

Optimal Transport for Measure and Image Interpolation and for Information Design

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

vorgelegt von
Jorge Andrés Justiniano Nava
aus
Santa Cruz de la Sierra, Bolivien

Bonn, 2025

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

Gutachter/Betreuer: Prof. Dr. Martin Rumpf
Gutachter: Prof. Dr. Matthias Erbar

Tag der Promotion: 20.03.2025
Erscheinungsjahr: 2025

Contents

List of Figures	VII
List of Tables	XI
1 Introduction	3
2 Preliminaries	13
2.1 Function Spaces on Euclidean Spaces	13
2.2 Convergence of Sequences in Function Spaces	14
2.2.1 Inequalities and Embeddings	15
2.3 Convergence of Measures and Related Theorems	17
2.4 Γ - and Mosco-convergence	18
3 Approximation of Splines in Wasserstein Spaces	21
3.1 Background	21
3.1.1 Review of Optimal Transport	21
3.1.2 Wasserstein spaces as a Riemannian Manifold	22
3.2 Splines in Wasserstein Spaces	23
3.2.1 Definition of splines	23
3.2.2 Variational discretization of splines	26
3.3 Gaussian E-splines	31
3.3.1 The case of general Gaussian distributions	32
3.3.2 The case of Gaussian distributions with diagonal covariance matrices	38
3.4 Convergence of discrete Gaussian E-splines	40
3.5 Fully discrete Wasserstein splines and numerical results	43
3.5.1 Algorithmic foundations	43
3.5.2 Numerical results	44
3.6 Generative texture synthesis based on Wasserstein spline interpolation of feature distributions	47
3.7 Generative spline interpolation of data aided by variational autoencoders	51
3.8 Variational-time Wasserstein regression	52
4 Unbalanced Transport Splines with Source Term	59
4.1 Review of the Flow of Diffeomorphisms model and Optimal Transport Model with Source Term	59
4.1.1 Flow of Diffeomorphisms	59
4.1.2 Generalized Optimal Transport Model with Source Term	61
4.2 Time Continuous Splines in the Generalized Model	63
4.3 Variational Time Discretization	69

4.4	Temporal Extension Operators	75
4.5	Convergence of Discrete Splines	78
4.6	Fully Discrete Metamorphosis Splines	90
4.7	Numerical Optimization Using the iPALM Algorithm	92
4.8	Applications	93
5	Entropy-Regularized Optimal Transport in Information Design	97
5.1	Information Design and the Moment Bayesian Persuasion Problem	97
5.2	The optimization task	99
5.3	Semi-discrete optimal transport revisited	104
5.4	Entropy regularization	105
5.5	Spatial discretization	112
5.6	Numerical experiments	116
5.7	Application to information design	118
5.8	Applications to Gerrymandering	121
6	Summary	127
	Bibliography	133

List of Figures

- 1.1 A comparison of the different spline models sampled at nine equidistant times in 1D for Gaussian probability distributions as interpolation constraints depicted in grey. Sampled random variables are drawn as black dots, and their optimal sample trajectories are depicted by the connecting black curves: top left: continuous P-spline (red); top middle: continuous E-spline/T-spline (orange) sampled at nine equidistant times; top right: standard deviations for both the P-spline (red) and E-spline/T-spline (orange). Orange dots represent the discrete values obtained with our method; bottom left: continuous T-spline (blue) sampled at nine equidistant times; bottom middle: continuous E-spline (orange) sampled at nine equidistant times; bottom right: standard deviations for both the T-spline (blue) and E-spline (orange). Orange dots represent the discrete values obtained with our method. 8

- 3.1 First two rows: A comparison of discrete piecewise geodesic interpolation (first row) and discrete spline interpolation (second row) for $\delta = 0$ is shown for key frame distributions framed in red. The optimization was done on $\mathcal{P}_2^{G,d}$. Third row: Same as second row, except the optimization was performed in the full space $\mathcal{P}_2(\mathbb{R}^d)$. Bottom row: Difference between spline and piece-wise geodesic interpolations. Top right: Plot of the center of masses as a polygonal curve in \mathbb{R}^2 . Bottom right: Plot of the standard deviations as a polygonal curve in \mathbb{R}^2 45



- 3.2 Discrete piecewise geodesic interpolation (top) and discrete spline interpolation for $\delta = 0$ (middle) of three key frames with constant density on an annulus for the first and constant density on a disk for the second and third (framed in red). Bottom: Contribution of each time-step $k = 1, \dots, K - 1$ to the spline energy, i.e. $\mathcal{W}^2(\mu_k, \text{Bar}(\mu_{k-1}, \mu_{k+1}))$ for the spline interpolation (orange) and piecewise geodesic interpolation (green). 46

- 3.3 Piecewise geodesic (top) and spline interpolation (middle) are shown for key frames (framed in red) consisting of a thin annulus-shaped distribution and two equal thin square-shaped distributions, using the color map $0 \text{ } \text{---} \text{ } 4\text{e-}4$. Bottom: Difference between the spline and piecewise geodesic interpolations, using the color map $-6\text{e-}5 \text{ } \text{---} \text{ } 6\text{e-}5$. 47

- 3.4 The key frames represent two Gaussians that are far apart from each other (first and fourth key frames) and close to each other (second and third key frames). Piecewise geodesic (top) and spline (bottom) interpolations are shown. 47

- 3.5 From left to right the key frames represent a single Gaussian, a pair of vertically displaced Gaussian of half the mass, and the vertically displaced configuration rotated by $-\frac{\pi}{4}$. Discrete piecewise geodesic (top) and spline (bottom) interpolations are shown. 48

3.6	Time t (in seconds) until convergence of the fully discrete spline interpolation problem is reached, for a series of five interpolation problems $P(x)$ depending on parameter $x \in \{0, 1, 2, 3, 4\}$ (x -axis). The interpolation problem $P(x)$ is defined as follows: The prescribed times are $\bar{t}_0 = 0$, $\bar{t}_1 = 0.5$ and $\bar{t}_2 = 1$, and the prescribed probability measures are given by $\bar{\mu}_0 = \mathcal{N}((0, 0), \sigma_0^2)$, $\bar{\mu}_1 = \mathcal{N}((x, x), \sigma_1^2)$, and $\bar{\mu}_2 = \mathcal{N}((0, 0), \sigma_2^2)$, for diagonal standard deviation matrices $\sigma_0 = \text{diag}(1, 2)$, $\sigma_1 = \text{diag}(1, 1)$ and $\sigma_2 = \text{diag}(2, 1)$. Dots denote the computation time of the algorithm solving problem $P(x)$, both with implementation of the decoupling of the means as described by equation (3.2.9) (orange), and without it (green).	48
3.7	Schematic plot of our algorithm to generate texture interpolations based on Wasserstein spline interpolation of feature distributions. The step from the first row to the second one corresponds to the extraction of empirical feature distributions. The step from the second row to the third describes the use of our interpolation algorithm to compute feature distributions in-between. The final step (from the third to the fourth row) represents mapping back features into textures.	49
3.8	Time discrete metamorphosis (top) and Wasserstein spline interpolations (bottom) with framed prescribed images/feature distributions for $K = 16$. For the metamorphosis spline the key frames are chosen identical to the synthesized texture of the Wasserstein splines. Due to symmetry, only the first half of the interpolations is shown.	50
3.9	Five spline interpolations in the MNIST-Dataset. Each of them has 11 frames; the fixed keyframes are at $k = 0$, $k = 5$, and $k = 10$ for each row.	50
3.10	Two realizations of a texture spline for different starting latent space samples, top and bottom of each panel respectively with parameters $K = 20$, $\delta = 0.01$. Let us remark that not only the actual spline interpolated textures but also the key frame textures differ as they are all different samples of the underlying spline probability distributions ν_k^K .	55
3.11	Top: A realization of a texture spline with parameters $K = 20$, $\delta = 0.01$. Bottom: Textures sampled at the interpolated times of the actual video from which the texture constraints were extracted from. As in Figure 3.10 the key frame textures from the spline interpolated path differ from the true images at the corresponding times as they are samples of the underlying spline probability distributions ν_k^K .	56
3.12	Schematic plot of our algorithm to generate data interpolations based on Wasserstein spline interpolation on the latent space of a VAE. The step from the first row to the second one corresponds to the encoder. The step from the second row to the third describes the use of our interpolation algorithm to compute Gaussian distributions in-between. The next step illustrates the sampling algorithm, and the final step (from the fourth to fifth row) represents the decoder part of the VAE.	57
4.1	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 $\theta^K[\Theta^K, \Phi^K](\cdot, x)$ along a path $(y_t^K(x))_{t \in [0, 1]}$, plotted against time. Dots represent the values ϑ_k^K , $k = 0, \dots, K$, and crosses the “half-way” values $\frac{1}{2}(\vartheta_k^K + \vartheta_{k-1}^K)$, $k = 1, \dots, K$, along the discrete transport path.	77
4.2	Top: Classic Wasserstein spline between three Gaussian distributions. Bottom: Unbalanced transport spline between three re-scaled Gaussian distributions.	94

4.3	Top: From left to right, the key frames represent a single Gaussian, a pair of vertically displaced Gaussians of half the mass, and the vertically displaced configuration rotated by $-\frac{\pi}{4}$. Bottom: From left to right, the key frames represent a single Gaussian, a pair of vertically displaced Gaussians of with each having the same mass as the left, and the vertically displaced configuration rotated by $-\frac{\pi}{4}$	95
4.4	Unbalanced transport spline (top) and classic Wasserstein spline (bottom) of three key frames with constant density on an annulus for the first and constant density on a disk for the second and third (framed in red). The keyframes at the top have different masses.	95
5.1	Convergence of the optimal power diagram for different entropy parameters $\varepsilon = 25, 5, 1, 0.2$. (four right-most plots), where the blur parameter values are given in units of $1/N$, for $N = 128$. The function Φ is plotted on the left-most panel.	117
5.2	Convergence of the optimal power diagram for a function Φ (column (i)) with global maxima at $(0.25, 0.75)$, $(0.75, 0.75)$ and $(0.5, 0.25)$ for regularization parameters given by $\eta = 1e-1, 1e-2, 1e-3, 1e-4, 1e-5$ (columns (ii)-(vi)).	117
5.3	Convergence of the optimal power diagram for a concave function Φ (left). One sees how the algorithm automatically pushes cells outside of the relevant unit square to enact the trivial solution. Plotted are the computed power diagrams after iterations $it = 1, 2, 4, 8, 16$ (five right-most plots).	118
5.4	Optimal configurations for the monopolist's problem with unit demand, with prices $p_1 = 1$ and $p_2 = 1, 1.25, 1.5, 1.75, 2$ (second row, from left to right) and quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$. The respective revenue function R for each case is plotted in the first row.	120
5.5	Optimal configurations for the monopolist's problem with unit demand, prices $p_1 = 1$ and $p_2 = 1.25$, upper quality bound $\bar{q} = 2$ and, from left to right, lower quality bounds $\underline{q} = 0.25, 0.5, 0.75, 1, 1.25$ (second row). The respective revenue function R (defined above) is plotted in the first row.	121
5.6	Optimal configurations for the monopolist's problem with additive valuations, with prices $p_1 = p_2 = 1$, quality bounds $\underline{q} = 0$ and $\bar{q} = 2$ and, from left to right, bundling surcharges/discounts $\Delta = -0.75, -0.5, -0.375, 0, 0.375$ (second row). The respective revenue function R (defined above) is plotted in the first row.	122
5.7	Plot of the Lloyd's centroidal power diagrams for $n = 1, 2, 4, 5, 7$ cells from left to right. The sites coincide with the barycenters of the cells.	123
5.8	Left: Population density ν map of Ohio. Middle-left: Political preference map \hat{P} , with values -1  $+1$. Middle: Map showing the current congressional map of Ohio with current majorities. Middle right: Optimal power diagram gerrymandering (from the point of view of the Republican party) of Ohio. Right: Optimal power diagram gerrymandering (from the point of view of the Democratic party) of Ohio. The colorbar represents the expected percentage of republican voters in the rightmost two figures: 0  100 . All computations were made with $n = 16$	126
6.1	Time discrete spline with framed fixed images (first and second row), first order material derivative slack variable \bar{z} (third and fourth row), second order material derivative \bar{w}_k (fifth and sixth row) and color-coded acceleration field \mathbf{a}_k (seventh and eighth row)	131

List of Tables

5.1	Monopolist's problem with unit demand, price $p_1 = 1$, quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$ for different prices p_2 (first row) under selected information policies.	120
5.2	Monopolist problem with unit demand, prices $p_1 = 1$, $p_2 = 1.25$, upper quality bound $\bar{q} = 2$ for different lower quality bounds $\underline{q} = 0.25, 0.5, 0.75, 1., 1.25$ (first row) under selected information policies.	121
5.3	Monopolist's problem with additive valuations for different bundling discounts Δ (first row), quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$ and prices $p_1 = p_2 = 1$, under selected information policies.	121
5.4	Monopolist's problem with additive valuations for different bundling surcharges Δ (first row), quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$ and prices $p_1 = p_2 = 1$ under selected information policies.	122

Acknowledgements

First and foremost, I wish to express my heartfelt gratitude to my advisor, Prof. Dr. Martin Rumpf, for introducing me to many fascinating topics in mathematics and for his exceptional guidance and unwavering support throughout my master's and doctoral studies.

I am deeply thankful to Prof. Dr. Matthias Erbar for being a great professional to work with and for his willingness to co-review this thesis, and to Prof. Dr. Sergio Conti and Prof. Dr. Sven Rady for serving on my doctoral committee.

I extend my sincere appreciation to my co-authors: Prof. Dr. Matthias Erbar, Dr. Marko Rajković, Prof. Dr. Andreas Kleiner, Prof. Dr. Benny Moldovanu, Prof. Dr. Martin Rumpf, and Prof. Dr. Philipp Strack. Their collaboration and insightful discussions greatly enriched this work.

I am particularly grateful to Prof. Dr. Philipp Strack and his team for warmly hosting me during my research stay in New Haven; without his support that would have not been possible.

Special thanks are due to Florine Hartwig and Christoph Smoch for proofreading parts of this thesis. I am also thankful to my current and former colleagues and the staff at the Institute for Numerical Simulation, University of Bonn: Dr. Martin Lenz, Dr. Behrend Heeren, Florine Hartwig, Christoph Smoch, Dr. Josua Sassen, and Carole Rossignol. Their valuable discussions, support, and the positive atmosphere they fostered in and out of the office have been a constant source of encouragement.

I am profoundly grateful to all of my teachers and everyone who has contributed to my education over the years.

Finally, I want to express my deepest gratitude to my family and friends for their unconditional love and enduring support. They are the foundation of all my achievements.

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, during my Doctoral studies, I was supported by the Bonn International Graduate School of Mathematics. During my Bachelor's studies, I was financially supported by the Undergraduate Excellence Scholarship of the DAAD, to whom I am also very grateful.

1 Introduction

The theory of optimal transport (OT) has experienced a surge in interest over recent years that can be attributed to several factors spanning new theoretical and computational advancements, the rise of data-driven applications and a growing applicability across diverse fields. Optimal transport is a mathematical framework initially introduced by Gaspard Monge in 1781 and later formalized by Leonid Kantorovich [Kan42] as an optimization tool for supply chain problems. Nowadays, it has become a cornerstone in various fields due to its ability to compare probability distributions and its connections to, among others, geometry, statistics, and machine learning.

One of the primary reasons for the growing interest in OT are the recent advances in computational methods. Aided by advances in GPU computing, efficient algorithms such as the Sinkhorn algorithm introduced by Cuturi [Cut13] and the sliced Wasserstein distance [KPT⁺17] have significantly extended its scope by making OT practical for high-dimensional data. Moreover, the proliferation of open-source software tools, such as POT [FC⁺21], OTT [CMPT⁺22] and GeomLoss [FSV⁺19], has made OT accessible to a broader audience. Some of the aforementioned tools are extensively used in this work.

Optimal transport has found extensive use in machine learning and data analysis, particularly in tasks involving distributional data. For instance, it serves as the foundation for Wasserstein generative adversarial networks (WGANs), where the Wasserstein distance improves training stability and convergence [ACB17]. In domain adaptation, OT facilitates the alignment of source and target distributions, enabling robust knowledge transfer across domains [CFTR17]. Additionally, it is applied in clustering [HNY17] and dimensionality reduction through barycenter computation [PC⁺19] and Wasserstein PCA [Big20]. This thesis further contributes to this field by leveraging our Wasserstein interpolation algorithm with the introduction of a variational Wasserstein regression algorithm in measure spaces (cf. Section 3.8) and a new data interpolation method aided by variational autoencoders (VAEs) (Section 3.7). In medical imaging, OT is used to compare and analyze cell distributions, such as mapping the spatial organization of gene expression patterns [S⁺19] and single-cell omics (cf. [WOS⁺10], [BSK⁺24]). Furthermore, it aids in tasks like shape analysis, where the Wasserstein distance quantifies differences between anatomical structures [HZTA04]. Owing it to the Benamou-Brenier formulation of OT [BB00] OT is used in physics to study phenomena like mass transport quantum mechanics, and density functional theory [B⁺12]. In statistics, OT provides a powerful tool for comparing distributions through metrics like the Wasserstein distance. It is extensively used in robust statistics and generative modelling [GPC18]. OT also aids in Bayesian inference by aligning posterior distributions [Z⁺21]. In natural language processing, OT helps in tasks such as word embedding alignment and topic modelling. The Wasserstein distance measures the similarity between word distributions, improving translation and semantic similarity tasks [KSKW15].

In computer vision, OT has gained prominence due to its ability to compare, interpolate, and transfer complex distributions. This versatility makes OT an effective tool for solving problems in image processing, shape analysis, and rendering, among other areas. It provides a principled way to transfer

color distributions [RDG14] and styles [PKD05] between images. In image segmentation, OT is employed by Courty et al. [CFT18] to model spatial dependencies and enhance segmentation performance in complex datasets. Similarly, the Wasserstein distance is used in texture mixing and color transfer between images [BRPP15]. It has also been utilized for texture synthesis by aligning feature distributions of textures. Peyré [Pey09] proposed a method using sliced Wasserstein distances for computational efficiency in high-dimensional texture applications. As shown in the work of Alaya et al. [AFC20], OT enables efficient and robust alignment of point clouds for applications in robotics and virtual reality. In motion estimation, OT models the transport of features across frames, aiding in tasks such as optical flow computation and video synthesis. Papadakis et al. [PMP14] explored OT-based motion estimation with promising results in video processing. OT has been employed in rendering applications to optimize sample distributions for Monte Carlo methods. Georgiev et al. [G⁺19] demonstrated how OT improves the efficiency of sampling strategies in physically-based rendering. Moreover, OT provides a natural framework for interpolating between images by finding geodesics in Wasserstein space. Bonneel et al. [BRPP15] applied OT to achieve transitions between images, widely used in computer graphics and animation. In this thesis, we contribute by introducing algorithms that leverage OT to compute discrete *interpolations of probability measures, textures and images* (cf. Chapters 3 and 4) that allow for smoother and more visually appealing transitions in time.

Optimal transport theory has also become a pivotal tool in economics, providing a rigorous framework for modelling resource allocation, market equilibria, and income inequality. By leveraging the mathematical foundation of OT, economists can address problems involving the transfer of resources and distributions in a variety of contexts: In matching theory, OT has been extensively applied to study markets where agents must be matched, such as job markets, marriage markets, and school assignments [Gal16]. It provides a robust method to analyze income inequality and intergenerational mobility. For instance, Ekeland et al. [EGH10] utilized OT to measure disparities and the redistribution of wealth across populations. By comparing distributions of wealth before and after policy interventions, OT aids in evaluating economic policies' effectiveness. In general equilibrium theory, OT facilitates the modelling of market equilibria by describing how resources are allocated across markets. Carlier et al. [CE03] demonstrated how OT can be used to compute equilibrium distributions in multi-market economies, providing insights into the efficiency and stability of resource allocations. Dynamic OT models have been employed to study international trade and economic growth. By modelling the evolution of resource distributions over time, Dessein and Szymanski [DS20] explored how OT can analyze trade patterns and the long-term effects of policy changes on economic growth. In urban economics, OT is used to study spatial equilibrium models, such as the distribution of housing, firms, and workers in cities. Ahlfeldt et al. [ARSW15] applied OT to model the spatial distribution of economic activity, helping to optimize urban planning and zoning policies. OT has applications in auction theory and mechanism design, where the goal is to design systems that allocate resources efficiently and fairly. For example, Ma [MF18] used OT to study revenue-maximizing auctions, showing how OT can address multi-objective optimization in auction mechanisms. Related to these topics, we offer in Chapter 5 computationally efficient algorithms aided by semi-discrete OT to compute optimal designs in *information design*, as well as giving a short additional overview of applications in gerrymandering.

Finally, the relevance of optimal transport does not limit itself to probability measures, and multiple generalizations of this framework have been proposed and have been gaining ground in the last few years. The quantum Wasserstein distance generalizes the classical Wasserstein distance to a setting where states are represented by density matrices instead of probability distributions. This framework is designed to account for the structure and dynamics of quantum systems, incorporating concepts such as entanglement

and non-commutativity, cf. [DPMTL21]. Adapted Optimal Transport (AOT) is a specialized extension of the classical optimal transport framework that incorporates temporal or sequential dependencies into the transport process. This generalization is particularly relevant in stochastic processes, financial mathematics, and control theory, where the transported measures evolve over time or are subject to dynamic constraints, cf. [BVBBE20]. The Wasserstein-Gromov distance, introduced in [Mém11] is a powerful mathematical framework that generalizes the optimal transport paradigm to compare structured datasets, such as graphs, networks, or metric spaces. While classical optimal transport measures the cost of aligning two probability distributions in a shared metric space, the Wasserstein-Gromov distance addresses situations where the underlying spaces of the distributions differ and no direct correspondence exists between their elements. These generalizations have greatly expanded the versatility of optimal transport, making it a powerful tool for applications in mathematics, physics, computer science, and beyond.

Measure and Image Interpolation

In recent years, advanced interpolation methods have gained significant attention for tasks like time-sequence interpolation or regression in data analysis. These methods are widely applied in fields such as computer graphics, computer vision, and medical imaging. The entities being interpolated are typically viewed as shapes within an infinite-dimensional manifold, equipped with a Riemannian metric tailored to the specific application.

One prominent approach is to extend the concept of path energy on a Riemannian manifold to a second-order spline energy functional. Given a sequence of objects (referred to as key frames) at different time points, the spline curve is determined by minimizing the spline energy while ensuring it passes through these key frames.

In Euclidean space, cubic splines $x : [0, 1] \rightarrow \mathbb{R}^d$ are well-known as the solutions that minimize the integral of squared acceleration, $\int_0^1 |\ddot{x}|^2 dt$, a result established by de Boor [dB63]. Our approach extends this concept to the Wasserstein space in discrete time. To achieve this, we approximate the integral using a rectangular quadrature rule and use second-order central differences to estimate the acceleration. This leads to the expression:

$$\int_0^1 |\ddot{x}|^2 dt \approx 4K^3 \sum_{k=1}^{K-1} \left| x_k - \frac{x_{k-1} + x_{k+1}}{2} \right|^2. \quad (1.0.1)$$

Numerous spline interpolation methods have been developed for non-linear spaces. Noakes *et al.* [NHP89] generalized de Boor’s result to finite-dimensional Riemannian manifolds by introducing Riemannian cubic splines. These are defined as paths that minimize the integrated squared covariant derivative of their velocity. Trouvé and Vialard [TV12] explored a second-order functional for shapes in landmark spaces, using an optimal control framework. Singh *et al.* [SVN15] proposed a functional measuring motion acceleration in diffeomorphic flows for image regression. Tahraoui and Vialard [TV19] introduced a second-order variational model for diffeomorphic flows, leading to a relaxed Fisher-Rao functional on the space of measures. Vialard [Via20] demonstrated the existence of a minimizer for the Riemannian acceleration energy on a group of diffeomorphisms with a high-order Sobolev metric.

For smooth interpolation of data distributions, Chen and Karlsson [CK18] studied an optimal control problem constrained by transport equations and interpolation conditions, particularly focusing on Gaussian distributions. Julien Clancy [Cla21] compared various spline techniques in probability measure spaces, incorporating entropy regularization and extending these methods to unbalanced measures.

Another higher-order approximation was proposed by Karimi and Georgiou [KG21]. They addressed a regression problem involving polynomial measure-valued curves, approximating distributional snapshots through a least-squares approach in Wasserstein space. This involved a multi-marginal optimal transport formulation and an efficient computation by using the Sinkhorn algorithm. Zhang and Noakes [ZN19] analyzed Riemannian cubic splines within the manifold of symmetric positive definite matrices, leveraging Lie algebra and the Riemannian geometry of Gaussian densities (Bures-Wasserstein manifold [Bur69], [FK16], [BJL17]) based on the Wasserstein distance [MMP18].

Rajković *et al.* [JRR23] investigated spline interpolation of images, adopting a model rooted in the metamorphosis framework. This approach distinguishes between the Eulerian flow acceleration and the second material derivative of image intensity. Unlike Riemannian splines [NHP89, TV19, Via20], this model does not minimize the squared covariant derivative of the velocity. Instead, it avoids intertwining different acceleration types. Their study included a rigorous Mosco convergence analysis [Mos69], a stronger form of Γ -convergence, for the time-discrete to time-continuous model. They also proved the existence of minimizers for the continuous spline energy.

On the front of measure interpolation, several methods have been introduced in recent years. Notably, Benamou *et al.* [BGV19] and Chen *et al.* [CCG18] independently developed an approach now referred to as P-splines (short for path splines). P-splines are stochastic processes $(X_t)_{t \in [0,1]}$ valued in \mathbb{R}^d , defined on a probability space (Ω, \mathbb{P}) . These processes are solutions to the following optimization problem:

$$\min_{(X_t)_t} \int_0^1 \int_{\Omega} \|\ddot{X}_t\|^2 d\mathbb{P} dt, \quad (1.0.2)$$

subject to a set of $I \geq 2$ prescribed marginal constraints, $X_{\bar{t}_i} \sim \bar{\mu}_i$, for specified times $0 = \bar{t}_1 < \dots < \bar{t}_I = 1$, where $\bar{\mu}_i$ are given probability measures for $i = 1, \dots, I$.

Numerically, solving this problem is computationally intensive. To address this, the problem is often relaxed using multi-marginal optimal transport with a quadratic cost and entropic regularization. However, this method has a notable limitation: solutions to the problem above may not always be deterministic. Specifically, there is no guarantee of the existence of a Monge map $\phi_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $X_t = \phi_t(X_0)$, even if the marginal constraints are smooth and well-behaved, as noted by Chewi *et al.* [CCLG⁺21].

To overcome this limitation, Chewi *et al.* [CCLG⁺21] introduced an alternative framework known as transport splines (T-splines). This method focuses on the smooth interpolation of probability measures within the optimal transport setting, using a particle flow approach. Their work also explores the connection between transport splines and energy splines, particularly for Gaussian distributions in one dimension.

From a theoretical perspective, both P-splines and T-splines adopt a Lagrangian view of optimal transport. Instead of working directly with probability measures, these approaches model stochastic processes X_t with C^2 smooth sample paths and corresponding laws μ_t . This Lagrangian viewpoint is advantageous in certain applications, as it allows for the tracking of particle trajectories in continuous time. The cited works also provide algorithms for computing sample trajectories of both P-splines and T-splines, further contributing to their practical applicability. On the other hand, E-splines are based on the Eulerian perspective of optimal transport, focusing on the evolution of densities and particle velocities through fixed points in time and space.

Splines in Wasserstein Spaces

In Chapter 3 we will study a time discrete variational model to compute spline paths in the space of probability measures equipped with the Wasserstein-2 metric. The spline paths are defined as measure-valued paths minimizing a spline energy subject to interpolation constraints and boundary conditions.

In light of (1.0.1), it is natural to replace the Euclidean norm $|\cdot|$ with the Wasserstein L^2 distance between probability measures, and use a notion of barycenter between measures μ_{k-1} and μ_{k+1} , denoted as $\text{Bar}(\mu_{k-1}, \mu_{k+1})$, instead of the middle point $\frac{x_{k-1}+x_{k+1}}{2}$. The proposed discrete spline functional studied in Chapter 3 will therefore be given by

$$4K^3 \sum_{k=1}^{K-1} \mathcal{W}^2(\mu_k, \text{Bar}(\mu_{k-1}, \mu_{k+1})). \quad (1.0.3)$$

To define continuous splines within the space of probability measures, as studied in this work, an analogous geometric framework is necessary. Fortunately, the renowned Benamou-Brenier formula [BB00] provides the foundation for formally equipping the Wasserstein space $\mathcal{P}_2(\mathbb{R}^d)$ with a Riemannian structure, as initially described in [Ott01]. However, extending this analysis to incorporate second-order path properties in Wasserstein spaces is essential. This has been systematically developed by Gigli in [Gig12], where the acceleration of a measure-valued curve is defined as the covariant derivative of its velocity.

Energy splines (E-splines) are then characterized as the minimizers of the total squared acceleration, as formalized in (3.2.2). However, this functional poses significant challenges: it is computationally intractable and, unlike the standard action functional (3.1.2), is non-convex.

In this chapter, we introduce a consistent variational time discretization for E-splines and outline an algorithm for their construction. We demonstrate that this approach aligns with the Riemannian geometry of the Wasserstein space in the Gaussian case. Additionally, we construct simple counterexamples in one dimension where E-splines differ from P-splines and/or T-splines. In both cases, our method accurately and precisely reproduces the theoretical E-spline values up to machine precision, as shown in Fig. 1.1.

Our discretization method relies on the general variational framework for time-discretized splines on Riemannian manifolds, as proposed in [HRW18]. The key components of this approach are a functional \mathcal{W} , approximating the squared Riemannian distance between nearby objects on the manifold, and a suitable notion of an approximate average. This methodology has been previously applied in [HRS⁺16] to achieve smooth interpolation of triangulated surfaces in computer graphics, leveraging the concept of discrete thin shells. In the context of probability measures, the local functional \mathcal{W} corresponds to the squared Wasserstein distance, while the approximate average is represented by the Wasserstein barycenter.

Chapter 3 is an extended version of the journal article [JRE24], published in the peer-reviewed journal ESAIM: Control, Optimisation and Calculus of Variations, to which the author of this thesis contributed as its main author. In addition, a few additional numerical examples are computed, and an application to discrete Wasserstein regression was added.

Organization. Chapter 3 is structured as follows: Section 4.1 provides a concise overview of the Wasserstein distance between probability measures, introduces the Riemannian framework underpinning Wasserstein spaces, and discusses the flow-based formulation of optimal transport. In Section 3.2, we derive the time-continuous spline energy through the Riemannian lens and introduce a variational time discretization of the continuous spline energy. Section 3.3 delves into the specific case of Gaussian

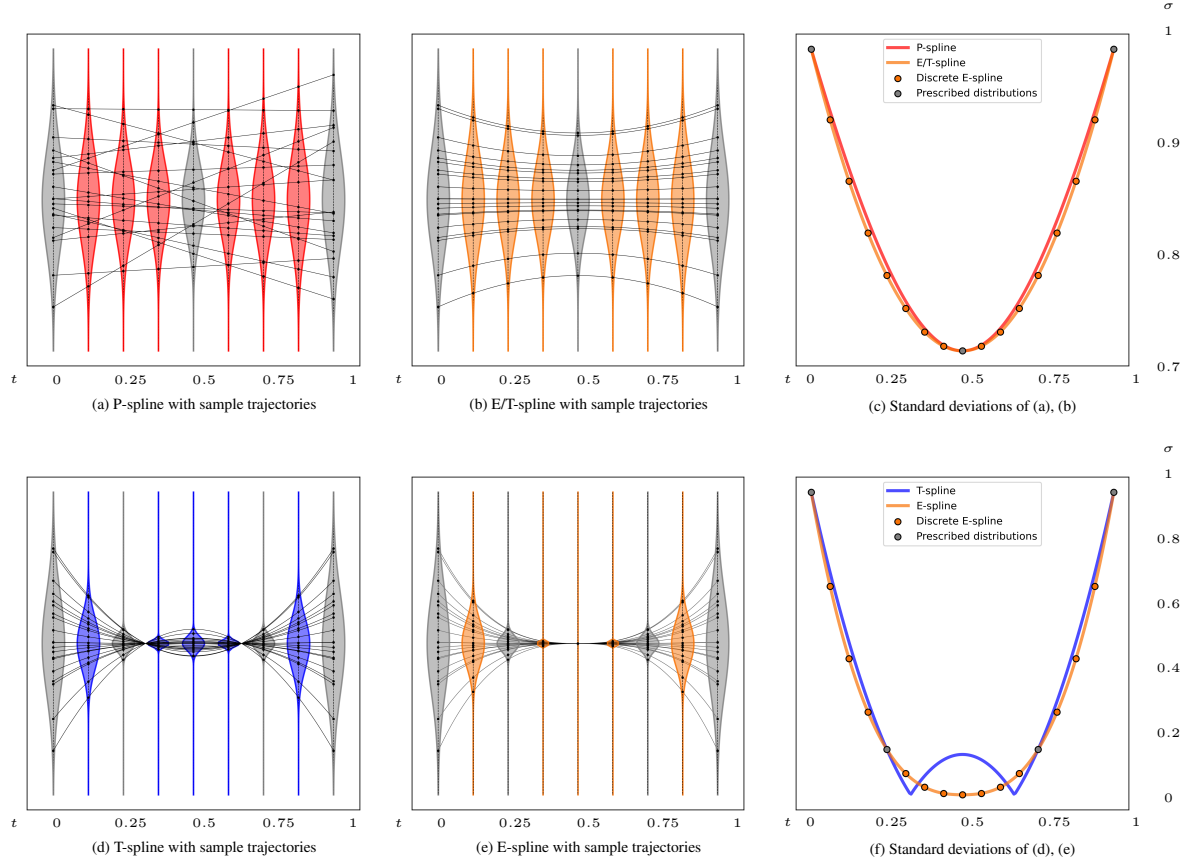


Figure 1.1: A comparison of the different spline models sampled at nine equidistant times in 1D for Gaussian probability distributions as interpolation constraints depicted in grey. Sampled random variables are drawn as black dots, and their optimal sample trajectories are depicted by the connecting black curves: top left: continuous P-spline (red); top middle: continuous E-spline/T-spline (orange) sampled at nine equidistant times; top right: standard deviations for both the P-spline (red) and E-spline/T-spline (orange). Orange dots represent the discrete values obtained with our method; bottom left: continuous T-spline (blue) sampled at nine equidistant times; bottom middle: continuous E-spline (orange) sampled at nine equidistant times; bottom right: standard deviations for both the T-spline (blue) and E-spline (orange). Orange dots represent the discrete values obtained with our method.

distributions, establishing the consistency of the discrete functional with the continuous one in this setting. Section 3.4 focuses on demonstrating the convergence of temporally extended discrete Wasserstein spline energies to their time-continuous counterparts. This is done in the sense of Mosco convergence, specifically for Gaussian distributions with diagonal covariance matrices. In Section 3.5, we outline the fully discrete scheme, which leverages the Sinkhorn algorithm and entropy-regularization to achieve computational efficiency. Additionally, we describe how to implement suitable variants of accelerated gradient algorithms [Nes05] to compute spline interpolations for a given set of key frames. This section also includes experimental results showcasing the application of the algorithm to probability measures. Section 3.6 illustrates the flexibility and robustness of our method through experiments, applying spline interpolation to image and texture interpolation tasks. In Section 3.7, we combine our interpolation method with variational autoencoders to tackle the problem of general data interpolation. Finally, we give a brief overview of discrete linear regression in Wasserstein spaces in Section 3.8.

Unbalanced Transport Splines

The main aim of Chapter 4 is to extend the results obtained in Chapter 3 for more general measure spaces, in regards to total mass. The motivation for this extension is rooted on two key observations in practical applications:

- *Handling mass differences in images:* In Chapter 3, a necessary condition for the spline interpolation was that all the measures have equal mass. In many cases, objects or structures within images that correspond to one another may have varying masses. Treating images strictly as probability distributions often proves too limiting on a global scale. This approach requires prior contrast adjustment, which can feel artificial. In classical optimal transport models, local mass imbalances result in artifacts, as excess mass must be arbitrarily redistributed without preserving structural correspondence. To address this, we relax the mass-preservation constraint by allowing for a source term in the path energy, explicitly accounting for density modulation.
- *Topological consistency and dissipative effects:* Unlike the flow of diffeomorphisms approach, optimal transport maps are not inherently homeomorphic. However, in many scenarios, maintaining topological consistency is critical. Simultaneously, the physical context of certain applications may necessitate incorporating a dissipative term into the path energy. To accommodate these requirements, we augment the classical transport cost framework with a weighted viscous dissipation term.

A few relaxations of the concept of optimal transport have been developed, which relax the mass conservation constraint. This extension is *unbalanced optimal transport*. This is achieved by introducing penalization terms for mass creation or destruction, allowing the model to handle measures with different total masses. For instance, the unbalanced transport framework has been formalized through the Kantorovich-Fisher-Rao distance [CPSV18], combining ideas from optimal transport and the Fisher-Rao metric.

In Chapter 4, we will recall the unbalanced optimal transport model originally introduced in [MRSS15]. Therein, a viscous dissipation term originating from the flow of diffeomorphisms model is used as a regularizer, which ensures that the maps remain orientation-preserving diffeomorphisms. On that basis, the basic optimal transport and the flow of diffeomorphisms models are then combined into a generalized image transport model for measures with different masses. We then derive the novel generalized spline energy functional as a natural extension of the path energy functional in [MRSS15]. Given a set of key

frames at disjoint times a spline curve is then given as a minimizer of this generalized spline energy respecting the key frame images as interpolation constraints. Furthermore, we will study a suitable time discrete variational model, which generalizes the time discrete first-order image transport model proposed in [MRSS15]. The central contribution of this paper is the consistency and Mosco convergence of this time discrete spline energy to the time continuous image transport spline energy. Finally, we discretize the model in space and we derive an efficient numerical scheme to solve for fully discrete image transport spline paths in the space of image densities, and present some numerical experiments.

Chapter 4 is not based on any publication, though we adhere to the same structure and incorporate multiple flow-related theoretical results obtained in [JRR21] and [JRR23], publications to which the author of this thesis contributed as one of the main authors.

Organization. In Section 4.1, we revisit the key properties of the optimal transport model and the flow of diffeomorphism model.

Section 4.2 is dedicated to deriving the time-continuous spline energy. It also explores the crucial interaction between the Lagrangian and Eulerian perspectives within the framework of the combined model of unbalanced transport and the flow of diffeomorphisms.

In Section 4.3, we introduce a variational time discretization of the continuous spline energy and investigate the existence of discrete splines.

Section 4.6 details the fully discrete scheme, while Section 4.7 describes the setup of an appropriate iPALM algorithm to numerically compute a spline interpolation for a given set of key frames.

Finally, Section 4.8 presents experimental results that illustrate the characteristics of the spline approach for optimal transport with source term and demonstrates applications and advantages of the proposed method.

Optimal Transport in Information Design

Mechanism Design is a fundamental branch of Game Theory and Economics that examines the creation of optimal institutions (often termed *game-forms*) to govern the interactions of self-interested, strategic agents possessing private information. The goal of an optimal mechanism is to resolve two critical challenges: the aggregation of information and the alignment of incentives, enabling the realization of a desired collective decision (see the introductory textbook [Bör15] for an overview). This field has found significant applications in social choice [Ser04] and voting systems [LW18], market design such as auctions [JMS99] and matching, and contracting [Wat07]. Theoretical and applied advancements in this area have garnered several Nobel Prizes, including recognition of Myerson’s seminal work on optimal auctions [Mye81].

In classical analyses of mechanism design, agents’ private information about the environment is treated as exogenously determined. A more recent line of research, termed *Information Design*, shifts the focus to endogenizing the information structure while considering the governing institution as fixed. In this framework, the information structure is strategically designed to achieve a predefined objective.

A prominent subfield within Information Design is the study of the *Bayesian Persuasion* problem. A notable and highly relevant subclass is the moment Bayesian persuasion problem, explored in works such as Kamenica and Gentzkow [KG11], Kolotilin [Kol18], and Dworczak and Martini [DM19]. Most research in this area has been constrained to one-dimensional settings, with only a few exceptions addressing multi-dimensional cases.

In Chapter 5, we address this gap by proposing an efficient algorithm to solve multi-dimensional moment Bayesian persuasion problems within certain subclasses. Specifically, we focus on settings where a sender designs an information policy to influence a receiver’s actions, such as making a purchase decision, thereby maximizing the sender’s utility.

In mathematical terms, the sender has full knowledge of the realization of a prior (probability measure) ν . She then decides how much information to give to an uninformed receiver to maximize her utility

$$\max_{\lambda \in F_\nu} \int_D \Phi(y) d\lambda(y). \quad (1.0.4)$$

This information policy can be encoded by another probability measure λ , carrying less information than ν , belonging to a subset of permissible measures F_ν . Due to the properties of the set F_ν , the functional in (1.0.4) and results by Klee [KJ58], it suffices to consider only exposed points of F_ν in the above formulation. The analysis in this chapter further builds upon the work of Kleiner et al. ([KMS21], [KMSW24]), which characterizes a large subset of finitely-supported, exposed information policies λ being induced by power diagram partitions of the domain. To solve this problem, we employ entropy-regularized optimal transport, enabling the development of an efficient computational framework. Numerical experiments are presented, showcasing the structural properties of the optimal configurations.

Furthermore, we extend the application of our framework to address monopolist problems involving multiple products, where the sender provides information about product qualities. Additionally, we explore the versatility of this approach by applying it to problems such as gerrymandering, yielding satisfactory results.

This chapter draws upon the preprint [JKM⁺25], to which the author of this thesis contributed as the main author. A condensed version of this work is currently under submission. Compared to the pre-print, this thesis includes additional complete proofs, extended numerical results, and an entirely new section dedicated to applying the framework to the gerrymandering problem, showing its versatility.

Organization. The structure of this chapter is as follows. Section 5.1 introduces the moment Bayesian persuasion model, emphasizing the role of optimizing over the closure of Lipschitz-exposed points in a subset of measures induced by information policies. This leads to a focus on measures represented by Laguerre partitions.

In Section 5.2, the optimization problem is reformulated over Laguerre partitions, accompanied by a relaxation strategy ensuring the existence of optimal designs. Section 5.3 reviews the connection between Laguerre diagrams and semi-discrete optimal transport, while Section 5.4 derives the entropy-regularized optimal transport formulation. In this setting, we prove the convergence of maximizers of the relaxed problems to maximizers of the non-relaxed problem. This formulation facilitates an efficient computational algorithm for solving the moment Bayesian persuasion problem.

Section 5.5 provides a detailed account of the spatial discretization process, ensuring reproducibility of computational results. Therein, the convergence of discrete maximizers to maximizers of the original functional is also studied. Section 5.6 illustrates the algorithm’s qualitative properties, and Section 5.7 applies the algorithm to compute optimal information policies for a multi-product monopolist. Finally, Section 5.8 briefly introduces the gerrymandering problem, along with a model thereof, and outlines ways to leverage our algorithm to solve that type of problems.

References

- [1] M. Erbar, J. Justiniano and M. Rumpf. Approximation of splines in Wasserstein spaces. In *ESAIM: Control, Optimisation and Calculus of Variations*, 30, p.64, 2024. <https://doi.org/10.1051/cocv/2024008>
- [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. Justiniano, A. Kleiner, B. Moldovanu, M. Rumpf and P. Strack. Entropy-Regularized Optimal Transport in Information Design. In *Proceedings of the 26th ACM Conference on Economics and Computation*, page 741–760, 2025. <https://doi.org/10.1145/3736252.3742617>

2 Preliminaries

Throughout this thesis, we assume an understanding of undergraduate calculus, i.e. multi-variate differential and integral calculus. Beyond this, we introduce in this chapter some basic notation and recall general definitions and relevant results. In this chapter, we recall the basic definitions of function classes on Euclidean spaces and their properties, outline some relevant results from measure theory, and introduce the notions of convergence that will be the subject of study in the central theoretical results of this thesis.

2.1 Function Spaces on Euclidean Spaces

In this section, we recall the definitions and fundamental properties of classical function spaces on Euclidean spaces. This includes key inequalities, embeddings, and convergence results. For detailed proofs, we refer the reader to classical texts on functional analysis, partial differential equations, and measure theory, such as [AF03, Alt16, Bre11, Leo24, Eva18].

Let $f : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be a sufficiently smooth function. For j -times partial derivatives in the direction of x_i , we write $\partial_{x_i}^j$. Using a multi-index $\beta = (\beta_1, \dots, \beta_n) \in \mathbb{N}^n$, we define:

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

For $k \in \mathbb{N}$, the *Banach space of continuous or continuously differentiable functions* $C^k(\overline{\Omega})$ on a compact set $\overline{\Omega} \subset \mathbb{R}^n$ is equipped with the norm $\|\cdot\|_{C^k(\overline{\Omega})}$ and seminorm $|\cdot|_{C^k(\overline{\Omega})}$:

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

We define $C^\infty(\overline{\Omega}) := \bigcap_{k \geq 0} C^k(\overline{\Omega})$. For $k \in \mathbb{N}$ and $\alpha \in (0, 1]$, the *space of Hölder continuously differentiable functions* $C^{k,\alpha}(\overline{\Omega})$ is a Banach space with the norm $\|\cdot\|_{C^{k,\alpha}(\overline{\Omega})}$ and seminorm $|\cdot|_{C^{k,\alpha}(\overline{\Omega})}$:

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

For simplicity, we set $C^{k,0}(\overline{\Omega}) = C^k(\overline{\Omega})$ for any $k \in \mathbb{N}$. If $k = 0$ and $\alpha = 1$, we call f a *Lipschitz function*, with $|\cdot|_{C^{0,1}}$ being the *Lipschitz constant*, often denoted L . If $\alpha > 1$, we interpret $C^{k,\alpha}$ as $C^{k+\lfloor \alpha \rfloor, \alpha - \lfloor \alpha \rfloor}$. For a deeper study of Hölder continuous functions, see [Fio17].

Definition 2.1.1 (Modulus of Continuity). A monotonically increasing function $\omega : [0, +\infty) \rightarrow [0, +\infty)$ is called a *modulus of continuity* for $f : \Omega \rightarrow \mathbb{R}$ if

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

The space of k -times differentiable functions on $\overline{\Omega}$ whose k -th derivative has modulus of continuity ω is denoted $C_{\omega}^k(\overline{\Omega})$. The subscript c always indicates compactly supported continuous or differentiable functions.

For *Lebesgue spaces*, the norm is given 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*, denoted $W^{m,p}(\Omega)$ with $m \in \mathbb{N}$ and $p \in [1, \infty]$, the norms and seminorms are:

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

We define $W_0^{m,p}(\Omega)$ as the closure of $C_c^\infty(\Omega)$ with respect to $\|\cdot\|_{W^{m,p}(\Omega)}$, and set $W^{0,p}(\Omega) = L^p(\Omega)$. For Sobolev spaces, we refer to [AF03, Alt16, Bre11].

Bochner Spaces

We extend the above notation to *Bochner spaces* of the form $X(I, Y)$, where I is an interval and X, Y are function spaces. The time derivative of a function f is denoted \dot{f} . For details, see [HvNV⁺16].

2.2 Convergence of Sequences in Function Spaces

In this section, we introduce fundamental concepts of convergence for sequences of functions in various function spaces. These notions are pivotal for the analysis throughout this thesis.

Definition 2.2.1 (Strong Convergence). Let $\{f^j\}_{j \in \mathbb{N}}$ be a sequence of functions in a normed space. We say that $\{f^j\}$ *converges strongly* to a function f if

$$\lim_{j \rightarrow \infty} \|f - f^j\| = 0,$$

where $\|\cdot\|$ is the norm defined on the function space.

Definition 2.2.2 (Weak Convergence). Let X be a function space and X' its dual space. A sequence $\{f^j\}_{j \in \mathbb{N}}$ is said to *weakly converge* to a function f , denoted by $f^j \rightharpoonup f$, if for every $\phi \in X'$,

$$\phi(f^j) \rightarrow \phi(f) \quad \text{as } j \rightarrow \infty.$$

For more details on weak convergence, we refer the reader to [Alt16, Chapter 8] and [Bre11, Chapter 3].

Theorem 2.2.3 (Reflexivity of Lebesgue and Sobolev Spaces). *Let $1 < p < \infty$. The function spaces $L^p(\Omega)$ and $W^{k,p}(\Omega)$ are reflexive, meaning that every bounded sequence in these spaces has a weakly convergent subsequence.*

Proof. See [Alt16, Section 8.8] and [Bre11, Section 3.5]. \square

Remark 2.2.4. A key application of compact embedding theorems (see Theorem 2.2.13) is the ability to transfer weak convergence in a function space to strong convergence in a larger space.

Definition 2.2.5 (Sequential Weak Lower Semicontinuity). A functional $F : X \rightarrow \mathbb{R}$ is said to be *sequentially weakly lower semicontinuous* if for every sequence $\{f^j\}_{j \in \mathbb{N}}$ with $f^j \rightharpoonup f$, it holds that

$$\liminf_{j \rightarrow \infty} F(f^j) \geq F(f).$$

Proposition 2.2.6. *Any functional that is strongly lower semicontinuous and convex is also weakly lower semicontinuous.*

Proof. See [Bre11, Proposition 3.5 and Corollary 3.9]. \square

In the subsequent sections, we recall some basic inequalities, embedding theorems, and convergence criteria that will play a significant role in the analysis of function spaces.

2.2.1 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.

Theorem 2.2.7 (Fatou's lemma, [Alt16, 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 2.2.8 (Dominated convergence theorem, [Alt16, 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 2.2.9 (Poincaré's inequality, [Alt16, 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 2.2.10 (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 2.2.11 (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*

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

where $(Df)^{sym} := \frac{1}{2}(Df + (Df)^T)$ 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 2.2.12 ([Alt16, 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 2.2.13 (Embedding theorem for Sobolev and Hölder functions, [Alt16, 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*

$$m_1 - \frac{n}{p_1} \geq m_2 - \frac{n}{p_2}, \quad m_1 \geq m_2, \quad (2.2.1)$$

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 (2.2.1) are strict, then the embedding is compact.

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

$$m - \frac{n}{p} \geq k + \alpha \quad \text{for any } \alpha \in (0, 1), k \in \mathbb{N} \quad (2.2.2)$$

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 $x \in \Omega$ and

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

If the inequality (2.2.2) 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 2.2.14 (Arzelà–Ascoli, [Alt16, 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.

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 2.2.15 ([Alt16, 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 $C^{l, \beta}([0, 1], C^{k, \alpha}(\overline{\Omega}))$.*

2.3 Convergence of Measures and Related Theorems

Definition 2.3.1. A *Polish space* is a topological space that is separable and completely metrizable.

Definition 2.3.2 (Narrow Convergence of Measures). Let $\{\mu_n\}_{n \in \mathbb{N}}$ be a sequence of Borel probability measures on a metric space (X, d) , and let μ be another Borel probability measure on X . The sequence $\{\mu_n\}_{n \in \mathbb{N}}$ is said to *converge narrowly* to μ (denoted $\mu_n \rightharpoonup \mu$) if:

$$\int_X f d\mu_n \rightarrow \int_X f d\mu \quad \text{for all } f \in C_b(X),$$

where $C_b(X)$ denotes the space of bounded continuous functions on X .

Theorem 2.3.3 (Portmanteau Theorem). *Let $\{\mu_n\}_{n \in \mathbb{N}}$ be a sequence of Borel probability measures on a metric space (X, d) , and let μ be another Borel probability measure on X . The following statements are equivalent:*

1. $\mu_n \rightharpoonup \mu$.
2. For every lower semi-continuous function $f : X \rightarrow \mathbb{R}$ that is bounded from below,

$$\liminf_{n \rightarrow \infty} \int_X f d\mu_n \geq \int_X f d\mu.$$

3. For every upper semi-continuous function $f : X \rightarrow \mathbb{R}$ that is bounded from above,

$$\limsup_{n \rightarrow \infty} \int_X f d\mu_n \leq \int_X f d\mu.$$

4. $\limsup_{n \rightarrow \infty} \mu_n(F) \leq \mu(F)$ for every closed set $F \subseteq X$.
5. $\liminf_{n \rightarrow \infty} \mu_n(G) \geq \mu(G)$ for every open set $G \subseteq X$.

Theorem 2.3.4 (Prokhorov's theorem, [Pro56]). *Let (X, τ) be a Polish space, and let $\{\mu_n\}_{n \in \mathbb{N}}$ be a sequence of probability measures on the Borel σ -algebra $\mathcal{B}(X)$. Then the following are equivalent:*

1. *The sequence $\{\mu_n\}_{n \in \mathbb{N}}$ is tight, i.e., for every $\varepsilon > 0$, there exists a compact set $K_\varepsilon \subset X$ such that*

$$\mu_n(X \setminus K_\varepsilon) < \varepsilon \quad \text{for all } n \in \mathbb{N}.$$

2. *The sequence $\{\mu_n\}_{n \in \mathbb{N}}$ has a subsequence that converges narrowly to a probability measure μ on $\mathcal{B}(X)$.*

2.4 Γ - and Mosco-convergence

A central theme across the first two main chapters of this thesis involves analyzing the convergence behaviour of functionals that are inherently discrete or non-local, towards a limiting functional that can be described as continuous or local in nature. To rigorously address this, we will introduce and formalize the specific definitions of convergence that underpin our analysis. Additionally, we will examine several key properties exhibited by sequences that demonstrate this type of convergence. For a deeper exploration of these concepts and their theoretical foundations, the reader is encouraged to consult the works in

Definition 2.4.1 (Γ - and Mosco-convergence). Let X be a metric vector space. Consider the sequence of functionals $\{\mathcal{F}^K\}_{K \in \mathbb{N}} : X \rightarrow \mathbb{R} \cup \{\infty\}$ and $\mathcal{F} : X \rightarrow \mathbb{R} \cup \{\infty\}$ that satisfies:

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

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

is valid,

- (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 “lim sup-inequality”

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

holds true.

Then $\{\mathcal{F}^K\}_{K \in \mathbb{N}}$ is said to Γ -converges to \mathcal{F} . In case X is a Banach space, some authors only require the recovery sequence to be weakly convergent, i.e. $x^K \rightharpoonup x$. On the other hand, if condition (ii) is left as above, and the first condition is changed to

- (i') for every sequence $\{x^K\}_{K \in \mathbb{N}} \subset X$ with $x^K \rightharpoonup x \in X$ the “weak lim inf-inequality”

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

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

Definition 2.4.2 (Equicoercivity, [Bra14, Definition 2.9]). Let X be a Banach space. A sequence of functionals $\{\mathcal{F}^K\}_{K \in \mathbb{N}} : X \rightarrow \mathbb{R} \cup \{\infty\}$ is said to be *equicoercive* if for all $r \in \mathbb{R}$ there exists a compact set $K_r \subset X$ such that $\{x \in X : \mathcal{F}^K[x] \leq r \ \forall K \in \mathbb{N}\} \subset K_r$.

Theorem 2.4.3 (Fundamental theorem of Γ -convergence, [Bra14, Theorem 2.10]). *Let X be a Banach space and $\{\mathcal{F}^K\}_{K \in \mathbb{N}} : X \rightarrow \mathbb{R} \cup \{\infty\}$ an equicoercive sequence of functionals that Γ -converges to $\mathcal{F} : X \rightarrow \mathbb{R} \cup \{\infty\}$. Then, it holds $\min_{x \in X} \mathcal{F}[x] = \lim_{K \rightarrow \infty} \inf_{x \in X} \mathcal{F}^K[x]$.*

As a result, for a sequence of minimizers $\{x^K\}_{K \in \mathbb{N}}$ of the functionals $\{\mathcal{F}^K\}_{K \in \mathbb{N}}$ satisfying the uniform energy bound $\mathcal{F}^K[x^K] \leq r$, the property of equi-coercivity ensures the existence of a subsequence that converges to some $x \in X$. Moreover, the Γ -convergence of \mathcal{F}^K to \mathcal{F} implies that x is a minimizer of \mathcal{F} . In the case of Mosco-convergence, the requirement of compactness for the set K_r can be replaced by weak compactness.

3 Approximation of Splines in Wasserstein Spaces

This paper investigates a time discrete variational model for splines in Wasserstein spaces to interpolate probability measures. Cubic splines in Euclidean space are known to minimize the integrated squared acceleration subject to a set of interpolation constraints. As generalization on the space of probability measures the integral of the squared acceleration is considered as a spline energy and regularized by addition of the usual action functional. Both energies are then discretized in time using local Wasserstein-2 distances and the generalized Wasserstein barycenter. The existence of time discrete regularized splines for given interpolation conditions is established. On the subspace of Gaussian distributions, the spline interpolation problem is solved explicitly and consistency in the discrete to continuous limit is shown. The computation of time discrete splines is implemented numerically, based on entropy regularization and the Sinkhorn algorithm. A variant of Nesterov's accelerated gradient descent algorithm is applied for the minimization of the fully discrete functional. A variety of numerical examples demonstrate the robustness of the approach and show striking characteristics of the method. As a particular application the spline interpolation for synthesized textures is presented.

3.1 Background

In this section, we briefly review the classical theory of optimal transport (OT), and the Riemannian structure of the Wasserstein space induced by this OT metric.

3.1.1 Review of Optimal Transport

Let Ω be a Polish space (separable, completely metrizable) that additionally satisfies the Heine-Borel property, i.e. its compact sets are exactly the closed and bounded ones. Moreover, we introduce the set of probability measures $\mathcal{P}(\Omega)$ on Ω . The subset of probability measures μ with finite second moment, i.e. $\int_{\Omega} d^2(x_0, x) d\mu < \infty$, for some (and any) $x_0 \in \Omega$ and a fixed metric $d(\cdot, \cdot)$ that completely metrizes Ω will be denoted as $\mathcal{P}_2(\Omega)$. For two probability measures $\mu, \nu \in \mathcal{P}(\Omega)$, we shall denote with $U(\mu, \nu)$ the set of couplings between them, that is, the set of (probability) measures $\Pi \in \mathcal{P}(\Omega^2)$ with $\Pi(A \times \Omega) = \mu(A)$ and $\Pi(\Omega \times A) = \nu(A)$ for all Borel sets A in Ω . For $\mu, \nu \in \mathcal{P}(\Omega)$, the set $U_o(\mu, \nu)$ is the set of all couplings Π between μ and ν that minimize $\int_{\Omega^2} d^2(x, y) d\Pi(x, y)$, i.e. the set of optimal couplings for the cost $d^2(\cdot, \cdot)$.

Definition 3.1.1 (Wasserstein distance). The squared (L^2 -)Wasserstein distance between two probability measures $\mu, \nu \in \mathcal{P}(\Omega)$ will be denoted by \mathcal{W}^2 , and is defined as

$$\mathcal{W}^2(\mu, \nu) := \inf_{\Pi \in U(\mu, \nu)} \int_{\Omega^2} d^2(x, y) d\Pi(x, y).$$

Note that an optimal coupling is guaranteed to exist, and hence the infimum is actually a minimum. Furthermore, restricting to the space $\mathcal{P}_2(\Omega) \times \mathcal{P}_2(\Omega)$ actually leads to a complete metric space, cf. [Vil09]. With this in mind, we define a $\mathcal{P}_2(\Omega)$ -valued curve $(\mu_t)_{t \in [0,1]}$ as absolutely continuous, if there exists $m \in L^1([0,1])$, so that $\mathcal{W}(\mu_t, \mu_s) \leq \int_s^t m(r) \, dr$ for all $0 \leq s \leq t \leq 1$.

Moreover, let $C_b^0(\Omega)$ be the set of continuous, bounded functions on Ω . We then say that the sequence of measures $(\mu_k)_k$ converges narrowly to some $\mu \in \mathcal{P}(\Omega)$, if

$$\int_{\Omega} f \, d\mu_k \rightarrow \int_{\Omega} f \, d\mu,$$

for all $f \in C_b^0(\Omega)$. This will be denoted by $\mu_k \rightharpoonup \mu$.

The concept of tightness of probability measures will play a key role in the sequel: A set $\mathcal{K} \subseteq \mathcal{P}(\Omega)$ is said to be tight, if for any $\varepsilon > 0$ there is a compact set $\Omega_\varepsilon \subseteq \Omega$, such that $\mu(\Omega \setminus \Omega_\varepsilon) \leq \varepsilon$ for all $\mu \in \mathcal{K}$. Prokhorov's theorem states that tightness of a set of measures is equivalent to relative compactness in the topology induced by the narrow convergence of measures, cf. [Pro56].

3.1.2 Wasserstein spaces as a Riemannian Manifold

In this section we consider the spline interpolation problem from a geometric perspective. To this end, we will rely on the formal definition of a Riemannian metric on $\mathcal{P}_2(\mathbb{R}^d)$ given in [Lot06] and chapter 8 of [AGS08]. We first introduce a characterization of absolutely continuous measure-valued curves $(\mu_t)_{t \in [0,1]}$. Indeed, absolute continuity of a curve $(\mu_t)_t$ is equivalent to the existence of a velocity field $v_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ for $t \in [0,1]$, satisfying certain estimates, and solving the continuity equation (CE):

$$\partial_t \mu_t + \nabla \cdot (v_t \mu_t) = 0 \quad \text{in } (0,1) \times \mathbb{R}^d, \quad (3.1.1)$$

encoding the conservation of mass (see [AGS08], Theorem 8.3.1. for a thorough proof). The above equation is to be understood in the sense of distributions. Moreover, due to the Benamou-Brenier formula (cf. [BB00], Proposition 1.1) one recovers the following definition of the Wasserstein distance in terms of the velocity field $(v_t)_t$:

$$\mathcal{W}^2(\mu_0, \mu_1) := \inf_{(\mu, v) \in CE(\mu_0, \mu_1)} \int_0^1 \int_{\mathbb{R}^d} |v_t|^2 \, d\mu_t \, dt,$$

where $CE(\bar{\mu}_0, \bar{\mu}_1)$ is the set of pairs (μ, v) , such that $\mu = (\mu_t)_t$ is an absolutely continuous curve in $\mathcal{P}_2(\mathbb{R}^d)$, and $v = (v_t)_t$ is a time-dependent vector field, such that it satisfies (3.1.1) in the distributional sense, with $\mu_0 = \bar{\mu}_0$ and $\mu_1 = \bar{\mu}_1$. For a fixed curve $(\mu_t)_t$, the optimal velocity field $(v_t)_t$ of the above problem can be characterized as belonging to the set

$$T_{\mu_t} := \overline{\{\nabla \varphi : \varphi \in C_c^\infty(\mathbb{R}^d)\}}^{L^2(\mu_t, \mathbb{R}^d)}$$

for almost every $t \in [0,1]$ (cf. [AGS08], Proposition 8.4.5), where $C_c^\infty(\mathbb{R}^d)$ is the set of all real-valued, smooth, compactly supported functions on \mathbb{R}^d , and the bar notation denotes the closure of a set with respect to the $L^2(\mu_t, \mathbb{R}^d)$ norm. This fact justifies the suggestive definition of the set T_μ as the tangent space of the Wasserstein space $\mathcal{P}_2(\mathbb{R}^d)$ at the point μ . The Riemannian metric on $\mathcal{P}_2(\mathbb{R}^d)$ at μ is then simply given by the L^2 product

$$\langle v, w \rangle_{T_{\mu_t}} := \int_{\mathbb{R}^d} \langle v, w \rangle \, d\mu_t,$$

where $\langle \cdot, \cdot \rangle$ on the right hand side represents the usual inner product on \mathbb{R}^d . Then, the path energy \mathcal{E} of the measure-valued curve $(\mu_t)_t$ can be expressed by

$$\mathcal{E}((\mu_t)_t) = \inf_{v: (\mu, v) \in CE(\mu_0, \mu_1)} \int_0^1 \int_{\mathbb{R}^d} |v_t|^2 d\mu_t dt. \quad (3.1.2)$$

In their landmark paper [BB00], Benamou and Brenier showed that the functional being minimised in the last line is convex in the variables μ and $w = v\mu$. In [dB63], classical splines are defined as minimizers of the squared acceleration, integrated over time. The Riemannian counterpart to the acceleration of a particle is the covariant derivative of its velocity field $(v_t)_t$. To define this let us call a curve $(\mu_t)_t$ in $\mathcal{P}_2(\mathbb{R}^d)$ *regular*, if it is absolutely continuous and the optimal velocity vector field $(v_t)_t$ satisfying the continuity equation is Lipschitz in space and satisfies

$$\int_0^1 \text{Lip}(v_t) dt < \infty,$$

where $\text{Lip}(v)$ denotes the Lipschitz constant of v . Then, by [AGS08], Proposition 8.1.8, there exists a unique family of flow maps $\mathcal{T}_s^t(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that satisfy

$$\frac{d}{dt} \mathcal{T}_s^t(x) = v_t(\mathcal{T}_s^t(x)), \quad \mathcal{T}_s^s(x) = x. \quad (3.1.3)$$

We have that $\mu_t = (\mathcal{T}_s^t)_\# \mu_s$ for all $s \leq t$. The total derivative of an absolutely continuous vector field $(w_t)_t$ along a regular curve $(\mu_t)_t$ on $\mathcal{P}_2(\mathbb{R}^d)$ is then defined for almost all $t \in (0, 1)$ as

$$\frac{D}{dt} w_t := \lim_{h \rightarrow 0} \frac{w_{t+h} \circ \mathcal{T}_t^{t+h} - w_t}{h},$$

in the sense of $L^2(\mu_t)$. For a smooth vector field $(w_t)_t$ along a regular measure curve $(\mu_t)_t$, we can use (3.1.3) to obtain explicitly

$$\frac{D}{dt} w_t = \partial_t w_t + \nabla w_t \cdot v_t.$$

Finally, the covariant derivative can be given by projecting onto the tangent space

$$\nabla_{v_t} w_t := P_{\mu_t}(\partial_t w_t + \nabla w_t \cdot v_t), \quad (3.1.4)$$

where P_μ is the orthogonal projection in $L^2(\mu)$ onto the tangent space T_μ . For a thorough derivation of the covariant derivative on $\mathcal{P}_2(\Omega)$, we refer to [AG13], chapter 6.

3.2 Splines in Wasserstein Spaces

3.2.1 Definition of splines

Based on the discussion in the previous section, for $v_t = \nabla \varphi_t \in T_{\mu_t}$ one may use $\nabla_{v_t} v_t$ as the acceleration of a regular measure-valued curve μ_t . This leads to

$$\nabla_{v_t} v_t = P_{\mu_t}(\partial_t v_t + \nabla v_t \cdot v_t) = P_{\mu_t}(\dot{v}_t + \frac{1}{2} \nabla |v_t|^2) = \dot{v}_t + \frac{1}{2} \nabla |v_t|^2, \quad (3.2.1)$$

where the last equality holds due to $\dot{v}_t = \nabla \dot{\phi}_t \in T_{\mu_t}$ and the second term already being in gradient-field form. This naturally leads to the following notion of a continuous-time spline energy functional. For a general curve $(\mu_t)_t : [0, 1] \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ we set

$$\mathcal{F}((\mu_t)_t) = \inf_v \int_0^1 \int_{\mathbb{R}^d} \left| \dot{v}_t + \frac{1}{2} \nabla |v_t|^2 \right|^2 d\mu_t dt, \quad (3.2.2)$$

where the infimum is taken over sufficiently regular time-dependent vector fields $v = (v_t)_t$, such that $(\mu, v) \in CE(\mu_0, \mu_1)$, and $v_t \in T_{\mu_t}$ for all $t \in (0, 1)$. The spline interpolation problem in the Wasserstein space is then to find a curve $(\mu_t)_t : [0, 1] \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ that minimizes the functional (3.2.2), subject to a set of $I > 2$ point-wise interpolation constraints

$$\mu_{\bar{t}_i} = \bar{\mu}_i, \quad i = 1, \dots, I, \quad (3.2.3)$$

for prescribed times $\bar{t}_i \in [0, 1]$, $i = 1, \dots, I$, with $\bar{t}_1 < \dots < \bar{t}_I$ and $\bar{\mu}_i \in \mathcal{P}_2(\mathbb{R}^d)$. As already discussed in [dB63] for the Euclidean case, and [HRW18] for the Riemannian case, we may impose one of the following boundary conditions (b.c.):

$$\text{natural b.c.:} \quad \text{no additional condition,} \quad (3.2.4)$$

$$\text{Hermite b.c.:} \quad v_0 = \bar{v}_0, v_1 = \bar{v}_1 \text{ for given } \bar{v}_0 \in T_{\mu_0} \text{ and } \bar{v}_1 \in T_{\mu_1}, \quad (3.2.5)$$

$$\text{periodic b.c.:} \quad \mu_0 = \mu_1, v_0 = v_1. \quad (3.2.6)$$

In the case of Hermite (also known as clamped) boundary conditions, we assume that $\bar{t}_1 = 0$ and $\bar{t}_I = 1$, so that μ_0 and μ_1 are prescribed as well.

From a theoretical point of view, it will be advantageous to regularize the above spline energy by adding the path energy \mathcal{E} multiplied by a regularization parameter $\delta > 0$. Hence, we introduce the regularized spline energy functional

$$\mathcal{F}^\delta := \mathcal{F} + \delta \mathcal{E}. \quad (3.2.7)$$

This will ensure tightness of all probability measures with finite energy, and consequently existence in the time-discrete case.

Definition 3.2.1. For given times $\bar{t}_i \in [0, 1]$ and prescribed probability distributions $\bar{\mu}_i \in \mathcal{P}_2(\mathbb{R}^d)$, $i = 1, \dots, I$, we define a (regularized) spline interpolation $(\mu_t)_t$ as a minimizer of the spline energy functional (3.2.2) (resp. (3.2.7)) subject to (3.2.3) and at most one of the boundary conditions (3.2.4)-(3.2.6).

Example 3.2.2 (Euclidean space). The Wasserstein distance between two delta distributions located at x and y is equal to the Euclidean distance $|x - y|$, and the associated Wasserstein geodesic is given by the curve of delta distributions at the locations of the Euclidean geodesic interpolating the end points. We now briefly check whether our definition is also consistent with cubic splines in \mathbb{R}^d when considering delta distributions.

Let $x : [0, 1] \rightarrow \mathbb{R}^d$ be a twice-differentiable curve, and define the measure-valued curve $\mu_t = \delta_{x_t}$. Then, one checks that with the choice $v_t \equiv \dot{x}_t$, (CE) is satisfied in distributional sense:

$$\int_0^1 \int_{\mathbb{R}^d} (\mu_t \partial_t \phi(t, x) + \mu_t v_t \cdot \nabla \phi(t, x)) dt dx = \int_0^1 (\partial_t \phi(t, x(t)) + \dot{x}_t \cdot \nabla \phi(t, x(t))) dt = 0,$$

for all $\phi \in C_c^\infty((0, 1) \times \mathbb{R}^d)$. Moreover, as v_t is constant in space, we have $Dv_t \equiv 0$. Due to (3.2.1), we obtain $\nabla_{v_t} v_t = \ddot{x}_t$, so using (3.2.2) one gets

$$\mathcal{F}((\mu_t)_t) = \int_0^1 \int_{\mathbb{R}^d} |\ddot{x}_t|^2 d\mu_t dt = \int_0^1 |\ddot{x}_t|^2 dt,$$

for which the minimizer is given by the cubic spline subject to the interpolation constraints [dB63].

The Wasserstein space $\mathcal{P}_2(\mathbb{R}^d)$ is isometrically isomorphic to $\mathbb{R}^d \times \mathcal{P}_2^0(\mathbb{R}^d)$, where the factor \mathbb{R}^d represents the center of mass and $\mathcal{P}_2^0(\mathbb{R}^d)$ is the space of probability distributions centered around 0. In this spirit the dynamic of spline paths can be split into the time evolution of the center the mass and the time evolution of the distribution around it, as we shall now demonstrate.

Let $(\mu, v) \in CE$ be a solution to the continuity equation, with $v = (v_t)_t$ being optimal. Hence, for all $t \in [0, 1]$, v_t is a gradient field, and in particular $Dv_t^T = Dv_t$. Let $m_t := \int x d\mu_t(x)$ be the center of mass and let $\tilde{\mu}_t(\cdot) := \mu_t(\cdot + m_t)$ be the re-centered distribution. Furthermore, we define the re-centered velocity field $\tilde{v}_t(x) := v_t(x + m_t) - \dot{m}_t$. Then one easily checks that $(\tilde{\mu}, \tilde{v}) \in CE$. Now, we first show a decoupling of the (first-order) action functional, i.e.

$$\begin{aligned} \int_0^1 \int_{\mathbb{R}^d} |\tilde{v}_t|^2 d\tilde{\mu}_t dt &= \int_0^1 \int_{\mathbb{R}^d} |v_t(\cdot + m_t) - \dot{m}_t|^2 d(\text{Id} - m_t)_\# \mu_t dt = \int_0^1 \int_{\mathbb{R}^d} |v_t - \dot{m}_t|^2 d\mu_t dt \\ &= \int_0^1 |v_t|^2 d\mu_t dt + \int_0^1 |\dot{m}_t|^2 dt - 2 \int_0^1 \int_{\mathbb{R}^d} \langle v_t, \dot{m}_t \rangle d\mu_t dt = \int_0^1 |v_t|^2 d\mu_t dt - \int_0^1 |\dot{m}_t|^2 dt, \end{aligned}$$

where we used that by the continuity equation

$$\dot{m}_t = \frac{d}{dt} \int_{\mathbb{R}^d} x d\mu_t(x) = \int_{\mathbb{R}^d} \nabla x \cdot v_t d\mu_t(x) = \int_{\mathbb{R}^d} v_t d\mu_t(x). \quad (3.2.8)$$

Next, we consider the decoupling of the (second-order) spline energy $\int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t + \frac{1}{2} \nabla |v_t|^2|^2 d\mu_t dt$. Taking into account

$$\begin{aligned} \dot{\tilde{v}}_t(x) &= \partial_t(v_t(x + m_t) - \dot{m}_t) = \dot{v}_t(x + m_t) + (Dv_t)(x + m_t) \cdot \dot{m}_t - \ddot{m}_t, \\ \nabla |\tilde{v}_t|^2 &= \nabla |\dot{m}_t|^2 + \nabla |v_t(x + m_t)|^2 - 2 \nabla \langle v_t(x + m_t), \dot{m}_t \rangle \\ &= 2(Dv_t)(x + m_t) \cdot v_t(x + m_t) - 2(Dv_t)(x + m_t) \dot{m}_t \end{aligned}$$

we obtain

$$\begin{aligned} &\int_0^1 \int_{\mathbb{R}^d} \left| \dot{\tilde{v}}_t + \frac{1}{2} \nabla |\tilde{v}_t|^2 \right|^2 d\tilde{\mu}_t dt \\ &= \int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t(x + m_t) + (Dv_t)(x + m_t) \cdot \dot{m}_t - \ddot{m}_t - (Dv_t)(x + m_t) \cdot \dot{m}_t \\ &\quad + (Dv_t)(x + m_t) \cdot v_t(x + m_t)|^2 d\tilde{\mu}_t(x) dt \\ &= \int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t(x + m_t) - \ddot{m}_t + (Dv_t)(x + m_t) \cdot v_t(x + m_t)|^2 d\tilde{\mu}_t(x) dt \\ &= \int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t - \ddot{m}_t + Dv_t(v_t)|^2 d\mu_t dt \end{aligned}$$

$$= \int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t + Dv_t(v_t)|^2 d\mu_t dt + \int_0^1 |\ddot{m}_t|^2 dt - 2 \int_0^1 \int_{\mathbb{R}^d} \langle \ddot{m}_t, \dot{v}_t + Dv_t(v_t) \rangle d\mu_t dt.$$

Now, differentiating (3.2.8) in time we achieve

$$\ddot{m}_t = \frac{d}{dt} \int_{\mathbb{R}^d} v_t d\mu_t = \int_{\mathbb{R}^d} \dot{v}_t + Dv_t(v_t) d\mu_t.$$

Finally, plugging this back into the previous computation we get

$$\int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t + \frac{1}{2} \nabla |\tilde{v}_t|^2|^2 d\tilde{\mu}_t dt = \int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t + \frac{1}{2} \nabla |v_t|^2|^2 d\mu_t dt - \int_0^1 |\ddot{m}_t|^2 dt \quad (3.2.9)$$

This decoupling is advantageous for the numerical implementation. In fact, it leads to a reduced computing time (cf. Figure 3.6).

3.2.2 Variational discretization of splines

The temporal discretization of (regularized) Wasserstein spline energies will be based on a variational problem. To motivate the proposed discrete spline energy functional, let us consider the situation in Euclidean spaces, in which the velocity field v of a smooth curve $x : [0, 1] \rightarrow \mathbb{R}^d$ coincides with \dot{x} . By sampling this curve uniformly, i.e. taking $x_k := x(t_k^K)$ for $t_k^K := k/K$, $k = 0, \dots, K$, we are able to approximate the velocity at a time t_k^K by finite differences, that is, $\dot{x}(t_k^K) \approx K(x_{k+1} - x_k)$. Therefore, we obtain

$$|\dot{x}(t_k^K)|^2 \approx K^2 |x_{k+1} - x_k|^2.$$

Similarly, in Euclidean spaces the covariant derivative of the velocity field coincides with the acceleration \ddot{x} . We approximate this by central second order difference quotients, i.e. $\ddot{x}(t_k^K) \approx K^2(x_{k+1} - 2x_k + x_{k-1})$. Thus, defining $\text{Bar}(x_{k+1}, x_{k-1}) := \frac{x_{k+1} + x_{k-1}}{2}$ one obtains

$$|\ddot{x}(t_k^K)|^2 \approx 4K^4 \left| x_k - \frac{x_{k+1} + x_{k-1}}{2} \right|^2 = 4K^4 |x_k - \text{Bar}(x_{k+1}, x_{k-1})|^2.$$

A simple rectangular quadrature rule $\int_0^1 f(t) dt \approx K^{-1} \sum_{k=1}^{K-1} f(t_k^K)$ for $t_k^K := \frac{k}{K}$ leads to the following approximations of the Euclidean velocity and acceleration functional, respectively:

$$\mathcal{E}(x) = \int_0^1 |\dot{x}_t|^2 dt \approx K \sum_{k=1}^K |x_{k+1} - x_k|^2, \quad (3.2.10)$$

$$\mathcal{F}(x) = \int_0^1 |\ddot{x}_t|^2 dt \approx 4K^3 \sum_{k=1}^{K-1} |x_k - \text{Bar}(x_{k+1}, x_{k-1})|^2. \quad (3.2.11)$$

Recall that the Euclidean barycenter is the solution to the following minimization problem:

$$\text{Bar}(x, y) = \operatorname{argmin}_{z \in \mathbb{R}^d} (|x - z|^2 + |y - z|^2),$$

for some $x, y \in \mathbb{R}^d$. Hence, it is intuitive to replace the Euclidean L^2 -norm with the Wasserstein distance, giving rise to the following discrete path energy

$$\mathbf{E}^K(\boldsymbol{\mu}^K) := K \sum_{k=0}^{K-1} \mathcal{W}^2(\mu_k^K, \mu_{k+1}^K), \quad (3.2.12)$$

for a $(K+1)$ -tuple of probability measures $\boldsymbol{\mu}^K := (\mu_0^K, \dots, \mu_K^K) \in \mathcal{P}_2(\Omega)^{K+1}$. Moreover, we will also give suitable definitions of a Wasserstein barycenter:

Definition 3.2.3. Let $\mu, \nu \in \mathcal{P}(\Omega)$, and $t \in [0, 1]$. The set of t -barycenters $\text{Bar}^t(\mu, \nu)$ between μ and ν is the set of solutions of the following minimization problem

$$\operatorname{argmin}_{\rho \in \mathcal{P}(\Omega)} (1-t)\mathcal{W}^2(\rho, \mu) + t\mathcal{W}^2(\rho, \nu). \quad (3.2.13)$$

For the sake of readability, we shall usually omit the t -index from both the notation and nomenclature when $t = \frac{1}{2}$.

Remark 3.2.4. If $\mu, \nu \in \mathcal{P}_2(\Omega)$, then we can guarantee the existence of a solution of (3.2.13) (cf. [Lot06]). Indeed, let $\Pi \in U_o(\mu, \nu)$, and let π^i be the projection operators onto the i -th coordinate. Then,

$$((1-t)\pi_1 + t\pi_2)_\# \Pi \in \text{Bar}^t(\mu, \nu). \quad (3.2.14)$$

Let $\Omega \subseteq \mathbb{R}^d$ and define $\mathcal{P}_2^{ac}(\Omega) \subset \mathcal{P}_2(\Omega)$ as the set of all absolutely continuous probability measures in $\mathcal{P}_2(\Omega)$ with respect to the Lebesgue measure on \mathbb{R}^d . If, in addition, at least one of μ or ν belong to the set $\mathcal{P}_2^{ac}(\Omega)$, then Brenier's theorem [Bre91] and McCann's interpolation [McC01] even guarantee uniqueness of the t -barycenter, given explicitly by

$$\text{Bar}^t(\mu, \nu) = \{((1-t)\mathbb{1} + tT_\mu^\nu)_\# \mu\}, \quad (3.2.15)$$

where T_μ^ν is the optimal transport map from μ to ν .

In Wasserstein spaces, there is another related notion of barycenter, which will be called generalized Wasserstein barycenter:

Definition 3.2.5. Let $\mu_1, \mu_2, \mu_3 \in \mathcal{P}(\Omega)$. Let now Π be a three-measure coupling between them, i.e. $\Pi \in \mathcal{P}(\Omega^3)$, and $\Pi(A \times \Omega \times \Omega) = \mu_1(A)$, $\Pi(\Omega \times A \times \Omega) = \mu_2(A)$, and $\Pi(\Omega \times \Omega \times A) = \mu_3(A)$ for all Borel sets $A \subseteq \Omega$. If furthermore, we have that $(\pi^1, \pi^2)_\# \Pi \in U_o(\mu_1, \mu_2)$, and $(\pi^2, \pi^3)_\# \Pi \in U_o(\mu_2, \mu_3)$, we say $\Pi \in U_o(\mu_1, \mu_2, \mu_3)$. A measure μ is in the set of generalized (Wasserstein) t -barycenters $\text{Bar}_{\mu_2}^t(\mu_1, \mu_3)$ between μ_1 and μ_3 with base point μ_2 , if it is of the form $\mu = ((1-t)\pi^1 + t\pi^3)_\# \Pi$ for a $\Pi \in U_o(\mu_1, \mu_2, \mu_3)$. When $t = \frac{1}{2}$, we shall omit t from the notation.

Remark 3.2.6. Similarly as above, if $\Omega \subseteq \mathbb{R}^d$ and $\mu_2 \in \mathcal{P}_2^{ac}(\Omega)$, then Brenier's theorem guarantees uniqueness of the generalized t -barycenter, given explicitly by

$$\text{Bar}_{\mu_2}^t(\mu_1, \mu_3) = \{((1-t)T_2^1 + tT_2^3)_\# \mu_2\}, \quad (3.2.16)$$

where T_2^i is the optimal transport map from μ_2 to μ_i , $i = 1, 3$. Note that $(1-t)T_2^1 + tT_2^3$ is again an optimal map (since it inherits the structure of being the gradient of a convex function from T_2^i).

In analogy to equation (3.2.11) we will define two notions of time discrete spline energies related to the different kinds of barycenters introduced above:

Definition 3.2.7 (Discrete spline energy). Let $\mu^K := (\mu_0^K, \dots, \mu_K^K) \in \mathcal{P}(\Omega)^{K+1}$ be a $(K+1)$ -tuple of probability measures. The discrete spline energy \mathbf{F}^K of μ^K is then defined as

$$\mathbf{F}^K(\mu^K) := \inf_{\tilde{\mu}^K} 4K^3 \sum_{k=1}^{K-1} \mathcal{W}^2(\mu_k^K, \tilde{\mu}_k^K), \quad (3.2.17)$$

where the infimum is taken over all $\tilde{\mu}^K = (\tilde{\mu}_k^K)_{k=1, \dots, K-1}$ with $\tilde{\mu}_k^K \in \text{Bar}(\mu_{k-1}^K, \mu_{k+1}^K)$. Similarly, one defines the generalized discrete spline energy \mathbf{F}_G^K of μ^K as

$$\mathbf{F}_G^K(\mu^K) := \inf_{\tilde{\mu}^K} 4K^3 \sum_{k=1}^{K-1} \mathcal{W}^2(\mu_k^K, \tilde{\mu}_k^K), \quad (3.2.18)$$

where the infimum is taken over all $\tilde{\mu}^K = (\tilde{\mu}_k^K)_{k=1, \dots, K-1}$ with $\tilde{\mu}_k^K \in \text{Bar}_{\mu_k}(\mu_{k-1}^K, \mu_{k+1}^K)$. The regularized discrete spline energies are given by

$$\mathbf{F}^{\delta, K} := \mathbf{F}^K + \delta \mathbf{E}^K, \quad \mathbf{F}_G^{\delta, K} := \mathbf{F}_G^K + \delta \mathbf{E}^K \quad (3.2.19)$$

for $\delta > 0$ (for $\delta = 0$ we retrieve the non-regularized spline energy). Computing a (regularized) time-discrete spline interpolation now consists in finding a tuple $\mu^K = (\mu_0^K, \dots, \mu_K^K)$ that minimizes the functional (3.2.19) in some sense to be defined, subject to a set of $I > 2$ point-wise interpolation constraints

$$\mu_{K\bar{t}_i}^K = \bar{\mu}_i, \quad i = 1, \dots, I, \quad (3.2.20)$$

for fixed prescribed times $\bar{t}_i \in [0, 1]$, which fulfil $K\bar{t}_i \in \mathbb{N}_0$, with $\bar{t}_1 < \dots < \bar{t}_I$ and $\bar{\mu}_i \in \mathcal{P}_2(\mathbb{R}^d)$ for $i = 1, \dots, I$.

The discrete counterparts of boundary conditions, one of which may be additionally imposed, can be written as follows:

$$\text{natural b.c.:} \quad \text{no additional condition}, \quad (3.2.21)$$

$$\text{Hermite b.c.:} \quad \mu_0^K = \bar{\mu}_0, \mu_1^K = \bar{\mu}_1, \mu_{K-1}^K = \bar{\mu}_{K-1}, \mu_K^K = \bar{\mu}_K, \quad (3.2.22)$$

$$\text{periodic b.c.:} \quad \mu_1^K = \mu_K^K, \mu_0^K = \mu_{K-1}^K. \quad (3.2.23)$$

Now we are in position to define regularized time-discrete spline interpolations:

Definition 3.2.8 (Regularized discrete spline interpolations). For $2 \leq I \leq K$, given data points $\bar{t}_i \in [0, 1]$ fulfilling $K\bar{t}_i \in \mathbb{N}_0$, $\delta > 0$ and fixed data $\bar{\mu}_i \in \mathcal{P}(\Omega)$ for $i = 1, \dots, I$, we define the tuple $\mu^K \in \mathcal{P}_2(\Omega)^{K+1}$ to be a regularized (generalized) discrete spline interpolation if it is a minimizer of the discrete spline energy functional $\mathbf{F}_{(G)}^{\delta, K}$ with $\delta > 0$ (cf. (3.2.19)) that satisfy the interpolation constraints (3.2.20) and one of the boundary conditions (3.2.21)-(3.2.23).

We will now show existence of a minimizer of the regularized spline energy functional introduced above, for all $\delta > 0$. First, let us show a technical lemma:

Lemma 3.2.9. *Let Ω be as in Subsection 3.1.1 and let $(\mu_n)_n \subseteq \mathcal{P}(\Omega)$ be tight, and $(\nu_n)_n \subseteq \mathcal{P}(\Omega)$. If $\sup_n \mathcal{W}^2(\mu_n, \nu_n) \leq C < \infty$, then $(\nu_n)_n$ is also tight.*

Proof. We will argue by contradiction: Assume that $(\nu_n)_n$ is not tight. Then, there is an $\varepsilon > 0$, so that for all $R > 0$ there is a $k = k(R) \in \mathbb{N}$ that fulfils $\nu_k(\Omega \setminus \overline{B_R(\omega)}) > \varepsilon$ for some fixed $\omega \in \Omega$.

Let $r > 0$ be chosen so that $R > r$, and $\mu_n(\Omega \setminus \overline{B_r(\omega)}) \leq \varepsilon/2$ for all $n \in \mathbb{N}$. This is possible due to the tightness of $(\mu_n)_n$. For any coupling $\Pi \in \mathcal{P}(\Omega^2)$ of $\mu_{k(R)}$ and $\nu_{k(R)}$ we have that

$$\Pi(\{(x, y) : d^2(x, y) \geq (R - r)^2\}) > \varepsilon/2.$$

Hence, we obtain

$$\mathcal{W}^2(\mu_k, \nu_k) > \frac{\varepsilon}{2}(R - r)^2.$$

Since ε and r are fixed, and k only depends on R , we can choose R big enough so that $\mathcal{W}^2(\mu_k, \nu_k) > C$, which leads to the desired contradiction. \square

Theorem 3.2.10. *For all $\delta > 0$, $K \in \mathbb{N}$, $2 \leq I \leq K$, given times $\bar{t}_i \in [0, 1]$ and prescribed probability measures $\bar{\mu}_i \in \mathcal{P}_2(\mathbb{R}^d)$ for all $i = 1, \dots, I$, there exists a discrete regularized (generalized) spline interpolation in the sense of Definition 3.2.8.*

Proof. Any choice of $\mu_k \in \mathcal{P}(\mathbb{R}^d)$ for $k = 0, \dots, K$ gives a finite regularized spline energy $\bar{F} := \mathbf{F}^{\delta, K}((\mu_k)_k)$. Let $(\mu^{(n)})_n$ be a minimizing sequence for $\mathbf{F}^{\delta, K}$ under the given constraints. In particular, $\sup_n \mathbf{F}^{\delta, K}(\mu^{(n)}) \leq \bar{F}$. Thus,

$$\bar{F} \geq \sup_n \mathbf{F}^{\delta, K}(\mu^{(n)}) \geq \sup_n \delta \mathcal{W}^2(\mu_k^{(n)}, \mu_{k+1}^{(n)}),$$

for any k . For $i = 1, \dots, I$, $\mu_{K\bar{t}_i}^{(n)} = \bar{\mu}_{K\bar{t}_i}$ for all n . Since any constant measure-valued sequence is tight, by the previous lemma the sequence $(\mu_{K\bar{t}_i+1}^{(n)})_n$ is also tight. We can use the previous lemma multiple times and "propagate" tightness by induction. Next, by Prokhorov's theorem we can choose a subsequence, so that for all $k \in \{0, \dots, K\}$ the sequence $(\mu_k^{(n)})_n$ is narrowly convergent to some $\mu_k \in \mathcal{P}(\mathbb{R}^d)$. In fact, we have by the triangle inequality

$$\mathcal{W}(\mu_k^{(n)}, \delta_{x_0}) \leq \mathcal{W}(\mu_k^{(n)}, \mu_{k-1}^{(n)}) + \mathcal{W}(\mu_{k-1}^{(n)}, \delta_{x_0}),$$

for a point x_0 in \mathbb{R}^d . If $k - 1 = K\bar{t}_i$ for some $i = 1, \dots, I$, then $\mu_{k-1} = \bar{\mu}_i \in \mathcal{P}_2(\mathbb{R}^d)$ and hence, the second term on the right hand side is uniformly bounded in n . Since the first term on the right hand side is one term in the discrete path energy contained in $\mathbf{F}^{\delta, K}$, it is uniformly bounded in n as well. We now use $\mu_k^{(n)} \rightharpoonup \mu_k$ and the lower semi-continuity of \mathcal{W} under narrow convergence to show that

$$\mathcal{W}^2(\mu_k, \delta_{x_0}) \leq \liminf_{n \rightarrow \infty} \mathcal{W}^2(\mu_k^{(n)}, \delta_{x_0}) < \infty.$$

Proceeding by induction, we obtain that $\mu_k \in \mathcal{P}_2(\mathbb{R}^d)$ for all $k = 0, \dots, K$. Let us now rewrite (3.2.17):

$$\mathbf{F}^K(\mu^{(n)}) = 4K^3 \sum_{k=1}^{K-1} \inf_{\tilde{\mu}_k^{(n)} \in \text{Bar}(\mu_{k-1}^{(n)}, \mu_{k+1}^{(n)})} \mathcal{W}^2(\mu_k^{(n)}, \tilde{\mu}_k^{(n)})$$

$$= 4K^3 \sum_{k=1}^{K-1} \inf_{\tilde{\Pi}_{k-1,k+1}^{(n)} \in U_o(\mu_{k-1}^{(n)}, \mu_{k+1}^{(n)})} \mathcal{W}^2 \left(\mu_k^{(n)}, \left(\frac{1}{2}\pi^1 + \frac{1}{2}\pi^2 \right)_{\#} \tilde{\Pi}_{k-1,k+1}^{(n)} \right).$$

Next, we denote the value of the inner infimum above $I_k^{(n)}$ and assume that $\Pi_{k-1,k+1}^{(n)} \in U_o((\mu_{k-1}^{(n)}, \mu_{k+1}^{(n)}))$ is chosen such that $\mathcal{W}^2 \left(\mu_k^{(n)}, \left(\frac{1}{2}\pi^1 + \frac{1}{2}\pi^2 \right)_{\#} \Pi_{k-1,k+1}^{(n)} \right) \leq I_k^{(n)} + 1/n$. By the stability of optimal couplings [AG13, Prop. 7.1.3], $\Pi_{k-1,k+1}^{(n)}$ converges (up to a subsequence) to an optimal coupling $\Pi_{k-1,k+1}$ of μ_{k-1} and μ_{k+1} . This entails narrow convergence of the barycenter

$$\text{Bar}(\mu_{k-1}^{(n)}, \mu_{k+1}^{(n)}) \ni \left(\frac{1}{2}\pi^1 + \frac{1}{2}\pi^2 \right)_{\#} \Pi_{k-1,k+1}^{(n)} \rightarrow \left(\frac{1}{2}\pi^1 + \frac{1}{2}\pi^2 \right)_{\#} \Pi_{k-1,k+1} =: \tilde{\mu}_k \in \text{Bar}(\mu_{k-1}, \mu_{k+1})$$

due to the continuity of the projections π^i (for $i = 1, 2$), and Remark 3.2.4.

Finally, we use the lower semi-continuity of the Wasserstein distance under narrow convergence and the fact that the spline energy contains a minimization over the choice of barycenters, to obtain

$$\mathbf{F}^{\delta,K}(\boldsymbol{\mu}) \leq 4K^3 \sum_{k=1}^{K-1} \mathcal{W}^2(\mu_k, \tilde{\mu}_k) + \delta \mathbf{E}^K(\boldsymbol{\mu}) \leq \liminf_{n \rightarrow \infty} \mathbf{F}^{\delta,K}(\boldsymbol{\mu}^{(n)}),$$

where $\boldsymbol{\mu} := (\mu_0, \dots, \mu_K)$. As the right-hand sequence was assumed to be a minimizing sequence, $\boldsymbol{\mu}$ is indeed a spline interpolation according to Definition 3.2.8. The proof of existence of generalized spline interpolations is by analogy. The only remarkable difference is to prove that generalized barycenters narrowly converge to a generalized barycenter, up to a subsequence. To see this, let $\Pi_{k-1,k+1}^{(n)} \in \mathcal{P}(\mathbb{R}^{3d})$ be a three-measure optimal transport plan between $\mu_{k-1}^{(n)}, \mu_k^{(n)}$ and $\mu_{k+1}^{(n)}$, i.e. $\Pi_{k-1,k+1}^{(n)} \in U_o(\mu_{k-1}^{(n)}, \mu_k^{(n)}, \mu_{k+1}^{(n)})$ (cf. Definition 3.2.5). Once again, as the marginals of $\Pi_{k-1,k+1}^{(n)}$ are tight, the sequence of optimal couplings $(\Pi_{k-1,k+1}^{(n)})_n$ is also tight, and due to the lower semicontinuity of \mathcal{W} , it narrowly converges (up to a subsequence) to an optimal coupling $\Pi_{k-1,k+1}$ of μ_{k-1}, μ_k and μ_{k+1} . From this, narrow convergence (up to a subsequence) of the sequence of generalized barycenters $\left(\frac{1}{2}\pi^1 + \frac{1}{2}\pi^3 \right)_{\#} \Pi_{k-1,k+1}^{(n)}$ to the generalized barycenter $\left(\frac{1}{2}\pi^1 + \frac{1}{2}\pi^3 \right)_{\#} \Pi_{k-1,k+1}$ follows, again due to the continuity of the projections π^i (for $i = 1, 3$), and Definition 3.2.5. \square

Sampling spline trajectories. As mentioned in the introduction, our interpolation method significantly differs from P-splines (cf. [BGV19] and [CCG18]) and T-splines (cf. [CCLG⁺21]) in the sense that the two aforementioned methods are primarily trajectory-based, while ours focuses on the densities themselves. Hence, it is a priori unclear whether sample spline trajectories can also be constructed that are consistent with our method. It turns out the answer to the previous question is positive. Next, we shall describe the different trajectory sampling methods for P-splines T-splines and E-splines (our method), and state algorithms to compute them:

P-splines: Any solution $(X_t)_t$ of (1.0.2) has spline trajectories $t \mapsto X_t(\omega)$ for \mathbb{P} -almost all $\omega \in \Omega$, which formally follows from rewriting (1.0.2):

$$\min_{\substack{(X_t)_t \\ X_{t_i} \sim \mu_i}} \int_0^1 \int_{\Omega} \|\ddot{X}_t\|^2 d\mathbb{P} dt = \min_{\substack{Q \in \mathcal{P}(\mathbb{R}^{d \cdot I}) \\ (\pi_i)_{\#} Q = \mu_i}} \int_{\mathbb{R}^{d \cdot I}} \min_{\substack{t \mapsto y_t \\ y_{t_i} = x_i}} \int_0^1 \|\ddot{y}_t\|^2 dt dQ(x_1, \dots, x_I)$$

$$= \min_{\substack{Q \in \mathcal{P}(\mathbb{R}^{d \cdot I}) \\ (\pi_i)_\# Q = \bar{\mu}_i}} \int_{\mathbb{R}^{d \cdot I}} \int_0^1 \|\ddot{s}_t\|^2 dt dQ(x_1, \dots, x_I), \quad (3.2.24)$$

where $t \mapsto s_t$ is the classical Euclidean spline interpolating the points (t_i, x_i) and $\pi_i : \mathbb{R}^{d \cdot I} \rightarrow \mathbb{R}^d$ is the projection onto the i -th d -sized batch of coordinates, i.e. $\pi_i(y_1, \dots, y_{d \cdot I}) = (y_{d \cdot i + 1}, \dots, y_{d \cdot i + d})$. Hence, any solution $(X_t)_t$ of (1.0.2) has spline trajectories $t \mapsto X_t(\omega)$ for \mathbb{P} -almost all $\omega \in \Omega$ (see Fig. 1.1, top left), cf. Algorithm 1

Algorithm 1: Algorithm for sampling P-spline trajectories.

```

1 Function PSplineSample  $((\bar{\mu}_1, \dots, \bar{\mu}_I), (\bar{t}_1, \dots, \bar{t}_I))$  :
2   for  $i = 1$  to  $I$  do
3     Sample  $X_i \sim \bar{\mu}_i$ 
4   return CubicSplineInterpolation  $((\bar{t}_1, X_1), \dots, (\bar{t}_I, X_I))$ 

```

T-splines: Samples are first drawn from one of the prescribed distributions, usually $X_1 \sim \bar{\mu}_1$, where X_1 is a random variable with distribution $\bar{\mu}_1$. These samples are then pushed by the Monge maps T_i between consecutive prescribed distributions $\bar{\mu}_i, \bar{\mu}_{i+1}$ and the resulting tuple of points $x_i := (T_{i-1} \circ \dots \circ T_1 \circ X_1)(\omega)$ at the prescribed times \bar{t}_i , for $i = 1, \dots, I$, is interpolated using classical cubic spline interpolation (see Fig. 1.1 top middle and bottom left). cf. Algorithm 2

Algorithm 2: Algorithm for sampling T-spline trajectories.

```

1 Function TSplineSample  $((\bar{\mu}_1, \dots, \bar{\mu}_I), (\bar{t}_1, \dots, \bar{t}_I))$  :
2   Sample  $X_1 \sim \bar{\mu}_1$ 
3   for  $i = 1$  to  $I - 1$  do
4      $T_i = \text{OptimalTransportMap}(\bar{\mu}_i, \bar{\mu}_{i+1})$ 
5      $X_{i+1} = T_i(X_i)$ 
6   return CubicSplineInterpolation  $((\bar{t}_1, X_1), \dots, (\bar{t}_I, X_I))$ 

```

E-splines: We assume we have already found a discrete-time interpolation (μ_0, \dots, μ_K) . Similarly to the T-spline method, samples are first drawn from one of the interpolated measures. Unlike before, we are allowed to choose any of the measures, not just one of the prescribed ones. Without loss of generality, assume $X_0 \sim \mu_0$ is the random variable to be sampled. Next, we compute the Monge maps T_i between measures μ_i and μ_{i+1} and push the sample $x_0 := X_0(\omega)$ iteratively through them, i.e. $x_i := T_{i-1}(x_{i-1})$. Finally, the tuple (x_0, \dots, x_K) is interpolated once again through Euclidean cubic spline interpolation, cf. Algorithm 3

3.3 Gaussian E-splines

In this section we will explicitly derive the continuous spline energy for measure-valued curves restricted to the space of Gaussian distributions, i.e. minimizers of the spline energy among Gaussian curves, and

Algorithm 3: Algorithm for sampling E-spline trajectories.

```

1 Function ESplineSample  $((\bar{\mu}_1, \dots, \bar{\mu}_I), (\bar{t}_1, \dots, \bar{t}_I))$  :
2    $(\mu_0, \dots, \mu_K) = \text{ESpline}((\bar{\mu}_1, \dots, \bar{\mu}_I), (\bar{t}_1, \dots, \bar{t}_I))$ 
3   Sample  $X_0 \sim \mu_0$ 
4   for  $i = 0$  to  $K - 1$  do
5      $T_i = \text{OptimalTransportMap}(\mu_i, \mu_{i+1})$ 
6      $X_{i+1} = T_i(X_i)$ 
7   return CubicSplineInterpolation  $((t_0, X_0), \dots, (t_K, X_K))$ 

```

show its consistency with the discrete spline energy notions we defined in the previous section. Let us first introduce some notation:

Definition 3.3.1. Let $\mathbb{R}_{+, \text{sym}}^{d \times d}$ be the space of symmetric, positive definite $d \times d$ matrices, and $\mathbb{R}_{+, \text{dia}}^{d \times d} \subset \mathbb{R}_{+, \text{sym}}^{d \times d}$ the space of diagonal, positive definite $d \times d$ matrices. Then, one can identify the space of Gaussian probability measures with the set $\mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d}$ through the bijective map

$$\Phi : \mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d} \longrightarrow \Phi(\mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d}) \subset \mathcal{P}_2(\mathbb{R}^d)$$

$$(m, \sigma) \mapsto \mathcal{N}(m, \sigma^2),$$

where $\mathcal{N}(m, \sigma^2)$ is the Gaussian probability measure with mean m and standard deviation matrix σ , i.e. the absolutely continuous probability measure with respect to the Lebesgue measure \mathcal{L} on \mathbb{R}^d with density $\frac{d\mathcal{N}(m, \sigma^2)}{d\mathcal{L}}$ given by $(2\pi)^{-\frac{d}{2}} \det(\sigma)^{-1} e^{-\frac{1}{2}(x-m)^T \sigma^{-2}(x-m)}$. Defining $\text{mean} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}^d$, $\mu \mapsto \int_{\mathbb{R}^d} x d\mu(x)$, and $\text{cov} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}_{+, \text{sym}}^{d \times d}$, $\mu \mapsto \int_{\mathbb{R}^d} (x - \text{mean}(\mu))(x - \text{mean}(\mu))^T d\mu(x)$ as the mean and covariance matrix of a probability measure μ , respectively, one can straightforwardly check that the inverse $\Phi^{-1} : \Phi(\mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d}) \longrightarrow \mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d}$ is explicitly given by $\mu \mapsto (\text{mean}(\mu), \text{std}(\mu))$, where the standard deviation matrix $\text{std}(\mu)$ is the unique element in $\mathbb{R}_{+, \text{sym}}^{d \times d}$ with $\text{std}^2(\mu) = \text{cov}(\mu)$.

3.3.1 The case of general Gaussian distributions

In what follows, we explicitly compute the spline energy for curves in the space of Gaussian distributions. To this end, we will first list some facts about optimal transport in this restricted setting. Since the space of Gaussian distributions is contained in $\mathcal{P}_2^{ac}(\mathbb{R}^d)$, we shall from now on abuse notation and denote with $\text{Bar}(\mu, \nu)$ the unique element in the set of barycenters, rather than the set itself.

Proposition 3.3.2. Let $m_1, m_2 \in \mathbb{R}^d$, and $\sigma_1, \sigma_2 \in \mathbb{R}_{+, \text{sym}}^{d \times d}$. Define $\mu_1 = \mathcal{N}(m_1, \sigma_1^2)$ and $\mu_2 = \mathcal{N}(m_2, \sigma_2^2)$. Then, the following statements hold:

1. The optimal transport map T from μ_1 to μ_2 is given by $x \mapsto T(x) = m_2 + \sigma_1^{-1}(\sigma_1 \sigma_2^2 \sigma_1)^{\frac{1}{2}} \sigma_1^{-1}(x - m_1)$. If σ_1 and σ_2 are simultaneously diagonalizable, T is simplified to $T(x) = m_2 + \sigma_1^{-1} \sigma_2(x - m_1)$.
2. The squared L^2 -Wasserstein distance between μ_1 and μ_2 is $\mathcal{W}^2(\mu_1, \mu_2) = |m_1 - m_2|^2 + B^2(\sigma_1, \sigma_2)$, where $B^2(\sigma_1, \sigma_2) := \text{tr}(\sigma_1^2 + \sigma_2^2 - 2(\sigma_1 \sigma_2^2 \sigma_1)^{1/2})$ is the squared Bures-Wasserstein metric defined in [Bur69]. If σ_1 and σ_2 are simultaneously diagonalizable, B^2 is given by $B^2(\sigma_1, \sigma_2) = \|\sigma_1 - \sigma_2\|_F^2$, where $\|A\|_F^2 := \text{tr}(A^T A)$ is the Frobenius norm of a matrix $A \in \mathbb{R}^{d \times d}$.

3. For all $t \in [0, 1]$, $\text{Bar}^t(\mu_1, \mu_2)$ is a Gaussian distribution with

$$\text{mean}(\text{Bar}^t(\mu_1, \mu_2)) = (1-t)m_1 + tm_2,$$

$$\text{std}(\text{Bar}^t(\mu_1, \mu_2)) = \left[\left((1-t)\sigma_1 + t\sigma_1^{-1}(\sigma_1\sigma_2^2\sigma_1)^{\frac{1}{2}} \right) \left((1-t)\sigma_1 + t\sigma_1^{-1}(\sigma_1\sigma_2^2\sigma_1)^{\frac{1}{2}} \right)^T \right]^{\frac{1}{2}}$$

4. Let $m \in \mathbb{R}^d$, $\sigma \in \mathbb{R}_{+, \text{sym}}^{d \times d}$, and $\mu := \mathcal{N}(m, \sigma^2)$. For all $t \in [0, 1]$, $\text{Bar}_\mu^t(\mu_1, \mu_2)$ is a Gaussian distribution with

$$\text{mean}(\text{Bar}_\mu^t(\mu_1, \mu_2)) = (1-t)m_1 + tm_2,$$

$$\text{std}(\text{Bar}_\mu^t(\mu_1, \mu_2))$$

$$= \left[\left((1-t)\sigma^{-1}(\sigma\sigma_1^2\sigma)^{\frac{1}{2}} + t\sigma^{-1}(\sigma\sigma_2^2\sigma)^{\frac{1}{2}} \right) \left((1-t)\sigma^{-1}(\sigma\sigma_1^2\sigma)^{\frac{1}{2}} + t\sigma^{-1}(\sigma\sigma_2^2\sigma)^{\frac{1}{2}} \right)^T \right]^{\frac{1}{2}}.$$

Proof. 1. : See [PC⁺19], equation (2.40).

2. : See [PC⁺19], equations (2.41)-(2.42).

3. and (4) : It is straightforward to prove that for $a, b \in \mathbb{R}^d$, $A \in \mathbb{R}^{d \times d}$ and $\Sigma \in \mathbb{R}_{+, \text{sym}}^{d \times d}$ then for the map $F : x \mapsto Ax + b$, it holds $F_\# \mathcal{N}(a, \Sigma) = \mathcal{N}(Aa + b, A\Sigma A^T)$. Plugging in the explicit expression for T given in (1), and using Remarks 3.2.4 and 3.2.6 respectively, one obtains the desired results. \square

Remark 3.3.3. The above proposition implies that Wasserstein geodesics $(\mu_t)_{t \in [0, 1]}$ between two Gaussian distributions are also Gaussian distributed for all $t \in [0, 1]$. However, at this point we are not able to prove an analogous statement for Wasserstein splines.

Proposition 3.3.4 (Consistency). *Let $(m_t, \sigma_t)_t$ be a curve in $C^3([0, 1], \mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d})$, and let $(\mu_t)_t := \mathcal{N}(m_t, \sigma_t^2)$ be the respective Gaussian-valued curve. Moreover, for $k = 0, \dots, K$, define $\mu_k^K := \mu_{t_k^K}$, with $t_k^K := k/K$. Then, we have*

$$\mathcal{E}((\mu_t)_t) = \int_0^1 \left\| \sigma_t^{-1} \frac{d}{dh} \Big|_{h=0} (\sigma_t \sigma_{t+h}^2 \sigma_t)^{\frac{1}{2}} \right\|_F^2 dt + \int_0^1 |\dot{m}_t|^2 dt = \mathbf{E}^K((\mu_0^K, \dots, \mu_K^K)) + \mathcal{O}(K^{-1}), \quad (3.3.1)$$

$$\mathcal{F}((\mu_t)_t) = \int_0^1 \left\| \sigma_t^{-1} \frac{d^2}{dh^2} \Big|_{h=0} (\sigma_t \sigma_{t+h}^2 \sigma_t)^{\frac{1}{2}} \right\|_F^2 dt + \int_0^1 |\ddot{m}_t|^2 dt = \mathbf{F}_G^K((\mu_0^K, \dots, \mu_K^K)) + \mathcal{O}(K^{-1}), \quad (3.3.2)$$

where the implicit constant in $\mathcal{O}(K^{-1})$ is independent of K .

Proof. Recall from Proposition 3.3.2 that the optimal transport map T_t^s from μ_t to μ_s is given by $T_t^s(x) = A_t^s(x - m_t) + m_s$ where $A_t^s := \sigma_t^{-1} (\sigma_t \sigma_s^2 \sigma_t)^{\frac{1}{2}} \sigma_t^{-1}$. Further we note that the optimal velocity field v_t in the continuity equation solved by μ_t is given for almost all $t \in (0, 1)$ by (see [AGS08, eq. (8.4.8)])

$$v_t(x) = \partial_s \Big|_{s=t} T_t^s(x) = \partial_s \Big|_{s=t} A_t^s(x - m_t) + \dot{m}_t.$$

By the assumptions on the curve $(\mu_t, \sigma_t)_t$, the matrix-valued function A_t^s is continuously-differentiable in t and s with all derivatives up to order 3 uniformly bounded for $s, t \in [0, 1]$. Hence also $T_t^s(x)$ is continuously differentiable in t and s and all derivatives up to order three can be bounded by $C\|x\|$ for a constant C independent of $s, t \in [0, 1]$ and $x \in \mathbb{R}^d$. In particular, $v_t(x)$ is continuously differentiable in t with derivatives up to order two bounded by $C\|x\|$.

Since $\frac{T_t^{t+h} + T_t^{t-h}}{2}$ is the optimal map between $\text{Bar}_{\mu_t}(\mu_{t-h}, \mu_{t+h})$ and μ_t , cf. Remark 3.2.6, we have that

$$4 \frac{\mathcal{W}^2(\mu_t, \text{Bar}_{\mu_t}(\mu_{t-h}, \mu_{t+h}))}{h^4} = \int_{\mathbb{R}^d} \frac{|T_t^{t+h} - 2 \cdot \text{Id} + T_t^{t-h}|^2}{h^4} d\mu_t. \quad (3.3.3)$$

By Taylor expansion, we have

$$\begin{aligned} T_t^{t+h} &= \text{Id} - h \partial_s|_{s=t+h} T_s^{t+h} + \frac{1}{2} h^2 \partial_s^2|_{s=t+h} T_s^{t+h} - \frac{1}{6} h^3 \partial_s^3|_{s=u} T_s^{t+h}, \\ T_t^{t-h} &= \text{Id} - h \partial_s|_{s=t} T_t^s + \frac{1}{2} h^2 \partial_s^2|_{s=t} T_t^s - \frac{1}{6} h^3 \partial_s^3|_{s=v} T_t^s, \end{aligned}$$

for some $u \in (t, t+h)$ and $v \in (t-h, t)$.

Taking into account the fact that $T_t^t \circ T_s^t = \text{Id}$ and taking first and second derivatives in s of this identity at $s = t$ we readily deduce that

$$\partial_s|_{s=t} T_t^t = -\partial_s|_{s=t} T_t^s = -v_t,$$

and

$$\partial_s^2|_{s=t} T_t^t = -\partial_s^2|_{s=t} T_t^s - 2(D\partial_s|_{s=t} T_t^s) \partial_s|_{s=t} T_t^s.$$

The last term evaluates to

$$-2D(\partial_s|_{s=t} T_t^s) \partial_s|_{s=t} T_t^t = 2(Dv_t)v_t = \nabla|v_t|^2.$$

Collecting these observations, we obtain

$$\begin{aligned} \frac{T_t^{t+h} - 2\text{Id} + T_t^{t-h}}{h^2} &= \frac{1}{h} \left[-\partial_s|_{s=t+h} T_s^{t+h} - \partial_s|_{s=t} T_t^s \right] + \frac{1}{2} \left[\partial_s^2|_{s=t+h} T_s^{t+h} + \partial_s^2|_{s=t} T_t^s \right] + O(h) \\ &= \frac{1}{h} (v_{t+h} - v_t) + \frac{1}{2} \left[-\partial_s^2|_{s=t+h} T_{t+h}^s + \partial_s^2|_{s=t} T_t^s + \nabla|v_{t+h}|^2 \right] + O(h) \\ &= \dot{v}_t + \frac{1}{2} \nabla|v_t|^2 + O(h), \end{aligned} \quad (3.3.4)$$

where the terms $O(h)$ are bounded by $Ch\|x\|$ for a constant C independent of t, h and x . Hence,

$$4 \frac{\mathcal{W}^2(\mu_t, \text{Bar}_{\mu_t}(\mu_{t-h}, \mu_{t+h}))}{h^4} = \int_{\mathbb{R}^d} \left| \dot{v}_t + \frac{1}{2} \nabla|v_t|^2 \right|^2 d\mu_t + \mathcal{O}(h),$$

with $\mathcal{O}(h)$ bounded by Ch for a uniform constant C . Finally, recall the rectangular quadrature rule $\int_0^1 f(t) dt = K^{-1} \sum_{k=1}^{K-1} f(t_k^K) + \mathcal{O}(K^{-1})$ with $t_k^K := k/K$ for a Lipschitz function f , where the implicit constant in the $\mathcal{O}(K^{-1})$ term depends only on $\text{Lip}(f)$. Setting $h = K^{-1}$ and defining $\mu_k^K := \mu_{t_k^K}$, we obtain for a uniform constant C :

$$\left| \int_0^1 \int_{\mathbb{R}^d} \left| \dot{v}_t + \frac{1}{2} \nabla|v_t|^2 \right|^2 d\mu_t dt - \mathbf{F}_G^K((\mu_0^K, \dots, \mu_K^K)) \right| \leq CK^{-1}.$$

The continuous spline energy on Gaussian distributions is given by

$$\begin{aligned} \int_0^1 \int_{\mathbb{R}^d} |\dot{v}_t + \frac{1}{2} \nabla |v_t|^2|^2 d\mu_t dt &= \int_0^1 \int_{\mathbb{R}^d} \left| \frac{d^2}{dh^2} \Big|_{h=0} \left(T_t^{t+h} \right) \right|^2 d\mu_t dt \\ &= \int_0^1 \text{tr} \left(\sigma_t^{-1} \left(\frac{d^2}{dh^2} \Big|_{h=0} (\sigma_t \sigma_{t+h}^2 \sigma_t)^{\frac{1}{2}} \right)^2 \sigma_t^{-1} \right) dt + \int_0^1 |\ddot{m}_t|^2 dt \\ &= \int_0^1 \left\| \sigma_t^{-1} \frac{d^2}{dh^2} \Big|_{h=0} (\sigma_t \sigma_{t+h}^2 \sigma_t)^{\frac{1}{2}} \right\|_F^2 dt + \int_0^1 |\ddot{m}_t|^2 dt, \end{aligned}$$

where we have used in the first step the expansion (3.3.4) and in the second step the decoupling of the energies from (3.2.9). The first equality in (3.3.1) can be proven once again by using (3.2.9) and the explicit form of v_t . Finally, the corresponding estimates for (3.3.1) can be proven similarly by Taylor approximation of the optimal map to first order. Namely, in place of (3.3.3) one uses that

$$\frac{1}{h^2} \mathcal{W}^2(\mu_t, \mu_{t+h}) = \int_{\mathbb{R}^d} \frac{|T_t^{t+h} - \text{Id}|^2}{h^2} d\mu_t,$$

as well as $T_t^{t+h} = \text{Id} + h v_t + \mathcal{O}(h^2)$ with $\mathcal{O}(h^2)$ being bounded by $C\|x\|$ with a uniform constant C , so that

$$\left| \int_0^1 \int_{\mathbb{R}^d} |v_t|^2 d\mu_t dt - \mathbf{E}^K((\mu_0^K, \dots, \mu_K^K)) \right| \leq CK^{-1}.$$

□

Remark 3.3.5. We expect the previous consistency result for the generalised discrete spline energy to hold for general curves $(\mu_t)_t$ with sufficiently regular densities and velocity fields. In fact, we note that the argument relies on the Gaussian structure essentially only for the explicit error estimates in the Taylor expansion of the optimal maps. In particular, we expect the identity $\nabla_{v_t} v_t = \frac{d^2}{dh^2} \Big|_{h=0} T_t^{t+h}$ to hold true for general sufficiently regular curves. However, obtaining a general consistency result for the discrete spline energy with the true barycenter seems much more delicate.

To show the consistency of the proposed non-generalized discrete spline energy functional $\mathbf{F}^K((\mu_0^K, \dots, \mu_K^K))$, we shall need the following lemmata (which are restricted to the case $d = 2$), which relate the barycenter with the generalized barycenter:

Lemma 3.3.6. *Let $m_1, m_2 \in \mathbb{R}^2$ and $\sigma_1, \sigma_2 \in \mathbb{R}_{+, \text{sym}}^{2 \times 2}$, such that $\|\sigma_1 - \sigma_2\|_F \leq 2h \leq 2$. Then, one obtains*

$$\left\| \text{std}(\text{Bar}(\mu_1, \mu_2)) - \frac{\sigma_1 + \sigma_2}{2} \right\|_F \leq Ch^2,$$

where $\mu_1 = \mathcal{N}(m_1, \sigma_1^2)$, $\mu_2 = \mathcal{N}(m_2, \sigma_2^2)$ and the constant C only depends on $\max\{\lambda_{\max}(\sigma_1), \lambda_{\max}(\sigma_2)\}$ and $\min\{\lambda_{\min}(\sigma_1), \lambda_{\min}(\sigma_2)\}$, where $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ denote the largest and smallest eigenvalue of a symmetric matrix A , respectively.

Proof. Recall that in 2D, the positive definite square root of a positive definite matrix $\sigma = \begin{pmatrix} a & c \\ c & b \end{pmatrix}$ is given by the following explicit formula:

$$\sigma^{\frac{1}{2}} = (\text{tr}(\sigma) + 2\sqrt{\det(\sigma)})^{-\frac{1}{2}} (\sigma + \sqrt{\det(\sigma)} \mathbb{1}).$$

Indeed,

$$\begin{aligned}
& (\operatorname{tr}(\sigma) + 2\sqrt{\det(\sigma)})^{-1}(\sigma + \sqrt{\det(\sigma)}\mathbf{1})^2 \\
&= (a + b + 2\sqrt{\det(\sigma)})^{-1}(\sigma^2 + 2\sigma\sqrt{\det(\sigma)} + \det(\sigma)\mathbf{1}) \\
&= (a + b + 2\sqrt{\det(\sigma)})^{-1}(\sigma + 2\sqrt{\det(\sigma)}\mathbf{1} + \det(\sigma)\sigma^{-1})\sigma \\
&= (a + b + 2\sqrt{\det(\sigma)})^{-1} \left(\begin{pmatrix} a & c \\ c & b \end{pmatrix} + 2\sqrt{\det(\sigma)}\mathbf{1} + \begin{pmatrix} b & -c \\ -c & a \end{pmatrix} \right) \sigma \\
&= (a + b + 2\sqrt{\det(\sigma)})^{-1}(a + b + 2\sqrt{\det(\sigma)})\sigma = \sigma.
\end{aligned}$$

Let us now define $\sigma := \frac{\sigma_1 + \sigma_2}{2}$, and $\sigma' := \frac{\sigma_2 - \sigma_1}{2h}$ and $\sigma_t = \sigma + t\sigma'$, for $t \in [-h, h]$. Moreover, we define $\mu_t := \mathcal{N}(0, \sigma_t^2)$ and $F : \mathbb{R} \rightarrow \mathbb{R}_{+, \text{sym}}^{2 \times 2}; t \mapsto F(t) = \operatorname{std}(\operatorname{Bar}(\mu_{-t}, \mu_t))$, which implies $\operatorname{std}(\operatorname{Bar}(\mu_1, \mu_2)) = F(h)$. Since F is actually smooth, we have by Taylor's theorem

$$F(h) = F(0) + \dot{F}(0)h + \frac{1}{2}\ddot{F}(s)h^2,$$

for some $s \in [0, h]$. It is straightforward to see that $F(0) = \operatorname{std}(\mu_0) = \frac{\sigma_1 + \sigma_2}{2}$. Next, we prove that $\dot{F}(0) = 0$. To this end, we show that if $\sigma : [0, 1] \rightarrow \mathbb{R}_{+, \text{sym}}^{2 \times 2}; t \mapsto \sigma_t$ is C^1 , with $\dot{\sigma}_t = 0$ for some $t \in [0, 1]$, then we have $\frac{d}{dt} \left[\sigma_t^{\frac{1}{2}} \right] = 0$. This is easily verified using the matrix square root formula above:

$$\begin{aligned}
\frac{d}{dt} \left[\sigma_t^{\frac{1}{2}} \right] &= -\frac{1}{2}(\operatorname{tr}(\sigma_t) + 2\sqrt{\det(\sigma_t)})^{-\frac{3}{2}}(\sigma_t + \sqrt{\det(\sigma_t)}\mathbf{1})(\operatorname{tr}(\dot{\sigma}_t) + \sqrt{\det(\sigma_t)}\operatorname{tr}(\dot{\sigma}_t\sigma_t^{-1})) \\
&\quad + (\operatorname{tr}(\sigma_t) + 2\sqrt{\det(\sigma_t)})^{-\frac{1}{2}}(\dot{\sigma}_t + \frac{1}{2}\sqrt{\det(\sigma_t)}\operatorname{tr}(\dot{\sigma}_t\sigma_t^{-1})\mathbf{1}) = 0.
\end{aligned}$$

Hence, it will be sufficient to prove $\frac{d}{dt}|_{t=0} F^2(t) = 0$. This is obviously true since F^2 is symmetric with respect to $t = 0$.

$$\ddot{F}(s) = \frac{d^2}{ds^2}(F^2)^{\frac{1}{2}}(s) = \frac{d^2}{ds^2} \left(\operatorname{tr}(F^2(s)) + 2\sqrt{\det(F^2(s))} \right)^{-\frac{1}{2}} (F^2(s) + \sqrt{\det(F^2(s))}\mathbf{1}).$$

Recall that

$$F^2(s) = \frac{1}{4} \left(\sigma_{-s}^2 + \sigma_s^2 + \sigma_s^{-1}(\sigma_s\sigma_{-s}^2\sigma_s)^{\frac{1}{2}}\sigma_s + \sigma_s(\sigma_s\sigma_{-s}^2\sigma_s)^{\frac{1}{2}}\sigma_s^{-1} \right).$$

Due to the formula of the matrix square root for 2×2 matrices, in order to find upper bounds of $\|\ddot{F}(s)\|_F$ it is enough to prove upper and (positive) lower bounds of $\|\sigma_s\|_F$, $\|\sigma_s^{-1}\|_F$, $\operatorname{tr}(\sigma_s)$, $\operatorname{tr}(\sigma_s^{-1})$, and $\det(\sigma_s)$ which are independent of $s \in (-h, h)$. For a matrix $A \in \mathbb{R}_{+, \text{sym}}^{2 \times 2}$, we have the following sequence of inequalities:

$$\begin{aligned}
\frac{1}{\sqrt{2}} \operatorname{tr}(A) &= \frac{1}{\sqrt{2}} \operatorname{tr}(A \cdot \mathbf{1}) \leq \frac{1}{\sqrt{2}} \|A\|_F \|\mathbf{1}\|_F = \|A\|_F = \sqrt{\operatorname{tr}(A^T A)} \\
&= \sqrt{\operatorname{tr}(A^2)} = \sqrt{\lambda_1^2 + \lambda_2^2} \leq \lambda_1 + \lambda_2 = \operatorname{tr}(A).
\end{aligned}$$

Hence, it will suffice to prove upper and (positive) lower bounds of $\|\sigma_s\|_F$, $\|\sigma_s^{-1}\|_F$ and $\det(\sigma_s)$, independent of $s \in (-h, h)$. Indeed, note that for $s \in (-h, h)$, we have that

$$\sigma_s = \left(1 - \frac{s+h}{2h} \right) \sigma_1 + \frac{s+h}{2h} \sigma_2,$$

is the sum of positive definite, symmetric matrices. By Weyl's inequality, we obtain that the smallest eigenvalue of σ_s , denoted as $\lambda_{\min}(\sigma_s)$ is bounded from below by $(1 - \frac{s+h}{2h}) \lambda_{\min}(\sigma_1) + \frac{s+h}{2h} \lambda_{\min}(\sigma_2)$ which is bounded from below by $\xi := \min(\lambda_{\min}(\sigma_1), \lambda_{\min}(\sigma_2))$, and is in particular independent of s . Similarly, one obtains $\lambda_{\max}(\sigma_s) \leq \zeta$, where $\zeta := \max(\lambda_{\max}(\sigma_1), \lambda_{\max}(\sigma_2))$. Hence, one finally obtains $\sqrt{2\xi} \leq \|\sigma_s\| \leq 2\zeta$, $\sqrt{2\zeta^{-1}} \leq \|\sigma_s^{-1}\| \leq 2\xi^{-1}$, $\xi^2 \leq \det(\sigma_s) \leq \zeta^2$ for all $s \in [-h, h]$. \square

Remark 3.3.7. This can be considered as the analogue (with a slight improvement on the exponent) of [RW15], Lemma 5.7.

Lemma 3.3.8. Let $m_i \in \mathbb{R}^2$ and $\sigma_i \in \mathbb{R}_{+, \text{sym}}^{2 \times 2}$ for $i = 1, 2, 3$ such that $\|\sigma_i - \sigma_2\|_F \leq 2h$ for $i = 1, 3$, and $\|\sigma_2 - \frac{\sigma_3 + \sigma_1}{2}\|_F \leq h^2$. Then, we can verify the estimates

$$\mathcal{W}^2(\mu_2, \text{Bar}_{\mu_2}(\mu_1, \mu_3)) = \mathcal{W}^2(\mu_2, \text{Bar}(\mu_1, \mu_3)) + \mathcal{O}(h^5), \quad (3.3.5)$$

and

$$\mathcal{W}^2(\mu_2, \text{Bar}_{\mu_2}(\mu_1, \mu_3)) = \frac{1}{2}\mathcal{W}^2(\mu_1, \mu_2) + \frac{1}{2}\mathcal{W}^2(\mu_3, \mu_2) - \frac{1}{4}\mathcal{W}^2(\mu_1, \mu_3) + \mathcal{O}(h^5), \quad (3.3.6)$$

where $\mu_i := \mathcal{N}(m_i, \sigma_i^2)$ for $i = 1, 2, 3$.

This lemma can be proven fully analogously to the previous one. However, as one needs to expand the terms up to the fifth order, the computations become extremely lengthy, so we will leave out the explicit computations and simply give a sketch of the proof:

Let $(\sigma(t))_{t \in [-h, h]}$ be the uniquely determined second-order $\mathbb{R}_{+, \text{sym}}^{2 \times 2}$ -valued curve, such that $\sigma(-h) = \sigma_1$, $\sigma(0) = \sigma_2$ and $\sigma(h) = \sigma_3$, and let $(\mu_t)_{t \in [-h, h]}$ be the respective measure-valued curve. Then, define

$$\begin{aligned} F_1(t) &:= \mathcal{W}^2(\mu_2, \text{Bar}_{\mu_2}(\mu_{-t}, \mu_t)), \\ F_2(t) &:= \mathcal{W}^2(\mu_2, \text{Bar}(\mu_{-t}, \mu_t)), \\ F_3(t) &:= \frac{1}{2}\mathcal{W}^2(\mu_{-h}, \mu_2) + \frac{1}{2}\mathcal{W}^2(\mu_h, \mu_2) - \frac{1}{4}\mathcal{W}^2(\mu_{-h}, \mu_h). \end{aligned}$$

Then, one has $F_1(h) = \mathcal{W}^2(\mu_2, \text{Bar}_{\mu_2}(\mu_1, \mu_3))$, $F_2(h) = \mathcal{W}^2(\mu_2, \text{Bar}(\mu_1, \mu_3))$, and $F_3(h) = \frac{1}{2}\mathcal{W}^2(\mu_1, \mu_2) + \frac{1}{2}\mathcal{W}^2(\mu_3, \mu_2) - \frac{1}{4}\mathcal{W}^2(\mu_1, \mu_3)$. Now, we can explicitly compute any derivatives of the F_i , and expand them at $t = 0$ up to the fifth order, i.e.

$$F_i(h) = \sum_{k=0}^4 \frac{1}{k!} F_i^{(k)}(0) h^k + \frac{1}{5!} F_i^{(5)}(s_i) h^5,$$

for some $s_i \in (0, h)$. Now, one checks that all derivatives up to the third order vanish for $i = 1, 2, 3$. The zeroth-order derivative vanishing is trivial, while the first and third order derivatives vanish at 0 due to the symmetry of the derivatives of lesser order. For the fourth order derivatives, one checks that they all coincide for $i = 1, 2, 3$. Now, it remains to show that the fifth-order derivative can be bounded by a constant independent of h . To this end, one uses Lemma 3.3.6 and the same estimation strategy as inside the proof of this lemma.

Remark 3.3.9. Let $t \mapsto (m_t, \sigma_t)$ for $t \in [0, 1]$ be a curve in $\mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d}$, such that $(\sigma_t)_{t \in [0, 1]}$ is simultaneously diagonalizable. After choosing a common diagonal basis, one may without loss of generality regard $(\sigma_t)_{t \in [0, 1]}$ as a curve in $\mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d}$ instead.

3.3.2 The case of Gaussian distributions with diagonal covariance matrices

Definition 3.3.10. Let Φ be the map from Definition 3.3.1, i.e. $\Phi(m, \sigma) = \mathcal{N}(m, \sigma^2)$ for $(m, \sigma) \in \mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d}$. Then, we define $\mathcal{P}_2^{G, d} := \Phi(\mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d})$ as the space of non-degenerate Gaussian distributions with diagonal covariance matrices.

Corollary 3.3.11. Let $(m_t, \sigma_t)_t$ be a curve in $C^3([0, 1], \mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d})$, and let $(\mu_t)_t := \mathcal{N}(m_t, \sigma_t^2)$ be the respective $\mathcal{P}_2^{G, d}$ -valued curve. Then, we have

$$\mathcal{E}((\mu_t)_t) = \int_0^1 \|\dot{\sigma}_t\|_F^2 dt + \int_0^1 |\dot{m}_t|^2 dt, \quad (3.3.7)$$

$$\mathcal{F}((\mu_t)_t) = \int_0^1 \|\ddot{\sigma}_t\|_F^2 dt + \int_0^1 |\ddot{m}_t|^2 dt, \quad (3.3.8)$$

where \mathcal{F} is the spline energy on the space of diagonal Gaussian distributions $\Phi(\mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d}) \subset \mathcal{P}_2(\mathbb{R}^d)$.

Proof. Using eq. (3.3.1) and (3.3.2), and assuming $\sigma_t \in \mathbb{R}_{+, \text{dia}}^{d \times d}$ for all $t \in [0, 1]$, we have

$$\begin{aligned} \sigma_t^{-1} \frac{d}{dh} \Big|_{h=0} (\sigma_t \sigma_{t+h}^2 \sigma_t)^{\frac{1}{2}} &= \sigma_t^{-1} \frac{d}{dh} \Big|_{h=0} \sigma_{t+h} \sigma_t = \frac{d}{dh} \Big|_{h=0} \sigma_{t+h} = \dot{\sigma}_t, \\ \sigma_t^{-1} \frac{d^2}{dh^2} \Big|_{h=0} (\sigma_t \sigma_{t+h}^2 \sigma_t)^{\frac{1}{2}} &= \sigma_t^{-1} \frac{d^2}{dh^2} \Big|_{h=0} \sigma_{t+h} \sigma_t = \frac{d^2}{dh^2} \Big|_{h=0} \sigma_{t+h} = \ddot{\sigma}_t, \end{aligned}$$

which proves the claim. \square

Alternatively, one can "brute-force" this:

Example 3.3.12. Let $U \subseteq \mathbb{R}^n$, $V \subseteq \mathbb{R}^m$ be open subsets. Then, $H^k(U, V) := W^{k, 2}(U, V)$ denotes the Sobolev space of functions $f : U \rightarrow V$, such that f and its weak derivatives up to order k have finite L^2 -norm. Let $t \mapsto (m_t, \sigma_t)$ for $t \in (0, 1)$ be a curve in $H^2((0, 1), \mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d})$. Then, by abusing notation we can define $G_t := (2\pi)^{-\frac{d}{2}} \det(\sigma_t)^{-1} e^{-\frac{1}{2}(x-m_t)^T \sigma_t^{-2} (x-m_t)}$, i.e. the Lebesgue density function of $\mu_t := \mathcal{N}(m_t, \sigma_t^2)$. We have

$$\begin{aligned} \partial_t G_t &= [-\text{tr}(\sigma_t^{-1} \dot{\sigma}_t) + \langle \dot{m}_t, \sigma_t^{-2} (x - m_t) \rangle + \langle x - m_t, \sigma_t^{-3} \dot{\sigma}_t (x - m_t) \rangle] G_t, \\ \nabla G_t &= [-\sigma_t^{-2} (x - m_t)] G_t. \end{aligned}$$

Take $\varphi_t(x) := \langle x, \dot{m}_t \rangle + \frac{1}{2} \langle x - m_t, \dot{\sigma}_t \sigma_t^{-1} (x - m_t) \rangle$. Then, we obtain

$$\begin{aligned} \nabla \varphi_t &= \dot{m}_t + \dot{\sigma}_t \sigma_t^{-1} (x - m_t), \\ \Delta \varphi_t &= \text{tr}(\dot{\sigma}_t \sigma_t^{-1}). \end{aligned} \quad (3.3.9)$$

Thus, the pair (μ_t, v_t) with $v_t = \nabla \varphi_t$ satisfies (CE), i.e.

$$\partial_t \mu_t + \nabla \cdot (\nabla \varphi_t \mu_t) = \partial_t \mu_t + \Delta \varphi_t \mu_t + \nabla \varphi_t \cdot \nabla \mu_t \equiv 0.$$

Moreover, we have that

$$\nabla \dot{\varphi}_t = \ddot{m}_t + \ddot{\sigma}_t \sigma_t^{-1} (x - m_t) - \dot{\sigma}_t \sigma_t^{-1} \dot{m}_t - \dot{\sigma}_t^2 \sigma_t^{-2} (x - m_t),$$

$$\frac{1}{2} \nabla |\nabla \varphi_t|^2 = \nabla^2 \varphi_t \nabla \varphi_t = \dot{\sigma}_t \sigma_t^{-1} (\dot{m}_t + \dot{\sigma}_t \sigma_t^{-1} (x - m_t)),$$

and hence we finally compute

$$\begin{aligned} \nabla \left(\dot{\varphi}_t + \frac{1}{2} |\nabla \varphi_t|^2 \right) &= \ddot{m}_t + \ddot{\sigma}_t \sigma_t^{-1} (x - m_t) - \dot{\sigma}_t \sigma_t^{-1} \dot{m}_t - \dot{\sigma}_t^2 \sigma_t^{-2} (x - m_t) \\ &\quad + \dot{\sigma}_t \sigma_t^{-1} (\dot{m}_t + \dot{\sigma}_t \sigma_t^{-1} (x - m_t)) = \ddot{m}_t + \ddot{\sigma}_t \sigma_t^{-1} (x - m_t). \end{aligned}$$

The continuous spline energy is then given by

$$\mathcal{F}(\mu) = \int_0^1 \int_{\mathbb{R}^N} |\ddot{m}_t + \ddot{\sigma}_t \sigma_t^{-1} (x - m_t)|^2 d\mu_t dt \quad (3.3.10)$$

$$= \int_0^1 \int_{\mathbb{R}^N} |\ddot{m}_t|^2 + |\ddot{\sigma}_t \sigma_t^{-1} (x - m_t)|^2 d\mu_t dt \quad (3.3.11)$$

$$= \int_0^1 |\ddot{m}_t|^2 + \text{tr}(\ddot{\sigma}_t^2) dt, \quad (3.3.12)$$

where in the second equality we used the fact that the μ_t -integral of an antisymmetric function (with respect to m_t) vanishes, and $\int_{\mathbb{R}^d} |A(x - m_t)|^2 d\mu_t = \text{tr}(A \sigma_t^2 A^T)$ in the last equality (cf. proof of Proposition 3.3.4). One can simplify the above expression even further:

$$\mathcal{F}(\mu) = \int_0^1 |\ddot{m}_t|^2 dt + \sum_{j=1}^d \int_0^1 |\ddot{\lambda}_t^j|^2 dt, \quad (3.3.13)$$

where $(\lambda_t^j)_{j=1,\dots,d}$ are the eigenvalues of σ_t . In view of (3.3.13), one might be tempted to assert that a spline interpolation on the space of Gaussian distributions with simultaneously diagonalizable covariances can be obtained by spline interpolating each eigenvalue independently (after choosing a fixed common eigenbasis). However, this assumption ignores the restriction that σ_t is required to be positive definite. Indeed, the spline interpolation of some given Gaussian key-frames amounts to solving the following minimization problem

$$\inf_{m_t \in \mathbb{R}^d, \lambda_t^j > 0} \int_0^1 |\ddot{m}_t|^2 dt + \sum_{j=1}^d \int_0^1 |\ddot{\lambda}_t^j|^2 dt,$$

together with some point-wise evaluation constraints. If the above minimization problem has a solution, then the spline interpolation results from a classical cubic spline interpolation of m_t and λ_t^j . The positivity constraint, however, implies that whenever the interpolating eigenvalue-spline becomes negative, it can not coincide with the Wasserstein E-spline.

The above equation (3.3.10) allows us to canonically identify any $\mathcal{P}_2^{G,d}$ -valued functional \mathcal{F} with a $(\mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d})$ -valued functional $\hat{\mathcal{F}}$ via

$$\hat{\mathcal{F}}((m_t, \sigma_t)_t) := \mathcal{F}((\mathcal{N}(m_t, \sigma_t^2))_t).$$

Lemma 3.3.13. *The regularized spline energy $\hat{\mathcal{F}}^\delta$ is lower semi-continuous under weak and continuous under strong convergence in $H^2((0, 1), \mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d})$.*

Proof. First, let us note that by (3.3.7) and (3.3.8), for any curve $(m_t, \sigma_t)_{t \in [0,1]}$ the path energy $\hat{\mathcal{E}}$ and the spline energy $\hat{\mathcal{F}}$ coincide with the squared semi-norms $|\cdot|_{H^1}^2$ and $|\cdot|_{H^2}^2$, respectively, both of which are weakly lower semi-continuous on H^2 . Thus,

$$\liminf_{n \rightarrow \infty} \hat{\mathcal{F}}^\delta((m_t^{(n)}, \sigma_t^{(n)})_t) \geq \hat{\mathcal{F}}^\delta((m_t, \sigma_t)_t),$$

for a $H^2((0,1); \mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d})$ -weakly convergent sequence $(m_t^{(n)}, \sigma_t^{(n)})_t \rightharpoonup (m_t, \sigma_t)_t$. Moreover, if the sequence converges strongly we also obtain $\liminf_{n \rightarrow \infty} \hat{\mathcal{F}}^\delta((m_t^{(n)}, \sigma_t^{(n)})_t) = \hat{\mathcal{F}}^\delta((m_t, \sigma_t)_t)$, since strong convergence implies convergence in norm. \square

3.4 Convergence of discrete Gaussian E-splines

In this section we will discuss the convergence of discrete Gaussian spline curves to continuous Gaussian E-splines. For the sake of presentation, we will at first only consider centered Gaussian curves, i.e. $m_t = 0$ for all $t \in [0,1]$. As the energy of the mean and standard deviation matrices decouples (cf. equations (3.2.9) and (3.3.8)), this is a very natural approach: for the general non-centered case, one can directly apply the results from [HRW18] to obtain Mosco convergence (cf. [Mos69]) of the mean term of the functionals. As the space of standard deviation matrices $\mathbb{R}_{+, \text{dia}}^{d \times d}$ has a non-trivial boundary the results of [HRW18] do not apply to the standard deviation matrix term. Nevertheless we follow the general procedure of the proof in this paper.

In the sequel, we will focus on natural boundary conditions with a comment on the periodic case below. We will now use a suitable interpolation to identify discrete curves with continuous ones to be able to rewrite the discrete energy as a functional on time-continuous curves. As in [HRW18] and [JRR23], this will be done via cubic Hermite interpolation at time interval midpoints. For a tuple $\sigma^K := (\sigma_0, \dots, \sigma_K) \in (\mathbb{R}_{+, \text{dia}}^{d \times d})^{K+1}$, we define the temporal extension η_{σ^K} of σ^K as

$$\eta_{\sigma^K}(t) := \begin{cases} \sigma_0 + (\sigma_1 - \sigma_0)Kt, & \text{if } t \in [0, t_{1/2}^K], \\ \frac{\sigma_{k-1} + \sigma_k}{2} + (\sigma_k - \sigma_{k-1})K(t - t_{k-1/2}^K) + (\sigma_{k+1} - 2\sigma_k + \sigma_{k-1})K^2 \frac{(t - t_{k-1/2}^K)^2}{2}, & \text{if } t \in [t_{k-1/2}^K, t_{k+1/2}^K], \\ \sigma_{K-1} + (\sigma_K - \sigma_{K-1})K(t - t_{K-1/2}^K), & \text{if } t \in [t_{K-1/2}^K, 1], \end{cases}$$

where $t_{k+1/2}^K := \frac{k+1/2}{K}$ for $k = 0, \dots, K-1$. For periodic boundary conditions, we can neglect the definition on the starting and final half-intervals, identifying $[t_{K-1/2}^K, t_{K+1/2}^K]$ with $[0, t_{1/2}^K] \cup [t_{K-1/2}^K, 1]$.

Next, we recall the following convergence of a piecewise cubic Hermite interpolation from [HRW18, Lemma 4.3].

Lemma 3.4.1 (Strong convergence of piecewise cubic Hermite curves to smooth Gaussian curves). *Let $\sigma = (\sigma(t))_{t \in [0,1]}$ be a C^3 curve in $\mathbb{R}_{+, \text{dia}}^{d \times d}$, and define $\sigma^K := (\sigma_j^K)_{j=0, \dots, K} = (\sigma(j/K))_{j=0, \dots, K}$, i.e. σ^K is an equidistant sampling of the continuous curve σ with $K+1$ samples. Then, η_{σ^K} converges strongly in $H^2([0,1], \mathbb{R}_{+, \text{dia}}^{d \times d})$ to σ for $K \rightarrow \infty$.*

The next two lemmas compare the discrete path and spline energy with the corresponding continuous counterpart evaluated on the piecewise cubic Hermite interpolation. To this end, we define the hat operator for discrete functionals: for a functional \mathbf{F}^K on $(\mathcal{P}_2^{G,d})^{K+1}$ we define $\hat{\mathbf{F}}^K$ on $(\mathbb{R}^d \times \mathbb{R}_{+, \text{dia}}^{d \times d})^{K+1}$

via

$$\hat{\mathbf{F}}^K((m_k, \sigma_k^2)_{k=0, \dots, K}) := \mathbf{F}^K((\mathcal{N}(m_k, \sigma_k^2))_{k=0, \dots, K}).$$

Lemma 3.4.2 (Path energy estimate). *Let $\sigma = (\sigma(t))_{t \in [0,1]} \in H^2((0,1), \mathbb{R}_{+,dia}^{d \times d})$, and define σ^K as above. Then, for K big enough, we have $|\hat{\mathcal{E}}[\eta_{\sigma^K}] - \hat{\mathbf{E}}^K[\sigma^K]| \leq CK^{-1}$, where \mathcal{E} and \mathbf{E}^K have been defined in (3.1.2) and (3.2.12), respectively, and the constant C depends only on the curve σ .*

Proof. By the definition of $\eta_{\sigma^K}(t)$, and using (3.3.7) and Proposition 3.3.2 (2) for the expressions of $\hat{\mathcal{E}}$ and $\hat{\mathbf{E}}$, respectively, we obtain

$$\begin{aligned} & |\hat{\mathcal{E}}[\eta_{\sigma^K}] - \hat{\mathbf{E}}^K[\sigma^K]| \\ &= \left| K \sum_{k=1}^{K-1} |\sigma_k^K - \sigma_{k-1}^K|^2 + \sum_{k=1}^{K-1} \int_{t_{k-1/2}^K}^{t_{k+1/2}^K} |\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K|^2 K^4 (t - t_{k-1/2}^K)^2 dt \right. \\ & \quad \left. + 2 \sum_{k=1}^{K-1} \int_{t_{k-1/2}^K}^{t_{k+1/2}^K} (\sigma_k^K - \sigma_{k-1}^K)(\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K) K^3 (t - t_{k-1/2}^K) dt - K \sum_{k=1}^K |\sigma_k^K - \sigma_{k-1}^K|^2 \right| \\ &\leq K \sum_{k=1}^{K-1} |\sigma_k^K - \sigma_{k-1}^K| |\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K| + \frac{K}{3} \sum_{k=1}^{K-1} |\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K|^2 \\ &\leq K \left(\sum_{k=1}^{K-1} |\sigma_k^K - \sigma_{k-1}^K|^2 \right)^{\frac{1}{2}} \left(\sum_{k=1}^{K-1} |\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K|^2 \right)^{\frac{1}{2}} + \frac{K}{3} \sum_{k=1}^{K-1} |\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K|^2 \\ &\leq C' K K^{-\frac{1}{2}} |\sigma|_{H^1} K^{-\frac{3}{2}} |\sigma|_{H^2} + C'' K^{-2} |\sigma|_{H^2}^2 \leq CK^{-1}, \end{aligned}$$

where we used $K \sum_{k=1}^{K-1} |\sigma_k^K - \sigma_{k-1}^K|^2 \leq C^* |\sigma|_{H^1}^2$, and $K^3 \sum_{k=1}^{K-1} |\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K|^2 \leq C'' |\sigma|_{H^2}^2$ (cf. proof of [HRW18, Lemma 4.2]). The final inequality holds for K chosen big enough. \square

Lemma 3.4.3 (Spline energy estimate). *Let $\sigma = (\sigma(t))_{t \in [0,1]} \in H^2((0,1), \mathbb{R}_{+,dia}^{d \times d})$, and define σ^K as above. Then, we have $\hat{\mathcal{F}}[\eta_{\sigma^K}] = \hat{\mathbf{F}}^K[\sigma^K]$, where \mathcal{F} and \mathbf{F}^K have been defined in (3.2.2) and (3.2.17), respectively.*

Proof. Recall that $\ddot{\eta}_{\sigma^K}(t) = (\sigma_{k+1}^K - 2\sigma_k^K + \sigma_{k-1}^K)K^2$ for $t \in [t_{k-1/2}^K, t_{k+1/2}^K]$, $k = 1, \dots, K-1$, and 0 otherwise. We then obtain by using (3.3.8) and Proposition 3.3.2

$$\hat{\mathcal{F}}[\eta_{\sigma^K}] - \hat{\mathbf{F}}^K[\sigma^K] = 4K^3 \sum_{k=1}^{K-1} \left| \sigma_k^K - \frac{\sigma_{k+1}^K + \sigma_{k-1}^K}{2} \right|^2 - \hat{\mathbf{F}}^K[\sigma^K] = 0.$$

\square

We are now in the position to prove the convergence of the discrete spline functional to the continuous one. To this end, we introduce two indicator functions to filter the constraints. For the continuous problem $\mathcal{I}[\sigma] = 0$ if $\mu_t := \mathcal{N}(0, \sigma_t^2)$ satisfies the evaluation constraints (3.2.3) (with $\bar{\mu}_i := \mathcal{N}(0, \bar{\sigma}_i^2)$) for given interpolation constraints $(\bar{t}_i, \bar{\sigma}_i)_{i=1, \dots, l}$ as well as the corresponding boundary condition from (3.2.4)-(3.2.6) and ∞ else. Furthermore, for the discrete problem $\mathcal{I}^K[\sigma] = 0$ if $\sigma = \eta_{(\sigma_0, \dots, \sigma_K)}$ for some

$(\sigma_0, \dots, \sigma_K) \in (\mathbb{R}_{+, \text{dia}}^{d \times d})^{K+1}$, where $\mathcal{N}(0, \sigma_i^2)_{i=0, \dots, K}$ satisfies the evaluation constraints (3.2.20) as well as the corresponding discrete boundary condition from (3.2.21)-(3.2.23), and ∞ else.

Regarding the compatibility (cf. equation (3.2.20)) of the given interpolation times \bar{t}_i and the number $K+1$ of points along a discrete curve, we shall in the following and without explicit mention always interpret $K \rightarrow \infty$ as a sequence of natural numbers approaching infinity, such that $K\bar{t}_i \in \mathbb{N}_0$ for $i = 1, \dots, I$.

We are now in position to state the main result of this section, which is concerned with the Mosco convergence (cf. Definition 2.4.1) in the restricted Gaussian case:

Theorem 3.4.4 (Mosco convergence and convergence of discrete minimizers). *Let $\hat{\mathcal{E}}^K[\sigma]$ be given by $\hat{\mathbf{E}}[\sigma^K]$ if $\sigma = \eta_{\sigma^K}$, and ∞ otherwise. Similarly, define $\hat{\mathcal{F}}^K[\sigma]$ be given by $\hat{\mathbf{F}}[\sigma^K]$ if $\sigma = \eta_{\sigma^K}$, and ∞ otherwise. Then, set $\hat{\mathcal{F}}^{\delta, K} := \hat{\mathcal{F}} + \delta \hat{\mathcal{E}}$. With respect to the weak topology in $H^2((0, 1); \mathbb{R}_{+, \text{dia}}^{d \times d})$ we have $\lim_{K \rightarrow \infty} \hat{\mathcal{F}}^{\delta, K} + \mathcal{I}^K = \hat{\mathcal{F}}^\delta + \mathcal{I}$ in the sense of Mosco, for $\delta > 0$. Moreover, any sequence $(\sigma^K)_K$ with $\hat{\mathcal{F}}^{\delta, K}[\sigma^K] + \mathcal{I}^K[\sigma^K]$ uniformly bounded contains a subsequence that converges weakly in $H^2((0, 1); \mathbb{R}_{+, \text{dia}}^{d \times d})$. As a consequence, any sequence of minimizers of $\hat{\mathcal{F}}^{\delta, K} + \mathcal{I}^K$ contains a subsequence converging weakly to a minimizer of $\hat{\mathcal{F}}^\delta + \mathcal{I}$.*

Proof. We have to show the weak *liminf* and the strong *limsup* inequalities defining Mosco convergence. Concerning the weak *liminf*-inequality, we need to show that for every sequence $(\sigma^K)_{K \in \mathbb{N}} \subset H^2((0, 1); \mathbb{R}_{+, \text{dia}}^{d \times d})$, such that $\sigma^K \rightharpoonup \sigma$, it holds that $\liminf_{K \rightarrow \infty} \hat{\mathcal{F}}^{\delta, K}[\sigma^K] + \mathcal{I}^K[\sigma^K] \geq \hat{\mathcal{F}}[\sigma] + \mathcal{I}[\sigma]$. Let $\sigma^K \rightharpoonup \sigma$ in $H^2((0, 1); \mathbb{R}_{+, \text{dia}}^{d \times d})$. Upon taking a subsequence, we may replace the *liminf* by an actual *lim* and may assume without loss of generality $\hat{\mathcal{F}}^{\delta, K}[\sigma^K] + \mathcal{I}^K[\sigma^K] \leq C$ for some constant $C < \infty$. Thus, we have $\sigma^K = \eta_{(\sigma_0^K, \dots, \sigma_K^K)}$ for some $(\sigma_0^K, \dots, \sigma_K^K) \in (\mathbb{R}_{+, \text{dia}}^{d \times d})^{K+1}$ and this estimate implies

$$d_K := \max_{k \in \{1, \dots, K\}} |\sigma_k^K - \sigma_{k-1}^K| \leq \sqrt{\sum_{k=1}^K |\sigma_k^K - \sigma_{k-1}^K|^2} = \sqrt{\hat{\mathcal{E}}^K[\sigma^K]/K} \leq \sqrt{\frac{C}{\delta K}},$$

which converges to zero as $K \rightarrow \infty$. Next, we show that $\mathcal{I}[\sigma] = 0$. It is straightforward to see that $\eta_{\sigma^K}(t)$ is in the convex hull of σ_{k-1} , σ_k , and σ_{k+1} for $t \in [t_{k-1/2}^K, t_{k+1/2}^K]$. Thus, the evaluation constraint is satisfied in the limit. To conclude, by the weak lower semi-continuity of $\hat{\mathcal{F}}^\delta$ due to Lemma 3.3.13, and by Lemmas 3.4.2 and 3.4.3 we obtain

$$\begin{aligned} \hat{\mathcal{F}}^\delta[\sigma] + \mathcal{I}[\sigma] &= \hat{\mathcal{F}}^\delta[\sigma] \leq \liminf_{K \rightarrow \infty} \hat{\mathcal{F}}^\delta[\sigma^K] \leq \liminf_{K \rightarrow \infty} \hat{\mathcal{F}}^{\delta, K}[\sigma^K] + \delta \frac{C}{K} |\sigma^K|_{H^1} |\sigma^K|_{H^2} \\ &\leq \liminf_{K \rightarrow \infty} \left(\hat{\mathcal{F}}^{\delta, K}[\sigma^K] + \mathcal{I}^K[\sigma^K] \right), \end{aligned}$$

where we used the uniform boundedness of $|\sigma^K|_{H^1}$ and $|\sigma^K|_{H^2}$ due to the weak convergence of σ^K .

Concerning the strong *limsup* inequality, we need to show that for every $\sigma \in H^2((0, 1); \mathbb{R}_{+, \text{dia}}^{d \times d})$ there is a sequence $(\sigma^K)_{K \in \mathbb{N}} \subset H^2((0, 1); \mathbb{R}_{+, \text{dia}}^{d \times d})$ with $\sigma^K \rightarrow \sigma$, such that $\hat{\mathcal{F}}^\delta[\sigma] + \mathcal{I}[\sigma] \geq \limsup_{K \rightarrow \infty} \hat{\mathcal{F}}^{\delta, K}[\sigma^K] + \mathcal{I}^K[\sigma^K]$. Let $\sigma \in C^3([0, 1], \mathbb{R}_{+, \text{dia}}^{d \times d})$ with finite energy (in particular, $\mathcal{I}(\sigma) = 0$), and choose $\sigma^K = \eta_{\sigma^K}$ as the recovery sequence. By definition, we have $\mathcal{I}^K[\sigma^K] = 0$. As $K \rightarrow \infty$, we have $d_K := \max_{\{1, \dots, K\}} |\sigma(t_k^K) - \sigma(t_{k-1}^K)| \rightarrow 0$, as well as $\sigma^K \rightarrow \sigma$ strongly in H^2 by Lemma 3.4.1. Thus, by the strong H^2 -continuity of $\hat{\mathcal{F}}^\delta$ from Lemma 3.3.13 and by Lemma 3.4.3 we have

$$\hat{\mathcal{F}}^\delta[\sigma] + \mathcal{I}[\sigma] = \hat{\mathcal{F}}^\delta[\sigma] = \lim_{K \rightarrow \infty} \hat{\mathcal{F}}^\delta[\eta_{\sigma^K}]$$

$$\geq \limsup_{K \rightarrow \infty} \hat{\mathcal{F}}^{\delta,K}[\sigma^K] - \delta C K^{-1} |\sigma^K|_{H^1} |\sigma^K|_{H^2} = \limsup_{K \rightarrow \infty} \left(\hat{\mathcal{F}}^{\delta,K}[\eta_\sigma^K] + \mathcal{I}^K[\sigma^K] \right),$$

where we again used the uniform boundedness of $|\sigma^K|_{H^1}$ and $|\sigma^K|_{H^2}$, now due to the strong convergence of σ^K . Thus, we obtain $\limsup \hat{\mathcal{F}}^{\delta,K} + \mathcal{I}^K \leq \hat{\mathcal{F}}^\delta + \mathcal{I}$ on $C^3([0, 1], \mathbb{R}_{+, \text{dia}}^{d \times d})$. By a density argument (C^3 functions fulfilling interpolation constraints are dense in the space of H^2 functions satisfying the same interpolation constraints, cf. [HRW18, Lemma 4.6]) and the strong H^2 continuity of $\hat{\mathcal{F}}^\delta$, we obtain

$$\hat{\mathcal{F}}^\delta[\sigma] + \mathcal{I}[\sigma] \geq \limsup_{K \rightarrow \infty} \hat{\mathcal{F}}^{\delta,K}[\sigma^K] + \mathcal{I}^K[\sigma^K].$$

To show the convergence of discrete minimizers it remains to establish equicoercivity. Here we will follow the same strategy as in the proof of [HRW18, Theorem 4.9]. Let $(\sigma^K)_K$ be a sequence with $\hat{\mathcal{F}}^{\sigma,K}[\sigma^K] + \mathcal{I}^K[\sigma^K]$ uniformly bounded. As before, we can assume without loss of generality that $\sigma^K = \eta_{(\sigma_0^K, \dots, \sigma_K^K)}$ for some $(\sigma_0^K, \dots, \sigma_K^K) \in (\mathbb{R}_{+, \text{dia}}^{d \times d})^{K+1}$. Following the proof of Lemma 3.4.2 and also recalling the estimate $K \sum_{k=1}^{K-1} |\sigma_k^K - \sigma_{k-1}^K|^2 \leq C^* |\sigma|_{H^1}^2$ we obtain a uniform bound of the H^1 -seminorm $|\sigma^K|_{H^1}$. By Poincaré's inequality, one even obtains uniform boundedness of the norm $\|\sigma^K\|_{H^1}$. It remains to show the uniform boundedness of the H^2 -seminorm $|\sigma^K|_{H^2}$. Indeed, we obtain the estimate

$$\begin{aligned} |\sigma^K|_{H^2}^2 &= |\eta_{(\sigma_0^K, \dots, \sigma_K^K)}|_{H^2}^2 = 4K^3 \sum_{k=1}^{K-1} \left| \sigma_k^K - \frac{\sigma_{k+1}^K + \sigma_{k-1}^K}{2} \right|^2 = 4K^3 \sum_{k=1}^{K-1} \mathcal{W}^2[\sigma_k^K, \text{Bar}(\sigma_{k+1}^K, \sigma_{k-1}^K)] \\ &= \hat{\mathbf{F}}^K[\sigma_0^K, \dots, \sigma_K^K] \leq \hat{\mathcal{F}}^{\delta,K}[\sigma^K]. \end{aligned}$$

The statement about the convergence of minimizers is now a standard consequence of the Mosco convergence from the previous theorem, cf. [Bra14]. \square

3.5 Fully discrete Wasserstein splines and numerical results

3.5.1 Algorithmic foundations

To implement Wasserstein splines numerically, we have to further discretize the time-discrete spline energy in space. With the application to images in mind, we consider $\Omega := [0, 1]^2$ ($d = 2$) and an 'Eulerian' discretization of probability measures: Let $\mu \in \mathcal{P}(\Omega)$. Here, the image intensities on each colour channel are encoded as Lebesgue densities. We can obtain the discretized version of a density μ by first defining the computational mesh

$$\Omega_{MN} = \left\{ \frac{0}{M-1}, \frac{1}{M-1}, \dots, \frac{M-1}{M-1} \right\} \times \left\{ \frac{0}{N-1}, \frac{1}{N-1}, \dots, \frac{N-1}{N-1} \right\} \quad \text{for } M, N \geq 3.$$

Next, we integrate the mass of μ on each cell $\Omega^{kl} := [\frac{k}{M-1}, \frac{k+1}{M-1}] \times [\frac{l}{N-1}, \frac{l+1}{N-1}]$, and define the weight

$$\omega_{kl} = \int_{\Omega^{kl}} d\mu,$$

obtaining the spatially discrete measure $\mu^D[\omega] := \sum_{k,l} \omega_{kl} \delta_{kl}$, where $\omega := (\omega_{kl})_{k,l} \in \Sigma_{MN} := \{\omega \in \mathbb{R}_+^{MN} : \sum_{m,n} \omega_{mn} = 1\}$, and δ_{kl} is defined as the delta distribution located at the center of the cell (pixel) Ω^{kl} . Alternatively, for other applications the 'Lagrangian' discretization might be more useful:

In this case, one considers $\Omega = \mathbb{R}^d$, and independently samples a measure $\mu \in \mathcal{P}_2(\Omega)$ a total of L times. One then defines the spatially discrete measure $\mu^D[x] := \frac{1}{L} \sum_{l=1}^L \delta_{x_l}$, where $x := (x_l)_l \in \mathbb{R}^{dL}$. Since obtaining the exact Wasserstein distance between two discrete measures with $L = MN$ atoms comes with a cost of $\mathcal{O}(M^3 N^3)$, we approximate the Wasserstein distance between two discrete measures μ and ν by the entropy-regularized Wasserstein distance \mathbf{W}_ϵ introduced originally in [Cut13], with regularization parameter $\epsilon > 0$. The loss \mathbf{W}_ϵ can be very efficiently computed in an auto-differentiable manner, i.e. the gradients of $\mathbf{W}_\epsilon(\cdot, \cdot)$ with respect to both the weights ω and locations x are obtained as a by-product of the evaluation of this function (cf. [CD14]) with state-of-the-art implementations of the Sinkhorn algorithm, such as in [Sch19] and [CFG⁺21]. Correspondingly, we take into account the entropy-regularized approximation $\mathbf{Bar}_{(\cdot)}^\epsilon$ (cf. [BCC⁺15]) of the (generalized) barycenter, which, once again, can be efficiently computed in an auto-differentiable fashion. The entropy-relaxed, regularized spline objective functional will then look as follows:

$$\begin{aligned} \mathbf{F}_{(G)}^{\delta, K, \epsilon}(\mu^D[\theta^0, \dots, \theta^K]) &:= 4K^3 \sum_{k=1}^{K-1} \mathbf{W}_\epsilon^2(\mu_k^D[\theta^k], \mathbf{Bar}_{(\mu_k^D[\theta^k])}^\epsilon(\mu_{k+1}^D[\theta^{k+1}], \mu_{k-1}^D[\theta^{k-1}])) \\ &\quad + \delta K \sum_{k=0}^{K-1} \mathbf{W}_\epsilon^2(\mu_k^D[\theta^k], \mu_{k+1}^D[\theta^{k+1}]), \end{aligned} \quad (3.5.1)$$

where we omitted the K super-index. For a fixed computational domain Ω_{MN} , temporal resolution K , entropy-regularization $\epsilon > 0$, regularizer $\delta > 0$, interpolation constraints and a chosen boundary condition, our aim in the first discretization variant is to minimize the previous functional with respect to the weights $\theta^k = \omega^k \in \Sigma_{MN}$ of the discrete probability measures $\mu_k^D = \sum_{y \in \Omega_{MN}} \omega_y^k \delta_y$ for indices k that are not fixed by the interpolation conditions. The explicit minimization of functional (3.5.1) as a function of weights is performed by Algorithm 4. For the second variant, one can instead fix a number of samples/locations L , and straightforwardly minimize the above functional over the positions $\theta^k = x^k \in \mathbb{R}^{dL}$ of atoms of the discrete probability measures $\mu_k^D := \frac{1}{L} \sum_{l=1}^L \delta_{x_l^k}$. This can be implemented completely analogously to Algorithm 4.

3.5.2 Numerical results

In what follows, we investigate and discuss qualitative properties of the spline interpolation in the space of probability distributions, being aware that the superior temporal smoothness of this interpolation is difficult to show with series of still images.

Figure 3.1 shows discrete piecewise geodesic and discrete spline interpolations of three two dimensional Gaussian distributions $\bar{\mu}_i = \mathcal{N}(m_i, \sigma_i^2)$ for $i = 1, 2, 3$ at the prescribed times $\bar{t}_1 = 0, \bar{t}_2 = \frac{1}{2}, \bar{t}_3 = 1$, where the interpolation is computed as a minimizer of \mathbf{F}^K over all Gaussian parameters $(m_k, \sigma_k) \in \mathbb{R}^d \times \mathbb{R}_{+, \text{sym}}^{d \times d}$ for $k = 0, \dots, K$. For the discrete splines the center of masses of spline interpolation correspond almost perfectly to the cubic spline interpolation of the center of masses of the key frames. The third row shows the spline interpolation for the same key frames. This time, we instead optimize functional (3.5.1) over all weights $\theta^k = \omega^k \in \Sigma_{MN}$ for $k = 0, \dots, K$ and $M = N = 128$. In particular, the solutions need not to be Gaussian distributions. The fourth row shows the difference between the first and second row, i.e. between the piecewise geodesic and spline interpolations.

The next example in Figure 3.2 investigates the interpolation of three key frames with constant density on an annulus for the first and constant density on a disk for the second and third (at times $\bar{t}_1 = 0, \bar{t}_2 = \frac{1}{2}, \bar{t}_3 = 1$). In case of the piecewise geodesic interpolation one observes a decreasing density on the

Algorithm 4: Algorithm for minimizing $\mathbf{F}_{(G)}^{\delta,K,\epsilon}$ as a function of weights $\omega^0, \dots, \omega^K \in \Sigma_{MN}$. The product \odot and the exponential function e act component-wise on vectors, and $\mathbf{1}_{MN} = (1, \dots, 1) \in \mathbb{R}^{MN}$.

```

1  $t = 0$ ;
2 for  $k = 0$  to  $K$  and  $k$  not fixed do
3    $\tilde{\omega}^k = \hat{\omega}^k = \mathbf{1}_{MN}/MN$ ;
4 while not converged do
5    $\beta = (t + 1)/2$ ;
6   for  $k = 0$  to  $K$  and  $k$  not fixed do
7     /* update weights (Nesterov's accelerated gradient update) */
7      $\omega^k = (1 - \beta^{-1})\hat{\omega}^k + \beta^{-1}\tilde{\omega}^k$ ;
7     /* compute gradient (Sinkhorn algorithm) */
8      $\text{grad}_k = \nabla_{\omega^k} \mathbf{F}_{(G)}^{\delta,K,\epsilon}(\boldsymbol{\mu}^D[\omega^0, \dots, \omega^K])$ ;
8     /* update weights */
9      $\tilde{\omega}^k = \tilde{\omega}^k \odot e^{-t\beta \text{grad}_k}$ ;
10     $\tilde{\omega}^k = \tilde{\omega}^k / (\tilde{\omega}^k)^T \mathbf{1}_{MN}$ ;
11     $\hat{\omega}^k = (1 - \beta^{-1})\hat{\omega}^k + \beta^{-1}\tilde{\omega}^k$ ;
12   $t = t + 1$ ;
```

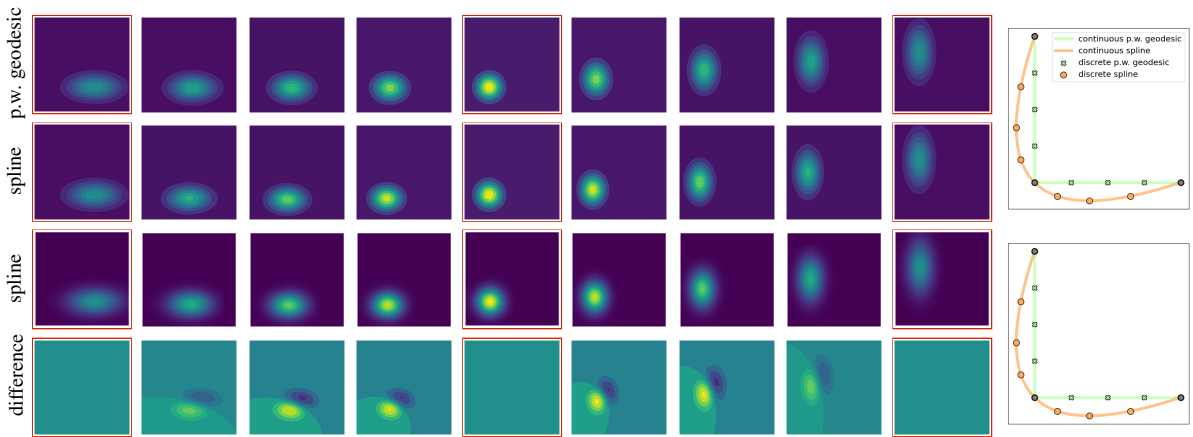


Figure 3.1: First two rows: A comparison of discrete piecewise geodesic interpolation (first row) and discrete spline interpolation (second row) for $\delta = 0$ is shown for key frame distributions framed in red. The optimization was done on $\mathcal{P}_2^{G,d}$. Third row: Same as second row, except the optimization was performed in the full space $\mathcal{P}_2(\mathbb{R}^d)$. Bottom row: Difference between spline and piece-wise geodesic interpolations. Top right: Plot of the center of masses as a polygonal curve in \mathbb{R}^2 . Bottom right: Plot of the standard deviations as a polygonal curve in \mathbb{R}^2 .

closing annulus in between the first two key frames and obviously constant interpolation in between the second and third key frames. In case of the spline interpolation ($\delta = 0$) the annulus also closes between the first and second key frames but shows a strong overshooting at the center between the second and third key frames.

In Figure 3.3 a thin annulus shaped distribution and two times an equal square shaped frame are taken into account as key frame distributions (at times $\bar{t}_1 = 0, \bar{t}_2 = \frac{1}{2}, \bar{t}_3 = 1$). Different from Figure 3 in [JRR23] in the case of spline interpolations in the metamorphosis model, one does not observe strongly inward pointing edges between the equal square shaped frames. Instead strong overshooting effects are visible at the corners of the squares.

In Figure 3.4 the key frames consist of pairs of Gaussians with constant mass and constant variance, which are far apart for the first and fourth key frame and close by for the second and third key frame (at times $\bar{t}_1 = 0, \bar{t}_2 = \frac{1}{3}, \bar{t}_3 = \frac{2}{3}, \bar{t}_4 = 1$). Piecewise geodesic interpolation leads to piecewise linear trajectories of the center of masses, whereas trajectories are curved in the spline case with a merger of the two bumps in between the second and third key frame.

In Figure 3.5, three key frames represent a single Gaussian, a pair of vertically displaced Gaussian of half the mass, and the vertically displaced configuration rotated by $-\frac{\pi}{4}$ (at times $\bar{t}_1 = 0, \bar{t}_2 = \frac{1}{2}, \bar{t}_3 = 1$). The piecewise geodesic shows the splitting of mass and approximately straight line trajectories between the pairs of key frames. For the spline interpolation one observes an overshooting with a positive rotation angle in between the first and the second key frame.

In Figure 3.6, we leverage the results of equation (3.2.9). An implementation of the decoupling leads to a significant decrease of the computing time and the number of iterations with no apparent loss of detail.

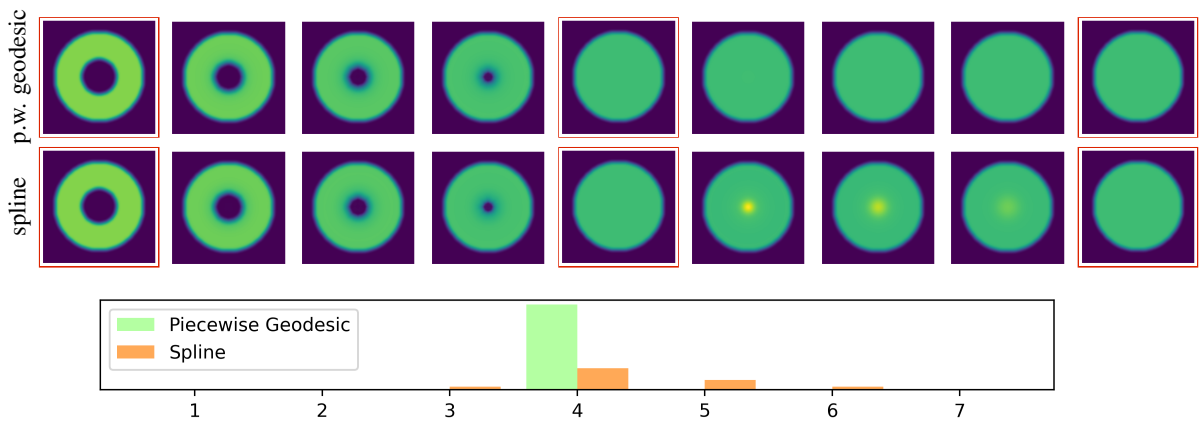


Figure 3.2: Discrete piecewise geodesic interpolation (top) and discrete spline interpolation for $\delta = 0$ (middle) of three key frames with constant density on an annulus for the first and constant density on a disk for the second and third (framed in red). Bottom: Contribution of each time-step $k = 1, \dots, K - 1$ to the spline energy, i.e. $\mathcal{W}^2(\mu_k, \text{Bar}(\mu_{k-1}, \mu_{k+1}))$ for the spline interpolation (orange) and piecewise geodesic interpolation (green).

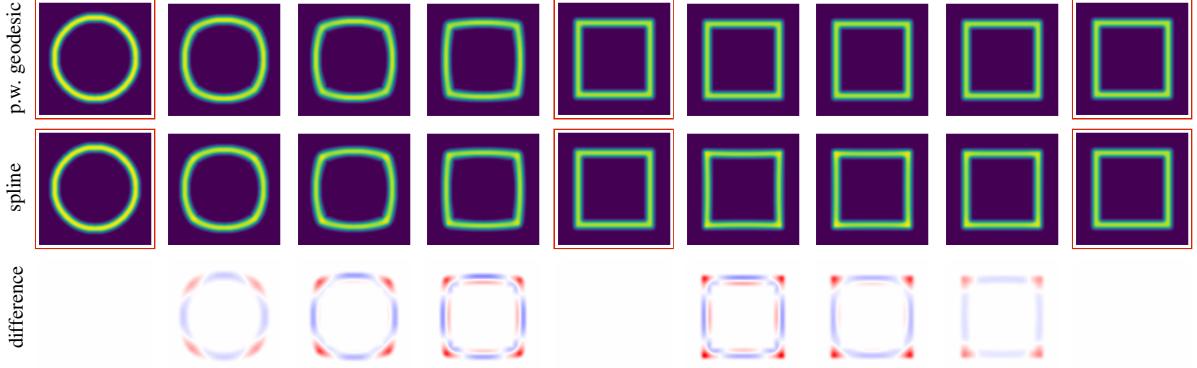


Figure 3.3: Piecewise geodesic (top) and spline interpolation (middle) are shown for key frames (framed in red) consisting of a thin annulus-shaped distribution and two equal thin square-shaped distributions, using the color map $0 \text{ } \text{---} \text{ } 4\text{e-}4$. Bottom: Difference between the spline and piecewise geodesic interpolations, using the color map $-6\text{e-}5 \text{ } \text{---} \text{ } 6\text{e-}5$.

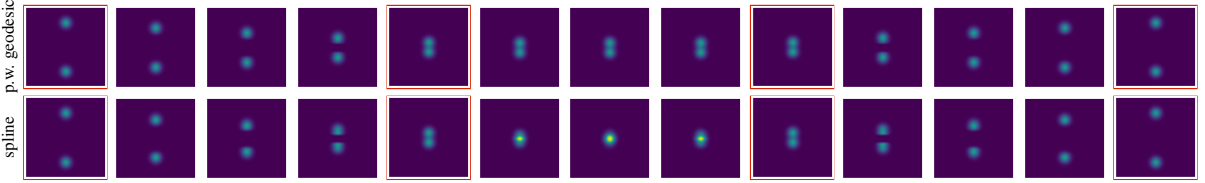


Figure 3.4: The key frames represent two Gaussians that are far apart from each other (first and fourth key frames) and close to each other (second and third key frames). Piecewise geodesic (top) and spline (bottom) interpolations are shown.

3.6 Generative texture synthesis based on Wasserstein spline interpolation of feature distributions

The flexibility of our model will be tested in this section to generate spline interpolations in the space of textures. Recently, Houdard *et al.* [HLPR20] proposed GOTEX, a generative model for texture synthesis from a single sample image. There, the parameters of the generator are chosen such that the distribution of features extracted from the generated textures is close in Wasserstein distance to the corresponding empirical feature distribution for the given sample image. In what follows, we shall outline how we leverage our spline interpolation model within the GOTEX framework. To this end, we proceed as follows:

- First, for a vector of feature maps, we compute empirical feature distributions $\bar{\nu}_i \in \mathcal{P}(\mathbb{R}^d)$ for all input images \bar{u}_i at times t_i for $i = 1, \dots, I$ and some $d \in \mathbb{N}$.
- Next, we use the discrete spline approach presented in the preceeding sections to compute a discrete spline interpolation $(\nu_k^K)_{k=0, \dots, K}$ for prescribed distributions $\bar{\nu}_i$ at times $k_i = \bar{t}_i K$ for $k_i \in \{0, \dots, K\}$.
- Finally, in a post processing, we train a generative texture model to obtain image representations $u_k = g_{\theta_k}(z)$ for the computed probability distributions ν_k^K , where θ_k is a set of optimal parameters

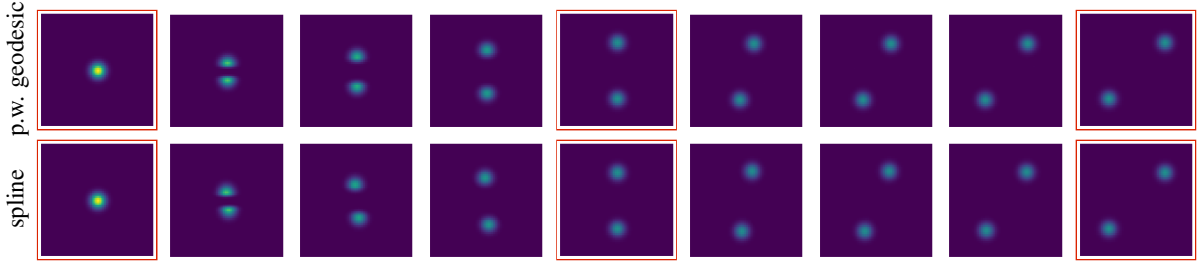


Figure 3.5: From left to right the key frames represent a single Gaussian, a pair of vertically displaced Gaussian of half the mass, and the vertically displaced configuration rotated by $-\frac{\pi}{4}$. Discrete piecewise geodesic (top) and spline (bottom) interpolations are shown.

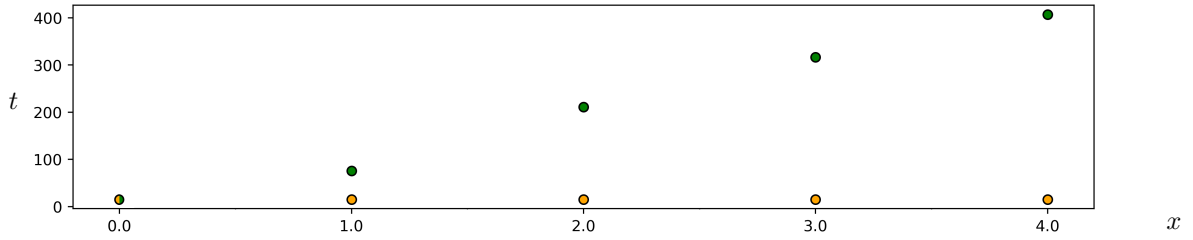


Figure 3.6: Time t (in seconds) until convergence of the fully discrete spline interpolation problem is reached, for a series of five interpolation problems $P(x)$ depending on parameter $x \in \{0, 1, 2, 3, 4\}$ (x -axis). The interpolation problem $P(x)$ is defined as follows: The prescribed times are $\bar{t}_0 = 0$, $\bar{t}_1 = 0.5$ and $\bar{t}_2 = 1$, and the prescribed probability measures are given by $\bar{\mu}_0 = \mathcal{N}((0, 0), \sigma_0^2)$, $\bar{\mu}_1 = \mathcal{N}((x, x), \sigma_1^2)$, and $\bar{\mu}_2 = \mathcal{N}((0, 0), \sigma_2^2)$, for diagonal standard deviation matrices $\sigma_0 = \text{diag}(1, 2)$, $\sigma_1 = \text{diag}(1, 1)$ and $\sigma_2 = \text{diag}(2, 1)$. Dots denote the computation time of the algorithm solving problem $P(x)$, both with implementation of the decoupling of the means as described by equation (3.2.9) (orange), and without it (green).

of a generative neural network g_{θ_k} applied to a sample z of a regular distribution.

In explicit, to generate a texture spline interpolation, we follow these steps cf. Fig. 3.7.

1. Extracting empirical feature distributions. Let $I \geq 2$ be fixed, and consider a vector $F = (F_m)_{m=1, \dots, M}$ of d -dimensional local feature maps $F_m : \mathbb{R}^N \rightarrow \mathbb{R}^d$ defined on images with N pixels. Each component F_m operates on small pixel neighbourhoods (patches). Given the images \bar{u}_i , $i = 1, \dots, I$ to be spline interpolated the associated empirical feature distributions are

$$\bar{\nu}_i := \frac{1}{M} \sum_{m=1, \dots, M} \delta_{F_m[\bar{u}_i]}$$

in $\mathcal{P}(\mathbb{R}^d)$ for $i = 1, \dots, I$, with δ_x being the Dirac measure at x in \mathbb{R}^d .

2. Computing discrete splines in the space of feature distributions. Given the set of feature distributions $\bar{\nu}_i$ with $i = 1, \dots, I$, obtained from the first step with associated interpolation times $0 \leq \bar{t}_0 < \dots, \bar{t}_I \leq 1$ and some $K \in \mathbb{N}$, we compute a discrete spline interpolation $(\nu_k^K)_{k=0, \dots, K}$

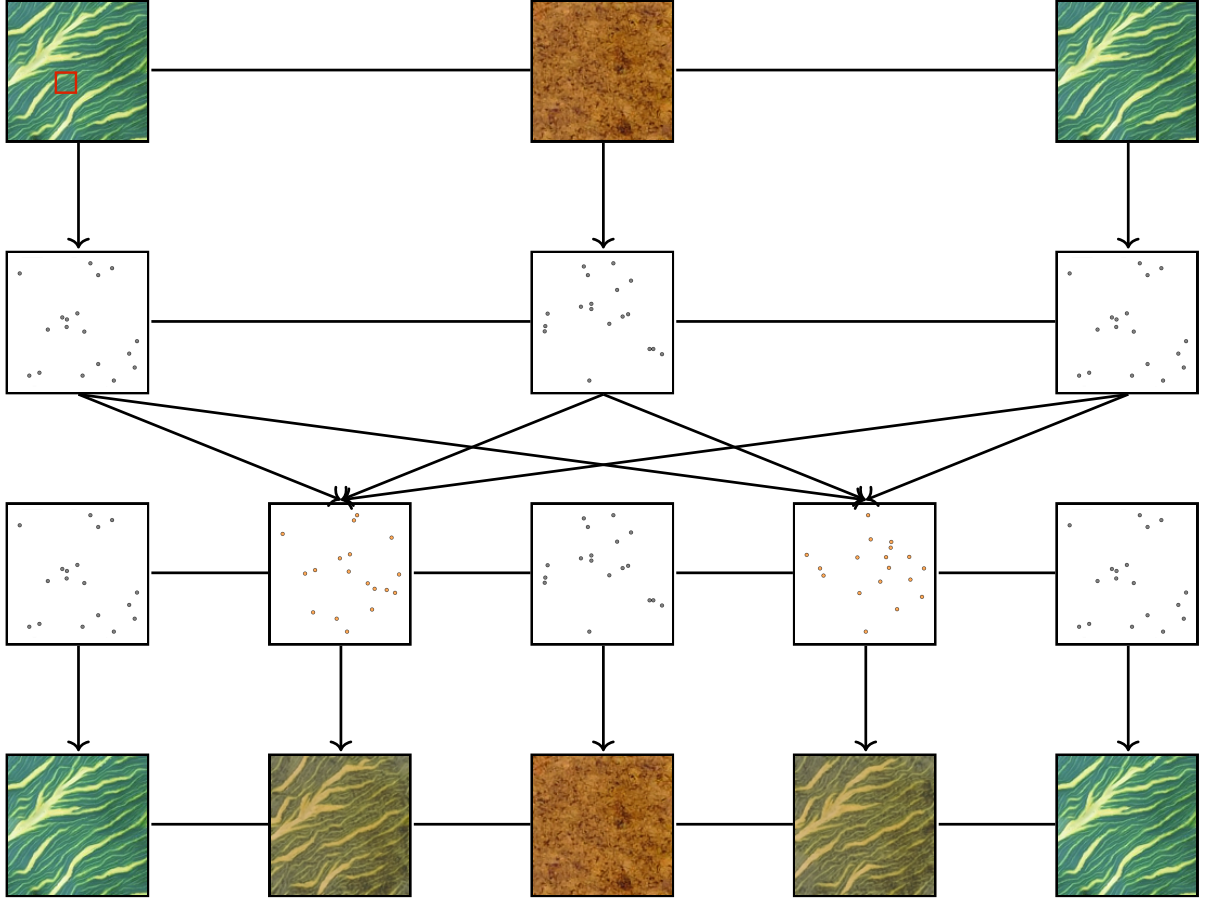


Figure 3.7: Schematic plot of our algorithm to generate texture interpolations based on Wasserstein spline interpolation of feature distributions. The step from the first row to the second one corresponds to the extraction of empirical feature distributions. The step from the second row to the third describes the use of our interpolation algorithm to compute feature distributions in-between. The final step (from the third to the fourth row) represents mapping back features into textures.

of the prescribed feature distributions $\bar{\nu}_i$ at times \bar{t}_i in $\mathcal{P}_2(\mathbb{R}^d)$ as described in the previous sections. Here, we constrain the discrete spline to lie in the space of feature distributions, i.e. each distribution ν_k^K must be represented as the sum of M delta distributions in \mathbb{R}^d with equal weights. To this end, we minimize the fully discrete spline energy functional (3.5.1) with respect to the locations x_m^k , where

$$\nu_k^K = \frac{1}{M} \sum_{m=1}^M \delta_{x_m^k},$$

and keep the distributions $\bar{\nu}_i$ at times $k_i = \bar{t}_i K$ for $k_i \in \{0, \dots, K\}$ fixed.

Retrieving image representations via a generative texture model based on neural networks. Different samples of a synthesized texture for a given feature distribution ν_k^K (obtained in the previous step) are regarded as samples of a probability distribution, which is de-

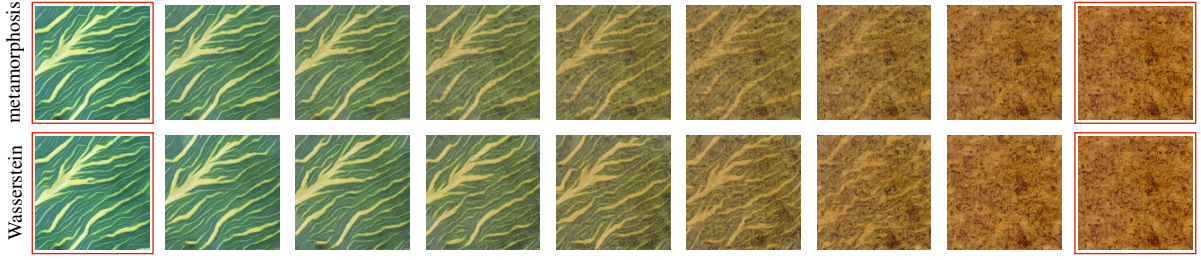


Figure 3.8: Time discrete metamorphosis (top) and Wasserstein spline interpolations (bottom) with framed prescribed images/feature distributions for $K = 16$. For the metamorphosis spline the key frames are chosen identical to the synthesized texture of the Wasserstein splines. Due to symmetry, only the first half of the interpolations is shown.

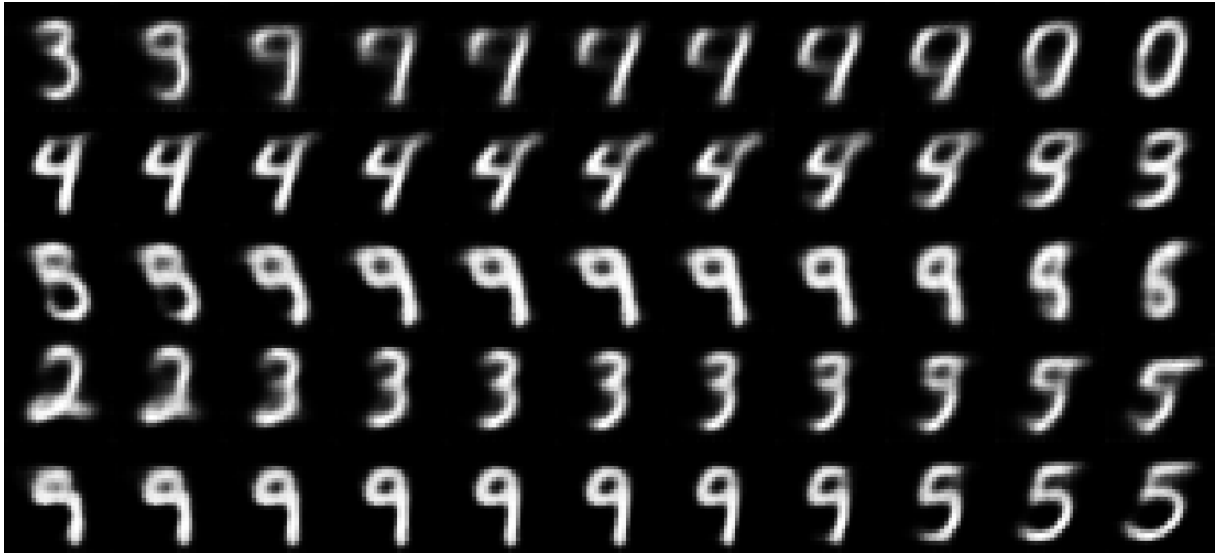


Figure 3.9: Five spline interpolations in the MNIST-Dataset. Each of them has 11 frames; the fixed keyframes are at $k = 0$, $k = 5$, and $k = 10$ for each row.

finned as the push-forward of a fixed distribution ζ defined on a latent space \mathcal{Z} , with a generator $g_\theta : \mathcal{Z} \rightarrow \mathbb{R}^N$ with parameter vector θ in a set of admissible parameters Θ . Typically, one may assume ζ to be the uniform distribution on the space $\mathcal{Z} = [0, 1]^N$ and g to be a feed-forward neural network with a fixed architecture. Hence, one is looking for an optimal parameter vectors $\theta_k \in \Theta$, which minimizes the Wasserstein distance $\mathcal{W}(\mu_{\theta_k}, \nu_k^K)$ of the resulting feature distribution

$$\mu_{\theta_k} := \frac{1}{M} \sum_{m=1, \dots, M} (F_m \circ g_{\theta_k}) \# \zeta$$

from the given feature distribution ν_k^K of the discrete spline.

The minimization of $\mathcal{W}(\mu_{\theta_k}, \nu_k^K)$ with respect to the parameter vector θ_k can be numerically realized via a stochastic gradient descent approach. After obtaining the optimal parameter vectors θ_k for $k = 0, \dots, K$, one then samples $z \sim \zeta$, and generates the resulting texture spline interpolation

$(u_k)_{k=0,\dots,K}$ as a set of images

$$u_k := g_{\theta_k}(z)$$

for $k = 0, \dots, K$. For different samples z one obtains different images $(u_k)_{k=0,\dots,K}$ representing the spline interpolation of the textures.

The multi-scale architecture of the generator network g is made up of chains of convolutional, non-linear activation and upsampling layers that take a noise sample z as an input and terminate by producing the final image, cf. [ULVL16]. Each convolution block in the generator network contains three convolutional layers followed by a non-linear ReLU activation layer. The convolutional layers contain 3×3 , 3×3 and 1×1 filters, respectively. Next, nearest-neighbour interpolation is used in the upsampling layers to obtain a tensor with the desired full resolution. For the last step, this tensor is mapped to an RGB-image by a batch of 1×1 filters.

Numerical results. In Figure 3.8 we compare our spline interpolation method with the alternative metamorphosis spline interpolation (middle) as described in [JRR23] with the same prescribed feature distributions/frames in each case. The key frames consist of a close-up of a leaf (at times $\bar{t}_1 = 0$, $\bar{t}_3 = 1$), and a close-up of a cork (at $t_2 = \frac{1}{2}$). Clearly, the leaf creases inherited from the first and last key frames are simply blended out in the first method. On the other hand, our approach ensures that the features are interpolated smoothly: Both the boundaries and the surface area covered by the creases change smoothly over time.

Fig. 3.10 shows discrete texture curves resulting from a discrete spline interpolation $(\nu_k^K)_k$ for $K = 20$ in the space of feature distributions as described above. The texture samples which have been synthesized from the prescribed feature distributions have been framed in red, and the prescribed times are given by $\bar{t}_i = i/4$, $i = 0, \dots, 4$. Two different texture realizations $u_k^1 := g_{\theta_k}(z_1)$, $u_k^2 := g_{\theta_k}(z_2)$ for $z_1, z_2 \sim \zeta$ are shown on top of each other for a normalized random distribution ζ and θ_k minimizing the entropy regularized Wasserstein distance between μ_{θ_k} and ν_k^K . The weights of the neural network used to generate these textures are kept unchanged between both interpolations (and along each interpolation between different time frames). Hence, even though for a fixed time step k the spatial arrangement of the texture pattern varies substantially between both samples, it becomes apparent that the texture characteristics described by the distribution of features ν_k^K coincide and vary smoothly along the curve.

Fig. 3.11 serves as a benchmark for our spline interpolation model on how well it can predict the texture patterns in comparison to the ground truth. Therein, the prescribed feature distributions are extracted from equally spaced still frames of a video showing the life cycle of a patch of mango skin. As the mango peel goes from green to ripe and eventually rots away, not only the colors but the texture of the peel changes significantly. In the frame of generative texture synthesis our method (top rows on each panel) matches both structure and coloring of the actual textures (bottom row on each panel) at corresponding times.

3.7 Generative spline interpolation of data aided by variational autoencoders

A variational autoencoder (VAE) is a type of neural network designed for generating new data similar to a given dataset (e.g., images or text) and learning meaningful data representations. Here we give a short breakdown of how a VAE works:

- 3• **Encoder:** The encoder ϕ maps input data x (like an image) to a lower-dimensional space called the latent space. In contrast to traditional autoencoders, encoders of VAEs do not map the input to a single point in the latent space. Instead, the encoder predicts a probability distribution (a Gaussian distribution with a mean and diagonal covariance matrix) in this space. This particular step makes VAEs specially suitable to combine with the spline interpolation method we have introduced in previous sections.
- **Latent Space Sampling:** A sample is drawn according to the predicted distribution. This sampling step adds randomness, making the VAE capable of generating diverse outputs.
- **Decoder:** The decoder ψ takes the sample from the latent space and reconstructs the input data, aiming to make it as close to the original input as possible.

The VAE is trained with two goals in mind: Minimize the reconstruction error (i.e. how well the output matches the input). Ensure the latent space distributions are close to a standard normal distribution by minimizing the KL divergence between them. After the learning process on a given dataset, VAEs can generate new data by sampling from the latent space and decoding it, while also providing structured representations of the input data.

Now we give a brief description on how to leverage our proposed interpolation method with VAEs to greatly increase its reach. Consider a VAE with encoder $\phi : \mathbb{R}^N \rightarrow \mathcal{P}_2^{G,d}$ and decoder $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^N$ that have been pre-trained on a dataset for which we want to generate similar data. For a given temporal resolution $K \in \mathbb{N}$ and prescribed data $\bar{x}_1, \dots, \bar{x}_I \in \mathbb{R}^N$ at prescribed times $\bar{t}_1, \dots, \bar{t}_I$, we follow the following steps:

1. Run the prescribed data through the encoder, i.e. compute the diagonal Gaussian measures $\bar{\mu}_i := \phi(\bar{x}_i) \in \mathcal{P}_2^{G,d}$ for all $i = 1, \dots, I$. This is the analogue of step 1 in the texture interpolation algorithm (cf. 3.6).
2. Next, apply our interpolation method (the version restricted to Gaussian measures with diagonal covariance matrices) to the prescribed measures $\bar{\mu}_1, \dots, \bar{\mu}_I \in \mathbb{R}^N$ at prescribed times $\bar{t}_1, \dots, \bar{t}_I$. This generates a discrete-time spline interpolation $(\mu_k)_{k=0, \dots, K} \subset (\mathcal{P}_2^{G,d})^{K+1}$. This is equivalent to step 2 in the texture interpolation algorithm.
3. Subsequently, use Algorithm 3 to sample a spline trajectory $y : [0, 1] \rightarrow \mathbb{R}^d$ in the latent space \mathbb{R}^d .
4. Finally, evaluate the sampled spline trajectory at times $t_k^K := \frac{k}{K}$ for $k = 0, \dots, K$ and apply the decoder ψ on these evaluations to translate them back to the input space \mathbb{R}^N . In explicit, compute $x_k := \psi(y(t_k^K)) \in \mathbb{R}^N$ for all $k = 0, \dots, K$. The resulting tuple (x_0, \dots, x_K) is the spline interpolation of the data.

To check the feasibility and robustness of our approach, in Fig. 3.9 we plotted a few sample splines on the MNIST dataset

3.8 Variational-time Wasserstein regression

This section focuses on deriving a straightforward (linear) regression model in Wasserstein spaces, which serves as a generalization of linear regression in Euclidean space. Specifically, for a given set of distributional data points, the goal is to find a geodesic measure-valued curve that best fits the data. This

approach is motivated by the fact that geodesics in Wasserstein space are the equivalent of linear curves in Euclidean space. Both linear and geodesics share two key properties:

- They are the shortest paths between two points.
- They have zero (Riemannian) acceleration.

To characterize geodesic regression curves, we use the vanishing acceleration property rather than the shortest-path property, as the latter can lead to degenerate solutions when the endpoints are not fixed. Instead of strictly enforcing a vanishing acceleration condition, we opt to use a penalty-based method, which augments the objective functional (the data term) with the spline energy term introduced in Section 3.2. This penalty-based approach makes use of the spline energy as a regularizer and thus allows control over the shape of the regression curve. The analogous method has been successfully applied in the context of viscous fluidic objects (cf. [BFH⁺13]). In this section, we present the continuous version of the Wasserstein regression model, followed by its variational-time formulation.

Let us consider sets of input probability measures $\{\bar{\mu}_k^i\}_{i=1,\dots,i_k} \subset \mathcal{P}_2(\Omega)$, for $k = 0, \dots, K$ that correspond to statistical measurements recorded at specific times $t_k \in [0, 1]$. To simplify the notation, we assume that all times at which input measures are given are multiples of the time step size $\tau := \frac{1}{K}$ of the discrete model introduced later. This assumption is not too restrictive, as extending the model to allow non-uniform time step sizes is straightforward.

To motivate this approach, we first consider the Euclidean case, equipped with the standard Euclidean metric. In this setting, the classical linear regression problem involves finding a linear function $y : [0, 1] \rightarrow \mathbb{R}^d$, typically expressed as $y(t) = mt + b$, with $m, b \in \mathbb{R}^d$ that minimizes the least squares error term (also called *data term*)

$$\sum_{k=0}^K \sum_{i=1}^{i_k} |y(t_k) - \bar{y}_k^i|^2,$$

where $\{\bar{y}_k^i\} \subset \mathbb{R}_{i=1,\dots,i_k}^d$, for $k = 0, \dots, K$ are the given (measured) data points.

In the Wasserstein framework, and similar to the motivation given in the introduction of this thesis, the linear function $y(t) = mt + b$ is instead replaced by a geodesic curve $\mu : [0, 1] \rightarrow \mathcal{P}_2(\Omega)$, and the Euclidean distance is substituted with the Wasserstein distance.

The geodesic regression problem in Wasserstein space involves minimizing the data term

$$\mathcal{D}[\mu] := \sum_{k=0}^K \sum_{i=1}^{i_k} \mathcal{W}^2(\mu(t_k) - \bar{\mu}_k^i), \quad (3.8.1)$$

while ensuring that μ is a geodesic curve. As discussed earlier, we shall relax this strict constraint using a penalty method. The penalty functional, given by the spline energy functional \mathcal{F} (cf. (3.2.2)), penalizes deviations from the geodesic equation. In the continuous-time Wasserstein regression model, the goal is to minimize the penalty-augmented objective

$$\mathcal{D}^\varepsilon[\mu] := \mathcal{D}[\mu] + \varepsilon^{-1} \mathcal{F}[\mu] \quad (3.8.2)$$

for a small penalty parameter $\epsilon > 0$.

We now derive a discrete-time version of this continuous variational problem, focusing on discrete geodesic regression curves. For this purpose, we consider discrete measure paths (μ_0, \dots, μ_K) and assume (possibly after re-indexing) that the input measures $\{\bar{\mu}_k^i\}$ are associated with times t_k . Using this notation, the discrete Wasserstein regression problem seeks a discrete path $\mu^K := (\mu_0^K, \dots, \mu_K^K) \in \mathcal{P}_2(\Omega)^{K+1}$ that minimizes the data term

$$\mathbf{D}[\mu^K] := \sum_{k=0}^K \sum_{i=1}^{i_k} \mathcal{W}^2(\mu_k^K - \bar{\mu}_k^i), \quad (3.8.3)$$

while ensuring that μ^K is a discrete geodesic, i.e., it minimizes the energy \mathbf{E} among all discrete paths with fixed endpoint measures μ_0^K and μ_K^K . As in the continuous case, this strict constraint is relaxed by replacing the continuous spline functional \mathcal{F} with its discrete counterpart \mathbf{F}^K , as defined in (3.2.17). The discrete (linear) Wasserstein regression problem then becomes a minimization of the objective

$$\mathbf{D}^\epsilon[\mu^K] := \mathbf{D}[\mu^K] + \epsilon^{-1} \mathbf{F}[\mu^K] \quad (3.8.4)$$

Alternatively, one could use the discrete path energy \mathbf{E} (cf. (3.2.12)) as a penalty term (cf. [BvTH16]). However, for decreasing ϵ , this approach results in degenerate curves because the discrete path energy penalizes the curve's total length, causing the resulting regression curve to collapse into the global Wasserstein barycenter of all the input measures $\bar{\mu}_k^i$ over all time steps. The discrete spline energy penalizes instead its discrete acceleration, and hence, the deviation from a geodesic.

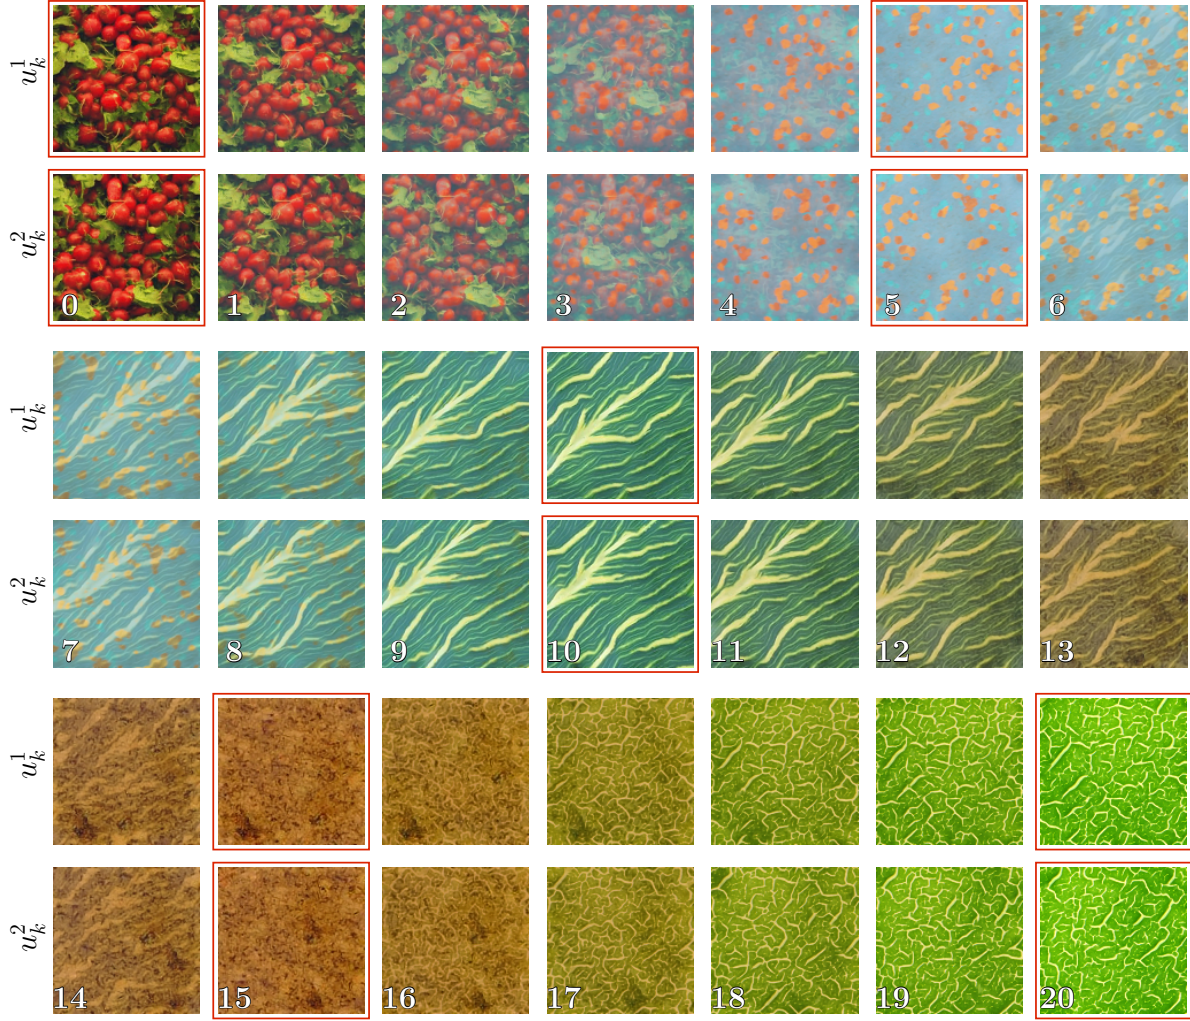


Figure 3.10: Two realizations of a texture spline for different starting latent space samples, top and bottom of each panel respectively with parameters $K = 20$, $\delta = 0.01$. Let us remark that not only the actual spline interpolated textures but also the key frame textures differ as they are all different samples of the underlying spline probability distributions ν_k^K .

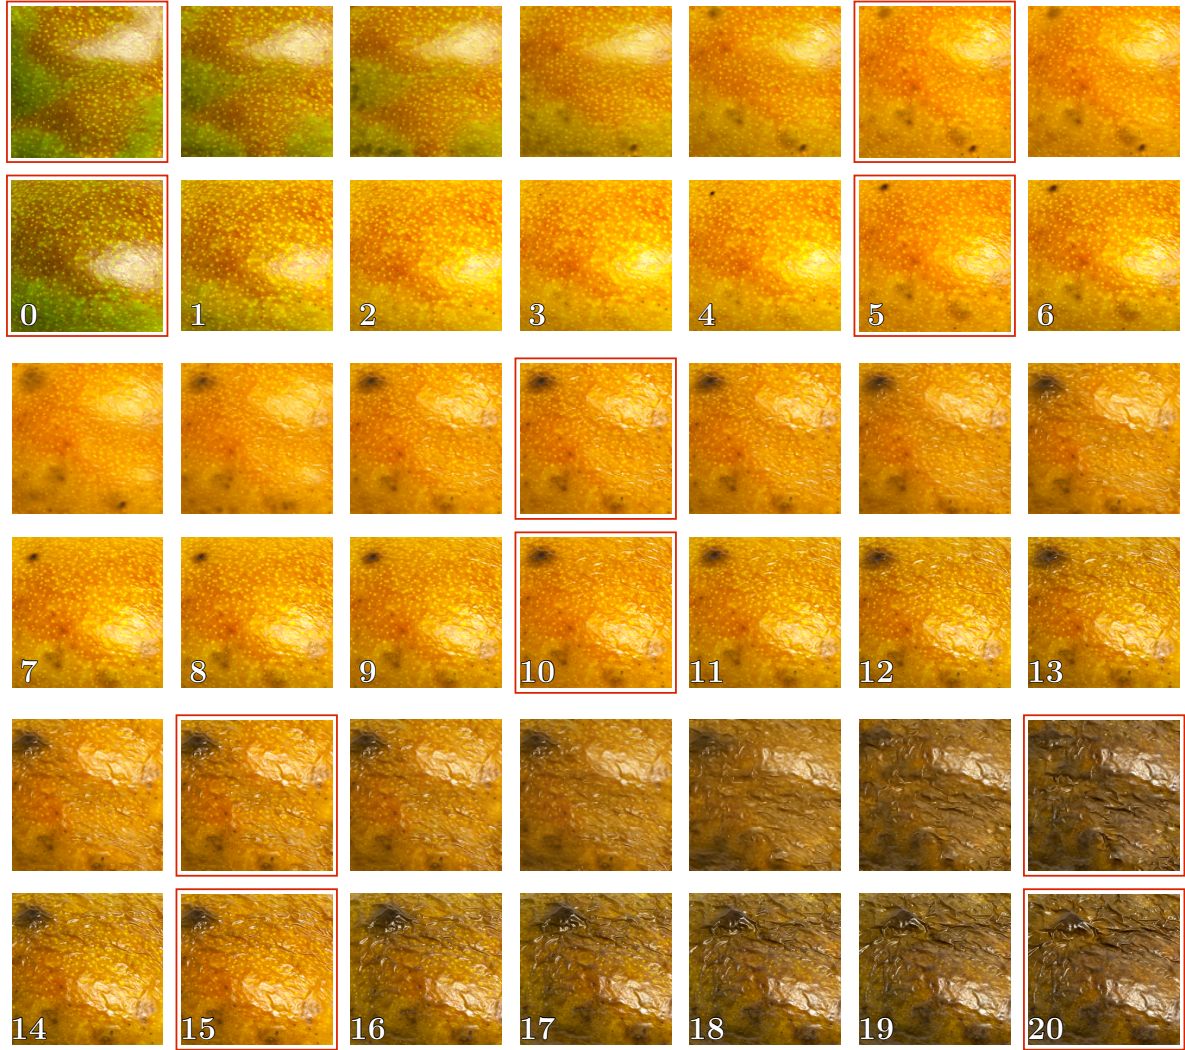


Figure 3.11: Top: A realization of a texture spline with parameters $K = 20$, $\delta = 0.01$. Bottom: Textures sampled at the interpolated times of the actual video from which the texture constraints were extracted from. As in Figure 3.10 the key frame textures from the spline interpolated path differ from the true images at the corresponding times as they are samples of the underlying spline probability distributions ν_k^K .

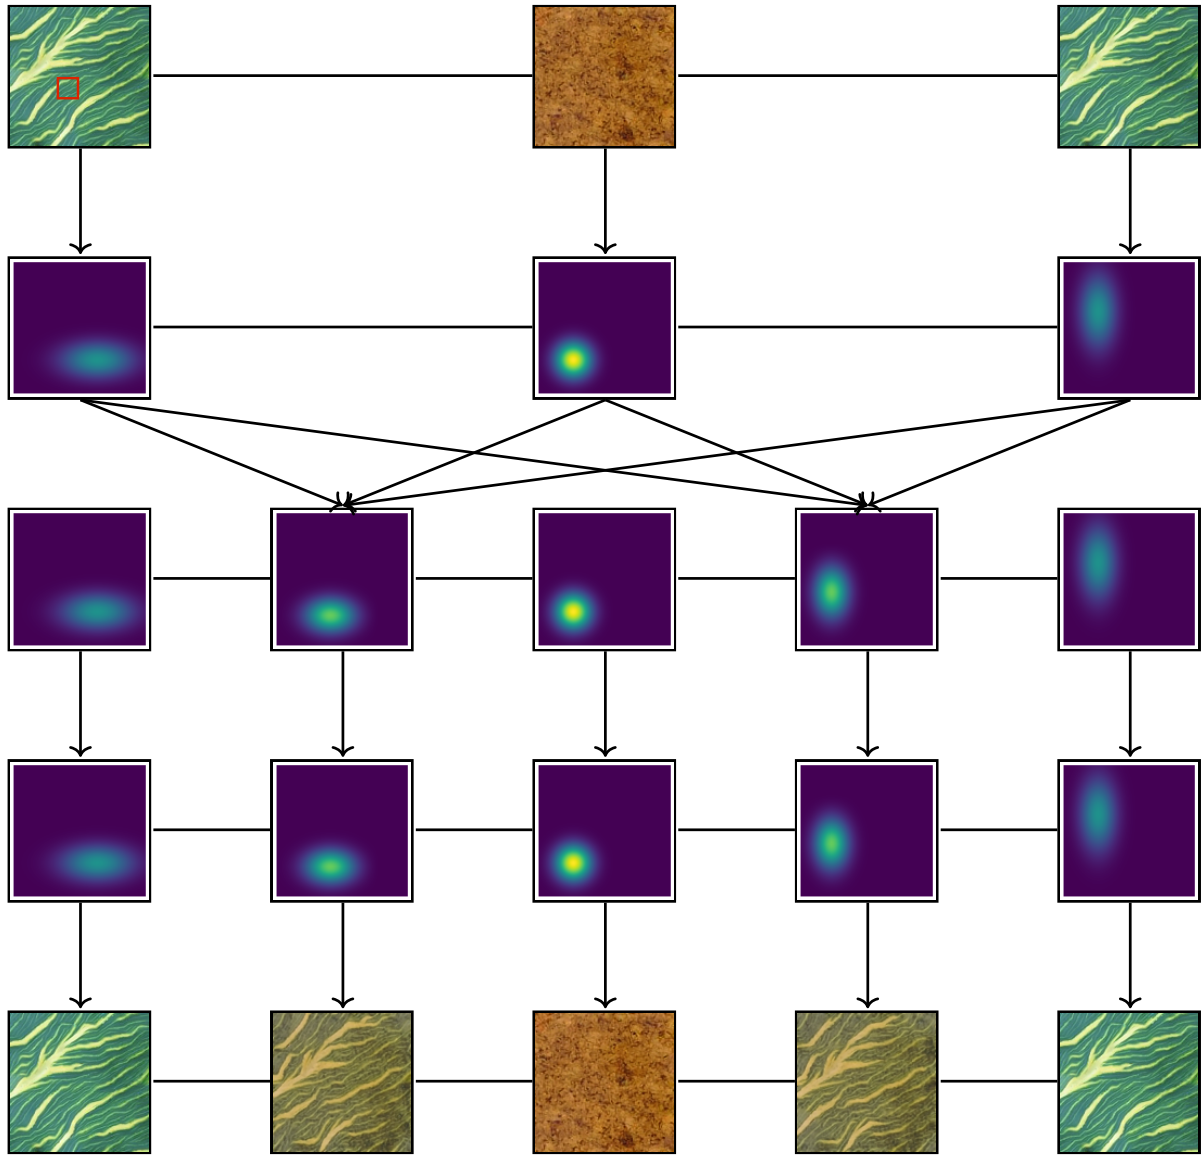


Figure 3.12: Schematic plot of our algorithm to generate data interpolations based on Wasserstein spline interpolation on the latent space of a VAE. The step from the first row to the second one corresponds to the encoder. The step from the second row to the third describes the use of our interpolation algorithm to compute Gaussian distributions in-between. The next step illustrates the sampling algorithm, and the final step (from the fourth to fifth row) represents the decoder part of the VAE.

4 Unbalanced Transport Splines with Source Term

This paper investigates a variational model for splines in the image metamorphosis model for the smooth interpolation of key frames in the space of images. The Riemannian manifold of images based on the metamorphosis model defines shortest geodesic paths interpolating two images as minimizers of the path energy which measures the viscous dissipation caused by the motion field and dissipation caused by the material derivative of the image intensity along motion paths. In this paper we aim at smooth interpolation of multiple key frame images picking up the general observation of cubic splines in Euclidean space which minimize the squared acceleration along the interpolation path. To this end, we propose the spline functional which combines quadratic functionals of the Eulerian motion acceleration and of the second material derivative of the image intensity as the proper notion of image intensity acceleration. We propose a variational time discretization of this model and study the convergence to a suitably relaxed time continuous model via Γ -convergence methodology. As a byproduct, this also allows to establish the existence of metamorphosis splines for given key frame images as minimizers of the time continuous spline functional. The time discretization is complemented by effective spatial discretization based on finite differences and a stable B-spline interpolation of deformed quantities. A variety of numerical examples demonstrates the robustness and versatility of the proposed method in applications. For the minimization of the fully discrete energy functional a variant of the iPALM algorithm is used.

4.1 Review of the Flow of Diffeomorphisms model and Optimal Transport Model with Source Term

In this section, we will briefly review the flow of diffeomorphisms model and the generalized optimal transport model with source term.

4.1.1 Flow of Diffeomorphisms

The flow of diffeomorphisms model [DGM98, BMTY05, JM00, MTY02] is based on Arnold's paradigm [Arn66] that studied the flow of ideal fluids transported by geodesics living in the space of orientation-preserving diffeomorphisms. Given a domain $\Omega \subset \mathbb{R}^d$, one considers a family of diffeomorphisms $(\psi_t)_{t \in [0,1]} : \bar{\Omega} \rightarrow \mathbb{R}^d$ determined by its time-dependent *Eulerian velocity* v_t via

$$v_t(\psi_t(x)) = \dot{\psi}_t(x), \quad \psi_0(x) = x, \quad (4.1.1)$$

where each ψ_t is an element of the space of *Sobolev diffeomorphisms* on Ω , denoted as $\mathcal{D}^m(\Omega)$, and defined by

$$\mathcal{D}^m(\Omega) := \{\psi \in H^m(\Omega, \Omega), \det(D\psi) > 0, \psi|_{\partial\Omega} = \mathbb{1}\}, \quad m > 1 + \frac{d}{2}. \quad (4.1.2)$$

Functions in this space are indeed $C^1(\Omega, \Omega)$ -orientation preserving diffeomorphisms since [IKT13, Lemma A.1] guarantees the existence of $\psi^{-1} \in \mathcal{D}^m(\Omega)$ and, by the Sobolev embedding theorem 2.2.13, we have $\mathcal{D}^m(\Omega) \subset C^1(\overline{\Omega}, \overline{\Omega})$. By [IKT13, Lemma 2.18], the set of all $H^m(\Omega, \Omega)$ functions which are orientation preserving diffeomorphisms is an open subset of the Hilbert space $H^m(\Omega, \Omega)$ and can be seen as C^∞ -manifold embedded in this space. For every $\psi \in \mathcal{D}^m(\Omega)$, the tangent space $T_\psi \mathcal{D}^m(\Omega) = \mathcal{V} := H^m(\Omega, \mathbb{R}^n) \cap H_0^1(\Omega, \mathbb{R}^n)$ at ψ (cf. ...) is the velocity space of smooth curves with origin at that point. The Riemannian structure on this space is defined via the metric

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

and the path energy

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

Here, L defines a quadratic form corresponding to a higher order elliptic operator. The particular choice used throughout this chapter is

$$L[v, v] := \text{tr}(\varepsilon[v]^2) + \gamma |D^m v|^2, \quad m > 1 + \frac{d}{2}, \gamma > 0. \quad (4.1.3)$$

In physical terms, the metric $g_{\psi_t}(\dot{\psi}_t, \dot{\psi}_t)$ describes the viscous dissipation in a multi-polar fluid model as investigated by Nečas and Šilhavý [NŠ91]. The first term of L represents the local dissipation density in a simple Newtonian fluid and the second term can be regarded as a higher order measure of the fluid friction. Using that the metric g_{ψ_t} is $H^m(\Omega)$ -coercive, [DGM98, Theorem 3.1] guarantees the existence of a flow of diffeomorphisms as a minimizer of the above path energy. This minimizer represents a geodesic path connecting two given diffeomorphisms.

In the context of image morphing, the flow of diffeomorphism consists in transporting densities along particle paths describing the temporal change of c -channel image densities $(\vartheta_t)_{t \in [0,1]} : \Omega \rightarrow \mathbb{R}^c$. This transport is given in terms of the equation $\vartheta_t(\cdot) := \vartheta_0 \circ \psi_t^{-1}(\cdot)$, also known as the *brightness constancy assumption* [HS81], which is equivalent to a vanishing material derivative $\frac{D}{Dt} \vartheta := \dot{\vartheta} + v \cdot D \vartheta$. Given two image density functions ϑ_A, ϑ_B , an associated geodesic path is a family of densities subject to the constraint $\vartheta_0 = \vartheta_A, \vartheta_1 = \vartheta_B$ and $\vartheta_t(\cdot) = \vartheta_A \circ \psi_t^{-1}(\cdot)$ where the underlying family of diffeomorphisms $(\psi_t)_{t \in [0,1]}$ minimizes the path energy above.

The Flow Equation

For the sake of self-containment, here we re-state some technical on the existence, uniqueness, and regularity of diffeomorphic flows given as a solution to (4.1.1), originally proven in [Raj23] and the textbook [You10]:

Theorem 4.1.1. *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 to the flow equation (4.1.1), in the sense that $\psi_t(x) = x + \int_0^t v_s \circ \psi_s(x) ds$ holds 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*

$$\|\psi\|_{C^0([0,1], C^{1,\alpha}(\overline{\Omega}))} + \|\psi^{-1}\|_{C^0([0,1], C^{1,\alpha}(\overline{\Omega}))} \leq G \left(\int_0^1 \|v_s\|_{C^{1,\alpha}(\overline{\Omega})} ds \right), \quad (4.1.4)$$

where $G(x) := C(x+1)\exp(Cx)$. The solution operator $L^2((0,1), \mathcal{V}) \rightarrow C^0([0,1], H^m(\Omega, \Omega))$, $v \mapsto \psi^v$, 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 $C^0([0,1] \times \overline{\Omega})$ -topology for ψ .

Proof. See [Raj23, Theorem 3.2.1.]. □

Remark 4.1.2. If \mathcal{V} is replaced by $C^{1,\alpha}(\overline{\Omega}, \mathbb{R}^n)$ with zero boundary condition [You10, Chapter 8], analogous existence results and estimates as in Theorem 4.1.1 hold. In addition, the mapping $v \mapsto \psi^v$ is Lipschitz continuous in v and it holds

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

where $C = \int_0^t \|v_s\|_{C^1(\overline{\Omega})} ds$.

Proof. See [You10, Proof of Theorem 8.10]. □

Remark 4.1.3. By the Sobolev embedding theorem 2.2.13, we have $\psi \in C^{0,\frac{1}{2}}([0,1], C^{1,\alpha}(\overline{\Omega}, \overline{\Omega}))$ and we obtain the following estimates for $|\psi|_{C^{0,\frac{1}{2}}([0,1], C^{1,\alpha}(\overline{\Omega}))}$:

$$\|\psi_t - \psi_s\|_{C^{1,\alpha}(\overline{\Omega})} \leq G_1 \left(\int_0^1 \|v_s\|_{C^{1,\alpha}(\overline{\Omega})} ds \right) \|v\|_{L^2((0,1), C^{1,\alpha}(\overline{\Omega}))} |t - s|^{\frac{1}{2}},$$

for a positive, monotonically increasing function G_1 . Analogous results hold for the inverse flow $|\psi^{-1}|_{C^{0,\frac{1}{2}}([0,1], C^{1,\alpha}(\overline{\Omega}))}$.

Proof. See [Raj23, Remark 3.2.3]. □

We refer the reader to [Raj23] for further details about the flow equation.

4.1.2 Generalized Optimal Transport Model with Source Term

In the sequel, we shall make no distinction between a measure $\vartheta \in \mathcal{M}(\Omega)$ or its Radon-Nikodym density, and we will use both terms interchangeably. In the flow formulation of the Wasserstein distance proposed by Benamou and Brenier (cf. [BB00]), the distance between two probability measures was computed by minimizing over measure-valued paths $(\vartheta_t)_t$ determined by velocity fields $(v_t)_t$. In this setting, mass particles of the initial measure move along the family of flow paths \mathcal{T} , which is uniquely determined by the velocity field via (3.1.3). This transport does not induce any change in the total mass of the involved measures, known as the *mass conservation assumption*. This assumption is equivalent to a vanishing source term \hat{z}_t in the continuity equation (cf. (3.1.1))

$$\hat{z}_t := \dot{\vartheta}_t + \nabla \cdot (v_t \vartheta_t), \tag{4.1.5}$$

for a path $(\vartheta_t)_{t \in [0,1]}$ in the more general space of signed measures (which do not necessarily have the same total mass).

In concrete applications, such as computer vision, objects of generally different masses are expected to be matched between images. This could be induced, for instance, by the expansion of an object without any change in its intensity, or alternatively by a change in lighting that affects the intensity without changing the size of the object in question. Hence, the assumption that different images are to

be encoded as probability measures is too restrictive, and leads to artifacts stemming from the attempt to compensate local mass surpluses/deficits in an artificial manner. The generalized optimal transport model with source term, originally proposed by Maas et al. in [MRSS15], relaxes the optimal transport model by defining a metric on the space of general (signed) measures that allows for a non-vanishing source term z_t , and penalizes its squared L^2 -norm instead, thus allowing for changes in mass along measure-valued paths. Under the assumption that the measure-valued path ϑ is sufficiently smooth, the metric and the path energy read as

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

for a penalization parameter $\delta > 0$. The term $|v|^2 \vartheta$ quantifies the underlying cost of the transport of mass, while the second term \hat{z} reflects the penalization of the changes in mass that cannot be explained by transport of mass. Hence, the original optimal transport model can be seen as the limit case of the generalized optimal transport model with source term for $\delta \rightarrow 0$.

Another issue that arises in the optimal transport model is that, as opposed to the setting in the flow of diffeomorphisms model, the optimal transport flow maps (3.1.3) are not necessarily homeomorphisms. In practice, however, one is interested in topological consistency. Due to the often times physically motivated background of applications, Maas et al. [MRSS15] propose to incorporate a dissipative term in the metric. This term not only endows the mass transport with a physical meaning, but also asserts some additional regularity on the resulting flow maps. In the combined optimal transport model with source term and a weighted viscous dissipation model, the metric and the path energy are explicitly given by

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

for penalty parameters $\delta, \lambda > 0$, and a local rate of viscous dissipation L given as in (4.1.3). Notably, this definition of the metric has two major drawbacks:

- In general, paths in the space of measures do not exhibit any smoothness properties (neither spatially nor temporally), and therefore the point-wise evaluation of the source term is not well-defined.
- Different pairs (v, \hat{z}) of velocity fields and source terms can induce the same temporal variation $\dot{\vartheta}$ of the measure-valued path.

To tackle both of these problems, Maas et al. proposed a non-linear geometric structure in the space of measures with L^2 Radon-Nikodym densities $\mathcal{I} := L^2(\Omega, \mathbb{R}^3)$ with respect to the Lebesgue measure. More specifically, for a given density path $\vartheta \in L^2([0, 1], \mathcal{I})$ 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 first-order source term* $\hat{z} \in L^2((0, 1), \mathcal{Z} := L^2(\Omega, \mathbb{R}^3))$ is incorporated in the model and it is implicitly given by

$$\int_0^1 \int_{\Omega} \eta \hat{z} dx dt = - \int_0^1 \int_{\Omega} (\partial_t \eta + v \cdot \nabla \eta) \vartheta dx dt \quad (4.1.6)$$

for all $\eta \in C_c^{\infty}((0, 1) \times \Omega)$. Here, we consider (v, \hat{z}) as a tangent vector in the tangent space of \mathcal{I} at ϑ . Indeed, (v, \hat{z}) represents a variation of the measure ϑ via transport and change of density. For a density

$\vartheta \in \mathcal{I}$, we define the non-linear structure on $\mathcal{V} \times \mathcal{Z}$ by

$$N_\vartheta := \left\{ (v, \hat{z}) \in \mathcal{V} \times \mathcal{Z} : \int_{\Omega} \eta \hat{z} + \vartheta v \cdot \nabla \eta \, dx = 0, \forall \eta \in C_c^\infty([0, 1] \times \overline{\Omega}) \right\}.$$

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

$$T_\vartheta \mathcal{I} := (\mathcal{V} \times \mathcal{Z}) / N_\vartheta,$$

giving rise to the associated tangent bundle $T\mathcal{I}$. The path energy can now be rigorously defined for a density-valued path ϑ by

$$\mathcal{E}^{\delta, \lambda}[\vartheta] := \int_0^1 \inf_{(v, \hat{z}) \in T_{\vartheta(t)} \mathcal{I}} \int_{\Omega} |v|^2 \vartheta + \frac{1}{\delta} |\hat{z}|^2 + \lambda L[v, v] \, dx \, dt. \quad (4.1.7)$$

We say ϑ is $H^1([0, 1], \mathcal{I})$ -regular if $\vartheta \in C^0([0, 1], \mathcal{I})$ and $\mathcal{E}^{\delta, \lambda}[\vartheta] < \infty$. Given this path energy, we can now define a (generalized) Riemannian Wasserstein distance $\mathcal{W}_{\delta, \lambda}(\vartheta_A, \vartheta_B)$ between two densities $\vartheta_A, \vartheta_B \in \mathcal{I}$ as

$$\mathcal{W}_{\delta, \lambda}^2(\vartheta_A, \vartheta_B) := \inf \{ \mathcal{E}^{\delta, \lambda}[\tilde{\vartheta}] : \tilde{\vartheta} \in H^1([0, 1], \mathcal{I}), \tilde{\vartheta}_0 = \vartheta_A, \tilde{\vartheta}_1 = \vartheta_B \} \quad (4.1.8)$$

A *geodesic curve in the space of densities* connecting $\vartheta_A, \vartheta_B \in \mathcal{I}$ is defined as a regular curve $\vartheta \in H^1([0, 1], \mathcal{I})$ with $\vartheta_0 = \vartheta_A, \vartheta_1 = \vartheta_B$, such that

$$\mathcal{E}^{\delta, \lambda}[\vartheta] = \mathcal{W}_{\delta, \lambda}^2(\vartheta_A, \vartheta_B).$$

The existence of a geodesic curve for the non-viscous case ($\lambda = 0$) is proven in [MRSS15, Theorem 2.5].

4.2 Time Continuous Splines in the Generalized Model

In this section, we expand on the unbalanced optimal transport model introduced by [MRSS15] and briefly explained in the previous section to study spline interpolations for a set of given key frame densities at given time stamps in a time continuous setting. In mathematical notation, given a set of $J \geq 2$ key frames $\vartheta_1^I, \dots, \vartheta_J^I \in \mathcal{I}$ and time stamps $0 \leq t_1 < \dots < t_J \leq 1$, we ask for a spline interpolation $(\vartheta_t)_{t \in [0, 1]} \in L^2([0, 1], \mathcal{I})$, which satisfies the constraints

$$\vartheta_{t_j} = \vartheta_j^I, \quad t_j \in [0, 1], \quad j = 1, \dots, J. \quad (4.2.1)$$

To this end, we recall that cubic splines in Euclidean space were introduced by de Boor [dB63] as minimizers of the integral over the squared acceleration, subject to position constraints, whereas linear interpolation is instead associated with the minimization of the integral over the squared velocity. In our case, density morphing via minimization of the path energy (4.1.7) corresponds to this linear interpolation. In this section, we introduce the relevant second-order quantities which will be penalized in the spline energy:

Flow acceleration.

In the original flow of diffeomorphisms model, we observed a diffeomorphic flow ψ determined by the Eulerian velocity $v \in L^2((0, 1), \mathcal{V})$, cf. (4.1.1). For our second order model, we now consider the pairs $(v, a) \in L^2((0, 1), \mathcal{V}) \times L^2((0, 1), \mathcal{V})$ which determine the system

$$a_t(\psi_t(x)) = \ddot{\psi}_t(x), \quad v_t(\psi_t(x)) = \dot{\psi}_t(x), \quad \psi_0(x) = x, \quad (4.2.2)$$

and the corresponding diffeomorphic flow $\psi \in L^2((0, 1), H^m(\Omega, \Omega))$. By using the same arguments as in Theorem 4.1.1, we obtain existence and uniqueness of a flow $\psi \in H^2((0, 1), H^m(\Omega, \Omega)) \subset C^{1, \frac{1}{2}}((0, 1), C^{1, \alpha}(\overline{\Omega}, \overline{\Omega}))$ as a solution to the above system. In this case, the $C^1([0, 1], C^{1, \alpha}(\overline{\Omega}))$ -norm of ψ depends on $\|a\|_{L^2((0, 1), \mathcal{V})}$ and $\|v\|_{L^2((0, 1), \mathcal{V})}$. In particular, it holds $v \in H^1((0, 1), \mathcal{V}) \subset C^{0, \frac{1}{2}}((0, 1), \mathcal{V})$. Notice that from the following interdependency between a and v

$$a_t \circ \psi_t = \ddot{\psi}_t = \frac{d}{dt}(v_t \circ \psi_t) = (\dot{v}_t + Dv_t v_t) \circ \psi_t, \quad (4.2.3)$$

one would expect the acceleration to have one fewer spatial derivative compared to the velocity. However, the approach we later take allows us to have the same number of derivatives. A similar gain of smoothness was observed in [Via20] and [JRR23].

Second order source term.

In order to derive a suitable formulation for the second-order source term, let us first prove a few technical lemmas in the smooth case, which will then be adapted for the second order calculus.

Lemma 4.2.1. *Let $(\vartheta_t)_{t \in [0, 1]} \in C^\infty([0, 1], C^\infty(\Omega))$ be a smooth family of densities, and $(\psi_t)_{t \in [0, 1]} \in C^\infty((0, 1), C^\infty(\overline{\Omega}, \overline{\Omega}))$ with $\det(D\psi_t) > 0$ for all $t \in (0, 1)$. Then, it holds*

$$\hat{z}_t = \det(D\psi_t^{-1}) \left(\frac{d}{dt} \det(D\psi_t) \vartheta_t(\psi_t) \right) \circ \psi_t^{-1}. \quad (4.2.4)$$

Proof. Using $v_t = \dot{\psi}_t \circ \psi_t^{-1}$, $\frac{d}{dt} \det(A_t) = \det(A_t) \operatorname{tr}(A_t^{-1} \dot{A}_t)$ and $\operatorname{tr}(Dv_t) = \operatorname{div}(v_t)$, we obtain

$$\begin{aligned} & \det(D\psi_t^{-1}) \left(\frac{d}{dt} \det(D\psi_t) \vartheta_t(\psi_t) \right) \circ \psi_t^{-1} \\ &= \det(D\psi_t^{-1}) \left(\det(D\psi_t) \operatorname{tr}((D\dot{\psi}_t)(D\psi_t)^{-1}) \vartheta_t(\psi_t) + \det(D\psi_t) \dot{\vartheta}_t(\psi_t) \right. \\ & \quad \left. + \det(D\psi_t) \nabla \vartheta_t(\psi_t) \dot{\psi}_t \right) \circ \psi_t^{-1} \\ &= \det(D\psi_t^{-1}) \det(D\psi_t)_{\psi_t^{-1}(\cdot)} \left(\operatorname{tr}(D(\dot{\psi}_t \circ \psi_t^{-1})) \vartheta_t + \dot{\vartheta}_t + \nabla \vartheta_t \cdot (\dot{\psi}_t \circ \psi_t^{-1}) \right) \\ &= \operatorname{div}(v_t) \vartheta_t + \dot{\vartheta}_t + \nabla \vartheta_t \cdot v_t = \dot{\vartheta}_t + \operatorname{div}(\vartheta_t v_t) = \hat{z}_t \end{aligned}$$

□

In light of Lemma 4.2.1, it is then natural to consider

$$\hat{w}_t := \det(D\psi_t^{-1}) \left(\frac{d^2}{dt^2} \det(D\psi_t) \vartheta_t(\psi_t) \right) \circ \psi_t^{-1}, \quad (4.2.5)$$

as the second order source term. For the smooth case, we obtain

Lemma 4.2.2. *Let $(\vartheta_t)_{t \in [0,1]} \in C^\infty([0,1], C^\infty(\Omega))$ be a smooth family of densities, and $(\psi_t)_{t \in [0,1]} \in C^\infty((0,1), C^\infty(\bar{\Omega}, \bar{\Omega}))$ with $\det(D\psi_t) > 0$ for all $t \in (0,1)$. Then, it holds*

$$\begin{aligned} \hat{w}_t = & \operatorname{div}(v_t)^2 \vartheta_t + \operatorname{div}(a_t \vartheta_t) - \operatorname{tr}((Dv_t)^2) \vartheta_t + 2 \operatorname{div}(v_t \dot{\vartheta}_t) + 2 \operatorname{div}(v_t) \nabla \vartheta_t \cdot v_t + \ddot{\vartheta}_t \\ & + \operatorname{Hess} \vartheta_t(v_t, v_t). \end{aligned} \quad (4.2.6)$$

Proof. Using the same facts as in the proof of Lemma 4.2.1, we have

$$\begin{aligned} \det(D\psi_t^{-1}) \left(\frac{d^2}{dt^2} \det(D\psi_t) \vartheta_t(\psi_t) \right) \circ \psi_t^{-1} &= \det(D\psi_t^{-1}) \left(\frac{d}{dt} \det(D\psi_t) \operatorname{tr}((D\dot{\psi}_t)(D\psi_t)^{-1}) \vartheta_t(\psi_t) \right. \\ &+ \det(D\psi_t) \dot{\vartheta}_t(\psi_t) + \det(D\psi_t) \nabla \vartheta_t(\psi_t) \dot{\psi}_t \left. \right) \circ \psi_t^{-1} \\ &= \det(D\psi_t^{-1}) \left(\det(D\psi_t) \operatorname{tr}((D\dot{\psi}_t)(D\psi_t)^{-1})^2 \vartheta_t(\psi_t) + \det(D\psi_t) \operatorname{tr}((D\ddot{\psi}_t)(D\psi_t)^{-1}) \vartheta_t(\psi_t) \right. \\ &- \det(D\psi_t) \operatorname{tr}((D\dot{\psi}_t)(D\psi_t)^{-1} (D\dot{\psi}_t)(D\psi_t)^{-1}) \vartheta_t(\psi_t) + 2 \det(D\psi_t) \operatorname{tr}((D\dot{\psi}_t)(D\psi_t)^{-1}) \dot{\vartheta}_t(\psi_t) \\ &+ 2 \det(D\psi_t) \operatorname{tr}((D\dot{\psi}_t)(D\psi_t)^{-1}) \nabla \vartheta_t(\psi_t) \dot{\psi}_t + \det(D\psi_t) \ddot{\vartheta}_t(\psi_t) + 2 \det(D\psi_t) \nabla \dot{\vartheta}_t(\psi_t) \dot{\psi}_t \\ &+ \det(D\psi_t) \nabla \vartheta_t(\psi_t) \ddot{\psi}_t + \operatorname{Hess} \vartheta_t(\psi_t)(\dot{\psi}_t, \dot{\psi}_t) \left. \right) \circ \psi_t^{-1} \\ &= \operatorname{tr}(D(\dot{\psi}_t \circ \psi_t^{-1}))^2 \vartheta_t + \operatorname{tr}(D(\ddot{\psi}_t \circ \psi_t^{-1})) \vartheta_t - \operatorname{tr}(D(\dot{\psi}_t \circ \psi_t^{-1})^2) \vartheta_t + 2 \operatorname{tr}(D(\dot{\psi}_t \circ \psi_t^{-1})) \dot{\vartheta}_t \\ &+ 2 \operatorname{tr}(D(\dot{\psi}_t \circ \psi_t^{-1})) \nabla \vartheta_t \cdot \dot{\psi}_t \circ \psi_t^{-1} + \ddot{\vartheta}_t + 2 \nabla \dot{\vartheta}_t \cdot \dot{\psi}_t \circ \psi_t^{-1} + \nabla \vartheta_t \cdot \ddot{\psi}_t \circ \psi_t^{-1} \\ &= \operatorname{div}(v_t)^2 \vartheta_t + \operatorname{div}(a_t) \vartheta_t - \operatorname{tr}((Dv_t)^2) \vartheta_t + 2 \operatorname{div}(v_t) \dot{\vartheta}_t + 2 \operatorname{div}(v_t) \nabla \vartheta_t \cdot v_t + \ddot{\vartheta}_t + 2 \nabla \dot{\vartheta}_t \cdot v_t \\ &+ \nabla \vartheta_t \cdot a_t + \operatorname{Hess} \vartheta_t(v_t, v_t) \\ &= \operatorname{div}(v_t)^2 \vartheta_t + \operatorname{div}(a_t \vartheta_t) - \operatorname{tr}((Dv_t)^2) \vartheta_t + 2 \operatorname{div}(v_t \dot{\vartheta}_t) + 2 \operatorname{div}(v_t) \nabla \vartheta_t \cdot v_t + \ddot{\vartheta}_t + \operatorname{Hess} \vartheta_t(v_t, v_t) \end{aligned}$$

□

Once again, (4.2.6) should be understood in weak form. Our proposed spline energy then consists of the integral over the three squared second order quantities:

$$\mathcal{F}^{\delta, \lambda}[\vartheta] := \min_v \int_0^1 \int_\Omega |a_t|^2 \vartheta_t + \frac{1}{\delta} |\hat{w}_t|^2 + \lambda L[a_t, a_t] \, dx \, dt,$$

where for simplicity we use the same elliptic operator L as in (4.1.3). As already mentioned in the introduction, the splitting of the spline energy into a flow acceleration and a second order source term is not fully consistent with the Riemannian structure defined in Subsection 4.1.2, since one would expect an interaction between the two terms in the second order calculus. However, we have seen the analogous approach in the metamorphosis setting (cf. [JRR23]) yield very good results.

As in the case of geodesic path energy, we shall now give rigorous formulations for general paths $\vartheta \in L^2([0,1], \mathcal{I})$ that may lack enough regularity to define the quantities (4.1.5) and (4.2.6) point-wise. By equations (4.2.4) and (4.2.5), the integral Lagrangian formulations of the first and second order source terms $\hat{z}, \hat{w} \in L^2((0,1), \mathcal{I})$ are formally given by

$$\begin{aligned} \int_s^t \det(D\psi_r) \hat{z}_r \circ \psi_r \, dr &= \det(D\psi_t) \vartheta_t \circ \psi_t - \det(D\psi_s) \vartheta_s \circ \psi_s, \\ \int_0^\tau \int_s^t \det(D\psi_{r+l}) \hat{w}_{r+l} \circ \psi_{r+l} \, dr \, dl &= \int_0^\tau \det(D\psi_{t+l}) \hat{z}_{t+l} \circ \psi_{t+l} - \det(D\psi_{s+l}) \hat{z}_{s+l} \circ \psi_{s+l} \, dl \end{aligned} \quad (4.2.7)$$

$$= \det(D\psi_{t+\tau})\vartheta_{t+\tau} \circ \psi_{t+\tau} - \det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_{s+\tau})\vartheta_{s+\tau} \circ \psi_{s+\tau} + \det(D\psi_s)\vartheta_s \circ \psi_s, \quad (4.2.8)$$

respectively 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 (4.2.8) is an integral version of the second order central difference. Because there is no differentiation involved in these definitions, they work for general density paths. Analogous to [JRR23], we introduce the scalar quantities $z, w \in L^2((0, 1) \times \Omega)$ as relaxations of the weak first and second order source terms, respectively:

$$|\det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_s)\vartheta_s \circ \psi_s| \leq \int_s^t \det(D\psi_r)z_r \circ \psi_r \, dr, \quad \forall s \leq t \in [0, 1], \quad (4.2.9)$$

$$\begin{aligned} & |\det(D\psi_{t+\tau})u_{t+\tau} \circ \psi_{t+\tau} - \det(D\psi_t)u_t \circ \psi_t - \det(D\psi_{s+\tau})u_{s+\tau} \circ \psi_{s+\tau} + \det(D\psi_s)u_s \circ \psi_s| \\ & \leq \int_0^\tau \int_s^t \det(D\psi_{r+l})w_{r+l} \circ \psi_{r+l} \, dr \, dl, \quad \forall \tau, s + \tau \leq t + \tau \in [0, 1]. \end{aligned} \quad (4.2.10)$$

for every $s \leq t \in (0, 1), \tau \geq 0, t + \tau \leq 1$. This relaxed Lagrangian approach is substantially more elegant in comparison to the Eulerian approach (4.1.5) and (4.2.6), which will be exploited in the proof of consistency of continuous and time discrete approaches.

We can show the equivalence of the two approaches corresponding to (4.2.8) and (4.2.10) (cf. [ENR20, Proposition 8]).

Proposition 4.2.3. *Let $\psi \in H^2((0, 1), H^m(\Omega, \Omega))$ be determined by the flow system (4.2.2). For every vector valued (\hat{z}, \hat{w}) satisfying (4.2.7) and (4.2.8) there exist scalar quantities (z, w) satisfying (4.2.9) and (4.2.10) with $z = |\hat{z}|$ and $w = |\hat{w}|$. Conversely, for every (z, w) satisfying (4.2.9) and (4.2.10) there exists (\hat{z}, \hat{w}) satisfying (4.2.7) and (4.2.8) with $z \geq |\hat{z}|$ and $w \geq |\hat{w}|$.*

Proof. The first claim is obvious by the triangle inequality. To prove the converse, let z satisfy (4.2.9). We take the $L^2(\Omega)$ -norm on both sides and use the triangle inequality to obtain

$$\|\det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_s)\vartheta_s \circ \psi_s\|_{L^2(\Omega)} \leq \int_s^t \|\det(D\psi_r)z_r \circ \psi_r\|_{L^2(\Omega)} \, dr.$$

By Theorem 4.1.1, we know that the function $r \mapsto \|\det(D\psi_r)z_r \circ \psi_r\|_{L^2(\Omega)}$ is in $L^2(0, 1)$ from where we conclude that the function $t \mapsto \det(D\psi_t)\vartheta_t \circ \psi_t$ is in $AC^2(0, 1; L^2(\Omega, \mathbb{R}^c))$, which implies the a.e. differentiability and the existence of a metric derivative $z' \in L^2((0, 1), L^2(\Omega, \mathbb{R}^c))$, such that

$$\det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_s)\vartheta_s \circ \psi_s = \int_s^t z'_r \, dr = \int_s^t \det(D\psi_r)\hat{z}_r \circ \psi_r \, dr$$

where $\hat{z}_r := \det(\psi_r^{-1})z'_r \circ \psi_r^{-1}$ (cf. [AGS08, Remark 1.1.3]). Next, by [AGS08, Theorem 1.1.2]), we obtain $|\det(D\psi_t)\hat{z}_t \circ \psi_t| \leq \det(D\psi_t)z_t \circ \psi_t$ for a.e. $t \in (0, 1)$. Since $\det(D\psi_t)$ is bounded from below by $c_v > 0$, and ψ_t is a diffeomorphism, we finally obtain $z \geq |\hat{z}|$ for a.e. $t \in (0, 1)$. Thus, we have verified (4.2.7). This now implies

$$\left| \int_0^\tau \det(D\psi_{t+l})\hat{z}_{t+l} \circ \psi_{t+l} - \det(D\psi_{s+l})\hat{z}_{s+l} \circ \psi_{s+l} \, dl \right| \leq \int_0^\tau \int_s^t \det(\psi_{r+l})w_{r+l} \circ \psi_{r+l} \, dr \, dl,$$

for every $s \leq t \in [0, 1]$ and $\tau > 0$ with $t + \tau \leq 1$, and the same holds for integration on interval $[-\tau, 0]$ for $s - \tau \geq 0$. Taking the limit as τ tends to zero and using Lebesgue's differentiation theorem [Fol99, Theorem 3.21] we conclude that for all $s \leq t \in [0, 1]$ and a.e. $x \in \Omega$ we have

$$\left| \det(D\psi_t)\hat{z}_t \circ \psi_t - \det(D\psi_s)\hat{z}_s \circ \psi_s \right| \leq \int_s^t \det(D\psi_r)w_r \circ \psi_r \, dr.$$

Thus, we can iterate the procedure from above to conclude the existence of $\hat{w} \in L^2((0, 1), L^2(\Omega, \mathbb{R}^c))$ satisfying

$$\det(D\psi_t)\hat{z}_t \circ \psi_t - \det(D\psi_s)\hat{z}_s \circ \psi_s = \int_s^t \det(D\psi_r)\hat{w}_r \circ \psi_r \, dr, \quad |\hat{w}| \leq w,$$

from where the claim directly follows. \square

In [HRW18] a regularization of the spline path energy by addition of weighted geodesic path energy was necessary for existence and further analysis of the splines. We follow an analogous approach which also seems to be indispensable in our model.

Definition 4.2.4 (Regularized spline energy). Let $\sigma > 0$, $m > 1 + \frac{d}{2}$ be an integer, and $\vartheta \in L^2([0, 1], \mathcal{I})$ be an image curve. Then, the regularized spline energy is defined by

$$\mathcal{F}^{\sigma, \delta, \lambda}[\vartheta] := \inf_{(v, a, z, w) \in \mathcal{C}[\vartheta]} \int_0^1 \int_{\Omega} |a_t|^2 \vartheta_t + \lambda L[a_t, a_t] + \frac{1}{\delta} w_t^2 + \sigma \left(|v_t|^2 \vartheta_t + \lambda L[v_t, v_t] + \frac{1}{\delta} z_t^2 \right) dx \, dt, \quad (4.2.11)$$

where $\mathcal{C}[\vartheta] \subset L^2((0, 1), \mathcal{V}) \times L^2((0, 1), \mathcal{V}) \times L^2((0, 1) \times \Omega) \times L^2((0, 1) \times \Omega)$ consists of tuples (v, a, z, w) satisfying

$$v_t(\psi_t(t, x)) = \dot{\psi}_t(x), \quad \psi_0(x) = x, \quad (4.2.12)$$

$$a_t(\psi_t(x)) = \ddot{\psi}_t(x), \quad \forall x \in \Omega, \quad t \in [0, 1], \quad (4.2.13)$$

$$|\det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_s)\vartheta_s \circ \psi_s| \leq \int_s^t \det(D\psi_r)z_r \circ \psi_r \, dr, \quad \forall s \leq t \in [0, 1], \quad (4.2.14)$$

$$\begin{aligned} & |\det(D\psi_{t+\tau})\vartheta_{t+\tau} \circ \psi_{t+\tau} - \det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_{s+\tau})\vartheta_{s+\tau} \circ \psi_{s+\tau} + \det(D\psi_s)\vartheta_s \circ \psi_s| \\ & \leq \int_0^\tau \int_s^t \det(D\psi_{r+l})w_{r+l} \circ \psi_{r+l} \, dr \, dl, \quad \forall \tau, s + \tau \leq t + \tau \in [0, 1]. \end{aligned} \quad (4.2.15)$$

Remark 4.2.5. One observes that a path $\vartheta \in L^2([0, 1], \mathcal{I})$ with finite energy $\mathcal{F}^{\sigma, \delta, \lambda}[\vartheta] < \infty$ exhibits additional smoothness properties in time. Indeed, from (4.2.14), by using Jensen's inequality, we have

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

where we further used (4.1.4), so that the constant C_v depends on $\|v\|_{L^2((0,1),\mathcal{V})}$. Thus, $\det(D\psi)\vartheta \circ \psi \in C_\omega^0([0,1],\mathcal{I})$, 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

$$\|\vartheta_t\|_{\mathcal{I}} \leq C_v \|\det(D\psi_t)\vartheta_t \circ \psi_t\|_{\mathcal{I}} \leq C_{v,z} (1 + \|\det(D\psi)\vartheta \circ \psi\|_{L^2([0,1],\mathcal{I})}), \quad (4.2.16)$$

so that $\vartheta \in L^\infty([0,1],\mathcal{I})$. 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 $t \mapsto \vartheta_t$, we can take functions $\tilde{\vartheta}^i \in L^\infty([0,1], C^\infty(\Omega, \mathbb{R}^c))$ such that $\|\vartheta - \tilde{\vartheta}^i\|_{L^\infty([0,1],\mathcal{I})} \leq 2^{-i}$ and $\|D\tilde{\vartheta}^i\|_{L^\infty([0,1], C^0(\bar{\Omega}))} \leq C_{\vartheta,i}$, where the constant $C_{\vartheta,i}$ depends on $\|\vartheta\|_{L^\infty([0,1],\mathcal{I})}$ and i (cf. Theorem 2.2.12). We have

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

where the corresponding constants depend on $\|v\|_{L^2((0,1),\mathcal{V})}$, $\|z\|_{L^2((0,1)\times\Omega)}$, $\|\vartheta\|_{L^\infty([0,1],\mathcal{I})}$, and i . Altogether, choosing a suitable i , we see that $\vartheta \in C_\omega^0([0,1],\mathcal{I})$ is uniformly continuous, and we have $\vartheta \in C_\omega^0([0,1],\mathcal{I})$, where the modulus of continuity ω depends on v and z .

In a completely analogous manner, we obtain by using (4.2.15) and Proposition 4.2.3

$$\|\det(D\psi_t)\hat{z}_t \circ \psi_t - \det(D\psi_s)\hat{z}_s \circ \psi_s\|_{\mathcal{I}} \leq C_{v,a} |t-s|^{\frac{1}{2}} \|w\|_{L^2((0,1)\times\Omega)}.$$

Moreover, we also obtain the analogous results $\det(D\psi)\hat{z} \circ \psi \in C_\omega^0([0,1],\mathcal{I})$ and $\hat{z} \in C_\omega^0([0,1],\mathcal{I})$ for a modulus of continuity ω depending on v, a, z, w .

Motivated by [dB63], we now define the continuous time spline interpolation for given key frames.

Definition 4.2.6 (Continuous time regularized spline interpolation). Let $J \geq 2$, $0 \leq \bar{t}_1 < \dots < \bar{t}_J \leq 1$, $\{\vartheta_j^I\}_{j=1,\dots,J} \in \mathcal{I}^J$ and $\delta, \lambda > 0$. We call a minimizer $\vartheta \in L^2((0,1),\mathcal{I})$ of $\mathcal{F}^{\sigma,\delta,\lambda}$ that satisfies (4.2.1) a continuous time regularized spline interpolation of $\{\vartheta_j^I\}_{j=1,\dots,J}$ with a regularization parameter $\sigma > 0$.

Boundary conditions. If we do not impose any additional conditions on top of (4.2.12)-(4.2.15), we say the continuous time spline interpolation has *natural boundary conditions*. In the case of *Hermite boundary conditions* (also called *clamped boundary conditions*) we additionally prescribe the values and tangent vectors of the interpolated curve at $\bar{t}_1 = 0$ and $\bar{t}_J = 1$, i.e. $\vartheta_1^I, \vartheta_J^I \in \mathcal{I}$ and $(\hat{z}_0, v_0) = (\bar{\hat{z}}_0, \bar{v}_0) \in T_{\vartheta_1^I}$ and $(\hat{z}_1, v_1) = (\bar{\hat{z}}_1, \bar{v}_1) \in T_{\vartheta_J^I}$. Note that, by Remark 4.2.5, the point-wise prescription of ϑ, v and z are well defined. Finally, imposing *periodic boundary conditions* is equivalent to defining the density curve ϑ_t on the sphere \mathbb{S}^1 instead of the interval $[0, 1]$.

4.3 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. We consider a discrete density curve $\Theta = (\vartheta_0, \dots, \vartheta_K) \in \mathcal{I}^{K+1}$ and define a set of admissible deformations

$$\mathcal{D} := \{\phi \in H^m(\Omega, \Omega), \det(D\phi) \geq \epsilon, \phi = \mathbb{1} \text{ on } \partial\Omega\}, \quad (4.3.1)$$

for a fixed (small) $\epsilon > 0$, which consists of $C^1(\Omega, \Omega)$ -diffeomorphisms [Cia88, Theorem 5.5-2].

Remark 4.3.1. The case $\epsilon = 0$ is discussed in Remark 4.3.8.

Considering $\Theta \in \mathcal{I}^{K+1}$ as time sampling at times $\frac{k}{K}$, $k = 0, \dots, K$, of a smooth curve $(\vartheta)_{t \in [0,1]}$ and $\Phi = (\phi_1, \dots, \phi_K) \in \mathcal{D}^K$ as relative flow $(\phi_k = \psi_{\frac{k}{K}} \circ \psi_{\frac{k-1}{K}}^{-1})$ and using forward finite difference approximations we obtain the discrete version of the Eulerian velocity $v_k := K(\phi_k - \mathbb{1})$ and $\hat{z}_k := K(\det(D\phi_k)\vartheta_k \circ \phi_k - \vartheta_{k-1})$ for the discrete first order source term. Furthermore, by using central finite differences we define the discrete acceleration

$$a_k := K^2(\phi_{k+1} \circ \phi_k - 2\phi_k + \mathbb{1}) \quad (4.3.2)$$

and

$$\hat{w}_k := K^2(\det(D(\phi_{k+1} \circ \phi_k))\vartheta_{k+1} \circ \phi_{k+1} \circ \phi_k - 2\det(D\phi_k)\vartheta_k \circ \phi_k + \vartheta_{k-1}) \quad (4.3.3)$$

as the discrete version of the second order source term. Following [BER15, EKP⁺21], we consider the discrete path energy

$$\mathbf{E}^{K,D,\delta,\lambda}[\Theta, \Phi] := K \sum_{k=1}^K \int_{\Omega} |\phi_k - \mathbb{1}|^2 \vartheta_{k-1} + \lambda (W_D(D\phi_k) + \gamma |D^m \phi_k|^2) + \frac{K^{-2}}{\delta} |\hat{z}_k|^2 dx,$$

where $W_D(B) := |B^{sym} - \mathbb{1}|^2$ is a simple elastic energy density, and $\delta, \lambda > 0$. Then, the discrete counterpart of the spline energy is defined as

$$\mathbf{F}^{K,D,\delta,\lambda}[\Theta, \Phi] := \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} |a_k|^2 \vartheta_{k-1} + \lambda (W_A(Da_k) + \gamma |D^m a_k|^2) + \frac{1}{\delta} |\hat{w}_k|^2 dx, \quad (4.3.4)$$

with the energy density $W_A(B) := |B^{sym}|^2$. Without loss of generality, and for the sake of clarity, we shall assume from now on that $\delta = \lambda = 1$, and omit both symbols from the indices. We note that

$a_k \in H^m(\Omega, \mathbb{R}^d)$ by [IKT13, Proposition 2.19]. Finally, the regularized discrete spline energy is given by

$$\mathbf{F}^{\sigma, K, D}[\Theta, \Phi] := \mathbf{F}^{K, D}[\Theta, \Phi] + \sigma \mathbf{E}^{K, D}[\Theta, \Phi].$$

As in the continuous time model we have interpolation constraints. 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 $\mathcal{I}_f^K = (\vartheta_1^I, \dots, \vartheta_{J^K}^I)$ and define the set of admissible image vectors

$$\mathcal{I}_{adm}^K := \{\Theta \in \mathcal{I}^{K+1}, \vartheta_{i_j} = \vartheta_j^I, j = 1, \dots, J^K\}. \quad (4.3.5)$$

We are now in a position to define discrete splines.

Definition 4.3.2 (Discrete time regularized spline interpolation). Let $\Theta = (\vartheta_0, \dots, \vartheta_K) \in \mathcal{I}_{adm}^K$. Then we set

$$\mathbf{F}^{\sigma, K}[\Theta] := \inf_{\Phi \in \mathcal{D}^K} \mathbf{F}^{\sigma, K, D}[\Theta, \Phi]. \quad (4.3.6)$$

A discrete time regularized spline interpolation of

$\{\vartheta_j^I\}_{j=1, \dots, J^K}$ is a discrete $(K+1)$ -tuple that minimizes $\mathbf{F}^{\sigma, K}$ over all discrete paths in \mathcal{I}_{adm}^K .

The presented discretization is valid in the case of natural boundary conditions, to which we will restrict ourselves in the sequel. We remark that in the case of periodic boundary conditions we make an identification $K \triangleq 0$, $K+1 \triangleq 1$ and the sum in (4.3.4) goes up to K . For the discrete version of 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 [EKP⁺21] in order to prove the existence of discrete spline interpolations. The following lemma is the analogous result to [EKP⁺21, Lemma 1] and we only state it for completeness.

Lemma 4.3.3. *There exists a constant C which only depends on Ω, m, d and γ , such that*

$$\|\phi - \mathbf{1}\|_{H^m(\Omega)} \leq C \left(\int_{\Omega} W_D(D\phi) + \gamma |D^m \phi|^2 dx \right)^{\frac{1}{2}}$$

for all $\phi \in \mathcal{D}$.

Proof. An application of the Gagliardo–Nirenberg inequality [Nir66] yields

$$\|\phi - \mathbf{1}\|_{H^m(\Omega)} \leq C(\|\phi - \mathbf{1}\|_{L^2(\Omega)} + |\phi - \mathbf{1}|_{H^m(\Omega)}). \quad (4.3.7)$$

The last term in (4.3.7) is bounded by

$$|\phi - \mathbf{1}|_{H^m(\Omega)} = |\phi|_{H^m(\Omega)} \leq \sqrt{\frac{C_\phi}{\gamma}}. \quad (4.3.8)$$

To estimate the lower order term on the right hand side we use Korn's and Poincaré's inequality and write

$$\|\phi - \mathbf{1}\|_{L^2(\Omega)} \leq C\|\varepsilon[\phi] - \mathbf{1}\|_{L^2(\Omega)} \leq C\sqrt{C_\phi}. \quad (4.3.9)$$

Thus, the lemma follows by combining (4.3.7), (4.3.8) and (4.3.9). \square

The analogous result holds for the boundedness of acceleration:

Lemma 4.3.4. *There exists a constant C which only depends on Ω, m, d and γ , such that*

$$\|a\|_{H^m(\Omega)} \leq C \left(\int_{\Omega} W_A(Da) + \gamma |D^m a|^2 dx \right)^{\frac{1}{2}}$$

for all $\phi \in \mathcal{D}$.

Proof. An application of the Gagliardo–Nirenberg inequality [Nir66] for bounded domains yields

$$\|a\|_{H^m(\Omega)} \leq C(\|a\|_{L^2(\Omega)} + |a|_{H^m(\Omega)}). \quad (4.3.10)$$

The last term in (4.3.10) is bounded by

$$|a|_{H^m(\Omega)} \leq \sqrt{\frac{C_a}{\gamma}}, \quad (4.3.11)$$

with $C_a := \int_{\Omega} W_A(Da) + \gamma |D^m a|^2 dx$. To estimate the lower order term on the right hand side we use Korn's and Poincaré's inequality and write

$$\|a\|_{L^2(\Omega)} \leq C\|\varepsilon[a]\|_{L^2(\Omega)} \leq C\sqrt{C_a}. \quad (4.3.12)$$

Thus, the lemma follows by combining (4.3.10), (4.3.11) and (4.3.12). \square

Now, we show the well-posedness of (4.3.6).

Proposition 4.3.5. *For every $K \in \mathbb{N}$ and every image vector $\Theta = (\vartheta_0, \dots, \vartheta_K) \in \mathcal{I}_{adm}^K$ there exists a deformation vector $\Phi = (\phi_1, \dots, \phi_K) \in \mathcal{D}^K$ such that*

$$\mathbf{F}^{\sigma, K, D}[\Theta, \Phi] = \inf_{\tilde{\Phi} \in \mathcal{D}^K} \mathbf{F}^{\sigma, K, D}[\Theta, \tilde{\Phi}].$$

Proof. Let $\{\Phi^j\}_{j \in \mathbb{N}} \subset \mathcal{D}^K$ be a sequence for which it holds $\mathbf{F}^{\sigma, K, D}[\Theta, \Phi^j] \leq \overline{\mathbf{F}^{\sigma, K}} := \mathbf{F}^{\sigma, K, D}[\Theta, \mathbb{1}^K]$ and $\lim_{j \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\Theta, \Phi^j] = \inf_{\tilde{\Phi} \in \mathcal{D}^K} \mathbf{F}^{\sigma, K, D}[\Theta, \tilde{\Phi}]$. By Lemma 4.3.3 we have

$$\|\phi_k^j - \mathbb{1}\|_{H^m(\Omega)} \leq C \left(\int_{\Omega} W_D(D\phi_k^j) + \gamma |D^m \phi_k^j|^2 dx \right)^{\frac{1}{2}} \leq C \sqrt{\frac{\mathbf{F}^{\sigma, K}}{\sigma K}}, \quad \forall j \in \mathbb{N}, k = 1, \dots, K,$$

where the constant C only depends on Ω, m, d, λ and γ . Thus, $\{\Phi^j\}_{j \in \mathbb{N}}$ is uniformly bounded in $H^m(\Omega, \Omega)^K$. Due to the reflexivity of this space, there exists a weakly convergent subsequence (for the sake of clarity, we will assume it has the same label) such that $\Phi^j \rightharpoonup \Phi$ in $H^m(\Omega, \Omega)^K$. By the compact Sobolev embedding, we have $\Phi^j \rightarrow \Phi$ in $C^{1, \alpha}(\overline{\Omega}, \overline{\Omega})^K$ for $\alpha \in (0, m - 1 - \frac{d}{2})$, which gives us that $\Phi \in \mathcal{D}^K$.

Analogously, we recall that by [IKT13, Proposition 2.19.], \mathbf{a}^j is in $H^m(\Omega, \Omega)^{K-1}$ for all $j \in \mathbb{N}$. By Lemma 4.3.4, we obtain the boundedness of $\{\mathbf{a}^j\}_{j \in \mathbb{N}}$ in $H^m(\Omega, \Omega)^{K-1}$, and similarly as before, a weakly convergent subsequence satisfying $\mathbf{a}^j \rightharpoonup \mathbf{a}$ 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 (4.3.2). 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 the weak lower semi-continuity of the H^m -seminorm and the continuity of energy densities we have for all k , as $j \rightarrow \infty$:

$$\liminf |\phi_k^j|_{H^m} \geq |\phi_k|_{H^m}, \quad \liminf |a_k^j|_{H^m} \geq |a_k|_{H^m}, \quad (4.3.13)$$

$$\|W_D(D\phi_k^j)\|_{L^1(\Omega)} \rightarrow \|W_D(D\phi_k^j)\|_{L^1(\Omega)}, \quad \|W_A(Da_k^j)\|_{L^1(\Omega)} \rightarrow \|W_A(Da_k)\|_{L^1(\Omega)}, \quad (4.3.14)$$

$$\int_{\Omega} |\phi_k^j - \mathbb{1}|^2 \vartheta_k \, dx \rightarrow \int_{\Omega} |\phi_k - \mathbb{1}|^2 \vartheta_k \, dx, \quad \int_{\Omega} |a_k^j|^2 \vartheta_k \, dx \rightarrow \int_{\Omega} |a_k|^2 \vartheta_k \, dx \quad (4.3.15)$$

To handle the rest of the terms we show that for all $\vartheta \in \mathcal{I}$ we have $\det(D\phi^j)\vartheta \circ \phi^j \rightarrow \det(D\phi)\vartheta \circ \phi$ in \mathcal{I} for $\{\phi^j\}_{j \in \mathbb{N}} \in \mathcal{D}$ with $\phi^j \rightarrow \phi$ in $C^1(\overline{\Omega}, \overline{\Omega})$ as $j \rightarrow \infty$. To see this, we approximate ϑ by smooth functions $\{\tilde{\vartheta}^i\}_{i \in \mathbb{N}} \in C^\infty(\overline{\Omega}, \overline{\Omega})$ with $\|\vartheta - \tilde{\vartheta}^i\|_{\mathcal{I}} \rightarrow 0$. Then, using the transformation formula we obtain

$$\begin{aligned} & \|\det(D\phi^j)\vartheta \circ \phi^j - \det(D\phi)\vartheta \circ \phi\|_{\mathcal{I}} \\ & \leq \|\det(D\phi^j)(\vartheta \circ \phi^j - \tilde{\vartheta}^i \circ \phi^j)\|_{\mathcal{I}} + \|\det(D\phi)(\tilde{\vartheta}^i \circ \phi - \vartheta \circ \phi)\|_{\mathcal{I}} \\ & \quad + \|\det(D\phi^j)\tilde{\vartheta}^i \circ \phi^j - \det(D\phi)\tilde{\vartheta}^i \circ \phi\|_{\mathcal{I}} \\ & \leq \|\vartheta - \tilde{\vartheta}^i\|_{\mathcal{I}} \left(\|\det(D\phi^j)\|_{L^\infty(\Omega)}^{\frac{1}{2}} + \|\det(D\phi)\|_{L^\infty(\Omega)}^{\frac{1}{2}} \right) + \|\det(D\phi^j)\tilde{\vartheta}^i \circ \phi^j - \det(D\phi)\tilde{\vartheta}^i \circ \phi\|_{\mathcal{I}} \\ & \quad + \|\det(D\phi^j)\tilde{\vartheta}^i \circ \phi - \det(D\phi)\tilde{\vartheta}^i \circ \phi\|_{\mathcal{I}} \\ & \leq \|\vartheta - \tilde{\vartheta}^i\|_{\mathcal{I}} \left(\|\det(D\phi^j)\|_{L^\infty(\Omega)}^{\frac{1}{2}} \right. \\ & \quad \left. + \|\det(D\phi)\|_{L^\infty(\Omega)}^{\frac{1}{2}} \right) + \|D\tilde{\vartheta}^i\|_{L^\infty(\Omega)} \|\det(D\phi^j)\|_{L^\infty(\Omega)} \|\phi^j - \phi\|_{L^2(\Omega)} \\ & \quad + \|\det(D\phi^j) - \det(D\phi)\|_{L^\infty(\Omega)} \|\det(D(\phi^{-1}))\|_{L^\infty(\Omega)}^{\frac{1}{2}} \|\tilde{\vartheta}^i\|_{\mathcal{I}}. \end{aligned} \quad (4.3.16)$$

By first fixing $\tilde{\vartheta}^i$ and then choosing a suitable j , we have that this expression converges to 0. Hence, for every $k = 1, \dots, K$ we have $\|\det(D\phi_k)\vartheta_k \circ \phi_k - \vartheta_{k-1}\|_{\mathcal{I}} = \lim_{j \rightarrow \infty} \|\det(D\phi_k^j)\vartheta_k \circ \phi_k^j - \vartheta_{k-1}\|_{\mathcal{I}}$. Furthermore, since $\phi_{k+1}^j \circ \phi_k^j \rightarrow \phi_{k+1} \circ \phi_k$ strongly in $C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})$, we can prove the analogous estimate $\|\det(D(\phi_{k+1} \circ \phi_k))\vartheta_{k+1} \circ \phi_{k+1} \circ \phi_k - 2\det(D\phi_k)\vartheta_k \circ \phi_k + \vartheta_{k-1}\|_{\mathcal{I}} = \lim_{j \rightarrow \infty} \|\det(D(\phi_{k+1}^j \circ \phi_k^j))\vartheta_{k+1} \circ \phi_{k+1}^j \circ \phi_k^j - 2\det(D\phi_k^j)\vartheta_k \circ \phi_k^j + \vartheta_{k-1}\|_{\mathcal{I}}$ for $k = 1, \dots, K-1$, which together with (4.3.13)-(4.3.15) finishes the proof. \square

In the next step, under suitable conditions, we prove that there exists a minimizing vector in \mathcal{I}_{adm}^K (see (4.3.5)) for a fixed deformation vector $\Phi \in \mathcal{D}^K$.

Proposition 4.3.6. *Let $K \geq 2$, \mathcal{I}_f^K and $\Phi \in \mathcal{D}^K$ be fixed. Assume that the deformations satisfy, for every $x \in \Omega$,*

$$C_{\det} \geq \det(D\phi_k(x)) \geq c_{\det} > 0, \quad k = 1, \dots, K. \quad (4.3.17)$$

Then there exists a vector of images $\Theta \in \mathcal{I}_{adm}^K$ such that

$$\mathbf{F}^{\sigma, K, D}[\Theta, \Phi] = \inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K, D}[\mathbf{v}, \Phi].$$

Proof. Let $\{\Theta^j\}_{j \in \mathbb{N}} \in \mathcal{I}_{adm}^K$ be an approximation sequence such that

$$\lim_{j \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\Theta^j, \Phi] = \inf_{\mathbf{v} \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K, D}[\mathbf{v}, \Phi] \leq \overline{\mathbf{F}^{\sigma, K, D}}.$$

Here, $\overline{\mathbf{F}^{\sigma, K, D}} := \mathbf{F}^{\sigma, K, D}[\bar{\vartheta}, \Phi] < \infty$ represents a finite upper bound for the energy with the vector of images $\bar{\vartheta}$ satisfying $\bar{\vartheta}_k = \vartheta_1^I$ for $0 \leq k \leq i_1$, $\bar{\vartheta}_k = \vartheta_j^I$ for $i_j \leq k \leq i_{j+1}$ with $1 \leq j \leq J^K - 1$ and

$\bar{\vartheta}_k = \vartheta_{j_K}^I$ for $i_{j_K} \leq k \leq K$.

$$\begin{aligned}
\|\vartheta_k^j\|_{\mathcal{I}} &\leq \|\det(D\phi_{k+1})\vartheta_{k+1}^j \circ \phi_{k+1} - \vartheta_k^j\|_{\mathcal{I}} + \|\det(D\phi_{k+1})\vartheta_{k+1}^j \circ \phi_{k+1}\|_{\mathcal{I}} \\
&\leq \sqrt{\frac{\delta \mathbf{F}^{\sigma, K, D}}{K}} + C_{\det}^{\frac{1}{2}} \|\vartheta_{k+1}^j\|_{\mathcal{I}}, \\
\|\vartheta_{k+1}^j\|_{\mathcal{I}} &\leq c_{\det}^{-\frac{1}{2}} \|\det(D\phi_{k+1})\vartheta_{k+1}^j \circ \phi_{k+1}\|_{\mathcal{I}} \leq c_{\det}^{-\frac{1}{2}} \left(\|\det(D\phi_{k+1})\vartheta_{k+1}^j \circ \phi_{k+1} - \vartheta_k^j\|_{\mathcal{I}} + \|\vartheta_k^j\|_{\mathcal{I}} \right) \\
&\leq c_{\det}^{-\frac{1}{2}} \left(\sqrt{\frac{\delta \mathbf{F}^{\sigma, K, D}}{K}} + \|\vartheta_k^j\|_{\mathcal{I}} \right),
\end{aligned} \tag{4.3.18}$$

from where we have, by induction, that $\{\vartheta_k^j\}_{j \in \mathbb{N}}$ is uniformly bounded in \mathcal{I} for every $k = 0, \dots, K$. By the reflexivity of \mathcal{I} , there exists a subsequence (not relabeled) such that $\vartheta_k^j \rightharpoonup \vartheta_k$ for some $\Theta \in \mathcal{I}_{adm}^K$. It remains to verify the weak lower semi-continuity of the matching functional, i.e. we have to show that

$$\|\det(D\phi_k)\vartheta_k \circ \phi_k - \vartheta_{k-1}\|_{\mathcal{I}}^2 \leq \liminf_{j \rightarrow \infty} \|\det(D\phi_k)\vartheta_k^j \circ \phi_k - \vartheta_{k-1}^j\|_{\mathcal{I}}^2, \tag{4.3.19}$$

$$\begin{aligned}
&\|\det(D(\phi_{k+1} \circ \phi_k))\vartheta_{k+1} \circ \phi_{k+1} \circ \phi_k - 2\det(D\phi_k)\vartheta_k \circ \phi_k + \vartheta_{k-1}\|_{\mathcal{I}}^2 \\
&\leq \liminf_{j \rightarrow \infty} \|\det(D(\phi_{k+1} \circ \phi_k))\vartheta_{k+1}^j \circ \phi_{k+1} \circ \phi_k - 2\det(D\phi_k)\vartheta_k^j \circ \phi_k + \vartheta_{k-1}^j\|_{\mathcal{I}}^2,
\end{aligned} \tag{4.3.20}$$

for every $k = 1, \dots, K$ and $k = 1, \dots, K-1$, respectively. To this end, we will first show $\det(D\phi_k)\vartheta_k^j \circ \phi_k \rightharpoonup \det(D\phi_k)\vartheta_k \circ \phi_k$ in \mathcal{I} . Indeed, for every $\eta \in \mathcal{I}$, the transformation formula yields

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

which converges to 0 since $\eta \circ \phi_k^{-1} \in \mathcal{I}$ due to (4.3.17) and the weak convergence $\vartheta_k^j \rightharpoonup \vartheta_k$. Hence, $\vartheta_k^j \circ \phi_k - \vartheta_{k-1}^j \rightharpoonup \vartheta_k \circ \phi_k - \vartheta_{k-1}$ in \mathcal{I} , which readily implies (4.3.19) and by applying the same technique in a nested fashion we get (4.3.20). Since $|\phi_k - \mathbb{1}|^2, |a_k|^2 \in L^\infty(\Omega) \subset \mathcal{I}$ for all $k = 0, \dots, K-1$, we also obtain

$$\int_{\Omega} |\phi_k - \mathbb{1}|^2 \vartheta_{k-1}^j \, dx \rightarrow \int_{\Omega} |\phi_k - \mathbb{1}|^2 \vartheta_{k-1} \, dx, \quad \int_{\Omega} |a_k|^2 \vartheta_{k-1}^j \, dx \rightarrow \int_{\Omega} |a_k|^2 \vartheta_{k-1} \, dx.$$

Altogether, we have

$$\mathbf{F}^{\sigma, K, D}[\Theta, \Phi] \leq \liminf_{j \rightarrow \infty} \mathbf{F}^{\sigma, K, D}[\Theta^j, \Phi],$$

from where the optimality follows. \square

We are now in position to show the existence of discrete time spline interpolations.

Theorem 4.3.7 (Existence of discrete time spline interpolations). *Let $K \geq 2$. Then for every \mathcal{I}_f^K there exists $\Theta \in \mathcal{I}_{adm}^K$ such that*

$$\mathbf{F}^{\sigma, K}[\Theta] = \inf_{\Theta \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma, K}[\bar{\Theta}].$$

Proof. Consider a sequence $\{\Theta^j\}_{j \in \mathbb{N}} \in \mathcal{I}_{adm}^K$ for which it holds $\lim_{j \rightarrow \infty} \mathbf{F}^{\sigma,K}[\Theta^j] = \inf_{\Theta \in \mathcal{I}_{adm}^K} \mathbf{F}^{\sigma,K}[\Theta] \leq \overline{\mathbf{F}^{\sigma,K}}$, where $\overline{\mathbf{F}^{\sigma,K}} := \mathbf{F}^{\sigma,K,D}[\overline{\Theta}, \mathbb{1}^K]$. Here, $\overline{\vartheta}_k = \overline{\vartheta}(\frac{k}{K}, \cdot)$ for $k = 0, \dots, K$, where $\overline{\vartheta}$ is a smooth in time curve with $\overline{\vartheta}(\frac{i_j}{K}, \cdot) = u_j^I$ for every $j = 1, \dots, J^K$. Then, we have

$$\begin{aligned} \mathbf{F}^{\sigma,K,D}[\overline{\Theta}^K, \mathbb{1}^K] &= \sigma K \sum_{k=0}^K \int_{\Omega} |\overline{\vartheta}_{k+1}^K - \overline{\vartheta}_k^K|^2 dx + K^3 \sum_{k=1}^{K-1} \int_{\Omega} |\overline{\vartheta}_{k+1}^K - 2\overline{\vartheta}_k^K + \overline{\vartheta}_{k-1}^K|^2 dx \\ &\leq C \left(\int_{\Omega} |\overline{\vartheta}|_{H^1((0,1))}^2 + |\overline{\vartheta}|_{H^2((0,1))}^2 dx + 1 \right), \end{aligned} \quad (4.3.21)$$

where the constant C is independent of K . Furthermore, for every j let $\mathbf{F}^{\sigma,K}[\Theta^j] = \mathbf{F}^{\sigma,K,D}[\Theta^j, \Phi^j]$, where $\Phi^j \in H^m(\Omega, \Omega)$ are the optimal deformation vectors guaranteed to exist by Proposition 4.3.5. Using the above uniform bound, we may proceed as in the proof of Proposition 4.3.5 to obtain weak convergence of $\{\Phi^j\}_{j \in \mathbb{N}}$ in $H^m(\Omega, \Omega)^K$ and strong convergence in $C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})^K$ to some $\Phi \in \mathcal{D}^K$. Furthermore, we once again have $a_k^j \rightharpoonup a_k$ in $H^m(\Omega, \Omega)$, and strongly in $C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})$, where $a_k = K^2(\phi_{k+1} \circ \phi_k - 2\phi_k + \mathbb{1})$ and estimates from (4.3.13) and (4.3.14) are satisfied.

By Proposition 4.3.6 we may replace Θ^j by an energy optimal image vector (which would not increase the spline energy) for every $j \in \mathbb{N}$. We assume without loss of generality that Θ^j are already the optimal image density vectors, and following the same arguments as in (4.3.18) we conclude that $\{\Theta^j\}_{j \in \mathbb{N}}$ is uniformly bounded in \mathcal{I} , containing a weakly converging sequence (not relabeled) that converges to some $\Theta \in \mathcal{I}$. By the weak convergence $\Theta^j \rightharpoonup \Theta$ and the strong convergence $\Phi^j \rightarrow \Phi$, we obtain the estimates in (4.3.15). Finally, we verify the estimates

$$\begin{aligned} \|\det(D\phi_k)\vartheta_k \circ \phi_k - \vartheta_{k-1}\|_{\mathcal{I}}^2 &\leq \liminf_{j \rightarrow \infty} \|\det(D\phi_k^j)\vartheta_k^j \circ \phi_k^j - \vartheta_{k-1}^j\|_{\mathcal{I}}^2, \\ \|\det(D(\phi_{k+1} \circ \phi_k))\vartheta_{k+1} \circ \phi_{k+1} \circ \phi_k - 2\det(D\phi_k)\vartheta_k \circ \phi_k + \vartheta_{k-1}\|_{\mathcal{I}}^2 \\ &\leq \liminf_{j \rightarrow \infty} \|\det(D(\phi_{k+1}^j \circ \phi_k^j))\vartheta_{k+1}^j \circ \phi_{k+1}^j \circ \phi_k^j - 2\det(D\phi_k^j)\vartheta_k^j \circ \phi_k^j + \vartheta_{k-1}^j\|_{\mathcal{I}}^2, \end{aligned}$$

for every $k = 1, \dots, K$ and $k = 1, \dots, K-1$, respectively. To that end, it is enough to show $\det(D\phi_k^j)\vartheta_k^j \circ \phi_k^j \rightharpoonup \det(D\phi_k)\vartheta_k \circ \phi_k$ and $\det(D(\phi_{k+1}^j \circ \phi_k^j))\vartheta_{k+1}^j \circ \phi_{k+1}^j \circ \phi_k^j \rightharpoonup \det(D(\phi_{k+1} \circ \phi_k))\vartheta_{k+1} \circ \phi_{k+1} \circ \phi_k$ in \mathcal{I} . To see this, we first take into account the decomposition

$$\det(D\phi_k^j)\vartheta_k^j \circ \phi_k^j - \det(D\phi_k)\vartheta_k \circ \phi_k = \det(D\phi_k^j)(\vartheta_k^j \circ \phi_k^j - \vartheta_k \circ \phi_k) + (\det(D\phi_k^j)\vartheta_k \circ \phi_k^j - \det(D\phi_k)\vartheta_k \circ \phi_k). \quad (4.3.22)$$

The second term is handled as in (4.3.16). It remains to consider the convergence properties of the first term. For a test function $\eta \in \mathcal{I}$ we obtain using the transformation rule

$$\int_{\Omega} \det(D\phi_k^j)(\vartheta_k^j \circ \phi_k^j - \vartheta_k \circ \phi_k) \cdot \eta dx = \int_{\Omega} (\vartheta_k^j - \vartheta_k) \cdot \eta((\phi_k^j)^{-1}(x)) dx.$$

The right hand side converges to 0 due to the strong convergence $\eta \circ (\phi_k^j)^{-1} \rightarrow \eta \circ (\phi_k)^{-1}$ in \mathcal{I} and $\vartheta_k^j \rightharpoonup \vartheta_k$ in \mathcal{I} for $j \rightarrow \infty$. Finally, one reiterates the above argument replacing ϑ_k^j by $\det(D\phi_{k+1}^j)\vartheta_{k+1}^j \circ \phi_{k+1}^j$, which yields $\det(D(\phi_{k+1}^j \circ \phi_k^j))\vartheta_{k+1}^j \circ \phi_{k+1}^j \circ \phi_k^j \rightharpoonup \det(D(\phi_{k+1} \circ \phi_k))\vartheta_{k+1} \circ \phi_{k+1} \circ \phi_k$ in \mathcal{I} . This proves our claim and finally proves the theorem. \square

Remark 4.3.8. The results of this section remain valid for any W_D satisfying conditions (W1) – (W2) from [EKP⁺21]. Furthermore, let us observe the case when $i_j = K \cdot t_j$ and $\vartheta_{i_j} = u_j^I$ for $t_j \in [0, 1]$ and

$\vartheta_j^I \in \mathcal{I}$ for every $j = 1, \dots, J$ (cf. (4.2.1) and (4.3.5)). Then for every large enough K (depending on u_j^I and t_j) one can show the existence of the discrete time spline interpolation even if $\epsilon = 0$ in the definition of the admissible set (4.3.1). Indeed, by (4.3.21) we have that $\overline{\mathbf{F}^{\sigma, K}}$ is a fixed finite upper bound for the discrete spline energy, independent of K . Then, using Lemma 4.3.3 and the Sobolev embedding theorem we have

$$\max_{k=1, \dots, K} \|\phi_k - \mathbb{1}\|_{C^{1, \alpha}(\overline{\Omega})} \leq C \max_{k=1, \dots, K} \|\phi_k - \mathbb{1}\|_{H^m(\Omega)} \leq C \sqrt{\frac{\mathbf{F}^{\sigma, K}}{K}}.$$

By the Lipschitz continuity of the determinant we have

$$\max_{k=1, \dots, K} \|\det(D\phi_k) - 1\|_{L^\infty(\Omega)} \leq C \sqrt{\frac{\mathbf{F}^{\sigma, K}}{K}} < 1,$$

for large enough K , which proves $\min_{k=1, \dots, K} \det(D\phi_k) \geq c_{\det} > 0$ and we can proceed as before.

4.4 Temporal Extension Operators

In this section, we define the suitable time extensions of the time discrete quantities from the previous section 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 $\Theta^K = (\vartheta_0^K, \dots, \vartheta_K^K) \in \mathcal{I}^{K+1}$ 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 $k = 1, \dots, K-1$

$$y_t^K(\cdot) := y_k^K(t, \cdot) := \mathbb{1} + \left(\frac{1}{2} + \frac{t - t_{k-\frac{1}{2}}^K}{\tau} \right) (\phi_k^K - \mathbb{1}) + \frac{(t - t_{k-\frac{1}{2}}^K)^2}{2\tau^2} (\phi_{k+1}^K \circ \phi_k^K - 2\phi_k^K + \mathbb{1}). \quad (4.4.1)$$

This can be seen as a cubic Hermite interpolation (see [de78, Chapter IV] for details) 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 4.1 left.

Next, we define the image extension operator $\theta^K[\Theta^K, \Phi^K] \in L^2([0, 1], \mathcal{I})$ as $\theta^K[\Theta^K, \Phi^K](t, x) = \vartheta^K(t, x)$, where

$$\det(Dy_t^K) \vartheta_t^K \circ y_t^K := \vartheta_0^K + \frac{t}{\tau} (\det(D\phi_1^K) \vartheta_1^K \circ \phi_1^K - \vartheta_0^K), \quad t \in [0, t_{\frac{1}{2}}^K],$$

$$\det(Dy_t^K)\vartheta_t^K \circ y_t^K := \vartheta_{K-1}^K + \frac{t - t_{K-\frac{1}{2}}^K}{\tau} (\det(D\phi_K^K)\vartheta_K^K \circ \phi_K^K - \vartheta_{K-1}^K), \quad t \in (t_{K-\frac{1}{2}}^K, 1],$$

and, for $k = 1, \dots, K-1$ and $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$

$$\begin{aligned} \det(Dy_t^K)\vartheta_t^K \circ y_t^K := & \vartheta_{k-1}^K + \left(\frac{1}{2} + \frac{t - t_{k-\frac{1}{2}}^K}{\tau} \right) (\det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K - \vartheta_{k-1}^K) \\ & + \frac{(t - t_{k-\frac{1}{2}}^K)^2}{2\tau^2} (\det(D(\phi_{k+1}^K \circ \phi_k^K))\vartheta_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K \\ & - 2\det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K + \vartheta_{k-1}^K). \end{aligned} \quad (4.4.2)$$

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 4.5 (see (4.5.3)).

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

$$a_t^K := \frac{1}{\tau^2} (\phi_{k+1}^K \circ \phi_k^K - 2\phi_k^K + \mathbb{1}).$$

We define the velocity and the acceleration along the incremental transport path by

$$\tilde{v}_t^K := v_t^K \circ x_t^K, \quad \tilde{a}_t^K := a_t^K \circ x_t^K. \quad (4.4.3)$$

Now, the *discrete flow* given as the map $(t, x) \mapsto \psi_t^K(x)$ is defined recursively by

$$\psi_t^K := y_t^K, \quad t \in [0, t_{\frac{1}{2}}^K], \quad \psi_t^K := y_t^K \circ \psi_{t_{k-\frac{1}{2}}^K}^K, \quad t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K], \quad \psi_t^K := y_t^K \circ \psi_{t_{K-\frac{1}{2}}^K}^K, \quad t \in (t_{K-\frac{1}{2}}^K, t_K^K]. \quad (4.4.4)$$

Following the same arguments as in [JRR23, Section 5], one shows that (4.4.4) is well-defined in the sense of equations (4.2.12) – (4.2.13), i.e., for every $x \in \Omega$ and $t \in [0, 1]$ we have

$$\dot{\psi}_t^K = \tilde{v}_t^K \circ \psi_t^K, \quad \psi_0^K(x) = x \quad (4.4.5)$$

$$\ddot{\psi}_t^K = \tilde{a}_t^K \circ \psi_t^K. \quad (4.4.6)$$

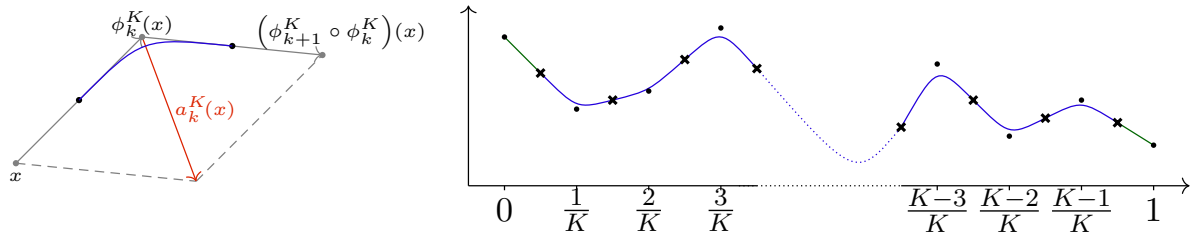


Figure 4.1: Left: Schematic drawing of the Hermite interpolation $y_k^K(x)$ (blue) on the time interval $[(k - \frac{1}{2})/K, (k + \frac{1}{2})/K]$ together with the discrete acceleration $a_k^K(x)$ (red). Right: Image extension $\theta^K[\Theta^K, \Phi^K](\cdot, x)$ along a path $(y_t^K(x))_{t \in [0,1]}$, plotted against time. Dots represent the values ϑ_k^K , $k = 0, \dots, K$, and crosses the “half-way” values $\frac{1}{2}(\vartheta_k^K + \vartheta_{k-1}^K)$, $k = 1, \dots, K$, along the discrete transport path.

Based on this, the *first order scalar weak source term* of $(\vartheta_t^K)_{t \in [0,1]}$ can be defined as the absolute value of the time derivative along the paths $t \mapsto \psi_t^K(x)$ with

$$\begin{aligned} \det(Dy_k^K)z_t^K \circ y_t^K &:= \frac{1}{\tau} |\det(D\phi_1^K)\vartheta_1^K \circ \phi_1^K - \vartheta_0^K|, \quad t \in [0, t_{\frac{1}{2}}^K], \\ \det(Dy_k^K)z_t^K \circ y_t^K &:= \frac{1}{\tau} |\det(D\phi_K^K)\vartheta_K^K \circ \phi_K^K - \vartheta_{K-1}^K|, \quad t \in (t_{K-\frac{1}{2}}^K, 1], \end{aligned}$$

and, for $t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$, $k = 1, \dots, K-1$,

$$\det(Dy_t^K)z_t^K \circ y_t^K := \left| \frac{1}{\tau} (\det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K - \vartheta_{k-1}^K) \right. \quad (4.4.7)$$

$$\left. + \frac{t - t_{k-\frac{1}{2}}^K}{\tau^2} (\det(D(\phi_{k+1}^K \phi_k^K))\vartheta_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K - 2\det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K + \vartheta_{k-1}^K) \right|. \quad (4.4.8)$$

$$\left. - 2\det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K + \vartheta_{k-1}^K \right|. \quad (4.4.9)$$

For the *second order scalar weak source term*, we have, for $t \in [t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K]$, $k = 1, \dots, K-1$,

$$\det(Dy_t^K)w_t^K \circ y_t^K := \frac{1}{\tau^2} |\det(D(\phi_{k+1}^K \circ \phi_k^K))\vartheta_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K - 2\det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K + \vartheta_{k-1}^K|, \quad (4.4.10)$$

and $w_t^K := 0$ elsewhere, which is the absolute value of the second time derivative of $\vartheta^K = \theta^K[\Theta^K, \Phi^K]$ along the path $t \mapsto \psi_t^K(x)$. Indeed, as in [ENR20, Proposition 9], 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 (4.2.14) and (4.2.15), i.e.,

$$\begin{aligned} |\det(D\psi_t^K)\vartheta_t^K \circ \psi_t^K - \det(D\psi_s^K)\vartheta_s^K \circ \psi_s^K| &\leq \int_s^t \det(D\psi_r^K)z_r^K \circ \psi_r^K \, dr, \quad \forall s \leq t \in [0, 1], \\ |\det(D\psi_{t+\tau}^K)\vartheta_{t+\tau}^K \circ \psi_{t+\tau}^K - \det(D\psi_t^K)\vartheta_t^K \circ \psi_t^K - \det(D\psi_{s+\tau}^K)\vartheta_{s+\tau}^K \circ \psi_{s+\tau}^K + \det(D\psi_s^K)\vartheta_s^K \circ \psi_s^K| \\ &\leq \int_0^\tau \int_s^{s+\tau} \det(D\psi_{r+l}^K)w_{r+l}^K \circ \psi_{r+l}^K \, dl \, dr, \quad \forall \tau, s + \tau \leq t + \tau \in [0, 1]. \end{aligned}$$

Let us finally notice that these quantities, together with the image extension operator (4.4.2), can be defined in an explicit manner by using the inverse transport path $(x_t^K)_{t \in [0,1]}$.

Remark 4.4.1. For periodic boundary conditions, 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 $K \triangleq 0, K+1 \triangleq 1$ 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

$$\mathcal{F}^{\sigma,K}[\vartheta] := \inf_{\Theta^K \in \mathcal{I}^{K+1}} \inf_{\Phi^K \in \mathcal{D}^K} \{\mathbf{F}^{\sigma,K,D}[\Theta^K, \Phi^K] : \mathcal{U}^K[\Theta^K, \Phi^K] = \vartheta\}, \quad (4.4.11)$$

if there exists such Θ^K, Φ^K and $+\infty$ otherwise.

Lemma 4.4.2. *If for a fixed image path $\vartheta \in L^2([0,1], \mathcal{I})$ there exist a vector of images $\tilde{\Theta}^K \in \mathcal{I}^{K+1}$ and a vector of deformations $\tilde{\Phi}^K \in \mathcal{D}^K$ such that $\theta^K[\tilde{\Theta}^K, \tilde{\Phi}^K] = \vartheta$, the infimum in (4.4.11) is attained for some $(\Theta^K, \Phi^K) \in \mathcal{I}^{K+1} \times \mathcal{D}^K$.*

Proof. The existence of an infimum follows from the continuity of the constraint $\theta^K[\Theta_j^K, \Phi^K] = \vartheta$ with respect to the weak convergence of Θ^K and strong convergence of Φ^K . □

4.5 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 4.5.1 (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 $\{\vartheta^K\}_{K \in \mathbb{N}} \subset L^2((0,1), \mathcal{I})$ which converges weakly to $\vartheta \in L^2((0,1), \mathcal{I})$ as $K \rightarrow \infty$, it holds $\liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma,K}[\vartheta^K] \geq \mathcal{F}^\sigma[\vartheta]$ (“lim inf-inequality”),*
- (ii) *for every $\vartheta \in L^2((0,1), \mathcal{I})$, there exists a sequence $\{\vartheta^K\}_{K \in \mathbb{N}}$ such that $\vartheta^K \rightarrow \vartheta$ in $L^2((0,1), \mathcal{I})$ as $K \rightarrow \infty$ and $\limsup_{K \rightarrow \infty} \mathcal{F}^{\sigma,K}[\vartheta^K] \leq \mathcal{F}^\sigma[\vartheta]$ (“the existence of a recovery sequence and lim sup-inequality”).*

Remark 4.5.2. The above result holds for any choice of W_D satisfying assumption (W1)–(W3) from [EKP⁺21], though we restrict ourselves to the special case $W_D = |A^{\text{sym}} - \mathbb{1}|^2$.

Remark 4.5.3. In the case of periodic boundary conditions the same applies in the topology $L^2(S^1, \mathcal{I})$. The proof only requires minor alterations.

Proof. We will prove the conditions separately and split the proofs into respective steps.

Proof of the \liminf -inequality

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

Suppose we have a sequence $\{\vartheta^K\}_{K \in \mathbb{N}} \subset L^2((0, 1), \mathcal{I})$ such that $\vartheta^K \rightharpoonup \vartheta$ in that space, as $K \rightarrow \infty$. To avoid the trivial case of an infinite limit, we suppose that $\mathcal{F}^{\sigma, K}[\vartheta^K] < \overline{\mathcal{F}} < \infty$. Then, by definition (4.4.11) and Lemma 4.4.2, for every K large enough, an optimal vector of images $\Theta^K \in \mathcal{I}^{K+1}$ and a corresponding optimal vector of deformations $\Phi^K \in \mathcal{D}^K$ exist, such that

$$\vartheta^K = \theta^K[\Theta^K, \Phi^K], \quad \mathcal{F}^{\sigma, K}[\vartheta^K] = \mathbf{F}^{\sigma, K, D}[\Theta^K, \Phi^K].$$

For the vectors of images Θ^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 source term of the first and the second order, \hat{z}_k^K and \hat{w}_k^K , respectively, as in the previous section. Using Lemmas 4.3.3 and 4.3.4, we obtain

$$\max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^{1, \alpha}(\overline{\Omega})} \leq CK^{-\frac{1}{2}}, \quad (4.5.1)$$

$$\max_{k=1, \dots, K-1} \|a_k^K\|_{C^{1, \alpha}(\overline{\Omega})} \leq CK^{\frac{1}{2}}, \quad (4.5.2)$$

and hence, by the definition of the discrete incremental transport path (4.4.1), we have

$$\begin{aligned} \max_{t \in [0, 1]} \|y_t^K - \mathbb{1}\|_{C^{1, \alpha}(\overline{\Omega})} &\leq \frac{3}{2} \max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^{1, \alpha}(\overline{\Omega})} + \frac{1}{2K^2} \max_{k=1, \dots, K-1} \|a_k^K\|_{C^{1, \alpha}(\overline{\Omega})} \\ &\leq C(K^{-\frac{1}{2}} + K^{-\frac{3}{2}}). \end{aligned} \quad (4.5.3)$$

This implies that y_t^K converges to the identity in $C^{1, \alpha}(\overline{\Omega}, \overline{\Omega})$, uniformly in t , and that it is invertible for all $t \in [0, 1]$ and every large enough K . Moreover, from (4.5.3) and the smoothness of the inversion, we have

$$\begin{aligned} \|x_t^K\|_{C^{1, \alpha}(\overline{\Omega})} &\leq 1 + C \left(\max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^{1, \alpha}(\overline{\Omega})} + \max_{k=1, \dots, K-1} \frac{1}{K^2} \|a_k^K\|_{C^{1, \alpha}(\overline{\Omega})} \right) \\ &\leq 1 + C(K^{-\frac{1}{2}} + K^{-\frac{3}{2}}). \end{aligned}$$

Thus, we are able to define all the temporal extended quantities introduced in Section 4.4.

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

Using the definition of the discrete acceleration \hat{w}_k^K (cf. (4.3.3)) and its time extension w_t^K (cf. (4.4.10)), together with (4.5.3) and the locally Lipschitz property of the determinant function, we get

$$\begin{aligned} \lim_{K \rightarrow \infty} \int_0^1 \int_{\Omega} (w_t^K)^2 \, dx \, dt &= \lim_{K \rightarrow \infty} \sum_{k=1}^{K-1} \int_{\Omega} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} |\det(Dx_t^K) \hat{w}_k^K \circ x_t^K|^2 \, dt \, dx \\ &= \lim_{K \rightarrow \infty} \sum_{k=1}^{K-1} \int_{\Omega} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} |\hat{w}_k^K|^2 \det(Dx_t^K) \, dt \, dx \end{aligned}$$

$$= \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} |\hat{w}_k^K|^2 dx.$$

Using the same ideas, together with $z_t^K \leq |\det(Dx_t^K) \hat{z}_k^K \circ x_t^K| + \frac{1}{K} |\det(Dx_t^K) \hat{w}_k^K \circ x_t^K|$ (cf. (4.4.7)) on every subinterval, we have

$$\begin{aligned} \lim_{K \rightarrow \infty} \int_0^1 \int_{\Omega} (z_t^K)^2 dx dt &\leq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \int_{\Omega} |\det(Dx_t^K) \hat{z}_k^K \circ x_t^K|^2 dx + CK^{-1} \\ &= \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \int_{\Omega} |\hat{z}_k^K|^2 dx + CK^{-\frac{1}{2}} \\ &= \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \int_{\Omega} |\hat{z}_k^K|^2 dx. \end{aligned}$$

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

$$\|z\|_{L^2((0,1) \times \Omega)}^2 \leq \liminf_{K \rightarrow \infty} \|z^K\|_{L^2((0,1) \times \Omega)}^2, \quad \|w\|_{L^2((0,1) \times \Omega)}^2 \leq \liminf_{K \rightarrow \infty} \|w^K\|_{L^2((0,1) \times \Omega)}^2. \quad (4.5.4)$$

Using Korn's inequality 2.2.11 and Poincaré's 2.2.9 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. Hence, by the Gagliardo-Nirenberg inequality, we have that $\{v^K\}_{K \in \mathbb{N}}$ and $\{a^K\}_{K \in \mathbb{N}}$ are uniformly bounded in $H^m((0,1), \mathcal{V})$, 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

$$\frac{1}{K^4} W_A(Da_k^K) = \frac{1}{2K^4} D^2 W_A(\mathbf{0})(Da_k^K, Da_k^K) + r_{a,k}^K = \frac{1}{K^4} \text{tr}(\varepsilon[a_k^K]^2) + r_{a,k}^K. \quad (4.5.5)$$

For the remainder term, we have $r_{a,k}^K = \mathcal{O}(K^{-6} |Da_k^K|^3)$, and by using Lemma 4.3.3 and (4.5.2), we obtain

$$\sum_{k=1}^{K-1} K^3 \int_{\Omega} r_{a,k}^K dx \leq \frac{C}{K^2} \max_{k=1, \dots, K-1} \|a_k^K\|_{C^1(\bar{\Omega})} \frac{1}{K} \sum_{k=1}^{K-1} \|a_k^K\|_{H^m(\Omega)}^2 \leq CK^{-\frac{3}{2}} \bar{\mathcal{F}}. \quad (4.5.6)$$

Then, we use weak lower semicontinuity of the energy to write

$$\liminf_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} W_A(Da_k^K) + \gamma |D^m a_k^K|^2 \, dx \quad (4.5.7)$$

$$\begin{aligned} &= \liminf_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} \operatorname{tr}(\varepsilon[a_k^K]^2) + \gamma |D^m a_k^K|^2 \, dx \\ &= \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} \operatorname{tr}(\varepsilon[a_t^K]^2) + \gamma |D^m a_t^K|^2 \, dx \, dt \\ &\geq \int_0^1 \int_{\Omega} \operatorname{tr}(\varepsilon[a]^2) + \gamma |D^m a|^2 \, dx \, dt. \end{aligned} \quad (4.5.8)$$

Analogously, the Taylor expansions of $W_D(\mathbb{1} + (t - t_{k-\frac{1}{2}})Dv_k^K(t, \cdot))$ around $t_{k-\frac{1}{2}}$, evaluated at $t_{k+\frac{1}{2}}$, for $k = 1, \dots, K-1$, and correspondingly for the intervals of size $\frac{1}{2K}$, give

$$\begin{aligned} &\liminf_{K \rightarrow \infty} K \sum_{k=1}^K \int_{\Omega} W_D(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 \, dx \\ &= \liminf_{K \rightarrow \infty} \int_{\Omega} \frac{1}{2K} (\operatorname{tr}(\varepsilon[v_1^K]^2) + \gamma |D^m v_1^K|^2) + \frac{1}{K} \sum_{k=1}^{K-1} \operatorname{tr}(\varepsilon[v_k^K]^2) + \gamma |D^m v_k^K|^2 \\ &\quad + \frac{1}{2K} (\operatorname{tr}(\varepsilon[v_1^K]^2) + \gamma |D^m v_1^K|^2) \, dx \\ &= \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} \operatorname{tr}(\varepsilon[v_t^K]^2) + \gamma |D^m v_t^K|^2 + \frac{1}{K} (\operatorname{tr}(\varepsilon[v_t^K]^2) + \gamma |D^m v_t^K|^2) \, dx \, dt \\ &\geq \int_0^1 \int_{\Omega} \operatorname{tr}(\varepsilon[v]^2) + \gamma |D^m v|^2 \, dx \, dt, \end{aligned} \quad (4.5.9)$$

where we once again used that $|v_t^K - v_k^K| \leq \frac{1}{K} |a_k^K|$ on every subinterval.

For the optimal transport terms, it holds

$$\begin{aligned} &\lim_{K \rightarrow \infty} \int_0^1 \int_{\Omega} |a_t^K|^2 \vartheta_t^K \, dx \, dt \\ &= \lim_{K \rightarrow \infty} \int_0^1 \int_{\Omega} |a_t^K \circ y_t^K|^2 \det(Dy_t^K) \vartheta_t^K \circ y_t^K \, dx \, dt \\ &\leq \lim_{K \rightarrow \infty} \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} |a_t^K \circ y_t^K|^2 (\vartheta_{k-1}^K + CK^{-1} \hat{z}_{k-1}^K + CK^{-2} \hat{w}_k^K) \, dx \, dt \\ &\leq \lim_{K \rightarrow \infty} (1 + CK^{-1}) \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} |a_t^K \circ y_t^K|^2 \vartheta_{k-1}^K \, dx \, dt \\ &\leq \lim_{K \rightarrow \infty} (1 + CK^{-1}) \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} (|a_t^K|^2 + C \|Da_t^K\|_{C^0} \|y_t^K - \mathbb{1}\|_{C^0}) \vartheta_{k-1}^K \, dx \, dt \end{aligned}$$

$$\begin{aligned}
&\leq \lim_{K \rightarrow \infty} (1 + CK^{-1}) \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} (|a_k^K|^2 + CK^{-\frac{1}{2}}) \vartheta_{k-1}^K dx dt \\
&\leq \lim_{K \rightarrow \infty} (1 + CK^{-\frac{1}{2}}) \sum_{k=1}^K \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} |a_k^K|^2 \vartheta_{k-1}^K dx dt \\
&\leq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \int_{\Omega} |a_k^K|^2 \vartheta_{k-1}^K dx dt,
\end{aligned}$$

where we used the definition of the image extension operator, the fact that $a_t^K \circ y_t^K \leq a_t^K + \|Da_t^K\|_{C^0} \|y_t^K - \mathbb{1}\|_{C^0}$ and equation (4.5.3). The analogous inequality for the first order term follows completely analogously.

Finally, it remains to show that the estimates

$$\begin{aligned}
\int_0^1 \int_{\Omega} |v_t|^2 \vartheta_t dx dt &\leq \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} |v_t^K|^2 \vartheta_t^K dx dt, \\
\int_0^1 \int_{\Omega} |a_t|^2 \vartheta_t dx dt &\leq \liminf_{K \rightarrow \infty} \int_0^1 \int_{\Omega} |a_t^K|^2 \vartheta_t^K dx dt
\end{aligned} \tag{4.5.10}$$

hold. This is a direct result of the weak \mathcal{I} -convergence $\vartheta^K \rightharpoonup \vartheta$ coupled with the strong $C^{1,\alpha}$ -convergence $v^K \rightarrow v$ and $a^K \rightarrow a$.

Altogether, (4.5.4), (4.5.8), (4.5.9) and (4.5.10) give

$$\begin{aligned}
\liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma,K}[\vartheta^K] &= \liminf_{K \rightarrow \infty} \mathbf{F}^{\sigma,K,D}[\Theta^K, \Phi^K] \\
&\geq \int_{\Omega} |a_t|^2 \vartheta + L[a_t, a_t] + \frac{1}{\delta} \hat{w}_t^2 + \sigma \left(|v_t|^2 \vartheta + L[v_t, v_t] + \frac{1}{\delta} \hat{z}_t^2 \right) dx dt.
\end{aligned}$$

(iii): *Proving the admissibility of the limit of the discrete velocities and acceleration fields, and the discrete first and second order source terms, respectively.*

In this step, we show that the limit objects v, a, z, w are indeed corresponding quantities for the image curve ϑ , i.e., $(v, a, z, w) \in \mathcal{C}[\vartheta]$ and they satisfy relations (4.2.12)–(4.2.15). Once again, we limit ourselves to the second order quantities here.

By [Fio17, Propositions 1.2.4 and 1.2.7], we have

$$\begin{aligned}
\|\tilde{v}_t^K\|_{C^{1,\alpha}(\overline{\Omega})} &\leq C \|v_t^K\|_{C^{1,\alpha}(\overline{\Omega})} (1 + \|x_t^K\|_{C^{1,\alpha}(\overline{\Omega})} + \|x_t^K\|_{C^{1,\alpha}(\overline{\Omega})}^{\alpha+1}), \\
\|\tilde{a}_t^K\|_{C^{1,\alpha}(\overline{\Omega})} &\leq C \|a_t^K\|_{C^{1,\alpha}(\overline{\Omega})} (1 + \|x_t^K\|_{C^{1,\alpha}(\overline{\Omega})} + \|x_t^K\|_{C^{1,\alpha}(\overline{\Omega})}^{\alpha+1}),
\end{aligned} \tag{4.5.11}$$

where \tilde{v}_t^K and \tilde{a}_t^K are defined by (4.4.3). In particular, this implies the uniform boundedness of $\{\tilde{v}^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1), C^{1,\alpha}(\overline{\Omega}, \mathbb{R}^n))$. Then, applying Remarks 4.1.2 and 4.1.3 to (4.4.5), we have that $\{\psi^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $H^1 \subset C^{0, \frac{1}{2}}([0, 1], C^{1,\alpha}(\overline{\Omega}, \overline{\Omega}))$. Furthermore, by the compact embedding of Hölder spaces (cf. Corollary 2.2.15), we have, for all $\min(\frac{1}{2}, \alpha) > \beta > 0$, that $\psi^K \rightarrow \psi$ in $C^{0,\beta}([0, 1], C^{1,\beta}(\overline{\Omega}, \overline{\Omega}))$. 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 4.1.1), we have $\psi^{v^K} \rightarrow \psi^v$ in $C^0([0, 1] \times \overline{\Omega})$. Furthermore, by Remark 4.1.2, we have the Lipschitz continuity of the solution operator, and using the spatial Lipschitz property of v_k^K together with (4.5.1), we have $\psi^K - \psi^{v^K} \rightarrow 0$ in $C^0([0, 1] \times \overline{\Omega})$, finally confirming that $\psi = \psi^v$.

To show that the equation $\ddot{\psi}_t = a_t \circ \psi_t$ is satisfied, first observe that (4.4.6), *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 $L^2((0, 1), C^{1,\alpha}(\overline{\Omega}, \overline{\Omega}))$. To this end, we used the uniform boundedness of $\{\tilde{a}_t^K\}_{K \in \mathbb{N}}$ in $L^2((0, 1), C^{1,\alpha}(\overline{\Omega}, \overline{\Omega}))$, following from (4.5.11), and the uniform boundedness of $\{\psi_t^K\}_{K \in \mathbb{N}}$ in $C^0([0, 1], C^{1,\alpha}(\overline{\Omega}, \overline{\Omega}))$, together with the estimate on composition of Hölder functions [Fio17, Propositions 1.2.4 and 1.2.7]. Together with the previous paragraph, we conclude that $\{\psi^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $H^2((0, 1), C^{1,\alpha}(\overline{\Omega}, \overline{\Omega}))$, and converges weakly to ψ in $H^2((0, 1), C^{1,\beta}(\overline{\Omega}, \overline{\Omega}))$, and strongly in $C^{1,\beta}([0, 1], C^{1,\beta}(\overline{\Omega}, \overline{\Omega}))$, for $\min(\frac{1}{2}, \alpha) > \beta > 0$. In particular, we have that $\psi \in H^2((0, 1), C^{1,\beta}(\overline{\Omega}, \overline{\Omega}))$ and $\tilde{a}^K \circ \psi^K = \ddot{\psi}^K \rightarrow \ddot{\psi}$ in $L^2((0, 1) \times \Omega)$. Hence, it suffices to verify $\tilde{a}^K \circ \psi^K \rightharpoonup a \circ \psi$ in $L^2((0, 1) \times \Omega)$. Since, by the proof of 4.3.7, we already have $a^K \circ \psi^K \rightharpoonup a \circ \psi$ in $L^2((0, 1) \times \Omega)$, we conclude the proof by checking that $\tilde{a}^K \circ \psi^K - a^K \circ \psi^K \rightarrow 0$ in $L^2((0, 1), C^0(\overline{\Omega}, \overline{\Omega}))$. Indeed,

$$\begin{aligned} & \|\tilde{a}^K \circ \psi^K - a^K \circ \psi^K\|_{L^2((0,1), C^0(\overline{\Omega}))}^2 \\ &= \sum_{k=1}^{K-1} \int_{t_{k-\frac{1}{2}}^K}^{t_{k+\frac{1}{2}}^K} \|\tilde{a}_t^K \circ \psi_t^K - a_t^K \circ \psi_t^K\|_{C^0(\overline{\Omega})}^2 dt \\ &\leq C \sum_{k=1}^{K-1} \int_{t_{k-1}^K}^{t_k^K} \|a_k^K(t, \cdot)\|_{C^1(\overline{\Omega})}^2 \|y_k^K(t, \cdot) - \mathbb{1}\|_{C^0(\overline{\Omega})}^2 dt \\ &\leq C \|a^K\|_{L^2((0,1), C^1(\overline{\Omega}))}^2 \left(\max_{k=1, \dots, K} \|\phi_k^K - \mathbb{1}\|_{C^0(\overline{\Omega})}^2 + \max_{k=1, \dots, K-1} K^{-2} \|a_k^K\|_{C^0(\overline{\Omega})}^2 \right) \\ &\leq CK^{-1} \|a^K\|_{L^2((0,1), C^1(\overline{\Omega}))}^2, \end{aligned}$$

where we used the Lipschitz property of a_t^K , the transformation formula, and finally, (4.5.1) and (4.5.2).

In order to show that z and w are indeed scalar weak source terms of ϑ , first observe that, from $\vartheta^K \rightharpoonup \vartheta$ and $\psi^K \rightarrow \psi$, by the proof of Theorem 4.3.7, we obtain $\det(D\psi^K)\vartheta^K \circ \psi^K \rightharpoonup \det(D\psi)\vartheta \circ \psi$ in $L^2((0, 1) \times \Omega)$, and analogously $\det(D\psi^K)z^K \circ \psi^K \rightharpoonup \det(D\psi)z \circ \psi$ and $\det(D\psi^K)w^K \circ \psi^K \rightharpoonup \det(D\psi)w \circ \psi$. Next, note that, for $s, t \in [0, 1]$, we have

$$\begin{aligned} & \|\det(D\psi_t^K)\vartheta_t^K \circ \psi_t^K - \det(D\psi_s^K)\vartheta_s^K \circ \psi_s^K\|_{\mathcal{I}}^2 \leq |t - s| \left| \int_s^t \int_{\Omega} (\det(D\psi_r^K)z_r^K \circ \psi_r^K)^2 dx dr \right| \\ & \leq C|t - s| \|z^K\|_{L^2((0,1) \times \Omega)}^2 \leq C|t - s|, \end{aligned}$$

where we used Jensen's inequality, the transformation formula, and, finally, the uniform boundedness of $\{\psi^K\}_{K \in \mathbb{N}}$ and $\{z^K\}_{K \in \mathbb{N}}$ in the spaces $C^{1,\beta}([0, 1], C^{1,\beta}(\overline{\Omega}, \overline{\Omega}))$ and $L^2((0, 1) \times \Omega)$, respectively. Thus, $\{\det(D\psi^K)\vartheta^K \circ \psi^K\}_{K \in \mathbb{N}}$ 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 . Then, by the weak closedness of this set we obtain $\det(D\psi)\vartheta \circ \psi \in C_{\omega}^0([0, 1], \mathcal{I})$. 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

$$\begin{aligned}
& \int_{\tilde{\Omega}} |\det(D\psi_{t+\tau})\vartheta_{t+\tau} \circ \psi_{t+\tau} - \det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_{s+\tau})\vartheta_{s+\tau} \circ \psi_{s+\tau} \\
& \quad + \det(D\psi_s)\vartheta_s \circ \psi_s| dx \\
& \leq \liminf_{K \rightarrow \infty} \int_{\tilde{\Omega}} |\det(D\psi_{t+\tau}^K)\vartheta_{t+\tau}^K \circ \psi_{t+\tau}^K - \det(D\psi_t^K)\vartheta_t^K \circ \psi_t^K - \det(D\psi_{s+\tau}^K)\vartheta_{s+\tau}^K \circ \psi_{s+\tau}^K \\
& \quad + \det(D\psi_s^K)\vartheta_s^K \circ \psi_s^K| dx \\
& \leq \liminf_{K \rightarrow \infty} \int_{\tilde{\Omega}} \int_0^\tau \int_s^t \det(D\psi_{r+l}^K)w_{r+l}^K \circ \psi_{r+l}^K dr dl dx \\
& = \int_{\tilde{\Omega}} \int_0^\tau \int_s^t \det(D\psi_{r+l})w_{r+l} \circ \psi_{r+l} dr dl dx.
\end{aligned}$$

Since this holds for any $\tilde{\Omega} \subset \Omega$, we have that w is the second (scalar) weak source term for ϑ . The proof of z being the first weak source term is analogous.

This finally finishes the proof of the lim inf-inequality. We now state a corollary of the preceding proof that shows that the infimum in (4.2.11) is actually attained.

Proposition 4.5.4. *For $\vartheta \in L^2([0, 1], \mathcal{I})$ with $\mathcal{F}^\sigma[\vartheta] < \infty$ there exists an optimal tuple $(v, a, z, w) \in \mathcal{C}[\vartheta]$ such that*

$$\mathcal{F}^\sigma[\vartheta] = \int_0^1 |a|^2 \vartheta + L[a, a] + \frac{1}{\delta} w^2 + \sigma \left(|v|^2 \vartheta + L[v, v] + \frac{1}{\delta} z^2 \right) dx dt.$$

Proof. The functional \mathcal{F}^σ is coercive by Korn's inequality 2.2.11 and Gagliardo-Nirenberg interpolation estimate 2.2.10, and it is clearly weakly lower semicontinuous. Since $\mathcal{C}[\vartheta]$ 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. \square

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

(i) *Construction of a recovery sequence.*

Consider an image curve $\vartheta \in L^2([0, 1], \mathcal{I})$ with $\mathcal{F}^\sigma[\vartheta] < \infty$. Then, the previous proposition guarantees the existence of an associated optimal velocity field, an acceleration field, and the first and second order weak (scalar) source terms, denoted by $(v, a, z, w) \in \mathcal{C}[\vartheta]$, respectively, such that

$$\mathcal{F}^\sigma[\vartheta] = \int_0^1 \int_{\Omega} |a|^2 \vartheta + L[a, a] + \frac{1}{\delta} w^2 + \sigma \left(|v|^2 \vartheta + L[v, v] + \frac{1}{\delta} z^2 \right) dx dt.$$

We define

$$\phi_k^K := \psi_{t_{k-1}^K, t_k^K} = \psi_{t_k^K} \circ \psi_{t_{k-1}^K}^{-1}, \quad k = 1, \dots, K, \quad (4.5.12)$$

where ψ is the flow associated with velocity v and $\psi_0 = \mathbf{1}$ (cf. (4.2.12)). We have

$$\max_{k=1, \dots, K} \|\phi_k^K - \mathbf{1}\|_{C^1(\bar{\Omega})} \leq \sup_{|t-s| \leq K^{-1}} \|\psi_{s,t} - \mathbf{1}\|_{C^1(\bar{\Omega})}$$

$$\begin{aligned}
&\leq \sup_{|t-s| \leq K^{-1}} C \left| \int_s^t \|v_r \circ \psi_r\|_{H^m(\Omega)} dr \right| \\
&\leq \sup_{|t-s| \leq K^{-1}} C \left| \int_s^t \|v_r\|_{H^m(\Omega)} dr \right| \\
&\leq CK^{-\frac{1}{2}} \sup_{|t-s| \leq K^{-1}} \left| \int_s^t \|v_r\|_{H^m(\Omega)}^2 dr \right|^{\frac{1}{2}} \leq CK^{-\frac{1}{2}} \sqrt{\mathcal{F}^\sigma[\vartheta]}, \quad (4.5.13)
\end{aligned}$$

by Lemma 4.3.3 and Cauchy's inequality. For the second inequality, we used [BV17, Lemma 3.5] which states that

$$\|v_r \circ \psi_r\|_{H^m(\Omega)} \leq C \|v_r\|_{H^m(\Omega)}. \quad (4.5.14)$$

Thereby, for K large enough, we have $\Phi^K \in \mathcal{D}^K$, and we are in a position to define

$$\vartheta^K := \theta^K[\Theta^K, \Phi^K], \quad \Theta^K := (\vartheta_{t_0^K}, \dots, \vartheta_{t_K^K}),$$

where the point evaluation in time is possible since $\vartheta \in C^1([0, 1], \mathcal{I})$ (cf. Remarks 4.2.5).

(ii) *Proof of the actual inequality.*

The arguments for the estimates and the convergence of the first order terms are the easier part of the proof, so we shall focus only on the second order terms; the first-order terms can be then trivially obtained. 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 source term $\hat{w}_k^K, k = 1, \dots, K-1$ given by (4.3.3) and its continuous counterpart w , via

$$\begin{aligned}
&\int_{\Omega} |\hat{w}_k^K|^2 dx \\
&= K^4 \int_{\Omega} |\det(D(\phi_{k+1}^K \circ \phi_k^K)) \vartheta_{t_{k+1}^K} \circ \phi_{k+1}^K \circ \phi_k^K - 2 \det(D\phi_k^K) \vartheta_{t_k^K} \circ \phi_k^K + \vartheta_{t_{k-1}^K}|^2 dx \\
&= K^4 \int_{\Omega} |\det(D\psi_{t_{k-1}^K, t_{k+1}^K}) \vartheta_{t_{k+1}^K} \circ \psi_{t_{k-1}^K, t_{k+1}^K} - 2 \det(D\psi_{t_{k-1}^K, t_k^K}) \vartheta_{t_k^K} \circ \psi_{t_{k-1}^K, t_k^K} + \vartheta_{t_{k-1}^K}|^2 dx \\
&= K^4 \int_{\Omega} |\det(D\psi_{t_{k+1}^K}) \vartheta_{t_{k+1}^K} \circ \psi_{t_{k+1}^K} - 2 \det(D\psi_{t_k^K}) \vartheta_{t_k^K} \circ \psi_{t_k^K} \\
&\quad + \det(D\psi_{t_{k-1}^K}) \vartheta_{t_{k-1}^K} \circ \psi_{t_{k-1}^K}|^2 \det(D\psi_{t_{k-1}^K})^{-1} dx \\
&\leq K^4 \int_{\Omega} \left(\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \det(D\psi_{r+s}) w_{r+s} \circ \psi_{r+s} dr ds \right)^2 \det(D\psi_{t_{k-1}^K})^{-1} dx \\
&\leq K^2 \int_0^{\frac{1}{K}} \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} w_{r+s}^2 \left(\det(D\psi_{t_{k-1}^K, r+s}) \circ \psi_{r+s, t_{k-1}^K} \right) dx ds dr \\
&\leq K^2 (1 + CK^{-\frac{1}{2}}) \int_0^{\frac{1}{K}} \int_{t_{k-1}^K}^{t_k^K} \int_{\Omega} w_{r+s}^2 dx ds dr.
\end{aligned}$$

Here, we first used (4.5.12) and the transformation formula for the second and third equalities, respectively, then the variational definition of the second order source term (4.2.10) for the first

inequality, and, finally, Jensen's inequality and (4.5.13) in the last two estimates. Averaging the above expressions over $k = 1, \dots, K - 1$, we obtain

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} |\hat{w}_k^K|^2 dx &\leq K(1 + CK^{-\frac{1}{2}}) \int_0^{\frac{1}{K}} \int_0^{1-\frac{1}{K}} \int_{\Omega} w_{r+s}^2 dx ds dr \\ &\leq K(1 + CK^{-\frac{1}{2}}) \int_0^{\frac{1}{K}} \int_0^1 \int_{\Omega} w_t^2 dx dt dr \\ &\leq (1 + CK^{-\frac{1}{2}}) \int_0^1 \int_{\Omega} w_t^2 dx dt. \end{aligned} \quad (4.5.15)$$

Completely analogously, we show

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^K \int_{\Omega} |\hat{z}_k^K|^2 dx &= K \sum_{k=1}^K \int_{\Omega} |\det(D\phi_k^K) \vartheta_k^K \circ \phi_k^K - \vartheta_{k-1}^K|^2 dx \\ &\leq (1 + CK^{-\frac{1}{2}}) \int_0^1 \int_{\Omega} z_t^2 dx dt. \end{aligned} \quad (4.5.16)$$

Next, we express the discrete acceleration $a_k^K, k = 1, \dots, K - 1$ in terms of its continuous counterpart a :

$$\begin{aligned} a_k^K &= K^2(\phi_{k+1}^K \circ \phi_k^K - 2\phi_k^K + \mathbb{1}) \\ &= K^2(\psi_{t_{k-1}^K, t_{k+1}^K} - 2\psi_{t_{k-1}^K, t_k^K} + \psi_{t_{k-1}^K, t_{k-1}^K}) \\ &= K^2 \left(\int_{t_k^K}^{t_{k+1}^K} \dot{\psi}_t \circ \psi_{t_k^K}^{-1} dt - \int_{t_{k-1}^K}^{t_k^K} \dot{\psi}_t \circ \psi_{t_{k-1}^K}^{-1} dt \right) \\ &= K^2 \left(\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \ddot{\psi}_{t+\tau} \circ \psi_{t_{k-1}^K}^{-1} d\tau dt \right) \\ &= K^2 \left(\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau} d\tau dt \right), \end{aligned} \quad (4.5.17)$$

where in the second equality we used (4.5.12), and in the last equality (4.2.13). Then, using the Cauchy-Schwarz inequality and (4.5.14), we obtain the following estimate

$$\max_{k=1, \dots, K-1} \|a_k^K\|_{C^1(\bar{\Omega})} \leq CK^{\frac{1}{2}} \sup_{t \in [0,1], 0 < \tau \leq K^{-1}} \left| \int_t^{t+2\tau} \|a_s\|_{H^m(\Omega)}^2 ds \right|^{\frac{1}{2}}. \quad (4.5.18)$$

The same Taylor expansion arguments as in (4.5.5) and (4.5.6) now imply, together with (4.5.18)

$$\begin{aligned} \int_{\Omega} W_A(Da_k^K) + \gamma |D^m a_k^K|^2 dx &\leq \int_{\Omega} L[a_k^K, a_k^K] + CK^{-2} |Da_k^K|^3 dx \\ &\leq \int_{\Omega} L[a_k^K, a_k^K] dx + CK^{-\frac{3}{2}} \sqrt{\mathcal{F}^\sigma[u]}. \end{aligned} \quad (4.5.19)$$

Applying Jensen's inequality twice on L , and taking into account (4.5.17), 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 $|\operatorname{tr}(AB)| \leq |\operatorname{tr}(A)| + |\operatorname{tr} A(B - \mathbb{1})|$, (4.5.13), and the transformation formula, to get

$$\begin{aligned} & \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \operatorname{tr} \left(\varepsilon [a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau}]^2 \right) d\tau dt dx \\ &\leq \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \operatorname{tr} \left((\varepsilon [a_{t+\tau}] \circ \psi_{t_{k-1}^K, t+\tau})^2 \right) \\ &\quad + \operatorname{tr} \left((\varepsilon [a_{t+\tau}] \circ \psi_{t_{k-1}^K, t+\tau})^2 (\varepsilon [\psi_{t_{k-1}^K, t+\tau}]^2 - \mathbb{1}) \right) d\tau dt dx \\ &\leq \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \operatorname{tr} (\varepsilon [a_{t+\tau}]^2) + CK^{-\frac{1}{2}} \operatorname{tr} (\varepsilon [a_{t+\tau}]^2) d\tau dt dx. \end{aligned}$$

For the second term, we use (4.5.13), (4.5.14) 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

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

By iterating this argument and applying a change of variables (under consideration of (4.5.13)), we obtain

$$\int_{t_{k-1}^K}^{t_k^K} |a_t \circ \psi_{t_{k-1}^K, t}|_{H^m(\Omega)}^2 dt \leq \int_{t_{k-1}^K}^{t_k^K} |a_t|_{H^m(\Omega)}^2 + C\|a_t\|_{H^m(\Omega)}^2 K^{-\frac{1}{2}} dt. \quad (4.5.20)$$

For the optimal transport term, we obtain

$$\begin{aligned} & \int_{\Omega} |a_k^K|^2 \vartheta_{k-1}^K dx \\ &= K^4 \int_{\Omega} \left(\int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau} d\tau dt \right)^2 \vartheta_{k-1}^K dx \end{aligned}$$

$$\begin{aligned}
&\leq K^2 \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \left(a_{t+\tau} \circ \psi_{t_{k-1}^K, t+\tau} \right)^2 \vartheta_{k-1}^K \, d\tau \, dt \, dx \\
&\leq K^2 \int_{\Omega} \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \det(D\psi_{t_{k-1}^K}) |a_{t+\tau} \circ \psi_{t+\tau}|^2 \vartheta_{t_{k-1}^K} \circ \psi_{t_{k-1}^K} \, d\tau \, dt \, dx \\
&\leq K^2 (1 + CK^{-\frac{1}{2}}) \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \int_{\Omega} \det(D\psi_{t+\tau}) |a_{t+\tau} \circ \psi_{t+\tau}|^2 \vartheta_{t+\tau} \circ \psi_{t+\tau} \, dx \, d\tau \, dt \\
&\leq K^2 (1 + CK^{-\frac{1}{2}}) \int_{t_{k-1}^K}^{t_k^K} \int_0^{\frac{1}{K}} \int_{\Omega} |a_{t+\tau}|^2 \vartheta_{t+\tau} \, dx \, d\tau \, dt,
\end{aligned}$$

where we used the fact that $\det(D\psi_t)\vartheta_t \circ \psi_t \in C_{\omega}([0, 1], \mathcal{I})$, and once again the local Lipschitz property of the determinant and (4.5.13). In combination with (4.5.18) and (4.5.19), we have

$$\begin{aligned}
&\frac{1}{K} \sum_{k=1}^{K-1} \int_{\Omega} |a_k^K|^2 \vartheta_{k-1}^K + W_A(Da_k^K) + \gamma |D^m a_k^K|^2 \, dx \\
&\leq K \int_0^{\frac{1}{K}} \int_0^{1-\frac{1}{K}} \int_{\Omega} |a_{t+\tau}|^2 \vartheta_{t+\tau} + L[a_{t+\tau}, a_{t+\tau}] + \mathcal{O}(K^{-\frac{1}{2}}) \, dx \, dt \, d\tau \\
&\leq \int_0^1 \int_{\Omega} |a_t|^2 \vartheta_t + L[a_t, a_t] + \mathcal{O}(K^{-\frac{1}{2}}) \, dx \, dt.
\end{aligned}$$

Altogether, taking into account (4.5.15), we obtain

$$\mathcal{F}^K[\vartheta^K] \leq \mathcal{F}[\vartheta] + \mathcal{O}(K^{-\frac{1}{2}}).$$

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

$$K \sum_{k=1}^K \int_{\Omega} W_D(D\phi_k^K) + \gamma |D^m \phi_k^K|^2 \, dx \leq \int_0^1 \int_{\Omega} L[v_t, v_t] + \mathcal{O}(K^{-\frac{1}{2}}) \, dx \, dt,$$

which, together with (4.5.16), gives

$$\mathcal{E}^K[\vartheta^K] \leq \mathcal{E}[\vartheta] + \mathcal{O}(K^{-\frac{1}{2}}),$$

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 $\vartheta^K \rightarrow \vartheta$ in $L^2((0, 1), \mathcal{I})$ as $K \rightarrow \infty$. To this end, we introduce the piecewise constant interpolation

$$\bar{\vartheta}_t^K := \begin{cases} \vartheta_0^K, & t \in [0, t_{\frac{1}{2}}^K], \\ \vartheta_{k-1}^K, & t \in (t_{k-\frac{1}{2}}^K, t_{k+\frac{1}{2}}^K], \, k = 1, \dots, K-1, \\ \vartheta_{K-1}^K, & t \in [t_{K-\frac{1}{2}}^K, 1]. \end{cases}$$

We will show

$$\vartheta^K - \bar{\vartheta}^K \rightarrow 0 \text{ in } L^\infty((0, 1), \mathcal{I}). \quad (4.5.21)$$

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

$$\begin{aligned} & \|\vartheta_t^K - \bar{\vartheta}_t^K\|_{\mathcal{I}}^2 \\ & \leq C \left(\|\det(Dx_t^K)(\vartheta_{k-1}^K - \det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K) \circ x_t^K\|_{\mathcal{I}}^2 + \|\det(Dx_t^K)\vartheta_{k-1}^K \circ x_t^K - \vartheta_{k-1}^K\|_{\mathcal{I}}^2 \right) \\ & \quad + \|\det(Dx_t^K)(\det(D(\phi_{k+1}^K \circ \phi_k^K))\vartheta_{k+1}^K \circ \phi_{k+1}^K \circ \phi_k^K - 2\det(D\phi_k^K)\vartheta_k^K \circ \phi_k^K + \vartheta_{k-1}^K) \circ x_t^K\|_{\mathcal{I}}^2 \\ & \leq C \left(K^{-2} \|\hat{z}_k^K\|_{L^2(\Omega)}^2 \|\det Dx^K\|_{L^\infty((0,1) \times \Omega)} + K^{-4} \|\hat{w}_k^K\|_{L^2(\Omega)}^2 \|\det Dx^K\|_{L^\infty((0,1) \times \Omega)} \right. \\ & \quad \left. + \|\det(Dx_t^K)\vartheta_{k-1}^K \circ x_t^K - \vartheta_{k-1}^K\|_{\mathcal{I}}^2 \right) \\ & \leq C \left(K^{-1} \|z\|_{L^2((0,1) \times \Omega)}^2 + CK^{-3} \|w\|_{L^2((0,1) \times \Omega)}^2 + \|\det(Dx_t^K)\vartheta_{k-1}^K \circ x_t^K - \vartheta_{k-1}^K\|_{\mathcal{I}}^2 \right) \end{aligned} \quad (4.5.22)$$

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. As before, we uniformly approximate the sequence $\{\vartheta_{k(K)}^K\}_{K \in \mathbb{N}}$ by smooth functions, and use (4.5.13) and (4.5.18), to prove

$$\|\det(Dx_t^K)\vartheta_{k(K)-1}^K \circ x_t^K - \vartheta_{k(K)-1}^K\|_{\mathcal{I}}^2 \rightarrow 0, \text{ uniformly in } t.$$

Plugging this back into (4.5.22), we get (4.5.21). Furthermore, as $\{\bar{\vartheta}^K\}_{K \in \mathbb{N}}$ is a sequence of piecewise constant approximations of $\vartheta \in C^1([0, 1], \mathcal{I})$, we have that $\bar{\vartheta}^K \rightarrow \vartheta$ in $L^2([0, 1], \mathcal{I})$ as $K \rightarrow \infty$. Thus, we can finally conclude $\vartheta^K \rightarrow \vartheta$ in $L^2((0, 1), \mathcal{I})$, as we wanted to show. This finally finishes the proof of Theorem 4.5.1. \square

As a corollary of the previous theorem, we are able to show the existence of the time continuous regularized spline interpolation (*cf.* Definition 4.2.6) 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 $(\vartheta_j^I)_{j=1, \dots, J} \subseteq \mathcal{I}$ be the set of constraint images at the corresponding constraint times (*cf.* (4.2.1) and (4.3.5)).

Theorem 4.5.5 (Convergence of discrete spline interpolations). *For every K that satisfies the above condition, let $\vartheta^K \in L^2([0, 1], \mathcal{I})$ be a minimizer of $\mathcal{F}^{\sigma, K}$ among the image curves satisfying $\vartheta^K = \theta^K[\Theta^K, \Phi^K]$ with $\vartheta_{i_j^K}^K = \vartheta_j^I$ for all $j = 1, \dots, J$. Then, a subsequence of $\{\vartheta^K\}_{K \in \mathbb{N}}$ 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 $\vartheta_{t_j} = \vartheta_j^I$ 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 $\eta_{t_i}^j = \delta_{ij}$. We define a smooth interpolating curve of the fixed images $\tilde{\vartheta}_t := \sum_{j=1}^J \eta_t^j \vartheta_j^I$. Let $\tilde{\Theta}^K := (\tilde{\vartheta}_{t_1^K}, \dots, \tilde{\vartheta}_{t_J^K})$ and define $\tilde{\vartheta}^K := \theta^K[\tilde{\Theta}^K, \mathbb{1}^K]$. This image curve gives an admissible candidate for a minimizer of the functional $\mathcal{F}^{\sigma, K}$. Indeed,

$$\mathcal{F}^{\sigma, K}[\tilde{\vartheta}^K] \leq \mathbf{F}^{\sigma, K, D}[\tilde{\Theta}^K, \mathbb{1}^K]$$

$$\begin{aligned}
&= \sigma K \sum_{k=1}^K \int_{\Omega} |\tilde{\vartheta}_k^K - \tilde{\vartheta}_{k-1}^K|^2 dx + K^3 \sum_{k=1}^{K-1} \int_{\Omega} |\tilde{\vartheta}_{k+1}^K - 2\tilde{\vartheta}_k^K + \tilde{\vartheta}_{k-1}^K|^2 dx \\
&\leq C \left(\int_{\Omega} |\tilde{\vartheta}|_{H^1((0,1))}^2 + |\tilde{\vartheta}|_{H^2((0,1))}^2 dx + 1 \right) := \overline{\mathcal{F}},
\end{aligned}$$

where the upper bound $\overline{\mathcal{F}}$ is independent of K . As defined above, let $\{\vartheta^K := \theta^K[\Theta^K, \Phi^K]\}_{K \in \mathbb{N}}$, where $\{(\Theta^K, \Phi^K)\}_{K \in \mathbb{N}}$ are the optimal pairs for the discrete spline (see Theorem 4.3.7). In particular, for K large enough, we have $\mathcal{F}^{\sigma,K}[\vartheta^K] = \mathbf{F}^{\sigma,K}[\Theta^K, \Phi^K] < \overline{\mathcal{F}}$. Then, from Lemma 4.3.3, we have the uniform boundedness of $\{\phi_k^K\}_{K \in \mathbb{N}}$ in $C^{1,\alpha}(\overline{\Omega}, \overline{\Omega})$, and, as in (4.3.18), we can show the uniform boundedness of $\{\vartheta_k^K\}_{K \in \mathbb{N}}$ in \mathcal{I} , for every $k = 0, \dots, K$. Furthermore, using the boundedness of the discrete incremental transport paths, following from (4.5.3), we can show that $\{\vartheta_t^K\}_{K \in \mathbb{N}}$ is uniformly bounded in \mathcal{I} , uniformly in $t \in [0, 1]$. Therefore, $\{\vartheta^K\}_{K \in \mathbb{N}}$ is uniformly bounded in $L^\infty([0, 1], \mathcal{I})$, and a subsequence converges weakly to some $\vartheta \in L^2([0, 1], \mathcal{I})$.

Let us show that $\vartheta_{t_j} = \vartheta_j^I$ for all $j = 1, \dots, J$. To this end, we can analogously to (4.5.21) show $\vartheta_t^K - \bar{\vartheta}_t^K \rightarrow 0$ in \mathcal{I} , uniformly in $t \in [0, 1]$. Together with $\vartheta_t^K \rightharpoonup \vartheta_t$ in \mathcal{I} for every $t \in [0, 1]$, we have the needed result since $\bar{\vartheta}_{t_j}^K = \vartheta_j^I$.

Now, we follow the usual argument and assume that there exists an image path $\hat{\vartheta} \in L^2([0, 1], \mathcal{I})$ with a finite energy, and $\hat{\vartheta}_{t_j} = \vartheta_j^I$ for all $j = 1, \dots, J$, such that $\mathcal{F}^\sigma[\hat{\vartheta}] < \mathcal{F}^\sigma[\vartheta]$. By the lim sup-part of Theorem 4.5.1, there exists a sequence $\{\hat{\vartheta}^K\}_{K \in \mathbb{N}} \subset L^2((0, 1), \mathcal{I})$ of time extensions of admissible vectors of images such that $\limsup_{K \rightarrow \infty} \mathcal{F}^{\sigma,K}[\hat{\vartheta}^K] \leq \mathcal{F}^\sigma[\hat{\vartheta}]$. Now, we apply the lim inf-part of Theorem 4.5.1 to obtain

$$\mathcal{F}^\sigma[\vartheta] \leq \liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma,K}[\vartheta^K] \leq \liminf_{K \rightarrow \infty} \mathcal{F}^{\sigma,K}[\hat{\vartheta}^K] \leq \mathcal{F}^\sigma[\hat{\vartheta}], \quad (4.5.23)$$

which is a contradiction to the above assumption. Hence, ϑ 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.*, $\lim_{K \rightarrow \infty} \mathcal{F}^{\sigma,K}[\vartheta^K] = \mathcal{F}^\sigma[\vartheta]$, which follows from (4.5.23) by using $\hat{\vartheta} = \vartheta$. \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 4.2.3, we have that for the optimal scalar quantities z, w it holds $z = |\hat{z}|$ and $w = |\hat{w}|$.

4.6 Fully Discrete Metamorphosis Splines

To numerically implement splines for image metamorphosis we have to further discretize the space. Here, we present a model for c image channels and a two dimensional image domain $\Omega := [0, 1]^2$ and follow the fully discrete version of image metamorphosis introduced in [EKP⁺21]. Before presenting the details of the space discretization, to avoid double warping in the second-order source term (4.3.3) and to further increase the robustness of the model we explicitly introduce a vector valued slack variable $\bar{z} \in L^2((0, 1), L^2(\Omega, \mathbb{R}^c))$. This leads to a relaxation of (4.2.11):

$$\mathcal{F}^\sigma[\vartheta] := \inf_{(v, a, \hat{z}, \bar{z}, w)} \int_0^1 \int_{\Omega} |a|^2 \vartheta + \lambda L[a, a] + \frac{1}{\delta} |w|^2 + \sigma(|v|^2 \vartheta + \lambda L[v, v] + \frac{1}{\delta} |\bar{z}|^2) + \frac{1}{\theta} |\bar{z} - \hat{z}|^2 dx dt,$$

with a penalty on the misfit of the new variable \bar{z} and the actual first-order source term \hat{z} , while w is the first-order source term of \bar{z} , *i.e.*,

$$\begin{aligned} \det(D\psi_t)\vartheta_t \circ \psi_t - \det(D\psi_s)\vartheta_s \circ \psi_s &= \int_s^t \det(D\psi_r)\hat{z}_r \circ \psi_r \, dr, \\ \det(D\psi_t)\bar{z}_t \circ \psi_t - \det(D\psi_s)\bar{z}_s \circ \psi_s &= \int_s^t \det(D\psi_r)w_r \circ \psi_r \, dr, \quad \forall s, t \in [0, 1]. \end{aligned}$$

The time discrete counterpart $\mathbf{F}^{\sigma, K}[\Theta]$ is defined by

$$\mathbf{F}^{\sigma, K}[\Theta] := \inf_{\bar{\mathbf{z}} \in L^2(\Omega, \mathbb{R}^c)^K, \Phi \in \mathcal{D}^K} \mathbf{F}^{\sigma, K, D}[\Theta, \bar{\mathbf{z}}, \Phi],$$

where, for $\bar{\mathbf{z}} = (\bar{z}_1, \dots, \bar{z}_K)$ we write

$$\begin{aligned} \mathbf{F}^{\sigma, K, D}[\Theta, \bar{\mathbf{z}}, \Phi] &= \int_{\Omega} \sum_{k=1}^{K-1} \frac{\lambda}{K} (W_A(Da_k) + \gamma |D^m a_k|^2) + \frac{K}{\delta} |\det(D\phi_k)\bar{z}_{k+1} \circ \phi_k - \bar{z}_k|^2 \\ &\quad + \sigma \left(\sum_{k=1}^K \lambda (KW_D(D\phi_k) + K\gamma |D^m \phi_k|^2) + \frac{1}{\delta K} |\bar{z}_k|^2 \right) \\ &\quad + \sum_{k=1}^K \frac{1}{\theta K} |K(\det(D\phi_k)\vartheta_k \circ \phi_k - \vartheta_{k-1}) - \bar{z}_k|^2 \, dx. \end{aligned}$$

Here, $\bar{w}_k := K(\det(D\phi_k)\bar{z}_{k+1} \circ \phi_k - \bar{z}_k)$, while the corresponding second order source term is $\hat{w}^k = K(\det(D\phi_k)\hat{z}_{k+1} \circ \phi_k - \hat{z}_k)$.

For $M, N \geq 3$ we define the computational domain

$$\Omega_{MN} := \left\{ \frac{0}{M-1}, \frac{1}{M-1}, \dots, \frac{M-1}{M-1} \right\} \times \left\{ \frac{0}{N-1}, \frac{1}{N-1}, \dots, \frac{N-1}{N-1} \right\},$$

with discrete boundary $\partial\Omega_{MN} := \Omega_{MN} \cap \partial([0, 1]^2)$ and $\|\vartheta\|_{L_{MN}^p}^p := \frac{1}{MN} \sum_{(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}} \|\vartheta(\mathbf{x}, \mathbf{y})\|_p^p$. The discrete image space is $\mathcal{I}_{MN} := \{\vartheta : \Omega_{MN} \rightarrow \mathbb{R}^c\}$ and the set of admissible deformations is

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

where the discrete Jacobian operator 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 the paper we used bold faced letters for fully discrete quantities. A spatial warping operator \mathbf{T} that approximates the pullback of an image channel $\vartheta^j \circ \phi$ at a point $(\mathbf{x}, \mathbf{y}) \in \Omega_{MN}$ is defined by

$$\mathbf{T}[\vartheta^j, \phi](\mathbf{x}, \mathbf{y}) := \sum_{(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \Omega_{MN}} s(\phi^1(\mathbf{x}, \mathbf{y}) - \tilde{\mathbf{x}}) s(\phi^2(\mathbf{x}, \mathbf{y}) - \tilde{\mathbf{y}}) \vartheta^j(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}),$$

where s is the third order B-spline interpolation kernel. This form of warping is also used for composition of deformations, *i.e.* we define the fully discrete acceleration as an approximation of (4.3.2) by

$$\mathbf{a}_k^j := K^2(\mathbf{T}[\phi_{k+1}^j - \mathbf{1}, \phi_k] - (\phi_k^j - \mathbf{1})), \quad j = 1, 2.$$

In summary, the fully discrete spline energy in the metamorphosis model for a $(K+1)$ -tuple $(\boldsymbol{\vartheta}_k)_{k=0}^K$ of discrete images, a K -tuple $(\bar{\mathbf{z}}_k)_{k=1}^K$ of discrete derivatives and a K -tuple $(\boldsymbol{\phi}_k)_{k=1}^K$ of discrete deformations reads as

$$\begin{aligned} \mathbf{F}_{MN}^{\sigma,K}[(\boldsymbol{\vartheta}_k)_{k=0}^K] &:= \inf_{\bar{\mathbf{z}} \in \mathcal{I}_{MN}^K, \boldsymbol{\Phi} \in \mathcal{D}_{MN}^K} \mathbf{F}_{MN}^{\sigma,K,D}[(\boldsymbol{\vartheta}_k)_{k=0}^K, (\bar{\mathbf{z}}_k)_{k=1}^K, (\boldsymbol{\phi}_k)_{k=1}^K] \\ &= \inf_{\bar{\mathbf{z}} \in \mathcal{I}_{MN}^K, \boldsymbol{\Phi} \in \mathcal{D}_{MN}^K} \sum_{k=1}^{K-1} \frac{1}{K} \|W_A(\nabla_{MN} \mathbf{a}_k)\|_{L_{MN}^1} + \frac{K}{\delta} \mathbf{D}_{MN}^s[\bar{\mathbf{z}}_k, \bar{\mathbf{z}}_{k+1}, \boldsymbol{\phi}_k] \\ &\quad + \sum_{k=1}^K \sigma \left(K \|W_D(\nabla_{MN} \boldsymbol{\phi}_k)\|_{L_{MN}^1} + \frac{1}{\delta K} \|\bar{\mathbf{z}}_k\|_{L_{MN}^2}^2 \right) \\ &\quad + \frac{1}{\theta K} \mathbf{D}_{MN}^g[\boldsymbol{\vartheta}_{k-1}, \boldsymbol{\vartheta}_k, \bar{\mathbf{z}}_k, \boldsymbol{\phi}_k], \end{aligned}$$

where

$$\begin{aligned} \mathbf{D}_{MN}^s[\mathbf{z}, \tilde{\mathbf{z}}, \boldsymbol{\phi}] &:= \frac{1}{2c} \sum_{j=1}^c \left\| \det(D\boldsymbol{\phi}) \mathbf{T}[\tilde{\mathbf{z}}^j, \boldsymbol{\phi}] - \mathbf{z}^j \right\|_{L_{MN}^2}^2, \\ \mathbf{D}_{MN}^g[\boldsymbol{\vartheta}, \tilde{\boldsymbol{\vartheta}}, \mathbf{z}, \boldsymbol{\phi}] &:= \frac{1}{2c} \sum_{j=1}^c \left\| K(\det(D\boldsymbol{\phi}) \mathbf{T}[\tilde{\boldsymbol{\vartheta}}^j, \boldsymbol{\phi}] - \boldsymbol{\vartheta}^j) - \mathbf{z}^j \right\|_{L_{MN}^2}^2. \end{aligned}$$

While in the spatially continuous context the compactness induced by the H^m -seminorm is indispensable, in this fully discrete model grid dependent regularity is ensured by the use of cubic B-splines. Thus, we dropped the higher order Sobolev norm terms in this fully discrete model.

To improve the robustness of the overall optimization, we take into account a multiresolution 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 $(\boldsymbol{\vartheta}_k)_{k=0}^K$ is computed as minimizer of $\mathbf{F}_{M_L N_L}^{\sigma,K}$ subject to given fixed images $\boldsymbol{\vartheta}_{i_j} = \boldsymbol{\vartheta}_j^I$, $j = 1, \dots, J$. In 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 a bilinear interpolation of the preceding coarse scale solutions.

4.7 Numerical Optimization Using the iPALM Algorithm

In this section, we discuss the numerical solution of the above fully discrete variational problem based on the application of a variant of the inertial proximal alternating linearized minimization algorithm (iPALM, [PS16]). Using this algorithm, effective optimization results were achieved for a wide range of non convex and non smooth problems. In particular, it was already used for numerical optimization in the context of the deep feature metamorphosis model [EKP⁺21]. Following [EKP⁺21], to enhance the stability the warping operation is linearized with respect to the deformation at $\boldsymbol{\phi}^{[\beta]} \in \mathcal{D}_{MN}$ coming from the previous iteration which leads to the modified energies

$$\begin{aligned} \tilde{\mathbf{D}}_{MN}^s[\mathbf{z}, \tilde{\mathbf{z}}, \boldsymbol{\phi}, \boldsymbol{\phi}^{[\beta]}] &:= \frac{1}{2c} \sum_{j=1}^c \left\| \det(\nabla_{MN} \boldsymbol{\phi}^{[\beta]}) \mathbf{T}[\tilde{\mathbf{z}}^j, \boldsymbol{\phi}^{[\beta]}] + \left\langle \Lambda_j(\mathbf{z}, \tilde{\mathbf{z}}, \boldsymbol{\phi}^{[\beta]}), \boldsymbol{\phi} - \boldsymbol{\phi}^{[\beta]} \right\rangle - \mathbf{z}^j \right\|_{L_{MN}^2}^2 \\ \tilde{\mathbf{D}}_{MN}^g[\boldsymbol{\vartheta}, \tilde{\boldsymbol{\vartheta}}, \mathbf{z}, \boldsymbol{\phi}, \boldsymbol{\phi}^{[\beta]}] &:= \frac{1}{2c} \sum_{j=1}^c \left\| K \det(\nabla_{MN} \boldsymbol{\phi}^{[\beta]}) \mathbf{T}[\tilde{\boldsymbol{\vartheta}}^j, \boldsymbol{\phi}^{[\beta]}] + \left\langle \Lambda_j(K\boldsymbol{\vartheta} + \mathbf{z}, K\tilde{\boldsymbol{\vartheta}}, \boldsymbol{\phi}^{[\beta]}), \boldsymbol{\phi} - \boldsymbol{\phi}^{[\beta]} \right\rangle \right\|_{L_{MN}^2}^2 \end{aligned}$$

$$-(K\boldsymbol{\vartheta}^j + \mathbf{z}^j)\|_{L^2_{MN}}^2,$$

with

$$\Lambda_j(\boldsymbol{\vartheta}, \tilde{\boldsymbol{\vartheta}}, \boldsymbol{\phi}^{[\beta]}) = \frac{1}{2}(\nabla_{MN} \det(\nabla_{MN} \boldsymbol{\phi}^{[\beta]}) \mathbf{T}[\tilde{\boldsymbol{\vartheta}}^j, \boldsymbol{\phi}^{[\beta]}] + \nabla \boldsymbol{\vartheta}^j).$$

To further stabilize the computation, the Jacobian operator applied to the images is approximated using a Sobel filter. Here, $\langle \cdot, \cdot \rangle$ represent the pointwise product of the involved matrices. We use the proximal mapping of a functional $f : \mathcal{D}_{MN} \rightarrow (-\infty, \infty]$ for $\tau > 0$ given as

$$\text{prox}_\tau^f[\boldsymbol{\phi}] := \underset{\tilde{\boldsymbol{\phi}} \in \mathcal{D}_{MN}}{\text{argmin}} \left(\frac{\tau}{2} \|\boldsymbol{\phi} - \tilde{\boldsymbol{\phi}}\|_{L^2_{MN}}^2 + f(\tilde{\boldsymbol{\phi}}) \right).$$

Then, with the function values on $\partial\Omega_{MN}$ remaining unchanged, the proximal operator we are interested in is given by

$$\begin{aligned} & \text{prox}_\tau^{\frac{K}{\delta} \tilde{\mathbf{D}}_{MN}^s + \frac{1}{K\theta} \tilde{\mathbf{D}}_{MN}^g}[\boldsymbol{\phi}_k^t] \\ &= \left(\mathbf{1} + \frac{K}{c\tau\delta} \sum_{j=1}^c |\Lambda_{jk}^s|^2 + \frac{1}{c\tau\theta K} \sum_{j=1}^c |\Lambda_{jk}^g|^2 \right) \left(\boldsymbol{\phi}_k^t - \frac{K}{c\tau\delta} \sum_{j=1}^c \Lambda_{jk}^s (\det(\nabla_{MN} \boldsymbol{\phi}_k^{[\beta]}) \mathbf{T}[\bar{\mathbf{z}}_{k+1}^j, \boldsymbol{\phi}_k^{[\beta]}] - (\Lambda_{jk}^s)^T \boldsymbol{\phi}_k^{[\beta]} \right. \\ & \quad \left. - \bar{\mathbf{z}}_k^j) - \frac{1}{c\tau\theta K} \sum_{j=1}^c \Lambda_{jk}^g (\det(\nabla_{MN} \boldsymbol{\phi}_k^{[\beta]}) \mathbf{T}[K\boldsymbol{\vartheta}_k^j, \boldsymbol{\phi}_k^{[\beta]}] - (\Lambda_{jk}^g)^T \boldsymbol{\phi}_k^{[\beta]} - K\boldsymbol{\vartheta}_{k-1}^j - \bar{\mathbf{z}}_k^j) \right), \end{aligned}$$

where $\Lambda_{jk}^s := \Lambda(\bar{\mathbf{z}}_k^j, \bar{\mathbf{z}}_{k+1}^j, \boldsymbol{\phi}_k^{[\beta]})$ and $\Lambda_{jk}^g := \Lambda(K\boldsymbol{\vartheta}_{k-1}^j + \bar{\mathbf{z}}_k^j, K\boldsymbol{\vartheta}_k^j, \boldsymbol{\phi}_k^{[\beta]})$. The first terms in both brackets are activated only for $k < K$.

The actual minimization $\mathbf{F}_{MN}^{\sigma, K}$ is performed by Algorithm 5. We used the following notation for the extrapolation with $\beta > 0$ of the k^{th} path element in the i^{th} iteration step

$$\begin{aligned} h_k^{[\beta, i]} &= h_k^{[i]} + \beta(h_k^{[i]} - h_k^{[i-1]}), \\ h^{[k, i]} &= (h_{1, \dots, k-1}^{[i+1]}, h_{k, \dots, K}^{[i]}), \\ h^{[\beta, k, i]} &= (h_1^{[i+1]}, \dots, h_{k-1}^{[i+1]}, h_k^{[\beta, i]}, h_{k+1}^{[i]}, \dots, h_K^{[i]}), \end{aligned}$$

while the acceleration $\mathbf{a}^{[\beta, i]}$ is computed with correspondingly updated $\boldsymbol{\phi}^{[\beta, i]}$ values. Furthermore, I^K is the set of fixed indices (cf. (4.3.5)) and we denote by $L[h]$ the Lipschitz constant of the gradient of the function h , which is determined by backtracking [BT09]. The discrete deformations are initialized by the identity deformation, the discrete images by piecewise linear interpolation of the key frame images, while the discrete derivatives are initialized as differences of two consecutive images in the sequence.

4.8 Applications

In what follows, we investigate and discuss qualitative properties of the spline interpolation in the space of images, being aware that the superior temporal smoothness of this interpolation is difficult to appreciate with sequences of still images. For all the examples we use $L=5$ levels in the multi-level approach and $I=250$ iterations of the iPALM algorithm on each level with the extrapolation parameter $\beta = \frac{1}{\sqrt{2}}$. For the first example $M=N=64$, while for the others $M=N=128$. Also, for all examples we used $K=8$. For plotting the resulting images we clip the values to $[0, 1]$.

Algorithm 5: Algorithm for minimizing $\mathbf{F}_{MN}^{\sigma,K}$.

```

1 for  $i = 1$  to  $I$  do
2   for  $k = 1$  to  $K$  do
3     /* update deformation */
4      $\phi_k^{i,t} = \phi_k^{[\beta,i]} - \frac{\lambda}{L[\phi_k^{[i]}]} \nabla \phi_k \left( \sigma K \|W_D(\nabla_{MN} \phi_k^{[\beta,i]})\|_{L_{MN}^1} + \frac{1}{K} \|W_A(\nabla_{MN} \mathbf{a}_k^{[\beta,i]}) + \right.$ 
5        $\left. W_A(\nabla_{MN} \mathbf{a}_{k-1}^{[\beta,i]})\|_{L_{MN}^1} \right)$ ;
6      $\phi_k^{[i+1]} = \text{prox}_{L[\phi_k^{[i]}]}^{\frac{K}{\delta} \tilde{\mathbf{D}}_{MN}^s + \frac{1}{K\theta} \tilde{\mathbf{D}}_{MN}^g}[\phi_k^{i,t}]$ ;
7     /* update derivative */
8      $\bar{\mathbf{z}}_k^{[i+1]} = \bar{\mathbf{z}}_k^{[\beta,i]} - \frac{\nabla_{\bar{\mathbf{z}}_k} \mathbf{F}_{MN}^{\sigma,K,D}[\vartheta^{[k,i]}, \bar{\mathbf{z}}^{[\beta,k,i]}, \phi^{[k+1,i]}]}{L[\bar{\mathbf{z}}_k^{[i]}]}$ ;
9     if  $k \notin I^K$  then
10      /* update image */
11       $\vartheta_k^{[i+1]} = \vartheta_k^{[\beta,i]} - \frac{\nabla_{\vartheta_k} \mathbf{F}_{MN}^{\sigma,K,D}[\vartheta^{[\beta,k,i]}, \bar{\mathbf{z}}^{[k+1,i]}, \phi^{[k+1,i]}]}{L[\vartheta_k^{[i]}]}$ ;

```

All examples in this chapter have been inspired by the numerical experiments from the previous chapter. In Fig. 4.2, we have three Gaussian distributions with different mass as keyframes. In the unbalanced spline case, mass is first annihilated in the first half, and then created again. This helps keep the intensity constant in time, unlike in the classic Wasserstein case. Fig. 4.3 depicts one Gaussian distribution that splits in two and drift apart in different directions. Unlike the experiment in the previous chapter, we are now allowed to have each Gauss distribution have the same mass (and hence, the latter frames will have more mass than the first frames). Once again, the unbalanced transport spline looks more natural, as the density modulation to keep mass constant is no longer necessary. Finally, Fig. 4.4 shows the equivalent "circle-annulus" example from the previous chapter. Once again, a spline is possible to be computed even if keyframes have different masses.

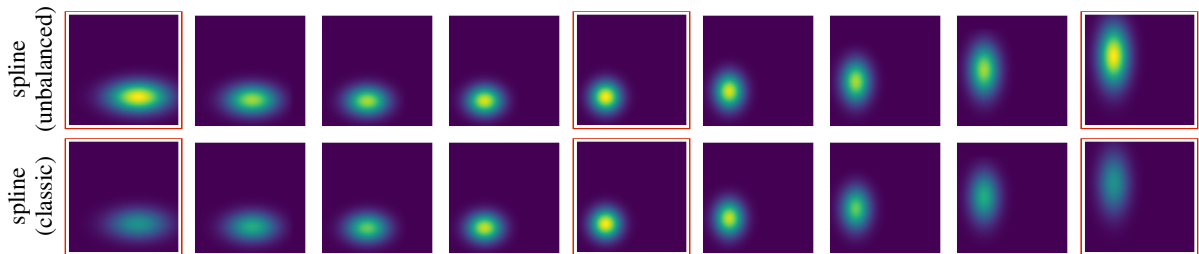


Figure 4.2: Top: Classic Wasserstein spline between three Gaussian distributions. Bottom: Unbalanced transport spline between three re-scaled Gaussian distributions.

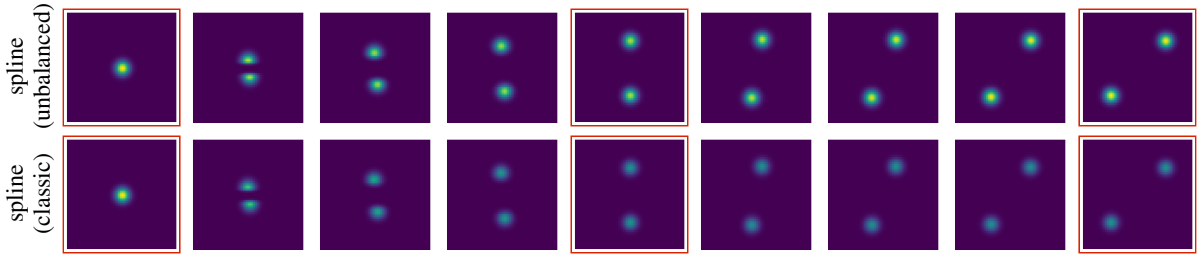


Figure 4.3: Top: From left to right, the key frames represent a single Gaussian, a pair of vertically displaced Gaussians of half the mass, and the vertically displaced configuration rotated by $-\frac{\pi}{4}$. Bottom: From left to right, the key frames represent a single Gaussian, a pair of vertically displaced Gaussians of with each having the same mass as the left, and the vertically displaced configuration rotated by $-\frac{\pi}{4}$.

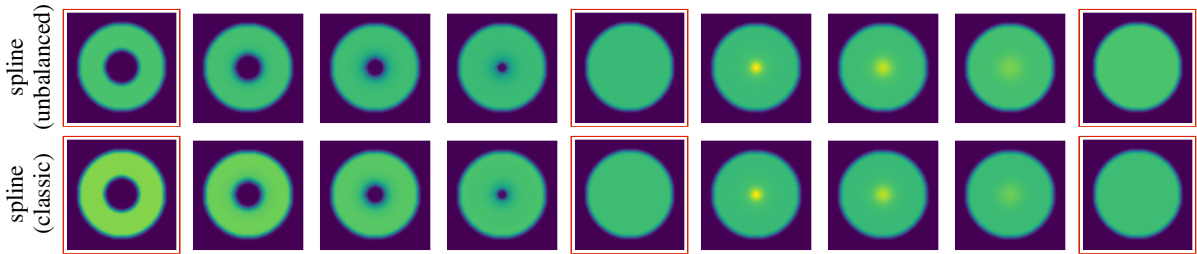


Figure 4.4: Unbalanced transport spline (top) and classic Wasserstein spline (bottom) of three key frames with constant density on an annulus for the first and constant density on a disk for the second and third (framed in red). The keyframes at the top have different masses.

Acknowledgements This work was partially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) via project 211504053 - Collaborative Research Center 1060 and project 390685813 - Hausdorff Center for Mathematics.

5 Entropy-Regularized Optimal Transport in Information Design

5.1 Information Design and the Moment Bayesian Persuasion Problem

In this section we will rigorously formulate the moment Bayesian persuasion problem. Let (D, \mathcal{D}, ν) be a probability space, where $D \subseteq \mathbb{R}^d$ is compact and convex, \mathcal{D} is its Borel σ -algebra, and ν is a probability measure which is assumed to be absolutely continuous with respect to the Lebesgue measure \mathcal{L} . For later usage, we assume throughout this paper that the Radon-Nikodym density $\frac{d\nu}{d\mathcal{L}} \in L^\infty(\mathcal{L})$. The *state of the world* ω is a realization of a random variable that is distributed according to the *prior* ν . An informed *sender* wants to persuade an uninformed *receiver* to take an action that the sender prefers. The receiver's optimal action depends on her information about the state, and initially she only knows that the state of the world is distributed according to ν . The sender, who knows ν and observes the realization of the state of the world ω , may reveal information to the receiver about the realized state. This revelation may often be strategic if the goals of sender and receiver are different. For example, if the state is represented by a one-dimensional random variable, the sender could reveal to the receiver if the realized state is above or below a certain threshold but provide no additional information — this may be better for the sender than revealing all information.

Specifically, the sender commits to an *information policy* (S, π) , where S is a measurable space, $\mathcal{P}(S)$ is the space of probability distributions on S , and $\pi : D \rightarrow \mathcal{P}(S)$. If the sender commits to information policy (S, π) and if the realized state of the world is ω , the receiver observes the information policy and the realization of a random variable with values in S that is drawn according to the probability distribution $\pi(\omega)$. The prior ν and the information policy π induce a joint probability distribution γ on $D \times S$ defined by $\gamma(E \times T) := \int_E \pi(\omega)(T) d\nu(\omega)$ for any measurable sets $E \subseteq D, T \subseteq S$. After seeing a realization $s \in S$, the receiver updates her beliefs about the state of the world according to Bayes' rule to a *posterior belief* given by $\gamma(\cdot \times \{s\} | D \times \{s\}) \in \mathcal{P}(D)$, which is a regular conditional probability measure of γ . This posterior belief determines the receiver's optimal behaviour and implicitly the sender's payoff. The expected value of the posterior belief is called the *posterior mean*; therefore, each information policy induces a probability distribution over posterior means.

As an example, consider the fully revealing information policy given by $S = (D, \mathcal{D})$ and $\pi(\omega) = \delta_\omega$, where δ_ω denotes the Dirac measure at ω . Under the fully revealing information policy, the random variable observed by the receiver equals the realized state of the world with probability 1 and hence the receiver perfectly learns the state of the world. For an example of a partially informative information policy, let $\{B_1, \dots, B_n\}$ be a partition of D , $S = (\{1, \dots, n\}, 2^{\{1, \dots, n\}})$, and $\pi(\omega) = \delta_i$ if $\omega \in B_i$. Under this information policy, the receiver learns in which partition element the realized state lies, and updates her prior belief accordingly.

In the moment Bayesian persuasion model it is assumed that the sender's payoff from inducing a posterior belief depends only on the posterior mean, and is given by an upper semicontinuous function

$\Phi : D \rightarrow \mathbb{R}$. Note that, in reduced form, this formulation allows for the sender's payoff to depend on an action taken by the receiver. To characterize the distributions of posterior means that can be induced by some information policy we use the following stochastic dominance concept.

Definition 5.1.1 (Shaked and Shanthikumar [SS07], Cartier et al. [CFM64] and Phelps [Phe01]). For measures ν and ρ , we say that ν *dominates* ρ in the *convex order* (or that ρ is a *fusion* of ν), denoted by $\nu \succeq \rho$, if $\int \psi d\nu \geq \int \psi d\rho$ for all convex functions ψ for which both integrals exist. We write $\nu \succ \rho$ if ν dominates ρ in the convex order and $\nu \neq \rho$. We denote by $F_\nu = \{\rho : \rho \preceq \nu\}$ the set of fusions of a given measure ν .

Any information policy necessarily reveals weakly less information than the state of the world and the generated distribution of posterior means is dominated in the convex order by the prior. Conversely, for any probability measure $\rho \preceq \nu$ there exists an information policy that generates ρ as its probability measure of posterior means (see Blackwell [Bla53] and Strassen [Str65]).

Therefore, instead of modeling the sender's choice of an information policy we can assume that the sender directly chooses a probability measure that is dominated in the convex order by ν . The sender chooses such a probability measure to maximize her expected payoff, and therefore solves the problem

$$\max_{\lambda \in F_\nu} \int_D \Phi(y) d\lambda(y). \quad (5.1.1)$$

This is the information design problem in the focus of this paper.

To deduce qualitative properties of solutions, let us recall some basic concepts of convex analysis. An *extreme* point of a convex set A is a point $y \in A$ that cannot be represented as a convex combination of two other points in A . I.e. $y \in A$ is an extreme point of A if $y = \alpha w + (1 - \alpha)z$, for $w, z \in A$ and $\alpha \in [0, 1]$ jointly imply that $y = w$ or $y = z$.

The *Krein–Milman Theorem* states that any convex and compact set A in a locally convex space is the closed, convex hull of its extreme points. In particular, such a set has extreme points. The usefulness of extreme points for optimization stems from *Bauer's Maximum Principle*, which states that a convex, upper semicontinuous functional on a non-empty, compact and convex set A of a locally convex space attains its maximum at an extreme point of A . An element y of a convex set A is called *exposed* if there exists a supporting hyperplane H such that $H \cap A = \{y\}$ or, equivalently, there is a continuous, linear functional that attains its unique maximum on A at y . Every exposed point is extreme, but the converse is not true in general.

The set of fusions F_ν appearing in the above maximization problem is convex and compact in the weak* topology of measures. As the objective is linear in the measure λ , a maximum is attained at one of the extreme points of F_ν . It is thus of interest to further explore the structure of extreme and exposed points, and we focus here on those measures that have a finite support.

For any measure ρ on D and a measurable set $B \subseteq D$, we denote by $\rho|_B$ the restriction of ρ to B . As stated in the following theorem a key feature of any extremal measure in F_ν with finite support is that there is a partition of the domain D into convex sets B such that all the original mass restricted to B , $\nu(B)$, remains within B and is fused into a measure $\rho|_B$ whose support is an affinely independent set of points.

Theorem 5.1.2 (Kleiner et al. [KMSW24]). *Let $D \subseteq \mathbb{R}^n$ be compact and convex, and let ν be an absolutely continuous probability measure on D . Suppose that ρ is an extreme point of F_ν with finite support. Then there exists a finite partition \mathcal{P} of D into convex sets such that, for each $B \in \mathcal{P}$, $\rho|_B \preceq \nu|_B$ and $\rho|_B$ has affinely independent support.*

The above result generalizes a unidimensional result found in Kleiner et al. [KMS21] and Arieli et al. [ABSY23]. In fact, more precise information is available about the geometric structure of the following subset of extreme points.

Definition 5.1.3. A measure $\rho \in F_\nu$ is a *Lipschitz-exposed* point of F_ν if there exists a Lipschitz-continuous function $\Phi: D \rightarrow \mathbb{R}$ such that ρ is the unique solution to the problem

$$\max_{\lambda \in F_\nu} \int \Phi(y) d\lambda(y).$$

The following theorem provides a characterization of Lipschitz-exposed points with finite support using Laguerre diagrams, also known as *power diagrams*, see Aurenhammer [Aur87]. Given an n -tuple $\mathbf{X} = (x_1, \dots, x_n)$ of pairwise distinct sites $x_i \in \mathbb{R}^d$ and a weight vector $\mathbf{g} = (g_1, \dots, g_n) \in \mathbb{R}^n$, the associated Laguerre cells $L_i[\mathbf{X}, \mathbf{g}] \in \mathcal{D}$ are the convex polyhedra

$$L_i[\mathbf{X}, \mathbf{g}] := \{y \in \Omega : |y - x_i|^2 - g_i \leq |y - x_j|^2 - g_j \ \forall \ 1 \leq j \leq n\}, \quad (5.1.2)$$

for $i = 1, \dots, n$. The ensemble of all Laguerre cells forms the Laguerre diagram.

Theorem 5.1.4 (Kleiner et al. [KMSW24]). *Let $D \subseteq \mathbb{R}^n$ be compact and convex, and let ν be a probability measure with full support on D that is absolutely continuous with respect to the Lebesgue measure. Let $\rho \in F_\nu$ have finite support. Then ρ is a Lipschitz-exposed point of F_ν if and only if there exists a Laguerre diagram \mathcal{P} of D such that, for all probability measures λ , if $\lambda|_L \preceq \nu|_L$ for all $L \in \mathcal{P}$ and $\text{supp}(\lambda) \subseteq \text{supp}(\rho)$ then $\lambda = \rho$.*

In particular, $\rho \in F_\nu$ is a finitely supported Lipschitz-exposed point if there is a Laguerre diagram such that for each Laguerre cell L with non-vanishing measure $\nu(L)$, the relation $\rho|_L \preceq \nu|_L$ holds and if the support of $\rho|_L$ is affinely independent. Indeed, in this case $\lambda|_L \preceq \nu|_L$ implies that $\lambda|_L$ and $\rho|_L$ have the same mean. In turn, this implies that if $\text{supp}(\lambda) \subseteq \text{supp}(\rho)$ and if the support of ρ is affinely independent then $\rho|_L = \lambda|_L$.

For any compact convex set in a normed vector space, the set of exposed points is dense in the set of extreme points of F_ν (Klee, [KJ58]). Therefore, to optimize a continuous linear objective functional on such a set, it would be sufficient to optimize over the closure of exposed points. In the following, we will focus on the subset of exposed points that are Lipschitz-exposed and where, on each cell L of the corresponding Laguerre diagram, the mass within the respective cell is fused to a unique mass point. These extreme points have received considerable attention in the Economics literature in one-dimensional settings (see for example Dworczak and Martini [DM19] or Ivanov [Iva21]). They are simple to use in practice because they can be implemented by only revealing, for a partition of the state space into convex sets, in which partition element the realized state lies. Finally, the relation to Laguerre diagrams makes these exposed points numerically tractable. Our computational method derived below focuses on solving the sender's problem among such exposed points.

5.2 The optimization task

Motivated by Theorem 5.1.4 we discuss in this section the optimization problem (5.1.1) of moment Bayesian persuasion as an optimization over Laguerre partitions along with a relaxation strategy. At first, we study the case where the sender commits to an information policy induced by general partitions.

To this end, we consider a probability space $(\Omega, \mathcal{D}, \nu)$ on a compact and convex domain $\Omega \subset \mathbb{R}^d$ for $d \geq 2$, equipped with a probability measure $\nu \in \mathcal{P}(\Omega)$ on the Borel- σ -algebra \mathcal{D} , which is assumed to be absolutely continuous with respect to the Lebesgue measure \mathcal{L} with Radon-Nikodym density $\frac{d\nu}{d\mathcal{L}} \in L^\infty(\mathcal{L})$. For fixed $n \in \mathbb{N}$, let $(B_i)_{i=1,\dots,n} \subset \mathcal{D}$ be a ν -partition of Ω , i.e. $\nu(B_i \cap B_j) = 0$ for $i \neq j$, and $\nu(\cup_{i=1}^n B_i) = 1$. Let $(S, \mathcal{S}) := (\{1, \dots, n\}, 2^{\{1, \dots, n\}})$, and assume the sender commits to the information policy $\pi : D \rightarrow \mathcal{P}(S)$, with $\pi(\omega) = \delta_i \in \mathcal{P}(S)$ if and only if $\omega \in B_i$. After receiving the signal realization $i \in S$, the receiver updates her belief to the posterior $\gamma(\cdot \times \{i\} \mid D \times \{i\})$. It holds that

$$\gamma(B \times \{i\} \mid D \times \{i\}) = \frac{\int_B \pi(\omega)(\{i\}) d\nu(\omega)}{\int_D \pi(\omega)(\{i\}) d\nu(\omega)} = \frac{\int_B \mathbb{1}_{B_i}(\omega) d\nu(\omega)}{\int_D \mathbb{1}_{B_i}(\omega) d\nu(\omega)} = \frac{\nu(B \cap B_i)}{\nu(B_i)} = \nu(B \mid B_i).$$

For any non- ν -null set $B \in \mathcal{D}$, we define the ν -barycenter

$$b[B] := \frac{\int_B y d\nu(y)}{\nu(B)}.$$

The ν -barycenter of B_i coincides then with the mean of the posterior $\nu(\cdot \mid B_i)$. To conclude, the posterior mean distribution generated by the information policy π is given by $\sum_{i=1}^n \nu(B_i) \delta_{b[B_i]} \in F_\nu$. In this case, for a continuous function $\Phi : \Omega \rightarrow \mathbb{R}$ the functional (5.1.1) takes the following explicit form:

$$\sum_{\substack{i=1,\dots,n \\ \nu(B_i) > 0}} \nu(B_i) \Phi(b[B_i]). \quad (5.2.1)$$

Now, in the light of Section 5.1 we focus on partitions described by Laguerre diagrams. In this case, we may optimize the cost functional (5.2.1) directly on the parameters (\mathbf{X}, \mathbf{g}) describing Laguerre cells

$$\mathcal{F}_n[\mathbf{X}, \mathbf{g}] := \sum_{\substack{i=1,\dots,n \\ m_i[\mathbf{X}, \mathbf{g}] > 0}} m_i[\mathbf{X}, \mathbf{g}] \Phi(b_i[\mathbf{X}, \mathbf{g}]), \quad (5.2.2)$$

where $m_i[\mathbf{X}, \mathbf{g}] := \nu(L_i[\mathbf{X}, \mathbf{g}])$, and $b_i[\mathbf{X}, \mathbf{g}] := b[L_i[\mathbf{X}, \mathbf{g}]]$ (cf. equation (5.1.2)), to be maximized over n -tuples of pairwise distinct sites $\mathcal{X}_n := \{\mathbf{X} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n : x_i \in \mathbb{R}^d, x_i \neq x_j \text{ for } i \neq j\}$ and weight vectors $\mathbf{g} \in \mathbb{R}^n$. Note that Laguerre cells might be sets of vanishing measure. Since Laguerre cells are by definition a ν -partition of Ω , the set of associated characteristic functions $\chi_{L_i[\mathbf{X}, \mathbf{g}]}$ of the Laguerre cells $L_i[\mathbf{X}, \mathbf{g}]$ for $i = 1, \dots, n$ forms a partition of unity of Ω , i.e. $\sum_{i=1,\dots,n} \chi_{L_i[\mathbf{X}, \mathbf{g}]} = 1$ a.e. in Ω . We define the power diagram associated with an n -tuple of sites \mathbf{X} and a weight vector \mathbf{g} as

$$L[\mathbf{X}, \mathbf{g}] := (\chi_{L_1[\mathbf{X}, \mathbf{g}]}, \dots, \chi_{L_n[\mathbf{X}, \mathbf{g}]}). \quad (5.2.3)$$

Along maximizing sequences for the functional $\mathcal{F}_n[\cdot, \cdot]$ it might happen that subsets of Laguerre cells collapse, or that sites as well as weights diverge. Here, the notion of power diagram as partitions of unity helps to deal with these degenerate cases. At first, we obtain the following relative compactness result:

Lemma 5.2.1. *The set $\mathcal{L} := L(\mathcal{X}_n \times \mathbb{R}^n)$ is relatively compact in $L^1(\nu)^n$ and any limit of a converging sequences in \mathcal{L} is a partition of unity a.e. on Ω .*

Proof. Let ν be trivially extended onto \mathbb{R}^d with density 0 outside Ω . Each non-empty Laguerre cell $L_i[\mathbf{X}, \mathbf{g}]$ is convex and its boundary in the interior of Ω is polygonal and consists of at most $n - 1$ planar

interfaces. Each of these interior interfaces has at most a \mathcal{H}^{d-1} measure $\text{diam}(\Omega)^{d-1}$ inside of Ω . Thus, for $h \in \mathbb{R}$ we observe that

$$\|\chi_{L_i[\mathbf{X}, \mathbf{g}]}(\cdot - h) - \chi_{L_i[\mathbf{X}, \mathbf{g}]\|_{L^1(\nu)} \leq (n-1) \left\| \frac{d\nu}{dx} \right\|_{\infty} \text{diam}(\Omega)^{d-1} h.$$

This, together with the compactness of Ω , implies by the Fréchet-Kolmogorov theorem the relative compactness of $\mathbf{L}(\mathcal{X}_n \times \mathbb{R}^n)$ in $L^1(\nu)$. Let $(\mathbf{X}^k)_k \subset \mathcal{X}_n$, and $(\mathbf{g}^k)_k \subset \mathbb{R}^n$. Since $\chi_{L_i[\mathbf{X}^k, \mathbf{g}^k]} = 0$ outside Ω any $L^1(\nu)$ -limit of $(\chi_{L_i[\mathbf{X}^k, \mathbf{g}^k]})_{i=1, \dots, n}$ for $k \rightarrow \infty$ is of the form $(\chi^1, \dots, \chi^n) \in L^1(\nu; \{0, 1\})^n$ and $\text{supp} \chi_i \subseteq \Omega$. Finally, $\|\sum_{i=1, \dots, n} \chi^i - 1\|_{L^1} = \lim_{k \rightarrow \infty} \|\sum_{i=1, \dots, n} \chi_{L_i[\mathbf{X}^k, \mathbf{g}^k]} - 1\|_{L^1} = 0$ and thus (χ^1, \dots, χ^n) is a partition of unity. \square

Let us denote by \mathcal{L}^c the closure of $\mathcal{L} = \mathbf{L}(\mathcal{X}_n \times \mathbb{R}^n) \in L^1(\nu; \{0, 1\})^n$. For $\mathbf{L}^c = (\chi^1, \dots, \chi^n) \in \mathcal{L}^c$ we define the relaxed functional

$$\mathcal{F}_n^c[\mathbf{L}^c] := \sum_{\substack{i=1, \dots, n \\ m[\chi^i] > 0}} m[\chi^i] \Phi(b[\chi^i]) \quad (5.2.4)$$

where we define with a slight misuse of notation the ν -mass $m[\chi] := \int_{\Omega} \chi \, d\nu$ of a characteristic function χ and its ν -barycenter $b[\chi] := m[\chi]^{-1} \int_{\Omega} x \chi \, d\nu$ for $m[\chi] > 0$. Let us remark that some of the χ^i in \mathbf{L}^c (but not all) might have zero mass $m[\chi^i]$. For this relaxed functional we obtain the following existence result of a maximum.

Theorem 5.2.2. *Assume Φ to be a upper semicontinuous function on Ω . Then \mathcal{F}_n^c attains its maximum on \mathcal{L}^c .*

Proof. At first, we recall that upper semicontinuous functions on compact domains are bounded from above. Let $\bar{\Phi}$ denote the maximum of Φ on Ω , which exists due to the upper semicontinuity of Φ . Let

$$(\mathbf{L}_k^c = (\chi_k^1, \dots, \chi_k^n))_{k \in \mathbb{N}} \subset \mathcal{L}^c$$

be a maximizing sequence of (5.2.4). Due to the compactness of \mathcal{L}^c , we obtain that, up to the selection of a subsequence, $(\chi_k^1, \dots, \chi_k^n)$ converges in $L^1(\nu)$ to a limit $\mathbf{L}^c = (\chi^1, \dots, \chi^n)$. For all $i \in \{1, \dots, n\}$ with $m[\chi^i] = 0$ it holds that $m[\chi_k^i] \rightarrow 0$ for $k \rightarrow \infty$ and thus $m[\chi_k^i] \Phi(b[\chi_k^i]) \leq m[\chi_k^i] \bar{\Phi} \rightarrow 0$. For $m[\chi^i] > 0$ the sequence of barycenters $(b[\chi_k^i])_{k \in \mathbb{N}}$ converges to $b[\chi^i]$. Taking into account the upper semicontinuity of Φ , this implies that $\limsup_{k \rightarrow \infty} \Phi(b[\chi_k^i]) \leq \Phi(b[\chi^i])$. Altogether, we obtain

$$\limsup_{k \rightarrow \infty} \mathcal{F}_n^c[\mathbf{L}_k^c] = \limsup_{k \rightarrow \infty} \sum_{\substack{i=1, \dots, n \\ m[\chi_k^i] > 0}} m[\chi_k^i] \Phi(b[\chi_k^i]) \leq \sum_{\substack{i=1, \dots, n \\ m[\chi^i] > 0}} m[\chi^i] \Phi(b[\chi^i]) = \mathcal{F}_n^c[\mathbf{L}^c].$$

Hence, the relaxed functional \mathcal{F}_n^c attains its maximum on \mathcal{L}^c at \mathbf{L}^c . \square

Example 5.2.3. Let $d = 2$, $D = [0, 1]^2$, $\nu = \mathcal{L}$ and $\Phi : D \rightarrow \mathbb{R}$, $y \mapsto \Phi(y) := -\prod_{i=1}^3 |y - y_i|^2$, where $y_1 = (\frac{1}{2}, \frac{1}{4})$, $y_2 = (\frac{3}{4}, \frac{3}{4})$, and $y_3 = (\frac{1}{4}, \frac{3}{4})$. Then, the unique maximizer of (5.2.4) lies in $\mathcal{L}^c \setminus \mathcal{L}$. Indeed, $\max_{y \in D} \Phi(y) = 0$, so we obtain

$$\mathcal{F}_n^c[\mathbf{L}^c] = \sum_{\substack{i=1, \dots, n \\ m[\chi^i] > 0}} m[\chi^i] \Phi(b[\chi^i]) \leq 0,$$

for all $\mathbf{L}^c \in \mathcal{L}^c$. On the other hand, let $B_1 := [0, 1] \times [0, \frac{1}{2}]$, $B_2 := [\frac{1}{2}, 1] \times [\frac{1}{2}, 1]$, and $B_3 := [0, \frac{1}{2}] \times [\frac{1}{2}, 1]$. Then, $\mathbf{B} := (\chi_{B_i})_{i=1,2,3} \in \mathcal{L}^c \setminus \mathcal{L}$. Clearly, as the interfaces $B_1 \cap B_2$ and $B_1 \cap B_3$ build an angle of π , then $\mathbf{B} \notin \mathcal{L}$. However, let $x_1 = (\frac{1}{2}, \frac{1}{4})$, $x_2^\epsilon = (\frac{1}{2} + \epsilon, \frac{3}{4})$, $x_3^\epsilon = (\frac{1}{2} - \epsilon, \frac{3}{4})$, and $\mathbf{g} = \mathbf{0}$. Then, $\mathbf{L}((x_1, x_2^\epsilon, x_3^\epsilon)) \rightarrow \mathbf{B}$ in $L^1(\nu)$ as ϵ approaches zero. Hence, $\mathbf{B} \in \mathcal{L}^c \setminus \mathcal{L}$, and by the definition of Φ , it holds $\mathcal{F}_n^c[\mathbf{B}] = 0$, so by the previous inequality it is a maximizer. Moreover, it is the unique maximizer, since any maximizing measure needs to be supported in $\{y_1, y_2, y_3\}$. Let $\mu = \sum_{i=1}^3 m_i \delta_{y_i}$ be any such measure. Then, it must hold

$$\begin{aligned} m_1 + m_2 + m_3 &= 1 \\ m_1 y_{11} + m_2 y_{21} + m_3 y_{31} &= \frac{1}{2} \\ m_1 y_{12} + m_2 y_{22} + m_3 y_{32} &= \frac{1}{2}, \end{aligned}$$

with the unique solution $m_1 = \frac{1}{2}, m_2 = m_3 = \frac{1}{4}$. The only element in \mathcal{L}^c satisfying the mass and barycenter conditions is then given by \mathbf{B} (up to ν -null sets).

Remark 5.2.4. The relation to the concept of stochastic dominance is as follows. As an upper semicontinuous function, Φ is a ν -measurable function on Ω . Given a polyhedral partition of Ω into Laguerre cells $(L_i[\mathbf{X}, \mathbf{g}])_{i=1, \dots, n}$, one may collapse the mass of each cell at its barycenter. This induces an atomic probability measure

$$\rho := \sum_{\substack{i=1, \dots, n \\ m_i[\mathbf{X}, \mathbf{g}] > 0}} m_i[\mathbf{X}, \mathbf{g}] \delta_{b_i[\mathbf{X}, \mathbf{g}]}$$

on Ω . Recall that a probability measure ν dominates a probability measure ρ in convex order if and only if $\mathbb{E}_{X \sim \nu}[\Phi(X)] \geq \mathbb{E}_{X \sim \rho}[\Phi(X)]$ for all convex functions $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}$ such that both expectations exist. By Jensen's inequality, this indeed holds for the initial probability measure ν and the atomic probability measure ρ considered here. Hence, in our ansatz we consider the expected value of the given function Φ with respect to an atomic measure associated with some power diagram \mathbf{L}^c and maximize this expected value over all atomic measures induced by power diagrams \mathbf{L}^c .

To avoid the relaxation, one might consider hard constraints to ensure that pairwise distances between sites and cell masses do not vanish in the limit along a maximizing sequence. Alternatively, a penalty formulation can be used as a more robust and effective alternative. To this end, we define for a penalty parameter $\eta > 0$

$$\mathcal{F}_n^\eta(\mathbf{X}, \mathbf{g}) = \mathcal{F}_n(\mathbf{X}, \mathbf{g}) - \eta \mathcal{R}_n(\mathbf{X}, \mathbf{g}), \quad (5.2.5)$$

where the penalty term \mathcal{R}_n is given as

$$\mathcal{R}_n(\mathbf{X}, \mathbf{g}) := \sum_{\substack{i=1, \dots, n \\ m_i[\mathbf{X}, \mathbf{g}] > 0}} \int_{\Omega} |y - x_i|^2 \chi_{L_i[\mathbf{X}, \mathbf{g}]}(y) d\nu(y) + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j \\ m_i[\mathbf{X}, \mathbf{g}], m_j[\mathbf{X}, \mathbf{g}] > 0}} \frac{m_i[\mathbf{X}, \mathbf{g}] m_j[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2}. \quad (5.2.6)$$

Here, the cell masses as scaling factors and the characteristic functions of cells as weight functions are to be understood as the natural scaling of the corresponding penalty terms. In particular, we observe a stronger penalization of the drift of sites away from the associated cell and the fusion of pairs of sites in case of larger cell masses.

It might happen that in the limit along a maximizing sequence $(\mathbf{X}^N, \mathbf{g}^N)_N \subset \mathcal{X}_l \times \mathbb{R}^l$ of \mathcal{F}_l^η for some $l \in \mathbb{N}$ Laguerre cells collapse and the effective number of cells decreases. The following existence theorem takes this into account.

Theorem 5.2.5. *Let us assume that $\Phi \in C(\Omega)$ with maximum $\bar{\Phi}$ on D . Then, for given number of sites $l \in \mathbb{N}$ there exists an $n \leq l$, such that a maximizer $(\mathbf{X}^*, \mathbf{g}^*)$ of \mathcal{F}_n^η exists with $m_i[\mathbf{X}^*, \mathbf{g}^*] > 0$ for $i = 1, \dots, n$ and $\mathcal{F}_n^\eta(\mathbf{X}^*, \mathbf{g}^*) \geq \sup_{(\mathbf{X}, \mathbf{g})} \mathcal{F}_l^\eta(\mathbf{X}, \mathbf{g})$.*

Proof. Consider a maximizing sequence for \mathcal{F}_l^η . Since $m_i[\mathbf{X}^N, \mathbf{g}^N] \in [0, 1]$, and $b_i[\mathbf{X}^N, \mathbf{g}^N] \in \Omega$ for $m_i[\mathbf{X}^N, \mathbf{g}^N] > 0$, we may assume that, up to the selection of a subsequence, the sequences of masses and barycenters converge to limits m_i^* and b_i^* , respectively, and that there exists an $n \leq l$ with $m_i^* > 0$ if and only if $i \leq n$. Then, we obtain

$$\begin{aligned}
\sup_{(\mathbf{X}, \mathbf{g})} \mathcal{F}_l^\eta(\mathbf{X}, \mathbf{g}) &= \lim_{N \rightarrow \infty} \mathcal{F}_l^\eta(\mathbf{X}^N, \mathbf{g}^N) \leq \limsup_{N \rightarrow \infty} \mathcal{F}_l(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_l(\mathbf{X}^N, \mathbf{g}^N) \\
&\leq \limsup_{N \rightarrow \infty} \mathcal{F}_n(\mathbf{X}^N, \mathbf{g}^N) + \limsup_{N \rightarrow \infty} \sum_{\substack{i=n+1, \dots, l \\ m_i[\mathbf{X}^N, \mathbf{g}^N] > 0}} m_i[\mathbf{X}^N, \mathbf{g}^N] \Phi(b_i[\mathbf{X}^N, \mathbf{g}^N]) \\
&\quad - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} (\mathcal{R}_l(\mathbf{X}^N, \mathbf{g}^N) - \mathcal{R}_n(\mathbf{X}^N, \mathbf{g}^N)) \\
&= \lim_{N \rightarrow \infty} \mathcal{F}_n(\mathbf{X}^N, \mathbf{g}^N) + 0 - \eta \lim_{N \rightarrow \infty} \mathcal{R}_n(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} (\mathcal{R}_l(\mathbf{X}^N, \mathbf{g}^N) \\
&\quad - \mathcal{R}_n(\mathbf{X}^N, \mathbf{g}^N)) \\
&\leq \lim_{N \rightarrow \infty} \mathcal{F}_n(\mathbf{X}^N, \mathbf{g}^N) - \eta \lim_{N \rightarrow \infty} \mathcal{R}_n(\mathbf{X}^N, \mathbf{g}^N) = \lim_{N \rightarrow \infty} \mathcal{F}_n(\mathbf{X}^N, \mathbf{g}^N) - \eta \mathcal{R}_n(\mathbf{X}^N, \mathbf{g}^N) \\
&= \lim_{N \rightarrow \infty} \mathcal{F}_n^\eta(\mathbf{X}^N, \mathbf{g}^N).
\end{aligned}$$

Hence, if it exists, a maximizer of \mathcal{F}_n^η attains a greater or equal objective value than a maximizer of \mathcal{F}_l^η .

Now, let $(\mathbf{X}^N, \mathbf{g}^N)_N \subset \mathcal{X}_n \times \mathbb{R}^n$ be a maximizing sequence for \mathcal{F}_n^η and in analogy to before assume that $m_i[\mathbf{X}^N, \mathbf{g}^N]$ and $b_i[\mathbf{X}^N, \mathbf{g}^N]$ have limits $m_i^* > 0$, and $b_i^* \in D$ for $i = 1, \dots, n$, that $\mathcal{F}_n^\eta(\mathbf{X}^N, \mathbf{g}^N)$ is monotonically increasing in N and that $m_i[\mathbf{X}^N, \mathbf{g}^N] \geq \frac{1}{2}m_i^*$. Consequently, for $\bar{\Phi}$ being the maximal value of Φ on D the estimate

$$\begin{aligned}
\mathcal{F}_n^\eta(\mathbf{X}^0, \mathbf{g}^0) &\leq \nu(D)\bar{\Phi} - \eta \left(\sum_{i=1, \dots, n} \int_{\Omega} |y - x_i^N|^2 \chi_{L_i[\mathbf{X}^N, \mathbf{g}^N]} d\nu(y) + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i[\mathbf{X}^N, \mathbf{g}^N] m_j[\mathbf{X}^N, \mathbf{g}^N]}{|x_i^N - x_j^N|^2} \right) \\
&\leq \bar{\Phi} - \eta \left(\frac{1}{2} \sum_{i=1, \dots, n} \text{dist}^2(x_i^N, \Omega) m_i^* + \frac{1}{4} \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^* m_j^*}{|x_i^N - x_j^N|^2} \right) \tag{5.2.7}
\end{aligned}$$

is obtained. This implies the following a priori bounds:

$$\text{dist}(x_i^N, \Omega) \leq \sqrt{\frac{2(\bar{\Phi} - \mathcal{F}_n^\eta(\mathbf{X}^0, \mathbf{g}^0))}{\eta m_i^*}}, \quad |x_i^N - x_j^N| \geq \sqrt{\frac{\eta m_i^* m_j^*}{4(\bar{\Phi} - \mathcal{F}_n^\eta(\mathbf{X}^0, \mathbf{g}^0))}}$$

for all N . For the uniform bound on \mathbf{g} , recall that \mathbf{g} and $\mathbf{g} + \lambda \mathbf{1}_n$ both induce the same Laguerre diagram. Hence, we may assume without loss of generality that $\mathbf{g}_1^N = 0$ for all $N \in \mathbb{N}$. Then, $\lim_{N \rightarrow \infty} \mathbf{g}_j^N = \infty$ for $j = 2, \dots, n$ would imply $\lim_{N \rightarrow \infty} m_1[\mathbf{X}^N, \mathbf{g}^N] = 0$, which contradicts our choice of n . Similarly, $\lim_{N \rightarrow \infty} \mathbf{g}_j = -\infty$ would imply $\lim_{N \rightarrow \infty} m_j[\mathbf{X}^N, \mathbf{g}^N] = 0$, which again is a contradiction. Hence, $|\mathbf{g}^N| \leq C$ for some $C > 0$ and all $N \in \mathbb{N}$. Finally, given these a priori bounds, the existence of a maximizer of \mathcal{F}_n^η follows directly from the Weierstraß extreme value theorem. \square

5.3 Semi-discrete optimal transport revisited

In the previous section we stated existence results of maximizers of the sender's revenue over the class of Laguerre partitions. To compute these optimal Laguerre partitions we recall in this section the connection of Laguerre diagrams to solutions of semi-discrete optimal transport problems and the associated dual formulation.

In optimal transport theory, one considers optimal couplings $\Pi \in \mathcal{P}(\Omega \times \Omega) \in U(\mu, \nu)$ of two probability measures $\mu, \nu \in \mathcal{P}(\Omega)$. Here, $\mathcal{U}(\mu, \nu)$ is the set $\Pi \in \mathcal{P}(\Omega \times \Omega)$ with $\Pi(A \times \Omega) = \mu(A)$ and $\Pi(\Omega \times A) = \nu(A)$ for all Borel sets $A \subset \Omega$. For the given cost function $(x, y) \mapsto |x - y|^2$ on $\Omega \times \Omega$ measuring the cost of transport from x to y on Ω , a coupling Π is optimal if it minimizes

$$\mathcal{W}^2[\mu, \nu] := \inf_{\Pi \in \mathcal{U}(\mu, \nu)} \int_{\Omega \times \Omega} |x - y|^2 d\Pi(x, y). \quad (5.3.1)$$

The functional $\mathcal{W}(\mu, \nu)$ is called the 2-Wasserstein distance and defines a metric between the probability measures μ and ν in $\mathcal{P}(\Omega)$ (cf. [San15, Chapter 5]). In semi-discrete optimal transport the measure μ is assumed to be a discrete (empirical) probability measure, i.e.

$$\mu := \sum_{i=1, \dots, n} m_i \delta_{x_i}$$

with $\sum_{i=1, \dots, n} m_i = 1$ and $m_i \geq 0$. The minimization of (5.3.1) is a constrained linear minimization problem and denoted the primal Kantorovich problem. As such, it can naturally be paired with a constrained linear maximization problem as the dual problem, the dual Kantorovich problem [AG13, Section 6.1]. We obtain

$$\mathcal{W}^2[\mu, \nu] = \sup_{(f, \mathbf{g}) \in \mathcal{R}} \mathbf{m} \cdot \mathbf{g} + \int_{\Omega} f(y) d\nu(y), \quad (5.3.2)$$

where $\mathbf{m} := (m_1, \dots, m_n)$, $\mathbf{g} := (g_1, \dots, g_n)$, and

$$\mathcal{R} := \{(f, \mathbf{g}) \in C(\Omega) \times \mathbb{R}^n : f(y) + g_i \leq |x_i - y|^2 \text{ for all } i = 1, \dots, n\}.$$

For given $\mathbf{g} \in \mathbb{R}^n$ we obtain for the optimal f which is consistent with the constrained condition $(f, \mathbf{g}) \in \mathcal{R}$ that $f(y) = \mathbf{g}^C(y)$ for all $y \in D$ with

$$\mathbf{g}^C(y) := \min_{1 \leq i \leq n} |y - x_i|^2 - g_i. \quad (5.3.3)$$

Given the C -transform $\mathbf{g}^C : \Omega \rightarrow \mathbb{R}^D$ one can reformulate the dual Kantorovich problem (5.3.2) as the unconstrained convex program

$$\mathcal{W}^2[\mu, \nu] = \max_{\mathbf{g} \in \mathbb{R}^n} \mathcal{D}[\mathbf{g}] \quad (5.3.4)$$

$$\mathcal{D}[\mathbf{g}] := \int_{\Omega} \mathbf{g}^C(y) d\nu(y) + \mathbf{m} \cdot \mathbf{g} = \sum_{i=1, \dots, n} \int_{L_i[\mathbf{X}, \mathbf{g}]} (|y - x_i|^2 - g_i) d\nu(y) + \mathbf{m} \cdot \mathbf{g}, \quad (5.3.5)$$

where $L_i[\mathbf{X}, \mathbf{g}]$ are the Laguerre cells associated with the weight vector \mathbf{g} and the vector of fixed sites \mathbf{X} defined in (5.1.2). Hence, solving (5.3.5) for ν and \mathbf{m} consists in finding a Laguerre cell partition of Ω

described via the weight vector $\mathbf{g} \in \mathbb{R}^n$ with cells centered at the given \mathbf{X} , and with $\mathbf{m} = (m_i)_{i=0,\dots,n}$ and $m_i = m(L_i[\mathbf{X}, \mathbf{g}])$ for all $i \in \{1, \dots, n\}$. In what follows, we will need to differentiate the functions \mathcal{F}_n^η defined in (5.2.2) with respect to the weights g_j . For the differentiation of $\mathcal{D}(\mathbf{g})$ we obtain

$$\partial_{g_j} \mathcal{D}[\mathbf{g}] = - \int_{\Omega} \chi_{L_j[\mathbf{X}, \mathbf{g}]} d\nu(y) + m_j. \quad (5.3.6)$$

If \mathbf{g}^* is optimal for (5.3.4), it holds

$$m_j = \int_{\Omega} \chi_{L_j[\mathbf{X}, \mathbf{g}]} d\nu(y). \quad (5.3.7)$$

5.4 Entropy regularization

Let us recall that our goal is to maximize the function \mathcal{F}_n^η via an optimization of the Laguerre cells $L_j[\mathbf{X}, \mathbf{g}]$ described in terms of the sites x_j and the weights g_j for $j = 1, \dots, n$. In general, changing the sites and the weights fosters topological changes in the diagram's topology, leading to a challenging combinatorial optimization problem. To avoid this, we introduce in this section the entropy relaxed optimal transport formulation [PC⁺19] of the associated semi-discrete optimal transport described in Section 5.2, which will allow us to derive a computationally efficient algorithm to optimize over Laguerre partitions and thus solve the moment Bayesian persuasion problem numerically. An extensive overview on entropy regularization of optimal transport is also given by Chewi et. al. in [CNWR24].

We begin by considering the regularized Wasserstein distance

$$\mathcal{W}_\varepsilon^2[\mu, \nu] := \inf_{\Pi \in \mathcal{U}(\mu, \nu)} \int_{\{x_1, \dots, x_n\} \times \Omega} c(x, y) d\Pi(x, y) + \varepsilon \text{KL}[\Pi|\xi] \quad (5.4.1)$$

with transport cost $c(x, y) = |x - y|^2$ from x to y and for a regularization parameter $\varepsilon > 0$. We consider a measure $\xi \in \mathcal{P}(\{x_1, \dots, x_n\} \times \Omega)$ with $\text{supp} \mu \otimes \nu \subseteq \text{supp} \xi$ for μ and ν as above, and we define the Kullback-Leibler divergence

$$\text{KL}[\Pi|\xi] := \int_{\{x_1, \dots, x_n\} \times \Omega} \log \left(\frac{d\Pi}{d\xi}(x, y) \right) d\Pi(x, y) + (d\xi(x, y) - d\Pi(x, y)) \quad (5.4.2)$$

between the measures Π and ξ in the case that Π is absolutely continuous with respect to ξ , and otherwise $\text{KL}[\Pi|\xi] := \infty$. The Kullback-Leibler divergence is a concave functional measuring the dissimilarity of the measures Π and ξ and acts here as a regularizing entropy functional. The standard choice for ξ is $\xi = \mu \otimes \nu$. In fact, as long as the above support property holds the functional in (5.4.1) only changes by an additive constant and hence, the minimizer remains the same. To simplify the optimization algorithm for an entropy regularization of the functional \mathcal{F}_l^η defined in (5.2.5) $\xi = \mathcal{N} \otimes \mathcal{L}$ is a particularly suitable choice where $\mathcal{N} = \sum_{j=1, \dots, n} \delta_{x_j}$ is the counting measure on the support of μ and \mathcal{L} the Lebesgue measure. For this choice we proceed as follows.

Associated with the constraint optimization problem (5.4.1) is the Lagrangian

$$\begin{aligned} \mathcal{L}^\varepsilon(\Pi, f, g) = & \int_{\Omega \times \Omega} c(x, y) d\Pi(x, y) + \varepsilon \int_{\Omega \times \Omega} \left(\log \frac{d\Pi}{d\xi} - 1 \right) (x, y) d\Pi(x, y) \\ & - \left(\int_{\Omega \times \Omega} f(y) d\Pi(x, y) - \int_{\Omega} f(y) d\nu(y) \right) - \left(\int_{\Omega \times \Omega} g(x) d\Pi(x, y) - \int_{\Omega} g(x) d\mu(x) \right). \end{aligned}$$

Now, assume that Π has the density p with respect to ξ , i.e. $d\Pi(x, y) = p(x, y) d\xi(x, y)$. Then, a necessary condition of a saddle point is that the derivative of \mathcal{L}^ε vanishes in all directions $q : (x, y) \mapsto q(x, y)$. Hence,

$$\begin{aligned} 0 = \left(\frac{\partial \mathcal{L}^\varepsilon}{\partial p}(p) \right) (q) &= \int_{\Omega \times \Omega} c(x, y) q(x, y) d\xi(x, y) + \varepsilon \int_{\Omega \times \Omega} (\log p(x, y) - 1) q(x, y) d\xi(x, y) \\ &\quad + \varepsilon \int_{\Omega \times \Omega} q(x, y) d\xi(x, y) \\ &\quad - \int_{\Omega \times \Omega} f(y) q(x, y) d\xi(x, y) - \int_{\Omega \times \Omega} g(x) q(x, y) d\xi(x, y). \end{aligned} \quad (5.4.3)$$

Hence, we obtain $c(x, y) + \varepsilon \log(p(x, y)) - f(y) - g(x) = 0$ pointwise, or equivalently

$$p(x, y) = \frac{d\Pi}{d\xi}(x, y) = \exp \left(\frac{-c(x, y) + g(x) + f(y)}{\varepsilon} \right).$$

and

$$\begin{aligned} \mathcal{L}^\varepsilon(\Pi(f, g), f, g) &= \int_{\Omega} f(y) d\nu(y) + \int_{\Omega} g(x) d\mu(x) - \varepsilon \int_{\Omega \times \Omega} \exp \left(\frac{-c(x, y) + g(x) + f(y)}{\varepsilon} \right) d\xi(x, y) \\ &= \int_{\Omega} f(y) d\nu(y) + \mathbf{m} \cdot \mathbf{g} - \varepsilon \int_{\Omega} \sum_{j=1, \dots, n} \exp \left(\frac{-c(x_j, y) + g(x_j) + f(y)}{\varepsilon} \right) d\mathcal{L}(y) \end{aligned}$$

Finally, we obtain a dual, entropy regularized, unconstrained formulation for the Wasserstein functional

$$\begin{aligned} \mathcal{W}_\varepsilon^2[\mu, \nu] &:= \sup_{\substack{f \in C(\Omega) \\ \mathbf{g} \in \mathbb{R}^n}} \mathcal{L}^\varepsilon(\Pi(f, g), f, g) \\ &= \sup_{\substack{f \in C(\Omega) \\ \mathbf{g} \in \mathbb{R}^n}} \left(\mathbf{m} \cdot \mathbf{g} + \int_{\Omega} f(y) d\nu(y) - \varepsilon \int_{\Omega} \sum_{j=1, \dots, n} \exp \left(\frac{-|y - x_j|^2 + f(y) + \mathbf{g}_j}{\varepsilon} \right) d\mathcal{L}(y) \right) \end{aligned} \quad (5.4.4)$$

Here, the exponent $\exp(\varepsilon^{-1}(f(y) + \mathbf{g}_j - |y - x_j|^2))$ acts as a soft penalty in place of the original hard constraint $f(y) + \mathbf{g}_j \leq |y - x_j|^2$. The optimal f for fixed g is characterized by

$$\begin{aligned} 0 = \partial_f \mathcal{L}^\varepsilon(\Pi(f, g), f, g)(r) &= \int_{\Omega} r(y) d\nu(y) - \int_{\Omega \times \Omega} \exp \left(\frac{-c(x, y) + g(x) + f(y)}{\varepsilon} \right) r(y) d\xi(x, y) \\ &= \int_{\Omega} r(y) \left(\frac{d\nu}{d\mathcal{L}} \right) (y) d\mathcal{L}(y) - \int_{\Omega \times \Omega} \exp \left(\frac{-c(x, y) + g(x) + f(y)}{\varepsilon} \right) r(y) d\mathcal{N}(x) \otimes d\mathcal{L}(y) \end{aligned}$$

for all directions r . Hence, we obtain

$$0 = \left(\frac{d\nu}{d\mathcal{L}} \right) (y) - \int_{\Omega} \exp \left(\frac{-c(x, y) + g(x) + f(y)}{\varepsilon} \right) d\mathcal{N}(x),$$

and finally

$$\mathbf{g}^{C, \varepsilon}(y) := \varepsilon \log \left(\frac{d\nu}{d\mathcal{L}} \right) (y) - \varepsilon \log \left(\sum_{j=1, \dots, n} \exp \left(\frac{\mathbf{g}_j - |y - x_j|^2}{\varepsilon} \right) \right) \quad (5.4.5)$$

defines the optimal f for given $\mathbf{g} \in \mathbb{R}^n$ and all $y \in \Omega$. Given this C -transform, the regularized dual formulation in (5.4.4) can be rewritten as $\mathcal{W}_\varepsilon^2[\nu, \mu] = \max_{\mathbf{g} \in \mathbb{R}^n} \mathcal{D}^\varepsilon(\mathbf{g})$ with

$$\mathcal{D}^\varepsilon[\mathbf{g}] := \mathcal{L}^\varepsilon(\Pi(\mathbf{g}^{C,\varepsilon}, g), \mathbf{g}^{C,\varepsilon}, g) \quad (5.4.6)$$

$$= \int_{\Omega} \mathbf{g}^{C,\varepsilon}(y) d\nu(y) + \mathbf{g} \cdot \mathbf{m} - \varepsilon \int_{\Omega} \left(\frac{d\nu}{d\mathcal{L}} \right) (y) d\mathcal{L}(y) = \int_{\Omega} \mathbf{g}^{C,\varepsilon}(y) d\nu(y) + \mathbf{g} \cdot \mathbf{m} - \varepsilon \quad (5.4.7)$$

A necessary and sufficient condition for a vector $\mathbf{g} \in \mathbb{R}^n$ to maximize (5.4.6) for given sites \mathbf{X} and masses \mathbf{m} is

$$0 = \partial_{\mathbf{g}_i} \mathcal{D}^\varepsilon[\mathbf{g}] = \partial_{\mathbf{g}_i} \left(\int_{\Omega} \mathbf{g}^{C,\varepsilon}(y) d\nu(y) + \mathbf{g} \cdot \mathbf{m} \right) = \int_{\Omega} \partial_{\mathbf{g}_i} \mathbf{g}^{C,\varepsilon}(y) d\nu(y) + m_i \quad (5.4.8)$$

$$\text{with } \partial_{\mathbf{g}_i} \mathbf{g}^{C,\varepsilon}(y) = - \frac{\exp\left(\frac{\mathbf{g}_i - |y - x_i|^2}{\varepsilon}\right)}{\sum_{j=1, \dots, n} \exp\left(\frac{\mathbf{g}_j - |y - x_j|^2}{\varepsilon}\right)} \quad (5.4.9)$$

for $i = 1, \dots, n$. The set of functions $\{\chi_i^\varepsilon[\mathbf{X}, \mathbf{g}]\}_{i=1, \dots, n}$ with $\chi_i^\varepsilon[\mathbf{X}, \mathbf{g}](y) = -\partial_{\mathbf{g}_i} \mathbf{g}^{C,\varepsilon}(y)$ for $y \in \Omega$ forms a partition of unity on Ω . Furthermore,

$$\lim_{\varepsilon \rightarrow 0} \chi_i^\varepsilon[\mathbf{X}, \mathbf{g}] = \chi_{L_i}[\mathbf{X}, \mathbf{g}] \quad (5.4.10)$$

for $i = 1, \dots, n$, where the convergence is in $L^1(\nu)$. For given $[\mathbf{X}, \mathbf{g}]$ with $\varepsilon > 0$, $\mathbf{X} = (x_1, \dots, x_n) \in \Omega^n$, $x_j \neq x_i$ for $j, i = 1, \dots, n$

$$m_i^\varepsilon[\mathbf{X}, \mathbf{g}] := \int_{\Omega} \chi_i^\varepsilon[\mathbf{X}, \mathbf{g}](y) d\nu(y), \quad b_i^\varepsilon[\mathbf{X}, \mathbf{g}] := \int_{\Omega} \frac{y \chi_i^\varepsilon[\mathbf{X}, \mathbf{g}](y) d\nu(y)}{m_i^\varepsilon[\mathbf{X}, \mathbf{g}]} \quad (5.4.11)$$

define regularized masses and regularized barycenters, respectively. By the L^1 -convergence of $\chi_i^\varepsilon[\mathbf{X}, \mathbf{g}]$ to χ_{L_i} and by the compactness of Ω , one gets

$$\lim_{\varepsilon \rightarrow 0} m_i^\varepsilon[\mathbf{X}, \mathbf{g}] = m_i[\mathbf{X}, \mathbf{g}], \quad \lim_{\varepsilon \rightarrow 0} b_i^\varepsilon[\mathbf{X}, \mathbf{g}] = b_i[\mathbf{X}, \mathbf{g}] \text{ for } m_i[\mathbf{X}, \mathbf{g}] > 0. \quad (5.4.12)$$

Consequently, we obtain the entropy-regularized cost functional

$$\mathcal{F}_n^\varepsilon[\mathbf{X}, \mathbf{g}] := \sum_{i=1, \dots, n} m_i^\varepsilon[\mathbf{X}, \mathbf{g}] \Phi(b_i^\varepsilon[\mathbf{X}, \mathbf{g}]) \quad (5.4.13)$$

as an approximation of the original cost functional \mathcal{F}_n on power diagrams defined in (5.2.2). Analogously, one finally obtains the entropy-regularized cost functional with penalty parameter $\eta > 0$:

$$\mathcal{F}_n^{\eta, \varepsilon}[\mathbf{X}, \mathbf{g}] := \mathcal{F}_n^\varepsilon[\mathbf{X}, \mathbf{g}] - \eta \mathcal{R}_n^\varepsilon[\mathbf{X}, \mathbf{g}], \quad (5.4.14)$$

where the regularized penalty term is defined as

$$\mathcal{R}_n^\varepsilon[\mathbf{X}, \mathbf{g}] := \sum_{i=1, \dots, n} \int_{\Omega} |y - x_i|^2 \chi_i^\varepsilon[\mathbf{X}, \mathbf{g}](y) d\nu(y) + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^\varepsilon[\mathbf{X}, \mathbf{g}] m_j^\varepsilon[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2}. \quad (5.4.15)$$

in analogy to (5.2.6).

We are now in the position to prove the convergence of maximizers of the entropy regularized functional $\mathcal{F}_n^{\eta, \varepsilon}$ given in (5.4.14) to a maximizer of the original functional defined in (5.2.5) for $\varepsilon \rightarrow 0$.

Proposition 5.4.1. *Let $(\varepsilon_N)_N \subset \mathbb{R}^+$, $(\mathbf{X}^N, \mathbf{g}^N)_N \subset \mathcal{X}_n \times \mathbb{R}^n$ so that $\varepsilon_N \rightarrow 0$ and $(\mathbf{X}^N, \mathbf{g}^N) \rightarrow (\mathbf{X}, \mathbf{g}) \in \mathcal{X}_n \times \mathbb{R}^n$ for $N \rightarrow \infty$. Then, we have that*

$$\lim_{N \rightarrow \infty} \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N] = \chi_{L_i}[\mathbf{X}, \mathbf{g}],$$

where the limit is taken in $L^1(\nu)$.

Proof. Let $y \in D$ be in the interior of $L_i(\mathbf{X}, \mathbf{g})$. Then, there is $\delta > 0$, such that $|y - x_i|^2 - \mathbf{g}_i \leq |y - x_j|^2 - \mathbf{g}_j - 3\delta$, for all $j \neq i$. and by the continuity of the function $(\mathbf{X}, \mathbf{g}) \mapsto |y - x_i|^2 - \mathbf{g}_i$ we obtain that $|y - x_i^N|^2 - \mathbf{g}_i^N \leq |y - x_j^N|^2 - \mathbf{g}_j^N - \delta$ for sufficiently large N . This implies that

$$\begin{aligned} \chi_i^{\varepsilon}[\mathbf{X}^N, \mathbf{g}^N](y) &= \frac{\exp\left(\frac{\mathbf{g}_i^N - |y - x_i^N|^2}{\varepsilon^N}\right)}{\sum_{j=1}^n \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon^N}\right)} \\ &\geq \frac{\exp\left(\frac{\mathbf{g}_i^N - |y - x_i^N|^2}{\varepsilon^N}\right)}{\exp\left(\frac{\mathbf{g}_i^N - |y - x_i^N|^2}{\varepsilon^N}\right) + (n-1) \exp\left(\frac{\mathbf{g}_i^N - |y - x_i^N|^2}{\varepsilon^N}\right) \exp(-\delta/\varepsilon^N)} \\ &= \frac{1}{1 + (n-1) \exp(-\delta/\varepsilon^N)} \rightarrow 1 \end{aligned}$$

for $N \rightarrow \infty$. Now, let $y \notin \overline{L_i(\mathbf{X}, \mathbf{g})}$. Then, there is $j \neq i$, and $\delta > 0$, such that $|y - x_i|^2 - \mathbf{g}_i \geq |y - x_j|^2 - \mathbf{g}_j + 3\delta$. and $|y - x_i^N|^2 - \mathbf{g}_i^N \geq |y - x_j^N|^2 - \mathbf{g}_j^N + \delta$ for sufficiently large N . From this it follows that

$$\begin{aligned} \chi_i^{\varepsilon}[\mathbf{X}^N, \mathbf{g}^N](y) &= \frac{\exp\left(\frac{\mathbf{g}_i^N - |y - x_i^N|^2}{\varepsilon^N}\right)}{\sum_{j=1}^n \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon^N}\right)} \leq \frac{\exp\left(\frac{\mathbf{g}_i^N - |y - x_i^N|^2}{\varepsilon^N}\right)}{\exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon^N}\right)} \\ &= \exp\left(\frac{\mathbf{g}_i^N - \mathbf{g}_j^N - |y - x_i^N|^2 + |y - x_j^N|^2}{\varepsilon^N}\right) \\ &\leq \exp\left(-\frac{\delta}{\varepsilon^N}\right) \rightarrow 0, \end{aligned}$$

as $N \rightarrow \infty$. By the absolute continuity of ν with respect to the Lebesgue measure, we have proven that $\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N] \rightarrow \chi_{L_i}[\mathbf{X}, \mathbf{g}]$, ν -almost everywhere, and by the dominated convergence theorem, we obtain the claim. \square

In preparation for the convergence of maximizers, we now state the entropy-regularized "pendant" of Theorem 5.2.5.

Theorem 5.4.2. *Let us assume that $\varepsilon > 0$, and $\Phi \in C(\Omega)$ with maximum $\bar{\Phi}$ on D . Then, for given number of sites $l \in \mathbb{N}$ there exists an $n \leq l$, such that a maximizer $(\mathbf{X}^*, \mathbf{g}^*)$ of $\mathcal{F}_n^{\eta, \varepsilon}$ exists with $m_i^{\varepsilon}[\mathbf{X}^*, \mathbf{g}^*] > 0$ for $i = 1, \dots, n$ and $\mathcal{F}_n^{\eta, \varepsilon}(\mathbf{X}^*, \mathbf{g}^*) \geq \sup_{(\mathbf{X}, \mathbf{g})} \mathcal{F}_l^{\eta, \varepsilon}(\mathbf{X}, \mathbf{g})$.*

Proof. The proof is analogous to Theorem 5.2.5: Consider a maximizing sequence for $\mathcal{F}_l^{\eta,\varepsilon}$. Since $m_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N] \in [0, 1]$, and $b_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N] \in \Omega$, we may assume that up to the selection of a subsequence the sequences of masses and barycenters converge to limits m_i^* and b_i^* , respectively, and that there exists an $n \leq l$ with $m_i^* > 0$ if and only if $i \leq n$. Then, we obtain

$$\begin{aligned}
\sup_{(\mathbf{X}, \mathbf{g})} \mathcal{F}_l^{\eta,\varepsilon}(\mathbf{X}, \mathbf{g}) &= \lim_{N \rightarrow \infty} \mathcal{F}_l^{\eta,\varepsilon}(\mathbf{X}^N, \mathbf{g}^N) \leq \limsup_{N \rightarrow \infty} \mathcal{F}_l^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_l^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) \\
&\leq \limsup_{N \rightarrow \infty} \mathcal{F}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) + \limsup_{N \rightarrow \infty} \sum_{i=n+1, \dots, l} m_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N] \Phi(b_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N]) \\
&\quad - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} (\mathcal{R}_l^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) - \mathcal{R}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N)) \\
&= \lim_{N \rightarrow \infty} \mathcal{F}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) + 0 - \eta \lim_{N \rightarrow \infty} \mathcal{R}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} (\mathcal{R}_l^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) \\
&\quad - \mathcal{R}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N)) \\
&\leq \lim_{N \rightarrow \infty} \mathcal{F}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) - \eta \lim_{N \rightarrow \infty} \mathcal{R}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) = \lim_{N \rightarrow \infty} \mathcal{F}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) - \eta \mathcal{R}_n^\varepsilon(\mathbf{X}^N, \mathbf{g}^N) \\
&= \lim_{N \rightarrow \infty} \mathcal{F}_n^{\eta,\varepsilon}(\mathbf{X}^N, \mathbf{g}^N).
\end{aligned}$$

Hence, a maximizer of $\mathcal{F}_n^{\eta,\varepsilon}$, if it exists, attains a greater or equal objective value than a maximizer of $\mathcal{F}_l^{\eta,\varepsilon}$. Now, let $(\mathbf{X}^N, \mathbf{g}^N)_N \subset \mathcal{X}_n \times \mathbb{R}^n$ be a maximizing sequence for $\mathcal{F}_n^{\eta,\varepsilon}$ and in analogy to before assume that $m_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N]$ and $b_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N]$ have limits $m_i^* > 0$, and $b_i^* \in D$ for $i = 1, \dots, n$, that $\mathcal{F}_n^{\eta,\varepsilon}(\mathbf{X}^N, \mathbf{g}^N)$ is monotonically increasing in N and that $m_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N] \geq \frac{1}{2}m_i^*$. Consequently, for $\bar{\Phi}$ being the maximal value of Φ on D the estimate

$$\begin{aligned}
\mathcal{F}_n^{\eta,\varepsilon}(\mathbf{X}^0, \mathbf{g}^0) &\leq \nu(D)\bar{\Phi} - \eta \left(\sum_{i=1, \dots, n} \int_{\Omega} |y - x_i^N|^2 \chi_{L_i[\mathbf{X}^N, \mathbf{g}^N]}^\varepsilon d\nu(y) \right. \\
&\quad \left. + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^\varepsilon[\mathbf{X}^N, \mathbf{g}^N] m_j^\varepsilon[\mathbf{X}^N, \mathbf{g}^N]}{|x_i^N - x_j^N|^2} \right) \\
&\leq \bar{\Phi} - \eta \left(\frac{1}{2} \sum_{i=1, \dots, n} \text{dist}^2(x_i^N, \Omega) m_i^* + \frac{1}{4} \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^* m_j^*}{|x_i^N - x_j^N|^2} \right)
\end{aligned} \tag{5.4.16}$$

is obtained. This implies the following a priori bounds:

$$\text{dist}(x_i^N, \Omega) \leq \sqrt{\frac{2(\bar{\Phi} - \mathcal{F}_n^{\eta,\varepsilon}(\mathbf{X}^0, \mathbf{g}^0))}{\eta m_i^*}}, \quad |x_i^N - x_j^N| \geq \sqrt{\frac{\eta m_i^* m_j^*}{4(\bar{\Phi} - \mathcal{F}_n^{\eta,\varepsilon}(\mathbf{X}^0, \mathbf{g}^0))}}$$

for all N . For the uniform bound on \mathbf{g} , recall that \mathbf{g} and $\mathbf{g} + \lambda \mathbb{1}_n$ both result on the same Laguerre diagram. Hence, we may assume without loss of generality that $\mathbf{g}_1^N = 0$ for all $N \in \mathbb{N}$. Then, $\lim_{N \rightarrow \infty} \mathbf{g}_j^N = \infty$ for $j = 2, \dots, n$ would imply $\lim_{N \rightarrow \infty} m_1^\varepsilon[\mathbf{X}^N, \mathbf{g}^N] = 0$, which contradicts our choice of n . Similarly, $\lim_{N \rightarrow \infty} \mathbf{g}_j = -\infty$ would imply $\lim_{N \rightarrow \infty} m_j^\varepsilon[\mathbf{X}^N, \mathbf{g}^N] = 0$, which again is a contradiction. Hence, $|\mathbf{g}^N| \leq C$ for some $C > 0$ and all $N \in \mathbb{N}$. Finally, given these a priori bounds the existence of a maximizer of $\mathcal{F}_n^{\eta,\varepsilon}$ follows directly from the Weierstraß extreme value theorem. \square

We are now in the position to prove the convergence of maximizers of the entropy regularized functional $\mathcal{F}_n^{\eta, \epsilon}$ given in (5.4.14) to a maximizer of the original functional \mathcal{F}_n^η defined in (5.2.5) for $\epsilon \rightarrow 0$.

Theorem 5.4.3. *Let $\Phi \in C(\Omega)$, $l \in \mathbb{N}$ and $\eta > 0$. For each $N \in \mathbb{N}$, consider maximizers $(\mathbf{X}^N, \mathbf{g}^N) \in \mathcal{X}_{\beta(N)} \times \mathbb{R}^{\beta(N)}$, where $\beta : \mathbb{N} \rightarrow \{1, \dots, n\}$, so that $\mathcal{F}_{\beta(N)}^{\eta, \epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) \geq \max_{k=1, \dots, n} \{\sup_{(\mathbf{X}, \mathbf{g}) \in \mathcal{X}_k \times \mathbb{R}^k} \mathcal{F}_k^{\eta, \epsilon_N}(\mathbf{X}, \mathbf{g})\}$. Then, there is $n \in \{1, \dots, l\}$ such that up to the selection of a subsequence $(\mathbf{X}^N, \mathbf{g}^N)_{N \in \mathbb{N}} \subset \mathcal{X}_n \times \mathbb{R}^n$ converges to a limit $(\mathbf{X}^*, \mathbf{g}^*) \in \mathcal{X}_n \times \mathbb{R}^n$ for which it holds*

$$\mathcal{F}_n^\eta(\mathbf{X}^*, \mathbf{g}^*) \geq \max_{k=1, \dots, l} \left\{ \sup_{(\mathbf{X}, \mathbf{g}) \in \mathcal{X}_k \times \mathbb{R}^k} \mathcal{F}_k^\eta(\mathbf{X}, \mathbf{g}) \right\}.$$

Proof. The proof of the first statement follows the proof of Theorem 5.2.5 very closely, albeit with a few key differences, which will be highlighted now: Let $n' \in \{1, \dots, l\}$, such that $|\{N : \beta(N) = n'\}| = \infty$. Choose a subsequence such that up to relabeling of indices $(\mathbf{X}^N, \mathbf{g}^N)_{N \in \mathbb{N}} \subset \mathcal{X}_{n'} \times \mathbb{R}^{n'}$, such that the sequence of masses $(m_1^{\epsilon_N}[\mathbf{X}^N, \mathbf{g}^N], \dots, m_{n'}^{\epsilon_N}[\mathbf{X}^N, \mathbf{g}^N])_{N \in \mathbb{N}}$ converges to some $(m_1^*, \dots, m_{n'}^*) \in [0, 1]^{n'}$. Moreover, assume without loss of generality that there is $n \in \{1, \dots, n'\}$ with $m_i^* > 0$ if and only if $i \leq n$, and $m_i^{\epsilon_N}[\mathbf{X}^N, \mathbf{g}^N] \geq \frac{1}{2}m_i^*$ for all $N \in \mathbb{N}$, $i \leq n$. It then holds for all $k \in \{1, \dots, l\}$ and all $(\bar{\mathbf{X}}, \bar{\mathbf{g}}) \in \mathcal{X}_k \times \mathbb{R}^k$ with $m_i[\bar{\mathbf{X}}, \bar{\mathbf{g}}] > 0$ for $i = 1, \dots, k$:

$$\begin{aligned} \mathcal{F}_k^\eta[\bar{\mathbf{X}}, \bar{\mathbf{g}}] &= \lim_{N \rightarrow \infty} \mathcal{F}_k^{\eta, \epsilon_N}[\bar{\mathbf{X}}, \bar{\mathbf{g}}] \leq \limsup_{N \rightarrow \infty} \mathcal{F}_{n'}^{\eta, \epsilon_N}[\mathbf{X}^N, \mathbf{g}^N] \\ &\leq \limsup_{N \rightarrow \infty} \mathcal{F}_{n'}^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_{n'}^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) \\ &\leq \limsup_{N \rightarrow \infty} \mathcal{F}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) + \limsup_{N \rightarrow \infty} \sum_{i=n+1, \dots, n'} m_i^{\epsilon_N}[\mathbf{X}^N, \mathbf{g}^N] \Phi(b_i^{\epsilon_N}[\mathbf{X}^N, \mathbf{g}^N]) \\ &\quad - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} (\mathcal{R}_{n'}^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) - \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N)) \\ &= \limsup_{N \rightarrow \infty} \mathcal{F}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) + 0 - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} (\mathcal{R}_{n'}^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) \\ &\quad - \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N)) \\ &\leq \limsup_{N \rightarrow \infty} \mathcal{F}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) \leq \bar{\Phi} - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N). \end{aligned}$$

We obtain $\liminf_{N \rightarrow \infty} \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) \leq \frac{\bar{\Phi} - \mathcal{F}_k^\eta[\bar{\mathbf{X}}, \bar{\mathbf{g}}]}{\eta}$. Without loss of generality, we can assume that for all $\delta > 0$, $\mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) < \frac{\bar{\Phi} - \mathcal{F}_k^\eta[\bar{\mathbf{X}}, \bar{\mathbf{g}}] + \delta}{\eta}$ holds for all $N \in \mathbb{N}$, $i = 1, \dots, l$. Then, we obtain the bounds

$$\text{dist}(x_i^N, \Omega) < \sqrt{\frac{2(\bar{\Phi} - \mathcal{F}_k^\eta[\bar{\mathbf{X}}, \bar{\mathbf{g}}] + \delta)}{\eta m_i^*}}, \quad |x_i^N - x_j^N| > \sqrt{\frac{\eta m_i^* m_j^*}{4(\bar{\Phi} - \mathcal{F}_k^\eta[\bar{\mathbf{X}}, \bar{\mathbf{g}}] + \delta)}},$$

for all pairwise different $i, j = 1, \dots, n$. Using the same arguments as in the proof of Theorem 5.2.5, one also finds analogous uniform bounds on $|\mathbf{g}^N|$. Hence, up to a subsequence, $(\mathbf{X}^N, \mathbf{g}^N)_{N \in \mathbb{N}}$ converges to some $(\mathbf{X}^*, \mathbf{g}^*) \in \mathcal{X}_n \times \mathbb{R}^n$. By the above computations, we finally obtain

$$\mathcal{F}_k^\eta[\bar{\mathbf{X}}, \bar{\mathbf{g}}] \leq \lim_{N \rightarrow \infty} \mathcal{F}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n^{\epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) = \lim_{N \rightarrow \infty} \mathcal{F}_n^{\eta, \epsilon_N}(\mathbf{X}^N, \mathbf{g}^N) = \mathcal{F}_n^\eta(\mathbf{X}^*, \mathbf{g}^*).$$

For the last equality, we used the fact that the masses $m_i^{\varepsilon_N}$ do not vanish in the limit and Proposition 5.4.1 for the convergence of masses $m_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N] \rightarrow m_i[\mathbf{X}^*, \mathbf{g}^*]$. For the convergence of barycenters $b_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N] \rightarrow b_i[\mathbf{X}^*, \mathbf{g}^*]$, we additionally use the fact that the function $y \mapsto y$ is bounded in Ω . Finally, for the convergence of

$$\int_{\Omega} |y - x_i^N|^2 \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y) \, d\nu(y) \rightarrow \int_{\Omega} |y - x_i^*|^2 \chi_i[\mathbf{X}^*, \mathbf{g}^*](y) \, d\nu(y),$$

we also use the fact that the family of functions $(y \mapsto |y - x_i^N|^2)_{N \in \mathbb{N}}$ is uniformly bounded in $L^\infty(\nu)$ and converges in $L^\infty(\nu)$ to $y \mapsto |y - x_i^*|^2$, and the uniform boundedness principle. \square

To numerically implement the maximization of 5.4.13 or 5.4.14 via a gradient ascent approach we need to compute the gradient of $\mathcal{F}_n^\varepsilon$. Using $z = x_k$ or $z = \mathbf{g}_k$ for $k = 1, \dots, n$ we obtain

$$\begin{aligned} \partial_z \mathcal{F}_n^\varepsilon[\mathbf{X}, \mathbf{g}] &= \sum_{j=1, \dots, n} \partial_z m_j^\varepsilon[\mathbf{X}, \mathbf{g}] \Phi(b_j^\varepsilon[\mathbf{X}, \mathbf{g}]) + m_j^\varepsilon[\mathbf{X}, \mathbf{g}] \partial_z \Phi(b_j^\varepsilon[\mathbf{X}, \mathbf{g}]), \\ \partial_{x_k} \mathcal{R}_n^\varepsilon[\mathbf{X}, \mathbf{g}] &= \int_{\Omega} 2(x_k - y) \chi_k^\varepsilon[\mathbf{X}, \mathbf{g}](y) \, d\nu(y) + \sum_{i=1, \dots, n} \int_{\Omega} |y - x_i|^2 \partial_{x_k} \chi_i^\varepsilon[\mathbf{X}, \mathbf{g}](y) \, d\nu(y) \\ &\quad + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \left(\frac{\partial_{x_k} m_i^\varepsilon[\mathbf{X}, \mathbf{g}] m_j^\varepsilon[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} + \frac{m_i^\varepsilon[\mathbf{X}, \mathbf{g}] \partial_{x_k} m_j^\varepsilon[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} \right. \\ &\quad \left. - 2 \frac{m_i^\varepsilon[\mathbf{X}, \mathbf{g}] m_j^\varepsilon[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^4} ((x_i - x_j) \delta_{ik} + (x_j - x_i) \delta_{jk}) \right), \\ \partial_{\mathbf{g}_k} \mathcal{R}_n^\varepsilon[\mathbf{X}, \mathbf{g}] &= \sum_{i=1, \dots, n} \int_{\Omega} |y - x_i|^2 \partial_{\mathbf{g}_k} \chi_i^\varepsilon[\mathbf{X}, \mathbf{g}](y) \, d\nu(y) \\ &\quad + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \left(\frac{\partial_{\mathbf{g}_k} m_i^\varepsilon[\mathbf{X}, \mathbf{g}] m_j^\varepsilon[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} + \frac{m_i^\varepsilon[\mathbf{X}, \mathbf{g}] \partial_{\mathbf{g}_k} m_j^\varepsilon[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} \right), \\ \partial_z \Phi(b_j^\varepsilon[\mathbf{X}, \mathbf{g}]) &= \nabla \Phi(b_j^\varepsilon[\mathbf{X}, \mathbf{g}]) \cdot \partial_z b_j^\varepsilon[\mathbf{X}, \mathbf{g}], \\ \partial_z m_j^\varepsilon[\mathbf{X}, \mathbf{g}] &= \int_{\Omega} \partial_z \chi_j^\varepsilon[\mathbf{X}, \mathbf{g}](y) \, d\nu(y), \\ \partial_z b_j^\varepsilon[\mathbf{X}, \mathbf{g}] &= \frac{1}{m_j^\varepsilon[\mathbf{X}, \mathbf{g}]} \int_{\Omega} y \partial_z \chi_j^\varepsilon[\mathbf{X}, \mathbf{g}](y) \, d\nu(y) - \frac{b_j^\varepsilon[\mathbf{X}, \mathbf{g}] \partial_z m_j^\varepsilon[\mathbf{X}, \mathbf{g}]}{m_j^\varepsilon[\mathbf{X}, \mathbf{g}]^2}, \\ \partial_{x_k} \chi_j^\varepsilon[\mathbf{X}, \mathbf{g}](y) &= \frac{2x_j}{\varepsilon} (\chi_j^\varepsilon[\mathbf{X}, \mathbf{g}](y) - \delta_{kj}) \chi_k^\varepsilon[\mathbf{X}, \mathbf{g}](y), \\ \partial_{\mathbf{g}_k} \chi_j^\varepsilon[\mathbf{X}, \mathbf{g}](y) &= -\frac{1}{\varepsilon} (\chi_j^\varepsilon[\mathbf{X}, \mathbf{g}](y) - \delta_{kj}) \chi_k^\varepsilon[\mathbf{X}, \mathbf{g}](y). \end{aligned}$$

Hence, since the gradient of \mathcal{F}^ε has the form of an integral with respect to ν , a stochastic gradient ascent is a very suitable method to numerically obtain an optimizer of this problem: since evaluating the integral of gradients becomes very expensive, one can save computational costs at every iteration by sampling the measure ν , and by approximating the integral by the sum of the integrand evaluated at the samples' values, cf. [BB07]. This method has the advantage of not having to discretize the measure ν in space, which might lead to further inaccuracies. Alternatively, to implement the maximization of 5.4.13 numerically, we have to further discretize the continuous given measure ν in space.

5.5 Spatial discretization

As the next step, we have to discretize our optimization problem in space. To this end, we restrict to the domain $D = [0, 1]^d$ and consider a dyadic mesh with grid size $h = 2^{-N}$ for $N \in \mathbb{N}$. The domain is subdivided into cells $D_h^\alpha = \times_{1,\dots,d} [\alpha_j h, (\alpha_j + 1)h]$ where α is a multi-index in $(\mathcal{I}^h)^d$ for the index set $\mathcal{I}^h = \{0, \dots, 2^N - 1\}$. We consider a discrete measure $\nu^h = \sum_{\alpha \in \mathcal{I}^d} \nu_\alpha \delta_{y_h^\alpha}$ with $\nu_\alpha = \nu(D_h^\alpha)$, where $y_h^\alpha = ((\alpha_j + \frac{1}{2})h)_{j=1,\dots,d}$ are the cell centers. This discretization ansatz gives rise to discrete counterparts of the continuous, regularized characteristic functions $\chi_i^\varepsilon[\mathbf{X}, \mathbf{g}]$, masses $m_i^\varepsilon[\mathbf{X}, \mathbf{g}]$, and barycenters $b_i^\varepsilon[\mathbf{X}, \mathbf{g}]$:

$$\chi_{i,\alpha}^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] := \chi_i^\varepsilon[\mathbf{X}, \mathbf{g}](y_h^\alpha) = \frac{\exp\left(\frac{\mathbf{g}_i - |y_h^\alpha - x_i|^2}{\varepsilon}\right)}{\sum_{j=1,\dots,n} \exp\left(\frac{\mathbf{g}_j - |y_h^\alpha - x_j|^2}{\varepsilon}\right)}, \quad (5.5.1)$$

$$m_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] := \sum_{\alpha \in \mathcal{I}^d} \chi_{i,\alpha}^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] \nu^\alpha, \quad b_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] := \sum_{\alpha \in \mathcal{I}^d} \frac{y_h^\alpha \chi_{i,\alpha}^{\varepsilon,h}[\mathbf{X}, \mathbf{g}]}{m_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}]} \nu^\alpha. \quad (5.5.2)$$

Based on this discretization, we obtain a discrete functional

$$\mathcal{F}_n^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] := \sum_{i=1,\dots,n} m_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] \Phi(b_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}]), \quad (5.5.3)$$

and the usual penalty-enhanced discrete functional

$$\mathcal{F}_n^{\eta,\varepsilon,h}[\mathbf{X}, \mathbf{g}] := \sum_{i=1,\dots,n} m_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] \Phi(b_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}]) - \eta \mathcal{R}_n^{\varepsilon,h}[\mathbf{X}, \mathbf{g}], \quad (5.5.4)$$

where the discrete penalty term is defined by

$$\mathcal{R}_n^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] := \sum_{i=1,\dots,n} \sum_{\alpha \in \mathcal{I}^d} |y_h^\alpha - x_i|^2 \chi_{i,\alpha}^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] \nu^\alpha + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] m_j^{\varepsilon,h}[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2}. \quad (5.5.5)$$

Now, in analogy to Proposition 5.4.1 we demonstrate the consistency of entropy regularized and discrete characteristic functions and the continuous counterpart. To this end, let us define the ν -almost surely piecewise continuous function $\chi_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}] \in L^1(\nu)$ as follows:

$$\chi_i^{\varepsilon,h}[\mathbf{X}, \mathbf{g}](y) = \chi_{i,\alpha}^{\varepsilon,h} \text{ for } \nu - \text{almost all } y \text{ in a cell interior } \overset{\circ}{\Omega}_h^\alpha.$$

Proposition 5.5.1. *For $(\varepsilon_N)_N \subset \mathbb{R}^+$, $(\mathbf{X}^N, \mathbf{g}^N)_N \subset \mathcal{X}_n \times \mathbb{R}^n$ with $\varepsilon_N \rightarrow 0$, $(\mathbf{X}^N, \mathbf{g}^N) \rightarrow (\mathbf{X}, \mathbf{g}) \in \mathcal{X}_n \times \mathbb{R}^n$ for $N \rightarrow \infty$ and for grid sizes $h_N = 2^{-N}$ we obtain that $\lim_{N \rightarrow \infty} \chi_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] = \chi_{L_i}[\mathbf{X}, \mathbf{g}]$, in $L^1(\nu)$.*

Proof. By the triangle inequality, we have that

$$\left\| \chi_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] - \chi_{L_i}[\mathbf{X}, \mathbf{g}] \right\|_{L^1(\nu)} \leq \left\| \chi_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] - \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N] \right\|_{L^1(\nu)}$$

$$+ \|\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N] - \chi_{L_i}[\mathbf{X}, \mathbf{g}]\|_{L^1(\nu)}.$$

In Proposition 5.4.1, we already showed that the second term converges to zero. To verify that the first term vanishes as well in the limit for $N \rightarrow \infty$, we estimate the norm of the gradient of $\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N]$ given by

$$\nabla_y \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y) = \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y) \left(-2 \frac{y - x_i^N}{\varepsilon_N} + \frac{\sum_{j=1}^n 2(y - x_j^N) \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon_N}\right)}{\varepsilon_N \sum_{j=1}^n \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon_N}\right)} \right).$$

Let $y \in D$ be in the interior of $L_i(\mathbf{X}, \mathbf{g})$. Then, as in the proof of Proposition 5.4.1, there is a $\delta > 0$, such that

$$|y - x_i^N|^2 - \mathbf{g}_i^N \leq |y - x_j^N|^2 - \mathbf{g}_j^N - \delta$$

for for $j \neq i$ and for sufficiently large N . Then, one obtains

$$\begin{aligned} & |\nabla_y \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y)| \\ &= \frac{2\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y)}{\varepsilon_N} \left| (y - x_i^N)(\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y) - 1) + \frac{\sum_{j \neq i} (y - x_j^N) \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon_N}\right)}{\sum_{j=1}^n \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon_N}\right)} \right| \\ &\leq \frac{2}{\varepsilon_N} \left(|y - x_i^N|(1 - \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y)) + \sum_{j \neq i} |y - x_j^N| \exp\left(\frac{-\delta}{\varepsilon_N}\right) \right) \\ &\leq \frac{2}{\varepsilon_N} \left(|y - x_i^N| \left(\frac{(n-1) \exp(-\delta/\varepsilon_N)}{1 + (n-1) \exp(-\delta/\varepsilon_N)} \right) + \sum_{j \neq i} |y - x_j^N| \exp\left(\frac{-\delta}{\varepsilon_N}\right) \right) \\ &\leq \frac{2}{\varepsilon_N} \exp(-\delta/\varepsilon_N) \left(|y - x_i^N|(n-1) + \sum_{j \neq i} |y - x_j^N| \right) \rightarrow 0, \end{aligned}$$

for $N \rightarrow \infty$. Next, let $y \notin \overline{L_i(\mathbf{X}, \mathbf{g})}$. Then, it holds

$$\begin{aligned} |\nabla_y \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y)| &= \frac{2\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y)}{\varepsilon_N} \left| -(y - x_i^N) + \frac{\sum_{j \neq i} (y - x_j^N) \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon_N}\right)}{\sum_{j=1}^n \exp\left(\frac{\mathbf{g}_j^N - |y - x_j^N|^2}{\varepsilon_N}\right)} \right| \\ &= \frac{2\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y)}{\varepsilon_N} |-(y - x_i^N) + \tilde{x}| \leq \frac{2}{\varepsilon_N} \exp(-\delta/\varepsilon_N) |x_i^N - y + \tilde{x}| \rightarrow 0, \end{aligned}$$

where $\tilde{x} = \sum_{j \neq i} (y - x_j^N) \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y)$ is an element in the convex hull of $(y - x_j^N)_{j=1, \dots, n}$. and thus due to the convergence of $(\mathbf{X}^N)_N$ uniformly bounded. This implies that $|x_i^N - y + \tilde{x}|$ is uniformly bounded.

For every $y \in \Omega$ in the interior of $L_i(\mathbf{X}, \mathbf{g})$ and sufficiently large N we observe that y is in the cell $\Omega_{h_N}^{\alpha^N}$ for some multi-indices α^N and $\Omega_{h_{N_0}}^{\alpha^{N_0}}$ is completely contained in the interior of $L_i(\mathbf{X}, \mathbf{g})$. There is

a $\delta > 0$ such that $|y - x_i^N|^2 - g_i^N \leq |y - x_j^N|^2 - g_j^N - \delta$ holds for $j \neq i$ and all $y \in \Omega_{h_N}^{\alpha^N}$. Hence, the restriction of $\chi_i^{\varepsilon_M}[\mathbf{X}^M, \mathbf{g}^M]$ onto this grid cell for $M \geq N$ is a family of uniformly Lipschitz functions, and in particular equicontinuous. It follows that

$$\begin{aligned} \left| \chi_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N](y) - \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y) \right| &= \left| \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y_{h_N}^{\alpha^N}) - \chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N](y) \right| \\ &\leq \left(\sup_{N \geq N_0} \text{Lip}(\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N]) \right) |y - y_{h_N}^{\alpha^N}| \rightarrow 0 \end{aligned}$$

The analogous statement holds for $y \notin \overline{L_i(\mathbf{X}, \mathbf{g})}$ as well. Thus, $\chi_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N]$ converges to $\chi_i^{\varepsilon_N}[\mathbf{X}^N, \mathbf{g}^N]$ ν -almost everywhere and, by the dominated convergence theorem, one obtains the desired convergence in $L^1(\nu)$. \square

We obtain the entropy-regularized, fully discrete analogue of Theorem 5.2.5.

Theorem 5.5.2. *Let us assume that $\varepsilon, h > 0$, and $\Phi \in C(\Omega)$ with maximum $\bar{\Phi}$ on D . Then, for given number of sites $l \in \mathbb{N}$ there exists an $n \leq l$, such that a maximizer $(\mathbf{X}^*, \mathbf{g}^*)$ of $\mathcal{F}_n^{\eta, \varepsilon, h}$ exists with $m_i^{\varepsilon, h}[\mathbf{X}^*, \mathbf{g}^*] > 0$ for $i = 1, \dots, n$ and $\mathcal{F}_n^{\eta, \varepsilon, h}(\mathbf{X}^*, \mathbf{g}^*) \geq \sup_{(\mathbf{X}, \mathbf{g})} \mathcal{F}_l^{\eta, \varepsilon, h}(\mathbf{X}, \mathbf{g})$.*

Proof. The proof is completely analogous to the proof of Theorem 5.2.5: Consider a maximizing sequence $(\mathbf{X}^N, \mathbf{g}^N)_N$ for $\mathcal{F}_l^{\eta, \varepsilon, h}$. Since $m_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] \in [0, 1]$, and $b_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] \in \Omega$, we may assume that up to the selection of a subsequence the sequences of masses and barycenters converge to limits m_i^* and b_i^* , respectively, and that there exists an $n \leq l$ with $m_i^* > 0$ if and only if $i \leq n$. Then, we obtain

$$\begin{aligned} \sup_{(\mathbf{X}, \mathbf{g})} \mathcal{F}_l^{\eta, \varepsilon, h}(\mathbf{X}, \mathbf{g}) &= \lim_{N \rightarrow \infty} \mathcal{F}_l^{\eta, \varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) \leq \limsup_{N \rightarrow \infty} \mathcal{F}_l^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_l^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) \\ &\leq \limsup_{N \rightarrow \infty} \mathcal{F}_n^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) + \limsup_{N \rightarrow \infty} \sum_{i=n+1, \dots, l} m_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] \Phi(b_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N]) \\ &\quad - \eta \liminf_{N \rightarrow \infty} \mathcal{R}_n^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) - \eta \liminf_{N \rightarrow \infty} \left(\mathcal{R}_l^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) - \mathcal{R}_n^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) \right) \\ &\leq \lim_{N \rightarrow \infty} \mathcal{F}_n^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) - \eta \lim_{N \rightarrow \infty} \mathcal{R}_n^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) \\ &= \lim_{N \rightarrow \infty} \mathcal{F}_n^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) - \eta \mathcal{R}_n^{\varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N) \\ &= \lim_{N \rightarrow \infty} \mathcal{F}_n^{\eta, \varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N), \end{aligned}$$

where for the sake of clarity, we denote by $(\mathbf{X}^N, \mathbf{g}^N)$ both itself and its truncation to the first n components. Hence, a maximizer of $\mathcal{F}_n^{\eta, \varepsilon, h}$, if it exists, attains a greater or equal objective value than a maximizer of $\mathcal{F}_l^{\eta, \varepsilon, h}$. Now, let $(\mathbf{X}^N, \mathbf{g}^N)_N \subset \mathcal{X}_n \times \mathbb{R}^n$ be a maximizing sequence for $\mathcal{F}_n^{\eta, \varepsilon, h}$ and in analogy to before assume that $m_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N]$ and $b_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N]$ have limits $m_i^* > 0$, and $b_i^* \in D$ for $i = 1, \dots, n$, that $\mathcal{F}_n^{\eta, \varepsilon, h}(\mathbf{X}^N, \mathbf{g}^N)$ is monotonically increasing in N and that $m_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] \geq \frac{1}{2} m_i^*$. Consequently, the estimate

$$\mathcal{F}_n^{\eta, \varepsilon, h}(\mathbf{X}^0, \mathbf{g}^0) \leq \nu(D) \bar{\Phi} - \eta \left(\sum_{\alpha \in (I^h)^d} \sum_{i=1, \dots, n} |y_h^\alpha - x_i^N|^2 \chi_{i, \alpha}^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] \nu^\alpha \right)$$

$$\begin{aligned}
& + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] m_j^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N]}{|x_i^N - x_j^N|^2} \Bigg) \\
& \leq \bar{\Phi} - \eta \left(\frac{1}{2} \sum_{i=1, \dots, n} \text{dist}^2(x_i^N, \Omega) m_i^* + \frac{1}{4} \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^* m_j^*}{|x_i^N - x_j^N|^2} \right)
\end{aligned} \tag{5.5.6}$$

is obtained. This implies the following a priori bounds:

$$\text{dist}(x_i^N, \Omega) \leq \sqrt{\frac{2(\bar{\Phi} - \mathcal{F}_n^{\eta, \varepsilon, h}(\mathbf{X}^0, \mathbf{g}^0))}{\eta m_i^*}}, \quad |x_i^N - x_j^N| \geq \sqrt{\frac{\eta m_i^* m_j^*}{4(\bar{\Phi} - \mathcal{F}_n^{\eta, \varepsilon, h}(\mathbf{X}^0, \mathbf{g}^0))}}$$

for all N . For the uniform bound on \mathbf{g} , recall that \mathbf{g} and $\mathbf{g} + \lambda \mathbb{1}_n$ both result on the same Laguerre diagram. Hence, we may assume without loss of generality that $\mathbf{g}_1^N = 0$ for all $N \in \mathbb{N}$. Then, $\lim_{N \rightarrow \infty} \mathbf{g}_j^N = \infty$ for $j = 2, \dots, n$ would imply $\lim_{N \rightarrow \infty} m_1^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] = 0$, which contradicts our choice of n . Similarly, $\lim_{N \rightarrow \infty} \mathbf{g}_j = -\infty$ would imply $\lim_{N \rightarrow \infty} m_j^{\varepsilon, h}[\mathbf{X}^N, \mathbf{g}^N] = 0$, which again is a contradiction. Hence, $|\mathbf{g}^N| \leq C$ for some $C > 0$ and all $N \in \mathbb{N}$. Finally, given these a priori bounds the existence of a maximizer of $\mathcal{F}_n^{\eta, \varepsilon, h}$ follows directly from the Weierstraß extreme value theorem. \square

Next, we investigate the convergence of maximizers of the fully discrete, entropy regularized functional $\mathcal{F}_n^{\eta, \varepsilon, h}$ given in (5.5.4) to a maximizer of the original functional \mathcal{F}_n^η defined in (5.2.5) for both $\varepsilon \rightarrow 0$ and $h \rightarrow 0$.

Theorem 5.5.3. *Let $\Phi \in C(\Omega)$, $l \in \mathbb{N}$ and $\eta > 0$. For each $N \in \mathbb{N}$, consider maximizers $(\mathbf{X}^N, \mathbf{g}^N) \in \mathcal{X}_{\beta(N)} \times \mathbb{R}^{\beta(N)}$, where $\beta : \mathbb{N} \rightarrow \{1, \dots, n\}$, so that $\mathcal{F}_{\beta(N)}^{\eta, \varepsilon_N, h_N}(\mathbf{X}^N, \mathbf{g}^N) \geq \max_{k=1, \dots, n} \{\sup_{(\mathbf{X}, \mathbf{g}) \in \mathcal{X}_k \times \mathbb{R}^k} \mathcal{F}_k^{\eta, \varepsilon_N, h_N}(\mathbf{X}, \mathbf{g})\}$. Then, there is $n \in \{1, \dots, l\}$ such that up to the selection of a subsequence $(\mathbf{X}^N, \mathbf{g}^N)_{N \in \mathbb{N}} \subset \mathcal{X}_n \times \mathbb{R}^n$ converges to a limit $(\mathbf{X}^*, \mathbf{g}^*) \in \mathcal{X}_n \times \mathbb{R}^n$. Furthermore, it holds*

$$\mathcal{F}_n^\eta(\mathbf{X}^*, \mathbf{g}^*) \geq \max_{k=1, \dots, l} \left\{ \sup_{(\mathbf{X}, \mathbf{g}) \in \mathcal{X}_k \times \mathbb{R}^k} \mathcal{F}_k^\eta(\mathbf{X}, \mathbf{g}) \right\}.$$

Proof. The proof of the first statement is along the same lines as the proof of Theorem 5.2.5. For the last statement, notice that

$$\begin{aligned}
b_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] &= \sum_{\alpha \in \mathcal{I}^d} \frac{y_h^\alpha \chi_{i, \alpha}^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N]}{m_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N]} \nu^\alpha = \int_{\Omega} \frac{g^N(y) \chi_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N]}{m_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N](y)} d\nu(y) \\
\mathcal{R}_n^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] &= \sum_{i=1, \dots, n} \sum_{\alpha \in \mathcal{I}^d} |y_h^\alpha - x_i^N|^2 \chi_{i, \alpha}^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] \nu^\alpha \\
&\quad + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] m_j^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N]}{|x_i^N - x_j^N|^2} \\
&= \sum_{i=1, \dots, n} \int_{\Omega} f^N(y) \chi_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N](y) d\nu(y)
\end{aligned}$$

$$+ \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \frac{m_i^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N] m_j^{\varepsilon_N, h_N}[\mathbf{X}^N, \mathbf{g}^N]}{|x_i^N - x_j^N|^2},$$

where $f^N(y) := y_h^\alpha$ for $y \in \Omega_h^\alpha$, and $g^N(y) := |y_h^\alpha - x_i^N|^2$ for $y \in \Omega_h^\alpha$ are the piecewise constant approximations of the functions $y \mapsto y$ and $y \mapsto |y - x_i^N|^2$ consistent with our grid cells. The rest of the proof is then completely analogous to the proof of the respective statement in Theorem 5.4.3. \square

To numerically implement the maximization of 5.4.13 via a gradient ascent approach we need to compute the gradient of $\mathcal{F}_n^\varepsilon$. We specifically obtain the following derivatives for spatially discrete quantities using $z = x_k$ or $z = \mathbf{g}_k$ for $k = 1, \dots, n$

$$\begin{aligned} \partial_z \mathcal{F}_n^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] &= \sum_{j=1, \dots, n} \partial_z m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \Phi(b_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]) + m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \partial_z \Phi(b_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]), \\ \partial_{x_k} \mathcal{R}_n^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] &= 2 \sum_{\alpha \in \mathcal{I}^d} (x_k - y_h^\alpha) \chi_{k, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \nu^\alpha + \sum_{i=1, \dots, n} \sum_{\alpha \in \mathcal{I}^d} |y_h^\alpha - x_i|^2 \partial_{x_k} \chi_{i, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \nu^\alpha \\ &\quad + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \left(\frac{\partial_{x_k} m_i^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} + \frac{m_i^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \partial_{x_k} m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} \right. \\ &\quad \left. - 2 \frac{m_i^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^4} ((x_i - x_j) \delta_{ik} + (x_j - x_i) \delta_{jk}) \right), \\ \partial_{\mathbf{g}_k} \mathcal{R}_n^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] &= \sum_{i=1, \dots, n} \sum_{\alpha \in \mathcal{I}^d} |y_h^\alpha - x_i|^2 \partial_{\mathbf{g}_k} \chi_{i, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \nu^\alpha \\ &\quad + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \left(\frac{\partial_{\mathbf{g}_k} m_i^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} + \frac{m_i^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \partial_{\mathbf{g}_k} m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]}{|x_i - x_j|^2} \right), \\ \partial_z \Phi(b_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]) &= \nabla \Phi(b_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]) \cdot \partial_z b_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}], \\ \partial_z m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] &= \sum_{\alpha \in \mathcal{I}^d} \partial_z \chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \nu^\alpha, \\ \partial_z b_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] &= \frac{1}{m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]} \sum_{\alpha \in \mathcal{I}^d} y_h^\alpha \partial_z \chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \nu^\alpha - \frac{b_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] \partial_z m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]}{m_j^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]^2}, \\ \partial_{x_k} \chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] &= \frac{2x_j}{\varepsilon} \left(\chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] - \delta_{kj} \right) \chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}], \\ \partial_{\mathbf{g}_k} \chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] &= -\frac{1}{\varepsilon} \left(\chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}] - \delta_{kj} \right) \chi_{j, \alpha}^{\varepsilon, h}[\mathbf{X}, \mathbf{g}]. \end{aligned}$$

5.6 Numerical experiments

Now, we will apply the above derived method to compute the optimal partition of the unit cube $[0, 1]^2$ for given n , a function Φ and a probability density ν . In the presented numerical results, we assume ν to be the Lebesgue measure on the unit cube. Moreover, for every numerical experiment in this and the next section, we made use of the sparse multi-scale algorithm by Schmitzer [Sch19] and the implementation

in [CFG⁺21] to efficiently compute the Sinkhorn iterations. To optimize the parameters of the power diagrams, we used the *Adam* optimizer (cf. [Kin14]) as a stochastic gradient ascent method. Unless otherwise explicitly stated, we use $\eta = 0$, $N = 256$ (i.e. grid size $h = 1/256$), and initialize our algorithm with $n = 12$ sites/weights for each of the experiments in this and the next section. On each plot in this section, the barycenters (triangles) and sites (circles) are plotted, for visualisation purposes the latter ones only if inside the unit cube. On the left of each plot, the respective values of the component of the g vector is also plotted.

In Figure 5.1 we show the dependence of the optimal numerical solution on the entropy parameter ε . The interfaces between the Laguerre cells become fuzzier with increasing ε , whereas the structure of the optimal solutions does not change much.

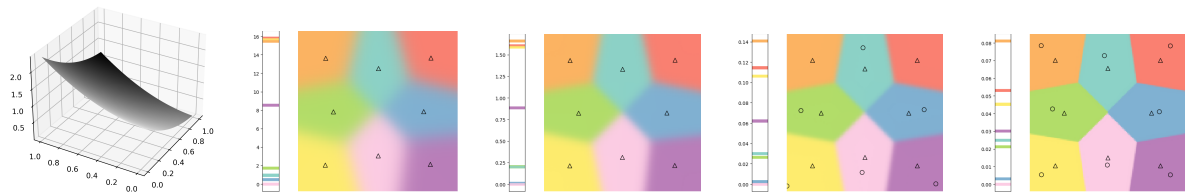


Figure 5.1: Convergence of the optimal power diagram for different entropy parameters $\varepsilon = 25, 5, 1, 0.2$. (four right-most plots), where the blur parameter values are given in units of $1/N$, for $N = 128$. The function Φ is plotted on the left-most panel.

Next, we plot in Figure 5.2 the optimal configurations for different values of η for a function Φ (left), which has global/local maxima of equal value at the points $(\frac{1}{2}, \frac{1}{4})$, $(\frac{3}{4}, \frac{3}{4})$, and $(\frac{1}{4}, \frac{3}{4})$. We know that the solution for $\eta = 0$ is not contained in the space of power diagrams, but we observe for $\eta \rightarrow 0$ the convergence of the solutions. Indeed, for large values of η a single cell has positive mass and no other cells contribute to the cost. For vanishing η , two additional cells appear as optimal configurations and the three cells meet in a triple-point converging to $(\frac{1}{2}, \frac{1}{2})$ with one angle converging to π .

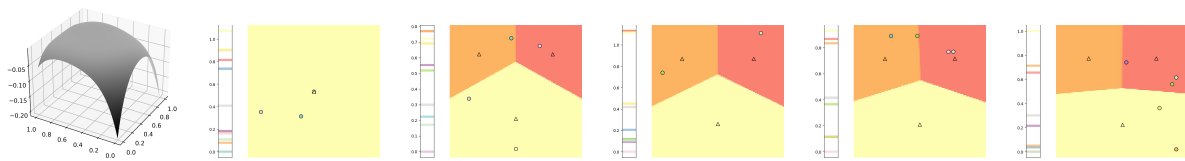


Figure 5.2: Convergence of the optimal power diagram for a function Φ (column (i)) with global maxima at $(0.25, 0.75)$, $(0.75, 0.75)$ and $(0.5, 0.25)$ for regularization parameters given by $\eta = 1e-1, 1e-2, 1e-3, 1e-4, 1e-5$ (columns (ii)-(vi)).

Finally, we check how our algorithm deals with fusing/pushing cells away, when the optimal solution requires a smaller number of sites/weights than what the algorithm was initialized with. To this end, in Figure 5.3 we consider a concave function Φ , that has the trivial partition as solution, and plot some iterations of our algorithm that show the proper recovery of this solution, i.e. cells disappear by either pushing the respective sites away (orange, green cells), or by increasing the difference between the g -values of the purple cell and those of the respective cells (beige, blue).

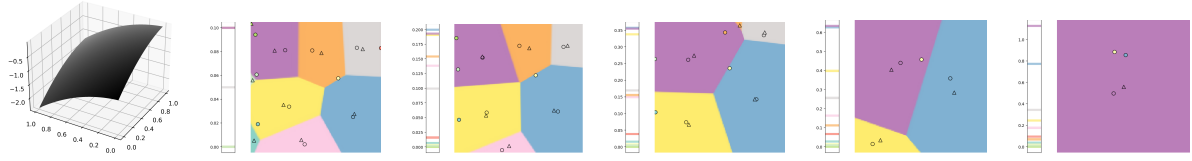


Figure 5.3: Convergence of the optimal power diagram for a concave function Φ (left). One sees how the algorithm automatically pushes cells outside of the relevant unit square to enact the trivial solution. Plotted are the computed power diagrams after iterations $it = 1, 2, 4, 8, 16$ (five right-most plots).

5.7 Application to information design

Finally, we consider the application of our algorithm to a particular problem in information design: the monopolist's problem. In this problem, a seller (sender) can disclose information about the qualities of the products she sells. There are two objects available individually at fixed prices p_1, p_2 , and the bundle of both objects can be bought at fixed price $p_3 = p_1 + p_2 + \Delta$ for some bundling “surcharge” $\Delta > 0$ or “discount” $\Delta < 0$. A consumer (receiver) has valuations $\mathbf{v} = (v_1, v_2)$ per unit of quality, distributed according to the Lebesgue measure on the unit square $[0, 1]^2$. The quality of the objects $\mathbf{q} = (q_1, q_2)$ is distributed on $[\underline{q}, \bar{q}]^2$ according to a measure μ that is absolutely continuous with respect to the Lebesgue measure. Throughout this section, we assume that μ is the (re-scaled) Lebesgue measure; the general case works analogously. The realized qualities are known to the seller but not to the buyer, and the realized valuations are known to the buyer but not to the seller.

We first consider a buyer with unit demand and then consider a buyer who demands more than one object and has additive valuations. A buyer has *unit demand* if she values at most one object. If the seller provides a signal about the qualities such that the buyer believes the expected qualities of the objects are (q_1, q_2) , then a buyer with unit demand buys only good $i = 1, 2$ if

$$q_i v_i - p_i > \max\{0, q_{-i} v_{-i} - p_{-i}\}$$

and buys nothing if

$$\max_i \{q_i v_i - p_i\} < 0.$$

Here, we ignore ties, which have probability zero of occurring. If a buyer with *additive valuations* buys only object i , her payoff is $q_i v_i - p_i$ but if she buys both objects her payoff is $q_1 v_1 + q_2 v_2 - p_3$. Therefore, a buyer with additive valuations buys only object i if

$$q_i v_i - p_i > \max\{0, q_{-i} v_{-i} - p_{-i}, q_1 v_1 + q_2 v_2 - p_3\},$$

buys both objects if

$$q_1 v_1 + q_2 v_2 - p_3 > \max\{0, q_1 v_1 - p_1, q_2 v_2 - p_2\},$$

and buys nothing otherwise.

For fixed expected qualities (q_1, q_2) , let $C_i(q_1, q_2)$ be the probability assigned to the set of valuations for which the buyer only buys object $i = 1, 2$, let $C_3(q_1, q_2)$ denote the probability assigned to the set of valuations for which the consumer buys both objects, and let $C_0(q_1, q_2)$ be the probability assigned to the set of valuations where the consumer buys nothing. Recall that these depend on whether we consider a consumer with unit demand or with additive valuations, and are computed from the uniform distribution

of valuations: these are the areas of the respective convex polygons defined by the above inequalities. The seller's revenue is then given by

$$R(q_1, q_2) = \sum_{i=1}^3 p_i C_i(q_1, q_2).$$

Recalling the characterization of exposed points, the seller chooses a Laguerre diagram $\pi = \{D^1, D^2, \dots, D^n\}$ of $[q, \bar{q}]^2$ with respective barycenters $\mathbf{q}^j = (q_1^j, q_2^j)$, $j = 1, 2, \dots, n$. In other words, for each realization of qualities the seller reveals to the consumer only to which cell the qualities belong. The consumer then updates her belief about the expected qualities to the barycenter \mathbf{q}^j of the cell D^j which contains the true qualities.

The designer's expected revenue is then given by

$$\sum_{j=1}^n \mu(D^j) R(q_1^j, q_2^j)$$

and the designer chooses a Laguerre diagram to maximize this expected revenue.

Unit demand. We first consider a buyer with unit demand. Figure 5.4 shows that the seller discloses only a coarse signal about the qualities of the products. Each cell corresponds to one expected quality pair of the products which we indicate by a triangle in the figure. There are at most 4 cells in the optimal signal: These cells can be roughly interpreted as corresponding to different quality pairs: (1) the orange cell corresponds to both products being of low expected quality, (2) the purple cell corresponds to product 1 being of low quality and product 2 being of high quality, (3) the red cell corresponds to product 1 being of high quality and product 2 being of low quality, and (4) the green cell corresponds to both products being of high quality. In the orange cells, the expected qualities are so low that the buyer never buys either of the products, independent of their valuations. In the red and purple cells, the expected quality of the lower-quality object is so low, that the buyer will never buy the lower-quality object; the buyer will either buy the higher-quality object or buy nothing at all. As the price of product 2 increases, in the optimal information policy, a signal indicating a high quality of product 2 becomes more informative in that it indicates a higher expected quality. This offsets the higher price and still induces some buyer types to purchase the more expensive product. If the price of product 2 becomes too high, no types will buy product 2 and the optimal Laguerre diagram has only two cells, as the right-most panel illustrates.

Under the optimal Laguerre diagram, the seller does not provide full information to the buyer even though that would raise efficiency. By revealing only imprecise information, the buyer's information rents are reduced and the seller's revenue raised. To evaluate the benefit of choosing an optimal information policy induced by a Laguerre diagram, we consider additional, non-optimal information policies, as benchmarks. Table 5.1 specifies the revenue generated under various information policies. It shows in row (i) the different values of the price p_2 , in row (ii) the revenue R_{opt} induced by the optimal Laguerre diagram partition computed with our algorithm, in row (iii) the revenue $R(1, 1)$ for an information policy where no information is given to the buyer, and in row (iv) the revenue for an additional benchmark information policy based on a partition generated with Lloyd's algorithm. Here, for the same number of cells as in the computed optimal Laguerre diagram, we generate with Lloyd's algorithm a Laguerre diagram partition which imposes the barycenter of each cell to coincide with the respective site, cf. [Llo82]. The respective Lloyd diagrams used for benchmarking are exemplified in Figure 5.7. Row (v) of Table 5.1

shows the revenue $\mathbb{E}(R)$ in case of the full information policy, and in row (vi) the percentage increase in revenue of the optimal power diagram policy compared to the full information policy is displayed. As one can see, optimal information design creates significant value to the seller: in this example it increases profits relative to full information in excess of 10%.

Figure 5.5 shows the optimal Laguerre diagrams for prices $p_1 = 1$ and $p_2 = 1.25$ and various values of the lower bound on the quality. As the lower bound increases, the optimal signals become less likely to contain significant information about the quality of object 2. The corresponding revenue and the revenue under alternative benchmarks is shown in Table 5.2.

p_2	1	1.25	1.5	1.75	2
R_{opt}	0.3153	0.2648	0.2175	0.1839	0.1716
$R(1, 1)$	0.0000	0.0000	0.0000	0.0000	0.0000
R_{Lloyd}	0.3056	0.2543	0.1667	0.1667	0.1667
$\mathbb{E}(R)$	0.2833	0.2407	0.1977	0.1657	0.1534
$pp(\%)$	11.30	10.01	10.01	10.98	11.86

Table 5.1: Monopolist's problem with unit demand, price $p_1 = 1$, quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$ for different prices p_2 (first row) under selected information policies.

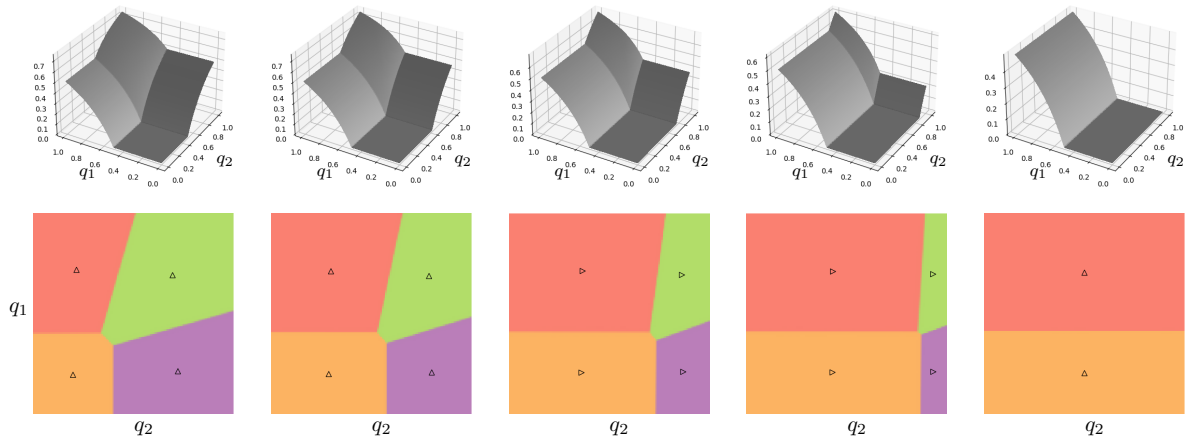


Figure 5.4: Optimal configurations for the monopolist's problem with unit demand, with prices $p_1 = 1$ and $p_2 = 1, 1.25, 1.5, 1.75, 2$ (second row, from left to right) and quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$. The respective revenue function R for each case is plotted in the first row.

Additive valuations. Figure 5.6 illustrates optimal information policies induced by Laguerre diagrams if the buyer has additive valuations. The optimal Laguerre diagrams become richer, with up to seven cells. Moreover, there is significant variation in the shape of the diagrams as the bundling surplus/discount varies. As Tables 5.3 and 5.4 show, there is again a significant benefit to only partially revealing information compared to fully revealing the qualities.

\underline{q}	0.25	0.5	0.75	1.	1.25
\mathbf{R}_{opt}	0.2999	0.3457	0.4074	0.4853	0.5687
$R(1, 1)$	0.1111	0.2000	0.3564	0.4757	0.5687
R_{Lloyd}	0.2838	0.3105	0.3604	0.4583	0.5687
$\mathbb{E}(R)$	0.2728	0.3146	0.3714	0.4528	0.5508
$pp(\%)$	11.30	10.01	10.01	10.98	11.86

Table 5.2: Monopolist problem with unit demand, prices $p_1 = 1, p_2 = 1.25$, upper quality bound $\bar{q} = 2$ for different lower quality bounds $\underline{q} = 0.25, 0.5, 0.75, 1., 1.25$ (first row) under selected information policies.

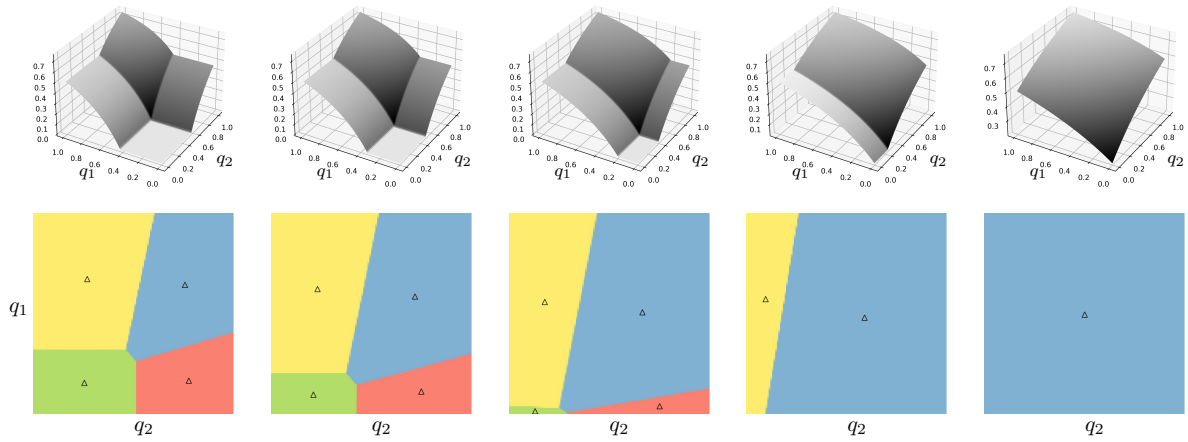


Figure 5.5: Optimal configurations for the monopolist's problem with unit demand, prices $p_1 = 1$ and $p_2 = 1.25$, upper quality bound $\bar{q} = 2$ and, from left to right, lower quality bounds $\underline{q} = 0.25, 0.5, 0.75, 1, 1.25$ (second row). The respective revenue function R (defined above) is plotted in the first row.

Δ	-1	-0.875	-0.75	-0.625	-0.5	-0.375	-0.25	-0.125
\mathbf{R}_{opt}	0.6577	0.5914	0.5215	0.4602	0.4196	0.3944	0.3754	0.3582
$R(1, 1)$	0.5000	0.4307	0.3516	0.2686	0.1875	0.1143	0.0547	0.0146
R_{Lloyd}	0.6528	0.5868	0.4824	0.3843	0.3510	0.3433	0.3533	0.3436
$\mathbb{E}(R)$	0.6417	0.5196	0.4578	0.4138	0.3803	0.3544	0.3343	0.3188
$pp(\%)$	2.49	13.82	13.91	11.21	10.33	11.29	12.29	12.36

Table 5.3: Monopolist's problem with additive valuations for different bundling discounts Δ (first row), quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$ and prices $p_1 = p_2 = 1$, under selected information policies.

5.8 Applications to Gerrymandering

As an additional application of the framework developed in this chapter, we present here an extension thereof to the optimal gerrymandering problem. When speaking of the problem of gerrymandering, one

Δ	0	0.125	0.25	0.375	0.5	0.625	0.75
\mathbf{R}_{opt}	0.3432	0.3303	0.3205	0.3154	0.3153	0.3153	0.3153
$R(1, 1)$	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
R_{Lloyd}	0.3333	0.3231	0.3142	0.3079	0.3056	0.3056	0.3056
$\mathbb{E}(R)$	0.3069	0.2979	0.2916	0.2875	0.2851	0.2839	0.2834
$pp(\%)$	11.83	10.88	9.91	9.70	10.59	11.06	11.26

Table 5.4: Monopolist's problem with additive valuations for different bundling surcharges Δ (first row), quality boundaries $\underline{q} = 0$ and $\bar{q} = 2$ and prices $p_1 = p_2 = 1$ under selected information policies.

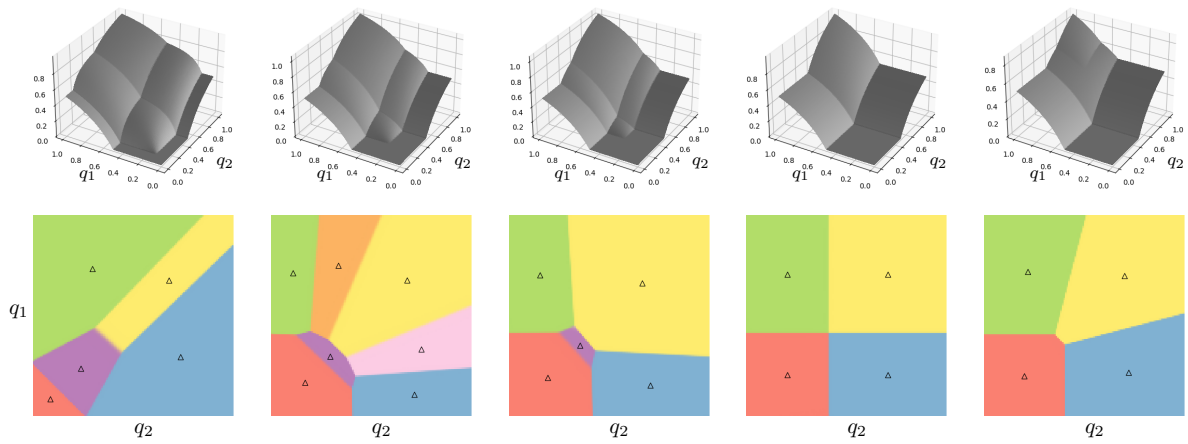


Figure 5.6: Optimal configurations for the monopolist's problem with additive valuations, with prices $p_1 = p_2 = 1$, quality bounds $\underline{q} = 0$ and $\bar{q} = 2$ and, from left to right, bundling surcharges/discounts $\Delta = -0.75, -0.5, -0.375, 0, 0.375$ (second row). The respective revenue function R (defined above) is plotted in the first row.

almost exclusively refers to the problem that arises in two-party systems, such as in the American democratic system. Therein, the process of redistricting, i.e. the drawing of electoral boundaries, is left to the state legislatures and governors themselves. History has shown that these institutions (independently of the political affiliation) are self-interested and have used the redistricting process to achieve political objectives. When the redistricting process turns partisan, one speaks instead of a gerrymandering process. In explicit, in the gerrymandering problem, a gerrymanderer observes a noisy signal from a continuous distribution of voter preferences, and, taking into account an aggregate preference shock, creates N contiguous districts with an equal amount of voters to maximize the number of districts (or seats) that she wins.

As of 1980, the main requirements of the drawn districts are both contiguity and population size equality across districts. The contiguity constraint, however, has been stretched to the limit in some cases, having multiple clusters of voters far away from each other, only joined by narrow strips of land between them. Numerous tests have been developed to prevent gerrymandering to take place; however, none of these tests were deemed to be workable or justiciable by Supreme Court decisions. For a deeper insight in the legal and institutional backdrop against which gerrymandering takes place, we refer to



Figure 5.7: Plot of the Lloyd's centroidal power diagrams for $n = 1, 2, 4, 5, 7$ cells from left to right. The sites coincide with the barycenters of the cells.

[FH08, Section I].

In this section, we will detail a voter distribution model, inspired by the examples in [FH08, Section II]. Furthermore, as a continuation of the discussion in this chapter, we will restrict the gerrymandering problem to partitions of the district (usually a federal state) induced by power diagrams. We will prove that under this framework, an optimal gerrymandering exists, and develop tests to detect and prevent said gerrymandering from happening.

Similarly as above, let (D, \mathcal{D}, ν) be a probability space, where $D \subseteq \mathbb{R}^d$ is a compact and convex domain, \mathcal{D} is its Borel σ -algebra, and ν a probability measure which is assumed to be absolutely continuous with respect to the Lebesgue measure \mathcal{L} , with Radon-Nikodym density $\frac{d\nu}{d\mathcal{L}} \in L^\infty(\mathcal{L})$. Moreover, let $n \in \mathbb{N}$ be a fixed number of districts for D to be partitioned in. Given a vector of sites $\mathbf{X} = (x_1, \dots, x_n) \in \mathcal{X}_n$, define the discrete measure

$$\mu[\mathbf{X}] := \frac{1}{n} \sum_{i=1, \dots, n} \delta_{x_i}$$

with equal mass concentrated on each site, and let $\mathbf{g}[\mathbf{X}] \in \mathbb{R}^n$ be an optimal vector for (5.3.4) for the L^2 -Wasserstein distance between ν and $\mu[\mathbf{X}]$. Then, $(L_i[\mathbf{X}], \mathbf{g}[\mathbf{X}])_{i=1, \dots, n}$ defines a partition of D induced by a power diagram in n subsets, each of which has equal ν -mass. Given a partition of D in n subsets of equal mass induced by an n -tuple of sites \mathbf{X} as above, define the *voter density in district i* by $\nu_i[\mathbf{X}] := \nu[\cdot | L_i[\mathbf{X}, \mathbf{g}[\mathbf{X}]]] = n \int_{L_i[\mathbf{X}, \mathbf{g}[\mathbf{X}]]} \chi_{[\cdot]} d\nu$.

Let $\hat{P} \in C(D; \mathbb{R})$ be a *political preference function*, i.e., it maps voters located at $x \in D$ to their political preference *bliss point* $\hat{P}(x) \in \mathbb{R}$. In particular, we make the simplifying assumption that all voters in the same location have the same bliss point. Given voter densities $\nu_i[X]$ for $i = 1, \dots, n$ as above, let $\lambda_i := \hat{P}_\# \nu_i$ be the political preference distribution of voters in the i -th district, and \hat{p}_i a random variable distributed according to λ_i . To account for the uncertainty with which the individual preferences are measured by the political parties (through polls, studies, etc.), we instead consider the true political preference random variable $p_i^\epsilon := \hat{p}_i + \eta_\epsilon$, where η_ϵ is a random variable with distribution $\mathcal{N}(s, \epsilon^2)$, and independent of \hat{p}_i . Let λ_i^ϵ be the law of p_i^ϵ . Then, λ_i^ϵ is absolutely continuous with respect to the Lebesgue measure on \mathbb{R} , and its Radon-Nikodym density $\tilde{\lambda}_i^\epsilon$ is explicitly given by

$$\tilde{\lambda}_i^\epsilon(s) = \int_{\mathbb{R}} \tilde{\eta}_\epsilon(s - t) d\lambda_i(t),$$

where $\tilde{\eta}_\epsilon$ is the density function of a centered one-dimensional Gaussian distribution with covariance equal to ϵ^2 . Note that we are not concerned with the limit $\epsilon \rightarrow 0$, as we always assume the voter preference is always measured with some degree of noise.

Define now the cumulative distribution function of λ_i^ϵ as $C_i^\epsilon(t) := \int_0^t \tilde{\lambda}_i^\epsilon(s) ds$. Recall that by the definition of λ_i^ϵ , it is absolutely continuous with respect to the Lebesgue measure on \mathbb{R} . Moreover, the Radon-Nikodym derivative of this measure is also strictly positive. As a result, C_i^ϵ is strictly increasing, with $\lim_{t \rightarrow -\infty} C_i^\epsilon(t) = 0$ and $\lim_{t \rightarrow \infty} C_i^\epsilon(t) = 1$. This means that the inverse function $Q_i^\epsilon := (C_i^\epsilon)^{-1} : (0, 1) \rightarrow \mathbb{R}$ is well defined.

In this framework, we assume that there are only two parties involved, L and R . Each party has a *party policy* $x^R > x^L$, and base party strengths $g^R, g^L \in \mathbb{R}$. In more rudimentary models, a voter with bliss point $P(x) \in \mathbb{R}$ will cast their vote for the party that minimizes the cost $|P(x) - x^j|^2 - g^j$ for $j \in \{L, R\}$ with probability 1. We instead assume that an aggregate preference shock between the drawing of districts and the actual vote might affect the party strengths g^L and g^R and thus influence the choice of said voter in the future. Hence, we assume that the realized party strengths \hat{g}^R, \hat{g}^L are non-deterministic random variables with mean given by g^R, g^L , respectively. Since only the difference $g^R - g^L$ play a role, we may assume without loss of generality that $\hat{g}^L \equiv 0$ and only \hat{g}^R varies. Examples of events that might induce an aggregate preference shock include political events, economics news, events of *force majeure* such as natural disasters, likeability of candidates or how the different campaigns are run. Accounting for this, we will explicitly assume that a voter with bliss point p_i^ϵ casts her vote for party $j \in \{L, R\}$ with probability

$$\frac{\exp\left(-\frac{|p_i^\epsilon - x^j|^2 - g^j}{\epsilon_v}\right)}{\sum_{k \in \{L, R\}} \exp\left(-\frac{|p_i^\epsilon - x^k|^2 - g^k}{\epsilon_v}\right)},$$

where ϵ_v is a regularization parameter encoding the uncertainty of an aggregate preference shock.

A peculiarity of this model is that if a voter with bliss point p_i^ϵ casts their vote for party R , then each voter with $p > p_i^\epsilon$ will cast their vote for the same party. Similarly, if said voter casts their vote for party L , then each voter with $p < p_i^\epsilon$ will cast their vote for party L as well. Hence, voter preferences satisfy the *single-crossing* property (cf. [FH08]). In particular, this means that whichever party receives the vote of the median voter in the i -th district (which has voting preference $Q_i^\epsilon(\frac{1}{2})$) will end up winning it. Hence, a gerrymanderer belonging to party $j \in \{R, L\}$ will try to maximize the expected value of total won seats/districts

$$\mathcal{G}_n[\mathbf{X}] := \sum_{i=1}^n \frac{\exp\left(-\frac{|Q_i^\epsilon(1/2)[\mathbf{X}] - x^j|^2 - g^j}{\epsilon_v}\right)}{\sum_{k \in \{L, R\}} \exp\left(-\frac{|Q_i^\epsilon[\mathbf{X}](1/2) - x^k|^2 - g^k}{\epsilon_v}\right)}. \quad (5.8.1)$$

among all permissible partitions of the domain D . In this chapter, we restrict ourselves to power diagram partitions, in order to leverage the algorithm introduced in previous sections. In contrast to the information design functional defined in (5.2.2), note that \mathcal{G}_n is no longer maximized over the weights \mathbf{g} . This is a result of the fact that in the gerrymandering problem, the cell masses are fixed. Hence, choosing a set of location sites \mathbf{X} already completely determines (up to an additive constant) the resulting weight vector \mathbf{g} that is required to realize the cell masses. To guarantee a resulting maximizer of (5.8.1), we proceed as before with the addition of a penalty term

$$\mathcal{G}_n^\eta[\mathbf{X}] := \mathcal{G}_n[\mathbf{X}] - \eta \mathcal{R}_n[\mathbf{X}, \mathbf{g}[\mathbf{X}]], \quad (5.8.2)$$

where \mathcal{R}_n is the penalty term defined in (5.2.6). We obtain the following existence result:

Theorem 5.8.1. *Let $D \subset \mathbb{R}^d$ be a compact an convex domain, $\nu \in \mathcal{P}(D)$ that is absolutely continuous with respect to the Lebesgue measure on D , and $n \in \mathbb{N}$ be a fixed number of districts. For $\mathbf{X} \in \mathcal{X}_n$, let $(L_i[\mathbf{X}, g[\mathbf{X}]])_{i=1, \dots, n}$ be the resulting power diagram with cells of equal mass n^{-1} , and $\nu_i[\mathbf{X}](\cdot) := n \int_{L_i[\mathbf{X}, g[\mathbf{X}]]} \chi_{[\cdot]} d\nu$. Moreover, let $\hat{P} \in C(D; \mathbb{R})$ be a continuous function and $\tilde{\lambda}_i^\epsilon[\mathbf{X}](s) = \int_D \tilde{\eta}_\epsilon(s - \hat{P}(y)) d\nu_i[\mathbf{X}](y)$, where η_ϵ is the Gaussian probability density function centered at 0 with covariance ϵ^2 . Define the cumulative distribution function $C_i^\epsilon[\mathbf{X}](t) := \int_{-\infty}^t \tilde{\lambda}_i^\epsilon[\mathbf{X}](s) ds$, and its inverse $Q_i^\epsilon[\mathbf{X}] := (C_i^\epsilon[\mathbf{X}])^{-1}$. For all party policies $x^L < x^R$ and base party strengths $g^L, g^R \in \mathbb{R}$ and for all $\eta > 0$ and $\epsilon_v > 0$, there exists a maximizer $\mathbf{X}^* \in \mathcal{X}_n$, of (5.8.2).*

Proof. Let $(\mathbf{X}^N)_{N \in \mathbb{N}}$ be a maximizing sequence for \mathcal{G}_n^η , and define $\bar{\mathcal{G}} := \mathcal{G}_n^\eta[\mathbf{X}^0] > -\infty$. Notice that $\bar{\mathcal{F}}_G < n$ by its definition. Moreover, it holds for sufficiently large $K \in \mathbb{N}$:

$$\bar{\mathcal{G}} \leq \mathcal{G}_n^\eta[\mathbf{X}^K] \leq n - \eta \mathcal{R}_n[\mathbf{X}^K, \mathbf{g}[\mathbf{X}^K]] \leq n - \eta \left(\sum_{\substack{i,j=1 \\ i \neq j}}^n \frac{1}{n^2 |x_i^K - x_j^K|^2} + \sum_{i=1}^n \text{dist}^2(x_i^K, D) \frac{1}{n} \right).$$

Hence, we obtain the estimates

$$\text{dist}(x_i^N, \Omega) \leq \sqrt{n \frac{n - \bar{\mathcal{G}}}{\eta}}, \quad |x_i^N - x_j^N| \geq \sqrt{\frac{\eta n^{-2}}{n - \bar{\mathcal{G}}}}.$$

As before, the sequence $(\mathbf{X}^N)_N$ is in a compact subset of \mathcal{X}_n , so up to a subsequence, it is convergent to some $\mathbf{X}^* \in \mathcal{X}_n$. First, note that \mathcal{R} is continuous in \mathbf{X} . Next, for $\mu[\mathbf{X}^N] := \sum_{i=1}^n \delta_{x_i^N}$ let $\mathbf{g}[\mathbf{X}^N]$ be optimal for (5.3.4) (for $\mu[\mathbf{X}^N]$ and ν). The optimal coupling $\Pi[\mathbf{X}^N] \in \mathcal{P}(\mathbb{R}^d \times D)$ for the primal problem (5.3.1) can then be explicitly given by

$$d\Pi[\mathbf{X}^N](x, y) = \frac{1}{n} \sum_{i=1}^n d\delta_{x_i^N}(x) \otimes d\nu_i[\mathbf{X}^N](y).$$

Recall that since \mathbf{X}^N is a convergent sequence, we have in particular that the measure sequence $(\mu[\mathbf{X}^N])_N$ narrowly converges to $(\mu[\mathbf{X}^*])_N$. In particular, this measure sequence is tight, and hence, the sequence of optimal couplings $\Pi[\mathbf{X}^N]$ is tight as well, and by Prokhorov's theorem (cf. [Pro56]) a subsequence is narrowly convergent to some $\Pi^* \in \mathcal{P}(\mathbb{R}^d \times D)$. By the stability of optimal couplings under narrow convergence of measures, we have that $\Pi^* = \Pi[\mathbf{X}^*]$, which implies the narrow convergence of

$$(\nu_i[\mathbf{X}^N])_N \rightharpoonup \nu_i[\mathbf{X}^*],$$

for all $i = 1, \dots, n$.

By the definition of narrow convergence and the continuity and boundedness of \hat{P} and $\tilde{\eta}_\epsilon$, we then obtain

$$\tilde{\lambda}_i^\epsilon[\mathbf{X}^N](s) = \int_D \tilde{\eta}_\epsilon(s - \hat{P}(y)) d\nu_i[\mathbf{X}^N](y) \rightarrow \int_D \tilde{\eta}_\epsilon(s - \hat{P}(y)) d\nu_i[\mathbf{X}^*](y) = \tilde{\lambda}_i^\epsilon[\mathbf{X}^*](s),$$

for all $i = 1, \dots, n$ and all $s \in \mathbb{R}$. Furthermore, we have

$$|\tilde{\lambda}_i^\epsilon[\mathbf{X}^N](s)| \leq \int_D \left| \tilde{\eta}_\epsilon(s - \hat{P}(y)) \right| d\nu_i[\mathbf{X}](y) \leq C(\epsilon) < \infty.$$

By the dominated convergence theorem, we obtain

$$\lim_{N \rightarrow \infty} C_i[\mathbf{X}^N](t) = \lim_{N \rightarrow \infty} \int_{-\infty}^t \tilde{\lambda}_i^\epsilon[\mathbf{X}^N](s) \, ds = \int_{-\infty}^t \tilde{\lambda}_i^\epsilon[\mathbf{X}^*](s) \, ds = C_i[\mathbf{X}^*](t)$$

for all $i = 1, \dots, n$ and all $t \in \mathbb{R}$. Since $C_i[\mathbf{X}^N]$ and $C_i[\mathbf{X}^*]$ are strictly increasing functions, their inverses $Q_i[\mathbf{X}^N] : (0, 1) \rightarrow \mathbb{R}$ and $Q_i[\mathbf{X}^*] : (0, 1) \rightarrow \mathbb{R}$ exist and are also continuous. Together with $\lim_{N \rightarrow \infty} C_i[\mathbf{X}^N](t) = C_i[\mathbf{X}^*](t)$, one finally obtains

$$\lim_{N \rightarrow \infty} Q_i[\mathbf{X}^N](t) = Q_i[\mathbf{X}^*](t).$$

Combined with the estimates shown in the beginning of this proof, the above implies the claim. \square

Using a slightly adapted version of our original algorithm for information design, we compute maximizers of (5.8.1) from the perspective of a republican gerrymanderer for the 2018 gubernatorial race in the US State of Ohio in Figure 5.8. In 2018, the republican party defeated the democratic party in the gubernatorial election, with a margin of 50.40% to 46.67%. Yet, due to the highly gerrymandered landscape of this state, the republican party won 12 out of the total 16 seats, cf. middle picture. In Fig. 5.8, our algorithm was able to find a power diagram re-districting of Ohio, such that this margin was even higher, 14 to 2, cf. fourth picture from the left (if the winner of each district is determined by the party with the higher probability of winning it). Tweaking our algorithm, one can instead optimize from the perspective of a democratic gerrymanderer. In this case, our algorithm was able to find a partition with a 7-9 democratic-favoured split, which is still surprisingly high accounting for the fact that they obtained fewer votes than their rival party.

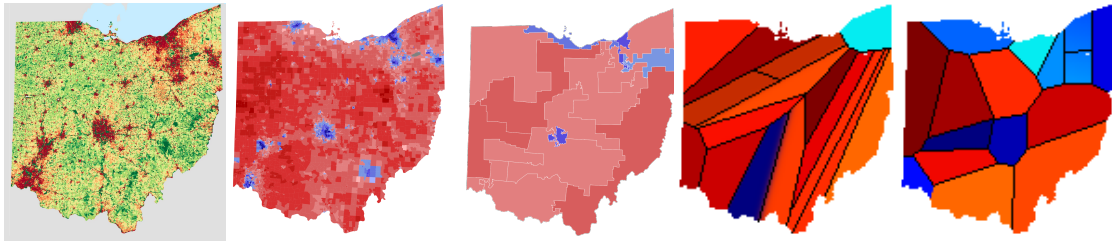

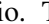


Figure 5.8: Left: Population density ν map of Ohio. Middle-left: Political preference map \hat{P} , with values -1  $+1$. Middle: Map showing the current congressional map of Ohio with current majorities. Middle right: Optimal power diagram gerrymandering (from the point of view of the Republican party) of Ohio. Right: Optimal power diagram gerrymandering (from the point of view of the Democratic party) of Ohio. The colorbar represents the expected percentage of republican voters in the rightmost two figures:  100. All computations were made with $n = 16$.

Acknowledgement

The work in this chapter was supported by the German Research Foundation (DFG) via Germany's Excellence Strategy project 390685813 – Hausdorff Center for Mathematics.

6 Summary

Approximation of Splines in Wasserstein Spaces

Conclusions

In Chapter 3, a time-discrete variational model was examined to compute spline paths within the space of probability measures equipped with the Wasserstein-2 metric. These spline paths were defined as measure-valued trajectories that minimized a spline energy while satisfying interpolation constraints and boundary conditions. Our method represents a significant improvement in terms of temporal smoothness of the interpolated frames, similar to how cubic splines in Euclidean spaces are a smoother interpolation method

To extend this framework to continuous splines on the space of probability measures, a compatible geometric structure was required. Leveraging the geometric structure of Wasserstein spaces as outlined in [BB00] and [Ott01] as well as the second-order analysis in this space (cf. [Gig12]), energy splines were defined as the minimizers of total squared acceleration (i.e. the covariant derivative of the velocity) (3.2.2) under additional boundary conditions, given in (3.2.4)-(3.2.6). However, this functional posed challenges, being computationally intractable and, unlike the action functional (3.1.2), non-convex.

To remedy this last issue, a consistent variational time discretization of E-splines was proposed instead, accompanied by an efficient and robust algorithm to construct them. The time discretization of energy splines relied on a general framework for variational time discretization of splines on Riemannian manifolds, originally introduced in [HRW18]. This framework's core elements included a functional \mathcal{W} , which approximated the squared Riemannian distance between nearby objects on the manifold. In our case, for probability measures the local functional \mathcal{W} corresponded to the squared Wasserstein distance, and the approximate average was naturally defined as the Wasserstein barycenter. In the variational-time setting, we were able to show existence of discrete splines, defined as minimizers of the discrete spline energy functional (cf. (3.2.17)) in the general case. A key role for this result was played by the addition of a small fraction of the usual discrete path energy functional to the pure spline functional, which seems to be indispensable. Furthermore, it was analytically demonstrated that this approximation aligned with the Riemannian geometry of Wasserstein spaces in the restricted Gaussian case, in the sense of Mosco convergence.

We validated our interpolation method through extensive numerical experiments, including synthetic examples and real-world applications. The results highlighted its versatility and capability to produce smooth, meaningful interpolations in scenarios ranging from image and measure morphing to synthesis of texture interpolants. Moreover, we were able to leverage recent advances in machine learning in the form of variational autoencoders to dramatically increase the scope of our interpolation method to any type of data that can be represented as realization of a VAE. Finally, by using our proposed spline functional as a penalty for the deviation from a geodesic, we proposed an efficient algorithm to construct discrete linear regressions on distributional data.

The computational framework, powered by entropy-regularized optimal transport and state-of-the-art implementations of the Sinkhorn algorithm, demonstrated scalability and efficiency, making it viable for large-scale problems.

Outlook and Extensions

Even though very important insights were won in this chapter, there are still some open questions that could be explored further in future research. For instance, the result on the Mosco convergence of discrete splines to continuous ones was only proven in the restricted case of diagonal Gaussian probability measures; a more general result would be desirable. Moreover, error estimates of the form

$$\mathcal{W}(\mu(t_k^K), \mu_k^K),$$

where μ is a continuous measure curve, and μ_k^K is the computed interpolant evaluated at time t_k^K are interesting from an analytical point of view.

An extension of our interpolation method would be interpolations of an even higher order (i.e. the equivalent of quintic, septic, etc. splines). Even though these kind of splines are sparsely used in the Euclidean setting, they could still be relevant in light of Section 3.8, since penalizing higher order derivatives would lead to more general polynomial regressions of higher order. In fact, using the fact that the third-order forward differences have the following form

$$\ddot{x} \approx K^3 (x_{k+2} - 3x_{k+1} + 3x_k - x_{k-1}) = 4K^3 \left(\frac{x_{k+2} + 3x_k}{4} - \frac{3x_{k+1} + x_{k-1}}{4} \right),$$

we conjecture that the corresponding third-order discrete spline energy might be given by

$$16K^5 \sum_{k=1}^{K-2} \mathcal{W}^2 \left(\text{Bar}^{\frac{1}{4}}(\mu_{k+2}, \mu_k), \text{Bar}^{\frac{1}{4}}(\mu_{k+1}, \mu_{k-1}) \right).$$

Unbalanced Transport Splines

Conclusions

In Chapter 4, we provided a detailed review of the foundational models underlying our study, including the optimal transport framework, the flow of diffeomorphisms, and the associated path energy functionals. These path energies are characterized by their minimizers, which represent geodesic paths within the respective shape spaces. Building on this theoretical groundwork, we combined the two models to develop a generalized image transport framework for measures that account for differences in mass, following the approach introduced in [MRSS15].

From this foundation, we derived a novel generalized spline energy functional, which serves as a natural extension of the existing framework. This functional enabled us to define spline curves that minimize the generalized spline energy while satisfying interpolation constraints imposed by a given set of key frames at distinct times. In addition to the continuous formulation, we also explored a time-discrete variational spline interpolation model as a generalization of the first-order time-discrete image transport model initially proposed in [MRSS15]. The central theoretical contribution of this chapter was demonstrating the consistency and Mosco convergence of the time-discrete spline energy to the continuous spline energy in the context of image transport.

To make this approach computationally practical, we further discretized the model in space and developed an efficient numerical scheme based on the iPALM algorithm [PS16]. This scheme enabled us to compute fully discrete image transport spline paths in the space of image densities. Finally, we were able to successfully test the robustness and flexibility of our method in the numerical experiments presented in Section 4.8.

The resulting methodology not only extended the existing theoretical framework but also provided a robust numerical tool for solving real-world problems involving image transport and interpolation. In doing so, we also managed to solve the two issues that we identified at the beginning of this thesis, and which were the main motivation for this chapter: difference in mass images, and topological consistency of the transport maps.

Outlook and Extensions

As mentioned multiple times throughout this thesis, the spline functional (4.2.11) we considered in this chapter is not intended nor expected to be Riemannian in nature. Even though the numerical experiments show an improved quality and visual appeal with respect to the piecewise geodesic interpolation, we ignore the interplay between different terms in the second order; as a result, there are no cross-terms in 4.2.11. Regarding the discrete spline functional, we conjecture it has the following form

$$\sum_{k=1}^{K-1} \mathcal{W}_{\delta,\lambda}^2(\vartheta_k, \text{Bar}(\vartheta_{k-1}, \vartheta_{k+1}))$$

where $\mathcal{W}_{\delta,\lambda}^2$ is the generalized Riemannian Wasserstein distance defined in (4.1.8). Moreover, as in the previous chapter one could consider higher order interpolation, which would be defined *mutatis mutandis*. The related work on metamorphosis splines (cf. Fig. 6.1) can also be considered an extension of this method.

Entropy-Regularized Optimal Transport in Information Design

Conclusions

Inspired by a challenging multi-dimensional information design problem, we embarked on developing a framework that utilizes the theoretical insights obtained in [KMSW24] regarding Lipschitz-exposed points of sets of fusions F_ν for a given probability measure ν . Building upon these findings, we reformulated an associated information design problem that focuses on optimizing over power diagrams, a mathematical representation well-suited to this context. This approach allowed us to connect the theoretical underpinnings of information design with practical optimization techniques.

To make the problem computationally tractable, we re-framed it using the perspective of semi-discrete optimal transport, a versatile and powerful tool in this domain. By further incorporating entropy regularization and a penalty term into this reformulation, we additionally enhanced the model's numerical stability and efficiency. In doing so, we also obtained theoretical guarantees of existence of optimal solutions. This enabled us to derive a practical and highly effective numerical, fully discrete algorithm, specifically tailored to solving the class of optimization problems at hand.

We subjected the proposed method to rigorous synthetic numerical evaluation to validate its effectiveness and robustness. Additionally, we applied it to a specific instance of an information design problem involving a multi-product monopolist. In this scenario, the monopolist strategically designs the information to disclose to buyers about the qualities of various products, with the goal of maximizing utility. Our

results demonstrated the potential of this approach to address complex information design challenges, offering both theoretical advancements and practical solutions, and allowing us to extract useful insights from the solutions of the problem at hand.

The method's relevance extends beyond information design, as optimization over power diagrams naturally appears in other mechanism design problems and gerrymandering problems, the latter of which we demonstrated in a short excerpt in the last chapter of this thesis. This suggests broader applications across economics and political science.

Outlook and Extensions

As mentioned before, the method we developed in this chapter can easily be adapted to other problems in which an optimization over power diagrams is required. In explicit, one can think about the following variation of the original functional (5.2.1)

$$\max_{\mathbf{g}, \mathbf{X}} \sum_{i=1}^n m_i[\mathbf{X}, \mathbf{g}] \Phi(x_i).$$

Additionally, one could be interested in more general partitions, such as

$$L_i^{(p)}[\mathbf{X}, \mathbf{g}] := \{y \in \Omega : |y - x_i|^p - g_i \leq |y - x_j|^p - g_j \ \forall \ 1 \leq j \leq n\},$$

which could enjoy similar properties as power diagrams and thus be suitable to optimize over. Finally, one could think about addressing the issue of convergence of maximizers in the $\eta \rightarrow 0$ case, where η was the penalty parameter, possibly using a different penalization term. This would completely close the gap between power diagrams and finitely-supported Lipschitz-exposed points of F_ν .

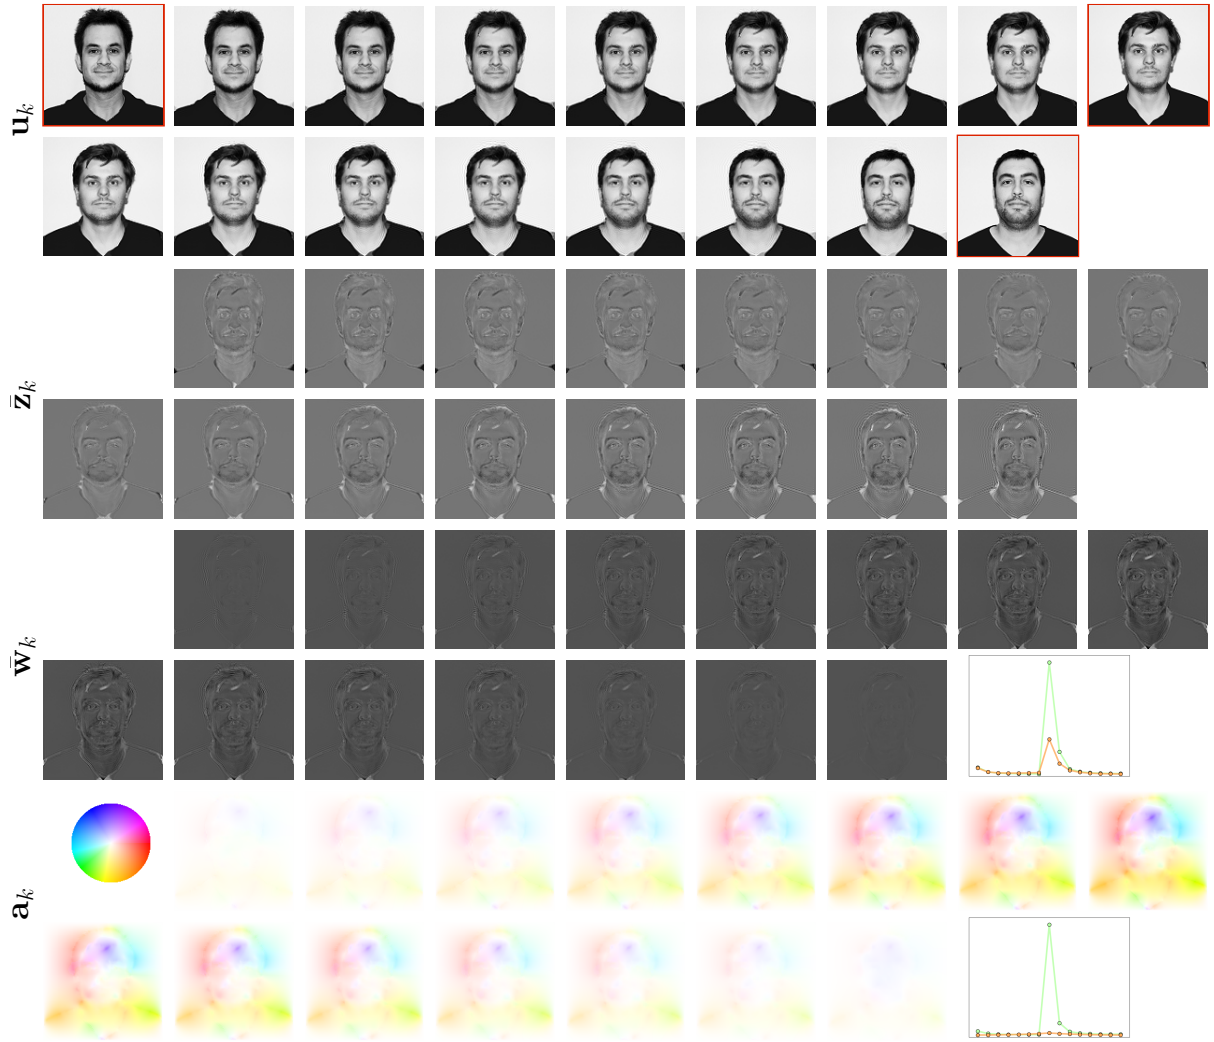


Figure 6.1: Time discrete spline with framed fixed images (first and second row), first order material derivative slack variable \bar{z} (third and fourth row), second order material derivative \bar{w}_k (fifth and sixth row) and color-coded acceleration field \mathbf{a}_k (seventh and eighth row)

Bibliography

- [ABSY23] I. Arieli, Y. Babichenko, R. Smorodinsky, and T. Yamashita. Optimal persuasion via bi-pooling. *Theoretical Economics*, 18(1):15–36, 2023.
- [ACB17] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [AF03] R. A. Adams and J. J. Fournier. *Sobolev spaces*. Elsevier, 2003.
- [AFC20] M. Alaya, R. Flamary, and N. Courty. Geometric optimal transport for 3d point cloud alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3464–3473, 2020.
- [AG13] L. Ambrosio and N. Gigli. A user’s guide to optimal transport. In *Modelling and optimisation of flows on networks*, pages 1–155. Springer, 2013.
- [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.
- [Alt16] H. W. Alt. Linear functional analysis. *An Application-oriented Introduction*, 2016.
- [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(fasc., fasc. 1):319–361, 1966.
- [ARSW15] G. M. Ahlfeldt, S. J. Redding, D. M. Sturm, and N. Wolf. The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189, 2015.
- [Aur87] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing*, 16(1):78–96, 1987.
- [B⁺12] G. Buttazzo et al. Optimal transportation and applications. *Lecture Notes in Mathematics*, 2071:1–25, 2012.
- [BB00] J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84:375–393, 01 2000.
- [BB07] L. Bottou and O. Bousquet. The tradeoffs of large scale learning. *Advances in neural information processing systems*, 20, 2007.
- [BCC⁺15] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- [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.
- [BFH⁺13] B. Berkels, P. T. Fletcher, B. Heeren, M. Rumpf, and B. Wirth. Discrete geodesic regression in shape space. In *Proc. of International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 8081 of *Lecture Notes in Computer Science*, pages 108–122. Springer, 2013.
- [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.

- [Big20] J. Bigot. Statistical data analysis in the wasserstein space. *ESAIM: Proceedings and Surveys*, 68:1–19, 2020.
- [BJL17] R. Bhatia, T. Jain, and Y. Lim. On the Bures-Wasserstein distance between positive definite matrices, 2017.
- [Bla53] D. Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272, 1953.
- [BMTY05] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.*, 61(2):139–157, 2005.
- [Bör15] T. Börgers. *An introduction to the theory of mechanism design*. Oxford University Press, USA, 2015.
- [Bra14] A. Braides. *Local minimization, variational evolution and Γ -convergence*, volume 2094. Springer, 2014.
- [Bre91] Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.
- [Bre11] H. Brezis. *Functional analysis, sobolev spaces and partial differential equations*, 2011.
- [BRPP15] N. Bonneel, J. Rabin, G. Peyré, and H. Pfister. Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1):22–45, 2015.
- [BSK⁺24] C. Bunne, G. Schiebinger, A. Krause, A. Regev, and M. Cuturi. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1):58, 2024.
- [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.
- [Bur69] D. Bures. An extension of Kakutani’s theorem on infinite product measures to the tensor product of semifinite w^* -algebras. *Transactions of the American Mathematical Society*, 135:199–212, 1969.
- [BV17] M. Bruveris and F.-X. Vialard. On completeness of groups of diffeomorphisms. *Journal of the European Mathematical Society*, 19(5):1507–1544, 2017.
- [BVBBE20] J. Backhoff-Veraguas, D. Bartl, M. Beiglböck, and M. Eder. Adapted wasserstein distances and stability in mathematical finance. *Finance and Stochastics*, 24(3):601–632, 2020.
- [BvTH16] C. Brandt, C. von Tycowicz, and K. Hildebrandt. Geometric flows of curves in shape space for processing motion of deformable objects. In *Computer Graphics Forum*, volume 35, pages 295–305. Wiley Online Library, 2016.
- [CCG18] Y. Chen, G. Conforti, and T. T. Georgiou. Measure-valued spline curves: An optimal transport viewpoint. *SIAM J. Numer. Anal.*, 50(6):5947–5968, 2018.
- [CCLG⁺21] S. Chewi, J. Clancy, T. Le Gouic, P. Rigollet, G. Stepaniants, and A. Stromme. Fast and smooth interpolation on Wasserstein space. In *International Conference on Artificial Intelligence and Statistics*, pages 3061–3069. PMLR, 2021.
- [CD14] M. Cuturi and A. Doucet. Fast computation of wasserstein barycenters. In *International conference on machine learning*, pages 685–693. PMLR, 2014.
- [CE03] G. Carlier and I. Ekeland. Optimal transportation and economic applications. *Economic Theory*, 21(4):957–978, 2003.
- [CFG⁺21] B. Charlier, J. Feydy, J. A. Glaunès, F.-D. Collin, and G. Durif. Kernel operations on the gpu, with autodiff, without memory overflows. *Journal of Machine Learning Research*, 22(74):1–6, 2021.
- [CFM64] P. Cartier, J. Fell, and P.-A. Meyer. Comparaison, des mesures portées par un ensemble convexe compact. *Bulletin de la Société Mathématique de France*, 92:435–445, 1964.

- [CFT18] N. Courty, R. Flamary, and D. Tuia. Optimal transport for structured data with application to image segmentation. *International Journal of Computer Vision*, 126(5):435–446, 2018.
- [CFTR17] N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy. Optimal transport for domain adaptation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 39, pages 1853–1865, 2017.
- [Cia88] P. G. Ciarlet. *Mathematical elasticity. Vol. I*, volume 20 of *Studies in Mathematics and its Applications*. North-Holland Publishing Co., Amsterdam, 1988.
- [CK18] Y. Chen and J. Karlsson. State tracking of linear ensembles via optimal mass transport. *IEEE Control Systems Letters*, 2(2):260–265, 2018.
- [Cla21] J. Clancy. *Interpolating Spline Curves of Measures*. PhD thesis, Massachusetts Institute of Technology, 2021.
- [CMPT⁺22] M. Cuturi, L. Meng-Papaxanthos, Y. Tian, C. Bunne, G. Davis, and O. Teboul. Optimal transport tools (ott): A jax toolbox for all things wasserstein. *arXiv preprint arXiv:2201.12324*, 2022.
- [CNWR24] S. Chewi, J. Niles-Weed, and P. Rigollet. Statistical optimal transport. *arXiv preprint arXiv:2407.18163*, 2024.
- [CPSV18] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation*, 87(314):2563–2609, 2018.
- [Cut13] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- [dB63] C. de Boor. Best approximation properties of spline functions of odd degree. *J. Math. Mech.*, 12:747–749, 1963.
- [de78] B. de. A practical guide to splines. 1978.
- [DGM98] P. Dupuis, U. Grenander, and M. I. Miller. Variational problems on flows of diffeomorphisms for image matching. *Quart. Appl. Math.*, 56(3):587–600, 1998.
- [DM19] P. Dworczak and G. Martini. The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019.
- [DPMTL21] G. De Palma, M. Marvian, D. Trevisan, and S. Lloyd. The quantum wasserstein distance of order 1. *IEEE Transactions on Information Theory*, 67(10):6627–6643, 2021.
- [DS20] W. Dessein and D. Szymanski. Optimal transport in dynamic trade models. *Journal of International Economics*, 127:103349, 2020.
- [EGH10] I. Ekeland, A. Galichon, and M. Henry. Optimal transportation and economic applications. *Economic Theory*, 42(2):333–364, 2010.
- [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(2):309–327, 2021.
- [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.
- [Eva18] L. Evans. *Measure theory and fine properties of functions*. Routledge, 2018.
- [FC⁺21] R. Flamary, N. Courty, et al. Pot: Python optimal transport library. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- [FH08] J. N. Friedman and R. T. Holden. Optimal gerrymandering: sometimes pack, but never crack. *American Economic Review*, 98(1):113–144, 2008.

- [Fio17] R. Fiorenza. *Hölder and locally Hölder Continuous Functions, and Open Sets of Class C^k , $C^{k,\lambda}$* . Birkhäuser, 2017.
- [FK16] P. J. Forrester and M. Kieburg. Relating the bures measure to the cauchy two-matrix model. *Communications in Mathematical Physics*, 342(1):151–187, 2016.
- [Fol99] G. B. Folland. *Real analysis: modern techniques and their applications*, volume 40. John Wiley & Sons, 1999.
- [FSV⁺19] J. Feydy, T. Séjourné, F.-X. Vialard, S.-i. Amari, A. Trounev, and G. Peyré. Interpolating distributions and point clouds with optimal transport. *arXiv preprint arXiv:1810.08278*, 2019.
- [G⁺19] I. Georgiev et al. Optimal transport for monte carlo rendering. *ACM Transactions on Graphics (SIGGRAPH)*, 38(4):123:1–123:13, 2019.
- [Gal16] A. Galichon. *Optimal Transport Methods in Economics*. Princeton University Press, 2016.
- [Gig12] N. Gigli. Second order analysis on $(p-2(m), w-2)$. *Memoirs of the American Mathematical Society*, 216:VII–+, 03 2012.
- [GPC18] A. Genevay, G. Peyré, and M. Cuturi. Learning generative models with sinkhorn divergences. *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1608–1617, 2018.
- [HLPR20] A. Houdard, A. Leclaire, N. Papadakis, and J. Rabin. Wasserstein generative models for patch-based texture synthesis. *CoRR*, abs/2007.03408, 2020.
- [HNY17] N. Ho, X. Nguyen, and M. Yurochkin. Multilevel clustering via wasserstein means. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 361–369, 2017.
- [HRS⁺16] B. Heeren, M. Rumpf, P. Schröder, M. Wardetzky, and B. Wirth. Splines in the space of shells. In *Computer Graphics Forum*, volume 35, pages 111–120. Wiley Online Library, 2016.
- [HRW18] B. Heeren, M. Rumpf, and B. Wirth. Variational time discretization of Riemannian splines. *IMA J. Numer. Anal.*, 39(1):61–104, 2018.
- [HS81] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981.
- [HvNV⁺16] T. Hytönen, J. van Neerven, M. Veraar, L. Weis, T. Hytönen, J. van Neerven, M. Veraar, and L. Weis. Bochner spaces. *Analysis in Banach Spaces: Volume I: Martingales and Littlewood-Paley Theory*, pages 1–66, 2016.
- [HZTA04] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent. Optimal mass transport for registration and warping. *International Journal of computer vision*, 60:225–240, 2004.
- [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.
- [Iva21] M. Ivanov. Optimal monotone signals in bayesian persuasion mechanisms. *Economic Theory*, 72(3):955–1000, 2021.
- [JKM⁺25] J. Justiniano, A. Kleiner, B. Moldovanu, M. Rumpf, and P. Strack. Entropy-regularized optimal transport in information design. In *Proceedings of the 26th ACM Conference on Economics and Computation*, page 741–760, 2025.
- [JM00] S. C. Joshi and M. I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Trans. Image Process.*, 9(8):1357–1370, 2000.

- [JMS99] P. Jehiel, B. Moldovanu, and E. Stacchetti. Multidimensional mechanism design for auctions with externalities. *Journal of economic theory*, 85(2):258–293, 1999.
- [JRE24] J. Justiniano, M. Rumpf, and M. Erbar. Approximation of splines in wasserstein spaces. *ESAIM: Control, Optimisation and Calculus of Variations*, 30:64, 2024.
- [JRR21] J. Justiniano, M. Rajković, and M. Rumpf. Splines for image metamorphosis. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 463–475. Springer, 2021.
- [JRR23] J. Justiniano, M. Rajković, and M. Rumpf. Consistent approximation of interpolating splines in image metamorphosis. *Journal of Mathematical Imaging and Vision*, 65(1):29–52, 2023.
- [Kan42] L. V. Kantorovich. On the translocation of masses. *C.R. (Doklady) Acad. Sci. URSS (N.S.)*, 37:199–201, 1942.
- [KG11] E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- [KG21] A. Karimi and T. T. Georgiou. Regression analysis of distributional data through multi-marginal optimal transport. *arXiv preprint arXiv:2106.15031*, 2021.
- [Kin14] D. P. Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [KJ58] V. L. Klee Jr. Extremal structure of convex sets. ii. *Mathematische Zeitschrift*, 69(1):90–104, 1958.
- [KMS21] A. Kleiner, B. Moldovanu, and P. Strack. Extreme points and majorization: Economic applications. *Econometrica*, 89(4):1557–1593, 2021.
- [KMSW24] A. Kleiner, B. Moldovanu, P. Strack, and M. Whitmeyer. Extreme points of multidimensional fusions. *discussion paper*, 2024.
- [Kol18] A. Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018.
- [KPT⁺17] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde. Optimal mass transport: Signal processing and machine-learning applications. *IEEE signal processing magazine*, 34(4):43–59, 2017.
- [KSKW15] M. J. Kusner, Y. Sun, N. I. Kolkin, and K. Q. Weinberger. From word embeddings to document distances. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 957–966, 2015.
- [Leo24] G. Leoni. *A first course in Sobolev spaces*, volume 181. American Mathematical Society, 2024.
- [Llo82] S. Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- [Lot06] J. Lott. Some geometric calculations on Wasserstein space. *arXiv preprint math/0612562*, 2006.
- [LW18] S. P. Lalley and E. G. Weyl. Quadratic voting: How mechanism design can radicalize democracy. In *AEA Papers and Proceedings*, volume 108, pages 33–37. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2018.
- [McC01] R. J. McCann. Polar factorization of maps on Riemannian manifolds. *Geometric & Functional Analysis GAFA*, 11(3):589–608, 2001.
- [Mém11] F. Mémoli. Gromov–wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11:417–487, 2011.
- [MF18] X. Ma and Z. Fang. Revenue-maximizing auctions with optimal transport theory. *Journal of Economic Theory*, 177:84–115, 2018.

- [MMP18] L. Malagò, L. Montrucchio, and G. Pistone. Wasserstein Riemannian geometry of Gaussian densities. *Information Geometry*, 1(2):137–179, 2018.
- [Mos69] U. Mosco. Convergence of convex sets and of solutions of variational inequalities. *Advances in Math.*, 3:510–585, 1969.
- [MRSS15] J. Maas, M. Rumpf, C. Schönlieb, and S. Simon. A generalized model for optimal transport of images including dissipation and density modulation. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1745–1769, 2015.
- [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.
- [Mye81] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [Nes05] Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103:127–152, 2005.
- [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 Cl. Sci. (3)*, 20:733–737, 1966.
- [NŠ91] J. Nečas and M. Šilhavý. Multipolar viscous fluids. *Quart. Appl. Math.*, 49(2):247–265, 1991.
- [Ott01] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. 2001.
- [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.
- [Pey09] G. Peyré. Texture synthesis with optimal transport. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 28(5):1–10, 2009.
- [Phe01] R. R. Phelps. *Lectures on Choquet’s theorem*. Springer, 2001.
- [PKD05] F. Pitié, A. C. Kokaram, and R. Doherty. N-dimensional probability density function transfer and its application to color transfer. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV)*, pages 1434–1439, 2005.
- [PMP14] N. Papadakis, E. Mémin, and J. Papadakis. Optimal transport for motion estimation. *Journal of Mathematical Imaging and Vision*, 48(1):163–178, 2014.
- [Pro56] Y. V. Prokhorov. Convergence of random processes and limit theorems in probability theory. *Theory of Probability & Its Applications*, 1(2):157–214, 1956.
- [PS16] T. Pock and S. Sabach. Inertial proximal alternating linearized minimization (iPALM) for nonconvex and nonsmooth problems. *SIAM J. Imaging Sci.*, 9(4):1756–1787, 2016.
- [Raj23] M. Rajkovic. *Geodesics, Splines, and Embeddings in Spaces of Images*. PhD thesis, 2023.
- [RDG14] J. Rabin, J. Delon, and Y. Gousseau. Adaptive color transfer with relaxed optimal transport. *SIAM Journal on Imaging Sciences*, 7(1):222–239, 2014.
- [RW15] M. Rumpf and B. Wirth. Variational time discretization of geodesic calculus. *IMA Journal of Numerical Analysis*, 35(3):1011–1046, 2015.
- [S⁺19] G. Schiebinger et al. Optimal transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- [San15] F. Santambrogio. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55(58-63):94, 2015.

- [Sch19] B. Schmitzer. Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM Journal on Scientific Computing*, 41(3):A1443–A1481, 2019.
- [Ser04] R. Serrano. The theory of implementation of social choice rules. *SIAM review*, 46(3):377–414, 2004.
- [SS07] M. Shaked and J. G. Shanthikumar. *Stochastic orders*. Springer, 2007.
- [Str65] V. Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.
- [SVN15] N. Singh, F.-X. Vialard, and M. Niethammer. Splines for diffeomorphisms. *Med. Image Anal.*, 25(1):56–71, 2015.
- [TV12] A. Trounev 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. Tahraoui and F.-X. Vialard. Minimizing acceleration on the group of diffeomorphisms and its relaxation. *ESAIM Control Optim. Calc. Var.*, 25, 2019.
- [ULVL16] D. Ulyanov, V. Lebedev, A. Vedaldi, and V. Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. *arXiv preprint arXiv:1603.03417*, 2016.
- [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.
- [Vil09] C. Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [Wat07] J. Watson. Contract, mechanism design, and technological detail. *Econometrica*, 75(1):55–81, 2007.
- [WOS⁺10] W. Wang, J. A. Ozolek, D. Slepčev, A. B. Lee, C. Chen, and G. K. Rohde. An optimal transportation approach for nuclear structure-based pathology. *IEEE transactions on medical imaging*, 30(3):621–631, 2010.
- [You10] L. Younes. *Shapes and diffeomorphisms*, volume 171 of *Applied Mathematical Sciences*. Springer-Verlag, Berlin, 2010.
- [Z⁺21] S. Zhang et al. Unifying bayesian inference and optimal transport for probabilistic modeling. In *International Conference on Machine Learning (ICML)*, pages 1235–1243, 2021.
- [ZN19] E. Zhang and L. Noakes. Riemannian cubics and elastica in the manifold $\text{spd}(n)$ of all $n \times n$ symmetric positive-definite matrices. *Journal of Geometric Mechanics*, 11:277–299, 01 2019.