et al.* [HSW89, HSW90]. In [HSW89], they show that neural networks with two layers with *squashing* activation function $\sigma (\sigma : \mathbb{R} \rightarrow [0, 1]$, non-decreasing, $\lim_{x \rightarrow \infty} \sigma(x) = 1$, $\lim_{x \rightarrow -\infty} \sigma(x) = 0$) can approximate any continuous function in L^∞ -norm on any compact set and any measurable function arbitrarily well, in a metric based on an arbitrary probability measure. This result is often known as the *universal approximation theorem*. The most widely used example of squashing function is the *sigmoid* or *logistic* function $\sigma(x) = \frac{1}{1 + \exp(-x)}$. In [HSW90], they further study the approximation of derivatives of a function. To this end, for a given $l \in \mathbb{N}_0$, we say that σ is *l-finite* if $\sigma \in C^l(\mathbb{R})$ and $0 < \int_{\mathbb{R}} |D^l \sigma| < \infty$. They show that if σ is *l-finite*, then for all $0 \leq m \leq l$, $R((A_l, b_l, \sigma)_{l=1,2})$ is dense in $C^\infty(\mathbb{R}^{n_0})$ with respect to $C^m(K)$ -norm for any compact set $K \subset \mathbb{R}^{n_0}$. Here, $C^\infty(\mathbb{R}^{n_0}) \supset C_c^\infty(\mathbb{R}^{n_0})$ is defined as the set of all functions in $C^\infty(\mathbb{R}^{n_0})$ such that for all multi-indices α, β it holds $x^\beta D^\alpha f(x) \rightarrow 0$ as $|x| \rightarrow \infty$. Then, using the density of the set C_c^∞ , the same result holds for L^p and $W^{k,p}$ functions with $1 \leq p < \infty$, also on an open subset $U \subset \mathbb{R}^{n_0}$. Besides the logistic function, widely used example of *l-finite* function for any $l \in \mathbb{N}_0$ is $\tanh(x) = \frac{\exp(x) - \exp(-x)}{\exp(x) + \exp(-x)}$. Furthermore, we notice that ReLU_β and $\text{LeakyReLU}_{\alpha,\beta}$ are not 0-, nor 1-finite, but are *l-finite* for any $l \geq 2$.

The above results are based on the application of the Stone-Weierstrass Theorem and the Fourier transform and its properties, respectively. These abstract methods, however, do not provide any upper bound on the number of neurons of Φ nor a bound on the norm of the weights. There are more recent results that provide such upper bounds, including those for neural networks using ReLU as the activation function. In particular, [GKP20, Theorem 4.1] provides an upper bound on the number of layers, neurons, and non-zero weights and biases that are enough to construct a neural network with ReLU as the activation function, whose realization approximates a given function $f \in W^{k,p}((0, 1)^d)$ with precision ϵ in $W^{s,p}((0, 1)^d)$ -norm, with $k \geq 2, p \in [1, \infty], s \in [0, 1]$. The upper bound is given as a function of $\epsilon, d, k, p, s, \|f\|_{W^{k,p}}$. This was further extended in [GR21, Proposition 4.8], where an upper bound for the number of layers, neurons, non-zero weights, and the norm of the network, depending on the same parameter as above, is given for the approximation of functions $f \in W^{l,p}((0, 1)^d)$ in $W^{k,p}$ -norm, $l \leq k - 1$, with a neural network whose activation function is in C^l on some closed interval, with some additional conditions.

These conditions include many widely used activation functions [GR21, Table 1], including ReLU, RePU and ReLU_β . To obtain the above results, one starts with the approximation of polynomials by first approximating monomials and then applying (an approximate) partition of unity. This allows for the approximation of localized Taylor polynomials and the final step is an application of the Bramble–Hilbert lemma.

There are also results that study the approximation properties of convolutional neural networks. In [PV20], it is shown that every fully-connected neural network can be expressed by a convolutional neural network with a comparable number of parameters, without pooling and with padding making the input periodic, and the same holds the other way around. [Zho20] establishes the universality of deep convolutional neural networks with ReLU activation function, without pooling layers and with padding which preserves the shape. [Zho20, Theorem A] deals with the approximation of continuous functions on compact $K \subset \mathbb{R}^d$ in $C(K)$ -norm, while [Zho20, Theorem B] shows that such neural networks can approximate a Sobolev function of order $r > \frac{d}{2} + 2$ in the L^∞ -norm, with a bound on the number of layers depending on the precision of the approximation and the Sobolev norm of the function to be approximated.

Finally, let us notice that the above results transfer to neural networks defined on a compact Riemannian manifold (\mathcal{M}, g) (cf. (2.14)) if the metric g is equivalent to the metric induced by immersion, as the Lebesgue and Sobolev spaces are staying the same under an equivalent norm [Heb96, Chapter 2, Proposition 2.2].

Further results on approximations properties of different kinds of neural networks, together with the comparison of deep versus shallow networks can be found in [GRK20] as a chapter in [GK22].

2.4 Γ - and Mosco-convergence

Some of the main results in both parts of the thesis will concern the convergence of functionals of discrete or nonlocal nature to a functional which can be characterized as being of continuous or local nature. To this end, we define the precise notions of convergence that we will use, and present some of the properties of the convergent sequences. For further details, we refer the reader to [Mos69, DM93, Bra06].

Definition 2.4.1 (Γ - and Mosco-convergence). Let X be a Banach space. Consider functionals $\{\mathcal{E}^K\}_{K \in \mathbb{N}}$ and \mathcal{E} from X to $\mathbb{R} \cup \{\infty\}$ that satisfy

- (i) for every sequence $\{x^K\}_{K \in \mathbb{N}} \subset X$ with $x^K \rightarrow x \in X$, the estimate

$$\liminf_{K \rightarrow \infty} \mathcal{E}^K[x^K] \geq \mathcal{E}[x]$$

holds true (“lim inf-inequality”),

- (ii) for every $x \in X$ there exists a recovery sequence $\{x^K\}_{K \in \mathbb{N}} \subset X$ satisfying $x^K \rightarrow x$ in X such that the estimate

$$\limsup_{K \rightarrow \infty} \mathcal{E}^K[x^K] \leq \mathcal{E}[x]$$

is valid (“the existence of a recovery sequence and lim sup-inequality”).

Then $\{\mathcal{E}^K\}_{K \in \mathbb{N}}$ Γ -converges to \mathcal{E} . The strong convergence in the above conditions can be changed to the weak convergence, yielding the Γ -convergence in the weak topology. If solely the first condition is changed to

- (i') for every sequence $\{x^K\}_{K \in \mathbb{N}} \subset X$ with $x^K \rightharpoonup x \in X$ it holds

$$\liminf_{K \rightarrow \infty} \mathcal{E}^K[x^K] \geq \mathcal{E}[x],$$

we say that $\{\mathcal{E}^K\}_{K \in \mathbb{N}}$ converges to \mathcal{E} in the sense of Mosco.

Our interest lies in determining some reasonable conditions under which we can obtain the equality in the lim inf–inequality conditions, at least for the sequence of minimizers of the functionals. This is given by the following definition and theorem.

Definition 2.4.2 (Equicoercive sequence, [Bra06, Definition 2.9]). Let X be a Banach space. A sequence of functionals $\{\mathcal{E}^K\}_{K \in \mathbb{N}}$ mapping from X to $\mathbb{R} \cup \{\infty\}$ is said to be *equicoercive* if for all $r \in \mathbb{R}$ there exists a compact set $K_r \subset X$ such that $\{x \in X : \mathcal{E}^K[x] \leq r \forall K \in \mathbb{N}\} \subset K_r$.

Theorem 2.4.3 (Fundamental theorem of Γ -convergence, [Bra06, Theorem 2.10]). *Let X be a Banach space and $\{\mathcal{E}^K\}_{K \in \mathbb{N}} : X \rightarrow \mathbb{R} \cup \{\infty\}$ an equicoercive sequence which Γ -converges to $\mathcal{E} : X \rightarrow \mathbb{R} \cup \{\infty\}$. Then $\min_{x \in X} \mathcal{E}[x] = \lim_{K \rightarrow \infty} \inf_{x \in X} \mathcal{E}^K[x]$.*

In particular, for a sequence $\{x^K\}_{K \in \mathbb{N}}$ of minimizers of $\{\mathcal{E}^K\}_{K \in \mathbb{N}}$ with uniformly bounded energy $\mathcal{E}^K[x^K] \leq r$, the equicoercivity implies convergence on a subsequence to some $x \in X$, and the Γ -convergence implies that x is a minimizer of \mathcal{E} . In the case of Mosco-convergence, one can replace the condition of compactness of the set K_r to the weak compactness. This will be the main idea of the approach we will take in Theorems 4.4.4, 5.4.5 and 6.3.2.

Chapter 3

Foundations of the Metamorphosis Model

IN this chapter, we first present a foundation for the study of geodesic and spline interpolation in the metamorphosis model. To this end, based on [RW15, HRW18], we introduce the framework for the study of both the time continuous and time discrete geodesic and spline calculus on possibly infinite-dimensional Riemannian manifolds. Furthermore, we observe the space of images as a Riemannian manifold and introduce the image metamorphosis model. Finally, we introduce the iPALM algorithm [PS16] that will be used for the numerical optimization of fully discrete models which are developed in later chapters.

3.1 Riemannian Geodesics and Splines

In this section, we define the basic setup for studying geodesics and splines on possibly infinite-dimensional Riemannian manifolds. We first study the time continuous setting and then define a variational time discretization. We present results showing the existence (and uniqueness) of the defined notions, and the convergence of the time discrete notions towards the continuous counterparts using the tool of Γ -convergence.

3.1.1 Time Continuous Setting

In Section 2.2, we introduced geodesics as one of the central concepts in the study of Riemannian manifolds in two different, but equivalent ways: as solutions to the geodesic equation (*cf.* Definition 2.2.10) and as minimizers of the path length for fixed boundary conditions (*cf.* Proposition 2.2.14). In this section, we introduce the basic framework for the study of possibly infinite dimensional Riemannian manifolds which will allow us to define the notion of a geodesic path on these objects.

To define the technical setup (*cf.* [RW15] for details), we observe \mathcal{M} as the closure of an open connected subset of a separable, reflexive Banach space \mathbf{V} that is compactly embedded in a Banach space \mathbf{Y} . This makes \mathcal{M} a Banach manifold. We also assume that \mathcal{M} is path-connected and that $\partial\mathcal{M}$ (if it is non-empty) is smooth. We identify the tangent space $T_x\mathcal{M}$, $x \in \mathcal{M}$ with \mathbf{V} and introduce a Riemannian metric as a mapping $g : \mathcal{M} \times \mathbf{V} \times \mathbf{V} \rightarrow \mathbb{R}$, such that g is uniformly bounded, \mathbf{V} -coercive, and uniformly continuous on \mathcal{M} . In particular, these conditions are satisfied for smooth m -dimensional manifolds, with $\mathbf{V} = \mathbf{Y} = \mathbb{R}^m$ and $g_x(v, w) = v^T D\varphi(x)^T D\varphi(x)w$, where φ is a local chart (*cf.* Definition 2.2.2). Furthermore, the coordinate-free definition of covariant derivative (Definition 2.2.17) is valid in this setting.

Given a sufficiently smooth path $(y(t))_{t \in [0,1]}$ on a Riemannian manifold (\mathcal{M}, g) , we define its length (*cf.* Definition 2.2.13)

$$\mathcal{L}[(y(t))_{t \in [0,1]}] := \int_0^1 \sqrt{g_{y(t)}(\dot{y}(t), \dot{y}(t))} dt, \quad (3.1)$$

and the path energy (cf. Definition 2.2.15)

$$\mathcal{E}[(y(t))_{t \in [0,1]}] := \int_0^1 g_{y(t)}(\dot{y}(t), \dot{y}(t)) dt. \quad (3.2)$$

Recall that in Remark 2.2.16 we discussed the fact that a minimizer of the path energy is also a minimizer of the path length. As the path energy is convex, it is, therefore, more convenient to define a *geodesic* connecting y_A and y_B as a minimizer of the path energy among all (regular) paths $y : [0, 1] \rightarrow \mathcal{M}$ with $y(0) = y_A$ and $y(1) = y_B$. For all $y_A, y_B \in \mathcal{M}$, the existence of a connecting geodesic $y \in H^1((0, 1), \mathcal{M})$ was shown in [RW15, Theorem 4.1], where the key step is the weak lower semicontinuity of the path energy (3.2), using the compact embedding of \mathbf{V} into \mathbf{Y} . The local uniqueness of a geodesic is shown under an additional assumption that g is twice continuously differentiable [RW15, Theorem 4.2]. Given the geodesic curve $(y(t))_{t \in [0,1]}$, the Riemannian distance between points y_A and y_B can be defined by $d_{\mathcal{M}}(y_A, y_B) := \sqrt{\mathcal{E}[y(t)]}$.

Geodesics offer a natural and smooth interpolation between two fixed points on a manifold. The next task one considers is interpolation between multiple points $p_1, \dots, p_J \in \mathcal{M}$ for a sequence of specified times $0 = t_1 < t_2 < \dots < t_J = 1$, i.e., finding an as smooth as possible curve $y : [0, 1] \rightarrow \mathcal{M}$ that fulfills the *interpolation constraints* $y(t_j) = p_j$ for $j = 1, \dots, J$. The first attempt would be to construct a path consisting of geodesics between every two consecutive points. However, the velocity field of such a curve exhibits discontinuities at the interpolation points (cf. Figure 1.2). In general, there is no curve solving the geodesic equation (2.3) which satisfies interpolation constraints. An approach to the construction of a smooth interpolation is to use the geodesic equation, not as a strict constraint, but to include it as a penalization term in a path energy. This leads to the functional [NHP89]

$$\mathcal{F}[(y(t))_{t \in [0,1]}] := \int_0^1 g_{y(t)} \left(\frac{D}{dt} \dot{y}(t), \frac{D}{dt} \dot{y}(t) \right) dt, \quad (3.3)$$

where $\frac{D}{dt}$ denotes the covariant derivative along y as defined by Definition 2.2.17. In the Euclidean context, this functional takes the simple form $\mathcal{F}_{\text{Euc}}[(y(t))_{t \in [0,1]}] = \int_0^1 |\dot{y}(t)|^2 dt$, and it was shown in [dB63] that the unique minimizer (with specified boundary conditions) is a third order spline, i.e., a twice continuously differentiable curve which is a piecewise cubic polynomial. This motivates us to call the above-defined energy *spline energy* and its minimizer *Riemannian (cubic) spline*.

For the rigorous study of Riemannian splines, one considers a Riemannian manifold (\mathcal{M}, g) where $\mathcal{M} = \mathbf{V}$ for a separable Hilbert space compactly embedded in a Banach space \mathbf{Y} (cf. [HRW18, Definition 2.1]) and a uniformly coercive Riemannian metric g which can be written as a sum of a compact and a bilinear part (cf. [HRW18, Definition 2.2]). These assumptions still do not guarantee the existence of a Riemannian spline (cf. [HRW18, Lemma 2.15] for an example where \mathcal{M} is cylinder in \mathbb{R}^3 , $t_1 = 0, t_2 = r \in (0, 1) \setminus \mathbb{Q}, t_3 = 1$, and $p_1, p_2, p_3 \in \mathcal{M}$, where p_1 and p_2 are opposite to each other on S^1). To this end, one instead observes *regularized spline energy*

$$\mathcal{F}^\sigma[(y(t))_{t \in [0,1]}] := \mathcal{F}[(y(t))_{t \in [0,1]}] + \sigma \mathcal{E}[(y(t))_{t \in [0,1]}], \quad \sigma > 0. \quad (3.4)$$

We refer to minimizers of this energy as *regularized splines* or *splines in tension* [Sch66]. The existence of regularized splines in space $H^2((0, 1), \mathcal{M})$, satisfying arbitrary interpolation constraints, was shown in [HRW18, Theorem 2.19], where one also considers one of the following boundary conditions

$$\begin{aligned} \frac{D}{dt} \dot{y}(0) = \frac{D}{dt} \dot{y}(1) = 0, & \quad (\text{natural b.c.}) \\ \dot{y}(0) = v_0, \quad \dot{y}(1) = v_1 \quad \text{for given } v_0 \in T_{y(0)}\mathcal{M}, v_1 \in T_{y(1)}\mathcal{M}, & \quad (\text{Hermite b.c.}) \\ y(0) = y(1), \quad \dot{y}(0) = \dot{y}(1), \quad \ddot{y}(0) = \ddot{y}(1). & \quad (\text{periodic b.c.}) \end{aligned} \quad (3.5)$$

The natural boundary conditions correspond to the case of not posing boundary conditions at all [HRW18, Remark 2.13], while periodic boundary conditions can be observed as if we define the problem on S^1 [HRW18, Remark 2.14].

3.1.2 Variational Time Discretization

The computation of geodesic and spline interpolations on infinite-dimensional Riemannian manifolds represents a numerically involved task, as working directly with definitions (3.2) and (3.4) is possible only in rare cases.

To this end, one discretizes these expressions both in time and space. In this section, we present a framework for a variational time discretization. For a smooth path on the manifold $y : [0, 1] \rightarrow \mathcal{M}$ we observe the points $(y_0, \dots, y_K) \in \mathcal{M}^{K+1}$, $K \geq 1$ obtained by sampling the path y at the times $t_k = \frac{k}{K}$, $k = 0, \dots, K$, i.e., $y_k = y(t_k)$. Taking into account the triangle inequality and Jensen's inequality, we have

$$\mathcal{L}[(y(t))_{t \in [0,1]}] \geq \sum_{k=1}^K d_{\mathcal{M}}(y_{k-1}, y_k), \quad \mathcal{E}[(y(t))_{t \in [0,1]}] \geq K \sum_{k=1}^K (d_{\mathcal{M}}(y_{k-1}, y_k))^2,$$

where equality holds for geodesic paths due to the constant speed property. As already said, a direct computation of the Riemannian distance is usually not possible. The main step in variational time discretization is an assumption on the existence of a functional $\mathcal{W} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ that locally approximates the squared Riemannian distance up to higher order terms

$$\mathcal{W}[y, \tilde{y}] = (d_{\mathcal{M}}(y, \tilde{y}))^2 + \mathcal{O}((d_{\mathcal{M}}(y, \tilde{y}))^3), \quad \text{for } y, \tilde{y} \in \mathcal{M}. \quad (3.6)$$

These observations motivate the definitions of the *discrete path length* and the *discrete path energy* associated with a *discrete* $(K+1)$ -path $(y_0, \dots, y_K) \in \mathcal{M}^{K+1}$ as

$$\mathbf{L}^K[(y_0, \dots, y_K)] := \sum_{k=1}^K \sqrt{\mathcal{W}[y_{k-1}, y_k]}, \quad \mathbf{E}^K[(y_0, \dots, y_K)] := K \sum_{k=1}^K \mathcal{W}[y_{k-1}, y_k]. \quad (3.7)$$

Then, a *time discrete geodesic of length* $K+1$ connecting y_A and y_B in \mathcal{M} is defined as a discrete path (y_0, \dots, y_K) that minimizes the discrete path energy (3.7) among all the paths with $y_0 = y_A, y_K = y_B$. For fixed $y_A, y_B \in \mathcal{M}$ and $K \geq 2$, the existence of discrete geodesic of length $K+1$ connecting y_A and y_B follows from [RW15, Theorem 4.3]. If \mathcal{W} is in addition twice differentiable, weakly lower semicontinuous, and coercive, and the approximation of the squared Riemannian distance is uniform locally, this geodesic is locally unique [RW15, Theorem 4.7].

In consistency with the time continuous approach, we defined geodesic paths as minimizers of the discrete path energy. In the time continuous setting, by Remark 2.2.16, minimization of path energy (3.2) was implying the minimization of the curve length energy (3.1). In the discrete setting, however, this is not the case as one can construct a counterexample on a manifold $\mathcal{M} \subset \mathbf{V}$ with two points $y_A, y_B \in \mathcal{M}$ which are close w.r.t. the norm induced by \mathbf{V} , but relatively far apart w.r.t. the Riemannian distance on \mathcal{M} (cf. [RW15, Remark 4.11]). However, one can show convergence of time discrete geodesics towards the time continuous geodesic curve as $K \rightarrow \infty$. To this end, one defines a time extension of a discrete path (y_0, \dots, y_K) constructed by piecewise geodesic interpolation. Furthermore, a time extension \mathcal{E}^K is defined such that $\mathcal{E}^K[y] = \mathbf{E}^K[(y_0, \dots, y_K)]$ if y is the time extension of a $(K+1)$ -path (y_0, \dots, y_K) , and $+\infty$ else. Then, one can show the Γ -convergence in $L^2((0, 1), \mathbf{Y})$ (cf. Definition 2.4.1) of \mathcal{E}^K to the time continuous path energy \mathcal{E} , and the convergence in $L^2((0, 1), \mathbf{Y})$ of a subsequence of extensions of discrete geodesics to the time continuous geodesic curve [RW15, Theorem 4.8 and Corollary 4.9].

To study the time discretization of Riemannian splines, let us again consider the Euclidean setting. Given a sufficiently smooth curve $(y(t))_{t \in [0,1]} \in \mathbb{R}^n$, using an approximation of acceleration $\ddot{y}(t)$ by the second order finite differences, we have $\ddot{y}(t_k) \approx 2K^2 \left(\frac{y(t_{k-1}) + y(t_{k+1})}{2} - y(t_k) \right)$. Using the piecewise constant approximation of the time integral, we then get

$$\mathcal{F}[(y(t))_{t \in [0,1]}] = \int_0^1 |\ddot{y}(t)|^2 dt \approx \sum_{k=1}^{K-1} 4K^3 \left| y(t_k) - \frac{y(t_{k-1}) + y(t_{k+1})}{2} \right|^2 dt,$$

where for $k = 0, \dots, K$ we have $t_k = \frac{k}{K}$. Having in mind the above approximation of a geodesic path, in order to translate this approximation to the discretization of the Riemannian spline energy (3.3), we can replace the Euclidean distance and geodesic average by (an approximation of) their Riemannian counterparts. This leads to the definition of *discrete spline energy* for a discrete path $(y_0, \dots, y_K) \in \mathcal{M}^{K+1}$ as

$$\mathbf{F}^K[(y_0, \dots, y_K)] := 4K^3 \sum_{k=1}^{K-1} \mathcal{W}[y_k, \tilde{y}_k], \quad (3.8)$$

subject to the constraint

$$\tilde{y}_k := \operatorname{argmin}_{y \in \mathcal{M}} \mathbf{E}^2[(y_{k-1}, y, y_{k+1})], \quad \text{for } k = 1, \dots, K-1.$$

As in the time continuous case (*cf.* (3.4)), we also define the regularized discrete spline energy $\mathbf{F}^{\sigma,K} := \mathbf{F}^K + \sigma \mathbf{E}^K$ for $\sigma > 0$. Given a sequence of points p_1, \dots, p_J and a sequence of indices $0 = i_1 < \dots < i_J = K$ we say that a discrete path $(y_0, \dots, y_K) \in \mathcal{M}^{K+1}$ is a *regularized (time) discrete spline interpolation* of p_1, \dots, p_J if it minimizes energy $\mathbf{F}^{\sigma,K}$ while satisfying interpolation constraints $y_{i_j} = p_j, j = 1, \dots, J$. The existence of a regularized discrete spline interpolation for the arbitrary interpolation constraints was shown in [HRW18, Theorem 3.10], under certain regularity and consistency conditions for \mathcal{W} (see [HRW18, Definition 3.1]). Here, one considers one of the following boundary constraints

$$\begin{aligned} & \text{no additional constraint, (natural b.c.)} \\ K(y_1 - y_0) = v_0, \quad K(y_K - y_{K-1}) = v_1 \quad & \text{for given } v_0, v_1 \in \mathbf{V}, \quad \text{(Hermite b.c.)} \\ y_0 = y_K, \quad y_{K+1} \sim y_1, \quad y_{-1} \sim y_{K-1}. & \quad \text{(periodic b.c.)} \end{aligned}$$

This identification of boundary points means that the summation in (3.8) is in range from $k = 0$ to $k = K - 1$.

To study the convergence properties as $K \rightarrow \infty$, one defines time extension of a $(K + 1)$ -path by the cubic Hermite interpolation on intervals $[(k - \frac{1}{2})\tau, (k + \frac{1}{2})\tau]$, for $\tau = \frac{1}{K}$ and $k = 1, \dots, K - 1$, and the affine interpolation on $[0, \frac{1}{2}\tau]$ and $(K - \frac{1}{2})\tau, 1]$. Then, one can define the time extension $\mathcal{F}^{\sigma,K}$ as being equal to $\mathbf{F}^{\sigma,K}$ on these extended curves and $+\infty$ elsewhere. [HRW18, Theorem 4.7] provides Mosco-convergence in $H^2((0, 1), \mathcal{M})$ of $\{\mathcal{F}^{\sigma,K}\}_{K \in \mathbb{N}}$ to \mathcal{F}^σ as $K \rightarrow \infty$. As a corollary, we have that any sequence of extensions of minimizers of $\mathbf{F}^{\sigma,K}$ contains a sequence that converges weakly in $H^2((0, 1), \mathcal{M})$ to a minimizer of \mathcal{F}^σ [HRW18, Theorem 4.9].

Finally, let us mention that the variational time discretization of some of the other notions introduced in Section 2.2 (exponential and logarithm map, parallel transport) was also studied in [RW15], but is out of the scope of this thesis.

3.2 Image Metamorphosis Model

In this section, we give a review of the image metamorphosis model, the flow of diffeomorphism model as its basis, and finally state some fundamental results on the flow equation that will be used throughout the first part of this thesis. Here and throughout the first part of this thesis, we assume that the domain $\Omega \subset \mathbb{R}^n$ is bounded and strongly Lipschitz (*cf.* Definition 2.1.1).

3.2.1 Flow of Diffeomorphism

In what follows, we present a short exposition of the flow of diffeomorphism model [DGM98, BMTY05, JM00, MTY02] as a primary stage of the metamorphosis model. This model is based on Arnold's approach [Arn66, AK98] to the study of the temporal evolution of ideal fluids by defining geodesics on the group of orientation preserving diffeomorphisms. To this end, we observe the set of *Sobolev diffeomorphisms*

$$\mathcal{D}^m(\Omega) := \{\psi \in H^m(\Omega, \Omega), \det(D\psi) > 0, \psi|_{\partial\Omega} = \mathbf{1}\}, \quad m > 1 + \frac{n}{2}. \quad (3.9)$$

Functions in this set are indeed $C^1(\Omega, \Omega)$ -orientation preserving diffeomorphisms since, by [IKT13, Lemma A.1], we have the existence of $\psi^{-1} \in \mathcal{D}^m(\Omega)$ and, by Sobolev embedding theorem 2.1.8, we have $\mathcal{D}^m(\Omega) \subset C^1(\bar{\Omega}, \bar{\Omega})$. By [IKT13, Lemma 2.18], the set of all $H^m(\Omega, \Omega)$ functions which are orientation preserving diffeomorphisms is an open subset of $H^m(\Omega, \Omega)$ and can be seen as C^∞ -manifold modeled on this Hilbert space. For every $\psi \in \mathcal{D}^m(\Omega)$, we have $T_\psi \mathcal{D}^m(\Omega) = \mathcal{V} := H^m(\Omega, \mathbb{R}^n) \cap H_0^1(\Omega, \mathbb{R}^n)$, as the velocity space of smooth curves with origin at that point (*cf.* Definition 2.2.3). We define the Riemannian metric by

$$g_{\psi_t}(\dot{\psi}_t, \dot{\psi}_t) := \int_{\Omega} L[v_t, v_t] \, dx,$$

and the associated path energy is

$$\mathcal{E}_{\psi_t}[(\psi_t)_{t \in [0,1]}] := \int_0^1 g_{\psi_t}(\dot{\psi}_t, \dot{\psi}_t) \, dt,$$

where $v_t = \dot{\psi}_t \circ \psi_t^{-1}$ is the *Eulerian velocity* and the quadratic form L is a higher order elliptic operator to be specified below. Following [DGM98, Theorem 2.5], given that the metric g_{ψ_t} is $H^3(\Omega)$ -coercive, paths with finite energy which connect two diffeomorphisms $\psi_0 = \psi_A$ and $\psi_1 = \psi_B$ are actually *one-parameter families of diffeomorphisms*, i.e., $\psi : [0, 1] \times \Omega \rightarrow \Omega$ such that for all $s, t \in [0, 1]$ with $s + t \in [0, 1]$ the map ψ_s is diffeomorphic and $\psi_s \circ \psi_t = \psi_{s+t}$. Moreover, by [DGM98, Theorem 3.1], any minimizing sequence of paths has a uniformly converging subsequence and the limiting minimizing path ψ , which is the geodesic path on $\mathcal{D}^m(\Omega)$, and its associated velocity field v satisfies $\dot{\psi}_t(x) = v(t, \psi(t, x))$ for every $t \in [0, 1]$ and $x \in \Omega$.

In the context of image morphing, the above minimizer determines the geodesic path $u_t = u_A \circ \psi_t^{-1}$, $t \in [0, 1]$ in the space of images $L^2(\Omega)$, connecting $u_A = u_0$ and $u_B = u_1 \circ \psi_B$.

For the rest of this chapter, the quadratic form L is the higher order elliptic operator

$$L[v, v] := \frac{\lambda}{2}(\text{tr } \varepsilon[v])^2 + \mu \text{tr}(\varepsilon[v]^2) + \gamma |D^m v|^2, \quad (3.10)$$

where $\varepsilon[v] := (Dv)^{\text{sym}} = \frac{1}{2}(Dv + Dv^T)$ and $\lambda, \mu, \gamma > 0$. The $H^3(\Omega)$ -coercivity of g_{ψ_t} is inferred by combining Korn's inequality 2.1.6 for the lower order terms with the Gagliardo–Nirenberg inequality 2.1.5 for the higher order term. Physically, the metric $g_{\psi_t}(\dot{\psi}_t, \dot{\psi}_t)$ describes the viscous dissipation in a multipolar fluid model as investigated by Nečas and Šilhavý [NŠ91]. The first two terms of the integrand represent the dissipation density in a Newtonian fluid and the third term can be regarded as a higher order measure of friction. Furthermore, the spatially integrated terms $\int_{\Omega} (\text{tr } \varepsilon[v])^2 dx$ and $\int_{\Omega} \text{tr}(\varepsilon[v]^2) dx$ roughly quantify the average change of volume and length induced by v , respectively (cf. [Cia88]). Let us also notice that, due to sole dependency on the symmetric part of the gradient in the lower order terms, a rotational flow, whose associated velocity field has a skew-symmetric gradient, does not lead to any friction, while for incompressible fluids we have $\text{tr } \varepsilon[v] = \text{div } v = 0$ so the first term in (3.10) vanishes. For more on the physical background, we refer the reader to [RW11].

3.2.2 Metamorphosis in Image Space

In the flow of diffeomorphism model, we observed image paths determined by an initial image and a family of diffeomorphisms. In this setting, each image particle moves along the motion paths without any change of intensity, which is known as the *brightness constancy assumption*. This assumption is equivalent to a vanishing material derivative

$$\hat{z}_t \circ \psi_t := \frac{d}{dt}(u_t \circ \psi_t) = (\dot{u} + v \cdot Du) \circ \psi_t := \frac{D}{\partial t} u_t \circ \psi_t$$

along a path $(u_t)_{t \in [0, 1]}$ in the space of images. The metamorphosis approach, originally proposed by Miller, Trounev, Younes and coworkers in [MY01, TY05b, TY05a], generalizes the flow of diffeomorphism model by defining a metric on the space of images that penalizes the squared material derivative and thus allows for image intensity variations along motion paths. Under the assumption that the image path u is sufficiently smooth, the metric and the path energy read as

$$g_u(\dot{u}, \dot{u}) := \min_{v: \Omega \rightarrow \mathbb{R}^n} \int_{\Omega} L[v, v] + \frac{1}{\delta} |\hat{z}|^2 dx, \quad \mathcal{E}[u] := \int_0^1 g_u(\dot{u}_t, \dot{u}_t) dt,$$

for a penalization parameter $\delta > 0$. The term $L[v, v]$ quantifies the cost of the underlying transport, while the penalization of the material derivative \hat{z} reflects the penalization of the changes in image intensity along the trajectories. Hence, the flow of diffeomorphism model can be seen as the limit case of the metamorphosis model for $\delta \rightarrow 0$.

This definition of the metric has two major drawbacks:

- In general, paths in the space of images do not exhibit any smoothness properties (neither in space nor time), and therefore the evaluation of the material derivative is not well-defined.
- Different pairs (v, \hat{z}) of velocity fields and material derivatives can imply the same time derivative of the image path \dot{u} .

To tackle these problems, Trounev and Younes [TY05a] proposed a nonlinear geometric structure in the space of RGB images $\mathcal{I} := L^2(\Omega, \mathbb{R}^3)$. In detail, for a given image path $u \in L^2([0, 1], \mathcal{I})$ and an associated velocity

field $v \in L^2((0, 1), \mathcal{V})$, where $\mathcal{V} := H^m(\Omega, \mathbb{R}^n) \cap H_0^1(\Omega, \mathbb{R}^n)$ denotes the velocity space, the *weak material derivative* $\hat{z} \in L^2((0, 1), \mathcal{Z} := L^2(\Omega, \mathbb{R}^3))$ is incorporated in the model and it is implicitly given by

$$\int_0^1 \int_{\Omega} \eta \hat{z} \, dx \, dt = - \int_0^1 \int_{\Omega} (\partial_t \eta + \operatorname{div}(v\eta)) u \, dx \, dt, \quad \forall \eta \in C_c^\infty([0, 1] \times \bar{\Omega}). \quad (3.11)$$

We consider (v, \hat{z}) as a tangent vector in the tangent space of \mathcal{I} at the image u . Indeed, (v, \hat{z}) represents a variation of the image u via transport and change of intensity. Furthermore, for an image $u \in \mathcal{I}$ we define the nonlinear structure on $\mathcal{V} \times \mathcal{Z}$ by

$$N_u := \left\{ (v, \hat{z}) \in \mathcal{V} \times \mathcal{Z} : \int_{\Omega} \hat{z} \eta + u \operatorname{div}(\eta v) \, dx = 0, \quad \forall \eta \in C_c^\infty([0, 1] \times \bar{\Omega}) \right\}.$$

This allows a restriction to the equivalence classes of pairs, where two pairs are equivalent if and only if they induce the same temporal change of the image path \dot{u} . Then the tangent space at u is defined as

$$T_u \mathcal{I} := \mathcal{V} \times \mathcal{Z} \setminus N_u,$$

giving rise to the associated tangent bundle $T\mathcal{I}$ (cf. Definition 2.2.3). The elements in this tangent space are denoted by $\gamma = \overline{(v, \hat{z})}$ and the space is endowed with the norm

$$\|\gamma\| := \inf \{ \|v\|_{\mathcal{V}} + \|\hat{z}\|_{\mathcal{Z}} : \gamma = \overline{(v, \hat{z})} \}.$$

Finally, the image curve u is defined to be *regular* in the space of images, denoted by $u \in H^1([0, 1], \mathcal{I})$, if $u \in C^0([0, 1], \mathcal{I})$ and exists a measurable path $\gamma : [0, 1] \rightarrow T\mathcal{I}$ with $\gamma_t \in T_{u_t} \mathcal{I}$, $t \in [0, 1]$, such that $\int_0^1 \|\gamma_t\|^2 \, dt < \infty$ and (3.11) is satisfied. The path energy can now be rigorously defined on $u \in H^1([0, 1], \mathcal{I})$ by

$$\mathcal{E}[u] := \int_0^1 \inf_{(v, \hat{z}) \in T_{u(t)} \mathcal{I}} \int_{\Omega} L[v, v] + \frac{1}{\delta} |\hat{z}|^2 \, dx \, dt, \quad (3.12)$$

and the *geodesic curve in the space of images* connecting $u_A, u_B \in \mathcal{I}$ is a regular curve $u \in H^1([0, 1], \mathcal{I})$ with $u_0 = u_A, u_1 = u_B$ such that

$$\mathcal{E}[u] := \inf \{ \mathcal{E}[\tilde{u}] : \tilde{u} \in H^1([0, 1], \mathcal{I}), \tilde{u}_0 = u_A, \tilde{u}_1 = u_B \}.$$

The existence of a geodesic curve is proven in [TY05a, Theorem 6]. Note that the infimum in (3.12) is attained along a measurable curve, which is shown in [TY05a, Proposition 1 and Theorem 2]. For further details, we refer the reader to [TY05a].

3.2.3 The Flow Equation

The relationship between diffeomorphic flows and their Eulerian velocity was of key importance for the definition of both the flow of diffeomorphism and image metamorphosis models. In this subsection, we state some results on the existence, uniqueness, and regularity of diffeomorphic flows given as a solution to the flow equation, to be used throughout this part of the thesis.

Theorem 3.2.1. *For every velocity field $v \in L^2((0, 1), \mathcal{V})$, there exists a unique flow $\psi \in H^1([0, 1], H^m(\Omega, \mathbb{R}^n))$ which is a solution of*

$$\begin{aligned} \dot{\psi}_t(x) &= v_t \circ \psi_t(x), \\ \psi_0(x) &= x, \end{aligned} \quad (3.13)$$

in the sense that $\psi_t(x) = x + \int_0^t v_s \circ \psi_s(x) \, ds$ for all $x \in \Omega$ and $t \in [0, 1]$. In particular, $\psi_t \in \mathcal{D}^m(\Omega)$ for all $t \in [0, 1]$. Furthermore, for $\alpha \in [0, m - 1 - \frac{n}{2}]$ the following estimate holds

$$\|\psi\|_{C^0([0, 1], C^{1, \alpha}(\bar{\Omega}))} + \|\psi^{-1}\|_{C^0([0, 1], C^{1, \alpha}(\bar{\Omega}))} \leq G \left(\int_0^1 \|v_s\|_{C^{1, \alpha}(\bar{\Omega})} \, ds \right), \quad (3.14)$$

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

Proof. The existence and regularity follow from [BV17, Theorem 4.4 and Lemma 3.5] by an additional observation that there exists a linear and continuous extension operator from $H^m(\Omega, \mathbb{R}^n)$ to $H^m(\mathbb{R}^n, \mathbb{R}^n)$ which is implied by Stein's extension theorem [Ste70]. The uniqueness follows from [You10, Theorem C.3.] which is a generalization of Picard–Lindelöf Theorem. The weak continuity follows from [TY05a, Theorem 9]. Finally, the estimate for the first term in (3.14) is proven in [ENR20, Theorem 6] based on [TY05a, Lemma 7] and relies on Gronwall's inequality, while the second term is estimated analogously by observing that ψ_t^{-1} is the flow associated with the motion field $-v_{1-t}$. \square

Remark 3.2.2. Analogous existence results and bounds as in Theorem 3.2.1 hold if \mathcal{V} is replaced by $C^{1,\alpha}(\bar{\Omega}, \mathbb{R}^n)$ with zero boundary condition [You10, Chapter 8]. In addition, the mapping $v \mapsto \psi^v$ is Lipschitz continuous in v and it holds

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

where $C = \int_0^t \|v_s\|_{C^1(\bar{\Omega})} ds$ [You10, (8.16)].

Remark 3.2.3. By the Sobolev embedding theorem 2.1.8, we have $\psi \in C^{0,\frac{1}{2}}([0,1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$ and we can obtain estimates for $|\psi|_{C^{0,\frac{1}{2}}([0,1], C^{1,\alpha}(\bar{\Omega}))}$ and $|\psi^{-1}|_{C^{0,\frac{1}{2}}([0,1], C^{1,\alpha}(\bar{\Omega}))}$:

$$\begin{aligned} \|\psi_t - \psi_s\|_{C^{1,\alpha}(\bar{\Omega})} &= \left\| \int_s^t v_r \circ \psi_r dr \right\|_{C^{1,\alpha}(\bar{\Omega})} \\ &\leq \int_s^t \|v_r \circ \psi_r\|_{C^{1,\alpha}(\bar{\Omega})} dr \\ &\leq \int_s^t \|v_r\|_{C^0(\bar{\Omega})} + C \left(\|Dv_r\|_{C^0(\bar{\Omega})} \|D\psi_r\|_{C^{0,\alpha}(\bar{\Omega})} + |Dv_r|_{C^{0,\alpha}(\bar{\Omega})} \|D\psi_r\|_{C^{0,\alpha}(\bar{\Omega})}^{\alpha+1} \right) dr \\ &\leq \int_s^t \|v_r\|_{C^0(\bar{\Omega})} dr + C \max \left(\|D\psi\|_{C^0([0,1], C^{0,\alpha}(\bar{\Omega}))}^{\alpha+1}, \|D\psi\|_{C^0([0,1], C^{0,\alpha}(\bar{\Omega}))} \right) \int_s^t \|Dv_r\|_{C^{0,\alpha}(\bar{\Omega})} dr \\ &\leq G_1 \left(\int_0^1 \|v_s\|_{C^{1,\alpha}(\bar{\Omega})} ds \right) \int_s^t \|v_r\|_{C^{1,\alpha}(\bar{\Omega})} dr \\ &\leq G_1 \left(\int_0^1 \|v_s\|_{C^{1,\alpha}(\bar{\Omega})} ds \right) \|v\|_{L^2((0,1), C^{1,\alpha}(\bar{\Omega}))} |t - s|^{\frac{1}{2}}, \end{aligned}$$

where we used the previous theorem to obtain a positive, monotonically increasing function G_1 , and in the last step we used the Cauchy-Schwarz inequality. Furthermore, to obtain the third line, we used that, by [Fio16, Propositions 1.2.4 and 1.2.7], for every $r \in [0, 1]$ we have $v_r \circ \psi_r \in C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$ with

$$|D(v_r \circ \psi_r)|_{C^{0,\alpha}(\bar{\Omega})} \leq C \left(\|Dv_r\|_{C^0(\bar{\Omega})} |D\psi_r|_{C^{0,\alpha}(\bar{\Omega})} + |Dv_r|_{C^{0,\alpha}(\bar{\Omega})} \|D\psi_r\|_{C^{0,\alpha}(\bar{\Omega})}^{\alpha+1} \right). \quad (3.15)$$

3.3 iPALM Algorithm

In this section, we introduce (a simplified version of) the Inertial Proximal Alternating Linearized Minimization (iPALM) algorithm [PS16] as the basis of the optimization algorithm for the fully discrete models in Sections 4.6 and 5.7.

We observe the following optimization problem

$$\min_{x \in \mathbb{R}^n} F(x) = F_1(x) + F_2(x), \quad (3.16)$$

where

- $F_1, F_2 : \mathbb{R}^n \rightarrow [0, \infty)$ are proper, convex functions,
- $F_2 \in C^{1,1}(\mathbb{R}^n)$ with the Lipschitz constant of DF_2 denoted by $L(F_2)$.

The first ingredient of the algorithm is *Moreau's proximal mapping* [Mor65], given by the following definition.

Definition 3.3.1 (Moreau's proximal mapping). For a convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we define its Moreau's proximal mapping by

$$\text{prox}_\tau^f(y) := \underset{x \in \mathbb{R}^n}{\text{argmin}} f(x) + \frac{\tau}{2}|x - y|^2.$$

Notice that minimizer is unique as we minimize a strictly convex function.

The PALM algorithm [BST14] is based on the combination of the standard gradient descent in F_2 with the proximal mapping of F_1 . In addition to this, the iPALM algorithm uses the extrapolation by an inertial term which adds part of the old direction to the new direction of the algorithm, motivated by the heavy ball approach by Polyak [Pol64]. Altogether, we have Algorithm 1.

Algorithm 1: iPALM Algorithm.

- 1 Choose $x^{[0]}, x^{[1]} \in \mathbb{R}^n, \beta \in [0, 1)$;
 - 2 **for** $i = 1, \dots$ **do**
 - 3 $x^{[i,\beta]} = x^{[i]} + \beta(x^{[i]} - x^{[i-1]})$;
 - 4 $x^{[i+1]} = \text{prox}_L^{F_1}(x^{[i,\beta]} - \frac{1}{L}DF_2(x^{[i,\beta]}))$;
-

There, L is a constant proportional to $L(F_2)$ (cf. [PS16, Lemma 4.1] for the exact relation). The convergence of the sequence $\{x^{[i]}\}_{i \in \mathbb{N}}$ generated by Algorithm 1 to a critical point of F is, by [PS16, Theorem 4.1], guaranteed if the sequence is bounded and F is a *Kurdyka–Lojasiewicz (KL) function*, i.e., subdifferential of F satisfies certain properties (see [PS16, Section 3] for more details and [BDLM10] for characterization of (convex) functions satisfying these properties).

Unfortunately, the exact Lipschitz constant $L(F_2)$ is usually unknown and one must resort to approximating it. To this end, we present a backtracking method based on an algorithm from the class of iterative shrinkage–thresholding algorithms (ISTA) (see [BT09] and references therein). One observes the quadratic approximation of $F(x)$ at a given point y (based on the quadratic approximation of F_2)

$$Q_L(x, y) := F_1(x) + (x - y) \cdot DF_2(y) + \frac{L}{2}|x - y|^2 + F_2(y).$$

Then, using the Cauchy–Schwarz inequality and the definition of the Lipschitz constant, we have

$$F(x) - Q_L(x, y) \leq \frac{L(F_2) - L}{2}|x - y|^2,$$

and thus, for every $L \geq L(F_2)$ and every $x, y \in \mathbb{R}^n$, it holds $F(x) \leq Q_L(x, y)$. There exists a unique

$$p_L(y) := \underset{x \in \mathbb{R}^n}{\text{argmin}} Q_L(x, y) = \text{prox}_L^{F_1}\left(y - \frac{1}{L}DF_2(y)\right).$$

In particular,

$$F(x) = Q_L(x, x) \geq Q(p_L(x), x) \geq F(p_L(x)), \quad (3.17)$$

leading to the update step $x^{[i]} = p_L(x^{[i-1]})$, for $L \geq L(F_2)$. Thus, to obtain a good estimate for L , we can increase the current estimate as long as (3.17) is not satisfied. Altogether, this leads to Algorithm 2.

From (3.17), we get that the sequence $F(x^{[i]})$ is non-increasing. If the minimization problem (3.16) has at least one optimal solution, then, for every optimal solution x^* , [BT09, Theorem 3.1] shows

$$F(x^{[i]}) - F(x^*) \leq \frac{\eta L(F_2)|x^{[0]} - x^*|^2}{2i},$$

which means that we have a sublinear rate of convergence. Finally, combining Algorithm 2 with Algorithm 1 we get Algorithm 3, which will be used in Sections 4.6 and 5.7.

Algorithm 2: ISTA Algorithm with backtracking.

-
- 1 Choose $x^{[0]} \in \mathbb{R}^n$, $\eta > 1$, $L^{[0]} > 0$;
 - 2 **for** $i = 1, \dots$ **do**
 - 3 Find the smallest non-negative integer n_i such that with $L = \eta^{n_i} L^{[i-1]}$:

$$F(p_L(x^{[i-1]})) \leq Q(p_L(x^{[i-1]}), x^{[i-1]});$$
 - 4 Set $L^{[i]} = \eta^{n_i} L^{[i-1]}$ and

$$x^{[i]} = p_{L^{[i]}}(x^{[i-1]});$$
-

Algorithm 3: iPALM Algorithm with the backtracking of Lipschitz constant.

-
- 1 Choose $x^{[0]}, x^{[1]} \in \mathbb{R}^n$, $\beta \in [0, 1)$, $\eta > 1$, $L^{[0]} > 0$;
 - 2 **for** $i = 1, \dots$ **do**
 - 3 $x^{[i,\beta]} = x^{[i]} + \beta(x^{[i]} - x^{[i-1]});$
 - 4 Find the smallest non-negative integer n_i such that with $L = \eta^{n_i} L^{[i-1]}$:

$$F(p_L(x^{[i,\beta]})) \leq Q(p_L(x^{[i,\beta]}), x^{[i,\beta]});$$
 - 5 Set $L^{[i]} = \eta^{n_i} L^{[i-1]}$ and

$$x^{[i+1]} = p_{L^{[i]}}(x^{[i]});$$
-

Remark 3.3.2. We finally comment on generalizations of the above-presented algorithms.

- (i) In Chapters 4 and 5 minimization problems of the form (3.16) will be observed over a set of functions instead of \mathbb{R}^n . In that case, one has to make an additional assumption on (lower semi)continuity of functions F_1 and F_2 .
- (ii) In Algorithm 1, one can do the update step with different inertial coefficients and step sizes in each step, taking the following form

$$x^{[i+1]} = \text{prox}_{L^{[i]}}^{F_1} \left(x^{[i,\alpha^{[i]}]} - \frac{1}{L^{[i]}} DF_2 \left(x^{[i,\beta^{[i]}]} \right) \right).$$

- (iii) The iPALM algorithm is a widely used and versatile optimization algorithm. It allows for the study of minimization problems of the form

$$\min_{x=(x_1, x_2) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}} F(x) = F_{11}(x_1) + F_{12}(x_2) + F_2(x),$$

where (cf. [PS16, Assumption A])

- $F_{11}, F_{12} : \mathbb{R}^{n_1, n_2} \rightarrow (-\infty, +\infty]$ are proper, lower semicontinuous, $\inf_{\mathbb{R}^{n_1, n_2}} F_{11}, F_{12} > -\infty$,
- $F_2 : \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \rightarrow \mathbb{R}$, $\inf_{\mathbb{R}^{n_1} \times \mathbb{R}^{n_2}} F_2 > -\infty$, is differentiable on $\mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$, in class $C^{1,1}$ on bounded subsets of $\mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$, and has globally Lipschitz partial gradients.

Chapter 4

Image Morphing in Deep Feature Space with Anisotropic Regularization

THE image metamorphosis model presented in the previous chapter, its time discrete counterpart, and the spatial discretization based on finite elements in [BER15] present a prominent color-based approach to image morphing. However, as the grayscale or color channels do not carry any semantic information, this leads to an unnatural morphing along the obtained geodesic paths. To overcome these problems, in this chapter, we advocate a metamorphosis model in a deep feature space, which amounts to replacing the input images with feature vectors combining image intensities and semantic information generated by a feature extraction operator. Furthermore, the standard metamorphosis model assumes the strong smoothness of the underlying flow and deformation fields, leading to improper behavior of shape interfaces along the geodesic paths. To this end, we introduce an anisotropic regularization of the variational energy, which allows for a weaker penalty of large gradients of the underlying velocity field in the proximity of interfacial structures, explicitly allowing for discontinuities in the deformation fields. Based on the approach from [BER15], we develop the variational time discretization of the model, prove the existence of discrete geodesic paths, and discuss its Mosco-convergence to the appropriate time continuous metamorphosis model in deep feature space. This in particular implies the convergence of time discrete to time continuous geodesic paths and establishes the existence of time continuous geodesics as minimizers of the time continuous metamorphosis model. The fully discrete model is based on a finite difference discretization of features and deformations and the discretization of the warping operation via cubic Hermite spline interpolation. We use the iPALM algorithm [PS16] for the optimization, which leads to efficient and robust computation of morphing sequences that visually outperform the results obtained by the color-based morphing. This chapter is an extended version of the conference proceeding [EKPR19] and the journal paper [EKP⁺21] (joint project with Alexander Effland, Erich Kobler, Thomas Pock, and Martin Rumpf).

This chapter is structured as follows. In Section 4.1, we introduce the time continuous metamorphosis model in deep feature spaces with an anisotropic regularization. Then, in Section 4.2, we introduce the time discrete model and prove the existence of discrete geodesic paths as minimizers of corresponding energy functionals. Section 4.3 is devoted to the introduction of the time extension operators, necessary for the study of convergence of the time discrete quantities towards their continuous counterparts, presented in Section 4.4. There, we show convergence of the extensions of the time discrete energies towards the time continuous energy, leading to the convergence of the discrete geodesic paths to a time continuous geodesic path. The fully discrete model and the optimization scheme using the iPALM algorithm are presented in Sections 4.5 and 4.6, respectively. In Section 4.7, several examples demonstrate the applicability of the proposed methods to real image data. Finally, in Section 4.8, we comment on two projects being influenced by the presented model: Bézier curves in deep feature space and variational time discretization of the flow of diffeomorphism model, while the spline interpolation in metamorphosis model is presented in Chapter 5.

4.1 Time Continuous Model

As briefly discussed in the introduction, the classical metamorphosis model, its time discrete counterpart, and the spatial discretization based on finite elements in [BER15] exhibit several drawbacks:

- The morphing of images in their original grayscale or color space is not invariant to natural radiometric transformations caused by lighting or material changes, shadows, etc. and hence might lead to a blending along the geodesic paths instead of flow-induced geometric transformations, destroying some vital parts of information important for a natural appearance of images.
- Sharp interfaces such as object boundaries, which frequently coincide with depth discontinuities of a scene, are in general not preserved along a geodesic path because of the strong smoothness implied by the homogeneous and isotropic variational prior for the deformation fields.

To this end, we advocate the following extensions of the previous approaches:

- We work in a deep feature space, which amounts to replacing the input images by feature vectors combining image intensities and semantic information generated by a feature extraction operator.
- To explicitly allow for discontinuities in the deformation fields, we further incorporate an anisotropic flow regularization.

In what follows, we give more details on our approach.

As we briefly mentioned in Section 2.3, deep convolutional neural networks have shown remarkable results for different image processing tasks. A part of this success lies in their ability to extract deep semantic features containing different properties of images that go far beyond the standard color channels. It is, therefore, natural to use a deep neural network as a feature extraction operator. We consider the mapping $F : \mathcal{I} \rightarrow L^2(\Omega, \mathbb{R}^c)$ for $c \geq 0$, which represents the deep convolutional neural network [SZ14] pretrained for classification on the ImageNet dataset [KSH12]. The structure of the network, which we will refer to as *VGG* network (stands for “Visual Geometry Group in Oxford”), is shown in Table 4.1. As already mentioned, as an input to the metamorphosis model, we consider vectors combining image intensities and deep features. To this end, we consider the *deep feature space* $\mathcal{F} := L^2(\Omega, \mathbb{R}^{3+c})$, and given fixed input images $u_A, u_B \in \mathcal{I}$, we want to define geodesic curve joining the deep features $f_A, f_B \in \mathcal{F}$ which are defined, for a fixed (small) $\eta > 0$, by

$$f_A := (\eta u_A, F(u_A)), \quad f_B := (\eta u_B, F(u_B)). \quad (4.1)$$

The parameter η is used to properly scale the RGB component mainly needed to compute the anisotropic regularization and to primarily focus on the actual VGG features when estimating the transport. This is due to the fact that the deep features are high-dimensional image patterns describing the local structure of the image as a superposition of different levels of a multiscale image approximation (*cf.* Figure 4.1). Let us point out that this form of pairing in general holds only for the initial and final feature on the feature path, and we do not restrict the space \mathcal{F} to such pairs. We comment on possibilities to do so in Remarks 4.2.6 and 4.2.9 (for the time discrete model), and Remark 4.4.5 (for the time continuous model).




Figure 4.1: An example of input image and the typical look of the corresponding feature channels on different levels (first to the fifth, from left to right, *cf.* Table 4.1) of multilevel semantic decomposition (plotted with the inverted values for a better visibility).

Next, to properly account for image structures such as sharp edges or corners, we include an anisotropic regularization in our model. We introduce an *anisotropy operator* $a : \mathcal{I} \rightarrow L^\infty(\Omega)$ as a scale factor for the elliptic

operator acting on the flow of diffeomorphisms, which nearly vanishes in the proximity of interfacial structures. Thus, large gradients of the velocity field are less penalized in these regions and consequently, sharp edges can be better preserved along geodesic paths. The need for such an operator was noticed by Perona and Malik [PM90], in the context of solving the problems of image deblurring and edge detection/enhancement by solving the diffusion equation of the type

$$\dot{u}_t = \operatorname{div}(a(|Du_t|^2)Du_t), \quad u(0) = u_0,$$

where the function a is non-negative and monotonically decreasing, with $a(0) \approx 1$. One of the functions they suggest is $a(s^2) = \exp(-\frac{s^2}{\rho})$, where $\rho > 0$ can be seen as an edge threshold or a contrast constant. Their idea is based on an observation that the regions with a higher value of Du have a large probability to be the interfacial regions. In [CLMC92], it was suggested to apply the function a to

$$u_\sigma := \mathcal{G}_\sigma * Du = D(\mathcal{G}_\sigma * u) = D\mathcal{G}_\sigma * u,$$

where \mathcal{G}_σ is the Gaussian kernel with standard deviation $\sigma > 0$ and $*$ is the convolution operation. This approach, besides the theoretical advantage typified in the existence of weak solutions for the diffusion equation, shows better results in the task of edge detection on blurred images, since it is insensitive to noises at scales smaller than σ . Thus, it avoids the shortcoming of the Perona–Malik approach, which misinterprets strong oscillations due to noise as edges that should be preserved or even enhanced. For an extensive discussion of the Perona–Malik model and its extensions, we refer the reader to [BL18, Section 5.3] and [Wei98, Section 1.3]. Taking into consideration the above discussion, our particular choice is the operator

$$a[u](x) := \exp\left(-\frac{|(\mathcal{G}_{\sigma_2} * D\mathcal{G}_{\sigma_1} * u)(x)|^2}{\rho}\right) + \xi, \quad (4.2)$$

for fixed (small) $\xi > 0$ which presents a strictly positive lower bound, necessary for the theoretical study. The study of the time continuous and the time discrete geodesic paths in our model is possible for all operators which fulfill the following technical assumptions

- (A1) *boundedness and coercivity*: $c_a < a[u](x) < C_a$ for $0 < c_a < C_a$ and all $u \in \mathcal{I}$ and a.e. $x \in \Omega$,
- (A2) *compactness*: $u_k \rightarrow u$ in \mathcal{I} implies $a[u_k] \rightarrow a[u]$ in $L^\infty(\Omega)$,
- (A3) *Lipschitz continuity*: for all neighborhoods $\mathcal{U} \subset \mathcal{I}$, there exists $L_a > 0$ such that $\|a[u] - a[\tilde{u}]\|_{L^\infty(\Omega)} \leq L_a \|u - \tilde{u}\|_{\mathcal{I}}$ for all $u, \tilde{u} \in \mathcal{U}$.

Note that (4.2), due to the smoothness of Gaussian kernels, indeed fulfills these conditions.



Figure 4.2: The visual representation of the anisotropy weight for the same input image as in Figure 4.1 and values of the parameters $\rho = 1000$, $\sigma_1 = 0.5$, $\sigma_2 = 2$, $\xi = 10^{-6}$. We observe that the lowest values (corresponding to the almost black grayscale intensity) are indeed obtained in the proximity of edges/discontinuities.

Before finally defining the path energy in our model, we introduce another novelty in comparison to [BER15]. Namely, with the diffeomorphic flow $\psi \in H^1((0, 1), H^m(\Omega, \Omega))$ given as the solution to the flow equation (3.13) for some velocity field $v \in L^2((0, 1), \mathcal{V})$, we could directly generalize the definition of regular paths and the notion of a weak material derivative presented in Subsection 3.2.2. As it was noticed in [ENR20], this approach

becomes technically involved for manifold-valued images/features since the proper tangent space and the norm for the penalization of the material derivative depend on the tangent spaces and Riemannian metrics on the space of values, which are in general different for every point of the domain. As a cure for these obstacles, that paper proposes the relaxed material derivative approach. We choose to follow this approach in the rest of this chapter as it also provides some technical simplifications in our framework. To details, the definition of the weak material derivative $\hat{z} \in L^2((0, 1) \times \Omega, \mathbb{R}^{3+c})$ which is equivalent to (3.11) (cf. [TY05a, Theorem 2]) for $f \in L^2((0, 1), \mathcal{F})$ is given by

$$f_t \circ \psi_t(x) - f_s \circ \psi_s(x) = \int_s^t \hat{z}_r \circ \psi_r(x) dr, \quad \text{for a.e. } x \in \Omega, s, t \in [0, 1]. \quad (4.3)$$

We replace it with the variational inequality

$$|f_t \circ \psi_t(x) - f_s \circ \psi_s(x)| \leq \int_s^t z_r \circ \psi_r(x) dr, \quad \text{for a.e. } x \in \Omega, \text{ for all } 1 \geq t > s \geq 0, \quad (4.4)$$

where the *scalar material derivative* $z \in L^2((0, 1) \times \Omega)$ replaces the actual vector-valued material derivative. This relaxed approach will turn out to be very natural when it comes to lower semicontinuity of the path energy in the context of the existence proof for geodesic paths in Section 4.4. The following proposition (see also [ENR20, Section 3]) shows the equivalence of these definitions.

Proposition 4.1.1. *For every \hat{z} fulfilling (4.3), there exists a z fulfilling (4.4) with $z = |\hat{z}|$. Conversely, for every z fulfilling (4.4), there exists a \hat{z} fulfilling (4.3) with $z \geq |\hat{z}|$.*

Proof. Let \hat{z} satisfy (4.3). Then, for $z = |\hat{z}|$, by the triangle inequality, we have

$$|f_t \circ \psi_t(x) - f_s \circ \psi_s(x)| \leq \int_s^t z_r \circ \psi_r(x) dr.$$

To prove the converse, for z satisfying (4.4), we take the norm on both sides, to get

$$\|f_t \circ \psi_t - f_s \circ \psi_s\|_{L^2(\Omega)} \leq \int_s^t \|z_r \circ \psi_r\|_{L^2(\Omega)} dr,$$

from where we have, by [AGS08, Remark 1.1.3], that the function $t \mapsto f_t \circ \psi_t$ is differentiable almost everywhere on $(0, 1)$ and there exists a derivative $z' \in L^2((0, 1), L^2(\Omega, \mathbb{R}^{3+c}))$ such that, for every $s, t \in [0, 1]$,

$$f_t \circ \psi_t(x) - f_s \circ \psi_s(x) = \int_s^t z'_r(x) dr = \int_s^t \hat{z}_r \circ \psi_r(x) dr,$$

where $\hat{z}_r(x) := z'_r \circ \psi_r^{-1}(x)$. Finally, by [AGS08, Theorem 1.1.2], for a.e. $t \in (0, 1)$, we have $z_t \circ \psi_t \geq |\hat{z}_t \circ \psi_t|$ which together with diffeomorphism property of ψ_t implies $z \geq |\hat{z}|$. \square

In fact, (4.4) defines a set $\mathcal{C}(f)$ of admissible pairs (v, z) given a path f in $L^2((0, 1), \mathcal{F})$ and one can define the path energy analogous to (3.12), with the infimum taken over the pairs in $\mathcal{C}(f)$. This is reasonable also from the fact that the infimum in (3.12) is attained along a measurable curve (cf. [TY05a, Proposition 1 and Theorem 2]).

We are now ready to define the path energy and the geodesic path between the input features given by (4.1).

Definition 4.1.2 (Continuous path energy). We consider the anisotropic elliptic operator

$$L[\tilde{a}, v, v] := \tilde{a} \left(\frac{\lambda}{2} (\text{tr} \varepsilon[v])^2 + \mu \text{tr}(\varepsilon[v]^2) \right) + \gamma |D^m v|^2$$

for an anisotropy weight $\tilde{a} \in L^\infty(\Omega)$, a velocity field $v \in \mathcal{V}$ and $\gamma, \mu, \lambda > 0$. Then, we define the *path energy*

$$\mathcal{E}[f] := \inf_{(v, z) \in \mathcal{C}(f)} \int_0^1 \int_\Omega L[a[\mathcal{P}[f]], v, v] + \frac{1}{\delta} z^2 dx dt \quad (4.5)$$

for a path $f \in L^2([0, 1], \mathcal{F})$, where \mathcal{P} is the projection onto the image component of a feature, i.e., $\mathcal{P}[(u, \tilde{f})] = u$, and

$$\mathcal{C}(f) \subset L^2((0, 1), \mathcal{V}) \times L^2((0, 1) \times \Omega)$$

denotes the set of admissible pairs of the velocity v and a scalar quantity z fulfilling

$$\dot{\psi}_t(x) = v_t \circ \psi_t(x), \quad \psi_0(x) = x, \quad \text{for all } x \in \Omega, t \in [0, 1], \quad (4.6)$$

$$|f_t \circ \psi_t(x) - f_s \circ \psi_s(x)| \leq \int_s^t z_r \circ \psi_r(x) \, dr, \quad \text{for a.e. } x \in \Omega, \text{ for all } t > s \in [0, 1]. \quad (4.7)$$

Let us stress that the anisotropy $\tilde{a} = a[\mathcal{P}[f]]$ solely takes into account local RGB values and not the actual VGG features with their discriminative multiscale characteristics.

Definition 4.1.3 (Continuous geodesic curve). Geodesic curves $f \in L^2([0, 1], \mathcal{F})$ in the deep feature space joining $f_A, f_B \in \mathcal{F}$ are defined as minimizers of the path energy (4.5) among all curves with the fixed boundary conditions $f_0 = f_A$ and $f_1 = f_B$.

The existence of a continuous geodesic curve will be proven in Theorem 4.4.4. However, the geodesic curve might not be unique and the reverse path is in general not a geodesic curve if the initial and the final point switch places.

Remark 4.1.4. One observes that a path $f \in L^2([0, 1], \mathcal{F})$ with finite energy $\mathcal{E}[f] < \infty$ exhibits additional smoothness properties in time. Indeed, from (4.4), by using the Cauchy–Schwarz inequality, we have

$$\begin{aligned} \|f_t \circ \psi_t - f_s \circ \psi_s\|_{\mathcal{F}} &\leq \left(\int_{\Omega} \left(\int_s^t z_r \circ \psi_r \, dr \right)^2 dx \right)^{\frac{1}{2}} \\ &\leq |t - s|^{\frac{1}{2}} \|z_r \circ \psi_r\|_{L^2((0,1) \times \Omega)} \\ &\leq |t - s|^{\frac{1}{2}} \|z\|_{L^2((0,1) \times \Omega)} \|\det D\psi^{-1}\|_{C^0([0,1] \times \bar{\Omega})}^{\frac{1}{2}} \leq C_v |t - s|^{\frac{1}{2}} \|z\|_{L^2((0,1) \times \Omega)}, \end{aligned}$$

where we further used (3.14), so that the constant C_v depends on $\|v\|_{L^2((0,1), \mathcal{V})}$. Thus, $f_t \circ \psi_t \in C^{0, \frac{1}{2}}([0, 1], \mathcal{F})$, or in a slightly different notation, $f_t \circ \psi_t \in C_{\omega}^0([0, 1], \mathcal{F})$, for the modulus of continuity $\omega(t) = C_v \|z\|_{L^2((0,1) \times \Omega)} |t|^{\frac{1}{2}}$. Furthermore, for every $t \in [0, 1]$, we have

$$\|f_t\|_{\mathcal{F}} \leq C_v \|f_t \circ \psi_t\|_{\mathcal{F}} \leq C_{v,z} (1 + \|f_t \circ \psi_t\|_{L^2([0,1], \mathcal{F})}), \quad (4.8)$$

so that $f \in L^{\infty}([0, 1], \mathcal{F})$. Here, $C_{v,z}$ depends on $\|v\|_{L^2((0,1), \mathcal{V})}$ and $\|z\|_{L^2((0,1) \times \Omega)}$. Then, to show the improved regularity for the feature curve $t \mapsto f_t$, we can take functions $\tilde{f}^i \in L^{\infty}([0, 1], C^{\infty}(\Omega, \mathbb{R}^{3+c}))$ such that $\|f - \tilde{f}^i\|_{L^{\infty}([0,1], \mathcal{F})} \leq 2^{-i}$ and $\|D\tilde{f}^i\|_{L^{\infty}([0,1], C^0(\bar{\Omega}))} \leq C_{f,i}$, where $C_{f,i}$ depends on $\|f\|_{L^{\infty}([0,1], \mathcal{F})}$ and i (cf. Theorem 2.1.7). We have

$$\begin{aligned} &\|f_t - f_s\|_{\mathcal{F}} \\ &\leq \|f_t \circ \psi_t \circ \psi_t^{-1} - f_s \circ \psi_s \circ \psi_t^{-1}\|_{\mathcal{F}} + \|f_s \circ \psi_s \circ \psi_t^{-1} - f_s \circ \psi_s \circ \psi_s^{-1}\|_{\mathcal{F}} \\ &\leq \|f_t \circ \psi_t \circ \psi_t^{-1} - f_s \circ \psi_s \circ \psi_t^{-1}\|_{\mathcal{F}} + \|f_s \circ \psi_s \circ \psi_t^{-1} - \tilde{f}_s^i \circ \psi_s \circ \psi_t^{-1}\|_{\mathcal{F}} \\ &\quad + \|\tilde{f}_s^i \circ \psi_s \circ \psi_t^{-1} - \tilde{f}_s^i \circ \psi_s \circ \psi_s^{-1}\|_{\mathcal{F}} + \|\tilde{f}_s^i \circ \psi_s \circ \psi_s^{-1} - f_s \circ \psi_s \circ \psi_s^{-1}\|_{\mathcal{F}} \\ &\leq \|f_t \circ \psi_t - f_s \circ \psi_s\|_{\mathcal{F}} \|\det D\psi_t\|_{C^0(\bar{\Omega})}^{\frac{1}{2}} + \|f_s - \tilde{f}_s^i\|_{\mathcal{F}} \|\det D(\psi_s \circ \psi_t^{-1})^{-1}\|_{C^0(\bar{\Omega})}^{\frac{1}{2}} \\ &\quad + \|D\tilde{f}_s^i\|_{C^0(\bar{\Omega})} \|D\psi_s\|_{C^0(\bar{\Omega})} |\psi^{-1}|_{C^{0, \frac{1}{2}}([0,1], C^0(\bar{\Omega}))} |t - s|^{\frac{1}{2}} + \|f_s - \tilde{f}_s^i\|_{\mathcal{F}} \\ &\leq C_v \|f_t \circ \psi_t - f_s \circ \psi_s\|_{\mathcal{F}} + C_v \|f_s - \tilde{f}_s^i\|_{\mathcal{F}} + C_v \|D\tilde{f}_s^i\|_{C^0(\bar{\Omega})} |t - s|^{\frac{1}{2}} + \|f_s - \tilde{f}_s^i\|_{\mathcal{F}} \\ &\leq C_{v,z} \left(|t - s|^{\frac{1}{2}} + 2^{-i} \right) + C_{v,f,i} |t - s|^{\frac{1}{2}}, \end{aligned}$$

where the corresponding constants depend, in a monotonically increasing manner, on $\|v\|_{L^2((0,1), \mathcal{V})}$, $\|z\|_{L^2((0,1) \times \Omega)}$, $\|f\|_{L^{\infty}([0,1], \mathcal{F})}$, and i . Altogether, choosing a suitable i , we see that $t \mapsto f_t \in C^0([0, 1], \mathcal{F})$ is uniformly continuous, and we have $f \in C_{\omega}^0([0, 1], \mathcal{F})$, where ω depends on $C_{v,z,f}$. Notice that this implies, by (A3), that $a[\mathcal{P}[f_t]]$ is well-defined, in $C^0([0, 1], L^{\infty}(\Omega))$, and uniformly continuous in time with the analogous modulus of continuity estimates.

4.2 Variational Time Discretization

In this section, we develop a variational time discretization of the above model. To this end, we are taking into account the approach presented in [BER15, Section 2] which is based on the discrete geodesic calculus framework presented in Subsection 3.1.2.

The first goal is to find a suitable time discretization of the quantities appearing in (4.5). First, we observe that, for $s < t \in [0, 1]$ close to each other, we have $\frac{\psi_t - \psi_s}{t-s} \approx v_s \circ \psi_s$, i.e., $v_s \approx \frac{\psi_t \circ \psi_s^{-1} - \mathbb{1}}{t-s}$, by applying the forward finite difference to (4.6). Similarly, from (4.7), we get $z_s \approx \frac{|f_t \circ \psi_t \circ \psi_s^{-1} - f_s|}{t-s}$. Motivated by these observations and the aforementioned approaches (cf. (3.6)), we define the *time discrete pairwise energy* for two feature maps $f, \tilde{f} \in \mathcal{F}$ by

$$\mathcal{W}[f, \tilde{f}] := \min_{\phi \in \mathcal{D}} \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi],$$

with $\mathcal{W}^D : L^\infty(\Omega) \times \mathcal{F} \times \mathcal{F} \times \mathcal{D} \rightarrow \mathbb{R}$ given by

$$\mathcal{W}^D[\tilde{a}, f, \tilde{f}, \phi] := \int_{\Omega} \tilde{a} \mathcal{W}(D\phi) + \gamma |D^m \phi|^2 + \frac{1}{\delta} |\tilde{f} \circ \phi - f|^2 dx. \quad (4.9)$$

The deformation ϕ can be seen as the relative flow $\psi_{\tilde{t}} \circ \psi_t^{-1}$, where \tilde{f}, f are seen as points on the feature curve at times \tilde{t}, t , respectively. The *set of admissible deformations* is

$$\mathcal{D} := \{\phi \in H^m(\Omega, \Omega) : \det(D\phi) > 0 \text{ a.e. in } \Omega, \phi|_{\partial\Omega} = \mathbb{1}\}. \quad (4.10)$$

We have already introduced this set as the set of Sobolev diffeomorphisms $\mathcal{D}^m(\Omega)$ (cf. (3.9)), and here, for convenience, we simplified the notation to $\mathcal{D} := \mathcal{D}^m(\Omega)$. We recall that this set is closed under the inverse [IKT13, Lemma A.1] and the composition operations [IKT13, Proposition 2.19]. Furthermore, by the Sobolev embedding theorem 2.1.8, we have that $\mathcal{D} \subset C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$ for $\alpha \in (0, m - \frac{n}{2} - 1]$, so that $\phi \in \mathcal{D}$ are indeed $C^1(\Omega, \Omega)$ -diffeomorphisms.

With $\text{GL}^+(n)$ being the set of n -dimensional matrices with the positive determinant, and $\mathbb{1}$ denoting both the identity mapping and the identity matrix, we make the following assumptions on the *energy density function* \mathcal{W} :

(W1) $\mathcal{W} : \mathbb{R}^{n,n} \rightarrow \mathbb{R}_0^+$, $\mathcal{W} \in C^4(\text{GL}^+(n))$, $\mathcal{W}(\mathbb{1}) = 0$, and $D\mathcal{W}(\mathbb{1}) = 0$.

(W2) There exist constants $C_{\mathcal{W},1}, C_{\mathcal{W},2}, r_{\mathcal{W}} > 0$ such that for all $A \in \text{GL}^+(n)$

$$\begin{aligned} \mathcal{W}(A) &\geq C_{\mathcal{W},1} |A^{\text{sym}} - \mathbb{1}|^2, & \text{if } |A - \mathbb{1}| < r_{\mathcal{W}}, \\ \mathcal{W}(A) &\geq C_{\mathcal{W},2}, & \text{if } |A - \mathbb{1}| \geq r_{\mathcal{W}}. \end{aligned}$$

(W3) For all $A \in \mathbb{R}^{n,n}$, it holds

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

The first assumption ensures that we indeed penalize the velocity term which vanishes when $\phi = \mathbb{1}$. We use (W1) and (W2) to prove existence of a minimizing deformation in (4.9), while the third condition is a consistency assumption with respect to the differential operator L (see (3.10)), required to guarantee that the below-defined discrete path energy is consistent with the time continuous path energy (4.5) (see proof of Theorem 4.4.1).

The particular energy density function

$$\mathcal{W}(D\phi) = \frac{\lambda}{2} \left(e^{(\log \det(D\phi))^2} - 1 \right) + \mu |\varepsilon[\phi] - \mathbb{1}|^2 \quad (4.11)$$

used for all numerical experiments, with $n = 2$, satisfies (W1)–(W3). On one hand, both terms penalize large deformations by penalizing the change of volume and change of length under the deformation, controlled by $\det(D\phi)$ and $|D\phi|$, respectively. On the other hand, the first term penalizes small deformations as it goes to infinity as $\det(D\phi)$ tends to 0. Thus, it favors a balance of shrinkage and growth as advocated in [DR04, BMR13]. Furthermore, it enforces the positivity of $\det(D\phi)$ which is a crucial condition for invertibility of deformation, i.e., prohibition of the interpenetration of matter [Bal81]. This function is not convex and, in general, there is

no a convex function such that $W(A) \rightarrow \infty$ as $\det(A) \rightarrow 0$ [Cia88, Theorem 4.8-1]. However, the function $\tilde{W} : \mathbb{R}^{2,2} \times \mathbb{R} \rightarrow \mathbb{R}$ with $\tilde{W}(A, \det(A)) := W(A)$ is convex. Then, we say that W is *polyconvex* and it gives rise to a lower semicontinuous functional $\int_{\Omega} W(D\phi) dx$ [Dac08, Theorems 5.3 and 8.11].

We proceed with the definition of the discrete path energy and the discrete geodesic path between two features $f_A = (\eta u_A, F(u_A)), f_B = (\eta u_B, F(u_B)) \in \mathcal{F}$.

Definition 4.2.1 (Discrete path energy and discrete geodesic path). Let $K \geq 1$ and $f_0 = f_A, f_K = f_B \in \mathcal{F}$. The *discrete path energy* \mathbf{E}^K , for a discrete $(K+1)$ -path $\mathbf{f} = (f_0, \dots, f_K) \in \mathcal{F}^{K+1}$, is defined by

$$\mathbf{E}^K[\mathbf{f}] := K \sum_{k=1}^K \mathcal{W}[f_{k-1}, f_k]. \quad (4.12)$$

A *discrete geodesic path* morphing $f_A \in \mathcal{F}$ into $f_B \in \mathcal{F}$ is a discrete $(K+1)$ -path that minimizes \mathbf{E}^K over all discrete paths $\mathbf{f} = (f_A, \hat{\mathbf{f}}, f_B) \in \mathcal{F}^{K+1}$ with $\hat{\mathbf{f}} = (f_1, \dots, f_{K-1}) \in \mathcal{F}^{K-1}$.

For arbitrary vectors $\mathbf{f} = (f_0, \dots, f_K) \in \mathcal{F}^{K+1}$ and $\Phi = (\phi_1, \dots, \phi_K) \in \mathcal{D}^K$ we set

$$\mathbf{E}^{K,D}[\mathbf{f}, \Phi] := K \sum_{k=1}^K \mathcal{W}^D[a[\mathcal{P}[f_k]], f_{k-1}, f_k, \phi_k].$$

These definitions come naturally if one observes \mathbf{f} as the time sampling of a feature curve at times $\frac{k}{K}$, $k = 0, \dots, K$, and the entries of Φ as the relative flows between the consecutive time points. Then, as already motivated, for $k = 1, \dots, K$ we have $v_k = K(\phi_k - \mathbb{1})$ and $z_k = K|f_k \circ \phi_k - f_{k-1}|$ as the discrete versions of the velocity and the scalar material derivative, respectively. Finally, the coefficient K in front of the sum comes from the quadratic growth assumption (W2) and the Riemannian approximation of the time integral. This discretization is consistent with the discrete path geodesic energy from Subsection 3.1.2 (cf. (3.7)).

In what follows, we will investigate the existence of discrete geodesic curves in this model. To this end, we combine the proofs of the local well-posedness of the pairwise energy \mathcal{W} with the existence result of a feature vector minimizing $\mathbf{E}^{K,D}$ for a fixed vector of deformations. The main differences of the following proofs in comparison to [BER15] are due to incorporating the anisotropy regularization.

We start with the following lemma, which provides an estimate for the H^m -norm of the displacement, is crucial for the well-posedness of the energy.

Lemma 4.2.2. *Let (W1)–(W2) and (A1) be satisfied. Then, there exists a continuous and monotonically increasing function $\theta : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ with $\theta(0) = 0$, which only depends on $\Omega, m, n, \gamma, c_a, C_{W,1}, C_{W,2}$, such that*

$$\|\phi - \mathbb{1}\|_{H^m(\Omega)} \leq \theta\left(\mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi]\right)$$

for all $f, \tilde{f} \in \mathcal{F}$ and all $\phi \in \mathcal{D}$. Furthermore, $\theta(x) \leq C(x + x^2)^{\frac{1}{2}}$ for a constant $C > 0$.

Proof. Set $\overline{\mathcal{W}} = \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi]$. An application of the Gagliardo–Nirenberg inequality for bounded domains (see Theorem 2.1.5) yields

$$\|\phi - \mathbb{1}\|_{H^m(\Omega)} \leq C(\|\phi - \mathbb{1}\|_{L^2(\Omega)} + |\phi - \mathbb{1}|_{H^m(\Omega)}). \quad (4.13)$$

The last term in (4.13) is bounded by

$$|\phi - \mathbb{1}|_{H^m(\Omega)} = |\phi|_{H^m(\Omega)} \leq \sqrt{\frac{\overline{\mathcal{W}}}{\gamma}}. \quad (4.14)$$

By using the embedding of $H^m(\Omega, \Omega)$ into $C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})$ (cf. Theorem 2.1.8) and the uniform boundedness of the minimizing sequence in $L^2(\Omega, \Omega)$, we get $\|\phi - \mathbb{1}\|_{C^{1,\alpha}(\overline{\Omega})} \leq C + C\sqrt{\overline{\mathcal{W}}}$. To control the lower order term appearing on the right-hand side of (4.13), we define $\mathcal{S} = \{x \in \Omega : |D\phi(x) - \mathbb{1}| < r_W\}$ and use (A1) and (W2), to obtain

$$|\Omega \setminus \mathcal{S}| c_a C_{W,2} \leq \int_{\Omega} a[\mathcal{P}[\tilde{f}]] W(D\phi) dx \leq \overline{\mathcal{W}},$$

which implies $|\Omega \setminus \mathcal{S}| \leq \frac{\overline{\mathcal{W}}}{c_a C_{\mathcal{W},2}}$. Then, using once again (A1), (W2), and the above observations, we infer

$$\begin{aligned} \int_{\Omega} |\varepsilon[\phi] - \mathbf{1}|^2 dx &\leq \int_{\mathcal{S}} \frac{\mathbf{W}(D\phi)}{C_{\mathcal{W},1}} dx + |\Omega \setminus \mathcal{S}| \left(C + C\sqrt{\overline{\mathcal{W}}} \right)^2 \\ &\leq \frac{\overline{\mathcal{W}}}{c_a C_{\mathcal{W},1}} + \frac{\overline{\mathcal{W}}}{c_a C_{\mathcal{W},2}} (C + C\overline{\mathcal{W}}). \end{aligned} \quad (4.15)$$

Finally, we observe that the inequality

$$\|\phi - \mathbf{1}\|_{L^2(\Omega)} \leq C \|\varepsilon[\phi] - \mathbf{1}\|_{L^2(\Omega)} \quad (4.16)$$

holds true, which follows from Korn's inequality (Theorem 2.1.6) and the Poincaré's inequality 2.1.4. Thus, the lemma follows by combining (4.13), (4.14), (4.15), and (4.16). \square

Proposition 4.2.3 (Well-posedness of \mathcal{W}). *Let $f \in \mathcal{F}$ be a fixed feature vector. Under the assumptions (W1)–(W2) and (A1), there exists a constant $C_{\mathcal{W}}$ (depending on $\Omega, m, n, \gamma, \delta, c_a, C_{\mathcal{W},1}, C_{\mathcal{W},2}$) such that, for every*

$$\tilde{f} \in \{g \in \mathcal{F} : \|f - g\|_{\mathcal{F}} < C_{\mathcal{W}}\}, \quad (4.17)$$

there exists a $\phi \in \mathcal{D}$ which minimizes $\mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \cdot]$.

Proof. For a fixed $f \in \mathcal{F}$, let \tilde{f} be a feature vector satisfying (4.17) for a constant $C_{\mathcal{W}}$ specified below. Let $\{\phi^j\}_{j \in \mathbb{N}} \in \mathcal{D}$ be any sequence such that $\mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi^j]$ converges to $\underline{\mathbf{W}} := \inf_{\phi \in \mathcal{D}} \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi] \geq 0$. Since $\mathbf{1} \in \mathcal{D}$, we can deduce, using (W1), that

$$\underline{\mathbf{W}} \leq \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi^j] \leq \overline{\mathbf{W}} := \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \mathbf{1}] = \frac{1}{\delta} \|\tilde{f} - f\|_{\mathcal{F}}^2 < \frac{C_{\mathcal{W}}^2}{\delta},$$

for all $j \in \mathbb{N}$. Using again the Gagliardo–Nirenberg inequality (see Theorem 2.1.5) in form of (4.13), we infer that $\{\phi^j\}_{j \in \mathbb{N}}$ is uniformly bounded in $H^m(\Omega, \Omega)$ because of the estimate $|\phi^j|_{H^m(\Omega)}^2 \leq \frac{\overline{\mathbf{W}}}{\gamma}$. Due to the reflexivity of $H^m(\Omega, \Omega)$, there exists a weakly convergent subsequence (not relabeled) such that $\phi^j \rightharpoonup \phi$ in $H^m(\Omega, \Omega)$. By using the Sobolev Embedding Theorem 2.1.8, as well as the Arzelà–Ascoli Theorem 2.1.9(i), we can additionally infer that, on a subsequence (again not relabeled), $\phi^j \rightarrow \phi$ in $C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})$ for $\alpha \in (0, m - 1 - \frac{n}{2}]$ holds true. Then, Lemma 4.2.2 implies

$$\|\phi^j - \mathbf{1}\|_{C^1(\overline{\Omega})} \leq C\theta(\overline{\mathbf{W}}) < C\theta(\delta^{-1}C_{\mathcal{W}}^2). \quad (4.18)$$

Thus, by choosing $C_{\mathcal{W}}$ sufficiently small and taking into account the Lipschitz continuity of the determinant, we obtain $\|\det(D\phi^j) - \mathbf{1}\|_{L^\infty(\Omega)} \leq C_{\det}$ for a constant $C_{\det} \in (0, 1)$ and all $j \in \mathbb{N}$, which implies $\det(D\phi^j) \geq C > 0$ for a constant C . Note that all the estimates remain valid for the limit deformation, so that $\{\phi^j\}_{j \in \mathbb{N}}$ and ϕ are in \mathcal{D} and thus $C^1(\Omega, \Omega)$ -diffeomorphisms. Finally, (W1) and the lower semicontinuity of the seminorm imply

$$\liminf_{j \rightarrow \infty} \int_{\Omega} a[\mathcal{P}[\tilde{f}]] \mathbf{W}(D\phi^j) + \gamma |D^m \phi^j|^2 dx \geq \int_{\Omega} a[\mathcal{P}[\tilde{f}]] \mathbf{W}(D\phi) + \gamma |D^m \phi|^2 dx.$$

Let us verify that

$$\lim_{j \rightarrow \infty} \|\tilde{f} \circ \phi^j - f\|_{\mathcal{F}} = \|\tilde{f} \circ \phi - f\|_{\mathcal{F}}. \quad (4.19)$$

To this end, we approximate \tilde{f} by smooth functions $\tilde{f}^i \in C^\infty(\Omega, \mathbb{R}^{3+c})$ with $\|\tilde{f} - \tilde{f}^i\|_{\mathcal{F}} \rightarrow 0$ (cf. Theorem 2.1.7). Then, using the transformation formula, we obtain

$$\begin{aligned} &\|\tilde{f} \circ \phi^j - \tilde{f} \circ \phi\|_{\mathcal{F}} \\ &\leq \|\tilde{f} \circ \phi^j - \tilde{f}^i \circ \phi^j\|_{\mathcal{F}} + \|\tilde{f}^i \circ \phi^j - \tilde{f}^i \circ \phi\|_{\mathcal{F}} + \|\tilde{f}^i \circ \phi - \tilde{f} \circ \phi\|_{\mathcal{F}} \\ &\leq \|\tilde{f} - \tilde{f}^i\|_{\mathcal{F}} \left(\|\det(D(\phi^j)^{-1})\|_{L^\infty(\Omega)}^{\frac{1}{2}} + \|\det(D\phi^{-1})\|_{L^\infty(\Omega)}^{\frac{1}{2}} \right) + \|D\tilde{f}^i\|_{L^\infty(\Omega)} \|\phi^j - \phi\|_{L^2(\Omega)}, \end{aligned} \quad (4.20)$$

where $\det(D(\phi^j)^{-1})$ and $\det(D\phi)^{-1}$ are pointwise estimated by $(1 - C_{\det})^{\frac{1}{2}}$. Finally, by first choosing i and then j , we obtain (4.19). Altogether,

$$\inf_{\phi \in \mathcal{D}} \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi] = \liminf_{j \rightarrow \infty} \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi^j] \geq \mathcal{W}^D[a[\mathcal{P}[\tilde{f}]], f, \tilde{f}, \phi],$$

which finishes the proof. \square

This proposition guarantees the existence of an admissible vector of deformations $\Phi \in \mathcal{D}^K$ for which it holds $\mathbf{E}^{K,D}[\mathbf{f}, \Phi] = \mathbf{E}^K[\mathbf{f}]$, provided that each pair of features (f_k, f_{k+1}) contained in $\mathbf{f} = (f_0, \dots, f_K) \in \mathcal{F}^{K+1}$ satisfies (4.17). The property (4.19) holds more generally, as stated in the following remark and used throughout this and the following chapter.

Remark 4.2.4. The estimate (4.20) shows that, for a sequence of diffeomorphisms $\{\phi^j\}_{j \in \mathbb{N}}$ with $\det D\phi^j \geq c_{\det} > 0$, for every (large enough) $j \in \mathbb{N}$, which converges to a diffeomorphism ϕ in $C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$ and every $f \in \mathcal{F}$, it holds $f \circ \phi^j \rightarrow f \circ \phi$ in \mathcal{F} . Analogously, $f \circ (\phi^j)^{-1} \rightarrow f \circ \phi^{-1}$ in \mathcal{F} . This straightforwardly extends to the Bochner spaces $L^2((0, 1), \mathcal{F})$ and $C^{0,\beta}([0, 1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$, $\beta \geq 0$, for f and ϕ , respectively.

In what follows, we prove the existence of an energy minimizing vector of features for a fixed vector of deformations.

Proposition 4.2.5. *Let $K \geq 2$, $f_A, f_B \in \mathcal{F}$ and $\Phi = (\phi_1, \dots, \phi_K) \in \mathcal{D}^K$ be fixed. Assume that the deformations satisfy*

$$\min_{k=1, \dots, K} \min_{x \in \Omega} \det(D\phi_k(x)) \geq c_{\det}, \quad (4.21)$$

for a constant $c_{\det} > 0$. Then, under the assumptions (W1) and (A1)–(A2), there exists a feature vector \mathbf{f} with $f_0 = f_A$ and $f_K = f_B$ such that

$$\mathbf{E}^{K,D}[\mathbf{f}, \Phi] = \inf \{ \mathbf{E}^{K,D}[(f_A, \hat{\mathbf{g}}, f_B), \Phi] : \hat{\mathbf{g}} \in \mathcal{F}^{K-1} \}.$$

Proof. We consider a minimizing sequence of features $\{\hat{\mathbf{f}}^j = (f_1^j, \dots, f_{K-1}^j)\}_{j \in \mathbb{N}} \in \mathcal{F}^{K-1}$ for the energy $\hat{\mathbf{g}} \mapsto \mathbf{E}^{K,D}[(f_A, \hat{\mathbf{g}}, f_B), \Phi]$. Then,

$$0 \leq \mathbf{E}^{K,D}[(f_A, \hat{\mathbf{f}}^j, f_B), \Phi] \leq \mathbf{E}^{K,D}[(f_A, (f_A, \dots, f_A), f_B), \Phi] =: \overline{\mathbf{E}^{K,D}}.$$

A straightforward computation reveals

$$\overline{\mathbf{E}^{K,D}} \leq K \sum_{k=1}^K C_a \|W(D\phi_k)\|_{L^1(\Omega)} + \gamma \|\phi_k\|_{H^m(\Omega)}^2 + \frac{CK^2}{\delta} ((1 + c_{\det}^{-1}) \|f_A\|_{\mathcal{F}}^2 + c_{\det}^{-1} \|f_B\|_{\mathcal{F}}^2),$$

where we used (A1), (4.21), and the transformation formula. Furthermore, again by (4.21), one obtains

$$\|f_k^j\|_{\mathcal{F}} \leq \|f_{k+1}^j \circ \phi_{k+1} - f_k^j\|_{\mathcal{F}} + \|f_{k+1}^j \circ \phi_{k+1}\|_{\mathcal{F}} \leq \sqrt{\frac{\delta \overline{\mathbf{E}^{K,D}}}{K}} + c_{\det}^{-\frac{1}{2}} \|f_{k+1}^j\|_{\mathcal{F}}. \quad (4.22)$$

Thus, an induction argument (starting from $k = K - 1$) shows that $\{\hat{\mathbf{f}}^j = (f_1^j, \dots, f_{K-1}^j)\}_{j \in \mathbb{N}}$ is uniformly bounded in \mathcal{F}^{K-1} , which implies, for a subsequence (not relabeled), $\hat{\mathbf{f}}^j \rightharpoonup \hat{\mathbf{f}}$ in \mathcal{F}^{K-1} .

In what follows, we prove the weak lower semicontinuity of the discrete path energy along the minimizing sequence. We observe that (A2) implies $a[\mathcal{P}[f_k^j]] \rightarrow a[\mathcal{P}[f_k]]$ in $L^\infty(\Omega)$, which yields

$$\lim_{j \rightarrow \infty} \int_{\Omega} a[\mathcal{P}[f_k^j]] W(D\phi_k) dx = \int_{\Omega} a[\mathcal{P}[f_k]] W(D\phi_k) dx,$$

for every $k = 1, \dots, K$. It remains to verify the weak lower semicontinuity of the matching functional, *i.e.*,

$$\|f_k \circ \phi_k - f_{k-1}\|_{\mathcal{F}}^2 \leq \liminf_{j \rightarrow \infty} \|f_k^j \circ \phi_k - f_{k-1}^j\|_{\mathcal{F}}^2, \quad (4.23)$$

for every $k = 1, \dots, K$. To this end, we show $f_k^j \circ \phi_k \rightharpoonup f_k \circ \phi_k$ in \mathcal{F} . Indeed, for every $g \in \mathcal{F}$, the transformation formula yields

$$\int_{\Omega} (f_k^j \circ \phi_k - f_k \circ \phi_k) \cdot g dx = \int_{\Omega} (f_k^j - f_k) \cdot (g(\det(D\phi_k))^{-1}) \circ \phi_k^{-1} dx,$$

which converges to 0 since $(g(\det(D\phi_k))^{-1}) \circ \phi_k^{-1} \in \mathcal{F}$ due to (4.21). Hence, $f_k^j \circ \phi_k - f_{k-1}^j \rightharpoonup f_k \circ \phi_k - f_{k-1}$ in \mathcal{F} , which readily implies (4.23). Therefore,

$$\inf_{\hat{\mathbf{g}} \in \mathcal{F}^{K-1}} \mathbf{E}^{K,D}[(f_A, \hat{\mathbf{g}}, f_B), \Phi] = \liminf_{j \rightarrow \infty} \mathbf{E}^{K,D}[(f_A, \hat{\mathbf{f}}^j, f_B), \Phi] \geq \mathbf{E}^{K,D}[(f_A, \hat{\mathbf{f}}, f_B), \Phi],$$

which proves the proposition. \square

Remark 4.2.6. If mapping $F : \mathcal{I} \rightarrow L^2(\Omega, \mathbb{R}^c)$ is compact in the sense that $u^j \rightarrow u$ in \mathcal{I} implies $F(u^j) \rightarrow F(u)$ in $L^2(\Omega, \mathbb{R}^c)$, the previous results also hold in the case where the space \mathcal{F} is restricted to the points of the type $(\eta u, F(u))$ for $u \in \mathcal{I}$ (cf. (4.1)). This, in particular, holds under the assumption on spatial smoothness of the convolutional weights and if the maximum pooling is replaced by the average pooling in the architecture of the network (cf. Table 4.1).

We can now combine both previous propositions to prove the existence of discrete geodesics.

Theorem 4.2.7 (Existence of discrete geodesic paths). *Let the assumptions (W1)–(W2) and (A1)–(A2) be satisfied, $K \geq 2$, and $f_A \in \mathcal{F}$. Then, there exists a constant $C_{\mathbf{E}} > 0$ which is independent of K , such that, for every*

$$f_B \in \left\{ g \in \mathcal{F} : \|g - f_A\|_{\mathcal{F}} < C_{\mathbf{E}} \sqrt{K} \right\}, \quad (4.24)$$

there exists $\hat{\mathbf{f}} \in \mathcal{F}^{K-1}$ such that

$$\mathbf{E}^K[(f_A, \hat{\mathbf{f}}, f_B)] = \inf_{\hat{\mathbf{g}} \in \mathcal{F}^{K-1}} \mathbf{E}^K[(f_A, \hat{\mathbf{g}}, f_B)].$$

Proof. For a fixed $f_A \in \mathcal{F}$, let f_B satisfy (4.24) for a constant $C_{\mathbf{E}}$ specified below. For $k = 0, \dots, K$, let $\bar{f}_k := \frac{k}{K} f_B + (1 - \frac{k}{K}) f_A \in \mathcal{F}$ be a convex combination of the input features. Then, we have that

$$\begin{aligned} \overline{\mathbf{E}^K} &:= \mathbf{E}^{K,D}[(\bar{f}_0, \bar{f}_1, \dots, \bar{f}_K), (\mathbf{1}, \dots, \mathbf{1})] \\ &= \frac{K}{\delta} \sum_{k=1}^K \|f_k - f_{k-1}\|_{\mathcal{F}}^2 = \frac{1}{\delta} \|f_B - f_A\|_{\mathcal{F}}^2 < \frac{C_{\mathbf{E}}^2 K}{\delta} \end{aligned} \quad (4.25)$$

is a finite upper bound for the energy. Consider a sequence

$$(\mathbf{f}^j, \Phi^j) = ((f_0^j, \dots, f_K^j), (\phi_1^j, \dots, \phi_K^j)) \in \mathcal{F}^{K+1} \times \mathcal{D}^K, \quad j \in \mathbb{N},$$

with $f_0^j = f_A$ and $f_K^j = f_B$, $\lim_{j \rightarrow \infty} \mathbf{E}^K[\mathbf{f}^j] = \lim_{j \rightarrow \infty} \mathbf{E}^{K,D}[\mathbf{f}^j, \Phi^j] = \inf_{\hat{\mathbf{g}} \in \mathcal{F}^{K-1}} \mathbf{E}^K[(f_A, \hat{\mathbf{g}}, f_B)]$, and with the finite upper bound $\overline{\mathbf{E}^K}$ for the energy \mathbf{E}^K . Following the same line of arguments as in Proposition 4.2.3, we obtain the boundedness of $\{\Phi^j\}_{j \in \mathbb{N}}$ in $H^m(\Omega, \Omega)$, which results in a weakly convergent subsequence (not relabeled) $\Phi^j \rightharpoonup \Phi$ in $H^m(\Omega, \Omega)$, and due to $H^m(\Omega, \Omega) \hookrightarrow C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$, one obtains $\Phi^j \rightarrow \Phi$ in $C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$, $\alpha \in (0, m - 1 - \frac{n}{2}]$, for a further subsequence (not relabeled). By taking into account Lemma 4.2.2, we get

$$\|\phi_k^j - \mathbf{1}\|_{C^{1,\alpha}(\bar{\Omega})} \leq C \|\phi_k^j - \mathbf{1}\|_{H^m(\Omega)} \leq C\theta(K^{-1} \overline{\mathbf{E}^K}) \leq C\theta(\delta^{-1} C_{\mathbf{E}}^2),$$

for every $j \in \mathbb{N}$ and every $k = 1, \dots, K$. By adapting $C_{\mathbf{E}}$ if necessary, we can assume

$$\inf_{j \in \mathbb{N}} \min_{k=1, \dots, K} \min_{x \in \Omega} \det(D\phi_k^j(x)) > c_{\det}$$

for a constant $c_{\det} > 0$. Thus, we can conclude that $\{\Phi^j\}_{j \in \mathbb{N}}$ and Φ are in \mathcal{D}^K and hence consist of $C^1(\Omega, \Omega)$ -diffeomorphisms. Using Proposition 4.2.5, we can replace \mathbf{f}^j by the energy minimizing feature vector associated with Φ^j , which possibly reduces the path energy. The features $\{f_k^j\}_{j \in \mathbb{N}}$ are uniformly bounded in \mathcal{F}^{K+1} , which follows from an analogous reasoning as (4.22). Thus, $\mathbf{f}^j \rightarrow \mathbf{f}$ holds true for a subsequence (not relabeled) in \mathcal{F}^{K+1} , which implies $a[\mathcal{P}[f_k^j]] \rightarrow a[\mathcal{P}[f_k]]$ in $L^\infty(\Omega)$, due to (A2). Combining this with (W1), for every $k = 1, \dots, K$, we obtain

$$\liminf_{j \rightarrow \infty} \int_{\Omega} a[\mathcal{P}[f_k^j]] W(D\phi_k^j) dx \geq \int_{\Omega} a[\mathcal{P}[f_k]] W(D\phi_k) dx.$$

Finally, we verify the lower semicontinuity estimate

$$\|f_k \circ \phi_k - f_{k-1}\|_{\mathcal{F}}^2 \leq \liminf_{j \rightarrow \infty} \|f_k^j \circ \phi_k^j - f_{k-1}^j\|_{\mathcal{F}}^2, \quad (4.26)$$

for every $k = 1, \dots, K$. To this end, we take into account the decomposition

$$f_k^j \circ \phi_k^j - f_k \circ \phi_k = (f_k^j \circ \phi_k^j - f_k \circ \phi_k^j) + (f_k \circ \phi_k^j - f_k \circ \phi_k).$$

The second term is estimated in (4.20), and for the first term we take a test function $g \in \mathcal{F}$ and using the transformation formula obtain

$$\int_{\Omega} (f_k^j \circ \phi_k^j - f_k \circ \phi_k^j) \cdot g \, dx = \int_{\Omega} (f_k^j - f_k) \cdot (g(\det(D\phi_k^j))^{-1}) \circ (\phi_k^j)^{-1} \, dx.$$

The right-hand side converges to 0 due to the convergence $(\det(D\phi_k^j))^{-1} \circ (\phi_k^j)^{-1} \rightarrow \det(D\phi_k)^{-1} \circ \phi_k^{-1}$ in $L^\infty(\Omega)$ and $f_k^j \rightarrow f_k$ in \mathcal{F} . Thus, $f_k^j \circ \phi_k^j \rightarrow f_k \circ \phi_k$ for $j \rightarrow \infty$, which, together with the lower semicontinuity of the L^2 -norm, proves (4.26). Altogether, we observe that $\mathbf{E}^K[\mathbf{f}] \leq \mathbf{E}^{K,D}[\mathbf{f}, \Phi] \leq \liminf_{j \rightarrow \infty} \mathbf{E}^{K,D}[\mathbf{f}^j, \Phi^j]$, which finishes the proof of the theorem. \square

The following remark, similarly to Remark 4.2.4, represents a generalization of (4.26), to be used throughout this and the following chapter.

Remark 4.2.8. The reasoning in the previous paragraph, together with Remark 4.2.4, shows that, for $f^j \rightarrow f$ in \mathcal{F} , and diffeomorphisms $\{\phi^j\}_{j \in \mathbb{N}}$ and ϕ such that $\det D\phi^j \geq c_{\det} > 0$, for every (large enough) $j \in \mathbb{N}$, and $\phi^j \rightarrow \phi$ in $C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$, we have $f^j \circ \phi^j \rightarrow f \circ \phi$ in \mathcal{F} . Furthermore, we have

$$\int_{\Omega} |f^j \circ \phi^j - f \circ \phi^j|^2 \, dx = \int_{\Omega} |f^j - f|^2 (\det(D\phi^j))^{-1} \, dx \leq C \|f^j - f\|_{\mathcal{F}}^2.$$

which shows, together with the previous arguments, that $f^j \rightarrow f$ implies $f^j \circ \phi^j \rightarrow f \circ \phi$ and $f^j \circ (\phi^j)^{-1} \rightarrow f^j \circ \phi^{-1}$ in \mathcal{F} . These results are straightforwardly extendable to the corresponding Bochner spaces.

The following remark considers other possible ways of combining the image intensity values and the deep feature vectors.

Remark 4.2.9. (i) If, in addition to the condition of Remark 4.2.6, we assume that the mapping F is Lipschitz continuous, in the sense that there exists a constant L_F such that

$$\|F(u) - F(\tilde{u})\|_{L^2(\Omega, \mathbb{R}^c)} \leq L_F \|u - \tilde{u}\|_{\mathcal{I}}$$

for every $u, \tilde{u} \in \mathcal{I}$ (again satisfied if the convolutional weights are sufficiently smooth and we use the average pooling in the network architecture (cf. Table 4.1)), then the discrete geodesic exists also if \mathcal{F} is restricted to points of the type $(\eta u, F(u))$ (cf. (4.1)). To this end, in the proof of the theorem above, we define $\bar{f}_k := (\bar{u}_k, F(\frac{\bar{u}_k}{\eta}))$ with $\bar{u}_k := \frac{k}{K} u_A + (1 - \frac{k}{K}) u_B$, so that the condition (4.24) can be expressed in terms of proximity of the image parts of the vectors of features and the constant $C_{\mathbf{E}}$ additionally depends on L_F .

(ii) In [EKPR19], one considers the optimization problems

$$\inf_{\hat{\mathbf{f}} \in \mathcal{F}^{K-1}, \Phi \in \mathcal{D}^K} K \sum_{k=1}^K \mathcal{W}^D[a[u_k], f_{k-1}, f_k, \phi_k], \quad \inf_{\hat{\mathbf{u}} \in \mathcal{I}^{K-1}} \sum_{k=1}^K \frac{K}{\delta} \int_{\Omega} |u_k \circ \phi_k - u_{k-1}|^2 \, dx,$$

in parallel, by using the optimal $\hat{\mathbf{u}}$ for the first functional and the optimal Φ for the second. Here, $u_0, u_K = u_A, u_B \in \mathcal{I}$ and $f_0, f_K = f_A, f_B = F(u_A), F(u_B) \in \mathcal{F} = L^2(\Omega, \mathbb{R}^c)$. This approach gave similar experimental results to the one we obtained with the approach we advocated in this section, presented in Section 4.7. This is not a surprise if we notice that our approach corresponds to the optimization of the sum of the above functionals (with an additional scaling of the contribution of the image intensity channels). Theoretical study of our approach, in the context of the study of the existence (and later convergence) of discrete geodesic paths, is considerably more handsome from the technical and bookkeeping perspective.

Remark 4.2.10. Note that Proposition 4.2.3 and Theorem 4.2.7 remain true even without assumptions on proximity of the features (4.17) and (4.24), respectively, if the definition of the admissible set of deformations given by (4.10) is changed to

$$\mathcal{D} = \{\phi \in H^m(\Omega, \Omega) : \det(D\phi) \geq \epsilon \text{ a.e. in } \Omega, \phi|_{\partial\Omega} = \mathbf{1}\},$$

for some (small) $\epsilon > 0$.

Finally, let us notice that, as in the continuous time case, we do not have the uniqueness of discrete geodesics and the discrete reverse path is not a discrete geodesic if boundary data switch places.

4.3 Temporal Extension Operators

In this section, we provide a framework for the temporal extension of the time discrete quantities defined in the previous section. This is a necessary step to allow the study of convergence of the time discrete model to its time continuous counterpart.

For $K \in \mathbb{N}$, let $\Phi^K = (\phi_1^K, \dots, \phi_K^K) \in \mathcal{D}^K$ be a vector of deformations associated with a vector of features $\mathbf{f}^K = (f_0^K, \dots, f_K^K) \in \mathcal{F}^{K+1}$. We define the step size $\tau = \frac{1}{K}$, and, for $k = 1, \dots, K$, we take $t_k^K := k\tau$. Then, for $t \in [t_{k-1}^K, t_k^K]$ and $x \in \Omega$, we define the *discrete transport map*

$$y_k^K(t, x) := x + \frac{t - t_{k-1}^K}{\tau} (\phi_k^K(x) - x). \quad (4.27)$$

Note that $y_k^K(t_{k-1}^K, x) = x$ and $y_k^K(t_k^K, x) = \phi_k^K(x)$. To ensure that $y_k^K(t, \cdot)$ is invertible, it is enough to have $\det(Dy_k^K(t, x)) > 0$ for every $x \in \Omega$ (cf. [IKT13, Lemma A.1], [Cia88, Theorem 5.5.2]). This is insured if $\|Dy_k^K(t, \cdot) - \mathbb{1}\|_{C^0(\bar{\Omega})} < c$ for a small enough constant $c > 0$ (see reasoning after (4.18)). As this condition will be satisfied in our exposition in Section 4.4 (cf. (4.37)), we assume it to hold in the rest of this section. Then, we have $y_k^K(t, \cdot) \in \mathcal{D}$ and we define the *inverse discrete transport paths* $x_k^K(t, \cdot) \in \mathcal{D}$ as its spatial inverse. By the definition (4.27), we have $x - x_k^K(t, x) = K(t - t_{k-1}^K)(\phi_k^K - \mathbb{1})(x_k^K(t, x))$ which after some simple manipulation gives

$$\|Dx_k^K(t, \cdot) - \mathbb{1}\|_{C^{0,\alpha}(\bar{\Omega})} \leq C \|D\phi_k^K - \mathbb{1}\|_{C^{0,\alpha}(\bar{\Omega})}.$$

Since Ω is bounded and $x_{K,k}(t, \cdot)$ is a diffeomorphism, we get

$$\|x_k^K(t, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} \leq C + \|Dx_k^K(t, \cdot)\|_{C^{0,\alpha}(\bar{\Omega})} \leq C \left(1 + \max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^{1,\alpha}(\bar{\Omega})}\right). \quad (4.28)$$

We introduce the *feature extension operator* $\mathcal{F}^K[\mathbf{f}^K, \Phi^K] \in L^2([0, 1], \mathcal{F})$, given by

$$\mathcal{F}^K[\mathbf{f}^K, \Phi^K](t, x) := \left(f_{k-1}^K + K(t - t_{k-1}^K)(f_k^K \circ \phi_k^K - f_{k-1}^K)\right)(x_k^K(t, x)). \quad (4.29)$$

This describes, for $t \in [t_{k-1}^K, t_k^K]$, $k = 1, \dots, K$, a linear blending between the features

$$f_{k-1}^K = \mathcal{F}^K[\mathbf{f}^K, \Phi^K](t_{k-1}^K, \cdot) \quad \text{and} \quad f_k^K = \mathcal{F}^K[\mathbf{f}^K, \Phi^K](t_k^K, \cdot), \quad (4.30)$$

along the affine transport path $\{(t, y_k^K(t, \cdot))\}$. The *discrete motion field* $v^K \in L^2((0, 1), \mathcal{V})$ corresponding to the discrete transport map (4.27) is given by

$$v^K|_{(t_{k-1}^K, t_k^K]} := v_k^K := K(\phi_k^K - \mathbb{1}). \quad (4.31)$$

Furthermore, we define the *discrete velocity field along the discrete transport path* by

$$\tilde{v}^K(t, x) := v_k^K(x_k^K(t, x)), \quad t \in (t_{k-1}^K, t_k^K], \quad x \in \Omega. \quad (4.32)$$

By [Fio16, Propositions 1.2.4 and 1.2.7], we have $\tilde{v}^K \in C^{1,\alpha}(\bar{\Omega}, \mathbb{R}^n)$ and for $t \in (t_{k-1}^K, t_k^K]$ it holds (see also (3.15))

$$\|\tilde{v}_t^K\|_{C^{1,\alpha}(\bar{\Omega})} \leq \|v_t^K\|_{C^{1,\alpha}(\bar{\Omega})} \left(1 + \|x_k^K(t, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})} + \|x_k^K(t, \cdot)\|_{C^{1,\alpha}(\bar{\Omega})}^{\alpha+1}\right). \quad (4.33)$$

The *discrete flow* $\psi^K : [0, 1] \times \bar{\Omega} \rightarrow \bar{\Omega}$ is defined as the concatenation of all small diffeomorphisms y_k^K along the motion path. In detail, we have, for all $x \in \Omega$

$$\psi^K(t, x) := y_1^K(t, x), \quad t \in [0, t_1^K]; \quad \psi^K(t, x) := y_k^K(t, \psi^K(t_{k-1}^K, x)), \quad t \in (t_{k-1}^K, t_k^K], \quad k = 2, \dots, K. \quad (4.34)$$

The *(scalar) material derivative* $z^K \in L^2((0, 1) \times \Omega)$ for $t \in (t_{k-1}^K, t_k^K]$ is defined as

$$z^K(t, x) := K \left| (f_k^K \circ \phi_k^K - f_{k-1}^K)(x_k^K(t, x)) \right|. \quad (4.35)$$

The following proposition shows that the temporal extensions of the features, the velocities, the material derivatives, and the discrete paths are indeed admissible for the problem, *i.e.*, that the conditions (4.6) and (4.7) are satisfied.

Proposition 4.3.1. For $\mathbf{f}^K \in \mathcal{F}^{K+1}$ and deformations $\Phi^K \in \mathcal{D}^K$, the quantities $\mathcal{F}^K[\mathbf{f}^K, \Phi^K]$, \tilde{v}^K , ψ^K , and z^K satisfy (4.6) and (4.7).

Proof. By the definition (4.34), we have $\psi^K(0, x) = x$ for all $x \in \Omega$. For $t \in (t_{k-1}^K, t_k^K)$ and $x \in \Omega$, we get

$$\frac{d}{dt}\psi^K(t, x) = \frac{d}{dt}y_k^K(t, \psi^K(t_{k-1}^K, x)) = K(\phi_k^K - \mathbf{1})(\psi^K(t_{k-1}^K, x)) = \tilde{v}^K(t, \psi^K(t, x)),$$

since $\psi^K(t_{k-1}^K, x) = x_k^K(t, \psi^K(t, x))$. Therefore, ψ^K is a solution of (4.6) in the weak sense according to Remark 3.2.2. By (4.35), for $s \leq t \in [t_{k-1}^K, t_k^K]$, we have

$$\begin{aligned} & |\mathcal{F}^K[\mathbf{f}^K, \Phi^K](t, \psi^K(t, x)) - \mathcal{F}^K[\mathbf{f}^K, \Phi^K](s, \psi^K(s, x))| \\ &= K(t-s)|f_{k-1}^K(\psi^K(t_{k-1}^K, x)) - f_k^K \circ \phi_k^K(\psi^K(t_{k-1}^K, x))| \\ &\leq \int_s^t z^K(r, \psi^K(r, x)) dr. \end{aligned}$$

If s and t are not in the same interval, we can use the triangle inequality multiple times, which concludes the proof. \square

As a corollary of this proposition, we note that, by Remark 4.1.4, we have $\mathcal{F}^K[\mathbf{f}^K, \Phi^K] \in C^0([0, 1], \mathcal{F})$ and $a^K(t, x) := a[\mathcal{P}[\mathcal{F}^K[\mathbf{f}^K, \Phi^K](t, \cdot)](x)$ is well-defined.

Finally, we define an extension $\mathcal{E}^K : L^2([0, 1], \mathcal{F}) \rightarrow [0, \infty]$ of the discrete path energy \mathbf{E}^K (cf. (4.12)) by

$$\mathcal{E}^K[f] := \inf_{\Phi^K \in \mathcal{D}^K} \left\{ \mathbf{E}^{K,D}[\mathbf{f}^K, \Phi^K] : \mathcal{F}^K[\mathbf{f}^K, \Phi^K] = f \right\}, \quad (4.36)$$

if there exist $\mathbf{f}^K \in \mathcal{F}^{K+1}$ and $\Phi^K \in \mathcal{D}^K$ such that $f = \mathcal{F}^K[\mathbf{f}^K, \Phi^K]$, and else we set $\mathcal{E}^K[f] = \infty$. Notice that, for given f and K , by (4.30), we have only one candidate for \mathbf{f}^K obtained by sampling of f at points $\frac{k}{K}, k = 0, \dots, K$. Thus, we only need to ensure the well-posedness of the defined energy w.r.t. the vector of deformations. To achieve this for every $K \in \mathbb{N}$ without additional assumptions on f , here and in the rest of this chapter we restrict the admissible set of deformations as in Remark 4.2.10.

Lemma 4.3.2. If, for a given fixed feature path $f \in L^2((0, 1), \mathcal{F})$, a discrete feature path $\mathbf{f}^K \in \mathcal{F}^{K+1}$ and a vector of deformations $\Phi^K \in \mathcal{D}^K$ exists such that $f = \mathcal{F}^K[\mathbf{f}^K, \Phi^K]$, then the infimum with respect to the vector of deformations in (4.36) is attained for some $\bar{\Phi}^K \in \mathcal{D}^K$.

Proof. Let $\{\Phi_j^K\}_{j \in \mathbb{N}} \subset \mathcal{D}^K$ be a minimizing sequence for $\Phi^K \mapsto \mathbf{E}^{K,D}[\mathbf{f}^K, \Phi^K]$, which satisfies the constraint $\mathcal{F}^K[\mathbf{f}^K, \Phi_j^K] = f$ for every $j \in \mathbb{N}$. With the arguments analogous to those from proofs of Lemma 4.2.2 and Proposition 4.2.3, we can prove that (on a subsequence) $\Phi_j^K \rightharpoonup \bar{\Phi}^K$ in $H^m(\Omega, \Omega)^K$ and $\Phi_j^K \rightarrow \bar{\Phi}^K$ in $C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})^K$, where $\bar{\Phi}^K \in \mathcal{D}^K$, by the new definition of \mathcal{D} . The weak lower semicontinuity and the coercivity of $\Phi^K \mapsto \mathbf{E}^{K,D}[\mathbf{f}^K, \Phi^K]$ are shown as in the proof of Proposition 4.2.3. Hence, it remains to prove the closedness of the equality constraint $f = \mathcal{F}^K[\mathbf{f}^K, \bar{\Phi}^K]$. In light of (4.28), we have that \bar{x}^K corresponding to $\bar{\Phi}^K$ exists and $x_j^K(t, \cdot) \rightarrow \bar{x}^K(t, \cdot)$ in $C^{1,\alpha}(\bar{\Omega}, \bar{\Omega})$, uniformly in t . Then, using Remarks 4.2.4 and 4.2.8, we have

$$\lim_{j \rightarrow \infty} \|\mathcal{F}^K[\mathbf{f}^K, \Phi_j^K](t, \cdot) - \mathcal{F}^K[\mathbf{f}^K, \bar{\Phi}^K](t, \cdot)\|_{\mathcal{F}} = 0, \quad \forall t \in [0, 1],$$

from where we conclude $\mathcal{F}^K[\mathbf{f}^K, \bar{\Phi}^K] = f$. \square

4.4 Convergence of Discrete Geodesic Paths

In this section, we prove the Mosco-convergence (cf. Definition 2.4.1) for $K \rightarrow \infty$ of the energy \mathcal{E}^K (as a temporal extension of the time discrete path energy \mathbf{E}^K) to the time continuous path energy \mathcal{E} , introduced in Definition 4.1.2. Furthermore, the convergence of the time discrete to time continuous geodesic paths is established, which in particular implies the existence of time continuous geodesics in the deep feature metamorphosis model with an anisotropic regularizer.

Theorem 4.4.1 (Mosco-convergence of the discrete path energies). *Let (W1)–(W3) and (A1)–(A3) be satisfied. Then, the sequence of time discrete path energies $\{\mathcal{E}^K\}_{K \in \mathbb{N}}$ Mosco-converges to \mathcal{E} in the $L^2([0, 1], \mathcal{F})$ -topology for $K \rightarrow \infty$, i.e.,*

(i) *for every sequence $\{f^K\}_{K \in \mathbb{N}} \subset L^2([0, 1], \mathcal{F})$ with $f^K \rightharpoonup f \in L^2([0, 1], \mathcal{F})$, the estimate*

$$\liminf_{K \rightarrow \infty} \mathcal{E}^K[f^K] \geq \mathcal{E}[f]$$

holds true (“lim inf-inequality”),

(ii) *for every $f \in L^2([0, 1], \mathcal{F})$, there exists a recovery sequence $\{f^K\}_{K \in \mathbb{N}} \subset L^2([0, 1], \mathcal{F})$ satisfying $f^K \rightarrow f$ in $L^2([0, 1], \mathcal{F})$ such that the estimate*

$$\limsup_{K \rightarrow \infty} \mathcal{E}^K[f^K] \leq \mathcal{E}[f]$$

is valid (“the existence of a recovery sequence and lim sup-inequality”).

Proof. We will prove the conditions separately.

Proof of the lim inf-inequality

To facilitate reading, we briefly outline the structure of the proof.

- (i) *Construction of the flow and feature intensity quantities.* We observe that we can restrict our proof to the sequences of the type $\{f^K = \mathcal{F}^K[\mathbf{f}^K, \Phi^K]\}_{K \in \mathbb{N}}$, with the uniformly bounded energy. Having these, we can construct the time extensions of the time discrete quantities corresponding to the vectors of features and deformations \mathbf{f}^K and Φ^K , respectively. In particular, this gives rise to the sequence of velocity fields $\{v^K\}_{K \in \mathbb{N}} \subset L^2((0, 1), \mathcal{V})$ and the discrete material derivatives $\{z^K\}_{K \in \mathbb{N}} \subset L^2((0, 1) \times \Omega)$.
- (ii) *Identification of the limit of the sequences of the discrete velocity fields and material derivatives.* We prove the uniform boundedness of these sequences in the corresponding spaces, leading to the existence of the weak limits v and z , respectively. We show that these quantities are indeed admissible for the feature curve f , i.e., $(v, z) \in \mathcal{C}(f)$. This proof is based on the properties of the flow equation (cf. Subsection 3.2.3) and the continuity of time point evaluation along the feature curves, which is due to the improved regularity (cf. Remark 4.1.4).
- (iii) *Proof of the actual inequality.* In this step, we prove the lim inf-inequality, by using the weak lower semi-continuity of the energy and the Taylor expansion of the energy density W (cf. (W3)). To finish this step, we need to handle the anisotropy terms, which is the task of the next step.
- (iv) *Identification of the limit of the discrete anisotropy weights.* In the final step, we show the convergence in $L^\infty([0, 1] \times \Omega)$ of $\{\bar{a}^K\}_{K \in \mathbb{N}}$, as the piecewise constant interpolation of the vector of discrete anisotropic weights $\{a_k^K\}_{k=1, \dots, K; K \in \mathbb{N}}$, to $a[\mathcal{P}[f]]$: the anisotropy regularization corresponding to the feature curve f . To this end, we again use the continuity of time point evaluation along the feature curves, together with the compactness and Lipschitz continuity of the anisotropy operator (cf. (A2)–(A3)).

ad (i): *Construction of the flow and feature intensity quantities.*

Let $\{f^K\}_{K \in \mathbb{N}} \subset L^2([0, 1], \mathcal{F})$ be a sequence weakly converging to a feature path $f \in L^2([0, 1], \mathcal{F})$. If we exclude the trivial case $\liminf_{K \rightarrow \infty} \mathcal{E}^K[f^K] = \infty$ and eventually pass to a subsequence (without relabeling), we may assume $\mathcal{E}^K[f^K] \leq \bar{\mathcal{E}} < \infty$, for all $K \in \mathbb{N}$. By definition of \mathcal{E}^K , this directly implies

$$f^K = \mathcal{F}^K[\mathbf{f}^K, \Phi^K], \quad \mathbf{f}^K = (f_0^K, \dots, f_K^K) \in \mathcal{F}^{K+1}, \quad \Phi^K = (\phi_1^K, \dots, \phi_K^K) \in \mathcal{D}^K,$$

where Φ^K is the minimizing deformation in (4.36), whose existence is proven in Lemma 4.3.2. Given this vector of deformations, we define the discrete transport path y^K and the discrete velocity field v^K as in (4.27) and (4.31), respectively. Using Lemma 4.2.2, for every $t \in (0, 1)$, we get

$$\begin{aligned} \max_{k=1, \dots, K} \left\{ \|y_k^K(t, \cdot) - \mathbb{1}\|_{C^{1,\alpha}(\bar{\Omega})}, \|\phi_k^K - \mathbb{1}\|_{C^{1,\alpha}(\bar{\Omega})} \right\} &\leq C \max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{H^m(\Omega)} \\ &\leq C\theta(\bar{\mathcal{E}}K^{-1}) \leq CK^{-\frac{1}{2}}. \end{aligned} \quad (4.37)$$

Furthermore, by using the Lipschitz continuity of the determinant function on the ball $B_r(\mathbb{1})$ with associated radius $r = \sup_K \max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{H^m(\Omega)} < \infty$, we have

$$\max_{k=1, \dots, K} \|\det(Dy_k^K) - 1\|_{L^\infty([t_{k-1}^K, t_k^K] \times \Omega)} \leq CK^{-\frac{1}{2}}, \quad (4.38)$$

where the right-hand side is strictly smaller than 1 for every large enough K . This implies the positivity of the determinant and thus the existence of the inverse transport path x^K . This means that the feature extension (4.29) is indeed well-defined and we can further define the quantities \tilde{v}^K , ψ^K and z^K , associated with Φ^K (see (4.32), (4.34) and (4.35), respectively).

ad (ii): *Identification of the limit of the sequences of the discrete material derivatives and velocity fields.*

For the weak material derivative given by (4.35), we have

$$\begin{aligned} \int_0^1 \int_\Omega (z^K)^2 dx dt &= \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_\Omega K^2 |f_{k-1}^K(x_k^K(t, x)) - f_k^K \circ \phi_k^K(x_k^K(t, x))|^2 dx dt \\ &= \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_\Omega K^2 |f_{k-1}^K(x) - f_k^K \circ \phi_k^K(x)|^2 \det(Dy_k^K(t, x)) dx dt. \end{aligned} \quad (4.39)$$

From (4.37) and $t_k^K - t_{k-1}^K = K^{-1}$, we obtain

$$\left| \sum_{k=1}^K K^2 \int_{t_{k-1}^K}^{t_k^K} \int_\Omega |f_{k-1}^K(x) - f_k^K \circ \phi_k^K(x)|^2 (\det(Dy_k^K(t, x)) - 1) dx dt \right| \leq \delta \bar{\mathcal{E}} CK^{-\frac{1}{2}}.$$

Plugging this back to (4.39), we get

$$\lim_{K \rightarrow \infty} \int_0^1 \int_\Omega (z^K)^2 dx dt = \lim_{K \rightarrow \infty} K \sum_{k=1}^K \int_\Omega |f_{k-1}^K(x) - f_k^K \circ \phi_k^K(x)|^2 dx. \quad (4.40)$$

This also shows the uniform boundedness of $\{z^K\}_{K \in \mathbb{N}} \subset L^2((0, 1) \times \Omega)$, which implies the existence of a weakly convergent subsequence with a limit $z \in L^2((0, 1) \times \Omega)$. Finally, using the weak lower semicontinuity of the norm we have

$$\int_0^1 \int_\Omega z^2 dx dt \leq \liminf_{K \rightarrow \infty} \int_0^1 \int_\Omega (z^K)^2 dx dt = \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_\Omega |f_{k-1}^K(x) - f_k^K \circ \phi_k^K(x)|^2 dx. \quad (4.41)$$

For the sequence of discrete velocity fields $\{v^K\}_{K \in \mathbb{N}}$, we show the uniform boundedness in $L^2((0, 1), \mathcal{V})$. To this end, we first assume that K is sufficiently large such that $\max_{k=1, \dots, K} \|D\phi_k^K - \mathbb{1}\|_{C^0(\bar{\Omega})} < r_W$ (see (W2)), which is possible due to (4.37). Then, using the Korn's inequality 2.1.6, the Poincaré's inequality 2.1.4, as well as (W2) and (A1), we obtain

$$\begin{aligned} \int_0^1 \int_\Omega |v^K|^2 dx dt &\leq C \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_\Omega K^2 |\varepsilon[\phi_k^K] - \mathbb{1}|^2 dx dt \\ &\leq CK \sum_{k=1}^K \int_\Omega \frac{W(D\phi_k^K)}{C_{W,1}} dx \leq \frac{C\bar{\mathcal{E}}}{C_{W,1}c_a}, \\ \int_0^1 \int_\Omega |D^m v^K|^2 dx dt &= \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_\Omega K^2 |D^m(\phi_k^K - \mathbb{1})|^2 dx dt \\ &= \sum_{k=1}^K K \int_\Omega |D^m \phi_k^K|^2 dx \leq \frac{\bar{\mathcal{E}}}{\gamma}. \end{aligned} \quad (4.42)$$

Finally, the uniform boundedness is implied by the Gagliardo–Nirenberg inequality 2.1.5. By passing to a subsequence (again labeled in the same way), we can deduce $v^K \rightharpoonup v \in L^2((0, 1), \mathcal{V})$ as $K \rightarrow \infty$.

The next step is to show the admissibility of the limit, *i.e.*, $(v, z) \in \mathcal{C}(f)$. Recall (*cf.* (4.6), (4.7)) that this means that

$$\dot{\psi}_t(x) = v_t \circ \psi_t(x) \text{ for } (t, x) \in [0, 1] \times \Omega, \quad \psi_0(x) = x \text{ for } x \in \Omega, \quad (4.43)$$

$$|f_s \circ \psi_s(x) - f_t \circ \psi_t(x)| \leq \int_s^t z_r \circ \psi_r(x) \, dr, \quad \forall s < t \in [0, 1], \text{ for a.e. } x \in \Omega. \quad (4.44)$$

The existence of a unique solution to (4.43), which for the moment we denote by ψ^v , is guaranteed by Theorem 3.2.1. What we show here is that it can be identified as the limit as $K \rightarrow \infty$ of the sequence of discrete flows $\{\psi^K\}_{K \in \mathbb{N}}$ which are given by (4.34) and present the flow associated to $\{\tilde{v}^K\}_{K \in \mathbb{N}}$, by Proposition 4.3.1. We start by observing that (4.28) and (4.33), together with (4.37), and the uniform boundedness of the sequence of $\{v^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1), \mathcal{V})$, imply that the sequence $\{\tilde{v}^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $L^2((0, 1), C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$. Then, incorporating Remark 3.2.2, we infer that $\{\psi^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $C^0([0, 1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$, and by Remark 3.2.3, uniformly bounded in $C^{0,\frac{1}{2}}([0, 1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$. Hence, by using the compact embedding of Hölder spaces (*cf.* Corollary 2.1.10), a subsequence of $\{\psi^K\}_{K \in \mathbb{N}}$ converges strongly to some ψ in $C^{0,\beta}([0, 1], C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$, for every $\beta \in (0, \min(\frac{1}{2}, \alpha))$. We will verify that $\psi = \psi^v$. To this end, we denote the solutions of (4.43) corresponding to v^K by ψ^{v^K} , where their existence again follows by Theorem 3.2.1. Then,

$$\|\psi - \psi^v\|_{C^0([0,1] \times \bar{\Omega})} \leq \|\psi - \psi^K\|_{C^0([0,1] \times \bar{\Omega})} + \|\psi^K - \psi^{v^K}\|_{C^0([0,1] \times \bar{\Omega})} + \|\psi^{v^K} - \psi^v\|_{C^0([0,1] \times \bar{\Omega})}.$$

Here, the first term converges to zero as shown above and the last term converges to zero by the continuous dependence of ψ^{v^K} on v^K , discussed in Theorem 3.2.1. We can estimate the remaining term as follows

$$\begin{aligned} \|\psi^K - \psi^{v^K}\|_{C^0([0,1] \times \bar{\Omega})} &\leq C \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \|v_k^K(s, x_k^K(s, \cdot)) - v_k^K(s, \cdot)\|_{C^0(\bar{\Omega})} \, ds \\ &\leq C \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \|v_k^K(s, \cdot)\|_{H^m(\Omega)} \|y_k^K(s, \cdot) - \mathbb{1}\|_{C^0(\bar{\Omega})} \, ds \\ &\leq C \|v^K\|_{L^2((0,1), H^m(\Omega))} \max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^0(\bar{\Omega})}. \end{aligned}$$

Here, the first inequality is deduced from Remark 3.2.2. To derive the second inequality, we exploit the Lipschitz property of $x \mapsto v_k^K(s, x_k^K(s, x)) - v_k^K(s, x)$, where the Lipschitz constant, by Sobolev embedding 2.1.8, is bounded by $C \|v_k^K(s, \cdot)\|_{H^m(\Omega)}$. Furthermore, we apply the coordinate transform $y_k^K(s, \cdot)$ and use the boundedness of the determinant of $Dy_k^K(s, \cdot)$, shown in (4.38). Using the uniform boundedness of $\{v^K\}_{K \in \mathbb{N}}$, (4.37), and plugging this back above, we get $\psi = \psi^v$.

With ψ identified as above, we want to show that (4.44) is fulfilled. We first note that, for $s, t \in [0, 1]$, we have

$$\begin{aligned} \|f_t^K \circ \psi_t^K - f_s^K \circ \psi_s^K\|_{\mathcal{F}}^2 &\leq \int_{\Omega} \left(\int_s^t z_r^K \circ \psi_r^K(x) \, dr \right)^2 dx \\ &\leq |t - s| \left| \int_{\Omega} \int_s^t (z_r^K \circ \psi_r^K(x))^2 \, dr \, dx \right| \leq C |t - s|, \quad (4.45) \end{aligned}$$

where we used the uniform boundedness of $\{z^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1) \times \Omega)$ (*cf.* (4.40)), and $\{(\psi^{-1})^K\}_{K \in \mathbb{N}}$ in $C^0([0, 1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$, following from Theorem 3.2.1. Thus, we conclude that there exists a positive constant L such that $\{f^K \circ \psi^K\}_{K \in \mathbb{N}} \subset C_{\omega}^0([0, 1], \mathcal{F})$, for $\omega(t) := Lt^{\frac{1}{2}}$, *i.e.*, all the elements of the sequence have a joint modulus of continuity. The space $C_{\omega}^0([0, 1], \mathcal{F})$ as a subset of $L^2([0, 1], \mathcal{F})$ has some closedness and regularity properties, as shown in the following lemma.

Lemma 4.4.2. *Let $\omega \in L^2([0, 1], \mathbb{R}^+)$ be monotonically increasing and continuous at 0 with $\omega(0) = 0$. Then, the point evaluation $f \mapsto f_t$ is well-defined on $C_\omega^0([0, 1], \mathcal{F})$ and it represents a bounded linear functional mapping from $L^2([0, 1], \mathcal{F})$ to \mathcal{F} . Furthermore, the set $C_\omega^0([0, 1], \mathcal{F})$ is closed with respect to the weak convergence in $L^2([0, 1], \mathcal{F})$.*

Proof. Let $f \in C_\omega^0([0, 1], \mathcal{F})$. Then, for every $t, s \in [0, 1]$, we have $\|f_t\|_{\mathcal{F}} \leq \omega(|t - s|) + \|f_s\|_{\mathcal{F}}$ and thus, for every $t \in [0, 1]$, we have

$$\sup_{\|f\|_{L^2([0, 1], \mathcal{F})} = 1} \|f_t\|_{\mathcal{F}} \leq 2\|\omega\|_{L^2([0, 1])} + 1,$$

showing the boundedness of point evaluation $f \mapsto f_t$.

To show the weak closedness, by [Bre11, Theorem 3.7], it is enough to show the convexity and strongly closedness of the set. The convexity follows by the triangle inequality. To prove closedness, take a sequence $\{f^j\}_{j \in \mathbb{N}} \subset C_\omega^0([0, 1], \mathcal{F})$ such that $f^j \rightarrow f$ in $L^2([0, 1], \mathcal{F})$. Then, by [Alt06, Theorem 3.22], we have that $\{f^j\}_{j \in \mathbb{N}}$ converges to f (on a subsequence with the same label) pointwise almost everywhere. Thus, for every $t \in [0, 1]$, we can choose an arbitrarily close point s such that $\|f_s^j - f_s\|_{\mathcal{F}} \rightarrow 0$ as $j \rightarrow \infty$. Then, for every $i, j \in \mathbb{N}$ sufficiently large, we have

$$\|f_t^i - f_t^j\|_{\mathcal{F}} \leq 2\omega(|t - s|) + \|f_s^i - f_s^j\|_{\mathcal{F}},$$

confirming that the sequence actually converges for every $t \in [0, 1]$. Then, for every $t, s \in [0, 1]$, we have

$$\|f_t - f_s\|_{\mathcal{F}} = \lim_{j \rightarrow \infty} \|f_t^j - f_s^j\|_{\mathcal{F}} \leq \omega(|t - s|),$$

showing that $f \in C_\omega^0([0, 1], \mathcal{F})$. □

To use these observations, we first note that, by Remark 4.2.8, $f^K \circ \psi^K \rightharpoonup f \circ \psi$ in $L^2((0, 1), \mathcal{F})$, so that $f \circ \psi \in C_\omega^0([0, 1], \mathcal{F})$. From the continuity of point evaluation in time, we have that, for every $s \leq t \in [0, 1]$ and every $\tilde{\Omega} \subset \Omega$, the functional $f \mapsto \int_{\tilde{\Omega}} |f_t(x) - f_s(x)| dx$ is continuous, and since it is also convex it is weak lower semicontinuous [Bre11, Corollary 3.9]. Altogether, we obtain

$$\begin{aligned} \int_{\tilde{\Omega}} |f_s \circ \psi_s(x) - f_t \circ \psi_t(x)| dx &\leq \liminf_{K \rightarrow \infty} \int_{\tilde{\Omega}} |f_t^K \circ \psi_t^K(x) - f_s^K \circ \psi_s^K(x)| dx \\ &\leq \liminf_{K \rightarrow \infty} \int_{\tilde{\Omega}} \int_s^t z_r^K \circ \psi_r^K(x) dr dx = \int_{\tilde{\Omega}} \int_s^t z_r \circ \psi_r(x) dr dx, \end{aligned} \quad (4.46)$$

which proves (4.44), as it holds for every $\tilde{\Omega} \subset \Omega$. Here, the last equality follows from the weak convergence of z^K , combined with the strong convergence of ψ^K , which, by Remark 4.2.8, implies the weak convergence of $z^K \circ \psi^K$ in $L^2((0, 1) \times \Omega)$.

ad (iii): *Proof of the actual inequality.*

Having (4.41), we are left to show

$$\int_0^1 \int_{\Omega} L[a, v, v] dx dt \leq \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} a_k^K W(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 dx, \quad (4.47)$$

where $a := a[\mathcal{P}[f]]$ is well-defined by Remark 4.1.4, and $a_k^K := a[\mathcal{P}[f_k^K]]$. We make the second order Taylor expansion around t_{k-1}^K of the function $t \mapsto W(\mathbb{1} + (t - t_{k-1}^K)Dv_k^K)$ and evaluate it at $t = t_k^K$, yielding

$$\begin{aligned} W(D\phi_k^K) &= W(\mathbb{1}) + K^{-1}DW(\mathbb{1})(Dv_k^K) + \frac{1}{2K^2}D^2W(\mathbb{1})(Dv_k^K, Dv_k^K) + r_k^K \\ &= K^{-2} \left(\frac{\lambda}{2} (\text{tr}(\varepsilon[v_k^K]))^2 + \mu \text{tr}(\varepsilon[v_k^K]^2) \right) + r_k^K, \end{aligned} \quad (4.48)$$

where r_k^K denotes the remainder. Here, the lower order terms vanish due to (W1), and the last equality follows from (W3). Then,

$$\begin{aligned} & K \sum_{k=1}^K \int_{\Omega} a_k^K \mathbb{W}(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 dx \\ &= K^{-1} \sum_{k=1}^K \int_{\Omega} a_k^K \left(\frac{\lambda}{2} (\text{tr}[\varepsilon[v_k^K]])^2 + \mu \text{tr}[\varepsilon[v_k^K]^2] \right) + \gamma |D^m v_k^K|^2 dx + K \sum_{k=1}^K \int_{\Omega} a_k^K r_k^K dx, \end{aligned}$$

where the remainder is of order $K^{-\frac{1}{2}}$. To see this, we apply (4.37), Lemma 4.2.2, (A1), and the uniform bound on the energy, to deduce

$$\begin{aligned} & K \sum_{k=1}^K \int_{\Omega} |a_k^K r_k^K| dx \leq CC_a K \sum_{k=1}^K \int_{\Omega} K^{-3} |Dv_k^K|^3 dx \\ & \leq CK \max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^1(\bar{\Omega})} \sum_{k=1}^K \|\phi_k^K - \mathbb{1}\|_{H^m(\Omega)}^2 \\ & \leq CK \theta(\bar{\mathcal{E}}K^{-1}) \sum_{k=1}^K \theta(\mathcal{W}^D(f_{k-1}^K, f_k^K, \phi_k^K))^2 \leq CK^{\frac{1}{2}} \sum_{k=1}^K \mathcal{W}^D(f_{k-1}^K, f_k^K, \phi_k^K) \leq C\bar{\mathcal{E}}K^{-\frac{1}{2}}. \end{aligned}$$

From here we see that, in the order to prove (4.47), it is enough to prove $\bar{a}^K \rightarrow a$ in $L^\infty([0, 1] \times \Omega)$, where

$$\bar{a}^K(t) := a[\mathcal{P}[\bar{f}^K]] = a_k^K \text{ for } t \in (t_{k-1}^K, t_k^K], \text{ where } \bar{f}^K(t) := f_k^K \text{ for } t \in (t_{k-1}^K, t_k^K]. \quad (4.49)$$

Indeed, then $\sqrt{\bar{a}^K} v^K \rightarrow \sqrt{a} v$ in $L^2((0, 1), \mathcal{V})$, and, as the functional is strongly continuous and convex, it is also weakly lower semicontinuous [Bre11, Proposition 3.5 and Corollary 3.9], implying

$$\begin{aligned} & \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} a_k^K \mathbb{W}(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 dx \\ &= \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} \bar{a}^K \left(\frac{\lambda}{2} (\text{tr}[\varepsilon[v^K]])^2 + \mu \text{tr}[\varepsilon[v^K]^2] \right) + \gamma |D^m v^K|^2 dx dt \\ &\geq \int_0^1 \int_{\Omega} a \left(\frac{\lambda}{2} (\text{tr}[\varepsilon[v]])^2 + \mu \text{tr}[\varepsilon[v]^2] \right) + \gamma |D^m v|^2 dx dt. \end{aligned} \quad (4.50)$$

ad (iv): *Identification of the limit of the discrete anisotropy weights.*

To this end, we will show

$$\bar{a}^K - a^K \rightarrow 0, \quad a^K \rightarrow a \text{ in } L^\infty([0, 1] \times \Omega), \text{ as } K \rightarrow \infty, \quad (4.51)$$

where $a^K := a[\mathcal{P}[f^K]]$ is well-defined by Remark 4.1.4. To prove the second claim, we first observe that, due to the uniform boundedness of $\{\tilde{v}^K\}_{K \in \mathbb{N}}$ and $\{z^K\}_{K \in \mathbb{N}}$, we can show that (cf. (4.8)) $\{f^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $L^\infty([0, 1], \mathcal{F})$. Then, by Remark 4.1.4, we have that there exists a modulus of continuity ω such that for every $K \in \mathbb{N}$ it holds $f^K \in C_\omega^0([0, 1], \mathcal{F})$. In particular, by Lemma 4.4.2, on this sequence the point evaluation $f \mapsto f_t$ is well-defined, linear, and continuous as a mapping from $L^2([0, 1], \mathcal{F})$ to \mathcal{F} . Hence, by [Bre11, Theorem 3.10], it is also weakly continuous, implying that

$$f_t^K \rightharpoonup f_t \text{ in } \mathcal{F}, \quad \forall t \in [0, 1], \quad (4.52)$$

which, together with (A2), gives

$$a_t^K \rightarrow a_t \text{ in } L^\infty(\Omega), \quad \forall t \in [0, 1]. \quad (4.53)$$

From the existence of joint modulus of continuity for $\{f^K\}_{K \in \mathbb{N}}$, together with (A3), we have the equicontinuity of the sequence of anisotropies, *i.e.*,

$$\|a_{t+t'}^K - a_t^K\|_{L^\infty(\Omega)} \rightarrow 0, \text{ as } t' \rightarrow 0, \text{ uniformly in } K \text{ and } t. \quad (4.54)$$

Hence, by Arzelà-Ascoli Theorem 2.1.9 (ii), we have $a^K \rightarrow a$ in $C^0([0, 1], L^\infty(\Omega))$.

To prove the second claim in (4.51), by (4.49) and (A3), it is enough to prove $\bar{f}^K - f^K \rightarrow 0$ in $L^\infty([0, 1], \mathcal{F})$. To this end, using the definition of feature extension, for $t \in [t_{k-1}^K, t_k^K]$, we write

$$|\bar{f}^K(t, x) - f^K(t, x)| \leq |f_k^K \circ \phi_k^K(x_k^K(t, x)) - f_{k-1}^K(x_k^K(t, x))| + |f_k^K \circ \phi_k^K(x_k^K(t, x)) - f_k^K(x)|.$$

For the first term, using uniform boundedness of Dy_k^K in $L^\infty([t_{k-1}^K, t_k^K] \times \Omega)$ (*cf.* (4.37) and (4.38)), we have

$$\begin{aligned} \int_{\Omega} |f_k^K \circ \phi_k^K(x_k^K(t, x)) - f_{k-1}^K(x_k^K(t, x))|^2 dx &= \int_{\Omega} |f_k^K \circ \phi_k^K(x) - f_{k-1}^K(x)|^2 \det(Dy_k^K(t, x)) dx \\ &\leq C\bar{\mathcal{E}}K^{-1}. \end{aligned} \quad (4.55)$$

To deal with the second term, we can take a sequence $f_{k,j}^K \in C^\infty(\Omega, \mathbb{R}^{3+c})$ such that $\|f_{k,j}^K - f_k^K\|_{\mathcal{F}} \leq 2^{-j}$ and $\|Df_{k,j}^K\|_{L^\infty(\Omega)} \leq C_{f,j}$ (*cf.* Theorem 2.1.7). The uniformity of the estimates follows from the uniform boundedness of $\{f_k^K\}_{k=1, \dots, K; K \in \mathbb{N}} \in \mathcal{F}$, which is due to the uniform boundedness of $f_K^K = f^K(1)$ deduced above and the induction procedure (*cf.* (4.22)). This gives an estimate

$$\begin{aligned} &\|f_k^K \circ \phi_k^K(x_k^K(t, \cdot)) - f_k^K(\cdot)\|_{\mathcal{F}} \quad (4.56) \\ &\leq \|f_k^K \circ \phi_k^K(x_k^K(t, \cdot)) - f_{k,j}^K \circ \phi_k^K(x_k^K(t, \cdot))\|_{\mathcal{F}} + \|f_{k,j}^K \circ \phi_k^K(x_k^K(t, \cdot)) - f_{k,j}^K\|_{\mathcal{F}} + \|f_{k,j}^K - f_k^K\|_{\mathcal{F}} \\ &\leq \|f_k^K - f_{k,j}^K\|_{\mathcal{F}} \|\det(D(\phi_k^K \circ x_k^K))\|_{L^\infty([t_{k-1}^K, t_k^K] \times \Omega)}^{-\frac{1}{2}} + \|f_{k,j}^K - f_k^K\|_{\mathcal{F}} \\ &\quad + C\|Df_{k,j}^K\|_{L^\infty(\Omega)} \|\phi_k^K - y_k^K\|_{L^\infty([t_{k-1}^K, t_k^K] \times \bar{\Omega})} \|\det(Dy_k^K)\|_{L^\infty([t_{k-1}^K, t_k^K] \times \bar{\Omega})} \\ &\leq C2^{-j}\epsilon^{-\frac{1}{2}} + C_{f,j}\|\phi_k^K - \mathbf{1}\|_{L^\infty([t_{k-1}^K, t_k^K] \times \bar{\Omega})} \leq C2^{-j} + C_{f,j}K^{-\frac{1}{2}}, \end{aligned}$$

where we used the transformation formula, (4.37), and (4.38). Choosing first j and then K , this proves the claim and finally finishes the proof of \liminf -estimate.

In what follows, we prove the existence of a recovery sequence and thus establish the Mosco-convergence. As a preparation, we prove that the infimum in (4.5) is actually attained, where we exploit some results from the previous proof.

Proposition 4.4.3. *For $f \in L^2([0, 1], \mathcal{F})$ with $\mathcal{E}[f] < \infty$, the infimum in (4.5) is attained, *i.e.*, there exists a tuple $(v, z) \in \mathcal{C}(f)$ satisfying (4.6) and (4.7) such that*

$$\mathcal{E}[f] = \int_0^1 \int_{\Omega} L[a[\mathcal{P}[f]], v, v] + \frac{1}{\delta} z^2 dx dt.$$

Proof. By assumption (A1) and Remark 4.1.4, the functional $(v, z) \mapsto \int_0^1 \int_{\Omega} L[a[\mathcal{P}[f]], v, v] + \frac{1}{\delta} z^2 dx dt$ is well defined and, by Korn's inequality 2.1.6 and Gagliardo–Nirenberg interpolation inequality 2.1.5, it is coercive and weakly lower semicontinuous on $\mathcal{C}(f)$ (see also [BK17]). Since $\mathcal{C}(f)$ is a subset of a reflexive Banach space, it suffices to prove the weak closedness of $\mathcal{C}(f)$, to obtain the existence of an optimal tuple $(v, z) \in \mathcal{C}(f)$.

Let $\{(v^j, z^j)\}_{j \in \mathbb{N}} \in \mathcal{C}(f)$ be a weakly convergent sequence with limit (v, z) . Due to Theorem 3.2.1 the corresponding flows $\{\psi^j\}_{j \in \mathbb{N}}$ and ψ exist and $\psi^j \rightarrow \psi$ in $C^0([0, 1] \times \bar{\Omega})$. Furthermore, as $\{v^j\}_{j \in \mathbb{N}}$ is uniformly bounded in $L^2((0, 1), C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$, we can apply the reasoning in the paragraph following (4.43) to conclude that a subsequence of $\{\psi^j\}_{j \in \mathbb{N}}$ converges strongly to ψ in $C^{0,\beta}([0, 1], C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$, for $\beta \in (0, \min(\frac{1}{2}, \alpha))$. Finally, by Remark 4.2.4, we have $f_t \circ \psi_t^j \rightarrow f_t \circ \psi_t$ in $L^2([0, 1], \mathcal{F})$ and, by Remark 4.2.8, it holds $z_t^j \circ \psi_t^j \rightarrow z_t \circ \psi_t$ in $L^2((0, 1) \times \Omega)$. Thus, we reason analogously to (4.46) to conclude that (4.7) is satisfied. \square

Proof of the existence of a recovery sequence and lim sup-inequality

We again provide an outline to facilitate reading.

- (i) *Construction of a recovery sequence.* The recovery sequence is constructed by applying the feature extension operator to the discrete vector of features obtained by sampling the given feature curve at time points $\frac{k}{K}$, $k = 0, \dots, K$, and the deformation vector obtained as the relative flow between the sampling points.
- (ii) *Proof of the actual inequality.* To bound the feature mismatch part of the energy, we use the definition of the vectors of features and deformations, and the variational inequality (4.7). For the dissipation part of the energy, we used the Taylor expansion of W , the Jensen inequality for the operator L , and, finally, the property of the trace of matrices, and an algebraic property of the set of Sobolev diffeomorphisms. The anisotropy terms are handled in a similar manner as in the above proof, upon having the result of the following step.
- (iii) *Proof of the convergence of the recovery sequence.* In this step, we prove the strong convergence in $L^2((0, 1), \mathcal{F})$ of the recovery sequence to the given feature curve. To this end, we use the variational inequality (4.7) and the strong convergence of the discrete global flows towards the optimal continuous flow.

ad (i): *Construction of a recovery sequence.*

Let $f \in L^2([0, 1], \mathcal{F})$ with $\mathcal{E}[f] < \infty$. Then, let $(v, z) \in \mathcal{C}(f)$ be a pair from the previous proposition with an associated flow ψ . We define the discrete feature and the discrete deformation vectors by

$$\begin{aligned} f_k^K(x) &:= f(t_k^K, x), \quad k = 0, \dots, K, \\ \phi_k^K(x) &:= \psi_{t_{k-1}^K, t_k^K}(x), \quad k = 1, \dots, K, \quad \text{where } \psi_{s,t}(\cdot) = \psi(t, \psi^{-1}(s, \cdot)), \end{aligned} \quad (4.57)$$

and the point evaluation in time of the feature curve is possible by Remark 4.1.4. To be in a position to apply the temporal extension construction from Section 4.3, we bound the $C^1(\bar{\Omega})$ -norm of the displacements as follows:

$$\begin{aligned} \max_{k \in \{1, \dots, K\}} \|\phi_k^K - \mathbf{1}\|_{C^1(\bar{\Omega})} &\leq \sup_{\substack{s, t \in [0, 1] \\ |t-s| \leq K^{-1}}} \|\psi_{s,t} - \mathbf{1}\|_{C^1(\bar{\Omega})} \\ &\leq \sup_{\substack{s, t \in [0, 1] \\ |t-s| \leq K^{-1}}} \|\psi\|_{C^0([0, 1], C^1(\bar{\Omega}))} |t-s|^{\frac{1}{2}} |\psi^{-1}|_{C^{0, \frac{1}{2}}([0, 1], C^{1, \alpha}(\bar{\Omega}))} \\ &\leq K^{-\frac{1}{2}} M (\|v\|_{L^2((0, 1), \nu)}), \end{aligned} \quad (4.58)$$

where M is some monotonically increasing function (cf. Remarks 3.2.2 and 3.2.3). Thus, choosing K sufficiently large ensures (cf. (4.37) and (4.38)) that $\Phi^K \in \mathcal{D}^K$ and we can indeed repeat the procedure from Section 4.3. Finally, the recovery sequence is defined by $f^K := \mathcal{F}^K [f^K, \Phi^K]$.

ad (ii): *Proof of the actual inequality.*

We start bounding the parts of $\mathcal{E}^K[f^K]$, given by (4.36), separately. For any $k = 1, \dots, K$, we infer, using (4.7), Jensen's inequality, and (4.58), that

$$\begin{aligned} \int_{\Omega} |f_{k-1}^K - f_k^K \circ \phi_k^K|^2 dx &= \int_{\Omega} |f_{k-1}^K \circ \psi_{t_{k-1}^K, t_k^K} - f_k^K \circ \psi_{t_{k-1}^K, t_k^K}|^2 \det(D\psi_{t_{k-1}^K, t_k^K}) dx \\ &\leq \int_{\Omega} \left(\int_{t_{k-1}^K}^{t_k^K} z_s \circ \psi_s ds \right)^2 \det(D\psi_{t_{k-1}^K, t_k^K}) dx \\ &\leq \frac{1}{K} \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} z_s^2 \det(D\psi_{s, t_{k-1}^K}) dx ds \\ &\leq \frac{1}{K} \left(1 + CK^{-\frac{1}{2}} \right) \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} z_s^2 dx ds. \end{aligned} \quad (4.59)$$

To control the deformation regularization terms, we recall that $v_k^K = K(\phi_k^K - \mathbf{1})$ and apply the same Taylor expansion argument as in (4.48), to get

$$\int_{\Omega} a_k^K W(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 dx \leq K^{-2} \int_{\Omega} L[a_k^K, v_k^K, v_k^K] dx + CK^{-3} \int_{\Omega} a_k^K |Dv_k^K|^3 dx. \quad (4.60)$$

Here, as usual, $a_k^K = a[\mathcal{P}[f_k^K]]$, for $k = 1, \dots, K$. Summing over the second term on the right-hand side and taking into account (4.58) and (A1), we obtain

$$\sum_{k=1}^K \int_{\Omega} a_k^K |Dv_k^K|^3 dx \leq C_a CK^3 \sum_{k=1}^K \|\phi_k^K - \mathbf{1}\|_{C^1(\bar{\Omega})}^3 \leq CK^{\frac{3}{2}}. \quad (4.61)$$

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

$$\begin{aligned} \int_{\Omega} L[a_k^K, v_k^K, v_k^K] dx &= \int_{\Omega} L \left[a_k^K, K \int_{t_{k-1}^K}^{t_k^K} v_t \circ \psi_{t_{k-1}^K, t} dt, K \int_{t_{k-1}^K}^{t_k^K} v_t \circ \psi_{t_{k-1}^K, t} dt \right] dx \\ &\leq \int_{\Omega} K \int_{t_{k-1}^K}^{t_k^K} L[a_k^K, v_t \circ \psi_{t_{k-1}^K, t}, v_t \circ \psi_{t_{k-1}^K, t}] dt dx. \end{aligned} \quad (4.62)$$

We will bound the terms appearing in L (cf. (3.10)) separately. By using (4.58), the transformation formula, and $|\text{tr}(AB)| \leq |\text{tr}(A)| + |\text{tr}(A(B - \mathbf{1}))|$, for $A, B \in \mathbb{R}^{n \times n}$, multiple times, we can estimate the part corresponding to the first summand of L as follows:

$$\begin{aligned} &\int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} a_k^K \text{tr} \left(D \left(v_t \circ \psi_{t_{k-1}^K, t} \right) \right)^2 dt dx \\ &= \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} a_k^K \text{tr} \left(Dv_t \circ \psi_{t_{k-1}^K, t} D\psi_{t_{k-1}^K, t} \right)^2 dt dx \\ &\leq \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} a_k^K \left[\text{tr} \left(Dv_t \circ \psi_{t_{k-1}^K, t} \right)^2 + \text{tr} \left(Dv_t \circ \psi_{t_{k-1}^K, t} \left(\mathbf{1} - D\psi_{t_{k-1}^K, t} \right) \right)^2 \right. \\ &\quad \left. + 2 \left| \text{tr} \left(Dv_t \circ \psi_{t_{k-1}^K, t} \right) \text{tr} \left(Dv_t \circ \psi_{t_{k-1}^K, t} \left(\mathbf{1} - D\psi_{t_{k-1}^K, t} \right) \right) \right| \right] dt dx \\ &\leq \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} a_k^K \text{tr} (\varepsilon[v])^2 + C \text{tr} (\varepsilon[v])^2 K^{-\frac{1}{2}} dt dx. \end{aligned} \quad (4.63)$$

The second term in L is estimated analogously:

$$\int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} a_k^K \text{tr} \left(\varepsilon \left[v_t \circ \psi_{t_{k-1}^K, t} \right] \right)^2 dt dx \leq \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} a_k^K \text{tr} (\varepsilon[v]^2) + C \text{tr} (\varepsilon[v]^2) K^{-\frac{1}{2}} dt dx. \quad (4.64)$$

It remains to bound the higher order term appearing in the definition of L . To this end, we use

$$\|v_t \circ \psi_t\|_{H^m(\Omega)} \leq C \|v_t\|_{H^m(\Omega)},$$

which follows from [BV17, Lemma 3.5], and the estimate $\|fg\|_{H^{\tilde{m}}} \leq C \|f\|_{H^m} \|g\|_{H^{\tilde{m}}}$ for $f \in H^m(\Omega)$, $g \in H^{\tilde{m}}(\Omega)$ and any $0 \leq \tilde{m} \leq m$ [IKT13, Lemma 2.3], resulting in the estimates

$$\begin{aligned} &|v_t \circ \psi_{t_{k-1}^K, t}|_{H^m(\Omega)} \\ &\leq |Dv_t \circ \psi_{t_{k-1}^K, t}|_{H^{m-1}(\Omega)} + \|Dv_t \circ \psi_{t_{k-1}^K, t} D(\psi_{t_{k-1}^K, t} - \mathbf{1})\|_{H^{m-1}(\Omega)} \\ &\leq |Dv_t \circ \psi_{t_{k-1}^K, t}|_{H^{m-1}(\Omega)} + C \|v_t\|_{H^m(\Omega)} \|\psi_{t_{k-1}^K, t} - \mathbf{1}\|_{H^m(\Omega)} \\ &\leq |Dv_t \circ \psi_{t_{k-1}^K, t}|_{H^{m-1}(\Omega)} + C \|v_t\|_{H^m(\Omega)} K^{-\frac{1}{2}}. \end{aligned}$$

By iterating this argument and applying a change of variables, we obtain

$$\int_{t_{k-1}^K}^{t_k^K} |v_t \circ \psi_{t_{k-1}^K, t}|_{H^m(\Omega)}^2 dt \leq \int_{t_{k-1}^K}^{t_k^K} |v_t|_{H^m(\Omega)}^2 + C \|v_t\|_{H^m(\Omega)}^2 K^{-\frac{1}{2}} dt. \quad (4.65)$$

By combining the estimate (4.59) with (4.60)–(4.65), we get

$$\begin{aligned} \mathcal{E}^K[f^K] &\leq \mathbf{E}^{K,D}[f^K, \Phi^K] = K \sum_{k=1}^K \int_{\Omega} a_k^K \mathbb{W}(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 + \frac{1}{\delta} |f_{k-1}^K - f_k^K \circ \phi_k^K|^2 dx \\ &\leq \sum_{k=1}^K \left(\int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} L[a_k^K, v, v] + CK^{-\frac{1}{2}} + CK^{-1} |Dv_k^K|^3 + \frac{1}{\delta} \left(1 + CK^{-\frac{1}{2}}\right) z_t^2 dx dt \right) \\ &\leq \int_0^1 \int_{\Omega} L[\bar{a}^K, v, v] + \frac{1}{\delta} z_t^2 dx dt + CK^{-\frac{1}{2}}, \end{aligned} \quad (4.66)$$

where \bar{a}^K is defined by (4.49). Thus, to prove the desired estimate, it is again enough to prove $\bar{a}^K \rightarrow a$ in $L^\infty([0, 1] \times \Omega)$. Having the estimates (4.58) and (4.59), we can repeat the same arguments as in the proof of lim inf-inequality (cf. (4.55), (4.56)) to prove that $\bar{a}^K - a^K \rightarrow 0$ in $L^\infty([0, 1] \times \Omega)$, while $a^K \rightarrow a$ in $C^0([0, 1], L^\infty(\Omega))$ will follow from the identification of the recovery sequence limit and continuity of point evaluation (cf. (4.52)), which is the last step of the proof.

ad (iii): *Proof of the convergence of the recovery sequence.*

As announced, we show $f^K \rightarrow f$ in $L^2([0, 1], \mathcal{F})$. To this end, with the discrete flow ψ^K given by (4.34), we estimate

$$\begin{aligned} &\int_0^1 \int_{\Omega} |f_s \circ \psi_s - f_s^K \circ \psi_s^K|^2 dx ds \\ &= \sum_{k=1}^K \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} |f_s \circ \psi_s - f_s^K \circ \psi_s^K|^2 ds dx \\ &\leq 2 \sum_{k=1}^K \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} |f_s \circ \psi_s - f_{t_{k-1}^K} \circ \psi_{t_{k-1}^K}|^2 + |f_{t_{k-1}^K} \circ \psi_{t_{k-1}^K} - f_s^K \circ \psi_s^K|^2 ds dx \\ &\leq C \sum_{k=1}^K \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} K^{-2} (z_s \circ \psi_s)^2 + |f_{t_{k-1}^K} \circ \psi_{t_{k-1}^K} - f_{t_{k-1}^K} \circ \psi_{t_{k-1}^K}^K|^2 + |f_s^K \circ \phi_s^K - f_{t_{k-1}^K}^K|^2 (\det D\psi_{t_{k-1}^K}^K)^{-1} ds dx \\ &\leq CK^{-2} \left(\|z_t \circ \psi_t\|_{L^2((0,1) \times \Omega)}^2 + \|z\|_{L^2((0,1) \times \Omega)}^2 \right). \end{aligned}$$

Here, to go from the third to the fourth line, we combined (4.7) with the Cauchy–Schwarz inequality for the first term on right-hand side, while for the second term we used the definition of feature extension (4.29), (4.59), the uniform boundedness of $\{(\psi^K)^{-1}\}_{K \in \mathbb{N}}$ following from (4.33), (4.58), Theorem 3.2.1, and finally

$$\psi_{t_k^K}^K = \phi_k^K \circ \psi_{t_{k-1}^K}^K = \phi_k^K \circ \dots \circ \phi_1^K = \psi_{t_k^K}, \text{ for all } k = 1, \dots, K, \quad (4.67)$$

which follows from (4.34) and (4.57). From here, by Remark 4.2.8, it follows that it is enough to prove that $\psi^K \rightarrow \psi$ in $C^{0,\beta}([0, 1], C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$, for some $\beta \geq 0$. This follows, for $\beta \in (0, \min(\frac{1}{2}, \alpha))$, from the uniform boundedness of $\{\psi^K\}_{K \in \mathbb{N}}$ in that space (cf. Theorem 3.2.1 and Remark 3.2.3), Corollary 2.1.10, and (4.67).

Finally, we can now repeat the arguments (4.53) and (4.54) to prove $a^K \rightarrow a$ in $C^0([0, 1], L^\infty(\Omega))$ which, plugged in (4.66), gives $\limsup_{K \rightarrow \infty} \mathcal{E}^K[f^K] \leq \mathcal{E}[f]$, as we wanted to show.

This finally finishes the proof of Mosco-convergence. \square

This theorem implies the existence of a geodesic path in the time continuous model as a limit of extensions of time discrete geodesic paths. This is studied in the following theorem.

Theorem 4.4.4 (Convergence of discrete geodesic paths). *Suppose that the assumptions (W1)–(W3) and (A1)–(A3) hold true. Let $f_A, f_B \in \mathcal{F}$ be fixed. For $K \in \mathbb{N}$ sufficiently large, let f^K be a minimizer of \mathcal{E}^K subject to $f_0^K = f_A$ and $f_1^K = f_B$. Then, a subsequence of $\{f^K\}_{K \in \mathbb{N}}$ weakly converges in $L^2([0, 1], \mathcal{F})$ to a minimizer of the continuous path energy \mathcal{E} as $K \rightarrow \infty$. Finally, the associated sequence of discrete path energies converges to the minimal continuous path energy.*

Proof. For every $K \geq 2$, let us observe the vector of features given by $\bar{f}_k^K := (1 - \frac{k}{K})f_A + \frac{k}{K}f_B$, for $k = 0, \dots, K$. Then, as in (4.25), we obtain that the path energies are, for every $K \in \mathbb{N}$, bounded by

$$\bar{\mathcal{E}} \leq \mathcal{E}^K[\mathcal{F}^K[\bar{\mathbf{f}}^K, \mathbf{1}^K]] \leq \frac{1}{\delta} \|f_A - f_B\|_{\mathcal{F}}^2.$$

For the optimal vectors of images \mathbf{f}^K and deformations Φ^K in the definition of \mathbf{E}^K , which exist, for K large enough, by Theorem 4.2.7, we can apply the temporal extension construction from Section 4.3 (cf. (4.37) and (4.38)). In particular, $\{f^K = \mathcal{F}^K[\mathbf{f}^K, \Phi^K]\}_{K \in \mathbb{N}}$ are minimizers of $\{\mathcal{E}^K\}_{K \in \mathbb{N}}$, and $\mathcal{E}^K[f^K] = \mathbf{E}^{K,D}[\mathbf{f}^K, \Phi^K] \leq \bar{\mathcal{E}}$ for all (large enough) $K \in \mathbb{N}$. Then, by Remark 3.2.2, together with (4.33), and (4.37), we have the uniform boundedness of $\{\psi^K\}_{K \in \mathbb{N}}$, $\{(\psi^{-1})^K\}_{K \in \mathbb{N}}$ in $C^0([0, 1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$. Furthermore, as in (4.40), we conclude that $\{z^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $L^2((0, 1) \times \Omega)$. Plugging this into Proposition 4.3.1, we obtain, for every $t \in [0, 1]$,

$$\|f_t^K\|_{\mathcal{F}} \leq C (\|f_t^K \circ \psi_t^K - f_A\|_{\mathcal{F}} + \|f_A\|_{\mathcal{F}}) \leq C (\|z^K\|_{L^2((0,1) \times \Omega)} + \|f_A\|_{\mathcal{F}}).$$

Therefore, $\{f^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $L^\infty([0, 1], \mathcal{F})$ and a subsequence converges weakly in $L^2([0, 1], \mathcal{F})$ to some $f \in L^2([0, 1], \mathcal{F})$. Assume that there exists an image path $\tilde{f} \in L^2([0, 1], \mathcal{F})$ with corresponding optimal tuple $(\tilde{v}, \tilde{\psi}, \tilde{z})$, which exists due to Proposition 4.4.3, satisfying (4.6) and (4.7), such that

$$\mathcal{E}[\tilde{f}] < \mathcal{E}[f]. \quad (4.68)$$

Using the lim sup-part of the previous theorem, there exists a sequence $\{\tilde{f}^K\}_{K \in \mathbb{N}} \subset L^2([0, 1], \mathcal{F})$ satisfying $\limsup_{K \rightarrow \infty} \mathcal{E}^K[\tilde{f}^K] \leq \mathcal{E}[\tilde{f}]$. However, applying the lim inf-inequality from the previous theorem, we obtain

$$\mathcal{E}[f] \leq \liminf_{K \rightarrow \infty} \mathcal{E}^K[f^K] \leq \limsup_{K \rightarrow \infty} \mathcal{E}^K[\tilde{f}^K] \leq \mathcal{E}[\tilde{f}], \quad (4.69)$$

which contradicts (4.68). Hence, f 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., $\lim_{K \rightarrow \infty} \mathcal{E}^K[f^K] = \mathcal{E}[f]$, which again follows from (4.69) by using $\tilde{f} = f$. \square

Let us conclude this section with a couple of properties of a geodesic curve $f \in L^2([0, 1], \mathcal{F})$.

Remark 4.4.5. (i) In the context of Proposition 4.1.1, for a geodesic curve $f \in L^2([0, 1], \mathcal{F})$ it holds $z = |\hat{z}|$, where \hat{z} is the vector material derivative.

(ii) Let $F : \mathcal{I} \rightarrow L^2(\Omega, \mathbb{R}^c)$ be compact and Lipschitz continuous, i.e., the assumptions of the Remarks 4.2.6 and 4.2.9 are satisfied. Then, the discrete geodesic paths can be assumed to be of the type $\mathbf{f}^K = (u_k^K, F(\frac{u_k^K}{\eta}))$ for $k = 1, \dots, K$ and every (large enough) $K \in \mathbb{N}$ (cf. (4.1)). This allows us to show that, with f being a geodesic curve, it holds

$$f_t = (u_t, F(\frac{u_t}{\eta})), \text{ for every } t \in [0, 1].$$

Indeed, f is obtained as a weak limit in \mathcal{F} of the extensions $f^K = \mathcal{F}^K[\mathbf{f}^K, \Phi^K]$, where $\{\mathbf{f}^K\}_{K \in \mathbb{N}}$ are the discrete geodesic paths. From the definition of feature extension, for every $t_k^K = \frac{k}{K}$, $k = 0, \dots, K$, it holds $f_{t_k^K}^K = f_k^K$ (cf. (4.30)). Thus, for every $t \in [0, 1] \cap \mathbb{Q}$, there exists infinitely many $K \in \mathbb{N}$ such that $f_t^K = (u_{k(K)}^K, F(\frac{u_{k(K)}^K}{\eta}))$. As we showed in (4.52), as $K \rightarrow \infty$, for every $t \in [0, 1]$ it holds $f_t^K \rightharpoonup f_t$ in \mathcal{F} . Thus, by the compactness of the mapping F , we get $f_t = (u_t, F(\frac{u_t}{\eta}))$ for every $t \in [0, 1] \cap \mathbb{Q}$. To show the same for all the points $t \in [0, 1]$, we take $t' \in [0, 1] \cap \mathbb{Q}$ and estimate

$$\|f_t - (u_t, F(\frac{u_t}{\eta}))\|_{\mathcal{F}} \leq \|f_t - f_{t'}\|_{\mathcal{F}} + \|(u_{t'}, F(\frac{u_{t'}}{\eta})) - (u_t, F(\frac{u_t}{\eta}))\|_{\mathcal{F}} \leq \left(2 + \frac{L_F}{\eta}\right) \omega(|t - t'|),$$

where ω is the modulus of continuity of f (cf. Remark 4.1.4) and L_F the Lipschitz constant of F . As t' can be chosen arbitrarily close to t , our claim is proven.

4.5 Fully Discrete Model

In this section, we present the fully discrete model on the image domain $\Omega = [0, 1]^2$. We use boldface letters to differentiate discrete points, feature maps, images, and deformations from their continuous counterparts. For $M, N \geq 3$, we define the computational domain (grid) and its boundary as follows

$$\Omega_{MN} := \Omega_M \times \Omega_N := \left\{ \frac{0}{M-1}, \dots, \frac{M-1}{M-1} \right\} \times \left\{ \frac{0}{N-1}, \dots, \frac{N-1}{N-1} \right\}, \quad \partial\Omega_{MN} := \Omega_{MN} \cap \partial([0, 1]^2).$$

The discrete image space and the discrete feature space are given by $\mathcal{I}_{MN} := \{\mathbf{u} : \Omega_{MN} \rightarrow \mathbb{R}^3\}$ and $\mathcal{F}_{MN} := \{\mathbf{f} : \Omega_{MN} \rightarrow \mathbb{R}^{3+c}\}$, respectively. The set of admissible deformations is given by

$$\mathcal{D}_{MN} := \left\{ \phi : \Omega_{MN} \rightarrow [0, 1]^2 : \phi = \mathbb{1} \text{ on } \partial\Omega_{MN}, \det(\nabla_{MN} \phi) > 0 \right\},$$

where the discrete Jacobian operator ∇_{MN} of ϕ is defined as the forward finite difference operator with Neumann boundary conditions. We define the discrete L^p -norm of a discrete object $\mathbf{g} \in \mathcal{F}_{MN}$ by

$$\|\mathbf{g}\|_{L^p(\Omega_{MN})}^p := \frac{1}{MN} \sum_{(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}} \sum_{j=1}^{3+c} |\mathbf{g}^j(\mathbf{x}, \mathbf{y})|^p.$$

In the fully discrete model, we have to define the spatial warping operator \mathbf{T} as a numerical approximation of the pullback of a feature channel $\mathbf{f}^j \circ \phi$ at a point $(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}$. The starting point of this approximation is the fact that if $\phi(\mathbf{x}, \mathbf{y}) = (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \Omega_{MN}$ then we want to have $\mathbf{T}[\mathbf{f}^j, \phi](\mathbf{x}, \mathbf{y}) = \mathbf{f}^j(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$. Between the grid points, we need to use a sufficiently smooth interpolation. One could simply make a linear interpolation of values at grid points, but this approach does not offer enough smoothness. We, therefore, turn to a higher order approach that uses piecewise cubic polynomial interpolation. To this end, in addition to values at the grid points, one needs to specify additional constraints. Here, we specify values of derivatives at grid points and use (*cellwise*) *cubic Hermite splines* interpolation (see [dB78, Chapter IV] for details). Our approach can be seen as a composition of one-dimensional interpolations on each interval. Thus, we first present details for the one-dimensional interpolations and then explore this for our two-dimensional case. To this end, for $\mathbf{g} : \Omega_M \rightarrow \mathbb{R}$ we define, for every $x \in [0, 1]$,

$$\mathbf{I}_M[\mathbf{g}](x) := s_{\mathbf{g}}(x), \quad (4.70)$$

where $s_{\mathbf{g}} \in C^1([0, 1], \mathbb{R})$ is the (intervalwise) cubic Hermite spline, such that, for $\mathbf{x}_i := \frac{i}{M-1}$, $i = 0, \dots, M-1$,

$$s_{\mathbf{g}}(\mathbf{x}_i) = \mathbf{p}_i := \mathbf{g}(\mathbf{x}_i), \quad s'_{\mathbf{g}}(\mathbf{x}_i) = \mathbf{m}_i := \frac{M-1}{2} (\mathbf{p}_{i+1} - \mathbf{p}_{i-1}),$$

with $\mathbf{p}_i := 0$ for $i < 0$ or $i > M-1$. The specific choice for the approximation of derivatives is called the *Catmull-Rom spline* [CR74]. On each subinterval $[\mathbf{x}_i, \mathbf{x}_{i+1}]$, for $i = 0, \dots, M-2$, cubic Hermite splines can be written in terms of *cubic Hermite basis* $\{h_{00}, h_{01}, h_{10}, h_{11}\}$, which allows for an efficient numerical implementation. We have

$$s_{\mathbf{g}}(x) = h_{00}(t)\mathbf{p}_{i_x} + \frac{h_{10}(t)}{M-1}\mathbf{m}_{i_x} + h_{01}(t)\mathbf{p}_{i_x+1} + \frac{h_{11}(t)}{M-1}\mathbf{m}_{i_x+1}, \quad (4.71)$$

where $i_x := \lfloor (M-1)x \rfloor$, $t := (M-1)(x - \mathbf{x}_{i_x})$, and

$$h_{00}(t) := 2t^3 - 3t^2 + 1, \quad h_{10}(t) := t^3 - 2t^2 + t, \quad h_{01}(t) := -2t^3 + 3t^2, \quad h_{11}(t) := t^3 - t^2.$$

In particular, we can also write this as

$$\begin{aligned} s_{\mathbf{g}}(x) &= \mathbf{H}_M[(\mathbf{g}(\mathbf{x}_{i_x+k}))_{k=-1, \dots, 2}, x] \\ &:= h_{00}(t)\mathbf{g}(\mathbf{x}_{i_x}) + \frac{h_{10}(t)}{2}(\mathbf{g}(\mathbf{x}_{i_x+1}) + \mathbf{g}(\mathbf{x}_{i_x-1})) + h_{01}(t)\mathbf{g}(\mathbf{x}_{i_x+1}) + \frac{h_{11}(t)}{2}(\mathbf{g}(\mathbf{x}_{i_x+2}) + \mathbf{g}(\mathbf{x}_{i_x})), \end{aligned}$$

where $\mathbf{g}(\mathbf{x}_i) := 0$ for $i < 0$ or $i > M-1$. For $\mathbf{g} : \Omega_N \rightarrow \mathbb{R}$ we can analogously define, by taking care of the size of the intervals, for every $y \in [0, 1]$,

$$\mathbf{I}_N[\mathbf{g}](y) := \mathbf{H}_N[(\mathbf{g}(\mathbf{x}_{i_y+l}))_{l=-1, \dots, 2}, y].$$

As already announced, the two dimensional interpolation is the composition of one dimensional interpolations. In details, for $\mathbf{g} : \Omega_{MN} \rightarrow \mathbb{R}$ and every $(x, y) \in [0, 1]^2$, we define

$$\mathbf{I}_{MN}[\mathbf{g}](x, y) := \mathbf{H}_N[(\mathbf{H}_M[(\mathbf{g}(\mathbf{x}_{i_x+k}, \mathbf{y}_{i_y+l}))_{k=-1, \dots, 2}, x])_{l=-1, \dots, 2}, y].$$

Finally, to define the warping operator \mathbf{T} we need to evaluate this interpolation at the deformed positions. To this end, for every $(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}$ and every $j = 1, \dots, 3+c$, we define

$$\mathbf{T}[\mathbf{f}^j, \phi](\mathbf{x}, \mathbf{y}) := \mathbf{I}_{MN}[\mathbf{f}^j](\phi(\mathbf{x}, \mathbf{y})).$$

Interpolation with cubic Hermite splines of prescribed function values on a discrete computational domain (derivatives can also be estimated from the function values as in our case) results in a C^1 -regular interpolant. In our application, this regularity reflects in the regularization of the deformations via discretization and thus allows us to neglect the H^m -seminorm term in the fully discretized energy.

The fully discrete mismatch functional \mathbf{D}_{MN} that approximates $\int_{\Omega} |\tilde{f} \circ \phi - f|^2 dx$ reads as

$$\mathbf{D}_{MN}[\mathbf{f}, \tilde{\mathbf{f}}, \phi] := \frac{1}{2(3+c)} \sum_{j=1}^{3+c} \left\| \mathbf{T}[\tilde{\mathbf{f}}^j, \phi] - \mathbf{f}^j \right\|_{L^2(\Omega_{MN})}^2.$$

Likewise, the discretization of the lower order anisotropic regularization functional $\int_{\Omega} aW(D\phi) dx$ is given by

$$\mathbf{R}_{MN}[\phi, \mathbf{a}] := \|\mathbf{a}W(\nabla_{MN}\phi)\|_{L^1(\Omega_{MN})}.$$

The energy density is given by (4.11) *i.e.*,

$$W(A) := \frac{\lambda}{2} \left(e^{(\log \det(A))^2} - 1 \right) + \mu |A^{\text{sym}} - \mathbb{1}|^2, \quad \text{for } A \in \text{GL}^+(2),$$

and the fully discrete anisotropic operator $\mathbf{a} : \mathcal{I}_{MN} \rightarrow L^\infty(\Omega_{MN})$ is the straightforward discretization of (4.2), *i.e.*, for every $\mathbf{u} \in \mathcal{I}_{MN}$ and $(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}$,

$$\mathbf{a}[\mathbf{u}](\mathbf{x}, \mathbf{y}) := \exp \left(- \frac{\sum_{j=1}^3 |(\mathcal{G}_{\sigma_2} \star \nabla_{MN}(\mathcal{G}_{\sigma_2} \star \mathbf{u}^j))(\mathbf{x}, \mathbf{y})|^2}{\rho} \right) + \xi,$$

where $\mathcal{G}_{\sigma_1, \sigma_2}$ are the discrete (truncated) two-dimensional Gaussian kernels with standard deviations σ_1 and σ_2 , respectively, and \star is the discrete cross correlation operation with stride equal to 1 and padding which ensures preservation of the dimensions (*cf.* (2.11)).

In summary, the fully discrete path energy in the deep feature metamorphosis model, for a $(K+1)$ -tuple $(\mathbf{f}_k)_{k=0}^K$ of discrete feature maps, a K -tuple $(\phi_k)_{k=1}^K$ of discrete deformations, and a K -tuple $(\mathbf{a}_k)_{k=1}^K$ of discrete anisotropies, reads as

$$\mathbf{E}_{MN}^{K,D}[(\mathbf{f}_k)_{k=0}^K, (\phi_k)_{k=1}^K, (\mathbf{a}_k)_{k=1}^K] := K \sum_{k=1}^K \mathbf{R}_{MN}[\phi_k, \mathbf{a}_k] + \frac{1}{\delta} \mathbf{D}_{MN}[\mathbf{f}_{k-1}, \mathbf{f}_k, \phi_k].$$

Finally, a discrete geodesic path $(\mathbf{f}_k)_{k=0}^K$ in feature space is a minimizer of \mathbf{E}_{MN}^K , subject to given discrete boundary data $\mathbf{f}_0 = \mathbf{f}_A$ and $\mathbf{f}_K = \mathbf{f}_B$, where

$$\mathbf{E}_{MN}^K[(\mathbf{f}_k)_{k=0}^K] := \inf_{(\phi_k)_{k=1}^K \in \mathcal{D}_{MN}^K} \mathbf{E}_{MN}^{K,D}[(\mathbf{f}_k)_{k=0}^K, (\phi_k)_{k=1}^K, (\mathbf{a}_k)_{k=1}^K].$$

As announced in the introduction of this section, we consider two substantially different options for \mathbf{f}_A and \mathbf{f}_B .

Simple RGB model. As a first model, we consider the simple image intensity-based feature space, *i.e.*, the feature space \mathcal{F}_{MN} coincides with the space of RGB images \mathcal{I}_{MN} and $\mathbf{f}_A = \mathbf{u}_A$ and $\mathbf{f}_B = \mathbf{u}_B$. Since a direct computation of the deformations on the full grid is numerically unstable, we incorporate a multilevel scheme. Initially, we start on the coarsest computational domain of size $M_L \times N_L$ with $M_L = 2^{-(L-1)}M$ and $N_L = 2^{-(L-1)}N$ for a given $L > 0$ and compute a time discrete geodesic path for bilinearly downsampled input images \mathbf{u}_A and \mathbf{u}_B . Then, in subsequent prolongation steps, the width and the height of the computational domain are successively doubled and the initial deformations and images are obtained via a bilinear interpolation of the preceding coarse scale solutions.

type	parameters	output size
$2 \times (\text{Conv2d} + \text{ReLU})$	3/1/1/64	$512 \times 512 \times 64$
MaxPool2d	2/2	$256 \times 256 \times 64$
$2 \times (\text{Conv2d} + \text{ReLU})$	3/1/1/128	$256 \times 256 \times 128$
MaxPool2d	2/2	$128 \times 128 \times 128$
$4 \times (\text{Conv2d} + \text{ReLU})$	3/1/1/256	$128 \times 128 \times 256$
MaxPool2d	2/2	$64 \times 64 \times 256$
$4 \times (\text{Conv2d} + \text{ReLU})$	3/1/1/512	$64 \times 64 \times 512$
MaxPool2d	2/2	$32 \times 32 \times 512$
$4 \times (\text{Conv2d} + \text{ReLU})$	3/1/1/512	$32 \times 32 \times 512$

Table 4.1: The architecture of VGG neural network (*cf.* Section 2.3 for the basic definitions and notation). The network consists of 16 convolutional layers followed by ReLU activation function. Max-pooling operation canonically yields a multilevel semantic decomposition of features into 5 levels. Parameters refer to convolutional kernel size/padding/stride/number of output channels for weight layers and padding/stride for pooling layers (all of the values are equal in all 4 directions). The names of the operations are as in the TensorFlow library [AAB⁺15]. The third column represents the output size of the features on different levels, for an initial input of size $M \times N \times c = 512 \times 512 \times 3$. This network is often referred to as VGG19, as there are 3 additional fully connected layers at the end, which are used for purposes of image classification and are not relevant to our application.

Deep feature space model. In the second model, $\mathbf{f}_A = (\eta \mathbf{u}_A, \mathbf{F}_{MN}(\mathbf{u}_A))$ and $\mathbf{f}_B = (\eta \mathbf{u}_B, \mathbf{F}_{MN}(\mathbf{u}_B))$, for (small) $\eta > 0$, where $\mathbf{F}_{MN} : \mathcal{I}_{MN} \rightarrow \{\mathbf{f} : \Omega_{MN} \rightarrow \mathbb{R}^c\}$ denotes the fully discrete feature extraction operator. The fully discrete images are again downsampled to match the corresponding grid size via bilinear interpolation. The feature extraction operator is represented by the VGG network as presented in [SZ14] to incorporate semantic information in image morphing. The VGG network is particularly designed for the localization and classification of objects in natural images and thus the feature decomposition of images is well-suited for semantic matching. Each feature map can be seen as a continuous map into some higher-dimensional feature space consisting of vectors in \mathbb{R}^c , where c ranges from 64 to 512 depending on the considered level associated with a certain network layer. Some of the features, together with the original image channels, are visualized in Figure 4.1. The architecture of the network is given in Table 4.1. In contrast to the simple RGB model, only the deformations are prolonged to the next level via a bilinear interpolation in the multilevel approach, since successive features on different levels are not necessarily related. To stabilize the optimization, the features on each level are first optimized using the upscaled deformations.

4.6 Numerical Optimization

In what follows, we present the numerical optimization scheme to compute geodesics for the fully discrete deep feature metamorphosis model. We use a variant of the iPALM algorithm [PS16] which we introduced in Section 3.3. While the regularization functional \mathbf{R}_{MN} is sufficiently well behaved w.r.t. the deformation variable, the data mismatch term \mathbf{D}_{MN} is non-convex and several numerical experiments indicated that a direct minimization of this term with respect to deformations is challenging due to the sensitivity of the warping operator to small perturbations of the deformations. Thus, to enhance the stability of the algorithm the warping operator is linearized around the extrapolated deformation $\phi^{[\beta]} \in \mathcal{D}_{MN}$, which is defined in terms of the value from the previous iteration step in the algorithm. The linearization is based on the gradient (*cf.* [WPZ⁺09])

$$\Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]}) := \frac{1}{2}(\nabla_{MN} \mathbf{T}[\tilde{\mathbf{f}}^j, \phi^{[\beta]}] + \nabla_{MN} \mathbf{f}^j),$$

where the Jacobian operator applied to the features is approximated using a Sobel filter [Sob90]. Here, $\langle \cdot, \cdot \rangle$ represent the pointwise product of the involved matrices. This yields the modified mismatch energy

$$\tilde{\mathbf{D}}_{MN}[\mathbf{f}, \tilde{\mathbf{f}}, \phi, \phi^{[\beta]}] := \frac{1}{2(3+c)} \sum_{j=1}^{3+c} \left\| \mathbf{T}[\tilde{\mathbf{f}}^j, \phi^{[\beta]}] + \langle \Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]}), \phi - \phi^{[\beta]} \rangle - \mathbf{f}^j \right\|_{L^2(\Omega_{MN})}^2.$$

The Moreau's proximal operator (*cf.* Definition 3.3.1) with respect to the deformation ϕ for a fixed $\tau > 0$ reads as

$$\text{prox}_{\tau \frac{K}{\delta} \tilde{\mathbf{D}}_{MN}}[\phi] := \left(\mathbf{1} + \frac{K}{\tau \delta (3+c)} \sum_{j=1}^{3+c} \Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]}) \Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]})^\top \right)^{-1} \left(\phi - \frac{K}{\tau \delta (3+c)} \sum_{j=1}^{3+c} \left(\Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]}) \mathbf{T}[\tilde{\mathbf{f}}^j, \phi^{[\beta]}] - \Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]}) \Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]})^\top \phi^{[\beta]} - \Lambda_j(\mathbf{f}, \tilde{\mathbf{f}}, \phi^{[\beta]}) \mathbf{f}^j \right) \right),$$

where we leave the function values on $\partial\Omega_{MN}$ unchanged. We are now ready to apply the iPALM algorithm with backtracking of Lipschitz constants (*cf.* Algorithm 3) to minimization of \mathbf{E}_{MN}^K . We initialize the vector of features by the linear interpolation of the fixed features, the vector of deformations by the identity vector, and choose some initial values for the step sizes (Lipschitz constants w.r.t. the features and deformations). Then, the alternating optimization is summarized in Algorithm 4, where the backtracking of the Lipschitz constants is achieved as in Algorithm 3.

Algorithm 4: Algorithm for minimizing \mathbf{E}_{MN}^K on one level.

```

1 for  $i = 1$  to  $I$  do
2   for  $k = 1$  to  $K$  do
3     /* update anisotropy */
4      $\mathbf{a}_k^{[i+1]} = \mathbf{a}[\mathcal{P}[\mathbf{f}_k^{[i]}]]$ ;
5     /* update deformation */
6      $\phi_k^{[i+1]} = \text{prox}_{\frac{K}{\delta} \tilde{\mathbf{D}}_{MN}} \left[ \phi_k^{[i,\beta]} - \frac{K}{L_{\phi_k}^{[i]}} D_{\phi_k} \mathbf{R}_{MN}[\phi_k^{[i,\beta]}, \mathbf{a}_k^{[i+1]}] \right]$ ;
7     if  $k < K$  then
8       /* update features */
9        $\mathbf{f}_k^{[i+1]} = \phi_k^{[i,\beta]} - \frac{1}{L_{\mathbf{f}_k}^{[i]}} D_{\mathbf{f}_k} \mathbf{E}_{MN}^{K,D}[\mathbf{f}^{[i,k,\beta]}, \phi^{[i,k+1]}, \mathbf{a}^{[i,k+1]}]$ ;

```

Here, for $\beta > 0$ we used the notation

$$\mathbf{h}_k^{[i,\beta]} := \mathbf{h}_k^{[i]} + \beta(\mathbf{h}_k^{[i]} - \mathbf{h}_k^{[i-1]}), \quad \mathbf{h}^{[i,k]} := (\mathbf{h}_{0,\dots,k-1}^{[i+1]}, \mathbf{h}_{k,\dots,K}^{[i]}), \quad \mathbf{h}^{[i,k,\beta]} := (\mathbf{h}_{0,\dots,k-1}^{[i+1]}, \mathbf{h}_k^{[i,\beta]}, \mathbf{h}_{k+1,\dots,K}^{[i]}).$$

The differentials of the corresponding fully discrete energies with respect to fully discrete deformations and features, in line 4 and 6 of the algorithm, respectively, are computed by the automatic differentiation package of TensorFlow library [AAB⁺15].

4.7 Numerical Results

In this section, we show numerical results for both the RGB and the deep feature model. Implementation was done in the Python programming language using the TensorFlow library [AAB⁺15], together with some C++ and Cuda extensions [NVF20] for additional acceleration. Let us note that the majority of the implementation tasks (including the warping operator and the optimization algorithm) were done by E. Kobler and A. Effland, while the final experimental results were obtained by the author of this thesis. All parameters used in the computation are specified in Table 4.2.

parameter	K	L	I	β	η	σ_1	σ_2	ρ	ξ	δ	μ	λ
RGB	15	5	250	$\frac{1}{\sqrt{2}}$	10^{-6}	0.5	2	1000	10^{-6}	1	0.025	0.1
deep								10^{-9}			0.002	0.002

Table 4.2: The parameter values for all examples. The choice of extrapolation parameter β is motivated by observations in [WCP17, MOPS20]. Notice that the different values of ρ result in the same contrasting due to the additional scaling of images by η in the deep feature model.

Figure 4.3 depicts the geodesic sequences for two self-portraits by van Gogh¹ ($M \times N = 496 \times 496$) for $k \in \{0, 3, 6, 9, 12, 15\}$ obtained with the RGB model (first row) and the deep feature model (fifth row). The superiority of the deep model, compared to the simple RGB model, is exemplarily visualized by the zoom (magnification factor 4) of the ear region depicted in the second and the sixth row. The remaining rows contain the corresponding sequences of anisotropy weights (third/seventh row) and color-coded displacement fields (fourth/eighth row), where the hue refers to the direction of the displacements and the intensity is proportional to its norm as indicated by the leftmost color wheel. Figure 4.4 presents analogous results for two photos of animals² for $M \times N = 512 \times 512$ with a zoom on the mouth region. Note that the deep model is capable of accurately deforming the carnassial teeth.

Figure 4.5 shows results of the deep feature model for two paintings of US presidents³ and two portraits of Catherine the Great⁴. In both cases, the input images have a resolution of $M \times N = 512 \times 512$.

Finally, we examine the effects of parameter changes of ρ and δ . Figure 4.6 visualizes the anisotropy weight and the deformation field in the RGB model for a value of the contrast parameter ρ fostering a significantly stronger anisotropy implying much more pronounced jumps in the deformation field (compare with Figure 4.3). Figure 4.7 illustrates the dependency of the resulting morphing sequences on δ for the RGB model (first to the third row) and the deep model (fourth to the sixth row). We notice that the larger values of δ lead to more blending. Furthermore, the generated geodesic paths using deep features are more robust to changes of δ than the RGB model, which can, for instance, be seen in the cheek or the eye regions.

In all numerical experiments, the displacement fields apparently evolve over time and the involved anisotropy promotes large deformation gradients in the proximity of image interfaces. These are indicated by the sharp interfaces in the color coding of the deformations. Both models fail to match image regions with no obvious correspondence of the input images, which can be seen in the cloth regions of the self-portraits, the presidents, and the empress examples, as well as on parts of the body region and the background in the animal example, where blending artifacts occur. The deep feature model clearly outperforms the simple RGB model in regions where the semantic similarity is not reflected by the RGB color features such as the cheek and the ear in the van Gogh example as well as the teeth of the animals. Moreover, to compute a visually appealing time discrete geodesic sequence, a fourth color channel representing a manual segmentation of image regions and a color adaptation of the van Gogh self-portraits was required in [BER15]. This is obsolete in the proposed deep feature based model due to the incorporation of semantic information.

¹public domain, https://commons.wikimedia.org/wiki/File:Vincent_Willem_van_Gogh_102.jpg; https://commons.wikimedia.org/wiki/File:SelbstPortrait_VG2.jpg

²first photo detail by Domenico Salvagnin (CC BY 2.0), [https://commons.wikimedia.org/wiki/File:Yawn!!!_\(331702223\).jpg](https://commons.wikimedia.org/wiki/File:Yawn!!!_(331702223).jpg); second photo detail by Eric Kilby (CC BY-SA 2.0), [https://commons.wikimedia.org/wiki/File:Panthera_tigris-Franklin_Park_Zoo,_Massachusetts,_USA-8a_\(2\).jpg](https://commons.wikimedia.org/wiki/File:Panthera_tigris-Franklin_Park_Zoo,_Massachusetts,_USA-8a_(2).jpg)

³first painting by Gilbert Stuart (public domain), https://commons.wikimedia.org/wiki/File:Gilbert_Stuart_Williamstown_Portrait_of_George_Washington.jpg; second painting by Rembrandt Peale (public domain), https://commons.wikimedia.org/wiki/File:Thomas_Jefferson_by_Rembrandt_Peale,_1800.jpg

⁴public domain, both portraits by J. B. Lampi [https://commons.wikimedia.org/wiki/File:Catherine_II_by_J.B.Lampi_\(Deutsches_Historisches_Museum\).jpg](https://commons.wikimedia.org/wiki/File:Catherine_II_by_J.B.Lampi_(Deutsches_Historisches_Museum).jpg); [https://commons.wikimedia.org/wiki/File:Catherine_II_by_J.B.Lampi_\(1780s,_Kunsthistorisches_Museum\).jpg](https://commons.wikimedia.org/wiki/File:Catherine_II_by_J.B.Lampi_(1780s,_Kunsthistorisches_Museum).jpg)

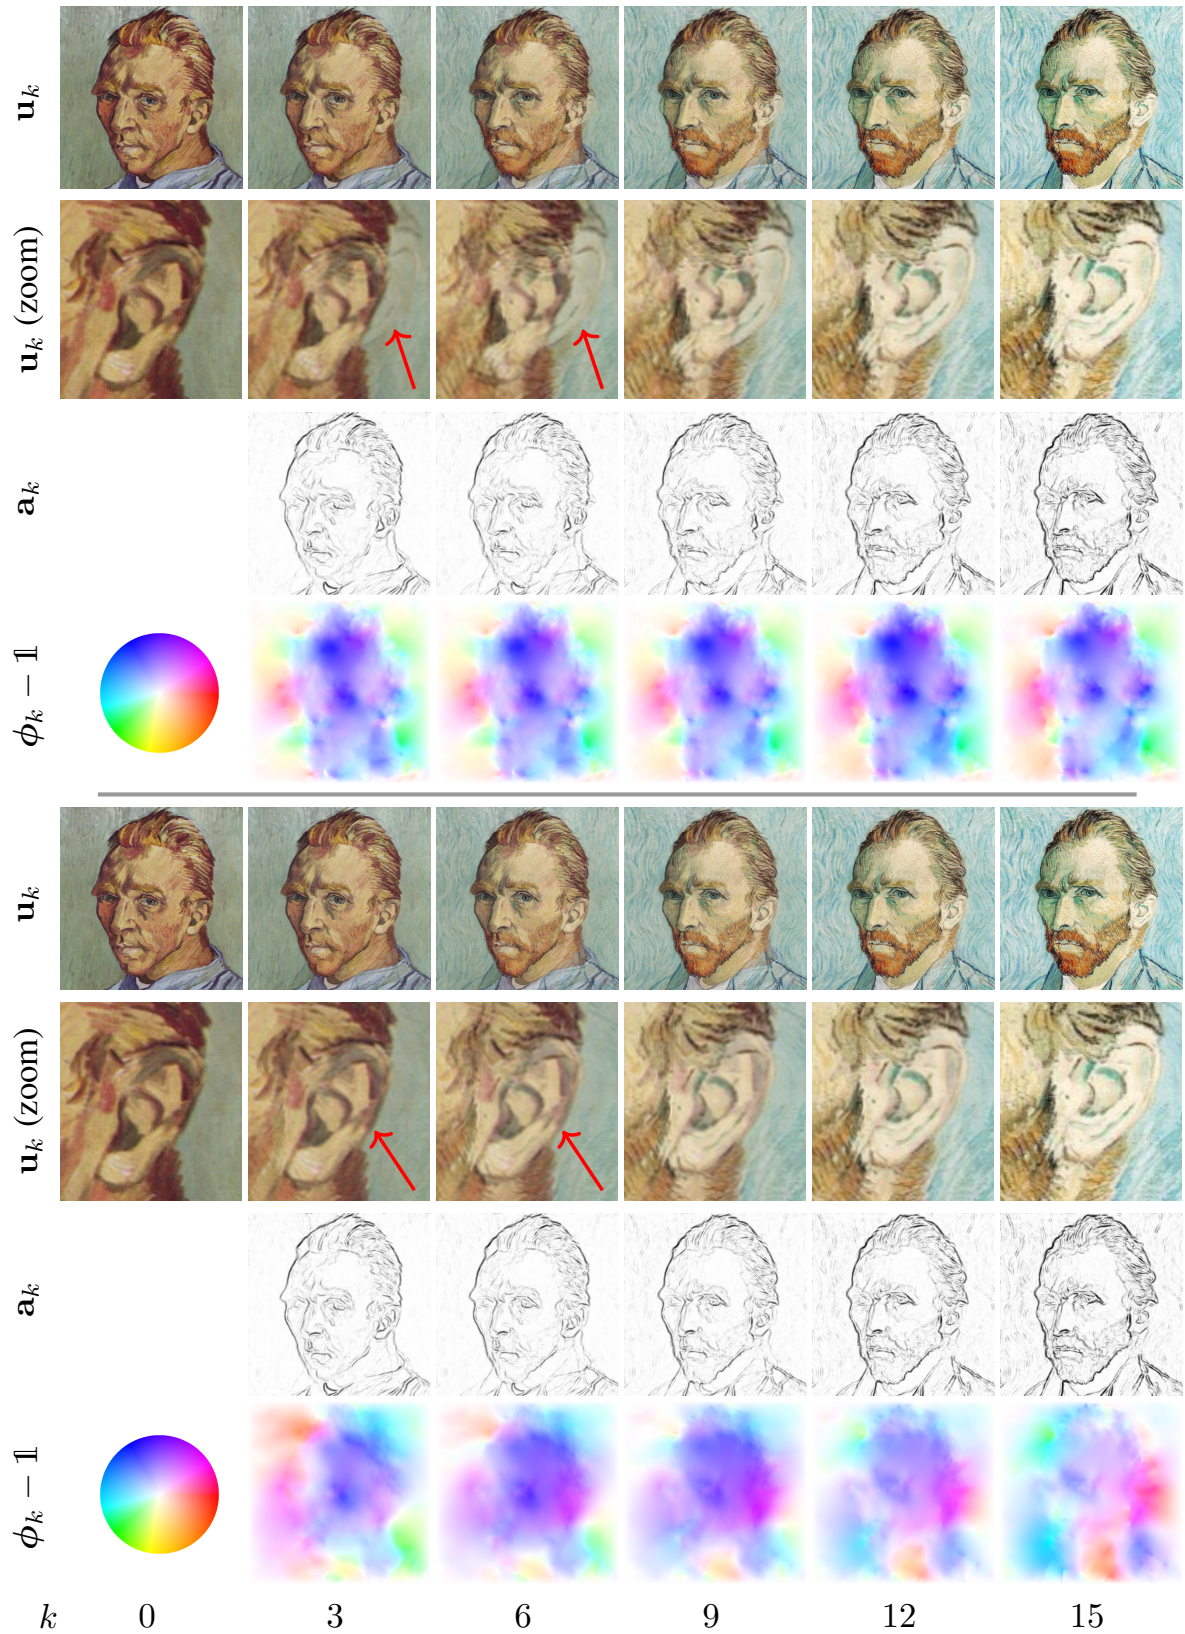


Figure 4.3: Time discrete geodesic sequences of self-portraits by van Gogh for the RGB model (first row) and deep feature model (fifth row), along with a zoom of the ear region with magnification factor 4 (second/sixth row), the associated sequences of anisotropy weights (third/seventh row), and color-coded displacement fields $\phi_k - \mathbb{1}$ (fourth/eighth row). Note that the intensity-based approach leads to blending artifacts indicated by the arrows, which are resolved in the deep feature-based model.

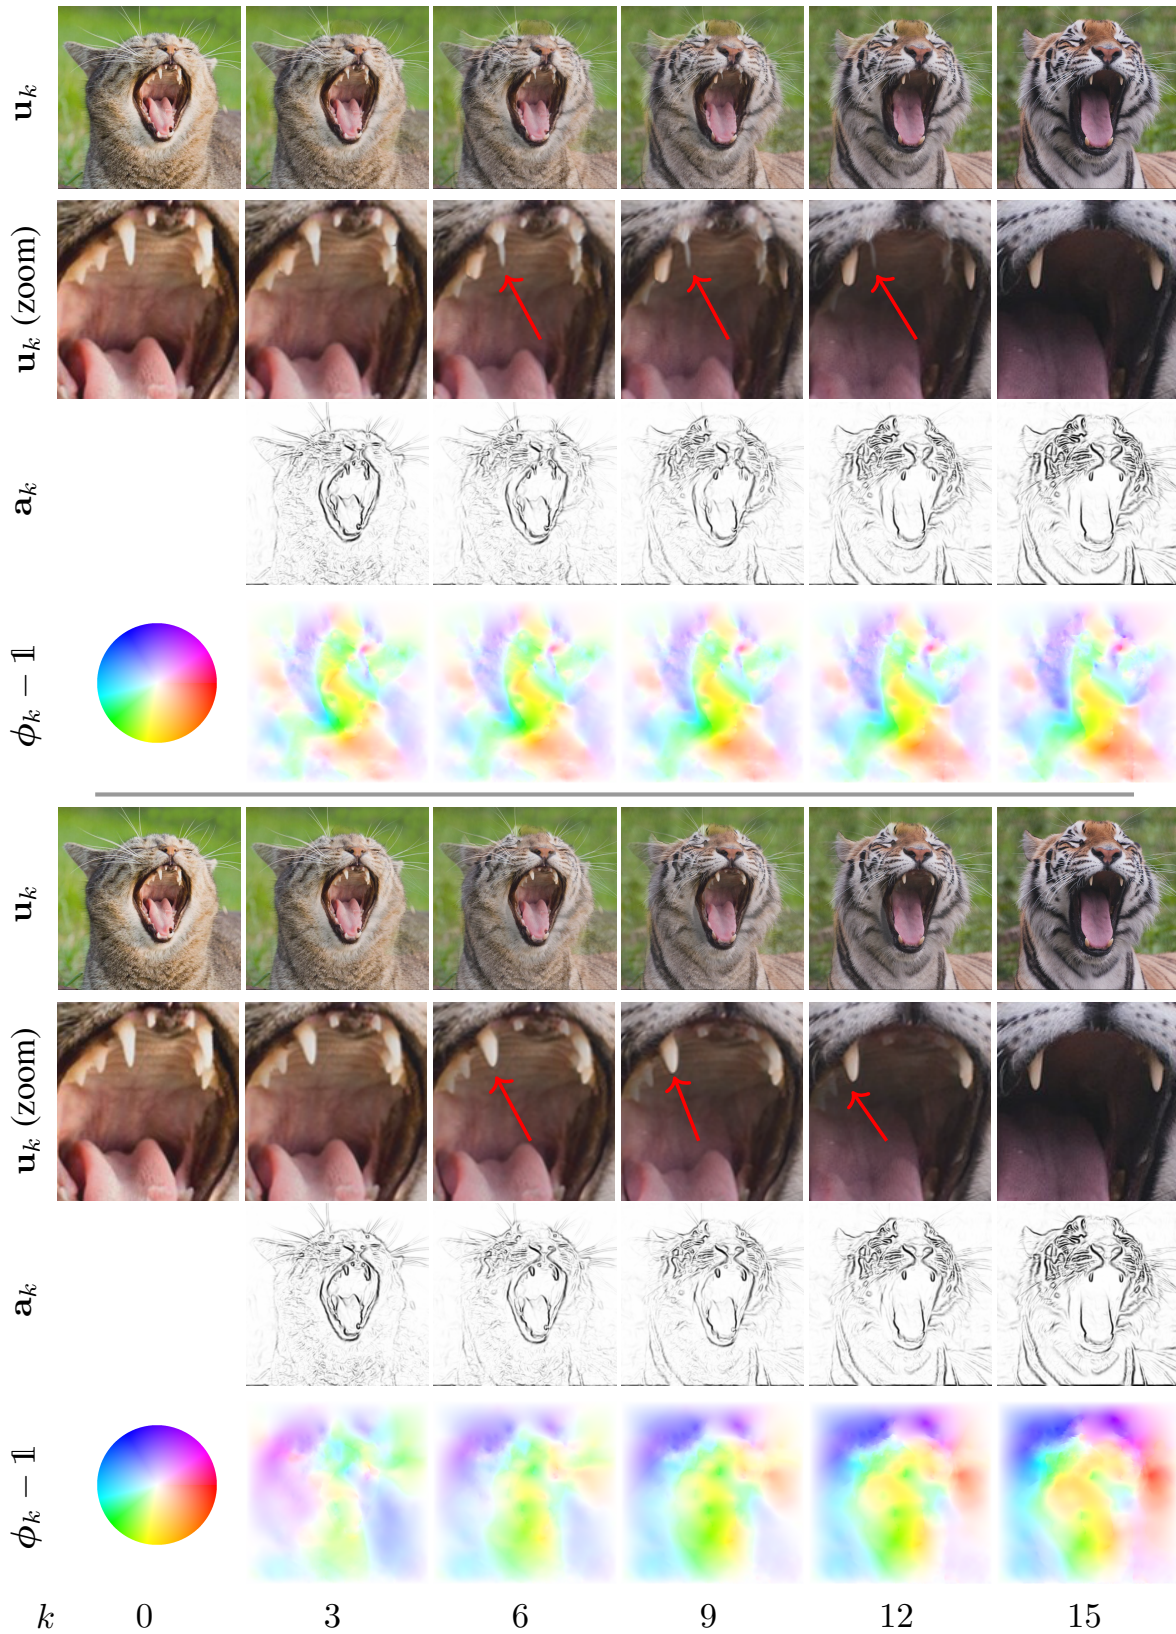


Figure 4.4: Time discrete geodesic sequences of animal photos for the RGB model (first row) and deep feature model (fifth row), along with a zoom of the mouth region with magnification factor 4 (second/sixth row), the associated sequences of anisotropy weights (third/seventh row), and color-coded displacement fields $\phi_k - \mathbb{1}$ (fourth/eighth row). Note that the novel deep feature-based model has significantly fewer blending artifacts as indicated by the arrows.



Figure 4.5: Pairs of time discrete geodesic paths using the deep feature model and corresponding color-coded displacement fields for paintings of US presidents (first/second row) as well as for paintings of Catherine the Great (third/forth row).



Figure 4.6: Visualization of the anisotropy in RGB model for a significant smaller value $\rho = 200$ compared to Figure 4.3: anisotropy weight (left) and color-coded displacement field (right) for $k = 12$.

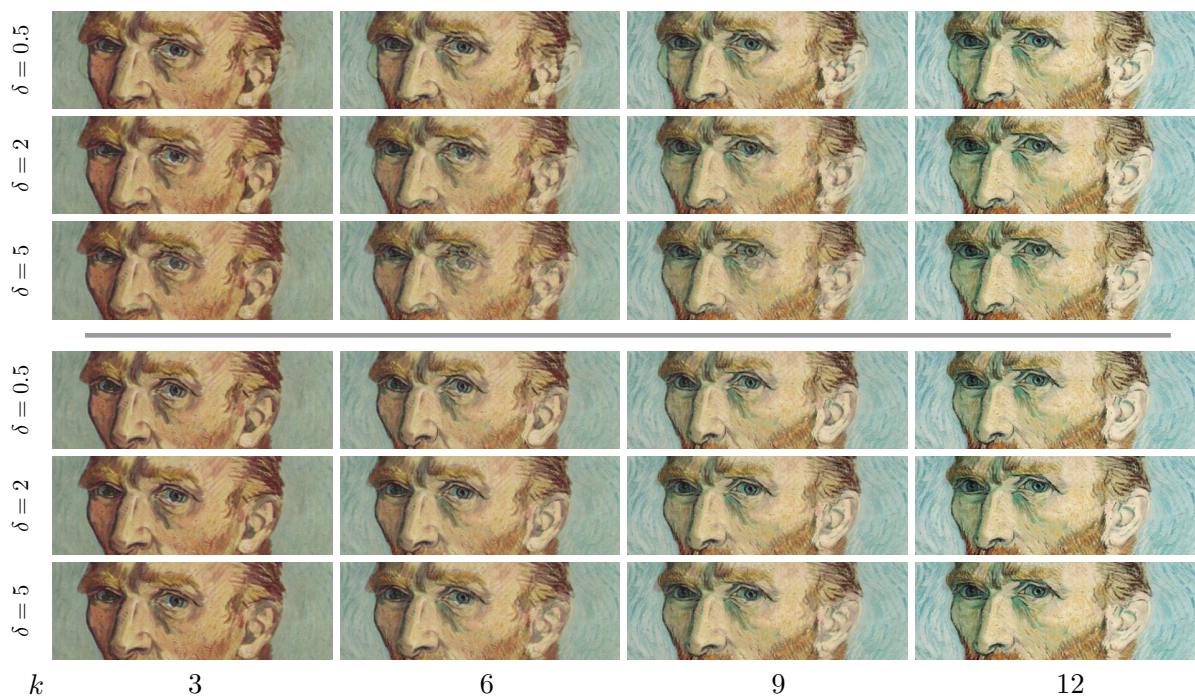


Figure 4.7: The part of the face region for the example with the van Gogh self-portraits and variation of the parameter δ for RGB model (first to third row) and deep feature model (fourth to sixth row).

4.8 Extensions, Conclusion and Outlook

Besides the spline interpolation for the image metamorphosis model which will be presented in detail in Chapter 5, there are two more projects where the author of the thesis was involved that can be considered as an extension of or being influenced by the model presented in this chapter.

– Bézier curves in deep feature space with anisotropic regularization were introduced in the Master’s thesis of Franziska Henrich [Hen19]. This work is based on the model for Bézier curves in the space of images introduced in [ERS⁺15]. Computation is done using a Riemannian version of de Casteljau’s algorithm where both time continuous and time discrete Riemannian distances are defined via the metamorphosis model presented in this chapter (*cf.* Sections 4.1 and 4.2, respectively). Numerical results, based on spatial discretization and numerical optimization analogous to the exposition in this chapter (*cf.* Sections 4.5 and 4.6) show advantages of the deep feature model in comparison with the simple RGB model, as already noticed in numerical experiments in Section 4.7. Rigorous theoretical study of the model with proofs of the existence of time continuous and time discrete Bézier curves in the space of images, together with the Mosco-convergence of the discrete towards the continuous model can be found in [ERS⁺15].

– Variational time discretization of the flow of diffeomorphism model (*cf.* Section 3.2.1) was introduced in the Master’s thesis of Mara Guastini [Gua22]. Together with a study of the weak form of geodesic equation, this study was published in the conference proceeding [GRRW23]. There, given a diffeomorphism $\psi_B \in \mathcal{D} = \mathcal{D}^m(\Omega)$ (*cf.* (4.10)), a discrete geodesic path $\Psi^K = (\psi_0 = \mathbb{1}, \psi_1, \dots, \psi_K = \psi_B)$ in the space of Sobolev diffeomorphisms \mathcal{D} is defined as a minimizer of the discrete path energy

$$\begin{aligned} \mathbf{E}^K[(\psi_0, \psi_1, \dots, \psi_K)] &:= K \sum_{k=1}^K \int_{\Omega} W(D(\psi_k \circ \psi_{k-1}^{-1})) + \gamma |D^m(\psi_k \circ \psi_{k-1}^{-1})|^2 dx \\ &= K \sum_{k=1}^K \int_{\Omega} W(D\phi_k) + \gamma |D^m \phi_k|^2 dx, \end{aligned}$$

where the energy density W satisfies conditions (W1)–(W3). Here, $\phi_k \circ \psi_{k-1} = \psi_k$, for every $k = 1, \dots, K$ giving $\psi_B = \phi_K \circ \dots \circ \phi_1$. Following along the lines of the proof of Theorem 4.2.7, the existence of a discrete geodesic path is shown given the proximity of ψ_B to $\mathbb{1}$ in $L^2(\Omega)$ (*cf.* [Gua22, Theorem 2.18]). Furthermore, upon constructing time extensions of the time discrete path and time discrete energy, based on an approach analogous to the one presented in Section 4.3, one can follow along the lines of proof of Theorem 4.4.1 to show the Γ -convergence (*cf.* Definition 2.4.1) as $K \rightarrow \infty$ in $C^{0,\beta}([0, 1], C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$, $0 < \beta < \min(\frac{1}{2}, \alpha)$ of the time extension of the energy \mathbf{E}^K towards the flow of diffeomorphism path energy

$$\mathcal{E}[(\psi_t)_{t \in [0,1]}] = \int_0^1 \int_{\Omega} L[v_t, v_t] dx dt,$$

where $v_t = \dot{\psi}_t \circ \psi_t^{-1}$ is the Eulerian velocity and the elliptic operator is given by (3.10) (*cf.* [Gua22, Theorem 2.20]). The time discrete geodesic path in the flow of diffeomorphism is then obtained by

$$\psi_k = \phi_k \circ \psi_{k-1} = \phi_k \circ \dots \circ \phi_1, \quad k = 1, \dots, K. \quad (4.72)$$

However, as one does not directly include the constraint $\phi_k \circ \psi_{k-1} = \psi_k$, $k = 1, \dots, K$, it can not be expected to have $\psi_K = \psi_B$, and thus the obtained results are not of a practical use. To this end, a relaxed approach is taken into account by introducing the discrete path energy

$$\mathbf{E}_{\delta}^K[\mathbf{Z}^K = (\zeta_0, \dots, \zeta_K)] := \inf_{\Phi \in \mathcal{D}^K} K \sum_{k=1}^K \int_{\Omega} W(D\phi_k) + \gamma |D^m \phi_k|^2 + \frac{1}{\delta} |\zeta_k \circ \phi_k - \zeta_{k-1}|^2 dx,$$

where, formally, one thinks of ζ_k as being the inverses of ψ_k and thus penalizes $\psi_k^{-1} \circ \phi_k - \psi_{k-1}^{-1}$. Then, for $\zeta_0 = \mathbb{1}$ and $\zeta_K = \zeta_B \in \mathcal{D}$, we define a time discrete geodesic path connecting these diffeomorphisms as a minimizer of the above energy. The discrete “forward flow” $\Psi^K = (\psi_0 = \mathbb{1}, \dots, \psi_K)$ is retrieved by (4.72), and can be thought of as a discrete geodesic connecting $\mathbb{1}$ and $\psi_K \approx \zeta_B^{-1}$ for $\delta \rightarrow 0$ (*cf.* Figure 4.10). One immediately observes

that the above energy is analogous to the time discrete metamorphosis energy introduced in Section 4.2, so we can apply the existence result (Theorem 4.2.7), the time extension operators (Section 4.3), and the Mosco-convergence result (Theorem 4.4.1), where the time continuous energy (whose minimizer exists by repeating the arguments of Theorem 4.4.4) is given by

$$\mathcal{E}_\delta[(\zeta_t)_{t \in [0,1]}] = \int_0^1 \int_\Omega L[v_t, v_t] + \frac{1}{\delta} z^2 \, dx \, dt,$$

for z being the (scalar) material derivative of ζ , *i.e.*,

$$|\zeta_t \circ \psi_t - \zeta_s \circ \psi_s| \leq \int_s^t z_r \circ \psi_r \, dr, \quad \forall s \leq t \in [0, 1].$$

For more details see [Gua22, Theorems 3.10 and 3.13]. We point out that, in this model, δ is meant as a penalization parameter, such that for $\delta \rightarrow 0$ one obtains the “original” flow of diffeomorphism model (*cf.* Figure 4.10). This should be countered to the role of δ in the metamorphosis model as a parameter determining a proper scaling between the flow velocity and the image intensity change energies.

The fully discrete model and the numerical optimization are obtained by an approach analogous to Sections 4.5 and 4.6. To obtain an input for numerical experiments, one first creates a time-independent spatially smooth velocity field $\tilde{\mathbf{v}}(\mathbf{x}, \mathbf{y}) = \mathbf{v}(\mathbf{x}, \mathbf{y})\mathbf{r}(\sqrt{\mathbf{x}^2 + \mathbf{y}^2})$, where multiplication with \mathbf{r} is necessary to guarantee that the field vanishes on the boundary. Then one defines

$$\zeta_B = \zeta(1) \quad \text{for} \quad \dot{\zeta}(t, \mathbf{x}) = \tilde{\mathbf{v}}(\zeta(t, \mathbf{x})),$$

where the solution is obtained using an ODE solver in Python (*cf.* [Gua22, Section 4.3]). Figure 4.8 shows an example for \mathbf{v} constructed as a composition of counterclockwise rotation in the left part of Ω_{MN} and clockwise rotation in the right part, together with the corresponding ζ_B , represented by the deformation of the grid Ω_{MN} . In Figure 4.9, we see the discrete geodesic path together with the intermediate deformations and the forward path together. Finally, Figure 4.10 shows that deviation of ψ_K from inverse of ζ_B indeed decreases as $\delta \rightarrow 0$.

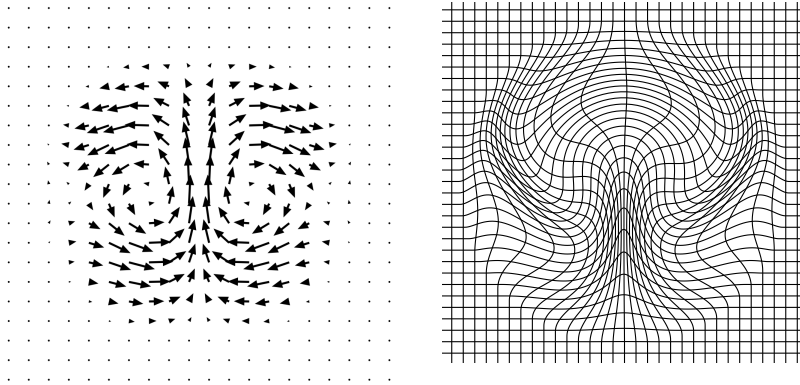


Figure 4.8: Representation of the velocity field $\tilde{\mathbf{v}}$ and the corresponding deformation ζ_B of Ω_{MN} [Gua22, Figure 2, Example 4].




Figure 4.9: Numerical results for input from Figure 4.8 with $K = 8$ and $\delta = 10^{-3}$: geodesic path ζ_k (first row), its zoom version on the central right part of Ω_{MN} (second row), deformations ϕ_k (third row), and forward flow ψ_k (fourth row) [Gua22, Figure 9].




Figure 4.10: Left: Values of $\|1 - \zeta_k \circ \psi_k\|_{L^2(\Omega_{MN})}^2$ for $k = 1, \dots, 8$ and $\delta = 1, 10^{-1}, 10^{-2}, 10^{-3}$ [Gua22, Figure 12 (d)]. Right: Deformation of Ω_{MN} under $\zeta_B \circ \psi_8$ for $\delta = 1, 10^{-1}, 10^{-2}, 10^{-3}$, from left to right [Gua22, Figure 12, column 4].

Conclusion and Outlook In this chapter, we have studied the metamorphosis model in a deep feature space obtained as a result of the application of a deep convolutional neural network to the original color channels. Furthermore, in order to allow for discontinuities in the flow/deformation fields, we introduced a regularization based on anisotropic diffusion operator proposed by Perona and Malik. This can be seen as an extension of the time continuous model proposed by Trouvé, Younes, and coworkers, while the variational time discretization was based on the work of Berkels, Effland, and Rumpf. We were able to show the existence of time discrete geodesic curves and the Mosco-convergence of a suitable extension of the time discrete path energy to the time continuous metamorphosis functional. As a corollary, we obtained the existence of time continuous geodesic curves as the limit of time discrete counterparts. Space discretization was based on finite differences and approximation of the warping operator via cubic Hermite splines. Experimental results, based on the iPALM algorithm developed by Pock and Sabach, show the superiority of the deep feature approach in comparison with the standard color-based model which is reflected in presence of significantly fewer blending artifacts. A possible improvement could be achieved by using a feature extraction operator which is trained on images that are used as input for the metamorphosis model. Furthermore, it would be interesting to compare the existing results with those obtained by restricting the optimization to the feature paths obtained as pairings of (scaled) images and their deep features.

Chapter 5

Spline Interpolation in Image Metamorphosis Model

THIS chapter investigates a variational model for spline interpolation in the image metamorphosis model, with the goal to obtain a smooth interpolation between several key frames in the space of images. The metamorphosis model, in more detail discussed in Section 3.2 and Chapter 4, is based on a simultaneous transport of image intensities and a modulation of intensities along motion trajectories. The underlying energy functional thus measures the motion velocity and the material derivative of the image intensity. In this chapter, motivated by the definition of cubic splines in Euclidean space as minimizers of the integral of the squared acceleration [dB63], and as stationary paths of the integrated squared covariant derivative of the velocity [NHP89] in a Riemannian context (see also Subsection 3.1.1 for a short introduction), we discuss a spline energy functional as a second order extension of the first order metamorphosis path energy. This functional combines quadratic functionals of the Eulerian motion acceleration and the second material derivative representing acceleration in the change of intensities along motion paths. In fact, our model separates in a physically intuitive way the Eulerian flow acceleration and the second material derivative of the image intensity. This implies that the model is not fully Riemannian in the above context, since the integral over the squared covariant derivative of the path velocity in the Riemannian metric would lead to an interwoven model of these different types of acceleration. Given a set of key frames at disjoint times, a spline interpolation path is then given as a minimizer of the spline energy additionally regularized by weighted geodesic energy, subject to the key frame interpolation constraint. Furthermore, a variational time discretization of this spline model is proposed as an extension of approaches from [BER15, EKP⁺21] (see also Sections 3.1.2 and 4.2), and the convergence to the time continuous model is discussed via Mosco-convergence methodology. As a byproduct, this also allows us to establish the existence of metamorphosis spline interpolation in the time continuous model as a limit of the corresponding time discrete spline interpolations. Spatial discretization is based on a finite difference discretization in space, combined with a stable B-spline approximation of the warping operator. A variety of numerical examples demonstrates the robustness and versatility of the proposed method in applications, where we use a variant of the iPALM algorithm [PS16] for the minimization of the fully discrete energy functional. Note that this chapter is an extended version of the previous conference proceeding [JRR21] and its journal version [JRR23] (joint project with Jorge Justiniano and Martin Rumpf).

This chapter is organized as follows. In Section 5.1, we introduce the appropriate (weak) versions of the flow and image intensity acceleration terms which leads to the derivation of the time continuous spline energy. Then, in Section 5.2, a variational time discretization of the continuous spline energy is introduced and the existence of discrete splines is studied. Section 5.3 represents the time extension of the time discrete quantities, which allows the study of convergence to the time continuous model, presented in Section 5.4. Before defining the fully discrete model in Section 5.6, we make a suitable relaxation of the previously defined energies in Section 5.5. This model allows an efficient numerical optimization, based on a version of the iPALM algorithm, in Section 5.7. Section 5.8 experimentally demonstrates properties of the spline interpolation of a set of key frames based on this model and shows applications of the proposed method. Finally, Section 5.9 provides a conclusion and outlooks of the chapter.

5.1 Time Continuous Model

In this section, we study the spline interpolation of key frames in the time continuous setting of the metamorphosis model. In mathematical notation, given a set of J key frames $u_j^I \in \mathcal{I} := L^2(\Omega, \mathbb{R}^c)$, we ask for spline interpolation $(u_t)_{t \in [0,1]}$ which satisfies the constraints

$$u_{t_j} = u_j^I, \quad t_j \in [0, 1], \quad j = 1, \dots, J. \quad (5.1)$$

To this end, we recall that cubic splines in Euclidean space minimize the integral over the squared motion acceleration subject to position constraints [dB63], whereas linear interpolation is associated with the minimization of the integral over the squared motion velocity (see Subsection 3.1.1 for more on geodesic and spline interpolation in time continuous setting). In the case of the metamorphosis model (cf. Section 3.2 for more background), for a sufficiently smooth image curve the linear interpolation would correspond to the minimization of the path energy

$$\mathcal{E}[u] = \int_0^1 \min_{v: \bar{\Omega} \rightarrow \mathbb{R}^n} \int_{\Omega} L[v, v] + \frac{1}{\delta} \left| \frac{D}{\partial t} u \right|^2 dx dt,$$

where $(v_t = \dot{\psi}_t \circ \psi_t^{-1})_{t \in [0,1]}$ is the Eulerian velocity of the underlying diffeomorphic flow $(\psi_t)_{t \in [0,1]}$. In this chapter, we use the elliptic operator (cf. (3.10))

$$L[v, v] = \text{tr}(\varepsilon[v]^2) + \gamma |D^m v|^2, \quad m > 1 + \frac{n}{2}, \gamma > 0. \quad (5.2)$$

This motivates the introduction of the following acceleration quantities that will be the ingredients of the spline energy. The *Eulerian flow acceleration* is defined by

$$a_t \circ \psi_t := \ddot{\psi}_t = \frac{d}{dt}(v_t \circ \psi_t) = (\dot{v}_t + Dv_t \cdot v_t) \circ \psi_t, \quad (5.3)$$

and for image paths with enough smoothness, the *second order material derivative* is given by

$$\frac{D^2}{\partial t^2} u_t \circ \psi_t := \frac{d}{dt}((\dot{u}_t + v_t \cdot Du_t) \circ \psi_t) = (\ddot{u}_t + v_t \cdot D\dot{u}_t + Du_t \cdot (\dot{v}_t + Dv_t v_t) + v_t \cdot (D\dot{u}_t + D^2 u_t v_t)) \circ \psi_t. \quad (5.4)$$

This splitting of acceleration term into flow acceleration and second order change of image intensity does not fully correspond to a Riemannian manifold approach (in the context of penalization of covariant derivative of velocity (cf. (3.3)) since the stronger interference of the two entities must be observed in the covariant derivative of the vector field $(v, \frac{D}{dt}u)$ (cf. Definition 2.2.17). Our model comes with the advantage that the splitting of the acceleration terms is a physically more intuitive way.

Since the derivation of a spline model requires minimizing integrals over quadratic acceleration quantities, one is naturally led to the following (still formal) spline energy:

$$\mathcal{F}[u] := \min_a \int_0^1 \int_{\Omega} L[a, a] + \frac{1}{\delta} \left| \frac{D^2}{dt^2} u \right|^2 dx dt,$$

where, for simplicity, we use the same elliptic operator (5.2). As in the case of geodesic path energy, we give rigorous formulations for general paths $u \in L^2((0, 1), \mathcal{I})$.

Let us first discuss the velocity and acceleration of the diffeomorphic flow $(\psi_t)_{t \in [0,1]}$. We observe pairs $(v, a) \in L^2((0, 1), \mathcal{V}) := H_0^1(\bar{\Omega}, \mathbb{R}^n) \cap H^m(\Omega, \mathbb{R}^n) \times L^2((0, 1), \mathcal{V})$ that determine the system

$$\begin{aligned} v_t \circ \psi_t(x) &= \dot{\psi}_t(x), \quad \psi_0(x) = x \\ a_t \circ \psi_t(x) &= \ddot{\psi}_t(x). \end{aligned}$$

The first equation of the system was studied in more detail in Subsection 3.2.3, with results on the existence and regularity. In particular, for every $v \in L^2((0, 1), \mathcal{V})$ there exists a unique solution to the first equation of the system $\psi \in H^1((0, 1), H^m(\Omega, \Omega)) \subset C^{0, \frac{1}{2}}([0, 1], C^{1, \alpha}(\bar{\Omega}, \bar{\Omega}))$, where $\alpha \in (0, m - \frac{n}{2} - 1]$, and ψ_t is a diffeomorphism for every $t \in [0, 1]$. Taking into consideration the second equation of the system, we have that any solution will satisfy

$\psi \in H^2((0, 1), H^m(\Omega, \Omega)) \subset C^{1, \frac{1}{2}}([0, 1], C^{1, \alpha}(\overline{\Omega}, \overline{\Omega}))$. Notice that from (5.3), one expects that the acceleration has one space derivative less in comparison to the velocity. However, the approach we later take allows us to have the same number of derivatives.

We can now turn to the study of the image intensity quantities. Let us recall that we can observe the vector material derivative $\hat{z} \in L^2((0, 1), L^2(\Omega, \mathbb{R}^c))$ given by

$$u_t \circ \psi_t - u_s \circ \psi_s = \int_s^t \hat{z}_r \circ \psi_r \, dr, \quad \forall s, t \in [0, 1], \quad (5.5)$$

or the scalar material derivative $z \in L^2((0, 1) \times \Omega)$ given in terms of

$$|u_t \circ \psi_t - u_s \circ \psi_s| \leq \int_s^t z_r \circ \psi_r \, dr, \quad \forall 1 \geq t > s \geq 0. \quad (5.6)$$

In analogy to (5.5), we can observe the material acceleration $\hat{w} \in L^2((0, 1), L^2(\Omega, \mathbb{R}^c))$ that satisfies

$$\hat{z}_t \circ \psi_t - \hat{z}_s \circ \psi_s = \int_s^t \hat{w}_r \circ \psi_r \, dr, \quad \forall s, t \in [0, 1].$$

By further using (5.5), we have

$$\begin{aligned} \int_0^\tau \int_s^t \hat{w}_{r+l} \circ \psi_{r+l} \, dr \, dl &= \int_0^\tau \hat{z}_{t+\tau} \circ \psi_{t+\tau} - \hat{z}_{s+\tau} \circ \psi_{t+\tau} \, dl \\ &= (u_{t+\tau} \circ \psi_{t+\tau} - u_t \circ \psi_t) - (u_{s+\tau} \circ \psi_{s+\tau} - u_s \circ \psi_s), \end{aligned} \quad (5.7)$$

for all $s, t \in (0, 1)$ and every τ , such that $t + \tau, s + \tau \in [0, 1]$. Observe that for $s = t - \tau$ the right-hand side of (5.7) is an integral version of the second order central difference. Because there is no differentiation involved in these definitions, they work for general image paths. Analogous to (5.6), we introduce the *scalar material acceleration* $w \in L^2((0, 1) \times \Omega)$ as a relaxation of the weak second order material derivative given by

$$\int_0^\tau \int_s^t w_{r+l} \circ \psi_{r+l} \, dr \, dl \geq |u_{t+\tau} \circ \psi_{t+\tau} - u_t \circ \psi_t - u_{s+\tau} \circ \psi_{s+\tau} + u_s \circ \psi_s|, \quad (5.8)$$

for every $s \leq t \in [0, 1]$ and $\tau \geq 0$, such that $t + \tau \in [0, 1]$. This relaxed Lagrangian approach is substantially more handsome in comparison to the Eulerian approach (5.4), which will be exploited in the proof of the consistency of continuous and time discrete approaches. The following proposition shows the equivalence of the two approaches corresponding to (5.7) and (5.8).

Proposition 5.1.1. *For every vector valued (\hat{z}, \hat{w}) satisfying (5.5) and (5.7), there exist scalar quantities (z, w) satisfying (5.6) and (5.8) with $z = |\hat{z}|$ and $w = |\hat{w}|$. Conversely, for every (z, w) satisfying (5.6) and (5.8), there exists (\hat{z}, \hat{w}) satisfying (5.5) and (5.7) with $z \geq |\hat{z}|$ and $w \geq |\hat{w}|$.*

Proof. The first claim follows from the triangle inequality. To prove the converse, recall that in Proposition 4.1.1, we showed that for z satisfying (5.6), there exists \hat{z} satisfying (5.5) with $z \geq |\hat{z}|$. This implies

$$\left| \int_0^\tau \hat{z}_{t+l} \circ \psi_{t+l} - \hat{z}_{s+l} \circ \psi_{s+l} \, dl \right| \leq \left| \int_0^\tau \int_s^t w_{r+l} \circ \psi_{r+l} \, dr \, dl \right|$$

for every $s \leq t \in [0, 1]$, and $\tau > 0$ such that $t + \tau \in [0, 1]$, and the same holds for integration on $[-\tau, 0]$ for $s - \tau \geq 0$. Taking the limit as τ tends to zero and using the Lebesgue differentiation theorem [Fol99, Theorem 3.21], we conclude that, for every $s \leq t \in [0, 1]$, we have

$$\left| \hat{z}_t \circ \psi_t - \hat{z}_s \circ \psi_s \right| \leq \int_s^t w_r \circ \psi_r \, dr.$$

From here, the function $t \mapsto \hat{z}_t \circ \psi_t$ is a.e. differentiable, with the derivative $w' \in L^2((0, 1), L^2(\Omega, \mathbb{R}^c))$ [AGS08, Remark 1.1.3] satisfying

$$\hat{z}_t \circ \psi_t - \hat{z}_s \circ \psi_s = \int_s^t w'_r \, dr = \int_s^t \hat{w}_r \circ \psi_r \, dr, \quad \forall s, t \in [0, 1]. \quad (5.9)$$

Finally, [AGS08, Theorem 1.1.2] implies $w \geq |\hat{w}|$, which finishes the proof. \square

In [HRW18], a regularization of the spline path energy by the addition of the weighted geodesic path energy was necessary for the existence and further analysis of the splines (*cf.* (3.4)). We follow the analogous approach which seems to be unavoidable also in our model. This finally leads to the following definitions.

Definition 5.1.2 (Time continuous regularized spline energy). Let $\sigma > 0$, $m > 1 + \frac{n}{2}$ be an integer, and $u \in L^2((0, 1), \mathcal{I})$ be an image curve. Then, the *regularized spline energy* is defined by

$$\mathcal{F}^\sigma[u] := \inf_{(v,a,z,w) \in \mathcal{C}[u]} \int_0^1 \int_\Omega L[a, a] + \frac{1}{\delta} w^2 + \sigma \left(L[v, v] + \frac{1}{\delta} z^2 \right) dx dt, \quad (5.10)$$

where $\mathcal{C}[u] \subset L^2((0, 1), \mathcal{V}) \times L^2((0, 1), \mathcal{V}) \times L^2((0, 1) \times \Omega) \times L^2((0, 1) \times \Omega)$ consists of tuples (v, a, z, w) satisfying

$$v_t \circ \psi_t(x) = \dot{\psi}_t(x), \quad \psi_0(x) = x, \quad (5.11)$$

$$a_t \circ \psi_t(x) = \ddot{\psi}_t(x), \quad \forall x \in \Omega, t \in [0, 1], \quad (5.12)$$

$$|u_t \circ \psi_t - u_s \circ \psi_s| \leq \int_s^t z_r \circ \psi_r dr, \quad \forall s \leq t \in [0, 1], \quad (5.13)$$

$$|u_{t+\tau} \circ \psi_{t+\tau} - u_t \circ \psi_t - u_{s+\tau} \circ \psi_{s+\tau} + u_s \circ \psi_s| \leq \int_0^\tau \int_s^{s+l} w_{r+l} \circ \psi_{r+l} dr dl, \quad \forall \tau, s + \tau \leq t + \tau \in [0, 1]. \quad (5.14)$$

The point evaluation in time of image path u is possible if the set $\mathcal{C}[u]$ is non-empty by Remark 4.1.4. Motivated by [dB63], we now define the continuous time spline interpolation for given key frames.

Definition 5.1.3 (Time continuous regularized spline interpolation). Let $\{u_j^I\}_{j=1, \dots, J} \in \mathcal{I}^J$. We call a minimizer $u \in L^2((0, 1), \mathcal{I})$ of \mathcal{F}^σ that satisfies (5.1) a *time continuous regularized spline interpolation* of $\{u_j^I\}_{j=1, \dots, J}$ with a regularization parameter σ .

Remark 5.1.4 (Improved regularity of image curves). As a corollary of Proposition 5.1.1, we have stronger regularity in the time variable of image curves with nonempty $\mathcal{C}[u]$. Namely, from (5.9), we have $\hat{z}_t \circ \psi_t \in H^1([0, 1], L^2(\Omega, \mathbb{R}^c))$. Plugging this back in (5.5), we have $u_t \circ \psi_t \in H^2([0, 1], \mathcal{I}) \subset C^{1, \frac{1}{2}}([0, 1], \mathcal{I})$, and by using the same technique as in Remark 4.1.4 together with $\psi \in C^{1, \frac{1}{2}}([0, 1], C^{1, \alpha}(\bar{\Omega}, \bar{\Omega}))$, we have $u \in C_\omega^1([0, 1], \mathcal{I})$, where the modulus of continuity depends on the corresponding norms of $(v, a, z, w) \in \mathcal{C}[u]$.

Remark 5.1.5 (Boundary constraints, *cf.* (3.5)). If we do not impose any additional constraints, we say the continuous time (regularized) spline interpolation has natural boundary constraints. Imposing periodic boundary conditions is equivalent to defining the image curve u_t on the sphere S^1 instead of the interval $[0, 1]$. In the case of Hermite boundary conditions, we ask for $v_0 = v_A, v_1 = v_B$ and, in light of Proposition 5.1.1 and differentiation of the left-hand side of (5.5), that $\hat{z}_0 = \hat{z}_A$, and $\hat{z}_1 = \hat{z}_B$. Note that in the case of Hermite boundary conditions (also called clamped boundary conditions), we implicitly require $t_1 = 0$ and $t_J = 1$, so that u_0 and u_1 are also prescribed.

5.2 Variational Time Discretization

In this section, we study the variational time discretization of the time continuous (regularized) spline energy. To this end, we pick up the approach of [BER15, EKP⁺21] for the variational time discretization of the geodesic energy, and thus build upon the exposition from Section 4.2. We consider a discrete image curve $\mathbf{u} = (u_0, \dots, u_K)$ with $u_k \in \mathcal{I}$, and define a set of admissible deformations

$$\mathcal{D} := \{\phi \in H^m(\Omega, \Omega), \det(D\phi) \geq \epsilon, \phi = \mathbb{1} \text{ on } \partial\Omega\}, \quad (5.15)$$

for a fixed (small) $\epsilon > 0$ (we comment on the case $\epsilon = 0$ in Remark 5.2.8). We recall that, considering $\mathbf{u} \in \mathcal{I}^{K+1}$ as time sampling of a sufficiently regular image curve at points $\frac{k}{K}, k = 0, \dots, K$, and $\Phi = (\phi_1, \dots, \phi_K) \in \mathcal{D}^K$ as the relative flow $\phi_k = \psi_{\frac{k}{K}} \circ \psi_{\frac{k-1}{K}}^{-1}$, and using the forward finite difference approximations, we obtained the discrete version of the Eulerian velocity $v_k := K(\phi_k - \mathbb{1})$ and $\hat{z}_k := K(u_k \circ \phi_k - u_{k-1})$ for the discrete material derivative,

with $k = 1, \dots, K$. By further taking this perspective and using the central finite difference approximation for the second derivative, we define the discrete acceleration by

$$a_k := K^2((\phi_{k+1} - \mathbb{1}) \circ \phi_k - (\phi_k - \mathbb{1})), \quad k = 1, \dots, K-1, \quad (5.16)$$

and the discrete version of the second order material derivative by

$$\hat{w}_k := K(\hat{z}_{k+1} \circ \phi_k - \hat{z}_k) = K^2(u_{k+1} \circ \phi_{k+1} \circ \phi_k - 2u_k \circ \phi_k + u_{k-1}), \quad k = 1, \dots, K-1. \quad (5.17)$$

Analogous to Section 4.2, we consider the discrete path energy

$$\mathbf{E}^{K,D}[\mathbf{u}, \Phi] := K \sum_{k=1}^K \int_{\Omega} W_D(D\phi_k) + \gamma |D^m \phi_k|^2 + \frac{1}{\delta} |u_k \circ \phi_k - u_{k-1}|^2 dx,$$

where $W_D(B) := |B^{\text{sym}} - \mathbb{1}|^2$ is an elastic energy density corresponding to the simplified elliptic operator given by (5.2). Then, in analogy with (5.10), the discrete counterpart of the spline energy is defined by

$$\mathbf{F}^{K,D}[\mathbf{u}, \Phi] := \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} W_A(Da_k) + \gamma |D^m a_k|^2 + \frac{1}{\delta} |\hat{w}_k|^2 dx, \quad (5.18)$$

with the energy density $W_A(B) := |B^{\text{sym}}|^2$. We notice that the discrete acceleration is indeed in $H^m(\Omega, \Omega)$, by the closedness of \mathcal{D} under the composition [IKT13, Proposition 2.19]. Finally, the regularized time discrete spline energy is, for $\sigma > 0$, given by

$$\mathbf{F}^{\sigma,K,D}[\mathbf{u}, \Phi] := \mathbf{F}^{K,D}[\mathbf{u}, \Phi] + \sigma \mathbf{E}^{K,D}[\mathbf{u}, \Phi].$$

In order to talk about spline interpolation, we must define the discrete version of the interpolation constraints (5.1). Let $I^K := (i_1, \dots, i_{J^K})$ be an index tuple with $2 \leq J^K \leq K$, $i_j \in \{0, \dots, K\}$ for $j = 1, \dots, J^K$. We consider a J^K -tuple $\mathcal{I}_f^K := (u_1^I, \dots, u_{J^K}^I)$ and define the set of admissible image vectors

$$\mathcal{I}_{adm}^K := \{\mathbf{u} \in \mathcal{I}^{K+1}, u_{i_j} = u_j^I, j = 1, \dots, J^K\}. \quad (5.19)$$

We are now in a position to define the time discrete spline interpolation.

Definition 5.2.1 (Time discrete regularized spline interpolation). Let $\sigma > 0$ and $\mathbf{u} = (u_0, \dots, u_K) \in \mathcal{I}_{adm}^K$. We set

$$\mathbf{F}^{\sigma,K}[\mathbf{u}] := \inf_{\Phi \in \mathcal{D}^K} \mathbf{F}^{\sigma,K,D}[\mathbf{u}, \Phi]. \quad (5.20)$$

A time discrete regularized spline interpolation of $\{u_j^I\}_{j=1, \dots, J^K}$ with a regularization parameter σ is a $(K+1)$ -tuple that minimizes $\mathbf{F}^{\sigma,K}$ over all discrete paths in \mathcal{I}_{adm}^K .

Remark 5.2.2 (Time discrete boundary constraints, cf. Remark 5.1.5). The presented discretization is valid in the case of natural boundary conditions, to which we restrict in further discussions. We remark that, in the case of periodic boundary conditions, we make an identification $K \hat{=} 0$, $K+1 \hat{=} 1$, and the sum in (5.18) goes up to K . For the discrete version of the Hermite boundary conditions, we prescribe $\phi_1 = \bar{\phi}_1$, $\phi_K = \bar{\phi}_K$, $u_0 = \bar{u}_0$, $u_K = \bar{u}_K$, and $\hat{z}_1 = \bar{\hat{z}}_1$, $\hat{z}_K = \bar{\hat{z}}_K$, for given $\bar{\phi}_1, \bar{\phi}_K \in \mathcal{D}$, $\bar{u}_0, \bar{u}_K \in \mathcal{I}$, and $\bar{\hat{z}}_1, \bar{\hat{z}}_K \in L^2(\Omega, \mathbb{R}^c)$.

Next, we follow ideas from Section 4.2 to prove the existence of a time discrete spline interpolation. Here, and throughout the chapter, we often omit the attribute ‘‘regularized’’ as we always consider that setting and no confusion can occur. We start with a lemma which is the analogous result to Lemma 4.2.2.

Lemma 5.2.3. *There exists a constant C which only depends on Ω, m, n, γ such that*

$$\|\phi - \mathbb{1}\|_{H^m(\Omega)} \leq C \sqrt{C_\phi}$$

for all $\phi \in \mathcal{D}$ satisfying $\int_{\Omega} W_D(D\phi) + \gamma |D^m \phi|^2 dx \leq C_\phi$.

Proof. An application of the Gagliardo–Nirenberg inequality 2.1.5 yields

$$\|\phi - \mathbb{1}\|_{H^m(\Omega)} \leq C(\|\phi - \mathbb{1}\|_{L^2(\Omega)} + |\phi - \mathbb{1}|_{H^m(\Omega)}). \quad (5.21)$$

The last term in (5.21) is bounded by

$$|\phi - \mathbb{1}|_{H^m(\Omega)} = |\phi|_{H^m(\Omega)} \leq \sqrt{\frac{C_\phi}{\gamma}}. \quad (5.22)$$

To estimate the lower order term on the right-hand side, we use Korn's inequality 2.1.6 and Poincaré's inequality 2.1.4 to obtain

$$\|\phi - \mathbb{1}\|_{L^2(\Omega)} \leq C\|\varepsilon[\phi - \mathbb{1}]\|_{L^2(\Omega)} \leq C\sqrt{C_\phi}. \quad (5.23)$$

Thus, the lemma follows by combining (5.21), (5.22) and (5.23). \square

Remark 5.2.4. The analogous result holds for the boundedness of acceleration: $\int_\Omega W_A(Da) + \gamma|D^m a|^2 dx \leq C_a$ implies $\|a\|_{H^m(\Omega)} \leq C\sqrt{C_a}$.

We now show the well-posedness of (5.20).

Proposition 5.2.5. *For every $K \in \mathbb{N}$ and every vector of images $\mathbf{u} = (u_0, \dots, u_K) \in \mathcal{I}^{K+1}$, there exists a vector of deformations $\Phi = (\phi_1, \dots, \phi_K) \in \mathcal{D}^K$ such that*

$$\mathbf{F}^{\sigma, K, D}[\mathbf{u}, \Phi] = \inf_{\tilde{\Phi} \in \mathcal{D}^K} \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \tilde{\Phi}].$$

Proof. Let $\{\Phi^j\}_{j \in \mathbb{N}} \in \mathcal{D}^K$ be a sequence for which it holds

$$\lim_{j \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \Phi^j] = \inf_{\tilde{\Phi} \in \mathcal{D}^K} \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \tilde{\Phi}], \quad \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \Phi^j] \leq \overline{\mathbf{F}^{\sigma, K}} := \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \mathbb{1}^K].$$

By Lemma 5.2.3, we have

$$\|\phi_k^j - \mathbb{1}\|_{H^m(\Omega)} \leq C\sqrt{\frac{\overline{\mathbf{F}^{\sigma, K}}}{K}}, \quad \forall j \in \mathbb{N}, k = 1, \dots, K.$$

Thus, $\{\Phi^j\}_{j \in \mathbb{N}}$ is uniformly bounded in $H^m(\Omega, \Omega)^K$. By reflexivity of this space, there exists a subsequence (with the same label) such that $\Phi^j \rightharpoonup \Phi$ in $H^m(\Omega, \Omega)^K$. By the compact Sobolev embedding (cf. Theorem 2.1.8), we have $\Phi^j \rightarrow \Phi$ in $C^{1, \alpha}(\overline{\Omega}, \overline{\Omega})^K$, $\alpha \in (0, m - \frac{n}{2} - 1]$, which gives us that $\Phi \in \mathcal{D}^K$. Analogously, by Remark 5.2.4, we have the boundedness of $\{\mathbf{a}^j\}_{j \in \mathbb{N}}$ in $H^m(\Omega, \Omega)^{K-1}$, and thus a convergent subsequence satisfying $\mathbf{a}^j \rightarrow \mathbf{a}$ in $H^m(\Omega, \Omega)^{K-1}$, and $\mathbf{a}^j \rightarrow \mathbf{a}$ in $C^{1, \alpha}(\overline{\Omega}, \overline{\Omega})^{K-1}$. Here, for every $j \in \mathbb{N}$ we have $\mathbf{a}^j = (a_1^j, \dots, a_{K-1}^j)$ given by (5.16). From the strong convergence of deformations, we have that $a_k = K^2(\phi_{k+1} \circ \phi_k - 2\phi_k + \mathbb{1})$ for all $k = 1, \dots, K-1$. Using weak lower semicontinuity of H^m -seminorm and continuity of energy densities, we have for all k , as $j \rightarrow \infty$:

$$\begin{aligned} \liminf |\phi_k^j|_{H^m(\Omega)} &\geq |\phi_k|_{H^m(\Omega)}, \quad \|\mathbb{W}_D(D\phi_k^j)\|_{L^1(\Omega)} \rightarrow \|\mathbb{W}_D(D\phi_k)\|_{L^1(\Omega)}, \\ \liminf |a_k^j|_{H^m(\Omega)} &\geq |a_k|_{H^m(\Omega)}, \quad \|\mathbb{W}_A(Da_k^j)\|_{L^1(\Omega)} \rightarrow \|\mathbb{W}_A(Da_k)\|_{L^1(\Omega)}. \end{aligned} \quad (5.24)$$

Finally, using Remark 4.2.4, we have $u \circ \phi_k^j \rightarrow u \circ \phi_k$ in \mathcal{I} for all $k = 1, \dots, K$, which finally gives

$$\begin{aligned} \|u_k \circ \phi_k^j - u_{k-1}\|_{\mathcal{I}} &\rightarrow \|u_k \circ \phi_k - u_{k-1}\|_{\mathcal{I}}, \quad k = 1, \dots, K, \\ \|u_{k+1} \circ \phi_{k+1}^j \circ \phi_k^j - 2u_k \circ \phi_k^j + u_{k-1}\|_{\mathcal{I}} &\rightarrow \|u_{k+1} \circ \phi_{k+1} \circ \phi_k - 2u_k \circ \phi_k + u_{k-1}\|_{\mathcal{I}}, \quad k = 1, \dots, K-1, \end{aligned}$$

which, together with (5.24), gives $\lim_{j \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \Phi^j] \geq \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \Phi]$, finally finishing the proof. \square

In the next step, under suitable conditions, we prove that there exists a minimizing vector in \mathcal{I}_{adm}^K (see (5.19)) for a fixed vector of deformations $\Phi \in \mathcal{D}^K$.

Proposition 5.2.6. *Let $K \geq 2$, \mathcal{I}_f^K and $\Phi \in \mathcal{D}^K$ be fixed. Assume that the deformations satisfy, for every $x \in \Omega$,*

$$C_{\det} \geq \det(D\phi_k(x)) \geq c_{\det} > 0, \quad k = 1, \dots, K. \quad (5.25)$$

Then, there exists a vector of images $\mathbf{u} \in \mathcal{I}_{adm}^K$ such that

$$\mathbf{F}^{\sigma, K, D}[\mathbf{u}, \Phi] = \inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K, D}[\mathbf{v}, \Phi].$$

Proof. Let $\{\mathbf{u}^j\}_{j \in \mathbb{N}} \in \mathcal{I}_{adm}^K$ be a sequence such that

$$\lim_{j \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\mathbf{u}^j, \Phi] = \inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K, D}[\mathbf{v}, \Phi] \leq \overline{\mathbf{F}^{\sigma, K, D}}.$$

Here, $\overline{\mathbf{F}^{\sigma, K, D}} := \mathbf{F}^{\sigma, K, D}[\bar{\mathbf{u}}, \Phi]$ represents a finite upper bound for the energy, with the vector of images $\bar{\mathbf{u}}$ satisfying $\bar{u}_k = u_1^I$ for $0 \leq k \leq i_1$, $\bar{u}_k = u_j^I$ for $i_j < k \leq i_{j+1}$ with $j = 1, \dots, K-1$, and $\bar{u}_k = u_{j_K}^I$ for $i_{j_K} \leq k \leq K$. Indeed, we have

$$\begin{aligned} \overline{\mathbf{F}^{\sigma, K, D}} &\leq \sum_{k=1}^{K-1} \frac{1}{K} \left(\|W_A(Da_k)\|_{L^1(\Omega)} + \gamma |a_k|_{H^m(\Omega)}^2 \right) \\ &\quad + \sigma \sum_{k=1}^K K \left(\|W_D(D\phi_k)\|_{L^1(\Omega)} + \gamma |\phi_k|_{H^m(\Omega)}^2 \right) + K^3 (1 + c_{\det}^{-1})^2 \sum_{j=1}^{j_K} \|u_j^I\|_{\mathcal{I}}^2, \end{aligned}$$

where we used the transformation formula and (5.25). By a further use of (5.25), we have

$$\begin{aligned} \|u_k^j\|_{\mathcal{I}} &\leq \|u_{k+1}^j \circ \phi_{k+1} - u_k^j\|_{\mathcal{I}} + \|u_{k+1}^j \circ \phi_{k+1}\|_{\mathcal{I}} \leq \sqrt{\frac{\delta \overline{\mathbf{F}^{\sigma, K, D}}}{K}} + c_{\det}^{-\frac{1}{2}} \|u_{k+1}^j\|_{\mathcal{I}}, \\ \|u_{k+1}^j\|_{\mathcal{I}} &\leq C_{\det}^{-\frac{1}{2}} \|u_{k+1}^j \circ \phi_{k+1}\|_{\mathcal{I}} \leq C_{\det}^{-\frac{1}{2}} \left(\|u_{k+1}^j \circ \phi_{k+1} - u_k^j\|_{\mathcal{I}} + \|u_k^j\|_{\mathcal{I}} \right) \leq C_{\det}^{-\frac{1}{2}} \left(\sqrt{\frac{\delta \overline{\mathbf{F}^{\sigma, K, D}}}{K}} + \|u_k^j\|_{\mathcal{I}} \right), \end{aligned} \quad (5.26)$$

from where we have, by induction, that $\{u_k^j\}_{j \in \mathbb{N}}$ is uniformly bounded in \mathcal{I} , for every $k = 0, \dots, K$. By reflexivity, there exists a subsequence (labeled in the same way) such that $u_k^j \rightharpoonup u_k$ for some $\mathbf{u} \in \mathcal{I}_{adm}^K$. Then, by Remark 4.2.8, we have $\mathbf{u}^j \circ \Phi \rightharpoonup \mathbf{u} \circ \Phi$ in \mathcal{I}^K , which, together with weak lower semicontinuity of the norm, gives

$$\begin{aligned} \|u_k \circ \phi_k - u_{k-1}\|_{\mathcal{I}}^2 &\leq \liminf_{j \rightarrow \infty} \|u_k^j \circ \phi_k - u_{k-1}^j\|_{\mathcal{I}}^2, \\ \|u_{k+1} \circ \phi_{k+1} \circ \phi_k - 2u_k \circ \phi_k + u_{k-1}\|_{\mathcal{I}}^2 &\leq \liminf_{j \rightarrow \infty} \|u_{k+1}^j \circ \phi_{k+1} \circ \phi_k - 2u_k^j \circ \phi_k + u_{k-1}^j\|_{\mathcal{I}}^2, \end{aligned}$$

for every $k = 1, \dots, K$ and $k = 1, \dots, K-1$, respectively. Altogether, we have

$$\liminf_{j \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\mathbf{u}^j, \Phi] \geq \mathbf{F}^{\sigma, K, D}[\mathbf{u}, \Phi],$$

from where the optimality follows. \square

We are now in a position to show the existence of a time discrete spline interpolation.

Theorem 5.2.7 (Existence of time discrete spline interpolations). *Let $K \geq 2$. Then, for every \mathcal{I}_f^K , there exists $\mathbf{u} \in \mathcal{I}_{adm}^K$ such that*

$$\mathbf{F}^{\sigma, K}[\mathbf{u}] = \inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K}[\mathbf{v}].$$

Proof. Consider a sequence $\{\mathbf{u}^j\}_{j \in \mathbb{N}} \in \mathcal{I}_{adm}^K$ for which it holds $\lim_{j \rightarrow \infty} \mathbf{F}^{\sigma, K}[\mathbf{u}^j] = \inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K}[\mathbf{v}] \leq \overline{\mathbf{F}^{\sigma, K}}$, where $\overline{\mathbf{F}^{\sigma, K}} := \mathbf{F}^{\sigma, K, D}[\bar{\mathbf{u}}, \mathbf{1}^K]$. Here, $\bar{u}_k = \bar{u}_{\frac{k}{K}}$ for $k = 0, \dots, K$, where $(\bar{u}_t)_{t \in [0,1]}$ is a curve smooth in

time such that $\bar{u}_{\frac{i_j}{K}} = u_j^I$ for every $j = 1, \dots, J^K$. Then, we have

$$\begin{aligned} & \mathbf{F}^{\sigma, K, D}[\bar{\mathbf{u}}^K, \mathbf{1}^K] \\ &= \sigma K \sum_{k=0}^K \int_{\Omega} |\bar{u}_{k+1}^K - \bar{u}_k^K|^2 dx + K^3 \sum_{k=1}^{K-1} \int_{\Omega} |\bar{u}_{k+1}^K - 2\bar{u}_k^K + \bar{u}_{k-1}^K|^2 dx \\ &\leq C \left(\int_{\Omega} |\bar{u}|_{H^1((0,1))}^2 + |\bar{u}|_{H^2((0,1))}^2 dx + 1 \right), \end{aligned} \quad (5.27)$$

where the constant is independent of K . Furthermore, for every $j \in \mathbb{N}$, let $\mathbf{F}^{\sigma, K}[\mathbf{u}^j] = \mathbf{F}^{\sigma, K, D}[\mathbf{u}^j, \Phi^j]$, where the existence of an optimal vector of deformations is shown in Proposition 5.2.5. By Lemma 5.2.3, we have the uniform boundedness of $\{\Phi^j\}_{j \in \mathbb{N}}$ in $H^m(\Omega, \Omega)^K$, and thus the weak convergence in this space and the strong convergence in $C^{1, \alpha}(\bar{\Omega}, \bar{\Omega})^K$ to some $\Phi \in \mathcal{D}^K$. Furthermore, by Remark 5.2.4, we have $a^j \rightharpoonup a$ in $H^m(\Omega, \Omega)^{K-1}$, and strongly in $C^{1, \alpha}(\bar{\Omega}, \bar{\Omega})^{K-1}$, where $a_k = K^2(\phi_{k+1} \circ \phi_k - 2\phi_k + \mathbf{1})$, and estimates from (5.24) are satisfied. By Proposition 5.2.6, we may replace \mathbf{u}^j by an energy optimal image vector. Keeping the same label and following the same arguments as in (5.26), we conclude that $\{\mathbf{u}^j\}_{j \in \mathbb{N}}$ is uniformly bounded in \mathcal{I} and it weakly converges to \mathbf{u} . Finally, the estimates

$$\begin{aligned} \|u_k \circ \phi_k - u_{k-1}\|_{\mathcal{I}}^2 &\leq \liminf_{j \rightarrow \infty} \|u_k^j \circ \phi_k^j - u_{k-1}^j\|_{\mathcal{I}}^2, \\ \|u_{k+1} \circ \phi_{k+1} \circ \phi_k - 2u_k \circ \phi_k + u_{k-1}\|_{\mathcal{I}}^2 &\leq \liminf_{j \rightarrow \infty} \|u_{k+1}^j \circ \phi_{k+1}^j \circ \phi_k^j - 2u_k^j \circ \phi_k^j + u_{k-1}^j\|_{\mathcal{I}}^2, \end{aligned}$$

for every $k = 1, \dots, K$ and $k = 1, \dots, K-1$, respectively, follow from Remark 4.2.8 and the weak lower semicontinuity of the norm. Altogether, $\inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K}[\mathbf{v}] = \liminf_{j \rightarrow \infty} \mathbf{F}^{\sigma, K}[\mathbf{u}^j] \geq \mathbf{F}^{\sigma, K}[\mathbf{u}]$, from where the optimality follows. \square

Remark 5.2.8. (i) The results of this section remain valid for any W_D satisfying conditions (W1)–(W2) from Section 4.2.

(ii) Let us observe the case when $i_j = K \cdot t_j$ and $u_{i_j} = u_j^I$ for every $j = 1, \dots, J$ (cf. (5.1) and (5.19)). Then, for every large enough K (depending on $\{u_j^I\}_{j=1, \dots, J}$ and $\{t_j\}_{j=1, \dots, J}$), one can show the existence of a discrete time spline interpolation even if $\epsilon = 0$ in the definition of the admissible set (5.15). Namely, by (5.27), we have that $\bar{\mathbf{F}}^{\sigma, K}$ is a fixed finite upper bound for the discrete spline energy, independent of K . Then, using Lemma 5.2.3 and the Sobolev embedding theorem 2.1.8, we have

$$\max_{k=1, \dots, K} \|\phi_k - \mathbf{1}\|_{C^{1, \alpha}(\bar{\Omega})} \leq C \max_{k=1, \dots, K} \|\phi_k - \mathbf{1}\|_{H^m(\Omega)} \leq C \sqrt{\frac{\bar{\mathbf{F}}^{\sigma, K}}{K}}.$$

By the Lipschitz continuity of the determinant, we have, for large enough K

$$\max_{k=1, \dots, K} \|\det(D\phi_k) - 1\|_{L^\infty(\Omega)} \leq C \sqrt{\frac{\bar{\mathbf{F}}^{\sigma, K}}{K}} < 1,$$

which proves $\min_{k=1, \dots, K} \det(D\phi_k) \geq c_{\det} > 0$ and we can proceed as before.

Remark 5.2.9. Variational time discretization of Riemannian splines was studied in [HRW18]. Following this work, a candidate for the time discretization of the spline energy is (cf. Subsection 3.1.2, (3.8))

$$\mathbf{F}^K[\mathbf{u}] = 4K^3 \sum_{k=1}^{K-1} \inf_{\phi_k \in \mathcal{D}} \int_{\Omega} W_D(D\phi_k) + |D^m \phi_k| + \frac{1}{\delta} |\tilde{u}_k \circ \phi_k - u_k|^2 dx$$

where \tilde{u}_k is the geodesic midpoint of u_{k-1} and u_{k+1} , i.e., $\tilde{u}_k \in \operatorname{argmin}_u \mathbf{E}^2[u_{k-1}, u, u_{k+1}]$. This approach is not suitable for our study, since, even after regularization with $\sigma > 0$, the geodesic midpoint is not continuous with respect to the weak convergence in \mathcal{I} , i.e., if we have $u_0^j, u_1^j \rightharpoonup u_0, u_1$ and for the sequence of geodesic midpoints $u_{\frac{1}{2}}^j \rightharpoonup u_{\frac{1}{2}}$, we do not necessarily have that $u_{\frac{1}{2}}$ is a geodesic midpoint of u_0 and u_1 .

5.3 Temporal Extension Operators

In this section, we define the suitable time extensions of the time discrete quantities from the previous section. We extend the approach presented in Section 4.3 in order to allow the study of spline interpolations.

Let $K \geq 2$, $\tau = \frac{1}{K}$, $t_k^K = k\tau$ for $k = 0, 1, \dots, K$ and $t_{k \pm \frac{1}{2}}^K = (k \pm \frac{1}{2})\tau$, for $k = 1, \dots, K$ and $k = 0, 1, \dots, K-1$, respectively. Consider a vector of images $\mathbf{u}^K = (u_0^K, \dots, u_K^K) \in \mathcal{I}^{K+1}$ and a vector of deformations $\Phi^K = (\phi_1^K, \dots, \phi_K^K) \in \mathcal{D}^K$. We first define the (incremental) discrete transport path y^K by $y_t^K = y_0^K(t, \cdot)$ for $t \in [0, t_{\frac{1}{2}}^K]$, $y_t^K = y_k^K(t, \cdot)$ for $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$ with $k = 1, \dots, K-1$, and $y_t^K = y_K^K(t, \cdot)$ for $t \in (t_{K-\frac{1}{2}}^K, 1]$, where

$$y_0^K(t, \cdot) := \mathbb{1} + \frac{t}{\tau}(\phi_1^K - \mathbb{1}), \quad y_K^K(t, \cdot) := \mathbb{1} + \frac{t - t_{K-1}^K}{\tau}(\phi_K^K - \mathbb{1}),$$

and, for $k = 1, \dots, K-1$

$$y_k^K(t, \cdot) := \frac{1}{2}(\mathbb{1} + \phi_k^K) + \frac{t - t_{k-\frac{1}{2}}^K}{\tau}(\phi_k^K - \mathbb{1}) + \frac{(t - t_{k-\frac{1}{2}}^K)^2}{2\tau^2}(\phi_{k+1}^K \circ \phi_k^K - 2\phi_k^K + \mathbb{1}). \quad (5.28)$$

This can be seen as a cubic Hermite interpolation (see [dB78, Chapter IV] for details, cf. (4.71)) on intervals $(t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$, and an affine interpolation on $[0, t_{\frac{1}{2}}^K]$ and $(t_{K-\frac{1}{2}}^K, 1]$. In particular, observe that $y_k^K(t_{k-\frac{1}{2}}^K, \cdot) = \frac{\mathbb{1} + \phi_k^K}{2}$ and $y_k^K(t_{k+\frac{1}{2}}^K, \cdot) = \frac{(\mathbb{1} + \phi_{k+1}^K) \circ \phi_k^K}{2}$, with the corresponding slopes $\frac{\phi_k^K - \mathbb{1}}{\tau}$ and $\frac{(\phi_{k+1}^K - \mathbb{1}) \circ \phi_k^K}{\tau}$, respectively. This is sketched on Figure 5.1 left.

Next, we define the image extension operator $\mathcal{U}^K[\mathbf{u}^K, \Phi^K] \in L^2([0, 1], \mathcal{I})$ as $\mathcal{U}^K[\mathbf{u}^K, \Phi^K](t, x) = u^K(t, x)$, where

$$\begin{aligned} u_t^K \circ y_t^K &:= u_0^K + \frac{t}{\tau}(u_1^K \circ \phi_1^K - u_0^K), \quad t \in [0, t_{\frac{1}{2}}^K], \\ u_t^K \circ y_t^K &:= u_{K-1}^K + \frac{t - t_{K-1}^K}{\tau}(u_K^K \circ \phi_K^K - u_{K-1}^K), \quad t \in (t_{K-\frac{1}{2}}^K, 1], \end{aligned}$$

and, for $k = 1, \dots, K-1$ and $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$

$$\begin{aligned} &u_t^K \circ y_t^K \tag{5.29} \\ &:= \frac{u_{k-1}^K + u_k^K \circ \phi_k^K}{2} + \frac{t - t_{k-\frac{1}{2}}^K}{\tau}(u_k^K \circ \phi_k^K - u_{k-1}^K) + \frac{(t - t_{k-\frac{1}{2}}^K)^2}{2\tau^2}(u_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K - 2u_k^K \circ \phi_k^K + u_{k-1}^K). \end{aligned}$$

This can be seen as a blending between the ‘‘half-way images’’, $\frac{u_{k-1}^K + u_k^K \circ \phi_k^K}{2}$ and $\frac{(u_k^K + u_{k+1}^K \circ \phi_{k+1}^K) \circ \phi_k^K}{2}$ along the incremental transport path $(y_t^K)_{t \in [0, 1]}$, which is depicted on Figure 5.1 right. It is well defined if the spatial inverse of y_t^K exists. The existence of this object, which we denote by x_t^K , follows from [Cia88, Chapter 5] if $\det(Dy_t^K) > 0$, which is guaranteed if $\|Dy_t^K - \mathbb{1}\|_{C^0(\Omega)} < c$ for a small enough constant $c > 0$, the latter being satisfied in Section 5.4 (see (5.39)).

The discrete velocity field $v^K \in L^2((0, 1), \mathcal{V})$ corresponding to the discrete transport path $(y_t^K)_{t \in [0, 1]}$ is given by

$$v_t^K := \frac{1}{\tau}(\phi_1^K - \mathbb{1}), \quad v_t^K := \frac{1}{\tau}(\phi_K^K - \mathbb{1}),$$

for $t \in [0, t_{\frac{1}{2}}^K]$ and $t \in (t_{K-\frac{1}{2}}^K, 1]$, respectively, and

$$v_t^K := \frac{1}{\tau}(\phi_k^K - \mathbb{1}) + \frac{t - t_{k-\frac{1}{2}}^K}{\tau^2}(\phi_{k+1}^K \circ \phi_k^K - 2\phi_k^K + \mathbb{1}),$$

for $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$ with $k = 1, \dots, K-1$. The corresponding discrete acceleration field $a^K \in L^2((0, 1), \mathcal{V})$ is given by $a_t^K = a_t^K := 0$ for $t \in [0, t_{\frac{1}{2}}^K]$ and $t \in (t_{K-\frac{1}{2}}^K, 1]$, respectively, and, for $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$ with $k = 1, \dots, K-1$, by

$$a_t^K := \frac{1}{\tau^2}(\phi_{k+1}^K \circ \phi_k^K - 2\phi_k^K + \mathbb{1}).$$




Figure 5.1: Left: Schematic drawing of the Hermite interpolation $y_k^K(x)$ (blue) on the time interval $[(k - \frac{1}{2})/K, (k + \frac{1}{2})/K]$ together with the discrete acceleration $a_k^K(x)$ (red). Right: Image extension $\mathcal{U}^K[\mathbf{u}^K, \Phi^K](\cdot, x)$ along a path $(y_t^K(x))_{t \in [0,1]}$, plotted against time. Dots represent the values u_k^K , $k = 0, \dots, K$, and crosses the “half-way” values $\frac{1}{2}(u_k^K + u_{k-1}^K)$, $k = 1, \dots, K$, along the discrete transport path.

We define the velocity and the acceleration along the incremental transport path by

$$\tilde{v}_t^K := v_t^K \circ x_t^K, \quad \tilde{a}_t^K := a_t^K \circ x_t^K. \quad (5.30)$$

Now, the *discrete flow* given as the map $(t, x) \mapsto \psi_t^K(x)$ is defined recursively by

$$\psi_t^K := y_t^K, \quad t \in [0, t_{\frac{1}{2}}^K], \quad \psi_t^K := y_t^K \circ \psi_{t_{k-\frac{1}{2}}^K}^K, \quad t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K], \quad \psi_t^K := y_t^K \circ \psi_{t_{K-\frac{1}{2}}^K}^K, \quad t \in (t_{K-\frac{1}{2}}^K, t_K^K]. \quad (5.31)$$

Following the same arguments as in Proposition 4.3.1, one shows that (5.31) is well-defined in the sense of equations (5.11) – (5.12), *i.e.*, for every $x \in \Omega$ and $t \in [0, 1]$ we have

$$\dot{\psi}_t^K = \tilde{v}_t^K \circ \psi_t^K, \quad \psi_0^K(x) = x \quad (5.32)$$

$$\ddot{\psi}_t^K = \tilde{a}_t^K \circ \psi_t^K. \quad (5.33)$$

Based on this, the *first order scalar weak material derivative* of $(u_t^K)_{t \in [0,1]}$ can be defined as the absolute value of the material derivative along the paths $t \mapsto \psi_t^K(x)$ with

$$z_t^K \circ y_t^K := \frac{1}{\tau} |u_1^K \circ \phi_1^K - u_0^K|, \quad t \in [0, t_{\frac{1}{2}}^K], \quad z_t^K \circ y_t^K := \frac{1}{\tau} |u_K^K \circ \phi_K^K - u_{K-1}^K|, \quad t \in (t_{K-\frac{1}{2}}^K, 1],$$

and, for $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$, $k = 1, \dots, K-1$,

$$z_t^K \circ y_t^K := \left| \frac{1}{\tau} (u_k^K \circ \phi_k^K - u_{k-1}^K) + \frac{t - t_{k-\frac{1}{2}}^K}{\tau^2} (u_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K - 2u_k^K \circ \phi_k^K + u_{k-1}^K) \right|. \quad (5.34)$$

For the *second order scalar weak material derivative*, we have, for $t \in [t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$, $k = 1, \dots, K-1$,

$$w_t^K \circ y_t^K := \frac{1}{\tau^2} |u_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K - 2u_k^K \circ \phi_k^K + u_{k-1}^K|, \quad (5.35)$$

and $w_t^K := 0$ elsewhere, which is the absolute value of the second time derivative of $u^K = \mathcal{U}^K[\mathbf{u}^K, \Phi^K]$ along the path $t \mapsto \psi_t^K(x)$. Indeed, as in Proposition 4.3.1, one verifies that $(z_t^K)_{t \in [0,1]}$ and $(w_t^K)_{t \in [0,1]}$ are admissible in the sense of equations (5.13) and (5.14), *i.e.*,

$$|u_t^K \circ \psi_t^K - u_s^K \circ \psi_s^K| \leq \int_s^t z_r^K \circ \psi_r^K \, dr, \quad \forall s \leq t \in [0, 1],$$

$$\left| u_{t+\tau}^K \circ \psi_{t+\tau}^K - u_t^K \circ \psi_t^K - u_{s+\tau}^K \circ \psi_{s+\tau}^K + u_s^K \circ \psi_s^K \right| \leq \int_0^\tau \int_s^{t+\tau} w_{r+l}^K \circ \psi_{r+l}^K \, dl \, dr, \quad \forall \tau, s + \tau \leq t + \tau \in [0, 1].$$

Let us finally notice that these quantities, together with the image extension operator (5.29), can be defined in an explicit manner by using the inverse transport path $(x_t^K)_{t \in [0,1]}$, as done in Section 4.3.

Remark 5.3.1. For the periodic boundary conditions (*cf.* Remark 5.2.2), we use the cubic interpolation definition on $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$ for $k = 1, \dots, K$, with the convention $K \hat{=} 0, K+1 \hat{=} 1$ and $(t_{K-\frac{1}{2}}^K, t_{K+\frac{1}{2}}^K] = (t_{K-\frac{1}{2}}^K, 1] \cup [0, t_{\frac{1}{2}}^K]$.

Finally, we define the *extension of the energy* $\mathbf{F}^{\sigma, K}$ to a functional $\mathcal{F}^{\sigma, K}$ by

$$\mathcal{F}^{\sigma, K}[u] := \inf_{\mathbf{u}^K \in \mathcal{I}^{K+1}} \inf_{\Phi^K \in \mathcal{D}^K} \{ \mathbf{F}^{\sigma, K, D}[\mathbf{u}^K, \Phi^K] : \mathcal{U}^K[\mathbf{u}^K, \Phi^K] = u \}, \quad (5.36)$$

if there exists such \mathbf{u}^K, Φ^K and $+\infty$, else. Note that, unlike for the extended energy in Section 4.3 (*cf.* (4.36)), we have more candidates for \mathbf{u}^K , so we have to prove the existence of the optimal one.

Lemma 5.3.2. *If for a fixed image path $u \in L^2([0, 1], \mathcal{I})$ there exist a vector of images $\mathbf{u}^K \in \mathcal{I}^{K+1}$ and a vector of deformations $\Phi^K \in \mathcal{D}^K$ such that $\mathcal{U}^K[\mathbf{u}^K, \Phi^K] = u$, the infimum in (5.36) is attained for some $(\mathbf{u}^K, \Phi^K) \in \mathcal{I}^{K+1} \times \mathcal{D}^K$.*

Proof. The existence of an optimal vector of deformations for a fixed vector of images can be proved analogously to Lemma 4.3.2. Thus, let us have a sequence $\{\mathbf{u}_j^K, \Phi_j^K\}_{j \in \mathbb{N}} \subset \mathcal{I}^{K+1} \times \mathcal{D}^K$ such that for every $j \in \mathbb{N}$ we have $\Phi_j^K \in \operatorname{argmin}\{\mathbf{F}^{\sigma, K, D}[\mathbf{u}_j^K, \Phi_j^K] : \mathcal{U}^K[\mathbf{u}_j^K, \Phi_j^K] = u\}$. We have that $\{\mathbf{F}^{\sigma, K, D}[\mathbf{u}_j^K, \Phi_j^K]\}_{j \in \mathbb{N}}$ is bounded by the energy of the pair from the assumption of the lemma. Then, by Lemma 5.2.3, we can prove the uniform boundedness of $\{\Phi_j^K\}_{j \in \mathbb{N}}$ in $H^m(\Omega, \Omega)^K$, and as in (5.26), we have the uniform boundedness of $\{\mathbf{u}_j^K\}_{j \in \mathbb{N}}$ in \mathcal{I}^{K+1} . This implies the existence of a (not relabeled) subsequence such that $\Phi_j^K \rightharpoonup \Phi^K$ in $H^m(\Omega, \Omega)^K$ and strongly in $C^{1, \alpha}(\overline{\Omega}, \overline{\Omega})$, and $\mathbf{u}_j^K \rightharpoonup \mathbf{u}^K$ in \mathcal{I} . The optimality of the energy for the pair \mathbf{u}^K, Φ^K is shown as in the proof of Theorem 5.2.7. Finally, the closedness of the image extension condition follows from the fact that $\Phi_j^K \rightarrow \Phi$ in $C^{1, \alpha}(\overline{\Omega}, \overline{\Omega})$ implies $x_j^K(t, \cdot) \rightarrow x^K(t, \cdot)$ in the same space, uniformly in t (see the proof of Lemma 4.3.2 for more details). Then, by Remarks 4.2.4 and 4.2.8, we have

$$\lim_{j \rightarrow \infty} \mathcal{U}^K[\mathbf{u}_j^K, \Phi_j^K](t, \cdot) \rightharpoonup \mathcal{U}^K[\mathbf{u}^K, \Phi^K](t, \cdot) \text{ in } \mathcal{I}, \forall t \in [0, 1],$$

from where we conclude $\mathcal{U}^K[\mathbf{u}^K, \Phi^K] = u$. □

5.4 Convergence of Discrete Splines

In this section, we study the Mosco-convergence (see Definition 2.4.1), as $K \rightarrow \infty$, of $\{\mathcal{F}^{\sigma, K}\}_{K \in \mathbb{N}}$, the extensions of the time discrete regularized spline energies, to the time continuous regularized spline energy \mathcal{F}^σ . As a corollary, we also obtain the convergence of the corresponding minimizers: the time discrete regularized spline interpolations converge to a time continuous regularized spline interpolation.

Theorem 5.4.1 (Mosco-convergence of the discrete spline energies). *Let $\sigma > 0$. Then, the time discrete spline energies $\{\mathcal{F}^{\sigma, K}\}_{K \in \mathbb{N}}$ converge to \mathcal{F}^σ in the sense of Mosco in the topology $L^2((0, 1), \mathcal{I})$ for $K \rightarrow \infty$. In explicit*

- (i) *for every sequence $\{u^K\}_{K \in \mathbb{N}} \subset L^2((0, 1), \mathcal{I})$ which converges weakly to $u \in L^2((0, 1), \mathcal{I})$ as $K \rightarrow \infty$, it holds $\liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma, K}[u^K] \geq \mathcal{F}^\sigma[u]$ (“lim inf-inequality”),*
- (ii) *for every $u \in L^2((0, 1), \mathcal{I})$, there exists a sequence $\{u^K\}_{K \in \mathbb{N}}$ such that $u^K \rightarrow u$ in $L^2((0, 1), \mathcal{I})$ as $K \rightarrow \infty$ and $\limsup_{K \rightarrow \infty} \mathcal{F}^{\sigma, K}[u^K] \leq \mathcal{F}^\sigma[u]$ (“the existence of a recovery sequence and lim sup-inequality”).*

Remark 5.4.2. The above result holds for any choice of W_D satisfying assumption (W1)–(W3) from Section 4.2, though we restrict ourselves to the special case $W_D = |A^{\text{sym}} - \mathbb{1}|^2$.

Remark 5.4.3. In the case of periodic boundary conditions the same applies in the topology $L^2(S^1, \mathcal{I})$. The proof requires minor alterations implied by the already indicated changes in the energy and interpolation (see Remarks 5.2.2 and 5.3.1).

Proof. We will prove the conditions separately. The parts of the proofs dealing with the first order terms (the velocity of the flow and material derivative) will follow along the lines of the proof of Theorem 4.4.1. We will split the proofs into steps analogous to the steps of the mentioned proof.

Proof of the lim inf-inequality

(i): *Construction of the flow and image intensity quantities.*

Suppose we have a sequence $\{u^K\}_{K \in \mathbb{N}} \subset L^2((0, 1), \mathcal{I})$ such that $u^K \rightharpoonup u$ in that space, as $K \rightarrow \infty$. To avoid the trivial case of an infinite limit, we suppose that $\mathcal{F}^{\sigma, K}[u^K] < \overline{\mathcal{F}} < \infty$. Then, by definition (5.36) and Lemma 5.3.2, for every K large enough, an optimal vector of images $\mathbf{u}^K \in \mathcal{I}^{K+1}$ and a corresponding optimal vector of deformations $\Phi^K \in \mathcal{D}^K$ exist, such that

$$u^K = \mathcal{U}^K[\mathbf{u}^K, \Phi^K], \quad \mathcal{F}^{\sigma, K}[u^K] = \mathbf{F}^{\sigma, K, D}[\mathbf{u}^K, \Phi^K].$$

For the vectors of images \mathbf{u}^K and the vectors of deformations Φ^K we define the discrete velocity and acceleration, v_k^K and a_k^K , respectively, and the discrete material derivative of the first and the second order, \hat{z}_k^K and \hat{w}_k^K , respectively, as in Section 5.2. Using Lemma 5.2.3 and Remark 5.2.4, we obtain

$$\max_{k=1, \dots, K} \|\phi_k^K - \mathbf{1}\|_{C^{1, \alpha}(\overline{\Omega})} \leq CK^{-\frac{1}{2}}, \quad (5.37)$$

$$\max_{k=1, \dots, K-1} \|a_k^K\|_{C^{1, \alpha}(\overline{\Omega})} \leq CK^{\frac{1}{2}}, \quad (5.38)$$

and hence, by the definition of the discrete incremental transport path (5.28), we have

$$\begin{aligned} \max_{t \in [0, 1]} \|y_t^K - \mathbf{1}\|_{C^{1, \alpha}(\overline{\Omega})} &\leq \frac{3}{2} \max_{k=1, \dots, K} \|\phi_k^K - \mathbf{1}\|_{C^{1, \alpha}(\overline{\Omega})} + \frac{1}{2K^2} \max_{k=1, \dots, K-1} \|a_k^K\|_{C^{1, \alpha}(\overline{\Omega})} \\ &\leq C(K^{-\frac{1}{2}} + K^{-\frac{3}{2}}). \end{aligned} \quad (5.39)$$

This implies that y_t^K converges to the identity in $C^{1, \alpha}(\overline{\Omega}, \overline{\Omega})$, uniformly in t , and that y_k^K , defined in the previous section, is invertible for any $k = 0, \dots, K$ and every large enough K . Thus, we are able to define all the temporal extended quantities introduced in Section 5.3.

(ii): *Proof of the actual inequality.*

Using the definition of the discrete acceleration \hat{w}_k^K (cf. (5.17)) and its time extension w_t^K (cf. (5.35)), together with (5.39) and the locally Lipschitz property of the determinant function, we get

$$\begin{aligned} \lim_{K \rightarrow \infty} \int_0^1 \int_{\Omega} (w_t^K)^2 dx dt &= \lim_{K \rightarrow \infty} \sum_{k=1}^{K-1} \int_{\Omega} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} |\hat{w}_k^K \circ x_{k,t}^K|^2 dt dx \\ &= \lim_{K \rightarrow \infty} \sum_{k=1}^{K-1} \int_{\Omega} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} |\hat{w}_k^K|^2 \det(Dy_{k,t}^K) dt dx \\ &= \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} |\hat{w}_k^K|^2 dx. \end{aligned}$$

Using the same ideas, together with $|z_t^K - \hat{z}_k^K \circ x_{k,t}^K| \leq \frac{1}{K} |\hat{w}_k^K \circ x_{k,t}^K|$ (cf. (5.34)) on every subinterval, we have

$$\begin{aligned} \lim_{K \rightarrow \infty} \int_0^1 \int_{\Omega} (z_t^K)^2 dx dt &= \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \int_{\Omega} |\hat{z}_k^K|^2 dx + \frac{1}{K^2} \sum_{k=1}^{K-1} \int_{\Omega} |\hat{w}_k^K|^2 dx \\ &= \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \int_{\Omega} |\hat{z}_k^K|^2 dx. \end{aligned}$$

This implies the uniform boundedness of $\{z^K\}_{K \in \mathbb{N}}$ and $\{w^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1) \times \Omega)$, and, by reflexivity of the space, the existence of weakly convergent subsequences (with the same labeling) to z and w , respectively. Then, by the weak lower semicontinuity, we have

$$\|z\|_{L^2((0, 1) \times \Omega)}^2 \leq \liminf_{K \rightarrow \infty} \|z^K\|_{L^2((0, 1) \times \Omega)}^2, \quad \|w\|_{L^2((0, 1) \times \Omega)}^2 \leq \liminf_{K \rightarrow \infty} \|w^K\|_{L^2((0, 1) \times \Omega)}^2. \quad (5.40)$$

Using Korn's inequality 2.1.6 and Poincaré's 2.1.4 inequality, we get

$$\begin{aligned} \int_0^1 \int_{\Omega} |a_t^K|^2 dx dt &= \sum_{k=1}^{K-1} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} \int_{\Omega} |a_k^K|^2 dx dt \leq \frac{C}{K} \sum_{k=1}^{K-1} \int_{\Omega} W_A(Da_k^K) dx \leq C\bar{\mathcal{F}}, \\ \int_0^1 \int_{\Omega} |D^m a_t^K|^2 dx dt &= \sum_{k=1}^{K-1} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} \int_{\Omega} |D^m a_k^K|^2 dx dt \leq \sum_{k=1}^{K-1} \frac{1}{K} \int_{\Omega} |D^m a_k^K|^2 dx \leq C\bar{\mathcal{F}}. \end{aligned}$$

The analogous estimates are obtained for v^K and $D^m v^K$, with the additional use of $|v_t^K - v_k^K| \leq \frac{1}{K}|a_k^K|$ on every subinterval (see also (4.42)). Hence, we have that $\{v^K\}_{K \in \mathbb{N}}$ and $\{a^K\}_{K \in \mathbb{N}}$ are uniformly bounded in $L^2((0, 1), \mathcal{V})$, and they have the corresponding weak limits v and a in that space.

We compute the Taylor expansion of $W_A((t - t_{k-\frac{1}{2}}^K)^2 Da_k^K(t, \cdot))$ around $t_{k-\frac{1}{2}}^K$, evaluated at $t = t_{k+\frac{1}{2}}^K$, to get

$$\frac{1}{K^4} W_A(Da_k^K) = \frac{1}{2K^4} D^2 W_A(\mathbf{0})(Da_k^K, Da_k^K) + r_{a,k}^K = \frac{1}{K^4} \text{tr}(\varepsilon[a_k^K]^2) + r_{a,k}^K. \quad (5.41)$$

For the remainder term, we have $r_{a,k}^K = \mathcal{O}(K^{-6}|Da_k^K|^3)$, and by using Lemma 5.2.3 and (5.38), we obtain

$$\sum_{k=1}^{K-1} K^3 \int_{\Omega} r_{a,k}^K dx \leq \frac{C}{K^2} \max_{k=1, \dots, K-1} \|a_k^K\|_{C^1(\bar{\Omega})} \frac{1}{K} \sum_{k=1}^{K-1} \|a_k^K\|_{H^m(\Omega)}^2 \leq CK^{-\frac{3}{2}}\bar{\mathcal{F}}. \quad (5.42)$$

Then, we use weak lower semicontinuity of the energy to write

$$\begin{aligned} \liminf_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} W_A(Da_k^K) + \gamma |D^m a_k^K|^2 dx &= \liminf_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} \text{tr}(\varepsilon[a_k^K]^2) + \gamma |D^m a_k^K|^2 dx \\ &= \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} \text{tr}(\varepsilon[a_t^K]^2) + \gamma |D^m a_t^K|^2 dx dt \\ &\geq \int_0^1 \int_{\Omega} \text{tr}(\varepsilon[a]^2) + \gamma |D^m a|^2 dx dt. \end{aligned} \quad (5.43)$$

Analogously, the Taylor expansions of $W(\mathbf{1} + (t - t_{k-\frac{1}{2}}^K)Dv_k^K(t, \cdot))$ around $t_{k-\frac{1}{2}}^K$, evaluated at $t_{k+\frac{1}{2}}^K$, for $k = 1, \dots, K-1$, and correspondingly for the intervals of size $\frac{1}{2K}$, give (see also (4.48)–(4.50))

$$\begin{aligned} \liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} W_D(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 dx &= \liminf_{K \rightarrow \infty} \int_{\Omega} \frac{1}{2K} (\text{tr}(\varepsilon[v_1^K]^2) + \gamma |D^m v_1^K|^2) + \frac{1}{K} \sum_{k=1}^{K-1} \text{tr}(\varepsilon[v_k^K]^2) + \gamma |D^m v_k^K|^2 \\ &\quad + \frac{1}{2K} (\text{tr}(\varepsilon[v_1^K]^2) + \gamma |D^m v_1^K|^2) dx \\ &= \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} \text{tr}(\varepsilon[v_t^K]^2) + \gamma |D^m v_t^K|^2 + \frac{1}{K} (\text{tr}(\varepsilon[v_t^K]^2) + \gamma |D^m v_t^K|^2) dx dt \\ &\geq \int_0^1 \int_{\Omega} \text{tr}(\varepsilon[v]^2) + \gamma |D^m v|^2 dx dt, \end{aligned} \quad (5.44)$$

where we once again used that $|v_t^K - v_k^K| \leq \frac{1}{K}|a_k^K|$ on every subinterval.

Altogether, (5.40), (5.43), and (5.44) give

$$\liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma, K}[u^K] = \liminf_{K \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\mathbf{u}^K, \Phi^K] \geq \mathcal{F}^{\sigma}[u].$$

(iii): *Proving the admissibility of the limit of the discrete velocities and acceleration fields, and the discrete material derivatives and accelerations, respectively.*

In this step, we show that the limit objects v, a, z, w are indeed corresponding quantities for the image curve u , i.e., $(v, a, z, w) \in \mathcal{C}[u]$ and they satisfy relations (5.11)–(5.14). This was already shown in Section 4.4 for the first order terms v and z , and we repeat just the most relevant of those arguments here.

First, let us observe that, from (5.39) and smoothness of the inversion, we have (cf. (4.28))

$$\|x_t^K\|_{C^{1,\alpha}(\bar{\Omega})} \leq C(1 + \max_{k=1,\dots,K} \|\phi_k^K - \mathbb{1}\|_{C^{1,\alpha}(\bar{\Omega})} + \max_{k=1,\dots,K-1} \frac{1}{K^2} \|a_k^K\|_{C^{1,\alpha}(\bar{\Omega})}),$$

while, by [Fio16, Propositions 1.2.4 and 1.2.7], we have (cf. (4.33))

$$\begin{aligned} \|\tilde{v}_t^K\|_{C^{1,\alpha}(\bar{\Omega})} &\leq C\|v_t^K\|_{C^{1,\alpha}(\bar{\Omega})}(1 + \|x_t^K\|_{C^{1,\alpha}(\bar{\Omega})} + \|x_t^K\|_{C^{1,\alpha}(\bar{\Omega})}^{\alpha+1}), \\ \|\tilde{a}_t^K\|_{C^{1,\alpha}(\bar{\Omega})} &\leq C\|a_t^K\|_{C^{1,\alpha}(\bar{\Omega})}(1 + \|x_t^K\|_{C^{1,\alpha}(\bar{\Omega})} + \|x_t^K\|_{C^{1,\alpha}(\bar{\Omega})}^{\alpha+1}), \end{aligned} \quad (5.45)$$

where \tilde{v}_t^K and \tilde{a}_t^K are defined by (5.30). In particular, this implies the uniform boundedness of $\{\tilde{v}_t^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1), C^{1,\alpha}(\bar{\Omega}, \mathbb{R}^n))$. Then, applying Remarks 3.2.2 and 3.2.3 to (5.32), we have that $\{\psi^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $H^1 \subset C^{0, \frac{1}{2}}([0, 1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$. Furthermore, by the compact embedding of Hölder spaces (cf. Corollary 2.1.10), we have, for all $\min(\frac{1}{2}, \alpha) > \beta > 0$, that $\psi^K \rightarrow \psi$ in $C^{0,\beta}([0, 1], C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$.

To show that ψ is indeed the solution corresponding to v , we consider ψ^{v^K} , the solution corresponding to v^K . By the weak continuity of the solution operator mapping velocities to the flows (Theorem 3.2.1), we have $\psi^{v^K} \rightarrow \psi^v$ in $C^0([0, 1] \times \bar{\Omega})$. Furthermore, by Remark 3.2.2, we have the Lipschitz continuity of the solution operator, and using the spatial Lipschitz property of v_k^K together with (5.37), we have $\psi^K - \psi^{v^K} \rightarrow 0$ in $C^0([0, 1] \times \bar{\Omega})$, finally confirming that $\psi = \psi^v$. For more details, see the discussion after (4.43).

To show that the equation $\ddot{\psi}_t = a_t \circ \psi_t$ is satisfied, first observe that (5.33), i.e., $\ddot{\psi}_t^K = \tilde{a}_t^K \circ \psi_t^K$, ensures the uniform boundedness of $\{\ddot{\psi}_t^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1), C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$. To this end, we used the uniform boundedness of $\{\tilde{a}_t^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1), C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$, following from (5.45), and the uniform boundedness of $\{\psi_t^K\}_{K \in \mathbb{N}}$ in $C^0([0, 1], C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$, together with the estimate on composition of Hölder functions [Fio16, Propositions 1.2.4 and 1.2.7]. Together with the previous paragraph, we conclude that $\{\psi_t^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $H^2((0, 1), C^{1,\alpha}(\bar{\Omega}, \bar{\Omega}))$, and converges weakly to ψ in $H^2((0, 1), C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$, and strongly in $C^{1,\beta}([0, 1], C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$, for $\min(\frac{1}{2}, \alpha) > \beta > 0$. In particular, we have that $\psi \in H^2((0, 1), C^{1,\beta}(\bar{\Omega}, \bar{\Omega}))$ and $\tilde{a}^K \circ \psi^K = \ddot{\psi}^K \rightarrow \ddot{\psi}$ in $L^2((0, 1) \times \bar{\Omega})$. Hence, it suffices to verify $\tilde{a}^K \circ \psi^K \rightarrow a \circ \psi$ in $L^2((0, 1) \times \bar{\Omega})$. Since, by Remark 4.2.8, we already have $a^K \circ \psi^K \rightarrow a \circ \psi$ in $L^2((0, 1) \times \bar{\Omega})$, we conclude the proof by checking that $\tilde{a}^K \circ \psi^K - a^K \circ \psi^K \rightarrow 0$ in $L^2((0, 1), C^0(\bar{\Omega}, \bar{\Omega}))$. Indeed,

$$\begin{aligned} &\|\tilde{a}^K \circ \psi^K - a^K \circ \psi^K\|_{L^2((0,1), C^0(\bar{\Omega}))}^2 \\ &= \sum_{k=1}^{K-1} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} \|\tilde{a}_t^K \circ \psi_t^K - a_t^K \circ \psi_t^K\|_{C^0(\bar{\Omega})}^2 dt \\ &\leq C \sum_{k=1}^{K-1} \int_{t_{k-1}^K}^{t_k^K} \|a_k^K(t, \cdot)\|_{C^1(\bar{\Omega})}^2 \|y_k^K(t, \cdot) - \mathbb{1}\|_{C^0(\bar{\Omega})}^2 dt \\ &\leq C \|a^K\|_{L^2((0,1), C^1(\bar{\Omega}))}^2 \left(\max_{k=1,\dots,K} \|\phi_k^K - \mathbb{1}\|_{C^0(\bar{\Omega})}^2 + \max_{k=1,\dots,K-1} K^{-2} \|a_k^K\|_{C^0(\bar{\Omega})}^2 \right) \\ &\leq CK^{-1} \|a^K\|_{L^2((0,1), C^1(\bar{\Omega}))}^2, \end{aligned}$$

where we used the Lipschitz property of a_t^K , the transformation formula, and finally, (5.37) and (5.38).

In order to show that z and w are indeed scalar weak material derivatives of u , first observe that, from $u^K \rightarrow u$ and $\psi^K \rightarrow \psi$, by Remark 4.2.8, we obtain $u^K \circ \psi^K \rightarrow u \circ \psi$ in $L^2((0, 1) \times \bar{\Omega})$, and analogously $z^K \circ \psi^K \rightarrow z \circ \psi$ and $w^K \circ \psi^K \rightarrow w \circ \psi$. Next, note that, for $s, t \in [0, 1]$, we have

$$\|u_t^K \circ \psi_t^K - u_s^K \circ \psi_s^K\|_{L^2}^2 \leq |t - s| \left| \int_s^t \int_{\Omega} (z_r^K \circ \psi_r^K)^2 dx dr \right| \leq C|t - s| \|z^K\|_{L^2((0,1) \times \bar{\Omega})}^2 \leq C|t - s|,$$

where we used the Cauchy–Schwarz inequality, the transformation formula, and, finally, the uniform boundedness of $\{(\psi^K)^{-1}\}_{K \in \mathbb{N}}$. Thus, $\{u^K \circ \psi^K\}_{K \in \mathbb{N}}$ is in $C_\omega^0([0, 1], \mathcal{I})$, where $\omega(t) = Ct^{\frac{1}{2}}$ for some constant C depending on the corresponding norms of v and z (cf. Remark 5.1.4). Then, by the weak closedness of this set following from Lemma 4.4.2, we obtain $u \circ \psi \in C_\omega^0([0, 1], \mathcal{I})$. Further using the properties of this set, we have that, for every $\tilde{\Omega} \subset \Omega$, the functional $b \mapsto \int_{\tilde{\Omega}} |b_{t+\tau}(x) - b_t(x) - b_{s+\tau}(x) + b_s(x)| dx$ is continuous on $C_\omega^0([0, 1], \mathcal{I})$, and convex, which implies weak lower semicontinuity. This finally gives

$$\begin{aligned} & \int_{\tilde{\Omega}} |u_{t+\tau} \circ \psi_{t+\tau} - u_t \circ \psi_t - u_{s+\tau} \circ \psi_{s+\tau} + u_s \circ \psi_s| dx \\ & \leq \liminf_{K \rightarrow \infty} \int_{\tilde{\Omega}} |u_{t+\tau}^K \circ \psi_{t+\tau}^K - u_t^K \circ \psi_t^K - u_{s+\tau}^K \circ \psi_{s+\tau}^K + u_s^K \circ \psi_s^K| dx \\ & \leq \liminf_{K \rightarrow \infty} \int_{\tilde{\Omega}} \int_0^\tau \int_s^t w_{r+l}^K \circ \psi_{r+l}^K dr dl dx \\ & = \int_{\tilde{\Omega}} \int_0^\tau \int_s^t w_{r+l} \circ \psi_{r+l} dr dl dx. \end{aligned}$$

Since this holds for any $\tilde{\Omega} \subset \Omega$, we have that w is the second (scalar) weak material derivative for u . The proof of z being the first weak material derivative is analogous (cf. (4.46)).

This finally finishes the proof of the lim inf-inequality.

The following proposition is a corollary of the preceding proof and it represents an analogue of Proposition 4.4.3, i.e., it shows that the infimum in (5.10) is actually attained.

Proposition 5.4.4. *For $u \in L^2([0, 1], \mathcal{I})$ with $\mathcal{F}^\sigma[u] < \infty$ there exists an optimal tuple $(v, a, z, w) \in \mathcal{C}[u]$ such that*

$$\tilde{\mathcal{F}}^\sigma[u] = \int_0^1 L[a, a] + \frac{1}{\delta} w^2 + \sigma \left(L[v, v] + \frac{1}{\delta} z^2 \right) dx dt.$$

Proof. The functional \mathcal{F}^σ is coercive by Korn's inequality 2.1.6 and Gagliardo-Nirenberg interpolation estimate 2.1.5, and it is clearly weakly lower semicontinuous. Since $\mathcal{C}[u]$ is a subset of a reflexive Banach space, then we just have to show the weak closedness of the set. This is verified as above. For more details, see the proof of Proposition 4.4.3. \square

Proof of the existence of a recovery sequence and the lim sup-inequality

(i) *Construction of a recovery sequence.*

Consider an image curve $u \in L^2([0, 1], \mathcal{I})$ with $\mathcal{F}^\sigma[u] < \infty$. Then, the previous proposition guarantees the existence of an associated optimal velocity field, an acceleration field, and the first and second order weak (scalar) material derivatives, denoted by $(v, a, z, w) \in \mathcal{C}[u]$, respectively, such that

$$\mathcal{F}^\sigma[u] = \int_0^1 \int_\Omega L[a, a] + \frac{1}{\delta} w^2 + \sigma \left(L[v, v] + \frac{1}{\delta} z^2 \right) dx dt.$$

We define

$$\phi_k^K := \psi_{t_{k-1}^K, t_k^K} = \psi_{t_k^K} \circ \psi_{t_{k-1}^K}^{-1}, \quad k = 1, \dots, K, \quad (5.46)$$

where ψ is the flow associated with velocity v and $\psi_0 = \mathbb{1}$ (cf. (5.11)). We have

$$\begin{aligned} \max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^1(\bar{\Omega})} & \leq \sup_{|t-s| \leq K^{-1}} \|\psi_{s,t} - \mathbb{1}\|_{C^1(\bar{\Omega})} \\ & \leq \sup_{|t-s| \leq K^{-1}} C \left| \int_s^t \|v_r \circ \psi_r\|_{H^m(\Omega)} dr \right| \\ & \leq \sup_{|t-s| \leq K^{-1}} C \left| \int_s^t \|v_r\|_{H^m(\Omega)} dr \right| \\ & \leq CK^{-\frac{1}{2}} \sup_{|t-s| \leq K^{-1}} \left| \int_s^t \|v_r\|_{H^m(\Omega)}^2 dr \right|^{\frac{1}{2}} \leq CK^{-\frac{1}{2}} \sqrt{\mathcal{F}^\sigma[u]}, \quad (5.47) \end{aligned}$$

by Lemma 5.2.3 and Cauchy's inequality. For the second inequality, we used [BV17, Lemma 3.5] which states that

$$\|v_r \circ \psi_r\|_{H^m(\Omega)} \leq C \|v_r\|_{H^m(\Omega)}. \quad (5.48)$$

Thereby, for K large enough, we have $\Phi^K \in \mathcal{D}^K$, and we are in a position to define

$$u^K := \mathcal{U}^K[\mathbf{u}^K, \Phi^K], \quad \mathbf{u}^K := (u_{t_0^K}, \dots, u_{t_K^K}),$$

where the point evaluation in time is possible since $u \in C^1([0, 1], \mathcal{I})$ (cf. Remarks 4.1.4 and 5.1.4).

(ii) *Proof of the actual inequality.*

The arguments for the estimates and the convergence of the first order terms are already represented in Section 4.4. In what follows, we state those for the reference and present more detailed arguments for the second order terms.

First, we are able to relate the discrete second order material derivative $\hat{w}_k^K, k = 1, \dots, K-1$ given by (5.17) and its continuous counterpart w , via

$$\begin{aligned} \int_{\Omega} |\hat{w}_k^K|^2 dx &= K^4 \int_{\Omega} |u_{t_{k+1}^K}^K \circ \phi_{t_{k+1}^K}^K \circ \phi_k^K - 2u_k^K \circ \phi_k^K + u_{t_{k-1}^K}^K|^2 dx \\ &= K^4 \int_{\Omega} |u_{t_{k+1}^K}^K \circ \psi_{t_{k-1}^K, t_{k+1}^K} - 2u_{t_k^K} \circ \psi_{t_{k-1}^K, t_k^K} + u_{t_{k-1}^K}^K|^2 dx \\ &= K^4 \int_{\Omega} |u_{t_{k+1}^K}^K \circ \psi_{t_{k+1}^K} - 2u_{t_k^K} \circ \psi_{t_k^K} + u_{t_{k-1}^K}^K \circ \psi_{t_{k-1}^K}|^2 \det(D\psi_{t_{k-1}^K}) dx \\ &\leq K^4 \int_{\Omega} \left(\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} w_{r+s} \circ \psi_{r+s} dr ds \right)^2 \det(D\psi_{t_{k-1}^K}) dx \\ &\leq K^2 \int_0^{\frac{1}{K}} \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} w_{r+s}^2 \det(D\psi_{r+s, t_{k-1}^K}) dx ds dr \\ &\leq K^2 (1 + CK^{-\frac{1}{2}}) \int_0^{\frac{1}{K}} \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} w_{r+s}^2 dx ds dr. \end{aligned}$$

Here, we first used (5.46) and the transformation formula for the first and the second equality, respectively, then the definition of the second order material derivative (5.8) for the fourth inequality, and, finally, the Cauchy-Schwarz inequality and (5.47) in the last two estimates. Summing the above expressions over $k = 1, \dots, K-1$, we obtain

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} |\hat{w}_k^K|^2 dx &\leq K(1 + CK^{-\frac{1}{2}}) \int_0^{\frac{1}{K}} \int_0^{1-\frac{1}{K}} \int_{\Omega} w_{r+s}^2 dx ds dr \\ &\leq K(1 + CK^{-\frac{1}{2}}) \int_0^{\frac{1}{K}} \int_0^1 \int_{\Omega} w_t^2 dx dt dr \\ &\leq (1 + CK^{-\frac{1}{2}}) \int_0^1 \int_{\Omega} w_t^2 dx dt. \end{aligned} \quad (5.49)$$

We analogously showed (cf. (4.59))

$$\frac{1}{K} \sum_{k=1}^K \int_{\Omega} |\hat{z}_k^K|^2 dx = K \sum_{k=1}^K \int_{\Omega} |u_k^K \circ \phi_k^K - u_{t_{k-1}^K}^K|^2 dx \leq (1 + CK^{-\frac{1}{2}}) \int_0^1 \int_{\Omega} z_t^2 dx dt. \quad (5.50)$$

Next, we express the discrete acceleration a_k^K , $k = 1, \dots, K-1$ in terms of its continuous counterpart a :

$$\begin{aligned}
a_k^K &= K^2(\phi_{k+1}^K \circ \phi_k^K - 2\phi_k^K + \mathbf{1}) \\
&= K^2(\psi_{t_{k-1}^K, t_{k+1}^K} - 2\psi_{t_{k-1}^K, t_k^K} + \psi_{t_{k-1}^K, t_{k-1}^K}) \\
&= K^2 \left(\int_{t_k^K}^{t_{k+1}^K} \dot{\psi}_t \circ \psi_{t_{k-1}^K}^{-1} dt - \int_{t_{k-1}^K}^{t_k^K} \dot{\psi}_t \circ \psi_{t_{k-1}^K}^{-1} dt \right) \\
&= K^2 \left(\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \ddot{\psi}_{t+\tau} \circ \psi_{t_{k-1}^K}^{-1} d\tau dt \right) \\
&= K^2 \left(\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau} d\tau dt \right), \tag{5.51}
\end{aligned}$$

where in the second equality we used (5.46), and in the last equality (5.12). Then, using the Cauchy-Schwarz inequality and (5.48), we obtain the following estimate

$$\max_{k=1, \dots, K-1} \|a_k^K\|_{C^1(\bar{\Omega})} \leq CK^{\frac{1}{2}} \sup_{t \in [0,1], 0 < \tau \leq K-1} \left| \int_t^{t+2\tau} \|a_s\|_{H^m(\Omega)}^2 ds \right|^{\frac{1}{2}}. \tag{5.52}$$

The same Taylor expansion arguments as in (5.41) and (5.42) now imply, together with (5.52)

$$\begin{aligned}
\int_{\Omega} W_A(Da_k^K) + \gamma |D^m a_k^K|^2 dx &\leq \int_{\Omega} L[a_k^K, a_k^K] + CK^{-2} |Da_k^K|^3 dx \\
&\leq \int_{\Omega} L[a_k^K, a_k^K] dx + CK^{-\frac{3}{2}} \sqrt{\mathcal{F}^\sigma[u]}. \tag{5.53}
\end{aligned}$$

Applying Jensen's inequality twice on L , and taking into account (5.51), we get

$$\begin{aligned}
\int_{\Omega} L[a_k^K, a_k^K] dx &= \int_{\Omega} K^4 L \left[\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau} d\tau dt, \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau} d\tau dt \right] dx \\
&\leq \int_{\Omega} K^2 \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} L \left[a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau}, a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau} \right] d\tau dt dx.
\end{aligned}$$

We now estimate the summands of L individually. For the first term, we use $|\operatorname{tr}(AB)| \leq |\operatorname{tr}(A)| + |\operatorname{tr}(B - \mathbf{1})|$, (5.47), and the transformation formula, to get

$$\begin{aligned}
&\int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \operatorname{tr} \left(\varepsilon [a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau}]^2 \right) d\tau dt dx \\
&\leq \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \operatorname{tr} \left((\varepsilon [a_{t+\tau}] \circ \psi_{t_{k-1}^K, t+\tau})^2 \right) + \operatorname{tr} \left((\varepsilon [a_{t+\tau}] \circ \psi_{t_{k-1}^K, t+\tau})^2 (\varepsilon [\psi_{t_{k-1}^K, t+\tau}]^2 - \mathbf{1}) \right) d\tau dt dx \\
&\leq \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \operatorname{tr} \left(\varepsilon [a_{t+\tau}]^2 \right) + CK^{-\frac{1}{2}} \|a_{t+\tau}\|_{H^m(\Omega)}^2 d\tau dt dx.
\end{aligned}$$

For the second term, we use (5.48) and the fact that for any $0 \leq \tilde{m} \leq m$ and $f \in H^m(\Omega, \mathbb{R}^n)$, $g \in H^{\tilde{m}}(\Omega, \mathbb{R}^n)$ we have $\|fg\|_{H^{\tilde{m}}(\Omega)} \leq C\|f\|_{H^m(\Omega)}\|g\|_{H^{\tilde{m}}(\Omega)}$ [IKT13, Lemma 2.3], implying

$$\int_{\Omega} |D^m(a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau})| dx \leq \int_{\Omega} |D^{m-1}a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau}| dx + CK^{-\frac{1}{2}} \|a_{t+\tau}\|_{H^m(\Omega)}.$$

Iterating this argument and using the transformation formula, we obtain

$$\int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} |D^m(a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau})|^2 d\tau dt dx \leq \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} |a_{t+\tau}|_{H^m(\Omega)}^2 + CK^{-\frac{1}{2}} \|a_{t+\tau}\|_{H^m(\Omega)}^2 d\tau dt.$$

In a combination with (5.52) and (5.53), we have

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} W_A(Da_k^K) + \gamma |D^m a_k^K|^2 dx &\leq K \int_0^{\frac{1}{K}} \int_0^{1-\frac{1}{K}} \int_{\Omega} L[a_{t+\tau}, a_{t+\tau}] + \mathcal{O}(K^{-\frac{1}{2}}) dx dt d\tau \\ &\leq \int_0^1 \int_{\Omega} L[a_t, a_t] + \mathcal{O}(K^{-\frac{1}{2}}) dx dt. \end{aligned}$$

Altogether, taking into account (5.49), we obtain

$$\mathcal{F}^K[u^K] \leq \mathcal{F}[u] + \mathcal{O}(K^{-\frac{1}{2}}).$$

This readily implies the lim sup-inequality for the pure spline part of the functional $\mathcal{F}^{\sigma, K}$. Analogously to the above, we have (cf. (4.60)–(4.65))

$$K \sum_{k=1}^K \int_{\Omega} W_D(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 dx \leq \int_0^1 \int_{\Omega} L[v_t, v_t] + \mathcal{O}(K^{-\frac{1}{2}}) dx dt,$$

which, together with (5.50), gives

$$\mathcal{E}^K[u^K] \leq \mathcal{E}[u] + \mathcal{O}(K^{-\frac{1}{2}}),$$

finally proving the lim sup-inequality.

(iii) *Proof of the convergence of the recovery sequence.*

As the final step, we are left to show that $u^K \rightarrow u$ in $L^2((0, 1), \mathcal{I})$ as $K \rightarrow \infty$. To this end, we introduce the piecewise constant interpolation

$$\bar{u}_t^K := \begin{cases} u_0^K, & t \in [0, t_{\frac{1}{2}}^K], \\ u_k^K, & t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K], \quad k = 1, \dots, K-1, \\ u_K^K, & t \in [t_{K-\frac{1}{2}}^K, 1]. \end{cases}$$

We will show

$$u^K - \bar{u}^K \rightarrow 0 \text{ in } L^\infty((0, 1), \mathcal{I}). \quad (5.54)$$

To this end, for $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$ with $k = 1, \dots, K-1$, we estimate

$$\begin{aligned} &\|u_t^K - \bar{u}_t^K\|_{\mathcal{I}}^2 \\ &\leq C \left(\| (u_{k-1}^K - u_k^K \circ \phi_k^K) \circ x_t^K \|_{\mathcal{I}}^2 + \| (u_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K - 2u_k^K \circ \phi_k^K + u_{k-1}^K) \circ x_t^K \|_{\mathcal{I}}^2 + \| u_k^K \circ \phi_k^K \circ x_t^K - u_k^K \|_{\mathcal{I}}^2 \right) \\ &\leq C \left(K^{-2} \| \hat{z}_k^K \|_{L^2(\Omega)}^2 \| \det Dy^K \|_{L^\infty((0,1) \times \Omega)} + K^{-4} \| \hat{w}_k^K \|_{L^2(\Omega)}^2 \| \det Dy^K \|_{L^\infty((0,1) \times \Omega)} \right. \\ &\quad \left. + \| u_k^K \circ \phi_k^K \circ x_t^K - u_k^K \|_{\mathcal{I}}^2 \right) \\ &\leq C \left(K^{-1} \| z \|_{L^2((0,1) \times \Omega)}^2 + CK^{-3} \| w \|_{L^2((0,1) \times \Omega)}^2 + \| u_k^K \circ \phi_k^K \circ x_{k,t}^K - u_k^K \|_{\mathcal{I}}^2 \right) \end{aligned} \quad (5.55)$$

For every $t \in (0, 1)$, we can find a sequence $\{k(K)\}_{K \in \mathbb{N}}$ such that $t \in (t_{k(K)-\frac{1}{2}}^K, t_{k(K)+\frac{1}{2}}^K]$, for any K large enough. Following the same procedure as in (4.56), we uniformly approximate the sequence $\{u_{k(K)}^K\}_{K \in \mathbb{N}}$ by smooth functions, and use (5.47) and (5.52), to prove

$$\|u_{k(K)}^K \circ \phi_{k(K)}^K \circ x_t^K - u_{k(K)}^K\|_{\mathcal{I}}^2 \rightarrow 0, \text{ uniformly in } t.$$

Plugging this back into (5.55), we get (5.54). Furthermore, as $\{\bar{u}^K\}_{K \in \mathbb{N}}$ is a sequence of piecewise constant approximations of $u \in C^1([0, 1], \mathcal{I})$, we have that $\bar{u}^K \rightarrow u$ in $L^2([0, 1], \mathcal{I})$ as $K \rightarrow \infty$. Thus, we can finally conclude $u^K \rightarrow u$ in $L^2((0, 1), \mathcal{I})$, as we wanted to show. This finally finishes the proof of Theorem 5.4.1. \square

As a corollary of the previous theorem, we are able to show the existence of the time continuous regularized spline interpolation (*cf.* Definition 5.1.3) as the corresponding limit of the time discrete spline interpolations. To this end, let $J \geq 2$ and $(t_1, \dots, t_J) \subset [0, 1] \cap \mathbb{Q}$ be a sequence of fixed times. Then, for infinitely many $K \in \mathbb{N}$, one can choose $i_j^K := K \cdot t_j \in \mathbb{N}$ for all $j = 1, \dots, J$. Let $(u_j^I)_{j=1, \dots, J} \subseteq \mathcal{I}$ be the set of constraint images at the corresponding constraint times (*cf.* (5.1) and (5.19)).

Theorem 5.4.5 (Convergence of discrete spline interpolations). *For every K that satisfies the above condition, let $u^K \in L^2([0, 1], \mathcal{I})$ be a minimizer of $\mathcal{F}^{\sigma, K}$ among the image curves satisfying $u^K = \mathcal{U}^K[\mathbf{u}^K, \Phi^K]$ with $u_{i_j^K}^K = u_j^I$ for all $j = 1, \dots, J$. Then, a subsequence of $\{u^K\}_{K \in \mathbb{N}}$ converges weakly in $L^2([0, 1], \mathcal{I})$ as $K \rightarrow \infty$ to a minimizer of the continuous spline energy \mathcal{F}^σ . This minimizer satisfies $u_{t_j} = u_j^I$ for all $j = 1, \dots, J$, and the associated sequence of discrete energies converges to the minimal continuous spline energy.*

Proof. For $j = 1, \dots, J$, let $\eta^j : [0, 1] \rightarrow \mathbb{R}$ be smooth functions with $\eta_{t_j}^j = \delta_{ij}$. We define a smooth interpolating curve of the fixed images $\tilde{u}_t := \sum_{j=1}^J \eta_t^j u_j^I$. Let $\tilde{\mathbf{u}}^K := (\tilde{u}_{t_1^K}, \dots, \tilde{u}_{t_J^K})$ and define $\tilde{u}^K := \mathcal{U}^K[\tilde{\mathbf{u}}^K, \mathbb{1}^K]$. This image curve gives an admissible candidate for a minimizer of the functional $\mathcal{F}^{\sigma, K}$. Indeed,

$$\begin{aligned} \mathcal{F}^{\sigma, K}[\tilde{u}^K] &\leq \mathbf{F}^{\sigma, K, D}[\tilde{\mathbf{u}}^K, \mathbb{1}^K] \\ &= \sigma K \sum_{k=1}^K \int_{\Omega} |\tilde{u}_k^K - \tilde{u}_{k-1}^K|^2 dx + K^3 \sum_{k=1}^{K-1} \int_{\Omega} |\tilde{u}_{k+1}^K - 2\tilde{u}_k^K + \tilde{u}_{k-1}^K|^2 dx \\ &\leq C \left(\int_{\Omega} |\tilde{u}|_{H^1((0,1))}^2 + |\tilde{u}|_{H^2((0,1))}^2 dx + 1 \right) := \bar{\mathcal{F}}, \end{aligned}$$

where the upper bound $\bar{\mathcal{F}}$ is independent of K . As defined above, let $\{u^K := \mathcal{U}^K[\mathbf{u}^K, \Phi^K]\}_{K \in \mathbb{N}}$, where $\{(\mathbf{u}^K, \Phi^K)\}_{K \in \mathbb{N}}$ are the optimal pairs for the discrete spline (see Theorem 5.2.7). In particular, for K large enough, we have $\mathcal{F}^{\sigma, K}[u^K] = \mathbf{F}^{\sigma, K}[\mathbf{u}^K, \Phi^K] < \bar{\mathcal{F}}$. Then, from Lemma 5.2.3, we have the uniform boundedness of $\{\phi_k^K\}_{K \in \mathbb{N}}$ in $C^{1, \alpha}(\bar{\Omega}, \bar{\Omega})$, and, as in (5.26), we can show the uniform boundedness of $\{u_k^K\}_{K \in \mathbb{N}}$ in \mathcal{I} , for every $k = 0, \dots, K$. Furthermore, using the boundedness of the discrete incremental transport paths, following from (5.39), we can show that $\{u_t^K\}_{K \in \mathbb{N}}$ is uniformly bounded in \mathcal{I} , uniformly in $t \in [0, 1]$. Therefore, $\{u^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $L^\infty([0, 1], \mathcal{I})$, and a subsequence converges weakly to some $u \in L^2([0, 1], \mathcal{I})$.

Let us show that $u_{t_j} = u_j^I$ for all $j = 1, \dots, J$. To this end, we can analogously to (5.54) show $u_t^K - \tilde{u}_t^K \rightarrow 0$ in \mathcal{I} , uniformly in $t \in [0, 1]$. Together with $u_t^K \rightarrow u_t$ in \mathcal{I} for every $t \in [0, 1]$, which can be shown as in (4.52), we have the needed result since $\tilde{u}_{t_j}^K = u_j^I$.

Now, we follow the usual argument and assume that there exists an image path $\hat{u} \in L^2([0, 1], \mathcal{I})$ with a finite energy, and $\hat{u}_{t_j} = u_j^I$ for all $j = 1, \dots, J$, such that $\mathcal{F}^\sigma[\hat{u}] < \mathcal{F}^\sigma[u]$. By the lim sup-part of Theorem 5.4.1, there exists a sequence $\{\hat{u}^K\}_{K \in \mathbb{N}} \subset L^2((0, 1), \mathcal{I})$ of time extensions of admissible vectors of images such that $\limsup_{K \rightarrow \infty} \mathcal{F}^{\sigma, K}[\hat{u}^K] \leq \mathcal{F}^\sigma[\hat{u}]$. Now, we apply the lim inf-part of Theorem 5.4.1 to obtain

$$\mathcal{F}^\sigma[u] \leq \liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma, K}[u^K] \leq \liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma, K}[\hat{u}^K] \leq \mathcal{F}^\sigma[\hat{u}], \quad (5.56)$$

which is a contradiction to the above assumption. Hence, u minimizes the continuous spline energy over all admissible image curves and the discrete spline energies converge to the limiting spline energy along a subsequence, *i.e.*, $\lim_{K \rightarrow \infty} \mathcal{F}^{\sigma, K}[u^K] = \mathcal{F}^\sigma[u]$, which follows from (5.56) by using $\hat{u} = u$. \square

The analogous result for arbitrary $(t_1, \dots, t_J) \subset [0, 1]$ follows from the density of \mathbb{Q} in $[0, 1]$. Let us remark that, in light of Proposition 5.1.1, we have that for the optimal scalar quantities z, w it holds $z = |\hat{z}|$ and $w = |\hat{w}|$ (*cf.* Remark 4.4.5).

5.5 Relaxation of the Warping Constraint

As the next step, we should proceed to the fully discrete model. In order to make comparisons of the piecewise geodesic interpolation with the spline interpolation, we want to follow the analogous discretization and optimiza-

tion techniques as in Sections 4.5 and 4.6. In particular, for $k = 2, \dots, K - 1$ we would have to determine

$$\begin{aligned} \phi_k \in \operatorname{argmin}_{\phi} \int_{\Omega} K^3 (W_A(D(\phi_{k+1} \circ \phi - 2\phi + \mathbf{1})) + W_A(D(\phi \circ \phi_{k-1} - 2\phi_{k-1} + \mathbf{1}))) \\ + \frac{K^3}{\delta} (|u_{k+1} \circ \phi_{k+1} \circ \phi - 2u_k \circ \phi + u_{k-1}|^2 + |u_k \circ \phi \circ \phi_{k-1} - 2u_{k-1} \circ \phi_{k-1} + u_{k-2}|^2) \\ + K\sigma \left(W_D(D\phi) + \frac{1}{\delta} |u_k \circ \phi - u_{k-1}|^2 \right) dx. \end{aligned}$$

Then, in order to use the linearization and the iPALM algorithm as proposed in Section 4.6, we would have to rewrite the second term in the second line as

$$\int_{\Omega} |u_k \circ \phi - 2u_{k-1} + u_{k-2} \circ \phi_{k-1}^{-1}|^2 \det(D\phi_{k-1})^{-1} \circ \phi_{k-1}^{-1} dx.$$

However, we do not have a stable and precise way for computation of the inverse of a given deformation. To solve this problem, we propose a relaxation of the model presented in Section 5.1 and its variational time discretization given in Section 5.2.

First, we explicitly introduce a vector valued material derivative $\bar{z} \in L^2((0, 1), \mathcal{Z} := L^2(\Omega, \mathbb{R}^c))$ and obtain a relaxation of (5.10):

$$\mathcal{F}_r^\sigma[u] := \inf_{(v, a, \bar{z}, \bar{z}, w)} \int_0^1 \int_{\Omega} L[a, a] + \frac{1}{\delta} |w|^2 + \sigma(L[v, v] + \frac{1}{\delta} |\bar{z}|^2) + \frac{1}{\theta} |\bar{z} - \hat{z}|^2 dx dt,$$

with a penalty on the misfit of the new variable \bar{z} and the actual material derivative \hat{z} (cf. (5.5)), while w is the first material derivative of \bar{z} , i.e.,

$$u_t \circ \psi_t - u_s \circ \psi_s = \int_s^t \hat{z}_r \circ \psi_r dr, \quad \bar{z}_t \circ \psi_t - \bar{z}_s \circ \psi_s = \int_s^t w_r \circ \psi_r dr, \quad \forall s, t \in [0, 1].$$

To adapt the time discrete counterpart, with $\bar{\mathbf{z}} = (\bar{z}_1, \dots, \bar{z}_K)$, we define

$$\begin{aligned} \mathbf{F}_r^{\sigma, K, D}[\mathbf{u}, \bar{\mathbf{z}}, \Phi] \\ := \int_{\Omega} \sum_{k=1}^{K-1} \frac{1}{K} (W_A(Da_k) + \gamma |D^m a_k|^2) + \frac{K}{\delta} |\bar{z}_{k+1} \circ \phi_k - \bar{z}_k|^2 + \sigma \left(\sum_{k=1}^K K W_D(D\phi_k) + K\gamma |D^m \phi_k|^2 + \frac{1}{\delta K} |\bar{z}_k|^2 \right) \\ + \sum_{k=1}^K \frac{1}{\theta K} |K(u_k \circ \phi_k - u_{k-1}) - \bar{z}_k|^2 dx, \end{aligned} \quad (5.57)$$

where $\bar{w}_k := K(\bar{z}_{k+1} \circ \phi_k - \bar{z}_k)$ is the discrete material derivative of \bar{z}_k , while $\hat{z}_k := K(u_k \circ \phi_k - u_{k-1})$ is the actual material derivative of u_k , with the corresponding second order derivative $(\hat{w}_k)_{k=1}^{K-1}$ given by $\hat{w}^k := K(\hat{z}_{k+1} \circ \phi_k - \hat{z}_k)$ (cf. (5.17)). We now define

$$\mathbf{F}_r^{\sigma, K}[\mathbf{u}] := \inf_{(\Phi, \bar{\mathbf{z}}) \in \mathcal{D}^K \times \mathcal{Z}^K} \mathbf{F}_r^{\sigma, K, D}[\mathbf{u}, \bar{\mathbf{z}}, \Phi].$$

This energy is well defined and has a minimizer in the set of admissible vectors of images (cf. (5.19)).

Theorem 5.5.1. *For every $\mathbf{u} \in \mathcal{I}_{adm}^K$, there exists a pair $(\Phi, \bar{\mathbf{z}}) \in \mathcal{D}^K \times \mathcal{Z}^K$ such that*

$$\mathbf{F}_r^{\sigma, K}[\mathbf{u}] = \mathbf{F}_r^{\sigma, K, D}[\mathbf{u}, \bar{\mathbf{z}}, \Phi].$$

Furthermore, for every $K \in \mathbb{N}$ and every choice of \mathcal{I}_f^K , there exists $\mathbf{u} \in \mathcal{I}_{adm}^K$ such that

$$\mathbf{F}_r^{\sigma, K}[\mathbf{u}] = \inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}_r^{\sigma, K}[\mathbf{v}].$$

Proof. The first claim follows along the same lines as Proposition 5.2.5, while for the second we follow the arguments analogous to those from the proof of Theorem 5.2.7. \square

5.6 Fully Discrete Model

In this section, we present (relaxed) fully discrete splines for the image metamorphosis model. We observe images with c channels and the two-dimensional image domain $\Omega := [0, 1]^2$. In a large part we follow the same approach as in Section 4.5. To this end, for $M, N \geq 3$ we define the computational domain

$$\Omega_{MN} := \left\{ \frac{0}{M-1}, \frac{1}{M-1}, \dots, \frac{M-1}{M-1} \right\} \times \left\{ \frac{0}{N-1}, \frac{1}{N-1}, \dots, \frac{N-1}{N-1} \right\}, \quad \partial\Omega_{MN} := \Omega_{MN} \cap \partial([0, 1]^2),$$

and the L^p -norm of the discrete objects via $\|\mathbf{u}\|_{L^p(\Omega_{MN})}^p = \frac{1}{MN} \sum_{(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}} \sum_{j=1}^c |\mathbf{u}^j(\mathbf{x}, \mathbf{y})|^p$. We define the discrete image and discrete derivative space by $\mathcal{I}_{MN} = \mathcal{Z}_{MN} := \{\mathbf{u}, \mathbf{z} : \Omega_{MN} \rightarrow \mathbb{R}^c\}$, and the set of admissible deformations by

$$\mathcal{D}_{MN} := \left\{ \phi = (\phi^1, \phi^2) : \Omega_{MN} \rightarrow [0, 1]^2, \phi = \mathbf{1} \text{ on } \partial\Omega_{MN}, \det(\nabla_{MN} \phi) > 0 \right\},$$

where the discrete Jacobian operator ∇_{MN} of ϕ at $(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}$ is defined as the forward finite difference operator with Neumann boundary conditions. Here, and in the rest of this chapter, we use boldfaced letters for the fully discrete quantities.

A spatial warping operator \mathbf{T} that approximates the pullback of an image channel $\mathbf{u}^j \circ \phi$ at a point $(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}$ is written in the following form

$$\mathbf{T}[\mathbf{u}^j, \phi](\mathbf{x}, \mathbf{y}) := \sum_{(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \Omega_{MN}} s(\phi^1(\mathbf{x}, \mathbf{y}) - \tilde{\mathbf{x}}) s(\phi^2(\mathbf{x}, \mathbf{y}) - \tilde{\mathbf{y}}) \mathbf{u}^j(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}),$$

where s is a sufficiently smooth function such that $s(0) = 1$ and it vanishes at other grid points $\Omega_M \setminus \{0\}$, where $\Omega_M := \{-\frac{M-1}{M-1}, \dots, 0, \dots, \frac{M-1}{M-1}\}$. (We focus our description on the first dimension, while the construction for the other dimension is completely analogous). We define s as the interpolation of these values by the *third order B-spline* (for details and proofs of all the claims below see [dB78, Chapter IX])

$$s(x) := \sum_{i=-M}^{M+1} \mathbf{w}^i b^i(x). \quad (5.58)$$

Here, $\mathbf{w}^i \in \mathbb{R}$ are weights which are determined from the interpolation constraints and the i -th basis function b^j is given by

$$b^i(x) := \beta_{-2}^i(x) \chi_{I_{i-2}}(x) + \beta_{-1}^i(x) \chi_{I_{i-1}}(x) + \beta_0^i(x) \chi_{I_i}(x) + \beta_1^i(x) \chi_{I_{i+1}}(x), \quad (5.59)$$

where, for $i \in \{-M, \dots, M+1\}$, we define $I_i := [\mathbf{x}_i, \mathbf{x}_{i+1} = \mathbf{x}_i + \frac{1}{M-1})$ with $\mathbf{x}_i := \frac{i}{M-1}$. The basis functions are determined as the elements with order 3 in the sequence determined by the recursion formula

$$b_0^i(x) = \chi_{I_i}(x),$$

$$b_{k+1}^i(x) = \omega_k^i(x) b_k^i(x) + (1 - \omega_k^{i+1}(x)) b_k^{i+1}(x), \quad \omega_k^i(x) = \begin{cases} \frac{x - \mathbf{x}_i}{\mathbf{x}_{i+k} - \mathbf{x}_i}, & \mathbf{x}_{i+k} \neq \mathbf{x}_i, \\ 0, & \text{else.} \end{cases}$$

This, in particular, ensures that at any point only a small number of basis functions does not vanish and the sum of values of all basis functions at any point is 1. In our setting, the basis functions, depicted on Figure 5.2 right, are given by

$$\beta_{-2}^i(x) := \frac{1}{6} (M-1)^3 (x - \mathbf{x}_{i-2})^3,$$

$$\beta_{-1}^i(x) := \frac{1}{6} (M-1)^3 \left((x - \mathbf{x}_{i-2})^2 (\mathbf{x}_i - x) + (x - \mathbf{x}_{i-1})^2 (\mathbf{x}_{i+2} - x) + (x - \mathbf{x}_{i-2}) (x - \mathbf{x}_{i-1}) (\mathbf{x}_{i+1} - x) \right),$$

$$\beta_0^i(x) := \frac{1}{6} (M-1)^3 \left((\mathbf{x}_{i+2} - x)^2 (x - \mathbf{x}_i) + (\mathbf{x}_{i+1} - x)^2 (x - \mathbf{x}_{i-2}) + (x - \mathbf{x}_{i-1}) (\mathbf{x}_{i+1} - x) (\mathbf{x}_{i+2} - x) \right),$$

$$\beta_1^i(x) := \frac{1}{6} (M-1)^3 (\mathbf{x}_{i+2} - x)^3.$$

We notice that the basis functions are translates of each other, *i.e.*, $b^i(x) = b^{i+k}(x + \frac{k}{M-1})$, and that the segments are scale invariant, *i.e.*, if we observe the segments $(\tilde{\beta}_k^i)_{k=-2,-1,0,1}$ on the domain with the grid size $\frac{1}{M-1}$ we have $\beta_k^i(x) = \tilde{\beta}_k^i(x \frac{M-1}{M-1})$.

From (5.58) and (5.59), for $x \in I_i$, we have

$$s(x) = \mathbf{w}^{i-1} \beta_1^{i-1}(x) + \mathbf{w}^i \beta_0^i(x) + \mathbf{w}^{i+1} \beta_{-1}^{i+1}(x) + \mathbf{w}^{i+2} \beta_{-2}^{i+2}(x),$$

which gives

$$s(\mathbf{x}_i) = \frac{1}{6} \mathbf{w}^{i-1} + \frac{2}{3} \mathbf{w}^i + \frac{1}{6} \mathbf{w}^{i+1}, \quad i = -(M-1), \dots, M-1,$$

so that we can write a linear tridiagonal system to determine the weights. However, we still have two degrees of freedom more than the number of equations. To this end, we impose the clamped boundary conditions $s''(-1) = s''(1) = 0$, which in our setting correspond to the condition $\mathbf{w}^{-(M-1)} = \mathbf{w}^{M-1} = 0$. Figure 5.2 left shows the resulting function s for $M = 16$. We notice that, due to the small support of the B-spline basis and the fact that only non-vanishing value to interpolate is $s(0) = 1$, the relevant contribution to the function s is only present on a support of length of several (four) discrete intervals around 0. Due to the scale invariance of the segments, the same observation holds for other relevant choices of M .

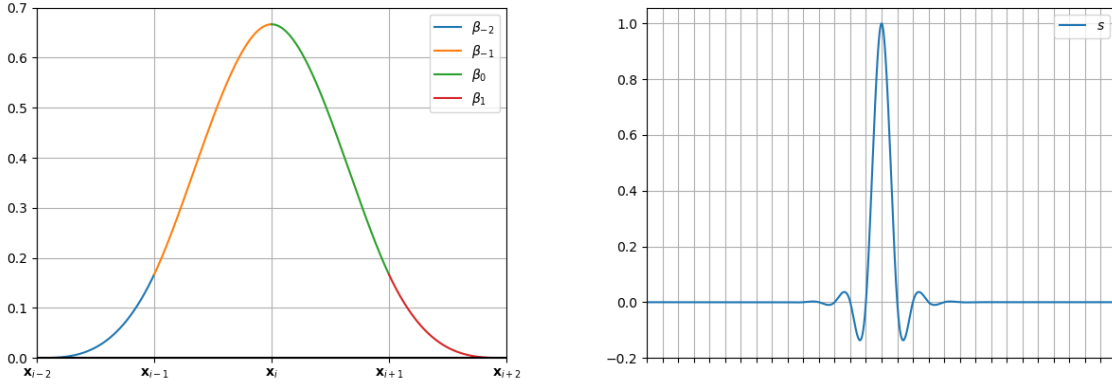


Figure 5.2: Right: B-spline basis b^i with segments colored in different colors. Left: Interpolation function s for $M = 16$.

From the construction, we clearly see that $s \in C^2([-1, 1])$, which allows us to obtain sufficient regularization of the deformations via discretization, and thus neglect the H^m -seminorm term in the fully discretized energy.

This form of warping is also used for the composition of deformations, *i.e.*, we define the fully discrete acceleration, as an approximation of (5.16), by

$$\mathbf{a}_k^j := K^2(\mathbf{T}[\phi_{k+1}^j - \mathbb{1}, \phi_k] - (\phi_k^j - \mathbb{1})), \quad j = 1, 2. \quad (5.60)$$

In summary, the fully discrete spline energy in the metamorphosis model, for a $(K+1)$ -tuple $(\mathbf{u}_k)_{k=0}^K$ of discrete images, a K -tuple $(\bar{\mathbf{z}}_k)_{k=1}^K$ of discrete derivatives, and a K -tuple of discrete deformations $(\phi_k)_{k=1}^K$, is obtained as a discretization of (5.57), and reads as

$$\begin{aligned} & \mathbf{F}_{r, MN}^{\sigma, K}[(\mathbf{u}_k)_{k=0}^K] \\ & := \inf_{(\mathbf{z}^K, \Phi) \in \mathcal{Z}_{MN}^K \times \mathcal{D}_{MN}^K} \mathbf{F}_{r, MN}^{\sigma, K, D}[(\mathbf{u}_k)_{k=0}^K, (\bar{\mathbf{z}}_k)_{k=1}^K, (\phi_k)_{k=1}^K] \\ & := \inf_{(\mathbf{z}^K, \Phi) \in \mathcal{Z}_{MN}^K \times \mathcal{D}_{MN}^K} \sum_{k=1}^{K-1} \frac{1}{K} \|\mathbf{W}_A(\nabla_{MN} \mathbf{a}_k)\|_{L^1(\Omega_{MN})} + \frac{K}{\delta} \mathbf{D}_{MN}^s[\bar{\mathbf{z}}_k, \bar{\mathbf{z}}_{k+1}, \phi_k] \\ & \quad + \sum_{k=1}^K \sigma \left(K \|\mathbf{W}_D(\nabla_{MN} \phi_k)\|_{L^1(\Omega_{MN})} + \frac{1}{\delta K} \|\bar{\mathbf{z}}_k\|_{L^2(\Omega_{MN})}^2 \right) + \frac{1}{\theta K} \mathbf{D}_{MN}^g[\mathbf{u}_{k-1}, \mathbf{u}_k, \bar{\mathbf{z}}_k, \phi_k], \end{aligned}$$

where

$$\mathbf{D}_{MN}^s[\mathbf{z}, \tilde{\mathbf{z}}, \phi] := \frac{1}{2c} \sum_{j=1}^c \left\| \mathbf{T}[\tilde{\mathbf{z}}^j, \phi] - \mathbf{z}^j \right\|_{L^2(\Omega_{MN})}^2, \quad \mathbf{D}_{MN}^g[\mathbf{u}, \tilde{\mathbf{u}}, \mathbf{z}, \phi] := \frac{1}{2c} \sum_{j=1}^c \left\| K(\mathbf{T}[\tilde{\mathbf{u}}^j, \phi] - \mathbf{u}^j) - \mathbf{z}^j \right\|_{L^2(\Omega_{MN})}^2.$$

To improve the robustness of the overall optimization, we take into account a multilevel strategy. In detail, on the coarse computational domain of size $M_L \times N_L$ with $M_L = 2^{-(L-1)}M$ and $N_L = 2^{-(L-1)}N$, for a given $L \geq 1$, a time discrete spline sequence $(\mathbf{u}_k)_{k=0}^K$ is computed as a minimizer of $\mathbf{F}_{r, M_L N_L}^{\sigma, K}$ subject to the given fixed images $\mathbf{u}_{i_j} = \mathbf{u}_j^I$, $j = 1, \dots, J^K$ (cf. (5.19)). In the subsequent prolongation steps, the width and the height of the computational domain are successively doubled, and the initial deformations, images, and derivatives are obtained via bilinear interpolation of the preceding coarse scale solutions.

5.7 Numerical Optimization

In this section, we discuss the numerical solution of the above-defined fully discrete variational problem, based on the application of a variant of the iPALM algorithm [PS16], presented in Section 3.3. Following the approach from Section 4.6, to enhance the stability, the warping operation is linearized with respect to the deformation at $\phi^{[\beta]} \in \mathcal{D}_{MN}$ coming from the previous iteration, leading to the modified energies

$$\tilde{\mathbf{D}}_{MN}^s[\mathbf{z}, \tilde{\mathbf{z}}, \phi, \phi^{[\beta]}] := \frac{1}{2c} \sum_{j=1}^c \left\| \mathbf{T}[\tilde{\mathbf{z}}^j, \phi^{[\beta]}] + \left\langle \Lambda_j(\mathbf{z}, \tilde{\mathbf{z}}, \phi^{[\beta]}), \phi - \phi^{[\beta]} \right\rangle - \mathbf{z}^j \right\|_{L^2(\Omega_{MN})}^2$$

$$\tilde{\mathbf{D}}_{MN}^g[\mathbf{u}, \tilde{\mathbf{u}}, \mathbf{z}, \phi, \phi^{[\beta]}] := \frac{1}{2c} \sum_{j=1}^c \left\| K\mathbf{T}[\tilde{\mathbf{u}}^j, \phi^{[\beta]}] + \left\langle \Lambda_j(K\mathbf{u} + \mathbf{z}, K\tilde{\mathbf{u}}, \phi^{[\beta]}), \phi - \phi^{[\beta]} \right\rangle - (K\mathbf{u}^j + \mathbf{z}^j) \right\|_{L^2(\Omega_{MN})}^2,$$

based on the gradient (cf. [WPZ⁺09]) $\Lambda_j(\mathbf{u}, \tilde{\mathbf{u}}, \phi^{[\beta]}) = \frac{1}{2}(\nabla_{MN} \mathbf{T}[\tilde{\mathbf{u}}^j, \phi^{[\beta]}] + \nabla_{MN} \mathbf{u}^j)$, where the Jacobian operator applied to the images is approximated using a Sobel filter [Sob90]. Here, $\langle \cdot, \cdot \rangle$ represents the pointwise product of the involved matrices. We use the Moreau's proximal mapping of a functional $f : \mathcal{D}_{MN} \rightarrow (-\infty, \infty]$ for $\tau > 0$ given as (cf. Definition 3.3.1) $\text{prox}_\tau^f[\phi] := \text{argmin}_{\tilde{\phi} \in \mathcal{D}_{MN}} \left(\frac{\tau}{2} \left\| \phi - \tilde{\phi} \right\|_{L^2(\Omega_{MN})}^2 + f[\tilde{\phi}] \right)$. Then, with the function values on $\partial\Omega_{MN}$ remaining unchanged, the proximal operator we are interested in is

$$\begin{aligned} & \text{prox}_{\frac{K}{c\tau\delta}}^{\tilde{\mathbf{D}}_{MN}^s + \frac{1}{K\theta} \tilde{\mathbf{D}}_{MN}^g}[\phi_k^t] \\ &= \left(\mathbf{1} + \frac{K}{c\tau\delta} \sum_{j=1}^c |\Lambda_j^s|^2 + \frac{1}{c\tau\theta K} \sum_{j=1}^c |\Lambda_j^g|^2 \right) \\ & \quad \left(\phi_k^t - \frac{K}{c\tau\delta} \sum_{j=1}^c \Lambda_j^s (\mathbf{T}[\tilde{\mathbf{z}}_{k+1}^j, \phi_k^{[\beta]}] - (\Lambda_j^s)^T \phi_k^{[\beta]} - \tilde{\mathbf{z}}_k^j) - \frac{1}{c\tau\theta K} \sum_{j=1}^c \Lambda_j^g (\mathbf{T}[K\tilde{\mathbf{u}}_k^j, \phi_k^{[\beta]}] - (\Lambda_j^g)^T \phi_k^{[\beta]} - K\mathbf{u}_{k-1}^j - \tilde{\mathbf{z}}_k^j) \right), \end{aligned}$$

where $\Lambda_j^s := \Lambda(\tilde{\mathbf{z}}_k^j, \tilde{\mathbf{z}}_{k+1}^j, \phi_k^{[\beta]})$ and $\Lambda_j^g := \Lambda(K\mathbf{u}_{k-1}^j + \tilde{\mathbf{z}}_k^j, K\tilde{\mathbf{u}}_k^j, \phi_k^{[\beta]})$. The first terms in both brackets are activated only for $k < K$.

We are now ready to apply the iPALM algorithm with the backtracking of the Lipschitz constant (cf. Algorithm 3) to minimization of $\mathbf{F}_{r, MN}^K$. We initialize the vector of images by the linear interpolation of the fixed images, the vector of derivatives as the scaled differences of the consecutive images, and the vector of deformations by the identity vector. Finally, we choose some initial values for the step sizes (the Lipschitz constants w.r.t. the images, derivatives and deformations). Then, the alternating optimization is summarized in Algorithm 5, while the backtracking of the Lipschitz constants is achieved as in Algorithm 3. For $\beta > 0$, we use the notation

$$\mathbf{h}_k^{[i, \beta]} := \mathbf{h}_k^{[i]} + \beta(\mathbf{h}_k^{[i]} - \mathbf{h}_k^{[i-1]}), \quad \mathbf{h}^{[i, k]} := (\mathbf{h}_{0, \dots, k-1}^{[i+1]}, \mathbf{h}_{k, \dots, K}^{[i]}), \quad \mathbf{h}^{[i, k, \beta]} := (\mathbf{h}_{0, \dots, k-1}^{[i+1]}, \mathbf{h}_k^{[i, \beta]}, \mathbf{h}_{k+1, \dots, K}^{[i]}),$$

and the acceleration $\mathbf{a}^{[i, \beta]}$ is computed from the formula (5.60), using the correspondingly updated $\phi^{[i, \beta]}$ values. Finally, I^K is the set of the indices with fixed images (cf. (5.19)), and the derivatives of the energies w.r.t. the deformations, material derivatives and images are computed with a help of the *Autograd* package of PyTorch library [PGC⁺17].

Algorithm 5: Algorithm for the minimization of $\mathbf{F}_{r,MN}^{\sigma,K}$ on one level.

```

1 for  $i = 1$  to  $I$  do
  /* start with  $k = 0$  for images                                     */
2 for  $k = 1$  to  $K$  do
  /* update deformation                                             */
3  $\phi_k^{i,t} =$ 
  
$$\phi_k^{[i,\beta]} - \frac{D_{\phi_k} \left( \sigma K \|\mathbf{W}_D(\nabla_{MN} \phi_k^{[i,\beta]})\|_{L^1(\Omega_{MN})} + \frac{1}{K} \|\mathbf{W}_A(\nabla_{MN} \mathbf{a}_k^{[i,\beta]}) + \mathbf{W}_A(\nabla_{MN} \mathbf{a}_{k-1}^{[i,\beta]})\|_{L^1(\Omega_{MN})} \right)}{L_{\phi_k}^{[i]}};$$

4  $\phi_k^{[i+1]} = \text{prox}_{L_{\phi_k}^{[i]}} \left[ \frac{K}{\delta} \tilde{\mathbf{D}}_{MN}^s + \frac{1}{K\sigma} \tilde{\mathbf{D}}_{MN}^g \right] [\phi_k^{i,t}];$ 
  /* update derivative                                             */
5  $\bar{\mathbf{z}}_k^{[i+1]} = \bar{\mathbf{z}}_k^{[i,\beta]} - \frac{D_{\bar{\mathbf{z}}_k} \mathbf{F}_{r,MN}^{\sigma,K,D} [\mathbf{u}^{[i,k]}, \bar{\mathbf{z}}^{[i,k,\beta]}, \phi^{[i,k+1]}]}{L_{\bar{\mathbf{z}}_k}^{[i]}};$ 
  /* fixed images remain unchanged                                  */
6 if  $k \notin I^K$  then
  /* update image                                                 */
7  $\mathbf{u}_k^{[i+1]} = \mathbf{u}_k^{[i,\beta]} - \frac{D_{\mathbf{u}_k} \mathbf{F}_{r,MN}^{\sigma,K,D} [\mathbf{u}^{[i,k,\beta]}, \bar{\mathbf{z}}^{[i,k+1]}, \phi^{[i,k+1]}]}{L_{\mathbf{u}_k}^{[i]}};$ 

```

5.8 Numerical Results

In what follows, we investigate and discuss the qualitative properties of the spline interpolation in the space of images, being aware that the superior temporal smoothness of this interpolation is difficult to show with a series of still images. For all the examples we use $L = 5$ levels in the multilevel approach and $I = 250$ iterations of iPALM algorithm on each level, with the extrapolation parameter $\beta = \frac{1}{\sqrt{2}}$. For the first two examples (Gaussians and circle-square) $M = N = 64$, while for the others (human portraits, letters and cells) $M = N = 128$. Also, for the first two and the final example $K = 8$, while for the others $K = 16$. For the plotting of the images we crop the values to $[0, 1]$, for the material derivatives the values are scaled to the interval $[0, 1]$ for plotting, while for the displacement and acceleration plots hue refers to the direction and the intensity is proportional to its norm, as indicated by the color wheel. Let us also notice that the displacement plots can also be observed as the plots of (scaled) discrete velocity fields, so we will use both notions.

Figure 5.3 shows a first test case. As key frames, we consider three images showing two-dimensional Gaussian distribution with a small variance at different positions, and of different masses. For the metamorphosis spline, the curve in the (x, y, m) -space (position, mass) corresponds almost perfectly to the cubic spline interpolation of the parameters of the Gaussian distribution on the key frames. Furthermore, the curve in (x, y, m) -coordinates obtained from the piecewise geodesic interpolation corresponds almost perfectly to the piecewise geodesic interpolation of the corresponding parameters. This, in particular, shows the smoothness of our proposed spline interpolation in comparison with the piecewise geodesic interpolation in the metamorphosis model.

Next, in Figure 5.4, we conceptually compare spline and piecewise geodesic interpolation in the image metamorphosis model. For this specific example, suggested and implemented by J. Justiniano, we considered the image of a circle and two identical squares as key frames. The influence of the circle's curvature on the spline segment between the two identical squares is still visible via concave 'edges', while for the piecewise geodesic interpolation any memory of the circle is lost between the squares.

We consider the spline interpolation between the three human portraits in Figure 5.5. The plots of the second material derivative and the acceleration show a strong concentration around the key frames, where the spline is expected to be smooth and the piecewise geodesic path at most Lipschitz continuous. The analogous observations hold for Figure 5.6 which shows the spline interpolation between different letters of alphabet as key frames.




Figure 5.3: Left: Time discrete spline with framed key frame images (first row), color-coded discrete displacement/velocity field (second row), discrete second order material derivative (third row) and color-coded discrete acceleration field (fourth row), for the Gaussians example and values of the parameters $\delta = 5 \cdot 10^{-3}$, $\sigma = 1$, $\theta = 5 \cdot 10^{-5}$. The colors and their intensities indicate the direction and the intensity of the field, as indicated by the color wheel on the left. Right: Euclidean spline and piecewise geodesic interpolation for the input parameters versus spline and piecewise geodesic interpolation in metamorphosis model, extracted from the numerical results in postprocessing. The plots are given in (x, y, m) -coordinates, with (x, y) denoting the center of mass and m the mass of the distribution.

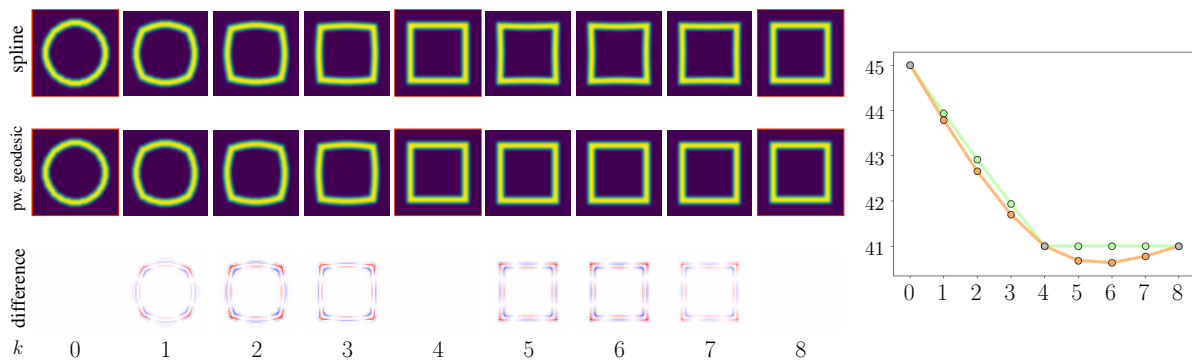


Figure 5.4: Left: Time discrete spline (top row) and piecewise geodesic (middle row) interpolation with framed key frames. The bottom row shows the difference in intensity between the different interpolations, using the color map -0.35 to 0.35 . Right: Width of the interpolated shape measured at the horizontal axis of symmetry (in number of pixels) for a spline interpolation (orange) and piecewise geodesic interpolation (green), showing the concavities in the spline interpolation ($\delta = 5 \cdot 10^{-3}$, $\sigma = 1$, $\theta = 5 \cdot 10^{-4}$).




Figure 5.5: Time discrete spline with framed fixed images (first and second row), first order material derivative variable $\bar{\mathbf{z}}$ (third and fourth row), second order material derivative with energies comparison (fifth and sixth row), and color-coded acceleration field with energies comparison (seventh and eighth row), for the values of the parameters $\delta = 2 \cdot 10^{-2}$, $\sigma = 2$, $\theta = 8 \cdot 10^{-4}$. The graphics on the right, in rows four and six, show the spline (orange) time plots of the L^2 -norm of the actual second order material derivative $\hat{\mathbf{w}}_k$ and the dissipation energy density $\|\mathbb{W}_A(\nabla_{MN} \mathbf{a}_k)\|_{L^1_{MN}}$, reflecting the motion acceleration, respectively. This is compared to the corresponding piecewise geodesic interpolation (green) (not visualized here, *cf.* Figure 5.8).



Figure 5.6: Left: Time discrete spline with framed fixed images (top and bottom row). Right: Energy density norm of acceleration flow $\|W_A(\nabla_{MN}\mathbf{a}_k)\|_{L^1_{MN}}$ (top), and L^2 -norm of the actual second order material derivative $\hat{\mathbf{w}}_k$ (bottom). Parameter values: $\delta = 10^{-3}$, $\sigma = 2$, $\theta = 2 \cdot 10^{-5}$.

The impact of the coloring of the key frame images on the geometry of the spline interpolation and the interplay between the Eulerian flow acceleration and the second order material derivative along the motion paths is depicted in Figure 5.7. Therein, we consider block-colored letters as key frame images. In our first example, the coloring is consistent with the interpolating flow shown in the black and white spline interpolation in Figure 5.6. Hence, in the corresponding spline interpolation color is mainly passively transported along this flow rather than blending it (top row). In the second example, the red patch in the middle key frame image is chosen to be located top right instead, further away from the red patches in the extremal key frame images (bottom row). Hence, the coloring is no longer consistent to the flow in the black and white example. In fact, the transport of the blue color is now strongly reconfigured, as seen in the first interval between the ‘P’ and the ‘A’. Furthermore, in the second interval between the ‘A’ and the ‘Q’, blending from blue to red and from red to blue occurs.




Figure 5.7: Top and bottom rows: Two time discrete splines with key-frames images differing in shape and color. Parameter values: $\delta = 8 \cdot 10^{-3}$, $\sigma = 2.5$, $\theta = 2 \cdot 10^{-4}$. For visualization purposes, only even-numbered frames \mathbf{u}_k are depicted.

In Figure 5.8, we compare particular images of the discrete splines and piecewise geodesic sequences. One particularly observes that for the faces the shown spline image is thicker and for the letters the spline image shows more round contours than for the piecewise geodesic counterpart. This is again the nonlocal impact of the key frames beyond those bounding the current interpolation. Furthermore, we compare the flow acceleration and the second order material derivative terms on which the jump in energies occurs (*cf.* Figures 5.5 and 5.6).

Finally, we ask for a reconstruction of frames given certain frames at selected time stamps extracted from a video. Here, we compare the resulting spline interpolation and the piecewise geodesic interpolation directly with corresponding frames of the original video as a benchmark for the both approaches. Indeed, Figure 5.9 shows this comparison of the original frames, the spline interpolation, and the piecewise geodesic interpolation for the images extracted from a video made by David Rogers from Vanderbilt University in the 1950s¹, which shows the interaction between white blood cells and bacteria. The spline interpolation clearly shows less blending artifacts and it is smoother in time in comparison with the piecewise geodesic interpolation.

¹https://embryology.med.unsw.edu.au/embryology/index.php/Movie_-_Neutrophil_chasing_bacteria. The images are courtesy of Robert A. Freitas, Institute for Molecular Manufacturing, California, USA rfreitas.com




Figure 5.8: Top row: Image \mathbf{u}_{11} for the human face example, second order material derivative $\hat{\mathbf{w}}_k$ and acceleration field \mathbf{a}_k for $k = 8$, for the time discrete piecewise geodesic (left image of each panel pair) and spline (right image of each panel). The pairs of material derivatives and the acceleration fields are jointly scaled to reflect the differences in intensities. Bottom row: Same visualization with image \mathbf{u}_4 from the letter example.




Figure 5.9: First row: The original frames extracted from the video showing white blood cell (neutrophil) chasing a *Staphylococcus aureus* bacterium. Second row: Time discrete spline with framed fixed images. Third row: Time discrete piecewise geodesic with framed fixed images. Fourth and fifth row: Fully discrete second order material derivative for the discrete spline and the the discrete piecewise geodesic interpolation, respectively. The values of parameters are $\delta = 4 \cdot 10^{-2}$, $\sigma = 2.5$, $\theta = 1.6 \cdot 10^{-4}$.

5.9 Conclusion and Outlook

In this chapter, we have proposed a spline interpolation approach for the image metamorphosis model. The approach is based on the penalization of the acceleration term, which is split into the flow acceleration and the second order change of image intensity. Variational time discretization is based on the approximation of the acceleration terms by the central finite difference. Both the time continuous and the time discrete energy are regularized by adding scaled geodesic energy. We were able to show the existence of time discrete spline interpolations as minimizers of the regularized energy. Furthermore, we showed the Mosco-convergence of the extension of time discrete energies, based on interpolation by cubic Hermite splines, towards the continuous time energy. As a corollary, we obtained the existence of a time continuous spline interpolation as a minimizer of this energy and the convergence of the time discrete interpolants towards their continuous counterpart. The fully discrete model is based on the finite difference approximation of derivatives and a cubic B-spline approximation of the warping operator, respectively, while the numerical optimization is done by using a version of the iPALM algorithm, proposed by Pock and Sabach. The experimental results show better smoothness in time of the proposed spline interpolation in comparison to the piecewise geodesic interpolation. This is, in particular, reflected in the stability of energy of both acceleration parts near the key frames. Furthermore, the influence of the preceding and/or succeeding key frame(s) is visible for the spline interpolation, which naturally relates to the larger support of the spline basis functions for Euclidean cubic splines. Although physically intuitive, this framework does not correspond to Riemannian splines introduced by Noakes *et al.*, as it does not represent penalization of the squared covariant derivative of the path velocity in the Riemannian metric. An attempt in building such a (time discrete) model could be based on the discretization of covariant derivative based on the framework by Effland, Heeren, Rumpf, and Wirth [EHRW21]. Furthermore, a challenge is the adaptation of our approach to images with textures and feature-based image representations.

Chapter 6

Learning Low Bending and Low Distortion Manifold Embeddings

AUTOENCODERS (see Section 1.2 for some introduction), consisting of an *encoder*, which embeds the input data manifold into a lower-dimensional latent space, and a *decoder*, which serves as an inverse map, providing a parametrization of the data manifold by the manifold in the latent space, are widely used in machine learning for dimension reduction of high-dimensional data. An appropriate regularity and structure of the embedded manifold may substantially simplify further processing of the data. In this chapter, we propose and analyze a novel regularization for learning the encoder component of an autoencoder by introducing a loss functional that prefers isometric and extrinsically flat embeddings. As training data, we use pairs of nearby points on the input manifold, together with their local Riemannian distance and Riemannian average. The loss functional is computed via Monte Carlo integration, where we use a few different sampling strategies. We first observe the limit as the number of samples tends to infinity, depending on the sampling strategy, leading to a nonlocal continuous sampling loss functional. We prove the existence of minimizers for this functional in a class of functions that can be obtained as a realization of a set of sufficiently well-behaved neural networks. We study the Mosco-limit of these functionals as the sampling radius, *i.e.*, the maximal distance of sampled point pairs, tends to zero. The limit is identified as a purely local continuous loss functional, which again promotes low bending and low distortion. Numerical experiments, using image data that encodes different explicitly given data manifolds, show that smooth manifold embeddings into latent space are obtained. Due to the promotion of extrinsic flatness, these embeddings are regular enough to ensure that the interpolation between not too distant points on the manifold is well approximated by the decoder image of the linear interpolation in the latent space. Note that this chapter is an extended version of the previous conference proceeding [BRRW21] and its journal version [BRRW22] (joint work with Juliane Brauns mann, Martin Rumpf and Benedikt Wirth).

This chapter is structured as follows. In Section 6.1, we provide the basic assumptions on the underlying manifold and introduce the discrete sampling loss functional. Furthermore, we obtain the continuous nonlocal regularization loss as the Monte Carlo limit for the dense sampling of input data. We prove the existence of a minimizer for this functional under the suitable assumptions on the class of embedding networks in Section 6.2. The Mosco-convergence of the nonlocal functionals towards the continuous local limit functional is demonstrated in Section 6.3. As a proof of concept, we train autoencoders on several image datasets representing a priori known data manifolds in Section 6.5. The neural network architecture of the encoders, together with the training procedure, is presented in Section 6.4. We discuss the numerical results, both from the perspective of smoothness, isometry, and flatness of obtained latent manifolds, and the quality of linear interpolations in the latent space. Finally, in Section 6.6, we provide a resume of the results of the chapter and comment on possible extensions of this work.

6.1 A Low Bending and Low Distortion Regularization for Encoders

We consider a smooth and compact m -dimensional Riemannian manifold $\overline{\mathcal{M}}$ with metric g , with or without boundary. Further assumptions on $\overline{\mathcal{M}}$ are given in the following subsection. For the moment, we assume that $\overline{\mathcal{M}}$ is embedded in some very high-dimensional space \mathbb{R}^N ($m \ll N$), where in our applications (see Section 6.5), this space is the space of grayscale or RGB images with N or $N/3$ pixels, respectively. Our aim is to compute the *latent manifold*, defined via an embedding (cf. Definition 2.2.5) of $\overline{\mathcal{M}}$ into Euclidean space \mathbb{R}^l , called the *latent space*. We want the *latent dimension* l to be only moderately larger than m , which is reasonable in light of Whitney’s embedding theorem [Whi92], which provides the existence of a smooth embedding for $l = 2m$. Formulated in terms of autoencoder framework, this amounts to learning a pair of maps

$$\phi_\theta : \overline{\mathcal{M}} \rightarrow \mathbb{R}^l, \psi_\xi : \mathbb{R}^l \rightarrow \mathbb{R}^N \text{ with } \psi_\xi(\phi_\theta(x)) \approx x \text{ for all } x \in \overline{\mathcal{M}}.$$

The *encoder*, ϕ_θ , and *decoder*, ψ_ξ , are implemented as deep neural networks with vectors of parameters θ and ξ , respectively (cf. Section 2.3). The parameters are typically optimized via the minimization of a loss functional measuring the difference between the original points x and their reconstructions $\psi_\xi(\phi_\theta(x))$. Minimization of the *reconstruction loss* is not enough to guarantee an appropriate structure and regularity of the *latent manifold* $\phi(\overline{\mathcal{M}})$. To this end, different versions of geometrically inspired regularization losses for the encoder (and decoder) are considered in addition to the regularization loss. The “regular” embeddings are of particular interest for downstream tasks such as classification [AGL20], Riemannian interpolation and extrapolation [BRRG18], clustering or anomaly detection [ZLC⁺21]. In the most simple case, $\overline{\mathcal{M}}$ is isometric to \mathbb{R}^m and thus can be embedded isometrically into an m -dimensional affine subspace of \mathbb{R}^l . Then, the most important and basic operations of computing distances and interpolations would become trivial in latent space, leading to simplicity in downstream tasks. Although such an embedding is usually prevented by the intrinsic or the global geometry of $\overline{\mathcal{M}}$, this motivates the search for an embedding as close as possible to isometric and flat, at least locally. To this end, we suggest the following two objectives for any two (not too distant) points $x, y \in \overline{\mathcal{M}}$:

Isometric (I) The intrinsic *Riemannian distance* between x and y in $\overline{\mathcal{M}}$ should differ as little as possible from the Euclidean distance between the latent codes $\phi(x)$ and $\phi(y)$.

Flat (F) The (weighted) *Riemannian average* between x and y in $\overline{\mathcal{M}}$ should deviate as little as possible from the (weighted) Euclidean average between $\phi(x)$ and $\phi(y)$.

The first objective (I) is, at least for infinitesimally close points $x, y \in \overline{\mathcal{M}}$, satisfied for *isometric embeddings*. This class of embeddings was extensively pursued in the literature (cf. Section 1.2 for related work). Isometry by itself is, however, insufficient to ensure regular embeddings, since the family of isometric embeddings is very large and contains quite irregular elements. In particular, Nash–Kuiper embeddings are in general only Hölder differentiable [Kui55b, Kui55a, Nas54]. Therefore we suggest (I) and (F) as stronger objectives:

- The isometry or low distortion objective (I) asks that the *intrinsic* distances between x and y in $\overline{\mathcal{M}}$ are approximated by the *extrinsic* distances between $\phi(x), \phi(y) \in \mathbb{R}^l$ in the latent space. This differs from an actual isometric embedding which would consider the intrinsic distances in the latent manifold $\phi(\overline{\mathcal{M}})$ with the metric induced by \mathbb{R}^l .
- The flatness or low bending objective (F) enforces some second order low bending regularity or flatness on ϕ by requiring that the *geodesic interpolation* between x and y in $\overline{\mathcal{M}}$ is well approximated by the extrinsic *linear interpolation* in the latent space \mathbb{R}^l .

In what follows, we give more details on the manifold $\overline{\mathcal{M}}$ and introduce the low bending and low distortion discrete sampling loss functional in Subsection 6.1.1, and study the Monte Carlo limit as the number of samples tends to infinity, depending of the sampling strategy, in Subsection 6.1.2.

6.1.1 A Low Bending and Low Distortion Loss Functional

In order to introduce the low bending and low distortion loss functional, we first give some more details on the structure of the manifold we are observing. For definitions of the below-mentioned notions on the manifold, we refer to Section 2.2 and the references therein.

We consider a smooth and compact m -dimensional Riemannian manifold $(\overline{\mathcal{M}}, g)$, with or without boundary, where in the latter case the boundary shall have dimension $m - 1$. As mentioned before, as the input data for the training of our encoder, we will also employ the Riemannian distance and average of point pairs on the manifold. In particular, the Riemannian average of two points on manifold is defined as the midpoint on the unique shortest geodesic connecting these two points. If the geodesic is not unique, the Riemannian average is not defined. One could slightly generalize the definition of the Riemannian mean as the midpoint of shortest connecting curves on manifold, which may include curves touching the boundary of $\overline{\mathcal{M}}$, but this way, for a given $x \in \overline{\mathcal{M}}$, there might exist several $y \in \overline{\mathcal{M}}$ with same Riemannian average. For this reason, we refrain from this generalization and, furthermore, decide to sample the point pairs only in the interior of the manifold. Let us denote by \mathcal{M} the relative interior of $\overline{\mathcal{M}}$, *i.e.*, $\mathcal{M} := \text{Int}(\overline{\mathcal{M}})$. By $T_x\mathcal{M}$ we denote the tangent space to \mathcal{M} at $x \in \mathcal{M}$. Let $d_{\mathcal{M}}(x, y)$ be the Riemannian distance between any two points $x, y \in \mathcal{M}$ and \exp_x the Riemannian exponential map defined on (a subset of) $T_x\mathcal{M}$. If $y = \exp_x v$, for $v \in T_x\mathcal{M}$ being the initial velocity of the unique shortest geodesic connecting x with y , then the Riemannian average is, due to the scaling properties of geodesics, given by $\text{av}_{\mathcal{M}}(x, y) := \exp_x \frac{v}{2}$.

In order for the above notions to be well-defined, we will require some natural conditions on the existence and behavior of the Riemannian exponential map associated with \mathcal{M} . To state these, for every $x \in \mathcal{M}$ let $V_x \subset T_x\mathcal{M}$ denote the largest open set on which \exp_x is defined. We will assume that there exist constants $r_0 > 0$ and $0 < \kappa < \frac{\pi}{2}$ such that the following holds:

- (M1) (injectivity condition) For $x \in \mathcal{M}$, let $B_{r_0}^{T_x\mathcal{M}}(0) := \{v \in T_x\mathcal{M} : g_x(v, v) < r_0^2\}$ denote the open ball in tangent space with radius r_0 . Then, the Riemannian exponential \exp_x is injective on $V_x \cap B_{r_0}^{T_x\mathcal{M}}(0)$, for all $x \in \mathcal{M}$.
- (M2) (cone condition) For every $x \in \mathcal{M}$, there exists an isometric isomorphism $\iota_x : \mathbb{R}^m \rightarrow T_x\mathcal{M}$ such that $x \mapsto \iota_x$ is Riemann–Lebesgue measurable on \mathcal{M} , and V_x contains $\iota_x(\mathcal{C}_{r_0, \kappa})$, where

$$\mathcal{C}_{r_0, \kappa} := \{w \in \mathbb{R}^m : 0 \leq |w| \leq r_0, \angle(w, e_1) \leq \kappa\}$$

is the cone of height r_0 and aperture angle κ with the first standard Euclidean basis vector $e_1 \in \mathbb{R}^m$.




Figure 6.1: A sketch of the cone condition (M2). The parametrization $\Pi_x = \exp_x \circ \iota_x$ maps the cone $\mathcal{C}_{r_0, \kappa} \subset U_x \subset \mathbb{R}^m$ into \mathcal{M} and the neighborhoods $U_x^\epsilon \subset U_x$ onto $D_\epsilon^{\mathcal{M}}(x)$.

In regard of (M1), for manifolds without boundary the largest r_0 coincides with the usual notion of injectivity radius, *i.e.*, the infimum over all $x \in \mathcal{M}$ of injectivity radius at x , as defined in Definition 2.2.12. In this case, we can replace the cone $\mathcal{C}_{r_0, \kappa}$ by $B_{r_0}^m(0) \subset \mathbb{R}^m$: the open ball with radius r_0 centered at the origin. For manifolds with boundary, the cone condition is not necessarily satisfied. For example, consider the submanifold $\overline{\mathcal{M}} = \{(x, y) \in \mathbb{R}^2 : y \geq \sqrt{|x|}, y \leq 1\}$ of \mathbb{R}^2 with the inherited metric. Then no cone can be positioned at $(0, 0)$. In general, the condition (M2) can be seen as a generalization of Definition 2.1.1, which allows us to define the linear isometric identification, for every $x \in \mathcal{M}$, of a subset of $V_x \subset T_x\mathcal{M}$ with a fixed and simple subset of \mathbb{R}^m (see Figure 6.1 for a sketch). The mapping ι_x , in a simple case, can be obtained by applying the orthonormalization procedure to the canonical basis of $T_x\mathcal{M}$, defined in Definition 2.2.3. Its isometry property is meant in the sense that, for every $x \in \mathcal{M}$, $w \in \mathbb{R}^m$, it holds $|w|^2 = g_x(\iota_x w, \iota_x w)$.

Taking the above assumptions, we can define a local parametrization of \mathcal{M} (cf. Figure 6.1)

$$\Pi_x : U_x \rightarrow \mathcal{M}, \quad \Pi_x := \exp_x \circ \iota_x, \quad \text{where } U_x := \iota_x^{-1}(V_x) \cap B_{r_0}^m(0). \quad (6.1)$$

This parametrization is known as the *normal coordinates around x* . For later notational convenience, we extend Π_x measurably (but arbitrarily) beyond U_x , and further introduce the notation

$$U_x^\epsilon := U_x \cap B_\epsilon^m(0), \quad D_\epsilon^{\mathcal{M}}(x) := \Pi_x(U_x^\epsilon), \quad \text{for } 0 < \epsilon < r_0. \quad (6.2)$$

In other words, the neighborhood $D_\epsilon^{\mathcal{M}}(x)$ of x consists of all points in the ϵ -neighborhood of x that can be reached from x via the Riemannian exponential map, and it coincides with the ϵ -neighborhood of x with respect to $d_{\mathcal{M}}$ if there exists a unique geodesic connecting every two points in \mathcal{M} , or for points x with the Riemannian distance to the boundary larger than ϵ . Thus, condition (M1) ensures that the Riemannian average of x and any $y \in D_\epsilon^{\mathcal{M}}(x)$ is well-defined.

We are now ready to define the loss functional and the corresponding sampling framework. As the input data to the training of the encoder, for a fixed *sampling radius* $\epsilon \in (0, r_0)$, we consider a finite *set of samples*

$$\mathcal{S}_\epsilon \subset \mathcal{D}_\epsilon := \{(x, y) \in M \times M : y \in D_\epsilon^{\mathcal{M}}(x)\}. \quad (6.3)$$

We assume that, for all the pairs in \mathcal{S}_ϵ , the Riemannian distance $d_{\mathcal{M}}$ and the Riemannian average $\text{av}_{\mathcal{M}}$ are well-defined and available. Then we have the following definition.

Definition 6.1.1 (Discrete sampling loss functional). Given the finite set of samples \mathcal{S}_ϵ , we define the *discrete sampling loss functional* by

$$E^{\mathcal{S}_\epsilon}[\phi] := \frac{1}{|\mathcal{S}_\epsilon|} \sum_{(x,y) \in \mathcal{S}_\epsilon} \left(\gamma(|\partial_{(x,y)}\phi|) + \lambda |\partial_{(x,y)}^2\phi|^2 \right),$$

where $\lambda > 0$, and $\gamma : [0, \infty) \rightarrow [0, \infty)$ is a function with a unique minimum $\gamma(1) = 0$. Furthermore, the *first and second order difference quotients* are defined as

$$\partial_{(x,y)}\phi := \frac{\phi(y) - \phi(x)}{d_{\mathcal{M}}(x, y)}, \quad \partial_{(x,y)}^2\phi := 8 \frac{\text{av}_{\mathbb{R}^l}(\phi(x), \phi(y)) - \phi(\text{av}_{\mathcal{M}}(x, y))}{d_{\mathcal{M}}(x, y)^2}, \quad (6.4)$$

where $\text{av}_{\mathbb{R}^l}(a, b) = (a + b)/2$ denotes the linear average in \mathbb{R}^l .

Notice that the first term in $E^{\mathcal{S}_\epsilon}$ has a strict minimum for $|\partial_{(x,y)}\phi| = 1$ and thus promotes $|\phi(x) - \phi(y)| \approx d_{\mathcal{M}}(x, y)$, and therefore low distortion and an approximate isometry (cf. (I)). In this chapter, we make the following choice

$$\gamma(s) := s^2 + \frac{(1 + c^2)^2}{s^2 + c^2} - 2 - c^2, \quad c > 0. \quad (6.5)$$

Let us emphasize that $\gamma(s)$ is chosen finite in $s = 0$.

The second term in $E^{\mathcal{S}_\epsilon}$ penalizes the deviation of the embedding of intrinsic averages on \mathcal{M} from the extrinsic averages of $\phi(\mathcal{M})$ in \mathbb{R}^l . Note that this does not only penalize bending or any extrinsic curvature of $\phi(\mathcal{M})$ in \mathbb{R}^l , but in addition it also penalizes deviation of the inplane parametrization of $\phi(\mathcal{M})$ from a linear one (cf. (F)).

Let us finally notice that the functional is rigid motion (affine isometry) invariant by construction, *i.e.*, composition of ϕ with a rigid motion does not change the energy. This possibly allows the non-uniqueness of the optimal embedding (up to a rigid motion invariance).

6.1.2 Monte Carlo Limit for Dense Sampling

We now discuss the Monte Carlo limit as the number of samples tends to infinity of the sampling loss functional given by Definition 6.1.1. This limit depends on the strategy used for obtaining the set \mathcal{S}_ϵ . To this end, we consider the pairs $(x, y) \in \mathcal{S}_\epsilon$ as realizations of a random variable (X, Y) taking values in \mathcal{D}_ϵ , with a corresponding probability distribution depending on the sampling strategy. Then, by the strong law of large numbers [Wil91,

Section 12.10], we have the almost sure convergence w.r.t. the same probability to a *continuous sampling loss functional*, i.e.,

$$\lim_{|\mathcal{S}_\epsilon| \rightarrow \infty} E^{\mathcal{S}_\epsilon}[\phi] = \mathbb{E} \left[\gamma(|\partial_{(X,Y)}\phi|) + \lambda |\partial_{(X,Y)}^2\phi|^2 \right] =: \mathcal{E}^\epsilon[\phi].$$

We consider three exemplary ways of sampling (later used in numerical experiments in Section 6.5), leading to three variants \mathcal{E}_i^ϵ , $i = 1, 2, 3$ of \mathcal{E}^ϵ . Here and in the rest of the chapter, V_g is the Riemann–Lebesgue volume measure on \mathcal{M} (cf. (2.9)). For the notation and the proofs of claims in this short excursion to probability theory, we refer the reader to the classical book [Wi91].

(S1) Sampling $x \in \mathcal{M}$ uniformly w.r.t. V_g , and then choosing y , again uniformly, in $D_\epsilon^{\mathcal{M}}(x)$. Thus, we have, for any measurable $A \subset \mathcal{M}$, $\mathbb{P}[X \in A] = \frac{V_g(A)}{V_g(\mathcal{M})}$, and the conditional law of Y is given by $\mathbb{P}[Y \in A|X = x] = \frac{V_g(A \cap D_\epsilon^{\mathcal{M}}(x))}{V_g(D_\epsilon^{\mathcal{M}}(x))}$. The limit is

$$\mathcal{E}_1^\epsilon[\phi] := \frac{1}{V_g(\mathcal{M})} \int_{\mathcal{M}} \frac{1}{V_g(D_\epsilon^{\mathcal{M}}(x))} \int_{D_\epsilon^{\mathcal{M}}(x)} \gamma(|\partial_{(x,y)}\phi|) + \lambda |\partial_{(x,y)}^2\phi|^2 dV_g(y) dV_g(x). \quad (6.6)$$

(S2) Sampling $x \in \mathcal{M}$ uniformly w.r.t. V_g , and then choosing $w \in U_x^\epsilon$ uniformly w.r.t. \mathcal{L} : the m -dimensional Lebesgue measure, and defining $y = \Pi_x(w)$. As above, $\mathbb{P}[X \in A] = \frac{V_g(A)}{V_g(\mathcal{M})}$, while Y is given via the conditional law $\mathbb{P}[Y \in A|X = x] = \frac{\mathcal{L}(\Pi_x^{-1}(A) \cap U_x^\epsilon)}{\mathcal{L}(U_x^\epsilon)}$. This gives

$$\mathcal{E}_2^\epsilon[\phi] := \frac{1}{V_g(\mathcal{M})} \int_{\mathcal{M}} \frac{1}{\mathcal{L}(U_x^\epsilon)} \int_{U_x^\epsilon} \gamma(|\partial_{(x,\Pi_x(w))}\phi|) + \lambda |\partial_{(x,\Pi_x(w))}^2\phi|^2 dw dV_g(x), \quad (6.7)$$

where dw represents the integration w.r.t. \mathcal{L} .

(S3) Sampling pairs $(x, y) \in \mathcal{D}_\epsilon$ uniformly. Then, with \otimes denoting the product of measures, $\mathbb{P}[(X, Y) \in A] = \frac{(V_g \otimes V_g)(A \cap \mathcal{D}_\epsilon)}{(V_g \otimes V_g)(\mathcal{D}_\epsilon)}$, for any measurable $A \subset \mathcal{M} \times \mathcal{M}$, so that

$$\mathcal{E}_3^\epsilon[\phi] := \frac{1}{\int_{\mathcal{M}} \int_{D_\epsilon^{\mathcal{M}}(x)} dV_g(y) dV_g(x)} \int_{\mathcal{M}} \int_{D_\epsilon^{\mathcal{M}}(x)} \gamma(|\partial_{(x,y)}\phi|) + \lambda |\partial_{(x,y)}^2\phi|^2 dV_g(y) dV_g(x). \quad (6.8)$$

We can observe these functionals as instances of a general continuous nonlocal loss functional which depends on the sampling probability density. This is the content of the following definition.

Definition 6.1.2 (Continuous nonlocal loss functional). Given the sampling radius $\epsilon > 0$, we define the *continuous nonlocal loss functional* by

$$\mathcal{E}^\epsilon[\phi] := \int_{\mathcal{M}} \int_{D_\epsilon^{\mathcal{M}}(x)} \left(\gamma(|\partial_{(x,y)}\phi|) + \lambda |\partial_{(x,y)}^2\phi|^2 \right) \rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, y) dV_g(y) dV_g(x), \quad (6.9)$$

where the density $\rho_{\mathcal{M} \times \mathcal{M}}^\epsilon \in L^1(\mathcal{M} \times \mathcal{M})$ satisfies the following conditions:

- (D1) $\int_{\mathcal{M}} \int_{D_\epsilon^{\mathcal{M}}(x)} \rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, y) dV_g(y) dV_g(x) = 1$ for all $\epsilon > 0$,
- (D2) There exist constants $c_\rho, C_\rho > 0$ with $C_\rho \geq \epsilon^m \rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, y) \geq c_\rho$ for a.e. $x \in \mathcal{M}$, $y \in B_\epsilon^{\mathcal{M}}(x)$, and $\epsilon > 0$,
- (D3) $\epsilon^m \rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, \Pi_x(\epsilon w))$ converges pointwise to some function $\rho(x, w)$ as $\epsilon \rightarrow 0$ for a.e. $x \in \mathcal{M}$, $w \in B_1^m(0)$.

While the form (6.9) is obvious for (6.6) and (6.8), for (6.7) it can be obtained by the converse of transformation defining the Riemann–Lebesgue volume measure (cf. (2.9)). This transformation (see below) will be also used for checking the conditions (D1)–(D3). Functionals of the type (6.9) were used in [BBM01] in the context of (fractional) Sobolev spaces, with applications in the study of variational problems, e.g., in [AK09, LPSS15].

As announced, we rewrite the functional (6.9) using normal coordinates around each x (cf. (6.1)), which will later simplify its analysis. To this end, we note that, by the definition of $D_\epsilon^{\mathcal{M}}(x)$ (cf. (6.2)) and the Riemann–Lebesgue measure (2.9), for any measurable function $f_x : \mathcal{M} \rightarrow \mathbb{R}$, we have

$$\int_{D_\epsilon^{\mathcal{M}}(x)} f_x(y) dV_g(y) = \int_{B_1^m(0)} w^\epsilon(x, w) f_x(\Pi_x(\epsilon w)) \sqrt{\det g_{\Pi_x}(\Pi_x(\epsilon w))} \epsilon^m dw, \quad (6.10)$$

where $w^\epsilon(x, w) := \chi_{U_x^\epsilon}(\epsilon w)$ for χ_B the characteristic function of a set B , and g_{Π_x} is the matrix whose elements are the components of g_x with respect to the normal coordinate system around x . Now, taking $f_x(y) = (\gamma(|\partial_{(x,y)}\phi|) + \lambda|\partial_{(x,y)}^2\phi|^2)\rho_{\mathcal{M}\times\mathcal{M}}^\epsilon(x, y)$, we obtain

$$\mathcal{E}^\epsilon[\phi] = \int_{\mathcal{M}} \int_{B_1^m(0)} \left(\gamma(|\partial_{(x, \Pi_x(\epsilon w))}\phi|) + \lambda|\partial_{(x, \Pi_x(\epsilon w))}^2\phi|^2 \right) \rho^\epsilon(x, w) dw dV_g(x), \quad (6.11)$$

for $\rho^\epsilon(x, w) := w^\epsilon(x, w) \sqrt{\det g_{\Pi_x}(\Pi_x(\epsilon w))} \epsilon^m \rho_{\mathcal{M}\times\mathcal{M}}^\epsilon(x, \Pi_x(\epsilon w))$. From the compactness of $\overline{\mathcal{M}}$ and (D2), we deduce that there exist $\tilde{c}_\rho, \tilde{C}_\rho > 0$ with

$$\tilde{C}_\rho \geq \rho^\epsilon(x, w) \geq \tilde{c}_\rho, \quad \text{for all } x \in \mathcal{M}, w \in U_x. \quad (6.12)$$

Furthermore, using the Taylor expansion of $g_{\Pi_x}(\Pi_x(\epsilon w))$ around $\epsilon = 0$, together with the local Lipschitz property of the determinant function, we have [Sak96, Chapter II, Proposition 3.1 and Lemma 3.5]

$$\det g_{\Pi_x}(\Pi_x(\epsilon w)) = 1 - \frac{1}{3} \text{Ric}(\iota_x(\frac{w}{|w|}), \iota_x(\frac{w}{|w|})) |w|^2 \epsilon^2 + \mathcal{O}(|w|^3 \epsilon^3), \quad (6.13)$$

where Ric is the Ricci curvature, which depends on the second derivatives of g_x with respect to x . Then, we conclude that, due to the smoothness of the metric and compactness of the manifold, the second order and the higher order terms are bounded and vanishing as $\epsilon \rightarrow 0$, uniformly in x . Together with assumption (D3) and the fact that $w^\epsilon(x, w)$ converges pointwise to 1, for every $x \in \mathcal{M}$ and $w \in B_1^m(0)$, this implies

$$\rho^\epsilon(x, w) \rightarrow \rho(x, w) \text{ as } \epsilon \rightarrow 0, \quad \text{for a.e. } x \in \mathcal{M}, w \in B_1^m(0). \quad (6.14)$$

Let us now check the conditions (D1)–(D3) for all three functionals (6.6), (6.7), (6.8). The limit in (D3) is the same, the constant function

$$\rho(x, w) = \frac{1}{V_g(\mathcal{M}) \mathcal{L}(B_1^m(0))}.$$

In particular, for (6.6) we have

$$\rho_{\mathcal{M}\times\mathcal{M}}^\epsilon(x, y) = \frac{1}{V_g(\mathcal{M}) V_g(D_\epsilon^{\mathcal{M}}(x))},$$

so that (D1) trivially holds. To confirm (D2) and (D3), we plug $f_x(y) = 1$ in (6.10) and use the same observations as above, to get $\frac{V_g(D_\epsilon^{\mathcal{M}}(x))}{\epsilon^m} \rightarrow \mathcal{L}(B_1^m(0))$ as $\epsilon \rightarrow 0$, for every $x \in \mathcal{M}$, from where the results directly follow. The same steps, together with the Dominated convergence theorem 2.1.3 (where the existence of a dominating function follows from the uniformity of the estimate (6.13)), are used in the case (6.8), where

$$\rho_{\mathcal{M}\times\mathcal{M}}^\epsilon(x, y) = \frac{1}{\int_{\mathcal{M}} V_g(D_\epsilon^{\mathcal{M}}(x)) dV_g(x)}.$$

Finally, for (6.7), by using the transformation formula analogous to (6.10), we have

$$\rho_{\mathcal{M}\times\mathcal{M}}^\epsilon(x, y) = \frac{1}{V_g(\mathcal{M})} \frac{(\det g_{\Pi_x}(y))^{-\frac{1}{2}}}{\int_{D_\epsilon^{\mathcal{M}}(x)} (\det g_{\Pi_x}(y))^{-\frac{1}{2}} dV_g(y)},$$

which straightforwardly confirms (D1). The nominator is bounded and converges pointwise to 1 by (6.13), while we handle the denominator directly, by rewriting it as in (6.10), showing that (D2) and (D3) are satisfied.

6.2 Existence of Minimizer to Nonlocal Energy

In this section, we study the existence of a minimizer for the nonlocal loss functional given by Definition 6.1.2. To this end, we restrict the function space in which we observe the possible minimizers of \mathcal{E}^ϵ , *i.e.*, we introduce ϵ -dependent function spaces $\{\mathcal{F}^\epsilon\}_{\epsilon>0}$, and the extension of the energies $\{\mathcal{E}^\epsilon\}_{\epsilon>0}$, given by

$$\mathcal{E}_{\mathcal{F}^\epsilon}^\epsilon[\phi] := \begin{cases} \mathcal{E}^\epsilon[\phi] & \text{if } \phi \in \mathcal{F}^\epsilon, \\ \infty & \text{else.} \end{cases} \quad (6.15)$$

We choose these spaces in such a way that we will be able to prove the existence of minimizers and their uniform boundedness, independent of ϵ . We first note that, since the energy \mathcal{E}^ϵ is translation invariant, functions in \mathcal{F}^ϵ are without loss of generality assumed to have zero mean. Furthermore, we will need an appropriate closedness of these spaces and their density, as ϵ tends to zero. Finally, to relate difference quotients (6.4) to the first and second order derivatives of the underlying functions, we will assume a sufficient smoothness of the functions in $\{\mathcal{F}^\epsilon\}_{\epsilon>0}$, together with a bound on growth of their derivatives as ϵ tends to zero. In short, with the definitions of the functions spaces introduced in Subsection 2.2.1, we impose the following conditions on $\{\mathcal{F}^\epsilon\}_{\epsilon>0}$:

(H1) *closedness in \dot{L}^2* : For every $0 < \epsilon < r_0$, we have that \mathcal{F}^ϵ is closed as a subset of $\dot{L}^2(\mathcal{M}, \mathbb{R}^l)$, the subset of $L^2(\mathcal{M})$ consisting of functions with zero mean.

(H2) *boundedness*: For every $0 < \epsilon < r_0$, it holds $\mathcal{F}^\epsilon \subset C^{2,1}(\mathcal{M}, \mathbb{R}^l)$, and there exists a constant $C_L \geq 0$ such that

$$\limsup_{\epsilon \rightarrow 0} \sup_{\phi^\epsilon \in \mathcal{F}^\epsilon} \epsilon (L_{\text{grad}}(\phi^\epsilon) + L_{\text{Hess}}(\phi^\epsilon)) \leq C_L.$$

(H3) *density in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$* : For every $\phi \in \dot{H}^2(\mathcal{M}, \mathbb{R}^l)$, there exists a sequence $\{\phi^\epsilon\}_{\epsilon>0}$ with $\phi^\epsilon \in \mathcal{F}^\epsilon$ for every $0 < \epsilon < r_0$ such that $\lim_{\epsilon \rightarrow 0} \|\phi^\epsilon - \phi\|_{\dot{H}^2(\mathcal{M}, \mathbb{R}^l)} = 0$.

Here, $L_{\text{grad}}(\phi)$ and $L_{\text{Hess}}(\phi)$ are the Lipschitz constants of the Riemannian gradient and Hessian of function ϕ , respectively (*cf.* (2.7) and (2.8)). The notation $\epsilon \rightarrow 0$ is short for a sequence $(\epsilon_k)_{k \in \mathbb{N}}$ with $\epsilon_k \rightarrow 0$ for $k \rightarrow \infty$.

These conditions can, in particular, be satisfied if the functions in $\{\mathcal{F}^\epsilon\}_{\epsilon>0}$ are realizations of deep neural networks with sufficiently smooth nonlinear activation functions (*e.g.* (2.13)) and ϵ -dependent bounds on the network architecture (*cf.* Section 2.3) and norm $\|\cdot\|_{\text{NN}}$ (*cf.* (2.15)). Indeed, applying the chain rule, bounds on the growth of derivatives of these realizations can be computed explicitly in terms of the numbers of layers and a priori bounds on the network norm given by (2.15). For this choice, $\{\mathcal{F}^\epsilon\}_{\epsilon>0}$ are not finite-dimensional. However, their closedness was shown in [PRV21, Proposition 3.5], while, for a suitable choice of the activation function, the density property was shown in an already classical article by Hornik *et al.* [HSW90]. Let us point out that, although all these results were actually shown for neural networks defined on Euclidean spaces, they can be transferred to neural networks defined on a smooth m -dimensional manifold (\mathcal{M}, g) immersed in \mathbb{R}^N , such that the metric g is equivalent to the metric induced by the embedding (*cf.* [BGKP19, Section 7]). In Section 6.4, we will provide more information on the specific network architecture that we use, with regard of satisfying (H1)–(H3). For more details on properties of functions spaces consisting of realizations of neural networks, confer to Section 2.3.1.

From an abstract numerical analysis point of view, the above approach can be seen as an instance of two-step approximation of infinite-dimensional optimization problems (see *e.g.* [FP04] for finite element discretizations of phase field approximations to sharp interface problems). First, the model or objective functional is approximated introducing an auxiliary small parameter ϵ (typically a length scale), yielding an optimization problem amenable to a numerical solution. Second, the new objective functional is restricted to a space or set of discretized functions, parametrized by finitely many parameters. To relate the final discrete problem to the original one, both the limit $\epsilon \rightarrow 0$ as well as the limit of increasing expressivity of the discrete function space have to be tackled. The parameter ϵ and the chosen discrete function space usually have to be compatible with each other. In our case, the discrete problem, *i.e.*, minimization for some fixed $\epsilon > 0$ represents the initial problem, while the convergence as $\epsilon \rightarrow 0$ and the obtained limit optimization problem are studied in Section 6.3, where the assumptions (H1)–(H3) will play the key role.

Before proceeding to the proof of the existence of minimizers, we will need two technical results. From now on we abbreviate

$$\bar{w} := \frac{w}{|w|}, \quad \text{for any } w \in \mathbb{R}^m.$$

The first result gives an estimate on the difference of the first and the second order difference quotients and the application of Riemannian gradient and Hessian at corresponding vectors, respectively.

Lemma 6.2.1. *Let $\phi \in C^{2,1}(\mathcal{M}, \mathbb{R}^l)$. Then, for all $x \in \mathcal{M}, w \in U_x$, we have*

$$\left| \partial_{(x, \Pi_x(w))} \phi - g_x(\text{grad } \phi(x), \iota_x \bar{w}) \right| \leq \frac{1}{2} L_{\text{grad}}(\phi) |w|, \quad (6.16)$$

$$\left| \partial_{(x, \Pi_x(w))}^2 \phi - g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w}) \right| \leq \frac{5}{6} L_{\text{Hess}}(\phi) |w|. \quad (6.17)$$

Proof. Given $v \in \iota_x(U_x)$, the first and the second order Taylor expansion of $r \mapsto \phi \circ \exp_x(rv)$ around $r = 0$, evaluated at $r = 1$, yield

$$\phi(\exp_x(v)) = \phi(x) + g(\text{grad } \phi(x), v) + R_1\phi(x, v), \quad (6.18)$$

$$= \phi(x) + g(\text{grad } \phi(x), v) + \frac{1}{2} g(\text{Hess } \phi(x)[v], v) + R_2\phi(x, v). \quad (6.19)$$

Here, $R_1\phi$ and $R_2\phi$ are corresponding remainder terms, such that [Bou23, Corollary 10.54 and Proposition 10.55]

$$|R_1\phi(x, v)| \leq \frac{1}{2} L_{\text{grad}}(\phi) g_x(v, v), \quad |R_2\phi(x, v)| \leq \frac{1}{6} L_{\text{Hess}}(\phi) g_x(v, v)^{\frac{3}{2}}.$$

Choosing $v = \iota_x w$ in (6.18), we get

$$\partial_{(x, \Pi_x(w))} \phi - g_x(\text{grad } \phi(x), \iota_x(\bar{w})) = \frac{R_1\phi(x, \iota_x w)}{|w|},$$

while (6.19), in a combination with the second order expansion of $r \mapsto \phi \circ \exp_x(r \frac{v}{2})$ around $r = 0$, computed at $r = 1$, yields

$$\partial_{(x, \Pi_x(w))}^2 \phi - g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w}) = \frac{4R_2\phi(x, \iota_x w) - 8R_2\phi(x, \iota_x \frac{w}{2})}{|w|^2},$$

which, together with the above estimates for the remainder terms, proves the lemma. \square

The following proposition provides a norm on $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$, equivalent to the standard $H^2(\mathcal{M}, \mathbb{R}^l)$ -norm

$$\|\phi\|_{H^2(\mathcal{M}, \mathbb{R}^l)}^2 := \int_{\mathcal{M}} |\phi|^2 + g(\text{grad } \phi, \text{grad } \phi) + g(\text{Hess } \phi, \text{Hess } \phi) dV_g, \quad (6.20)$$

where we assume the summation convention from Subsection 2.2.1. The norm is obtained by replacing, in this expression, the standard metric based norms on $T_x\mathcal{M}$ and $L_{\text{sym}}(T_x\mathcal{M}, T_x\mathcal{M})$ (the space of symmetric endomorphisms on $T_x\mathcal{M}$), respectively, with the norms uniformly equivalent to them.

Proposition 6.2.2. *Let $V = \frac{1}{r_0} \mathcal{C}_{r_0, \kappa}$ for $\mathcal{C}_{r_0, \kappa}$ defined in (M2). Then,*

$$|W|_{av} := \left(\int_V g(W, \iota_x \bar{w})^2 dw \right)^{\frac{1}{2}}$$

defines a norm on $T_x\mathcal{M}$, which is uniformly equivalent to the standard norm based on the metric. Analogously, a norm on $L_{\text{sym}}(T_x\mathcal{M}, T_x\mathcal{M})$ is given by

$$|A|_{av} := \left(\int_V g(A[\iota_x \bar{w}], \iota_x \bar{w})^2 dw \right)^{\frac{1}{2}},$$

and it is uniformly equivalent to the standard norm on $L_{\text{sym}}(T_x\mathcal{M}, T_x\mathcal{M})$. As a consequence, the norm (with the summation convention)

$$\|\phi\|_{H_{av}^2(\mathcal{M}, \mathbb{R}^l)}^2 := \int_{\mathcal{M}} |\text{grad } \phi(x)|_{av}^2 + |\text{Hess } \phi(x)|_{av}^2 dV_g(x)$$

on $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$ is equivalent to the standard $H^2(\mathcal{M}, \mathbb{R}^l)$ -norm given by (6.20).

Proof. It is obvious that $W \mapsto |W|_{\text{av}}$ is 1-homogeneous and satisfies the triangle inequality, while its positive definiteness will follow from the norm equivalence shown below. By the Cauchy–Schwarz inequality, we have $|W|_{\text{av}}^2 \leq \mathcal{L}(V)g_x(W, W)$. Furthermore, using the isometry of ι_x , we observe

$$|W|_{\text{av}}^2 = \int_V (\iota_x^{-1}W \cdot \bar{w})^2 dw \geq c|\iota_x^{-1}W|^2 = cg_x(W, W),$$

where \cdot is the Euclidean dot product on \mathbb{R}^m , and $c > 0$ is independent of x . This is due to the norm equivalence on \mathbb{R}^m , as the second term above is a norm on \mathbb{R}^m , due to the fact that $v \cdot w = 0$ for all $w \in V$ implies $v = 0$ since V contains an open set. The same arguments apply for the norm on $L_{\text{sym}}(T_x\mathcal{M}, T_x\mathcal{M})$, upon noticing that $\iota_x^{-1} \circ A \circ \iota_x \in L_{\text{sym}}(\mathbb{R}^m, \mathbb{R}^m)$, and that $Bw \cdot w = 0$ for every $w \in V$ implies that B is skew symmetric, which, with $B \in L_{\text{sym}}(\mathbb{R}^m, \mathbb{R}^m)$, implies $B = 0$. From here, it immediately follows that the norm $\|\cdot\|_{H_{\text{av}}^2(\mathcal{M}, \mathbb{R}^l)}$ is equivalent to the standard $H^2(\mathcal{M}, \mathbb{R}^l)$ -seminorm $\int_{\mathcal{M}} g(\text{grad } \phi, \text{grad } \phi) + g(\text{Hess } \phi, \text{Hess } \phi) dV_g$, which is by Poincaré’s inequality 2.2.18 on $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$ equivalent to the standard norm (6.20). \square

The following theorem finally provides the existence of a minimizer for $\mathcal{E}_{\mathcal{F}}^\epsilon$, and, in addition, the uniform boundedness of all minimizers in $H^2(\mathcal{M}, \mathbb{R}^l)$, independent of ϵ .

Theorem 6.2.3 (Existence of minimizers for nonlocal continuous energies). *Let conditions (H1)–(H2) be satisfied. Then, for every $\epsilon \in (0, r_0)$, there exists a minimizer ϕ^ϵ of energy (6.15), and $\|\phi^\epsilon\|_{H^2(\mathcal{M}, \mathbb{R}^l)} \leq C_\mathcal{E}$, for a constant $C_\mathcal{E} > 0$ independent of ϵ .*

Proof. Let us observe the reformulation (6.11) of \mathcal{E}^ϵ . By condition (M2), for every $x \in \mathcal{M}$, we have $\Pi_x(\mathcal{C}_{r_0, \kappa}) \subset B_{r_0}^{\mathcal{M}}(x)$, and thus, $\Pi_x(\mathcal{C}_{\epsilon, \kappa}) \subset D_\epsilon^{\mathcal{M}}(x)$, for all $\epsilon < r_0$. Recalling the definition of $w^\epsilon(x, w) = \chi_{U_\epsilon^\epsilon}(\epsilon w)$, we see

$$w^\epsilon(x, w) = 1 \text{ for all } x \in \mathcal{M}, w \in V = \frac{1}{r_0} \mathcal{C}_{r_0, \kappa}.$$

Using first (6.12) (which comes from (D2)), and then the inequality $|a|^2 \geq \frac{1}{2}|b|^2 - |a - b|^2$, for $a = \partial_{(x, \Pi_x(\epsilon w))} \phi$ and $b = g_x(\text{grad } \phi(x), \iota_x(w))$, we get

$$\begin{aligned} & \int_{\mathcal{M}} \int_{D_\epsilon^{\mathcal{M}}(x)} |\partial_{(x, y)} \phi|^2 \rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, y) dV_g(y) dV_g(x) \\ &= \int_{\mathcal{M}} \int_{B_1^m(0)} |\partial_{(x, \Pi_x(\epsilon w))} \phi|^2 \rho^\epsilon(x, w) dw dV_g(x) \\ &\geq \frac{\tilde{c}_\rho}{2} \int_{\mathcal{M}} \int_V |g_x(\text{grad } \phi(x), \iota_x \bar{w})|^2 dw dV_g(x) - \tilde{c}_\rho \int_{\mathcal{M}} \int_V |g_x(\text{grad } \phi(x), \iota_x \bar{w}) - \partial_{(x, \Pi_x(\epsilon w))} \phi|^2 dw dV_g(x) \\ &\geq \frac{\tilde{c}_\rho}{2} \int_{\mathcal{M}} |\text{grad } \phi|_{\text{av}}^2 dV_g - C\epsilon^2 L_{\text{grad}}(\phi)^2, \end{aligned}$$

for a constant $C > 0$, where in the last step we used Lemma 6.2.1. We analogously obtain

$$\int_{\mathcal{M}} \int_{D_\epsilon^{\mathcal{M}}(x)} |\partial_{(x, y)}^2 \phi|^2 \rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, y) dV_g(y) dV_g(x) \geq \frac{\tilde{c}_\rho}{2} \int_{\mathcal{M}} |\text{Hess } \phi|_{\text{av}}^2 dV_g - C\epsilon^2 L_{\text{Hess}}(\phi)^2.$$

Summing both terms and applying the norm equivalence from Proposition 6.2.2, we get that there exists a constant $C > 0$ such that

$$\|\phi\|_{H^2(\mathcal{M}, \mathbb{R}^l)} \leq C \left(\mathcal{E}^\epsilon[\phi] + \epsilon^2 (L_{\text{grad}}(\phi) + L_{\text{Hess}}(\phi))^2 \right) \leq C (\mathcal{E}^\epsilon[\phi] + C_L^2),$$

for every $\phi \in \mathcal{F}^\epsilon$, where we also applied assumption (H2). Obviously, $\mathcal{E}^\epsilon[0]$ is a finite upper bound for the energy of the minimizer. Thus, there exists a constant $C_\mathcal{E} > 0$, such that, for any fixed $\epsilon > 0$, there exists a sequence $\{\phi_j^\epsilon\}_{j \in \mathbb{N}}$ with $\|\phi_j^\epsilon\|_{H^2(\mathcal{M}, \mathbb{R}^l)} \leq C_\mathcal{E}$ and $\lim_{j \rightarrow \infty} \mathcal{E}^\epsilon[\phi_j^\epsilon] = \inf_{\phi \in \mathcal{F}^\epsilon} \mathcal{E}^\epsilon[\phi]$. Hence, there exists a weakly convergent subsequence (after relabeling) $\phi_j^\epsilon \rightharpoonup \phi^\epsilon$ in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$. By the weak lower semicontinuity of the norm, we have $\|\phi^\epsilon\|_{H^2(\mathcal{M}, \mathbb{R}^l)} \leq C_\mathcal{E}$. Furthermore, by Rellich’s compact embedding 2.2.19, we have $\phi_j^\epsilon \rightarrow \phi^\epsilon$ in $\dot{L}^2(\mathcal{M}, \mathbb{R}^l)$,

so that (H1) implies $\phi^\epsilon \in \mathcal{F}^\epsilon$. Extracting a further subsequence, we even obtain pointwise almost everywhere convergence [Alt06, Theorem 3.22], so that, by Fatou's lemma 2.1.2, we have

$$\liminf_{j \rightarrow \infty} \mathcal{E}^\epsilon[\phi_j^\epsilon] \geq \mathcal{E}^\epsilon[\phi^\epsilon] = \mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon],$$

which finishes the proof. \square

The uniqueness of the minimizer of (6.15) cannot be expected. Indeed, the energy \mathcal{E}^ϵ is invariant under composition from the left with respect to a rigid motion or a reflection, but even apart from this invariance, the nonconvexity of the first integrand in (6.9), which is unavoidable when promoting isometries, may prevent uniqueness. However, whenever $\overline{\mathcal{M}}$ is intrinsically flat and homeomorphic to $\overline{D^m}$: the m -dimensional closed disc (as a representative of compact and (locally) convex subsets of \mathbb{R}^m), there is a unique minimizer of the original nonlocal energy \mathcal{E}^ϵ (cf. (6.11)), up to a rigid transformation, *i.e.*, an affine isomorphism.

Proposition 6.2.4 (The unique optimal embedding of intrinsically flat discs). *If \mathcal{M} is the flat m -dimensional open disc D^m , the unique minimizer (up to a rigid motion) of \mathcal{E}^ϵ is $\phi : \mathcal{M} \ni x \mapsto (x, 0, \dots, 0) \in \mathbb{R}^l$.*

Proof. Since $\mathcal{E}^\epsilon[\phi] = 0$ and $\mathcal{E}^\epsilon \geq 0$, we have that ϕ is a global minimizer. Let $\tilde{\phi}$ be any other minimizer. We claim that one can find a rigid motion $R : \mathbb{R}^l \mapsto \mathbb{R}^l$ such that $\tilde{\phi}(x) = R\phi(x)$ for all $x \in \mathcal{M}$. To determine R , fix points $x_1, \dots, x_m \subset B_\epsilon(x)$ such that $\phi(x_1), \dots, \phi(x_m)$ are linearly independent. Then, since the image of ϕ has dimension m , R is uniquely determined by $R\phi(x_i) = \tilde{\phi}(x_i)$ for all $i = 1, \dots, m$ and $R\phi(0) = R(0) = \tilde{\phi}(0)$. Next, we show that $\tilde{\phi} = R\phi$ on the convex hull of $\{x_1, \dots, x_m\}$. To this end, observe that each z in the convex hull of $\{x_1, \dots, x_m\}$ can be written as the limit of iterated averages of x_1, \dots, x_m , *i.e.*, there exists a sequence $\{z_n\}_{n \in \mathbb{N}}$ converging to z such that $z_n = \frac{z_{n_1} + z_{n_2}}{2}$, where $n_1, n_2 < n$ and z_1, z_2 are among points x_1, \dots, x_m . Indeed, this follows from the fact that the convex hull is the closure of the union of sets $\{S_n\}_{n \in \mathbb{N}}$ where $S_0 = \{x_1, \dots, x_m\}$ and $S_{n+1} = \{\frac{a+b}{2} | a, b \in S_n\}$ for $n \geq 0$ [BV04, Chapter 2]. Since $\partial_{\tilde{x}, \tilde{y}}^2 \tilde{\phi} = 0$ for all $\tilde{x}, \tilde{y} \in B_\epsilon(x)$, we know that

$$\tilde{\phi}(z_n) = \tilde{\phi}\left(\frac{z_{n_1} + z_{n_2}}{2}\right) = \frac{\tilde{\phi}(z_{n_1}) + \tilde{\phi}(z_{n_2})}{2}.$$

Then,

$$\tilde{\phi}(z_3) = \frac{\tilde{\phi}(z_{1,2}) + \tilde{\phi}(z_{1,2})}{2} = \frac{R\phi(z_{1,2}) + R\phi(z_{1,2})}{2} = R\phi\left(\frac{z_{1,2} + z_{1,2}}{2}\right) = R\phi(z_3),$$

and, by induction, we have $\tilde{\phi}(z_n) = R\phi(z_n)$, for all $n \geq 1$. Since $\partial_{\tilde{x}, \tilde{y}}^2 \tilde{\phi} = 0$ for all $\tilde{x}, \tilde{y} \in B_\epsilon(x)$, we have that $\tilde{\phi}$ is continuous on $\overline{B_\epsilon(x)}$, and thus, we have

$$\tilde{\phi}(z) = \tilde{\phi}\left(\lim_{n \rightarrow \infty} z_n\right) = \lim_{n \rightarrow \infty} \tilde{\phi}(z_n) = \lim_{n \rightarrow \infty} R\phi(z_n) = R\phi(z).$$

Proceeding in this fashion, we can cover D^m with convex sets contained in balls of size ϵ , obtaining a rigid motion on each of these sets. Choosing the sets in a way that their intersections are m -dimensional, we conclude that all the rigid motions have to be identical, which finally shows that $\tilde{\phi} = R\phi$ on D^m . \square

Being able to find flat embeddings may actually be quite relevant in applications as it was noticed in [SKTF18] that generative image manifolds have almost no curvature.

We conclude this section with the following remark on the choice of the function spaces for studying the nonlocal energy \mathcal{E}^ϵ .

Remark 6.2.5. One way to approach to the problem of minimization of the energy \mathcal{E}^ϵ , in the form (6.11), would be to study it on the function space naturally associated with this energy: the completion of $C^{2,1}(\mathcal{M}, \mathbb{R}^l)$ under the Hilbert space norm $\|\phi\|_\epsilon^2 := \|\phi\|_{L^2(\mathcal{M}, \mathbb{R}^l)}^2 + |\phi|_\epsilon^2$, with

$$|\phi|_\epsilon^2 := \int_{\mathcal{M}} \int_{D_\epsilon^{\mathcal{M}}(x)} (|\partial_{(x,y)} \phi|^2 + |\partial_{(x,y)}^2 \phi|^2) \rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, y) dV_g(y) dV_g(x).$$

On this space, the energy would be coercive and likely admit a minimizer. Indeed, for our choice of γ (see (6.5)), there exists $C > 0$, independent of ϵ , such that $-C + \frac{1}{C} |\phi|_\epsilon^2 \leq \mathcal{E}^\epsilon[\phi] \leq C + C |\phi|_\epsilon^2$. This function space is closely

related to $H^2(\mathcal{M}, \mathbb{R}^l)$. In [BBM01, Bor07], it is shown, in the case $\mathcal{M} = \mathbb{R}^m$, that for $\phi \in L^2(\mathcal{M}, \mathbb{R}^l)$ satisfying $\liminf_{\epsilon \rightarrow 0} |\phi|_\epsilon < \infty$ it holds $\phi \in H^2(\mathcal{M}, \mathbb{R}^l)$. Furthermore, there are also equicoercivity or compactness results for sequences of $\{\phi^\epsilon\}_{\epsilon > 0}$ with uniformly bounded $\{|\phi^\epsilon|_\epsilon\}_{\epsilon > 0}$ (cf. [BBM01, Corollary 6]), which would ultimately allow us to prove that minimizers of \mathcal{E}^ϵ converge to a minimizer of a local limit functional.

However, for two reasons we did not follow this approach and instead only look for minimizers on a more restricted ϵ -dependent set \mathcal{F}^ϵ of functions which will be realized later via neural networks. Namely, first we would need to translate the theory of [BBM01, Bor07] to functions on manifolds, and second, the compactness results are sensitive to the particular choice of the sampling weight $\rho_{\mathcal{M} \times \mathcal{M}}^\epsilon$. In particular, while equicoercivity will probably hold for our examples (6.6)–(6.8), one might also think of sampling weights $\rho_{\mathcal{M} \times \mathcal{M}}^\epsilon(x, \cdot)$ that are supported on an annulus $D_\epsilon^{\mathcal{M}}(x) \setminus B_\epsilon^{\mathcal{M}}(x)$ (in our numerical experiments we will make use of this idea to improve numerical stability), in which case equicoercivity might potentially be lost (cf. [BBM01, Counterexample 1]), while assumptions (H1)–(H3) still hold. In contrast, all our other results readily extend to that setting.

6.3 Convergence for Vanishing Sampling Radius

In this section, we study convergence of the functional (6.15) as the sampling radius ϵ tends to 0. We will prove Mosco-convergence (cf. Definition 2.4.1) towards the local loss functional given by

$$\mathcal{E}[\phi] := \int_{\mathcal{M}} \Gamma(\text{grad } \phi(x) + \lambda \|\text{Hess } \phi(x)\|^2) dV_g(x), \quad \text{where} \quad (6.21)$$

$$\Gamma(W) := \int_{B_1^m(0)} \gamma(|g_x(W, \iota_x \bar{w})|) \rho(x, w) dw, \quad \text{and} \quad (6.22)$$

$$\|A\|^2 := \int_{B_1^m(0)} |g_x(A[\iota_x \bar{w}], \iota_x \bar{w})|^2 \rho(x, w) dw, \quad (6.23)$$

for all $W \in (T_x \mathcal{M})^l$, $A \in L_{\text{sym}}(T_x \mathcal{M}, (T_x \mathcal{M})^l)$, and ρ from (D3).

This loss functional promotes low distortion and low bending embeddings. Indeed,

- The isometry or low distortion objective **(I)** is being promoted since, by the definition of Riemannian gradient (2.5), an isometric embedding $\phi : \mathcal{M} \rightarrow \mathbb{R}^l$ is characterized by $\text{grad } \phi(x)$ being orthogonal in any point $x \in \mathcal{M}$ or equivalently $|g_x(\text{grad } \phi(x), \iota_x w)| = 1$, for all $w \in S^{m-1} := \partial B_1^m(0)$. Therefore, by the choice of γ (cf. (6.5)), the first term of (6.21) penalizes deviation from an isometric embedding, as it is zero if ϕ is locally isometric and strictly positive otherwise.
- The extrinsic bending of the embedding $\phi : \mathcal{M} \rightarrow \mathbb{R}^l$ is manifested by a non-vanishing Riemannian Hessian (cf. (2.6)). Since the term (6.23) defines a squared norm on the space $L_{\text{sym}}(T_x \mathcal{M}, (T_x \mathcal{M})^l)$ by Proposition 6.2.2, the second term in (6.21) penalizes extrinsic bending and thus promotes flatness and low bending objective **(F)**.

Theorem 6.3.1 (Mosco-convergence of the nonlocal energy functionals). *Let (H2)–(H3) be satisfied with $C_L = 0$. Then, the nonlocal energy functionals $\{\mathcal{E}_{\mathcal{F}}^\epsilon\}_{\epsilon > 0}$, given by (6.15), converge as $\epsilon \rightarrow 0$ to the continuous local energy \mathcal{E} given by (6.21), in the sense of Mosco in the $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$ -topology, i.e.,*

- (i) *for every sequence $\{\phi^\epsilon\}_{\epsilon > 0} \subset H^2(\mathcal{M}, \mathbb{R}^l)$ converging weakly to ϕ in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$ as $\epsilon \rightarrow 0$, it holds $\liminf_{\epsilon \rightarrow 0} \mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon] \geq \mathcal{E}[\phi]$ (“lim inf-inequality”),*
- (ii) *for every $\phi \in \dot{H}^2(\mathcal{M}, \mathbb{R}^l)$, there exists a sequence $\{\phi^\epsilon\}_{\epsilon > 0}$ converging strongly to ϕ in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$ as $\epsilon \rightarrow 0$ with $\limsup_{\epsilon \rightarrow 0} \mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon] \leq \mathcal{E}[\phi]$ (“existence of a recovery sequence and lim sup-inequality”).*

Proof. As in Theorems 4.4.1 and 5.4.1, we will check the conditions separately. Unlike for those proofs, the only bigger task is proving the inequalities, so we are not splitting the proofs into further steps. For ease of notation, below we will sometimes use an expression of the form $f(x, w)$ to actually indicate the map $\mathcal{M} \times B_1^m(0) \ni (x, w) \mapsto f(x, w)$. It will always be clear from the context when this meaning is intended.

Proof of the lim inf-inequality

Let $\phi^\epsilon \rightharpoonup \phi$ in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$. Putting aside the trivial case of $\liminf_{\epsilon \rightarrow 0} \mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon] = +\infty$, we may assume that $\phi^\epsilon \in \mathcal{F}^\epsilon$, for every $\epsilon > 0$ along a subsequence. We first estimate

$$\begin{aligned} & \limsup_{\epsilon \rightarrow 0} \left\| \partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon - g_x(\text{grad } \phi(x), \iota_x \bar{w}) \right\|_{L^2(\mathcal{M} \times B_1^m(0))} \\ & \leq \limsup_{\epsilon \rightarrow 0} \left\| \partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon - g_x(\text{grad } \phi^\epsilon(x), \iota_x \bar{w}) \right\|_{L^2(\mathcal{M} \times B_1^m(0))} \\ & \quad + \limsup_{\epsilon \rightarrow 0} \left\| g_x(\text{grad } \phi^\epsilon(x), \iota_x \bar{w}) - g_x(\text{grad } \phi(x), \iota_x \bar{w}) \right\|_{L^2(\mathcal{M} \times B_1^m(0))}. \end{aligned}$$

By (6.16), the first summand can be estimated by

$$\left\| \partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon - g_x(\text{grad } \phi^\epsilon(x), \iota_x \bar{w}) \right\|_{L^2(\mathcal{M} \times B_1^m(0))} \leq C\epsilon L_{\text{grad}}(\phi^\epsilon),$$

which, by (H2), converges to 0. To show the same for the second summand, we first observe that $\text{grad } \phi^\epsilon \rightarrow \text{grad } \phi$ in $L^2(\mathcal{M}, (T\mathcal{M})^l)$, by the compact embedding of $H^2(\mathcal{M}, \mathbb{R}^l)$ into $H^1(\mathcal{M}, \mathbb{R}^l)$, which is due to Theorem 2.2.19. Then, we can conclude that

$$g_x(\text{grad } \phi^\epsilon(x), \iota_x \bar{w}) \rightarrow g_x(\text{grad } \phi(x), \iota_x \bar{w}) \quad \text{in } L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l),$$

as this represents a bounded linear function from $L^2(\mathcal{M}, (T\mathcal{M})^l)$ to $L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)$. Indeed, for every $\phi \in H^1(\mathcal{M}, \mathbb{R}^l)$, by Cauchy–Schwarz inequality, we have

$$\|g_x(\text{grad } \phi(x), \iota_x \bar{w})\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)} \leq \sqrt{\mathcal{L}(B_1^m(0))} \|\text{grad } \phi\|_{L^2(\mathcal{M}, (T\mathcal{M})^l)}.$$

Altogether, we have

$$\partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon \rightarrow g_x(\text{grad } \phi(x), \iota_x \bar{w}) \quad \text{in } L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l). \quad (6.24)$$

We can now pass to a pointwise convergent subsequence (without relabeling) [Alt06, Theorem 3.22], use (6.14), and continuity of the distortion function γ given by (6.5), to conclude that, for a.e. $x \in \mathcal{M}$, $w \in B_1^m(0)$, we have

$$\gamma(|\partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon|) \rho^\epsilon(x, w) \rightarrow \gamma(|g_x(\text{grad } \phi(x), \iota_x \bar{w})|) \rho(x, w).$$

Applying Fatou's lemma 2.1.2, we obtain

$$\begin{aligned} & \liminf_{\epsilon \rightarrow 0} \int_{\mathcal{M}} \int_{B_1^m(0)} \gamma(|\partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon|) \rho^\epsilon(x, w) \, dw dV_g(x) \\ & \geq \int_{\mathcal{M}} \int_{B_1^m(0)} \gamma(|g_x(\text{grad } \phi(x), \iota_x \bar{w})|) \rho(x, w) \, dw dV_g(x). \end{aligned} \quad (6.25)$$

To handle the second order term in (6.11), we estimate

$$\begin{aligned} & \liminf_{\epsilon \rightarrow 0} \left\| \partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon \sqrt{\rho^\epsilon(x, w)} \right\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)} \\ & \geq \liminf_{\epsilon \rightarrow 0} \left\| g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \sqrt{\rho^\epsilon(x, w)} \right\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)} \\ & \quad - \limsup_{\epsilon \rightarrow 0} \left\| \left(\partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon - g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \right) \sqrt{\rho^\epsilon(x, w)} \right\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)}. \end{aligned}$$

By (6.17), the last term can be estimated as

$$\left\| \left(\partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon - g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \right) \sqrt{\rho^\epsilon(x, w)} \right\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)} \leq C\epsilon L_{\text{Hess}}(\phi^\epsilon), \quad (6.26)$$

which vanishes in the limit by (H2). To bound the first term, we first observe that, analogously to the above,

$$g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \rightarrow g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w}) \quad \text{in } L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l),$$

since this represents a bounded linear function on $L^2(\mathcal{M}, L_{\text{sym}}(T\mathcal{M}, (T\mathcal{M})^l))$. Furthermore, using the uniform boundedness (6.12) and pointwise convergence (6.14) of ρ^ϵ , and the Dominated convergence theorem 2.1.3

$$g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \sqrt{\rho^\epsilon(x, w)} \rightarrow g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w}) \sqrt{\rho(x, w)} \quad \text{in } L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l). \quad (6.27)$$

Using the weak lower semicontinuity of the $L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)$ -norm, this implies

$$\liminf_{\epsilon \rightarrow 0} \|\partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon \sqrt{\rho^\epsilon(x, w)}\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)} \geq \|g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w}) \sqrt{\rho(x, w)}\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)}.$$

We thus proved

$$\liminf_{\epsilon \rightarrow 0} \int_{\mathcal{M}} \int_{B_1^m(0)} \lambda |\partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon|^2 \rho^\epsilon(x, w) \, dw dV_g(x) \geq \int_{\mathcal{M}} \int_{B_1^m(0)} \lambda |g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w})|^2 \rho(x, w) \, dw dV_g(x),$$

that, together with (6.25), proves the desired inequality.

Proof of the existence of a recovery sequence and lim sup-inequality

By the density assumption (H3), for every $\phi \in \dot{H}^2(\mathcal{M}, \mathbb{R}^l)$, there exists a sequence $\{\phi^\epsilon\}_{\epsilon > 0}$ such that $\phi^\epsilon \in \mathcal{F}^\epsilon$ and $\phi^\epsilon \rightarrow \phi$ in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$ as $\epsilon \rightarrow 0$. We can now repeat the same arguments as for (6.24), to show

$$\partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon \rightarrow g_x(\text{grad } \phi(x), \iota_x \bar{w}) \quad \text{in } L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l).$$

Then, using the uniform boundedness (6.12) and pointwise convergence (6.14) of ρ^ϵ , together with the Dominated convergence theorem 2.1.3, we get

$$\partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon \sqrt{\rho^\epsilon(x, w)} \rightarrow g_x(\text{grad } \phi(x), \iota_x \bar{w}) \sqrt{\rho(x, w)} \quad \text{in } L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l),$$

and thus pointwise a.e. on a (not relabeled) subsequence. Taking into account the splitting of $\gamma(s)$ (cf. (6.5)) into $|s|^2$ and the uniformly bounded and continuous part $\gamma(s) - |s|^2$, we use the convergence in norm for the first term, and the Dominated convergence theorem 2.1.3 for the second term, to get

$$\limsup_{\epsilon \rightarrow 0} \int_{\mathcal{M}} \int_{B_1^m(0)} \gamma(|\partial_{(x, \Pi_x(\epsilon w))} \phi^\epsilon|) \rho^\epsilon(x, w) \, dw dV_g(x) = \int_{\mathcal{M}} \int_{B_1^m(0)} \gamma(|g_x(\text{grad } \phi(x), \iota_x \bar{w})|) \rho(x, w) \, dw dV_g(x).$$

For the second order term, we write

$$\begin{aligned} & \limsup_{\epsilon \rightarrow 0} \left\| \partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon \sqrt{\rho^\epsilon(x, w)} \right\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)} \\ & \leq \limsup_{\epsilon \rightarrow 0} \left\| g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \sqrt{\rho^\epsilon(x, w)} \right\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)} \\ & \quad + \limsup_{\epsilon \rightarrow 0} \left\| \left(\partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon - g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \right) \sqrt{\rho^\epsilon(x, w)} \right\|_{L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l)}, \end{aligned}$$

where the last term vanishes analogously to (6.26), while for the first term, analogously to (6.27), we have

$$g_x(\text{Hess } \phi^\epsilon(x)[\iota_x \bar{w}], \iota_x \bar{w}) \sqrt{\rho^\epsilon(x, w)} \rightarrow g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w}) \sqrt{\rho(x, w)} \quad \text{in } L^2(\mathcal{M} \times B_1^m(0), \mathbb{R}^l).$$

Thus, we finally have

$$\limsup_{\epsilon \rightarrow 0} \int_{\mathcal{M}} \int_{B_1^m(0)} \lambda |\partial_{(x, \Pi_x(\epsilon w))}^2 \phi^\epsilon|^2 \rho^\epsilon(x, w) \, dw dV_g(x) \leq \int_{\mathcal{M}} \int_{B_1^m(0)} \lambda |g_x(\text{Hess } \phi(x)[\iota_x \bar{w}], \iota_x \bar{w})|^2 \rho(x, w) \, dw dV_g(x),$$

which, together with the above estimate for the isometry promoting term, proves the lim sup-inequality, and finishes the proof of the theorem. \square

The existence of a minimizer for (6.21) can be established using the direct method of the calculus of variations. Indeed, the zero function provides an upper bound for the value of the energy and Proposition 6.2.2 implies the uniform boundedness of an approximation sequence in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$. Finally, reflexivity and weak lower semicontinuity of the energy provides the desired result. Notice that, as for $\mathcal{E}_{\mathcal{F}}^\epsilon$, the minimizer (modulo a rigid motion or reflection) is in general not unique due to the isometry promoting term. The following result establishes this minimizer as the limit of minimizers of nonlocal energies $\{\mathcal{E}_{\mathcal{F}}^\epsilon\}_{\epsilon > 0}$ as $\epsilon \rightarrow 0$, based on the uniform boundedness of the minimizers in $H^2(\mathcal{M}, \mathbb{R}^l)$, shown in Theorem 6.2.3.

Theorem 6.3.2 (Convergence of minimizers of nonlocal energies). *Let (H1)–(H3) hold with $C_L = 0$. Then, for every sequence $\{\phi^\epsilon\}_{\epsilon>0}$ of minimizers of $\{\mathcal{E}_{\mathcal{F}}^\epsilon\}_{\epsilon>0}$, there exists a subsequence converging weakly in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$ to a minimizer ϕ of the energy \mathcal{E} . Furthermore, the corresponding function values $\mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon]$ converge to $\mathcal{E}[\phi]$.*

Proof. Let $\{\phi^\epsilon\}_{\epsilon>0}$ be a sequence of minimizers of the energies $\mathcal{E}_{\mathcal{F}}^\epsilon$. By Theorem 6.2.3, we have that $\{\phi^\epsilon\}_{\epsilon>0}$ is uniformly bounded in $\dot{H}^2(\mathcal{M}, \mathbb{R}^l)$. By reflexivity, there exists a weakly convergent subsequence (with the same label) such that $\phi^\epsilon \rightharpoonup \phi$ in $H^2(\mathcal{M}, \mathbb{R}^l)$ as $\epsilon \rightarrow 0$. By the lim inf-inequality part of Theorem 6.3.1, we have $\liminf_{\epsilon \rightarrow 0} \mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon] \geq \mathcal{E}[\phi]$. Let us suppose that there exists $\tilde{\phi} \in H^2(\mathcal{M}, \mathbb{R}^l)$ such that $\mathcal{E}[\phi] > \mathcal{E}[\tilde{\phi}]$. Without loss of generality we can suppose that $\tilde{\phi} \in \dot{H}^2(\mathcal{M}, \mathbb{R}^l)$. Then, by the lim sup-part of Theorem 6.3.1, there exists a sequence $\{\tilde{\phi}^\epsilon\}_{\epsilon>0}$ such that $\tilde{\phi}^\epsilon \in \mathcal{F}^\epsilon$ for every $\epsilon > 0$ and $\limsup_{\epsilon \rightarrow \infty} \mathcal{E}_{\mathcal{F}}^\epsilon[\tilde{\phi}^\epsilon] \leq \mathcal{E}[\tilde{\phi}]$. However, $\mathcal{E}[\phi] \leq \liminf_{\epsilon \rightarrow 0} \mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon] \leq \limsup_{\epsilon \rightarrow 0} \mathcal{E}_{\mathcal{F}}^\epsilon[\tilde{\phi}^\epsilon] \leq \mathcal{E}[\tilde{\phi}]$, which proves contradiction of the assumption. Thus, ϕ is a minimizer of \mathcal{E} and, taking $\tilde{\phi} = \phi$ in the above inequality chain, we have $\lim_{\epsilon \rightarrow 0} \mathcal{E}_{\mathcal{F}}^\epsilon[\phi^\epsilon] = \mathcal{E}[\phi]$. \square

6.4 Autoencoder Architecture and Training Procedure

In this section, we describe the architecture of deep neural networks (*cf.* Section 2.3 for the basic definitions and notation), used for the representation of encoder and decoder functions. Furthermore, we will comment on the training procedure, used for the optimization of weights of encoders and decoders, in order to obtain the corresponding identification of high-dimensional data with low-dimensional latent manifolds.

The architecture of deep convolutional neural networks representing autoencoders is motivated by the approach in [BRRG18]. The changes are made to allow using larger input images and smaller latent space dimension. Both encoder and decoder consist of blocks of two convolutions, each followed by a nonlinear activation function, and a subsequent average pooling (encoder) or upsampling (decoder). We determine the number of blocks such that the final output image of the encoder consists of 4×4 pixels, resulting in a latent code of size 16. The full architectures are shown in Table 6.1.

As the activation function, we used LeakyReLU (*cf.* (2.12)). Although this activation function prevents us from satisfying the assumption (H2), we noticed that replacing LeakyReLU by a smooth approximation (2.13) did not have much effect in practice. Networks with (leaky) ReLU nonlinearity also provide nice approximation properties with a provable growth rate of the network architecture, though only in Sobolev spaces of order $0 \leq s \leq 1$ [GKP20]. Furthermore, in view of (H3), the architecture (number of layers and weights, bound on weight growth or penalty on weights) should depend on the sampling radius ϵ . However, since our experiments are performed with ϵ in a fixed range, we always keep the same architecture.

The training data will consist of triplets of images (two images and their Riemannian average), together with the Riemannian distance of the extremal images: $(x, y, \text{av}_{\mathcal{M}}(x, y), d_{\mathcal{M}}(x, y))$ with $x, y \in \mathcal{S}_\epsilon \subset \mathcal{M} \times \mathcal{M}$ (*cf.* (6.3)). This input allows us to compute the ingredients $\partial_{(x,y)}\phi$ and $\partial_{(x,y)}^2\phi$ of the discrete sampling loss functional $E^{\mathcal{S}_\epsilon}[\phi]$, defined by Definition 6.1.1. Our method allows us to train the encoder map separately by minimizing

$$E^{\mathcal{S}_\epsilon}[\phi_\theta] := \frac{1}{|\mathcal{S}_\epsilon|} \sum_{(x,y) \in \mathcal{S}_\epsilon} \left(\gamma(|\partial_{(x,y)}\phi_\theta|) + \lambda \|\partial_{(x,y)}^2\phi_\theta\|_{L^2}^2 \right),$$

where θ are the weights of the encoder, and to train the decoder map ψ_ξ subsequently, by minimizing, for fixed ϕ_θ , the reconstruction loss

$$R[\phi_\theta, \psi_\xi] := \frac{1}{2|\mathcal{S}_\epsilon|} \sum_{(x,y) \in \mathcal{S}_\epsilon} \|\psi_\xi(\phi_\theta(x)) - x\|_{L^2}^2 + \|\psi_\xi(\phi_\theta(y)) - y\|_{L^2}^2$$

where ξ are the weights of the decoder. Here, $\|\cdot\|_{L^2}$ is the discrete L^2 -norm on pixel images. Alternatively, encoder and decoder can be trained simultaneously by minimizing the loss functional, for some $\kappa > 0$,

$$E^{\mathcal{S}_\epsilon}[\phi_\theta] + \kappa R[\phi_\theta, \psi_\xi].$$

For training, we use the Adam optimization algorithm [KB15], with the learning rate $lr = 0.0001$ and (default) values of momenta weights $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\varepsilon = 10^{-8}$. We further employ a so-called weight decay: penalty on the squares of the weights, with a scale $\delta = 10^{-5}$. The pseudocode is given in Algorithm 6.

type	parameters	output size	type	parameters	output size
Conv2d	1/1/1/16	$66 \times 66 \times 16$	$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/128	$4 \times 4 \times 128$
$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/16	$66 \times 66 \times 16$	Upsample	2 / nearest	$8 \times 8 \times 128$
AvgPool2d	2/2	$33 \times 33 \times 16$	$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/64	$8 \times 8 \times 64$
$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/32	$33 \times 33 \times 32$	Upsample	2 / nearest	$16 \times 16 \times 64$
AvgPool2d	2/2	$16 \times 16 \times 32$	$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/32	$16 \times 16 \times 32$
$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/64	$16 \times 16 \times 64$	Upsample	2 / nearest	$32 \times 32 \times 32$
AvgPool2d	2/2	$8 \times 8 \times 64$	$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/16	$32 \times 32 \times 16$
$2 \times (\text{Conv2d} + \text{LeakyReLU}_{0.01})$	3/1/1/128	$8 \times 8 \times 128$	Upsample	2 / nearest	$64 \times 64 \times 16$
AvgPool2d	2/2	$4 \times 4 \times 128$	Conv2d+LeakyReLU _{0.01}	3/1/1/16	$64 \times 64 \times 16$
Conv2d+LeakyReLU _{0.01}	3/1/1/256	$4 \times 4 \times 256$	Conv2d	3/1/1/1	$64 \times 64 \times 1$
Conv2d	3/1/1/1	$4 \times 4 \times 1$			

Table 6.1: Encoder (left) and decoder (right) architectures for 64×64 grayscale images. Parameters for convolutional layers (Conv2d) are filter size / stride / padding / output channels, and for average pooling layers (AvgPool2d) filter size / stride, equal in each direction. Parameters for upsampling layers (Upsample) are scale and mode. The names of the operations are as in the PyTorch library [PGM⁺19]. Output sizes are given with channels in the last dimension. The architecture for color images with the input size $128 \times 128 \times 3$ is similar, except that there is one more block, thus the maximum number of output channels of the encoder is 512.

Algorithm 6: Adam Algorithm for minimization of energy E with respect to the weights θ .

- 1 Choose $lr, \beta_1, \beta_2, \delta, \varepsilon > 0$;
 - 2 Initialize θ_0, m_0, v_0 ;
 - 3 **for** $i = 1, \dots$ **do**
 - 4 $g_i = D_\theta E(\theta_{i-1}) + \delta \theta_{i-1}$;
 - 5 $m_i = \beta_1 m_{i-1} + (1 - \beta_2) g_i$;
 - 6 $v_i = \beta_2 v_{i-1} + (1 - \beta_2) g_i^2$;
 - 7 $\hat{m}_i = \frac{m_i}{1 - \beta_1}$;
 - 8 $\hat{v}_i = \frac{v_i}{1 - \beta_2}$;
 - 9 $\theta_i = \theta_{i-1} - \frac{lr \hat{m}_i}{\sqrt{\hat{v}_i + \varepsilon}}$;
-

We use Kaiming initialization [HZRS15], *i.e.*, all convolutional weights are initialized as zero-mean Gaussian random variables with standard deviation $\sqrt{2}/\sqrt{\text{fan_in}(1 + 0.01^2)}$ for fan_in the number of input channels of the layer times the number of entries in the filter, and all biases are initialized with zeros.

The computation of the derivative in line 4 is not performed on the entire input dataset, but only on stochastically chosen batches. Once the entire dataset is processed, we say we trained for one epoch. Even though in our applications we generate random samples on the fly, we simulate epochs of size around 10000, split into batches of size 128, to make the training graphs look more familiar and less noisy. The training is finished once a chosen (small) threshold for the value of the energy we optimize on the training or the test set is achieved.

6.5 Numerical Results

In this section, we present the numerical results which support our theoretical investigations from the earlier chapters. To this end, we first create image manifolds where the hidden manifold is explicitly known and $\text{av}_{\mathcal{M}}$ and $d_{\mathcal{M}}$ can be explicitly computed. We then apply the training procedures discussed above to obtain a low bending and low distortion embedding into a latent space. We visualize these embeddings and study their regularity and smoothness. Furthermore, we observe the quality of image interpolation obtained by an application of the decoder mapping to the linear interpolation in the latent space. We point out that all of the experiments were performed

by J. Braunsmann, in whose thesis (yet to appear) one can find more numerical results with the corresponding discussions.

6.5.1 Input Manifolds

As already announced, we create image manifolds where the underlying low-dimensional manifold is explicitly known. This approach, already taken in [DG05], allows a fast on the fly generation of random samples. Furthermore, it allows us to study the obtained results by visually comparing them with the expected result: the given input manifold. Contrary to this approach, on the given (high-dimensional) image manifold, we could use approximations of Riemannian distance and average with means of the flow of diffeomorphism or metamorphosis model (cf. Subsections 3.2.1 and 3.2.2), optimal transport [PC⁺19] or other general shape manifolds with application in imaging. However, this would take significantly more time for the computations of input data. In addition, these approaches mostly correspond to an infinite-dimensional manifold structure, and thus, it is an open question whether to expect a perfect recovery of the underlying low-dimensional structure.

We consider image data that implicitly represent four different manifolds:

Sundial Shadows (S). To create this dataset, inspired by [OYHO20], we consider the shadows of a vertical rod on a plane from a lightsource positioned on the upper hemisphere $S^2 \cap \{x_3 \geq 0\}$. Contrary to [OYHO20], we do not render these images with a 3D engine, but approximate the shadows by Gaussians (cf. Figure 6.2): a point $x \in S^2$ is first mapped onto the plane by drawing the line through x and the rod tip, intersecting the plane at some $x_p \in \mathbb{R}^2$. We then use a Gaussian function with variance $|x_p|^2$ in direction x_p and a fixed small variance in the orthogonal direction, centered at $x_p/2$. These images form our input manifold $\overline{\mathcal{M}}$ with the upper hemisphere as its underlying hidden manifold. The metric is induced from S^2 , so that the distance on \mathcal{M} is the geodesic distance on S^2 : $d_{\mathcal{M}}(x, y) := \arccos(x^T y)$. We use a resolution of 64×64 pixels and sampling strategy (S3).

Rotating Object (R). We generate images forming the manifold \mathcal{M} by projecting an arbitrarily rotated three-dimensional object (here a toy cow model¹) to the viewing plane and use Pytorch3D [RRN⁺20] to render the images for the training (cf. Figure 6.4a). As distance on \mathcal{M} , which has the group of rotations $SO(3)$ as the hidden manifold and thus has no boundary, we use the induced distance from $SO(3)$, computed from the quaternions representation $x \mapsto q_x$, determined by the angle and the axis of rotation, i.e., given points $x, y \in \mathcal{M}$ we define $d_{\mathcal{M}}(x, y) := \arccos(|q_x \cdot q_y|)$ [Huy09]. We consider RGB color images with resolution 128×128 pixels and sampling strategy (S1).

Arc rotating around two specified axes (A). To generate this image manifold, we use Blender² to render bent and capped tubes, colored using a transition from purple inside to yellow outside, representing rotations of the great circle $[0, 1] \ni t \mapsto (\cos(t\pi), \sin(t\pi), 0)$. Rotations are determined by the two angles $\alpha = \pi x_1$ and $\beta = \pi(2x_2 + \frac{1}{2})$ (cf. Figure 6.5a), where (x_1, x_2) is a point in the Klein bottle $[0, 1]^2 / \sim$, where \sim represents the glueing operation on the boundary of $[0, 1]^2$ defining the Klein bottle (the hidden closed manifold for this dataset), i.e., we identify $(x_1, 0)$ and $(x_1, 1)$ as well as $(0, x_2)$ and $(1, 1 - x_2)$ (cf. [GHL90]). To determine the geodesic distance between two images, we use the Euclidean distance between points generating these images. Taking into consideration the relation \sim , we have $d_{\mathcal{M}}(x, y) := \min\{|\bar{x} - y| : \bar{x} \in \{(x_1, x_2), (x_1, x_2 - 1), (x_1, x_2 + 1), (x_1 + 1, 1 - x_2), (x_1 + 1, 2 - x_2), (x_1 + 1, -x_2), (x_1 - 1, 1 - x_2), (x_1 - 1, 2 - x_2), (x_1 - 1, -x_2)\}\}$. We use sampling strategy (S3) on $[0, 1]^2 / \sim$ and color images with the resolution 128×128 .

Ellipses (E). This image manifold is created by considering rotations and translations of an ellipse with a fixed aspect ratio $E(\alpha, c) := \{z \in \mathbb{R}^2 : (z - c)^T A(\alpha)^{-1}(z - c) \leq 1\}$ (cf. [Gha14, Section 2.5.3]), with

$$A(\alpha) := \frac{1}{10} \begin{pmatrix} 1 + 2 \cos(\alpha)^2 & 2 \cos(\alpha) \sin(\alpha) \\ 2 \cos(\alpha) \sin(\alpha) & 1 + 2 \sin(\alpha)^2 \end{pmatrix}$$

with parameters (α, c) in $[0, \pi) \times [-1, 1]^2$ and periodic identification of 0 and π . The hidden manifold is then $S^1 \times [-1, 1]^2$, with the corresponding distance $d_{\mathcal{M}}((\alpha_1, c_1), (\alpha_2, c_2))^2 := \min(|\alpha_1 - \alpha_2|, |\alpha_1 + \pi - \alpha_2|, |\alpha_1 - \pi - \alpha_2|)^2 + |c_1 - c_2|^2$. We used the sampling strategy (S1) for obtaining the corresponding parameters. The ellipses are discretized as 64×64 images by evaluating, on a grid $[-1, 1]^2$, either the characteristic function of $E(\alpha, c)$ or the smoothed variant (cf. Figure 6.6)

$$f_{\alpha, c, k}(z) := (1 + \exp(-k(1 - (z - c)^T A^{-1}(\alpha)(z - c))))^{-1}, \quad k > 0. \quad (6.28)$$

¹<https://www.cs.cmu.edu/kmcrane/Projects/ModelRepository/>

²<https://www.blender.org/>

6.5.2 Visualization of the Embeddings

In this section, we present the results of the numerical experiments. The obstacle to the visualization of the results is the fact that the latent space dimension is $l = 16$, so we need to find a way to represent the obtained latent vectors in three dimensions in a reasonable way. To this end, we use the *principal component analysis (PCA)* [Jol02]. This approach is based on solving the problem of finding the best d -dimensional representation of l -dimensional data $\{x_1, \dots, x_n\} \in \mathbb{R}^l$, by observing the minimization problem

$$\min_{\mu \in \mathbb{R}^l, \lambda_i \in \mathbb{R}^d, V_d \in \mathbb{R}^{l,d}} \sum_{i=1}^n \|x_i - (\mu + V_d \lambda_i)\|^2, \text{ subject to } V_d^T V_d = \mathbf{1}_d.$$

A solution to this problem is given in terms of $\mu = \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$, $\lambda_i = V_d^T (x_i - \bar{x})$, $i = 1, \dots, n$, and V_d consisting of d orthogonalized eigenvectors of $X^T X$ corresponding to d largest eigenvalues σ_i^2 , $i = 1, \dots, d$, where $X \in \mathbb{R}^{n,l}$ is the matrix with rows $(x_i - \bar{x})^T$, $i = 1, \dots, n$. Besides giving us the best d -dimensional representation, PCA is also used to determine how large is d supposed to be in order to have a proper low-dimensional representation of the given data. This is determined by measuring the percentage of the *explained variance* (name coming from statistical perspective of PCA):

$$\frac{\sum_{i=1}^d \sigma_i^2}{\sum_{i=1}^l \sigma_i^2},$$

where σ_i^2 , $i = 1, \dots, l$ are all eigenvalues of $X^T X$, ordered in non-increasing order. We will refer to the number of dimensions necessary for reaching a threshold of 99% of explained variance as the *number of relevant dimensions*.

The columns of V_d are called the *principal directions* for the data $\{x_1, \dots, x_n\}$. The vectors λ_i , $i = 1, \dots, n$ are called *principal components* of the data and represent the projection of the data to the space spanned by the principal directions. In what follows, we visualize the projections of the point clouds representing $\phi(\mathcal{M})$ onto three-dimensional subspaces spanned by three selected principal directions. Let us mention that, for better numerical stability in the computation of the finite differences $\partial_{(x,y)} \phi$ and $\partial_{(x,y)}^2 \phi$, we rejected pairs with distance below a certain threshold for the datasets (S), (R), and (A). We will express this distance in terms of the maximal possible distance between data points on the manifold.

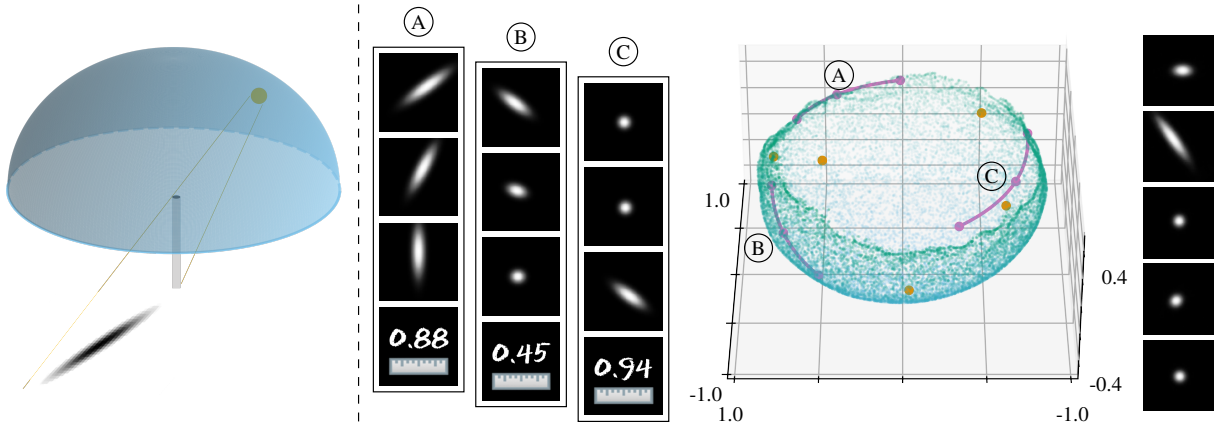


Figure 6.2: Results for dataset (S), for $\lambda = 0$ and $\epsilon = \pi/8$ (with maximal distance being π). From left to right: a sketch of the sundial configuration, training data (image pairs with their geodesic average and distance), the latent manifold $\phi(\mathcal{M})$ projected into \mathbb{R}^3 via PCA, decoder outputs for the orange points in latent space. The encoder was trained separately from the decoder and the training was stopped after the value of the functional $E^{S^\epsilon}[\phi]$ evaluated on a test set did not decrease for 10 epochs. Pairs with distance below 0.01π were rejected. The decoder was trained until an accuracy of 10^{-5} was reached on the same test set.

For dataset (S), our approach and its result at one glance are summarized in Figure 6.2 for $\lambda = 0$. We observe that the embedding is smooth and the geometry of the hemisphere is reproduced. In comparison, Figure 6.3 shows the result for a higher bending penalty $\lambda = 5$, resulting in a much flatter embedding of the hemisphere. This

can also be observed from the significantly smaller variance contribution of the third component. The number of relevant dimensions in both cases is three.




Figure 6.3: Latent manifold $\phi(\mathcal{M})$ for dataset **(S)** for $\lambda = 5$ (colored points as in Figure 6.2). The other parameters and the training procedure are the same as for Figure 6.2.

In Figure 6.4a, we show results for dataset **(R)**. The embedding is smooth, revealing the topology and geometry of the manifold $SO(3)$. For this dataset, in Figure 6.4b, we show that the isometry, flatness, and reconstruction parts of the loss function decrease continuously and monotonically (up to the usual stochastic variations) during the training. The same observations hold for other datasets.


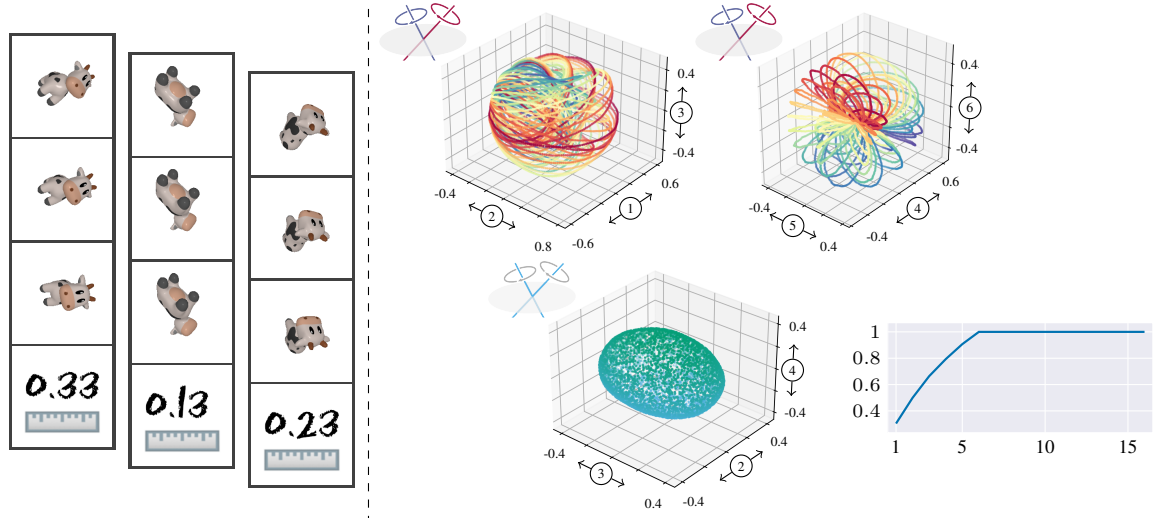
In all our experiments, the latent space dimensionality l is chosen substantially larger than the minimum value required for a smooth embedding – in real applications the intrinsic dimensionality m is unknown beforehand anyway, so one needs to pick l rather larger than smaller. Furthermore, the encoder might achieve flatter embeddings when more dimensions are available. For instance, $S^1 \times S^1 \subset \mathbb{R}^4$ provides an embedding of the flat torus, possibly better than the standard embedding in \mathbb{R}^3 . The encoder makes use of that freedom for dataset **(R)**, which can be read from the graph of the total amount of explained variance as a function of increasing subspace dimension, in the bottom right part of Figure 6.4a. The embedding uses six Euclidean dimensions even though $l = 5$ would be enough. Indeed, $SO(3) \cong \mathbb{R}P^3$, where $\mathbb{R}P^3$ is the 3-dimensional real projective space, which embeds into \mathbb{R}^5 , but not in \mathbb{R}^4 [Hop40, Han38].

In Figure 6.5b, we can read off the used number of dimensions for dataset **(A)**. To the best of our knowledge, it is not known whether there exists a smooth isometric embedding into less than five dimensions. On the other hand, isometric immersions (with self-intersections) exist in four dimensions [Tom41]. Since our encoder loss functional cannot distinguish between immersions and embeddings (there is no term preventing self-intersections), the minimum dimensionality would thus be no larger than four, but our latent manifold uses five Euclidean dimensions in the case $\lambda > 0$, while without bending regularization, *i.e.*, for $\lambda = 0$, the embedding uses eleven dimensions. As discussed above, the isometry term does not make this necessary, but rather a manifestation of the strong degeneracy of the set of isometric embeddings, which are possible using arbitrarily many Euclidean dimensions.


It is well-known that no three-dimensional embedding without self-intersections exists [Gra71], explaining the self-intersections in each of the visualized projections in Figure 6.5a. Denoting the composition of ϕ with the projection onto a 3-dimensional subspace V by ϕ_V , a self-intersection is characterized by $\phi_V(x) = \phi_V(y)$ for $x \neq y$. Thus, a suitable indicator of a self-intersection at x is

$$\min_{y \in \mathcal{M}, y \neq x} \frac{|\phi_V(x) - \phi_V(y)|}{d_{\mathcal{M}}(x, y)}. \quad (6.29)$$

If this quantity is small, we expect a self-intersection at x . Recall that when training the encoder separately, our method in principle does not prevent self-intersections. In Figure 6.5b, we plot the true distances on the Klein bottle against Euclidean distances of the corresponding embedded points in latent space, clearly indicating the absence of any self-intersections. While this might be due to the joint training with the decoder which promotes the injectivity of the embedding, we observed similar results when training the encoder separately. This plot further shows that for $\lambda = 0$, the distances are more or less preserved up to the value of ϵ used in the training, whereas for $\lambda = 1$, the isometry property is sacrificed for an embedding with less bending.

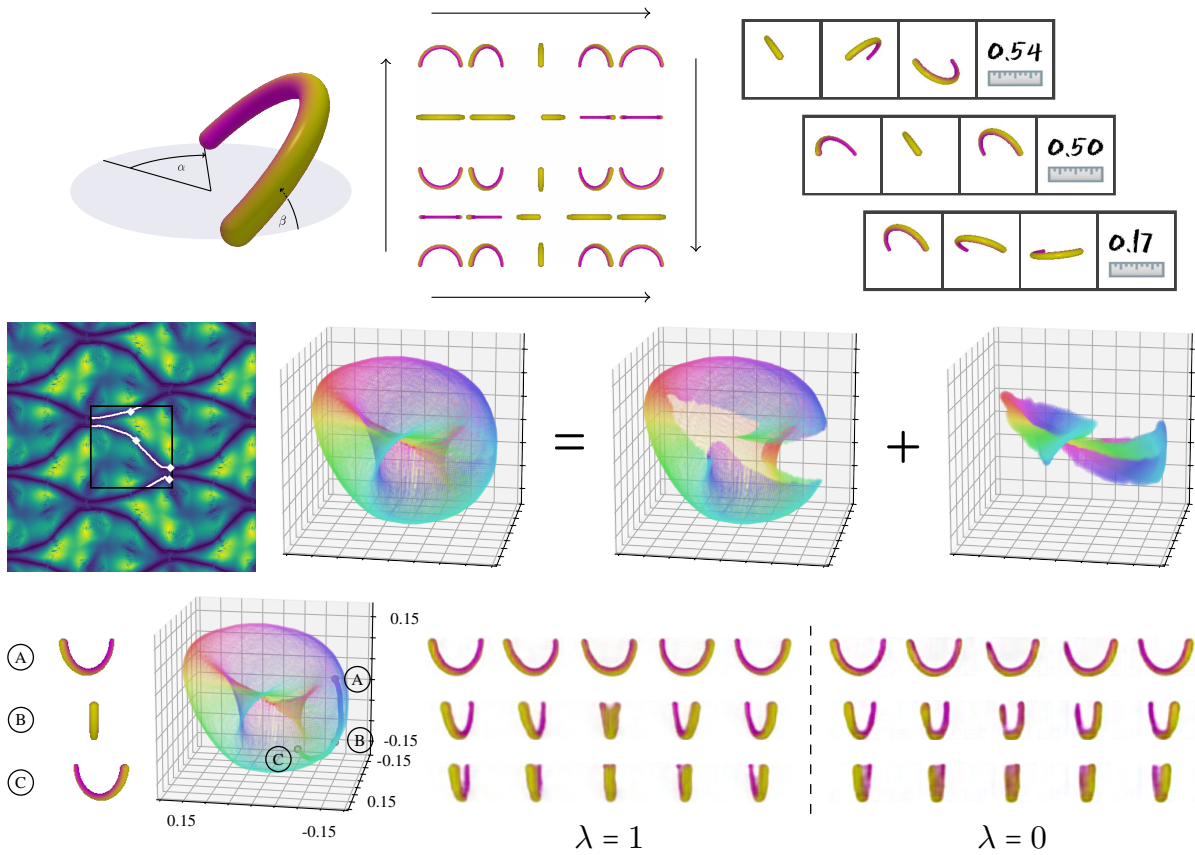


(a) Visualization of the results for dataset (\mathbf{R}) , for $\lambda = 10$ (if not stated otherwise). Top left: Three input triples and their distances. Top right: Visualizations of the latent manifold $\phi(\mathcal{M})$, where the number on the axis corresponds to the order of the principal component. For the top row visualizations, we regularly sample rotation axes from different latitudes of S^2 , and then show points corresponding to rotations around each axis in the same color (color corresponds to z -coordinate of the rotation axis). For the bottom row visualization, we fixed a rotation angle and randomly sampled the rotation axis from S^2 . In components 2, 3, 4, a sphere-like structure can be observed, suggesting that these three components encode the rotation axis. The graph on the bottom right shows the total amount of explained variance as a function of increasing subspace dimension, with a threshold of 99% reached for 6 dimensions. Below the dashed line: Examples of interpolations generated by interpolating linearly in latent space and subsequent decoding. For vanishing bending regularization, $\lambda = 0$, those interpolations are unreliable, while they look very reasonable for $\lambda = 10$, even though the decoder was neither trained on linear interpolations in latent space nor was it regularized.




(b) Bottom figure: Comparison of the temporal evolution of the three loss components, for $\lambda = 10$ and $\lambda = 0$ (logarithmic y -axis). In one epoch (optimization step), around 10000 images are processed in batches of 128.

Figure 6.4: Visualization of the results for dataset (\mathbf{R}) . We used $\epsilon = \frac{\pi}{8}$ (with maximal distance being $\frac{\pi}{2}$). Encoder and decoder were trained at the same time with $\kappa = 0.1$, and training was stopped once a reconstruction error R on a test set reached a threshold of $3.5 \cdot 10^{-3}$. Pairs with distance below $\frac{1}{20}$ of the maximal distance were rejected.



(a) Visualization of the results for dataset (A) and $\lambda = 1$ (unless indicated otherwise). First row: Illustration of the data generation. The first image shows how an arc is rotated by angles α and β . The second image shows the variation of α in horizontal and variation of β in vertical direction, making visible the mirroring leading to the topology of the Klein bottle. On the right, three input triples and their distances are shown. Second row: Projection of the latent manifold $\phi(\mathcal{M})$ onto the first three PCA components. The points are colored according to the original α coordinate, using a cyclic color map. The first image visualizes, in the black frame, the quantity (6.29) evaluated on a regular grid on $[0, 1]^2$. The image is extended in all directions according to the identification of the sides. The white lines were computed using the watershed algorithm [Beu79]. Each one of the lines (circles) is folded together, with the diamond markers showing the folding points. The folding region is visible in the second “summand” of the visualization. Third row: the embedding of a geodesic curve (computed by the interpolation of the input parameters) between A and C, where B is their geodesic midpoint, as well as the images obtained through linear interpolation in latent space for $\lambda = 0$ and $\lambda = 1$ using different starting points symmetric around B.



(b) Comparison of the results for $\lambda = 0$ and $\lambda = 1$. Left: The total amount of explained variance as a function of increasing subspace dimension. The threshold of 99% explained variance was reached for 5 and 11 dimensions for $\lambda = 1$ and $\lambda = 0$, respectively. Right: The extrinsic latent space distance $|\phi(x) - \phi(y)|$ versus the intrinsic manifold distance $d_{\mathcal{M}}(x, y)$ for random point pairs $(x, y) \in \mathcal{M}$.

Figure 6.5: Visualization of results for dataset (A). We used $\epsilon = \frac{1}{4\sqrt{2}}$ (with maximal distance being $\frac{1}{\sqrt{2}}$). Encoder and decoder were trained at the same time with $\kappa = 0.1$, and training was stopped once a reconstruction error R on a test set reached a threshold of $2 \cdot 10^{-3}$. Pairs with distance below $\frac{1}{20}$ of the maximal distance were rejected.

To illustrate the regularization properties of our loss functional we tried to keep the amount of noise in our experiments as low as possible (*i.e.*, the only noise is due to discretization as pixel images), as noisy images would require additional tailored regularization. However, the process of generating images from low dimensional manifolds may introduce noise. We illustrate this effect on the dataset **(E)** by comparing smooth and binary images. Indeed, while the set of smooth ellipse images is infinite, the set of binary ellipse images is finite due to quantization. Thus, while the underlying manifold is still the same, the corresponding set in image space is different. For instance, the same pair of binary ellipse images may be sampled with different distances, leading to a type of noise Figure 6.6 shows that the cylindrical structure of the resulting latent manifold $\phi(\mathcal{M})$ is only barely observable in this case. Increasing the resolution of the images would reduce the disparity between the underlying manifold and the image manifold, leading to a reduction of the observed effect.

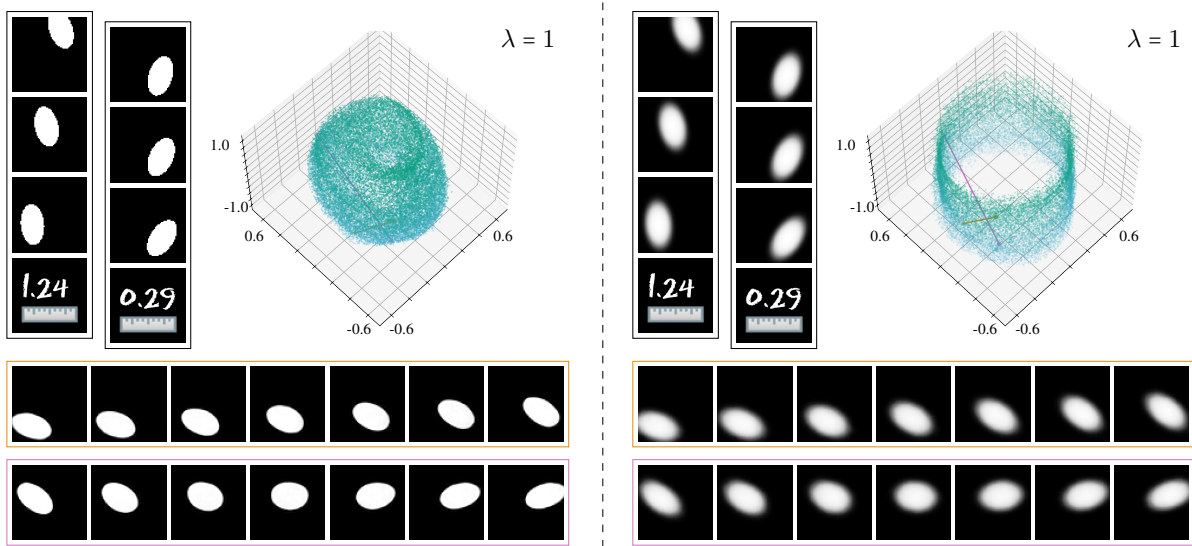


Figure 6.6: Visualization of the results for dataset **(E)** for $\lambda = 1$ and $\epsilon = \frac{\pi}{2}$ for different ways of rendering ellipses: as the characteristic function (left) and as its smoothed version (6.28) with $k = 3$ (right). In each case, two input triples and their distances (with maximal distance ≈ 4.3) are shown. In both cases, there are 4 relevant directions. PCA coordinates 1, 3, 4 and 2, 3, 4, respectively, are used to visualize the latent manifolds. The choice is taken for consistency of visualization, where the change of the relevant dimension is due to the symmetry of the center position parameters c in the dataset. On the bottom, example interpolations are visualized. Their corresponding interpolation path in latent space is shown in the color of the frame. The encoder and decoder were trained at the same time with $\kappa = 0.1$, and training was stopped after 500 epochs.

6.5.3 Linear Interpolation in Latent Space

As already several times mentioned in this chapter, the flatness of the latent manifold is expected to improve the usefulness of time discrete linear interpolation across moderate distances. Recall that this linear interpolation is, for given points $x_A, x_B \in \mathcal{M}$ and some $K \geq 2$, defined by

$$\psi(\tilde{x}_k), \text{ for } \tilde{x}_k = \left(1 - \frac{k}{K}\right)\phi(x_A) + \frac{k}{K}\phi(x_B), \quad k = 0, \dots, K. \quad (6.30)$$

In this section, we study this perspective of our approach, based on the results of the numerical experiments.

The obvious question to ask is the one about the influence of the bending energy parameter λ on the quality of linear interpolations. This is qualitatively illustrated in Figure 6.4a for dataset **(R)**. For $\lambda = 0$, the decoder output of linear interpolations in latent space does not at all reproduce continuously rotating objects and the interpolation objects have a large amount of blending artifacts. The results are considerably better for $\lambda = 10$ and one can follow the rotation of the object. We point out that, since our focus is on the regularizing properties of our encoder loss functional E^{S_ϵ} , the decoder was intentionally not trained on linear interpolation, and not additionally regularized.

An additional training of the decoder on linear interpolants in the latent space would improve the interpolation quality even more, and perhaps even for $\lambda = 0$ allow the decoder to compensate the deficiencies visible in Figure 6.4a. As illustrated in Figure 6.6 by the line representing the linear interpolation path in the latent space, due to the extrinsic curvature of the latent manifold, the linear interpolation only makes sense for sufficiently close endpoints.

To quantify the error of the linear interpolation in latent space, we measured the L^2 -error to the ground truth geodesic interpolation using the following error measure on a test sample set $\mathcal{S}'_\epsilon \subset \{(x, y) \in \mathcal{M} \times \mathcal{M} : y \in D_\epsilon^{\mathcal{M}}(x)\}$ as a function of interpolation distance

$$\text{err}(\delta)^2 = \frac{1}{|\mathcal{S}'_\epsilon|} \sum_{\substack{(x,y) \in \mathcal{S}'_\epsilon \\ d_{\mathcal{M}}(x,y) \leq \delta}} \text{err}_i(x, y)^2 - \text{err}_b(x, y)^2, \quad \text{where} \quad (6.31)$$

$$\text{err}_i(x, y) = \|\text{av}_{\mathcal{M}}(x, y) - \psi(\text{av}_{\mathbb{R}^l}(\phi(x), \phi(y)))\|_{L^2},$$

$$\text{err}_b(x, y) = \|\text{av}_{\mathcal{M}}(x, y) - \psi(\phi(\text{av}_{\mathcal{M}}(x, y)))\|_{L^2}.$$

Here, err_i is the error due to linear interpolation, and err_b is the base reconstruction error that occurs independently of interpolation.

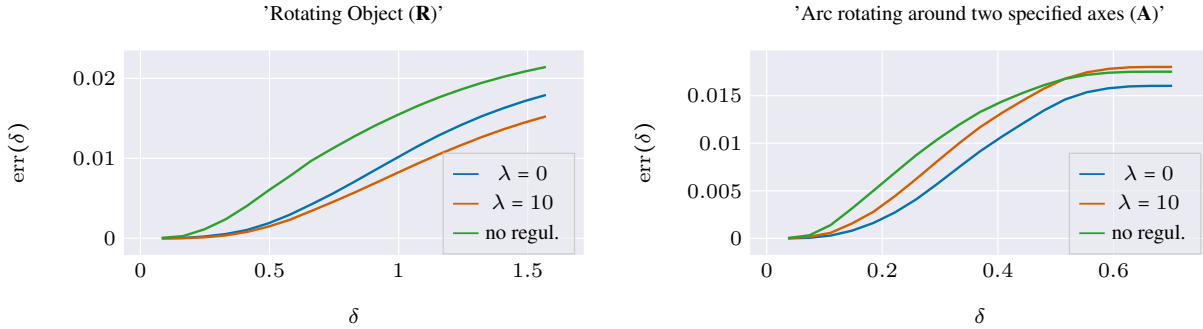


Figure 6.7: Comparison of the interpolation error, for different options for the regularization, for datasets **(R)** (left) and **(A)** (right). The encoder was trained separately from the decoder for 500 epochs, and the weights yielding the best value for the functional $E^{\mathcal{S}_\epsilon}[\phi]$ evaluated on a test set were chosen. Pairs with distance below $\frac{1}{20}$ of the respective maximal distance were rejected. The decoder was trained until an accuracy of $3.5 \cdot 10^{-3}$ for **(R)** and $2 \cdot 10^{-3}$ for **(A)** was reached.

For datasets **(R)** and **(A)**, Figure 6.7 compares $\text{err}(\delta)$ for different options for the regularization. Without encoder regularization, *i.e.*, for minimization of solely the reconstruction loss, the error is the highest and increases quickly already for small interpolation distances. Our low distortion regularization reduces the error, particularly for small interpolation distances. However, the effect of the low bending regularization, *i.e.*, the bending weight λ is ambiguous: Higher λ improves the interpolation quality for dataset **(R)**, but not for dataset **(A)**, which is also reflected in the interpolation examples from Figures 6.4a and 6.5a, respectively.

The other point in our study of the linear interpolation is the interplay of the sampling radius ϵ and the bending parameter λ . For dataset **(S)**, we evaluated, in Figure 6.8, the interpolation error $\text{err}(\pi)$ from (6.31) on a test set of arbitrarily chosen point pairs, using sampling strategy (S3). For $\epsilon = \pi$, *i.e.*, using triples with arbitrarily large distances in the training process, we observe a low interpolation error independent of the choice of the bending weight. This indicates that a very large sampling radius ϵ also has some regularizing effect. For smaller ϵ , however, increasing the bending parameter λ strongly decreases the interpolation error. This speaks in favor of our hypothesis that the bending penalty produces latent manifolds with a flatter and smoother structure which is expected to reduce the *generalization error* in postprocessing tasks. The name “generalization” comes from the fact that, even though the encoder was only trained on point pairs with distance below ϵ , the interpolation error is small for arbitrary point pairs.




Figure 6.8: Average interpolation error for training with different sampling radii ϵ , computed across a test set consisting of random pairs with arbitrary distance on the upper hemisphere. The encoder was trained separately from the decoder, and the training was stopped after the value of the functional $E^{S_\epsilon}[\phi]$, evaluated on a different test set, did not decrease for 10 epochs. Pairs with the distance below 0.01π were rejected. The decoder was trained until an accuracy of 10^{-5} was reached on the test set. This was repeated for three different initializations of the network weights, obtained as outputs of training the encoder.

6.6 Extensions, Conclusion and Outlook

In the previous subsection, we saw that the decoded images of the linear interpolation in the latent space (*cf.* (6.30)), obtained as the result of low distortion and low bending manifold embeddings, could produce both visually and quantitatively decent results in certain examples, while in other examples this was not the case. Therefore, this method is not so reliable. To this end, in this section, we present an idea for a more reliable way for obtaining the geodesic interpolation on the manifold based on the geodesic interpolation on the latent manifold. In addition to the already mentioned collaborators, this project also included Florine Hartwig, who, in particular, implemented the discrete geodesic calculus.

The idea of our approach is based on an implicit description of the latent manifold, *i.e.*, we would like to have a function which is regular and zero on the latent manifold and nonzero everywhere else. To this end, let us suppose that we learned the latent manifold $\phi(\mathcal{M}) \subset \mathbb{R}^l$. Inspired by [AB14], we first learn a projection onto the manifold of the points in \mathbb{R}^l obtained by adding the normally distributed noise to the latent manifold points. Subsequently, we use the difference between the projection and the identity as the function which gives the implicit representation of the latent manifold. To this end, we observe the minimization problem

$$\min_{\theta} \frac{1}{|S_{\phi(\mathcal{M})}|} \sum_{x \in S_{\phi(\mathcal{M})}} \frac{1}{|S_n|} \sum_{\delta \in S_n} |x - p_{\theta}(x + \delta)|^2,$$

where θ is the vector of parameters of the neural network whose realization is the projection function $p_{\theta} : \mathbb{R}^l \rightarrow \mathbb{R}^l$, $S_{\phi(\mathcal{M})}$ is a finite sample of points from $\phi(\mathcal{M})$, and S_n a finite sample of points with the normal distribution $\mathcal{N}(0, \sigma \mathbb{1}_l)$. One can observe the limit as the number of samples tends to infinity and study the properties of the obtained optimal function p , along the same lines as in [AB14]. More on this and the architecture of the underlying network will be presented in the thesis of J. Braunsman. For our discussion, it is important to observe that the (optimal) function p tends to project (thus the name) points in the neighborhood of the manifold to their nearest points on the manifold.

Having the projection function p , we can observe the latent manifold $\phi(\mathcal{M})$ as the zero set of the function $\tilde{p} : \mathbb{R}^l \rightarrow \mathbb{R}^l, x \mapsto p(x) - x$. For determining a discrete geodesic path between two points on the manifold, we use the framework developed in Subsection 3.1.2. In particular, in the context of (3.6), we use the standard Euclidean distance between the points as the approximation functional, *i.e.*, $\mathcal{W}[x, \tilde{x}] := |x - \tilde{x}|^2$. This is a reasonable choice in the context of penalizing bending of the embedding, which, in particular, penalizes deviation of geometry of $\phi(\mathcal{M})$ from the standard linear one. Then, in the context of (3.7), the path energy of a discrete path (x_0, \dots, x_K) is given by

$$\mathbf{E}^K[(x_0, \dots, x_K)] := K \sum_{k=1}^K |x_k - x_{k-1}|^2.$$


However, it is straightforward to check that the minimizer of this energy, for fixed x_0 and x_K , is nothing else but the linear interpolation. What we miss are the terms that constrain the discrete interpolation path to stay on the

manifold. To this end, let us have the fixed initial and the final points $x_0 = x_A$ and $x_K = x_B$, respectively. Then, in order to find the discrete geodesic path between these points on the latent manifold, among the several possible approaches, we used the so called *augmented Lagrangian approach* [NW06], which combines the Lagrangian approach and quadratic penalty method. Altogether, we have the following optimization problem

$$\min_{X \in (\mathbb{R}^l)^{K-1}} \max_{\Lambda \in \mathbb{R}^{l, K-1}} \mathbf{E}^K[(x_0, X, x_K)] - \Lambda : \tilde{p}(X) + \frac{\mu}{2} \|\tilde{p}(X)\|^2,$$

where $\lambda : \tilde{p}(X) := \sum_{i=1}^l \sum_{k=1}^{K-1} \Lambda_{ik} \tilde{p}_i(X_k)$ and $\|\tilde{p}(X)\|^2 = \sum_{k=1}^{K-1} |\tilde{p}(X_k)|^2$. The elements of Λ are called *Lagrangian multipliers*. Solving the problem is based on an alternating algorithm, which includes a Newton scheme for optimization of X , and the gradient descent for optimization of Λ . The penalty parameter μ is increased at each step, until we reach a sufficient attainment of the equality constraints $\tilde{p}(X) = 0$ [NW06, Chapter 17].

For numerical experiments, we used the dataset consisting of ellipses obtained as anisotropic Gaussians with fixed aspect ratio, similar to the DSprites dataset in [MHHL17]. The Riemannian distance was computed analogous to the one for the dataset (**E**), we use sampling strategy (S1), and the images are of resolution 64×64 .



(a) Left: Visualization of the learned manifold by projection onto the first three PCA components. Green dots represent the geodesic interpolation obtained from the augmented Lagrangian approach, with red points being the given initial and final points. Blue points represent linear interpolation. Right: Visualization of the values of the norm of the function \tilde{p} on a two-dimensional slice through the learned “cylinder”.



(b) Decoder outputs of the above geodesic interpolation with $K = 10$. The fixed images (embedded points on the manifold) are framed in red.



(c) Decoder outputs of the above linear interpolation.

Figure 6.9: The visualization of the results of the embedding of Gaussian ellipses and geodesic interpolation on the obtained latent manifold. We trained the encoder and the decoder jointly on around 10000 images with $\kappa = 0.1$. The bending parameter is $\lambda = 1$, while the standard deviation parameter σ is $\frac{1}{10}$ of the maximal distance of the training points. Visualization is obtained by the projection to the first 3 PCA dimensions (out of 4 relevant). The results clearly show the advantages of the geodesic interpolation on the manifold in comparison with the simple linear interpolation in the latent space. In particular, the shape of the ellipse is preserved for the geodesic interpolation, which is not the case for the linear interpolation.

Conclusion and outlook In this chapter, we have proposed a regularization energy functional for the encoder mapping of the autoencoder framework which promotes low bending and low distortion embeddings to the latent spaces. The discrete sampling loss functional is computed via Monte Carlo integration, where the training data consists of pairs of not too distant images, measured in terms of the sampling radius, together with their Riemannian distance and Riemannian average. We studied the limit as the number of samples tends to infinity leading to the continuous nonlocal loss functional. For this functional, which depends on the strategy used for obtaining the discrete samples, we proved the existence of a minimizer in the class of functions which are representable by classes of neural networks and satisfy certain regularity and density properties. The limit as the sampling radius tends to zero is studied with the tool of Mosco-convergence, where the fully local limit functional again promotes low bending and low distortion embeddings. In the presented numerical experiments, the input images were realizations of handcrafted low-dimensional manifolds, allowing us to efficiently compute the underlying Riemannian distances and averages. The obtained results, visualized by projections to principal component spaces, show that the obtained latent manifolds are smooth and low-dimensional. The role of the part of the energy promoting low bending is studied, both qualitatively and quantitatively, by observing the decoder images of linear interpolation in the latent space. While for some datasets the quality of the results of interpolation was considerably higher for a larger weighting of the bending energy, in others the quality was almost the same or even lower and the measured error was bigger. In fact, for more complicated latent manifolds, various factors might additionally impact the interpolation error such as the sign of the intrinsic curvature, an increased tangential distortion of the encoder map for flatter latent manifolds, or the interplay of the differently trained encoder and decoder map. Indeed, simultaneous training of encoder and decoder via minimization of the (weighted) sum of encoder regularization loss and reconstruction loss may have additional regularizing effects since some embeddings may be more favorable for the decoder than others. As a continuation of the presented work, such effects yet have to be understood as well as whether and how the decoder itself should be regularized on top. A possible direction to this study is the extension proposed above. Furthermore, as an extension of the theoretical results, one should consider studying the (non)local energies in a less restricted setting. Finally, this approach should be applied to some more realistic shape spaces arising from the classical image processing applications.

Bibliography

- [AAB⁺15] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [AAS22] B. B. Amor, S. Arguillere, and L. Shao. ResNet-LDDMM: Advancing the LDDMM framework using deep residual networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [AB14] G. Alain and Y. Bengio. What regularized auto-encoders learn from the data-generating distribution. *The Journal of Machine Learning Research*, 15(1):3563–3593, 2014.
- [ACPA06] V. Arsigny, O. Commowick, X. Pennec, and N. Ayache. A log-euclidean framework for statistics on diffeomorphisms. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 924–931. Springer, 2006.
- [AF03] R. A. Adams and J. J. F. Fournier. *Sobolev spaces*, volume 140 of *Pure and Applied Mathematics (Amsterdam)*. Elsevier/Academic Press, Amsterdam, second edition, 2003.
- [AGL20] M. Atzmon, A. Gropp, and Y. Lipman. Isometric autoencoders. *arXiv preprint arXiv:2006.09289*, 2020.
- [AGS08] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.
- [AK98] V. Arnold and B. Khesin. *Topological methods in hydrodynamics*. Springer, 1998.
- [AK09] G. Aubert and P. Kornprobst. Can the nonlocal characterization of Sobolev spaces by Bourgain et al. be useful for solving variational problems? *SIAM Journal on Numerical Analysis*, 47(2):844–860, 2009.
- [AKAD22] R. Almoheisen, M. R. Keaton, D. A. Adjeroh, and G. Doretto. Generative probabilistic novelty detection with isometric adversarial autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2003–2013, 2022.
- [Alt06] H. W. Alt. *Lineare Funktionalanalysis*. Springer Verlag, fifth edition, 2006.
- [AMS09] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [Arn66] 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 (Grenoble)*, 16(1):319–361, 1966.
- [BAC19] A. Bock, A. Arnaudon, and C. Cotter. Selective metamorphosis for growth modelling with applications to landmarks. In *International Conference on Geometric Science of Information*, pages 39–48. Springer, 2019.
- [Bal81] J. M. Ball. Global invertibility of Sobolev functions and the interpenetration of matter. *Proc. Roy. Soc. Edinburgh*, 88A:315–328, 1981.
- [BBE⁺17] 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*, Informatik aktuell, pages 68–73. Springer, 2017.
- [BBM01] J. Bourgain, H. Brezis, and P. Mironescu. Another look at Sobolev spaces. *Optimal Control and Partial Differential Equations*, A Volume in Honour of A. Bensoussan’s 60th Birthday:439–455, 2001.
- [BC20] A. Bock and C. Cotter. Space-time metamorphosis. *arXiv preprint arXiv:2005.08743*, 2020.
- [BCV13] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

- [BDLM10] J. Bolte, A. Daniilidis, O. Ley, and L. Mazet. Characterizations of Łojasiewicz inequalities: subgradient flows, talweg, convexity. *Transactions of the American Mathematical Society*, 362(6):3319–3363, 2010.
- [BER15] 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.
- [Beu79] S. Beucher. Use of watersheds in contour detection. In *Proceedings of the International Workshop on Image Processing*. CCETT, 1979.
- [BGKP19] H. Bolcskei, P. Grohs, G. Kutyniok, and P. Petersen. Optimal approximation with sparsely connected deep neural networks. *SIAM Journal on Mathematics of Data Science*, 1(1):8–45, 2019.
- [BGV19] J.-D. Benamou, T. O. Gallouët, and F.-X. Vialard. Second-order models for optimal transport and cubic splines on the Wasserstein space. *Found. Comput. Math.*, 19(5):1113–1143, 2019.
- [BK17] B. Benešová and M. Kružík. Weak lower semicontinuity of integral functionals and applications. *SIAM Rev.*, 59(4):703–766, 2017.
- [BL18] K. Bredies and D. Lorenz. *Mathematical Image Processing*. Vieweg+Teubner Verlag, 1 edition, 2018.
- [BMR13] M. Burger, J. Modersitzki, and L. Ruthotto. A hyperelastic regularization energy for image registration. *SIAM J. Sci. Comput.*, 35(1):B132–B148, 2013.
- [BMTY02] M. F. Beg, M. Miller, A. Trouvé, and L. Younes. Computational anatomy: Computing metrics on anatomical shapes. In *Proceedings of 2002 IEEE ISBI*, pages 341–344, 2002.
- [BMTY05] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Inter. J. Comput. Vision*, 61(2):139–157, 2005.
- [BN01] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Advances in neural information processing systems*, 14, 2001.
- [Bor07] R. Borghol. Some properties of Sobolev spaces. *Asymptot. Anal.*, 51(3-4):303–318, 2007.
- [Bou23] N. Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
- [Bra06] A. Braides. *A Handbook of Γ -Convergence*, volume III of *Handbook of differential equations: stationary partial differential equations*. Elsevier/North-Holland, Amsterdam, 2006.
- [Bre11] H. Brezis. *Functional analysis, Sobolev spaces and partial differential equations*, volume 2. Springer, 2011.
- [BRRG18] D. Berthelot, C. Raffel, A. Roy, and I. Goodfellow. Understanding and improving interpolation in autoencoders via an adversarial regularizer. *arXiv preprint arXiv:1807.07543*, 2018.
- [BRRW21] J. Braunsmann, M. Rajković, M. Rumpf, and B. Wirth. Learning low bending and low distortion manifold embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 4416–4424, 2021.
- [BRRW22] J. Braunsmann, M. Rajković, M. Rumpf, and B. Wirth. Convergent autoencoder approximation of low bending and low distortion manifold embeddings. 2022.
- [BST14] J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1):459–494, 2014.
- [BT09] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.
- [BV04] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [BV17] M. Bruveris and F.-X. Vialard. On completeness of groups of diffeomorphisms. *J. Eur. Math. Soc. (JEMS)*, 19(5):1507–1544, 2017.
- [BVCD20] A. Bône, P. Vernhet, O. Colliot, and S. Durrleman. Learning joint shape and appearance representations with metamorphic auto-encoders. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 202–211. Springer, 2020.
- [CCG18] Y. Chen, G. Conforti, and T. T. Georgiou. Measure-valued spline curves: An optimal transport viewpoint. *SIAM Journal on Mathematical Analysis*, 50(6):5947–5968, 2018.
- [CCG⁺21] S. Chewi, J. Clancy, T. L. Gouic, P. Rigollet, G. Stepaniants, and A. J. Stromme. Fast and smooth interpolation on Wasserstein space. In *International Conference on Artificial Intelligence and Statistics*, pages 3061–3069. PMLR, 2021.
- [CCT18] N. Charon, B. Charlier, and A. Trouvé. Metamorphoses of functional shapes in Sobolev spaces. *Foundations of Computational Mathematics*, 18(6):1535–1596, jan 2018.

- [CGÖ19] C. Chen, B. Gris, and O. Öktem. A new variational model for joint image reconstruction and motion estimation in spatiotemporal imaging. *SIAM Journal on Imaging Sciences*, 12(4):1686–1719, 2019.
- [Cia88] P. G. Ciarlet. *Mathematical elasticity. Vol. I*, volume 20 of *Studies in Mathematics and its Applications*. North-Holland Publishing Co., Amsterdam, 1988. Three-dimensional elasticity.
- [CKF⁺20] N. Chen, A. Klushyn, F. Ferroni, J. Bayer, and P. Van Der Smagt. Learning flat latent manifolds with VAEs. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1587–1596, 2020.
- [CKK⁺18] N. Chen, A. Klushyn, R. Kurle, X. Jiang, J. Bayer, and P. Smagt. Metrics for deep generative models. In *International Conference on Artificial Intelligence and Statistics*, pages 1540–1550. PMLR, 2018.
- [CL06] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.
- [CLMC92] F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal on Numerical analysis*, 29(1):182–193, 1992.
- [CÖ18] C. Chen and O. Öktem. Indirect image registration with large diffeomorphic deformations. *SIAM Journal on Imaging Sciences*, 11(1):575–617, 2018.
- [CP17] Y. Chen and T. Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1256–1272, 2017.
- [CR74] E. Catmull and R. Rom. A class of local interpolating splines. In *Computer aided geometric design*, pages 317–326. Elsevier, 1974.
- [CRM96] G. Christensen, R. D. Rabbitt, and M. I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447, 1996.
- [CRM⁺22] E. Cramer, F. Rauh, A. Mitsos, R. Tempone, and M. Dahmen. Nonlinear isometric manifold learning for injective normalizing flows. *arXiv preprint arXiv:2203.03934*, 2022.
- [Dac08] B. Dacorogna. *Direct methods in the calculus of variations*, volume 78 of *Applied Mathematical Sciences*. Springer, New York, second edition, 2008.
- [dB63] C. de Boor. Best approximation properties of spline functions of odd degree. *J. Math. Mech.*, 12:747–749, 1963.
- [dB78] C. de Boor. *A practical guide to splines*, volume 27. Springer verlag New York, 1978.
- [DBGS19] A. V. Dalca, G. Balakrishnan, J. Guttag, and M. R. Sabuncu. Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical image analysis*, 57:226–236, 2019.
- [dC92] M. P. do Carmo. *Riemannian Geometry*. Birkhäuser, Boston, 1992.
- [DG03] D. L. Donoho and C. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003.
- [DG05] D. L. Donoho and C. Grimes. Image manifolds which are isometric to Euclidean space. *Journal of mathematical imaging and vision*, 23(1):5–24, 2005.
- [DGM98] P. Dupuis, U. Grenander, and M. I. Miller. Variational problems on flows of diffeomorphisms for image matching. *Quart. Appl. Math.*, 56:587–600, 1998.
- [DM93] G. Dal Maso. *An introduction to Γ -convergence*. Progress in Nonlinear Differential Equations and their Applications, 8. Birkhäuser Boston, Inc., Boston, MA, Boston, 1993.
- [DR04] M. Droske and M. Rumpf. A variational approach to non-rigid morphological registration. *SIAM Journal on Applied Mathematics*, 64(2):668–687, 2004.
- [EHRW21] A. Effland, B. Heeren, M. Rumpf, and B. Wirth. Consistent curvature approximation on Riemannian shape spaces. *IMA J. Numer. Anal.*, 2021.
- [EKP⁺21] A. Effland, E. Kobler, T. Pock, M. Rajković, and M. Rumpf. Image morphing in deep feature spaces: Theory and applications. *J. Math. Imaging Vis.*, 63:309–327, 2021.
- [EKPR19] A. Effland, E. Kobler, T. Pock, and M. Rumpf. Time discrete geodesics in deep feature spaces for image morphing. In J. Lellmann, M. Burger, and J. Modersitzki, editors, *Scale Space and Variational Methods in Computer Vision*, pages 171–182, Cham, 2019. Springer International Publishing.
- [ENR20] 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.
- [ERS⁺15] A. Effland, M. Rumpf, S. Simon, K. Stahn, and B. Wirth. Bézier curves in the space of images. In *Proc. of International Conference on Scale Space and Variational Methods in Computer Vision*, volume 9087 of *Lecture Notes in Computer Science*, pages 372–384. Springer, Cham, 2015.

- [ERS18] 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.
- [FGG21] A. François, P. Gori, and J. Glaunès. Metamorphic image registration using a semi-Lagrangian scheme. In *International Conference on Geometric Science of Information*, pages 781–788. Springer, 2021.
- [Fio16] R. Fiorenza. *Hölder and locally Hölder continuous functions, and open sets of class $C^k, C^{k,\lambda}$* . Frontiers in Mathematics. Birkhäuser/Springer, Cham, 2016.
- [FMO⁺22] A. François, M. Maillard, C. Oppenheim, J. Pallud, P. Gori, and J. A. Glaunès. Weighted metamorphosis for registration of images with different topology. In *10th International Workshop on Biomedical Image Registration*, 2022.
- [Fol99] G. B. Folland. *Real analysis: modern techniques and their applications*, volume 40. John Wiley & Sons, 1999.
- [FP04] X. Feng and A. Prohl. Analysis of a fully discrete finite element method for the phase field and approximation of its sharp interface limits. *Mathematics of computation*, 73(246):541–567, 2004.
- [GBB11] X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323. JMLR Workshop and Conference Proceedings, 2011.
- [GBC16] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [GCÖ20] B. Gris, C. Chen, and O. Öktem. Image reconstruction through metamorphosis. *Inverse Problems*, 36(2):025001, 2020.
- [Gha14] L. E. Ghaoui. Hyper-textbook: Optimization models and applications, 2014.
- [GHL90] S. Gallot, D. Hulin, and J. Lafontaine. *Riemannian geometry*, volume 2. Springer, 1990.
- [GK22] P. Grohs and G. Kutyniok. *Mathematical Aspects of Deep Learning*. Cambridge University Press, 2022.
- [GKP20] I. Gühring, G. Kutyniok, and P. Petersen. Error bounds for approximations with deep ReLU neural networks in $W^{s,p}$ norms. *Analysis and Applications*, 18(05):803–859, 2020.
- [GPAM⁺20] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [GR21] I. Gühring and M. Raslan. Approximation rates for neural networks with encodable weights in smoothness spaces. *Neural Networks*, 134:107–130, 2021.
- [Gra71] A. Gramain. *Topologie des surfaces*, volume 7. Presses universitaires de Frances, 1971.
- [GRK20] I. Gühring, M. Raslan, and G. Kutyniok. Expressivity of deep neural networks. *arXiv preprint arXiv:2007.04759*, 2020.
- [GRRW23] M. Guastini, M. Rajković, M. Rumpf, and B. Wirth. The variational approach to the flow of Sobolev-diffeomorphisms model. In L. Calatroni, M. Donatelli, S. Morigi, M. Prato, and M. Santacesaria, editors, *Scale Space and Variational Methods in Computer Vision*, pages 551–564, Cham, 2023. Springer International Publishing.
- [Gua22] M. Guastini. Time-discretization of the flow of diffeomorphisms model. Master’s thesis, University of Bonn, 2022.
- [Han38] W. Hantzsche. Einlagerung von Mannigfaltigkeiten in euklidische Räume. *Math. Z.*, 43(1):38–58, 1938.
- [HC22] H.-W. Hsieh and N. Charon. Weight metamorphosis of varifolds and the LDDMM-Fisher-Rao metric. *Calculus of Variations and Partial Differential Equations*, 61(5):1–41, 2022.
- [HDY⁺12] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6):82–97, 2012.
- [Heb96] E. Hebey. *Sobolev spaces on Riemannian manifolds*, volume 1635. Springer Science & Business Media, 1996.
- [Hen19] F. Henrich. Kernel-based methods for image metamorphosis. Master’s thesis, University of Bonn, 2019.
- [HJE18] J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [HJS⁺12] 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*, volume 7512 of *Lecture Notes in Computer Science*, pages 197–205, 2012.

- [Hop40] H. Hopf. Systeme symmetrischer Bilinearformen und euklidische Modelle der projektiven Räume. *Vierteljschr. Naturforsch. Ges. Zürich*, 85(Beiblatt (Festschrift Rudolf Fueter)):165–177, 1940.
- [HRS⁺16] B. Heeren, M. Rumpf, P. Schröder, M. Wardetzky, and B. Wirth. Splines in the space of shells. *Comput. Graph. Forum*, 35(5):111–120, 2016.
- [HRW18] B. Heeren, M. Rumpf, and B. Wirth. Variational time discretization of Riemannian splines. *IMA J. Numer. Anal.*, 39(1):61–104, 2018.
- [HSW89] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [HSW90] K. Hornik, M. Stinchcombe, and H. White. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural networks*, 3(5):551–560, 1990.
- [HSW⁺12] J. Hinkle, M. Szegedi, B. Wang, B. Salter, and S. Joshi. 4D CT image reconstruction with diffeomorphic motion model. *Medical image analysis*, 16(6):1307–1316, 2012.
- [HTY09] D. Holm, A. Trounev, and L. Younes. The Euler-Poincaré theory of metamorphosis. *Quart. Appl. Math.*, 67:661–685, 2009.
- [Huy09] D. Q. Huynh. Metrics for 3D rotations: Comparison and analysis. *J. Math. Imaging Vis.*, 35(2):155–164, 2009.
- [HvNVW16] T. Hytönen, J. van Neerven, M. Veraar, and L. Weis. *Analysis in Banach spaces. Vol. I. Martingales and Littlewood-Paley theory*, volume 63. Springer, Cham, 2016.
- [HZN09] G. L. Hart, C. Zach, and M. Niethammer. An optimal control approach for deformable registration. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2009.
- [HZRS15] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [HZRS16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [IKT13] H. Inci, T. Kappeler, and P. Topalov. On the regularity of the composition of diffeomorphisms. *Mem. Amer. Math. Soc.*, 226(1062):vi+60, 2013.
- [JM00] S. C. Joshi and M. I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Trans. Image Process.*, 9(8):1357–1370, 2000.
- [Jol02] I. T. Jolliffe. *Principal component analysis for special types of data*. Springer, 2002.
- [JRE23] J. Justiniano, M. Rumpf, and M. Erbar. Approximation of splines in Wasserstein spaces. *arXiv preprint arXiv:2302.10682*, 2023.
- [JRR21] J. Justiniano, M. Rajković, and M. Rumpf. Splines for Image Metamorphosis. In A. Elmoataz, J. Fadili, Y. Quéau, J. Rabin, and L. Simon, editors, *Scale Space and Variational Methods in Computer Vision: 8th International Conference, SSVM 2021, Virtual Event, May 16–20, 2021, Proceedings*, Lecture Notes in Computer Science, pages 463–475. Springer Nature, Springer International Publishing, 2021.
- [JRR23] J. Justiniano, M. Rajković, and M. Rumpf. Consistent approximation of interpolating splines in image metamorphosis. *J. Math. Imaging Vis.*, 65(1):29–52, 2023.
- [KB15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, 2015.
- [KCM21] D. Kohli, A. Cloninger, and G. Mishne. LDLE: low distortion local eigenmaps. *Journal of Machine Learning Research*, 22(282):1–64, 2021.
- [Kli95] W. P. A. Klingenberg. *Riemannian geometry*, volume 1 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, second edition, 1995.
- [KSH12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
- [Kui55a] N. H. Kuiper. On C^1 isometric embeddings. II. In *Nederl. Akad. Wetensch. Proc. Ser. A*, volume 58, pages 683–689, 1955.
- [Kui55b] N. H. Kuiper. On C^1 -isometric imbeddings. I. In *Indagationes Mathematicae (Proceedings)*, volume 58, pages 545–556. Elsevier, 1955.

- [KW13] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [KZSN20] K. Kato, J. Zhou, T. Sasaki, and A. Nakagawa. Rate-distortion optimization guided autoencoder for isometric embedding in Euclidean latent space. In *International Conference on Machine Learning*, pages 5166–5176. PMLR, 2020.
- [L⁺89] Y. LeCun et al. Generalization and network design strategies. *Connectionism in perspective*, 19(143-155):18, 1989.
- [Lan95] S. Lang. *Differential and Riemannian Manifolds*, volume 160 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, third edition, 1995.
- [Lee03] J. M. Lee. *Introduction to smooth manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2003.
- [LNÖS20] L. F. Lang, S. Neumayer, O. Öktem, and C.-B. Schönlieb. Template-based image reconstruction from sparse tomographic data. *Applied Mathematics & Optimization*, 82(3):1081–1109, 2020.
- [LPSS15] J. Lellmann, K. Papafitsoros, C. Schönlieb, and D. Spector. Analysis and application of a nonlocal Hessian. *SIAM Journal on Imaging Sciences*, 8(4):2161–2202, 2015.
- [MF11] Y. Ma and Y. Fu. *Manifold learning theory and applications*. CRC press, 2011.
- [MFG⁺22] M. Maillard, A. François, J. Glaunès, I. Bloch, and P. Gori. A deep residual learning implementation of metamorphosis. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–4. IEEE, 2022.
- [MHHL17] L. Matthey, I. Higgins, D. Hassabis, and A. Lerchner. dSprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- [Mod09] J. Modersitzki. *FAIR: flexible algorithms for image registration*. SIAM, 2009.
- [MOPS20] M. C. Mulkamala, P. Ochs, T. Pock, and S. Sabach. Convex-concave backtracking for inertial Bregman proximal gradient algorithms in nonconvex optimization. *SIAM Journal on Mathematics of Data Science*, 2(3):658–682, 2020.
- [Mor65] J.-J. Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France*, 93:273–299, 1965.
- [Mos69] U. Mosco. Convergence of convex sets and of solutions of variational inequalities. *Advances in Math.*, 3:510–585, 1969.
- [MP43] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, 1943.
- [MR17] A. Mang and L. Ruthotto. A Lagrangian Gauss–Newton–Krylov solver for mass-and intensity-preserving diffeomorphic image registration. *SIAM Journal on Scientific Computing*, 39(5):B860–B885, 2017.
- [MTY02] M. I. Miller, A. Trounev, and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annu. Rev. Biomed. Eng.*, 4(1):375–405, 2002.
- [MTY06] M. I. Miller, A. Trounev, and L. Younes. Geodesic shooting for computational anatomy. *J. Math. Imaging Vision*, 24(2):209–228, 2006.
- [MTY15] M. I. Miller, A. Trounev, 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.
- [MY01] M. I. Miller and L. Younes. Group actions, homeomorphisms, and matching: a general framework. *Int. J. Comput. Vis.*, 41(1–2):61–84, 2001.
- [Nas54] J. Nash. C^1 isometric imbeddings. *Annals of mathematics*, pages 383–396, 1954.
- [NHP89] L. Noakes, G. Heinzinger, and B. Paden. Cubic splines on curved spaces. *IMA J. Math. Control Inform.*, 6(4):465–473, 1989.
- [Nir66] L. Nirenberg. An extended interpolation inequality. *Ann. Scuola Norm. Sup. Pisa*, 20:733–737, 1966.
- [NKS21] A. Nakagawa, K. Kato, and T. Suzuki. Quantitative understanding of VAE as a non-linearly scaled isometric embedding. In *International Conference on Machine Learning*, pages 7916–7926. PMLR, 2021.
- [NPS18] 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.
- [NPS19] S. Neumayer, J. Persch, and G. Steidl. Regularization of inverse problems via time discrete geodesics in image spaces. *Inverse Problems*, 35(5):055005, 28, 2019.
- [NŠ91] J. Nečas and M. Šilhavý. Multipolar viscous fluids. *Quart. Appl. Math.*, 49(2):247–265, 1991.

- [NT21] S. Neumayer and A. Topalovic. Template based image reconstruction facing different topologies. *arXiv preprint arXiv:2108.13777*, 2021.
- [NVF20] NVIDIA, P. Vingelmann, and F. H. Fitzek. Cuda, release: 10.2.89, 2020.
- [NW06] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2 edition, 2006.
- [OYHO20] A. Oring, Z. Yakhini, and Y. Hel-Or. Autoencoder image interpolation by shaping the latent space, 2020.
- [PC⁺19] G. Peyré, M. Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [PGC⁺17] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch, 2017.
- [PGM⁺19] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [PLD⁺20] E. Peterfreund, O. Lindenbaum, F. Dietrich, T. Bertalan, M. Gavish, I. G. Kevrekidis, and R. R. Coifman. Local conformal autoencoder for standardized data coordinates. *Proceedings of the National Academy of Sciences*, 117(49):30918–30927, 2020.
- [PM90] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- [PNVM20] T. Polzin, M. Niethammer, F.-X. Vialard, and J. Modersitzki. A discretize–optimize approach for LDDMM registration. In *Riemannian Geometric Statistics in Medical Image Analysis*, pages 479–532. Elsevier, 2020.
- [Pol64] B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR computational mathematics and mathematical physics*, 4(5):1–17, 1964.
- [PPS17] J. Persch, F. Pierre, and G. Steidl. Exemplar-based face colorization using image morphing. *Journal of Imaging*, 3(4), 2017.
- [PRV21] P. Petersen, M. Raslan, and F. Voigtlaender. Topological properties of the set of functions generated by neural networks of fixed size. *Foundations of computational mathematics*, 21(2):375–444, 2021.
- [PS16] T. Pock and S. Sabach. Inertial proximal alternating linearized minimization (iPALM) for nonconvex and nonsmooth problems. *SIAM Journal on Imaging Sciences*, 9(4):1756–1787, 2016.
- [PTBK19] G. Pai, R. Talmon, A. Bronstein, and R. Kimmel. Dimal: Deep isometric manifold learning using sparse geodesic sampling. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 819–828. IEEE, 2019.
- [PV20] P. Petersen and F. Voigtlaender. Equivalence of approximation by convolutional neural networks and fully-connected networks. *Proceedings of the American Mathematical Society*, 148(4):1567–1581, 2020.
- [QLC⁺19] S. Qian, G. Li, W.-M. Cao, C. Liu, S. Wu, and H.-S. Wong. Improving representation learning in autoencoders via multidimensional interpolation and dual regularizations. In *IJCAI*, pages 3268–3274, 2019.
- [RHW86] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- [RPC⁺07] M. Ranzato, C. Poultney, S. Chopra, Y. LeCun, et al. Efficient learning of sparse representations with an energy-based model. *Advances in neural information processing systems*, 19:1137, 2007.
- [RRN⁺20] N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari. Accelerating 3D deep learning with PyTorch3D. *arXiv:2007.08501*, 2020.
- [RS00] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.
- [RVM⁺11] S. Rifai, P. Vincent, X. Muller, X. Glorot, and Y. Bengio. Contractive auto-encoders: Explicit invariance during feature extraction. In *Icml*, 2011.
- [RW11] M. Rumpf and B. Wirth. Variational methods in shape analysis. In O. Scherzer, editor, *Handbook of Mathematical Methods in Imaging*, pages 1363–1401. Springer, 2011.
- [RW15] M. Rumpf and B. Wirth. Variational time discretization of geodesic calculus. *IMA J. Numer. Anal.*, 35(3):1011–1046, 2015.

- [RY13] C. L. Richardson and L. Younes. Computing metamorphoses between discrete measures. *J. Geom. Mech.*, 5(1):131–150, 2013.
- [RY16] C. L. Richardson and L. Younes. Metamorphosis of images in reproducing kernel Hilbert spaces. *Adv. Comput. Math.*, 42(3):573–603, 2016.
- [Sak96] T. Sakai. *Riemannian Geometry*. Amer. Math. Soc., 1996.
- [Sch66] D. G. Schweikert. An interpolation curve using a spline in tension. *J. Math. and Phys.*, 45:312–317, 1966.
- [SHCV22] R. Shankar, H.-W. Hsieh, N. Charon, and A. Venkataraman. A diffeomorphic flow-based variational framework for multi-speaker emotion conversion. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2022.
- [SHJF13] N. Singh, J. Hinkle, S. Joshi, and P. T. Fletcher. A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In *2013 IEEE 10th International Symposium on Biomedical Imaging*, pages 1219–1222. IEEE, 2013.
- [SHM⁺16] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [Sim87] J. Simon. Compact sets in the space $L^p(0, T; B)$. *Annali di Matematica Pura ed Applicata*, 146:65–96, 1987.
- [SKTF18] H. Shao, A. Kumar, and P. Thomas Fletcher. The Riemannian geometry of deep generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 315–323, 2018.
- [Sob90] I. Sobel. An isotropic 3×3 gradient operator, machine vision for three-dimensional scenes. *Freeman, H., Academic Pres, NY*, 376:379, 1990.
- [ST19] A. Schwartz and R. Talmon. Intrinsic isometric manifold learning with application to localization. *SIAM Journal on Imaging Sciences*, 12(3):1347–1391, 2019.
- [Ste70] E. Stein. *Singular Integrals and Differentiability Properties of Functions*. Princeton University Press, 1970.
- [STT⁺18] T. Sainburg, M. Thielk, B. Theilman, B. Migliori, and T. Gentner. Generative adversarial interpolative autoencoding: adversarial training on latent space interpolations encourage convex latent distributions. *arXiv preprint arXiv:1807.06650*, 2018.
- [SVN15] N. Singh, F.-X. Vialard, and M. Niethammer. Splines for diffeomorphisms. *Med. Image Anal.*, 25(1):56–71, 2015.
- [SWTM22] K. Stouffer, M. Witter, D. Tward, and M. Miller. Projective diffeomorphic mapping of molecular digital pathology with tissue MRI. *Communications Engineering*, 1(1):44, 2022.
- [SZ14] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [TDSL00] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- [Tom41] C. Tompkins. A flat Klein bottle isometrically embedded in Euclidean 4-space. *Bulletin of the American Mathematical Society*, 47(6):508, 1941.
- [Tro95] A. Trouvé. An infinite dimensional group approach for physics based models in pattern recognition. In *International Journal of Computer Vision*, 1995.
- [Tro98] A. Trouvé. Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision*, 28 (3):213–221, 1998.
- [TV12] A. Trouvé and F.-X. Vialard. Shape splines and stochastic shape evolutions : A second order point of view. *Quart. Appl. Math.*, 70(2):219–251, 2012.
- [TV19] R. Tahaoui and F.-X. Vialard. Minimizing acceleration on the group of diffeomorphisms and its relaxation. *ESAIM Control Optim. Calc. Var.*, 25, 2019.
- [TY05a] A. Trouvé and L. Younes. Local geometry of deformable templates. *SIAM J. Math. Anal.*, 37(1):17–59, 2005.
- [TY05b] A. Trouvé and L. Younes. Metamorphoses through Lie group action. *Found. Comput. Math.*, 5(2):173–198, 2005.
- [Via20] F.-X. Vialard. Variational Second-Order Interpolation on the Group of Diffeomorphisms with a Right-Invariant Metric. In *Mathematics Of Shapes And Applications*, pages 1–14. World Scientific, 2020.

- [VLBM08] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008.
- [VRRRC12] F.-X. Vialard, L. Risser, D. Rueckert, and C. J. Cotter. Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. *International Journal of Computer Vision*, 97:229–241, 2012.
- [VRRH12] F.-X. Vialard, L. Risser, D. Rueckert, and D. D. Holm. Diffeomorphic atlas estimation using geodesic shooting on volumetric images. *Annals of the BMVA*, 2012:1–12, 2012.
- [War83] F. W. Warner. *Foundations of differentiable manifolds and Lie groups*, volume 94 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1983. Corrected reprint of the 1971 edition.
- [WCP17] B. Wen, X. Chen, and T. K. Pong. Linear convergence of proximal gradient algorithm with extrapolation for a class of nonconvex nonsmooth minimization problems. *SIAM Journal on Optimization*, 27(1):124–145, 2017.
- [Wei98] J. Weickert. *Anisotropic diffusion in image processing*. European Consortium for Mathematics in Industry. B. G. Teubner, Stuttgart, 1998.
- [Whi92] H. Whitney. Collected papers, Volume II, (James Eells, Domingo Toledo, eds.), 1992.
- [Wil91] D. Williams. *Probability with martingales*. Cambridge university press, 1991.
- [WJW⁺22] Y. Wu, T. Z. Jiahao, J. Wang, P. A. Yushkevich, M. A. Hsieh, and J. C. Gee. NODEO: A neural ordinary differential equation based optimization framework for deformable image registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20804–20813, 2022.
- [WPZ⁺09] A. Wedel, T. Pock, C. Zach, H. Bischof, and D. Cremers. An improved algorithm for $TV - L^1$ optical flow. In *Statistical and geometrical approaches to visual motion analysis*, pages 23–45. Springer, 2009.
- [Y⁺18] B. Yu et al. The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
- [YHPC18] T. Young, D. Hazarika, S. Poria, and E. Cambria. Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, 13(3):55–75, 2018.
- [YKSN17] X. Yang, R. Kwitt, M. Styner, and M. Niethammer. Quicksilver: Fast predictive image registration—a deep learning approach. *NeuroImage*, 158:378–396, 2017.
- [You07] L. Younes. Jacobi fields in groups of diffeomorphisms and applications. *Q. Appl. Math*, pages 113–134, 2007.
- [You10] L. Younes. *Shapes and diffeomorphisms*, volume 171 of *Applied Mathematical Sciences*. Springer-Verlag, Berlin, 2010.
- [YYSP21] L. Yonghyeon, S. Yoon, M. Son, and F. C. Park. Regularized autoencoders for isometric representation learning. In *International Conference on Learning Representations*, 2021.
- [ZF19] M. Zhang and P. T. Fletcher. Fast diffeomorphic image registration via Fourier-approximated Lie algebras. *International Journal of Computer Vision*, 127(1):61–73, 2019.
- [Zho20] D.-X. Zhou. Universality of deep convolutional neural networks. *Applied and computational harmonic analysis*, 48(2):787–794, 2020.
- [ZLC⁺21] C. Zhang, J. Liu, W. Chen, J. Shi, M. Yao, X. Yan, N. Xu, and D. Chen. Unsupervised anomaly detection based on deep autoencoding and clustering. *Security and Communication Networks*, 2021, 2021.
- [ZYY⁺18] Y. Zhan, J. Yu, Z. Yu, R. Zhang, D. Tao, and Q. Tian. Comprehensive distance-preserving autoencoders for cross-modal retrieval. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1137–1145, 2018.