Mathematisches Forschungsinstitut Oberwolfach

Report No. 38/2022

DOI: 10.4171/OWR/2022/38

Mathematical Imaging and Surface Processing

Organized by
Mirela Ben-Chen, Haifa
Antonin Chambolle, Paris
Martin Rumpf, Bonn
Peter Schröder, Pasadena

21 August – 27 August 2022

Abstract. This workshop was gathering applied mathematicians and computer scientists interested in image and geometry processing. These topics have developed tremendously in the past few years with the rise of artificial intelligence, parallel hardware and strong needs for real-world applications (3D scene reconstruction, architecture, medical imaging and data analysis, etc). These research fields are at the intersection of many mathematical disciplines, from geometry, calculus of variations and optimization to the analysis and numerical analysis of PDEs. The almost 50 participants to this workshop, including many young researchers, had many fruitful exchanges, interested in common issues and speaking a common language, yet often coming from different backgrounds and with different knowledge.

This volume collects the abstracts for all the presentations covering this wide spectrum of tools and application domains.

Mathematics Subject Classification (2020): 65Dxx, 68Uxx.

Introduction by the Organizers

The workshop *Mathematical Imaging and Surface Processing*, coming after a previous event organized in January 2016 (workshop # 1604), gathered nearly 50 participants, applied mathematicians, experts in discrete geometry, computer scientists, all interested in problems related to image processing (shape and information retrieval, inverse problems) and geometry processing (modeling, discrete surfaces and their analysis, computational issues...) The midsummer atmosphere was relaxing and encouraged fruitful discussions between researchers embracing

different aspects of the problems presented during the talks, and a strong participation of young researchers (advanced PhD students, post-docs) was noticed, which fostered many exchanges of great interest to all participants.

The talks addressed many issues related to shape and image processing, including mathematical and numerical analysis, computational issues, learning techniques. During the first day, we could for instance hear a seminar on a representation of the Willmore energy on triangular meshes (embedded in $\mathbb{R}^3$, consistent in the sense of Γ-convergence), as well as a nice companion talk, more computer science oriented, on discrete "distortion" energies for comparing shapes and their symmetric counterparts. In particular, the ARAP criterion ("as rigid as possible"), which is a sort of local distance to rotations, was introduced and this proved very useful for many subsequent talks relying on this family of energies and addressing similar problems. We also had the same day two talks on geometric learning and one impressive simplification of the deep learning (usually adversarial) based methods for image generation, based on score-matching. All in all, it appeared quickly that neural network based methods and similar learning techniques, not much discussed in 2016, have become important for the applications addressed in this conference, whether it be as a main tool, a subject of study, or more often as being of great help to improve part of the processes. They appeared in some talks as auto-encoders to reduce the dimension and learn low-dimensional manifolds from images of rotating/moving figures, as a tool for encoding implicit surfaces, or appropriately tweaked so as to represent and approximate one-Lipschitz functions (the latter being possibly an interesting approach for the former)...

A few very computational talks impressed the audience with movies of self-avoiding deformed shells (well introduced first theoretically), or efficient learning methods for shape matching, also, various convex representation to compute minimal surfaces or similar problems were presented.

PDEs, numerical analysis and variational calculus were also a strong common language for many participants. In addition to some already mentioned issues, a few talks were addressing polygonal meshes, from the point of view of designing finite elements on arbitrary shapes, meshing surfaces with hexagons or other shapes, or numerical criteria for building developpable surfaces, notably for applications to architecture and design. PDEs oriented talks were also a few, most manifold-oriented, whether it be the study of geodesic on shape manifolds, the optimality conditions for Fréchet means, or of the W_2-Wasserstein flow as a Riemannian flow for image quantization. The Wasserstein optimal transport problems appeared in two talks as a very interesting tool for multi-dimensional sorting, whether it be for quantile regression and learning or from a computational point of view.

It appears that this workshop was quite a success, as the apparent variety of the participants was nicely balanced by a common language and culture and shared interests, and all the talks were appreciated by most, as was obvious from the discussions afterwards and during the coffee breaks. The afternoon dedicated to young researchers (on the Tuesday) allowed them to feel more quickly part of the

group and be included in the exchanges, this interaction was very valuable for all participants.

It is hoped that a next similar conference will be organized in the future on these topics, with hopefully as many discussions and exchanges, fruitful for all, and to start with many young researchers.

Workshop: Mathematical Imaging and Surface Processing

Table of Contents

Abstracts

GLASS: Geometric Latent Augmentation for Learned Shape Spaces

Niloy J. Mitra
(joint work with Sanjeev Muralikrishnan and Luca Morreale)

Figure 1. Starting from just very few shapes, GLASS iteratively augments the collection by alternating between training a VAE, and exploring random perturbations in its low-dimensional latent space guided by a purely geometric deformation energy. Despite facial expressions not being perfectly locally rigid, we show above that GLASS generates plausible novel expressions on the COMA dataset from just three landmarks.

Building a parameterized shape space (e.g., for humans) is desirable for many downstream tasks. Simplistically, the problem requires the handling of three subproblems: establishing correspondence across the observed shapes, building an underlying deformation model (or energy), and discovering a lower-dimensional parameterization of the templatized shapes. Although the problem has been investigated by many hundreds of papers over multiple decades, the subproblems have largely been handled separately. We revisit the problem with deep learning tools to simultaneously solve all three tasks. In this talk, we report our current findings and show initial results that indicate that one may be able to tackle this problem in an unsupervised setup. Technically, we analyze the Hessian of the as-rigid-as-possible (ARAP) energy to adaptively sample from and project to the underlying (local) shape space and use the augmented dataset to train a variational autoencoder (VAE). We iterate the process of building latent spaces of VAE and augmenting the associated dataset to progressively reveal a richer and more expressive generative space for creating geometrically and semantically valid samples. We demonstrate applications in surface map extraction, data augmentation, and model building.

References

[1] S. Muralikrishnan, S. Chaudhuri, N. Aigerman, V. Kim, M. Fisher, N. Mitra, *GLASS: Geometric Latent Augmentation for Shape Spaces*, Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), (2022).

[2] L. Morreale, S. Chaudhuri, N. Aigerman, V. Kim, N. Mitra, *Neural Surface Maps*, Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), (2021).

Consistent discrete approximations of the Willmore energy

Peter Gladbach
(joint work with Heiner Olbermann)

We consider two discretizations of the Willmore energy for C^2 compact surfaces $M \subset \mathbb{R}^3$, which can be written as

$$W(M) := \int_M |Dn(x)|^2 \, dA, \tag{1}$$

where $n : M \to S^2$ is the (locally oriented) *unit normal field* and its derivative $Dn(x) \in \mathrm{Lin}(T_xM; T_xM)$ is the *shape operator*.

The discrete approximations of W act on triangular surfaces instead of C^2 surfaces, where the true Willmore energy is typically infinite. A *triangular surface* is a finite collection of triangles $\mathcal{T} := \{K_1, \dots, K_N\}$, $K_i = \mathrm{conv}(x_i, y_i, z_i) \subset \mathbb{R}^3$, with each triangle sharing precisely one edge with each of its three neighbors.

The first discrete approximation we consider is the discrete bending energy from [1]

$$E(\mathcal{T}) := \frac{1}{2} \sum_{K \in \mathcal{T}} \sum_{L \sim K} \frac{l_{KL}}{d_{KL}} |n(K) - n(L)|^2. \tag{2}$$

Here $l_{KL}; = \mathrm{Len}(K \cap L)$ is the length of the common edge of two neigbhoring triangles, and $d_{KL} := |q_K - q_L|$ is the euclidean distance of the circumcenters, whereas $n(K), n(L) \in S^2$ are the like-oriented normals of K and L respectively.

In [3] we show the following theorem:

Theorem 1.

(i) *Whenever M is a $W^{2,2} \cap W^{1,\infty}$ graph and whenever $\mathcal{T}_k$ is a sequence of triangular surfaces that are (a) uniformly Lipschitz graphs (b) Delaunay (c) uniformly inner regular (d) $\max_{K \in \mathcal{T}_k} \mathrm{diam}(K) \to 0$ (e) converge in Hausdorff distance to M, then*

$$W(M) \leq \liminf_{k \to \infty} E(\mathcal{T}_k).$$

(ii) *Whenever M is a $W^{2,2} \cap W^{1,\infty}$ graph, there is a sequence $\mathcal{T}_k$ of triangular surfaces satisfying (a),(b),(c),(d),(e) such that*

$$W(M) = \lim_{k \to \infty} E(\mathcal{T}_k).$$

If M is C^3, we may choose all vertices on M.

We note here that condition (a) is a technical assumption that guarantees convergence of the measures in normal, whereas condition (b) is necessary to guarantee the lower bound (i). The construction in (ii) is very delicate, since the only triangular surfaces guaranteeing convergence in energy have edges almost parallel to

curvature nets, i.e. almost right angles. However, right triangles barely satisfy the Delaunay condition and their circumcenters lie on the hypotenuse, where $d_{KL} = 0$.

In an upcoming article, we show that a different discrete energy, first seen in [2], has better convergence properties. The energy uses so called *edge directors*, which are vector fields $n : \mathcal{E}(\mathcal{T}) \to S^2$ defined on the edges of a triangular surface $\mathcal{T}$ and satisfying $n(K \cap L) \cdot \tau(K \cap L) = 0$, where $\tau(K \cap L) \in \mathbb{R}^3$ is the tangent vector. Given a triangular surface and an edge director, we define the energy

$$F(\mathcal{T}, n) := \sum_{K \in \mathcal{T}} \mathrm{Area}(K) |Dn_K|^2, \tag{3}$$

where $n_K : K \to \mathbb{R}^3$ is the affine interpolation of the edge director field sampled at the edge midpoints. We show the following theorem:

Theorem 2.

(i) *M is a $W^{2,2} \cap W^{1,\infty}$ graph and whenever $\mathcal{T}_k$ is a sequence of triangular surfaces satisfying (a), (c), (d), (e), n_k a sequence of edge directors for $\mathcal{T}_k$, then*

$$W(M) \leq \liminf_{k \to \infty} F(\mathcal{T}_k, n_k).$$

(ii) *Whenever M is a $W^{2,2} \cap W^{1,\infty}$ graph, there is a sequence $\mathcal{T}_k$ of triangular surfaces satisfying (a),(b),(c),(d),(e) and a sequence of edge directors n_k for $\mathcal{T}_k$ such that*

$$W(M) = \lim_{k \to \infty} F(\mathcal{T}_k, n_k).$$

If M is C^3, any triangular surface $\mathcal{T}$ with vertices in M satisfying (c),(d) has an edge director n such that

$$|W(M) - F(\mathcal{T}, n)| \leq C \max_{K \in \mathcal{T}} \mathrm{diam}|K|.$$

References

[1] E. Grinspun, A. Hirani M. Desbrun, P. Schröder, *Discrete shells*, Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation (2003), 62–67.

[2] M. Wardetzky, M. Bergou, D. Harmon, D. Zorin, E. Grinspun, *Discrete quadratic curvature energies*, Computer Aided Geometric Design **24** (2007), 499–518.

[3] P. Gladbach, H. Olbermann, *Approximation of the Willmore energy by a discrete geometry model*, Advances in Calculus of Variations (2021).

Score based diffusion models for conditional image generation

Carola-Bibiane Schönlieb

(joint work with Georgios Batzolis, Jan Stanczuk, Christian Etmann)

Score-based diffusion models have emerged as one of the most promising frameworks for deep generative modelling. In [1] we conduct a systematic comparison and theoretical analysis of different approaches to learning conditional probability distributions with score-based diffusion models. In particular, we prove results which provide a theoretical justification for one of the most successful estimators of

the conditional score. Moreover, we introduce a multi-speed diffusion framework, which leads to a new estimator for the conditional score, performing on par with previous state-of-the-art approaches. Our theoretical and experimental findings are accompanied by an open source library `MSDiff` which allows for application and further research of multi-speed diffusion models.

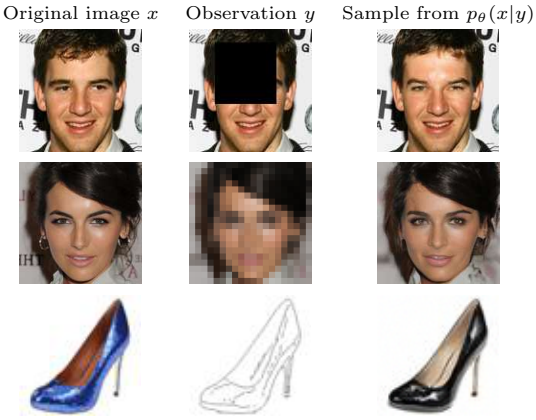

FIGURE 1. Results from our conditional multi-speed diffusive estimator.

Generative models based on score based diffusion modelling. In a recent work [2] score-based [3, 4] and diffusion-based [5, 6] generative models have been unified into a single continuous-time score-based framework with diffusion driven by stochastic differential equations. This continuous-time score-based diffusion technique relies on Anderson's Theorem [7], which states that (under certain assumptions on $\mu : \mathbb{R}^{n_x} \times \mathbb{R} \to \mathbb{R}^{n_x}$ and $\sigma : \mathbb{R} \to \mathbb{R}$) a forward diffusion process

$$dx = \mu(x,t)dt + \sigma(t)dw \tag{1}$$

has a reverse diffusion process governed by the following SDE:

$$dx = [\mu(x,t) - \sigma(t)^2 \nabla_x \ln p_{X_t}(x)]dt + \sigma(t)d\bar{w}, \tag{2}$$

where $\bar{w}$ is a standard Wiener process in reverse time.

The forward diffusion process transforms the *target distribution* $p(x_0)$ to a *diffused distribution* $p(x_T)$. By appropriately selecting the drift and the diffusion coefficients of the forward SDE, we can make sure that after sufficiently long time T, the diffused distribution $p(x_T)$ approximates a simple distribution, such as $\mathcal{N}(0, I)$. We refer to this simple distribution as the *prior distribution*, denoted by π.

If we have access to the score of the marginal distribution, $\nabla_{x_t} \ln p(x_t)$, for all t, we can derive the reverse diffusion process and simulate it to map p_T to p_0. In practice, we approximate the score of the time-dependent distribution by a

neural network $s_\theta(x_t, t) \approx \nabla_{x_t} \ln p(x_t)$ and map the prior distribution $\pi \approx p(x_T)$ to $p_\theta(x) \approx p(x_0)$ by solving the reverse-time SDE from time T to time 0. One can integrate the reverse SDE using standard numerical SDE solvers such Euler-Maruyama or other discretisation strategies. The authors propose to couple the standard integration step with a fixed number of Langevin MCMC steps to leverage the knowledge of the score of the distribution at each intermediate timestep. The MCMC correction step improves sampling; the combined algorithm is known as a predictor-corrector scheme.

In order to fit a neural network model $s_\theta(x_t, t)$ to approximate the score $\nabla_{x_t} \ln p(x_t)$, we minimize the weighted Fisher's divergence

$$\mathcal{L}_{SM}(\theta) := \frac{1}{2} \mathbb{E}_{\substack{t \sim U(0,T) \\ x_t \sim p(x_t)}} [\lambda(t) \, \| \nabla_{x_t} \ln p(x_t) - s_\theta(x_t, t) \|_2^2] \tag{3}$$

where $\lambda : [0, T] \to \mathbb{R}_+$ is a positive weighting function. We refer to [2] for details. Conditional score based generation. The continuous score-matching framework can be extended to conditional generation, as shown in [2]. Suppose we are interested in $p(x|y)$, where x is a *target image* and y is a *condition image*. Again, we use the forward diffusion process (Equation 1) to obtain a family of diffused distributions $p(x_t|y)$ and apply Anderson's Theorem to derive the *conditional reverse-time SDE*

$$dx = [\mu(x, t) - \sigma(t)^2 \nabla_x \ln p_{X_t}(x|y)]dt + \sigma(t) d\tilde{w}. \tag{4}$$

Now we need to learn the score $\nabla_{x_t} \ln p(x_t|y)$ in order to be able to sample from $p(x|y)$ using reverse-time diffusion.

In [1] we propose a new approach to estimate the conditional score, called multi-speed diffusive estimator. Here, x_t and y_t diffuse according to SDEs with the same drift but different diffusion rates,

$$\begin{aligned} dx &= \mu(x, t)dt + \sigma^x(t)dw \\ dy &= \mu(y, t)dt + \sigma^y(t)dw. \end{aligned}$$

In order to learn $p(x_t|y_t)$, observe that

$$\nabla_{x_t} \ln p(x_t|y_t) = \nabla_{x_t} \ln p(x_t, y_t) = \nabla_{z_t} \ln p(z_t)[: n_x],$$

where $z_t := (x_t, y_t)$ and n_x is the dimensionality of x. Therefore we can learn the (unconditional) score of the joint distribution $p(x_t, y_t)$ using the denoising score matching objective just like as in the unconditional case but with a positive definite weighting matrix $\Lambda(t) : \mathbb{R} \to \mathbb{R}^{(n_x+n_y)\times(n_x+n_y)}$. Hence, the new training objective becomes

$$\frac{1}{2} \mathbb{E}_{\substack{t \sim U(0,T) \\ z_0 \sim p_0(z_0) \\ z_t \sim p(z_t|z_0)}} [v^T \Lambda(t) v], \tag{5}$$

where $v = \nabla_{z_t} \ln p(z_t|z_0) - s_\theta(z_t, t)$, $z_t = (x_t, y_t)$. We can then extract our approximation for the conditional score $\nabla_{x_t} \ln p(x_t|y_t)$ by simply taking the first n_x components of $s_\theta(x_t, y_t, t)$. Moreover, we can prove the following result which provides a guidance on how to choose Λ which ensures that the objective of the

score-based model upper-bounds the negative log-likelihood of the data, thus enabling approximate maximum likelihood training of score-based diffusion models.

Theorem 1. *Let $\mathcal{L}(\theta)$ be the CMDE training objective (Equation 5) with the following weighting:*

$$\Lambda_{i,j}^{MLE}(t) = \begin{cases} \sigma^x(t)^2, \text{ if } i = j, \ i \leq n_x \\ \sigma^y(t)^2, \text{ if } i = j, \ n_x < i \leq n_y \\ 0, \text{ otherwise} \end{cases}$$

Then the joint negative log-likelihood is upper bounded (up to a constant in θ) by the training objective of CMDE

$$-\mathbb{E}_{(x,y)\sim p(x,y)}[\ln p_\theta(x,y)] \leq \mathcal{L}(\theta) + C.$$

For more details on these results and a wider discussion on theoretical guarantees for conditional score based diffusion models in the literature, as well as more results from numerical experiments see our preprint [1].

References

[1] G. Batzolis, J. Stanczuk, C.-B. Schönlieb, C. Etmann, *Conditional image generation with score-based diffusion models*, arXiv:2111.13606.

[2] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, B. Poole, *Score-Based Generative Modeling through Stochastic Differential Equations*, International Conference on Learning Representations (2021).

[3] A. Hyvärinen, *Estimation of Non-Normalized Statistical Models by Score Matching*, Journal of Machine Learning Research **6**(24) (2005), 695–709.

[4] Y. Song, S. Ermon, *Generative Modeling by Estimating Gradients of the Data Distribution*, Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (2019).

[5] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, S. Ganguli, *Deep Unsupervised Learning using Nonequilibrium Thermodynamics*, Proceedings of the 32nd International Conference on Machine Learning (2015).

[6] J. Ho, A. Jain, P. Abbeel, *Denoising Diffusion Probabilistic Models*, Advances in Neural Information Processing Systems 33 (2020).

[7] B. D.O. Anderson, *Reverse-time diffusion equation models*, Stochastic Processes and their Applications **12**(3) (1982), 313–326.

DeltaConv: Anisotropic Convolutions for Learning from Point Clouds

Klaus Hildebrandt

(joint work with Ruben Wiersma, Ahmad Nasikun, and Elmar Eisemann)

Due to advances in 3D acquisition and modeling, geometric 3D-data and 3D-data collections are becoming more and more commonplace and are used in a number of application areas. This has resulted in a need for computational methods to make use of these data. In image processing and analysis, learning from image collections using convolutional neural networks has proven to be very successful and has led to many breakthroughs. This motivates the generalization of these methods to geometric data such as point clouds or triangle meshes.

The main differences between images and discrete surfaces are that images are represented by a regular grid in a flat Euclidean domain whereas discrete surface representations are irregularly sampled from curved domains. To build convolutional neural networks for geometric data, the main operation that needs to be generalized are the convolutions. Various generalized convolutions for point clouds and other surface representations have been proposed [1]. One difficulty in generalizing the convolutions is that, unlike in the Euclidean plane, there is no global coordinate system on surfaces. This makes the construction of convolution layers for surfaces that consistently extract and process directional information a challenge.

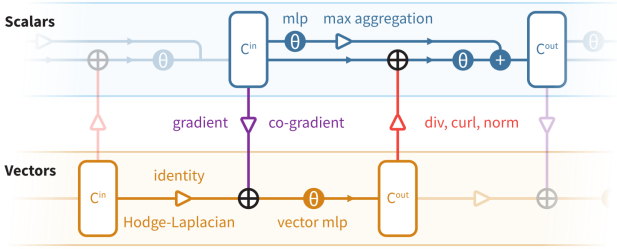

FIGURE 1. Schematic overview of a DeltaConv block.

We present *DeltaConv*, a novel convolutional layer for learning from point clouds. DeltaConv builds convolutions as compositions and linear combinations of a selected set of geometric differential operators and pointwise nonlinearities. The parameters of these combinations are optimized during training. To facilitate the extraction and processing of directional information, the input and output of a DeltaConv layer consist of scalar-valued and (tangential) vector-valued features. Within a DeltaConv block, scalar-valued features are aggregated locally using the Laplace–Beltrami operator or maximum aggregation and vector-valued features are aggregated using the Hodge–Laplace operator. In addition, the gradient and co-gradient are used to map from scalar-valued features to vector-valued features, and divergence, curl, and the vector norm are used to map vector-valued features to scalar-valued features. A schematic overview of a DeltaConv block is shown in Figure 1.

By design, DeltaConv layers can construct anisotropic differential operators and explicitly represent directional information in the vector features. Furthermore, DeltaConv layers benefit from the properties of the geometric differential operators. For example, DeltaConv layers are invariant under isometric deformations of a surface and robust to different samplings of the same surface.

For an extended description of DeltaConv, a discussion of its properties, and experimental results, we refer to [2].

References

[1] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, Pierre Vandergheynst *Geometric Deep Learning: Going beyond Euclidean data*, IEEE Signal Process. Mag. **34** (2017), 18–42.

[2] Ruben Wiersma, Ahmad Nasikun, Elmar Eisemann, Klaus Hildebrandt, *DeltaConv: anisotropic operators for geometric deep learning on point clouds*, ACM Trans. Graph. **41(4)** (2022).

Grid-Free Monte Carlo Methods for Partial Differential Equations

Keenan Crane

Partial differential equations (PDEs) provide fundamental models for phenomena throughout science and engineering. Yet even after decades of research, traditional methods for solving PDEs such as *finite element methods (FEM)* still struggle to capture the immense complexity encountered in nature. Directly simulating the inner workings of a human cell, or accurately predicting energy usage from a detailed building model (including insulation, ventilation, plumbing, etc.), is still in the realm of science fiction. A perennial challenge is the need to spatially discretize the domain, i.e., dice up space into a grid or mesh that can be used to approximate the solution in a finite basis. This process is expensive, error prone, and is often the weakest link in terms of building reliable systems on top of PDEs. A major revolution in computer graphics was to replace FEM with stochastic Monte Carlo methods-enabling photorealistic image generation for scenes of real-world complexity. This talk likewise explores a radical departure from traditional PDE solvers, by connecting deep insights from photorealistic Monte Carlo rendering to a broad class of PDEs fundamental in scientific computing.

Walk on Spheres (WoS). The basic principle is to solve PDEs by simulating a large collection of random walks. For instance, *Kakutani's principle* asserts that the steady-state distribution of heat in a domain Ω is equal to the expected boundary value "seen" by a Brownian random walk. But how do you simulate random walks? Rather than dice up space and take random walks on a discrete graph, Muller [1] had a brilliant idea: since a walk starting at $x \in \Omega$ has equal probability of exiting through any point on a sphere around x, just uniformly sample a point on the largest empty sphere $S(x) \subset \Omega$ around x, and repeat until you reach the boundary. This *walk on spheres (WoS)* algorithm avoids all the major headaches of traditional PDE solvers: no mesh needs to be generated, there is no spatial discretization error, and random walks are trivial to distribute in parallel. However, Muller's approach has received almost no attention over the past 60 years, lagging far behind the capabilities of modern methods like FEM in terms of the class of PDEs it can handle.

Next Generation WoS Methods. This talk explores strategies for expanding the scope of problems that can be solved via grid-free Monte Carlo methods. A major part of this effort has been "modernizing" WoS by connecting

it to tools developed for Monte Carlo rendering [2]—including both mathematical techniques (e.g., variance reduction strategies), as well as system engineering (e.g., high-performance geometric queries). A more fundamental question is how to construct Monte Carlo estimators for larger classes of PDEs. For instance, by making connections to volume rendering, we have expanded this set to include all linear-elliptic PDEs that fall under the umbrella of the Feynman-Kac formula, including those with spatially-varying coefficients [3]. These more general PDEs are essential for modeling complex heterogeneous behavior found in real systems (such as spatially-varying material density or conductivity). More recently, we have established a procedure for simulating reflecting Brownian motion on arbitrary, non-convex domains, which extends WoS to derivative (Neumann) boundary conditions critical for most scientific problems. Note that these components represent only the most basic elements of a PDE solver: just as FEM has taken over 80 years to reach its current level of maturity, WoS offers rich new questions (and opportunities) for many years to come.

References

[1] M. Muller. Some Continuous Monte Carlo Methods for the Dirichlet Problem. *Ann. Math. Statist.*, 27(3), 09 1956.

[2] Rohan Sawhney and Keenan Crane. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-Based Methods on Volumetric Domains. *ACM Trans. Graph.*, 39(4), 2020.

[3] Rohan Sawhney, Dario Seyb, Wojciech Jarosz, and Keenan Crane. Grid-Free Monte Carlo for PDEs with Spatially Varying Coefficients. *ACM Trans. Graph.*, 2022.

Optical Flow On Manifolds and For Elastography

Otmar Scherzer

(joint work with Clemens Kirisits, Ekaterina Sherina)

The motivating example is the embryogenesis of a zebrafish, which is observed via volumetric microscopy images. Shortly after fertilization cells move on the yolk and then they start to split and merge forming different organs, and eventually the whole fish; see Figure 1.

We consider three topics:

(1) We generalize the concept of **optical flow** to a dynamic non-Euclidean setting in order to analyze the image data (for instance, the cell movement) on the manifold (the yolk), thereby assuming that the cells move on the surface of the yolk, which may itself change its shape during time. The motion of the cells and the yolk can **only** be separated if **biological** and **physical priors** are taken into account, which is therefore advocated. Let us assume that the yolk can be segmented, then the estimation of the yolk's movement can be reduced to a shape registration problem (see [1]).

(2) Our method of choice for solving **shape registration problems** is minimizing elastic thin shell energies, where the shapes are represented as zero-level-sets [2, 3]. The outcome of the algorithm is depicted in Figure 2.

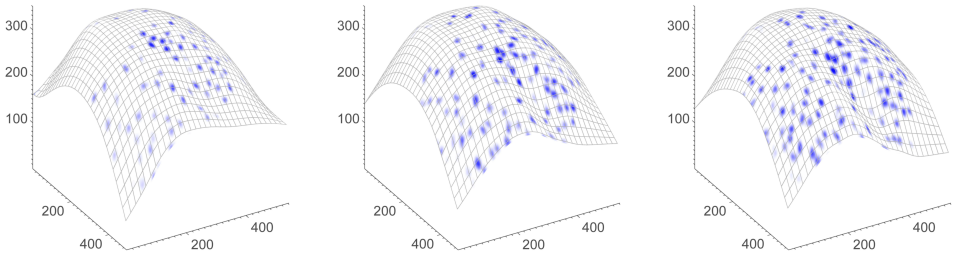

FIGURE 1. Images from [1]: Sequence of embryonic zebrafish images. The curved mesh represents a section of the yolk's surface and the blue dots are marked cells. Original microscopic data by Robin Kimmel and Dirk Meyer from the University of Innsbruck, Austria.

FIGURE 2. Images from [3]: From left to right: Textured hand shape M_1, resulting deformed shape, comparison of target and obtained shapes.

(3) **Computational Elastography** requires extraction of mechanical parameters from tomographic imaging data during compression of a probe or body. Typically, the extraction process splits into two steps, consisting of the strain reconstruction first and then subsequently estimating the mechanical properties, like the Young's modulus. The main observation is that strain rate estimated via standard Computer Vision tools, like the optical flow, is too inaccurate, and therefore needs to be adapted to the elastographic experiments. In other words, we require again **prior information on physical quantities** (like a background material) to extract accurate strain estimates. In particular, novel computational methods have to be developed, which combine tracking and flow estimation (see

[4, 5]). The subsequent estimation of mechanical properties is for instance discussed in [6].

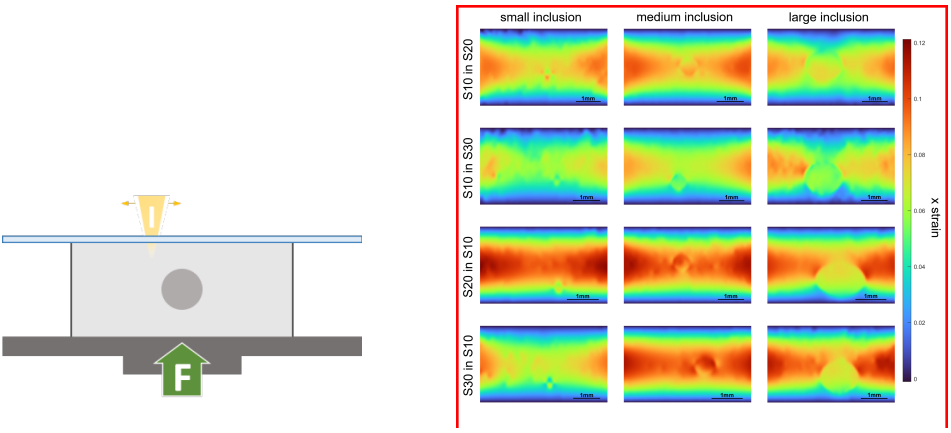

FIGURE 3. Images from [7]: Schematic depiction of an elastographic experiment. Typical strain images reconstructed from imaging data.

REFERENCES

[1] C. Kirisits, L.F. Lang, and O. Scherzer, *Optical flow on evolving surfaces with an application to the analysis of 4D microscopy data*, In: A. Kuijper, K. Bredies, T. Pock, H. Bischof (eds) Scale Space and Variational Methods in Computer Vision. SSVM 2013. Lecture Notes in Computer Science **7893** (2013), 246–257.

[2] J.A. Iglesias, B. Berkels, M. Rumpf, and O. Scherzer, *A thin shell approach to the registration of implicit surfaces*, In: M. Bronstein, J. Favre, K. Hormann (eds) Vision, Modeling & Visualization (2013), 89–96.

[3] J.A. Iglesias, M. Rumpf, and O. Scherzer, *Shape-aware matching of implicit surfaces based on thin shell energies*, Found. Comput. Math. **18(4)** (2018), 891–927.

[4] E. Sherina, L. Krainz, S. Hubmer, W. Drexler, and O. Scherzer, *Displacement field estimation from OCT images utilizing speckle information with applications in quantitative elastography*, Inverse Probl. **36(12)** (2020), 124003.

[5] E. Sherina, L. Krainz, S. Hubmer, W. Drexler, and O. Scherzer, *Challenges for optical flow estimates in elastography*, In: A. Elmoataz, J. Fadili, Y. Quéau, J. Rabin, L. Simon (eds) Scale Space and Variational Methods in Computer Vision. SSVM 2021. Lecture Notes in Computer Science **12679** (2021), 128–139.

[6] S. Hubmer, E. Sherina, A. Neubauer, and O. Scherzer, *Lamé parameter estimation from static displacement field measurements in the framework of nonlinear inverse problems*, SIAM J. Imaging Sciences **11(2)** (2018), 1268–1293.

[7] L. Krainz, E. Sherina, S. Hubmer, M. Liu, W. Drexler and O. Scherzer, *Quantitative Optical Coherence Elastography: A novel Intensity-based Inversion Method versus Strain-based Reconstructions*, IEEE J. Sel. Top. Quantum Electron (2022), submitted.

Designing Symmetric Distortion Energies

Justin Solomon
(joint work with S. Mazdak Abulnaga, Oded Stein, Polina Golland)

Many algorithms in geometry processing aim to find a low-distortion map from one shape onto another, often by optimizing an energy measuring how well the map preserves local geometry. In applications where there is no clear "source" and "target" shape, we might reasonably expect a mapping algorithm to be *symmetric*, i.e., to have the property that the algorithm outputs the same map (up to inversion) whether we map from domain M_1 to domain M_2 or from M_2 to M_1. Motivated by some recent research in our group at MIT, my talk explores the counterintuitive consequences the symmetry assumption has on the behavior and design of distortion energies, opening new questions for numerical shape analysis.

In more detail, suppose $M_1, M_2 \subseteq \mathbb{R}^3$ are compact volumes bounded by smooth surfaces. For applications like texture/annotation transfer and deformation, many algorithms in geometry processing attempt to compute a (nearly) bijective map $\phi : M_1 \to M_2$ satisfying various desirable properties, like smoothness and feature preservation. A commonly-used optimization objective evaluates the *distortion* of a candidate map ϕ via the energy

$$E_f[\phi] := \int_{M_1} f(J_\phi(x))\, dV(x). \tag{1}$$

Here, f (sometimes called a *constitutive model*) is some function of the signed singular values of the Jacobian $J_\phi(x) \in \mathbb{R}^{3\times 3}$. For example, the choice $f_{\text{Dirichlet}}(J) = \|J\|^2_{\text{Fro}}$ corresponds to the Dirichlet energy of ϕ, while $f_{\text{ARAP}}(J) = \min_{R\in\text{SO}(3)} \|J - R\|^2_{\text{Fro}}$ gives the popular as-rigid-as-possible deformation energy.

Many applications have no clear choice of "source" and "target" shape for the computed map, i.e., the labeling of M_1 vs. M_2 is irrelevant to the mapping problem. This situation occurs, for example, when mapping between anatomical scans of two subjects in a medical study. In this case, to avoid bias we might reasonably expect that our algorithm produces the same map—up to inversion—regardless of the labeling of M_1 and M_2. This concept of symmetry leads us to consider energy densities $f(\cdot)$ leading to the property that $E_f[\phi] = E_f[\phi^{-1}]$ for invertible maps ϕ.

We can transform any constitutive model $f(\cdot)$ into a symmetric energy $f_{\text{Sym}}(\cdot)$ by summing $E_f[\phi] + E_f[\phi^{-1}]$, equivalent after integral change of variables to substituting the following constitutive model into (1):

$$f_{\text{Sym}}(J) := f(J) + |\det J| f(J^{-1}). \tag{2}$$

While intuitively one might predict that $f_{\text{Sym}}(\cdot)$ would promote bijectivity or even elasticity of ϕ, we find that symmetrizing many common distortion energies does *not* have that effect. For example, the symmetrized counterpart of the Dirichlet energy $f_{\text{Dirichlet}}(J)$ is still minimized at $J = 0$, preferring maps that collapse rather than bijective maps. After considering a number of alternatives, we find that symmetrizing $f_{\text{ARAP}}(J)$ leads to a distortion energy with many desirable

properties for the volume correspondence problem. For example, the symmetrized counterpart is minimized when maps are isometric and blows up for singular maps.

Beyond exploring the mathematical consequences of designing symmetric distortion energies, we show that our calculations have direct bearing on performance of algorithms for volumetric correspondence. We refer readers to [1] for details and experiments.

References

[1] S.M. Abulnaga, O. Stein, P. Golland, and J. Solomon, *Symmetric volume maps*, ArXiv 2202.02568 (2022).

A multi-metric perspective on non-rigid surfaces with learning and spectra twists

Ron Kimmel
(joint work with David BenSaid, Amit Bracha)

In the talk I have shown how utilizing the spectra and corresponding eigen-spaces extracted from various LBOs (w.r.t. different metrics) provide better ways to treat non-rigid surface matching, reconstruction, identification, and processing. The focus of my lecture was our recent efforts in this domain jointly with David BenSaid and Amit Bracha. We started with identifying scale invariant arc length of planar curves by modulating the Euclidean arc length by the magnitude of the Euclidean curvature. This pseudo-metric gives rise to a corresponding second order differential operator, which eigenfunctions are the Fourier functions w.r.t. that measure. At the other end, the regular Fourier w.r.t. the second derivative w.r.t. the regular arc length can also be computed. Based on O. Halimi's self-functional-maps result, we argued that the inner product of these two spaces provides an algebraic descriptor of planar curves.

In a similar way, we extend the discussion to surfaces, with corresponding regular and scale invariant laplacians. Then, we used these measures as natural inputs to Litani's functional-map-network and Halimi's unsupervised version of that, ending up with Bracha's version which optimizes for the unitary relation between the scale invariant Eigen-spaces of two different surfaces we would like to compare. Finally, we also argued that the corresponding spectra w.r.t. the two laplacians could serve as measures for comparing partial shapes to the whole following Rampini et al. efforts. That last result was lead by David BenSaid.

Wasserstein Steepest Descent Flows of Discrepancies with Distance Kernels

Gabriele Steidl

(joint work with Robert Beinert, Manuel Gräf, Johannes Hertich)

Let $\big(\mathcal{P}_2(\mathbb{R}^d), W_2\big)$ denote the separable, complete metric space of all Borel probability measures with finite second moment with the Wasserstein distance $W_2 : \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^d) \to [0,\infty)$ defined by

$$W_2^2(\mu,\nu) := \min_{\boldsymbol{\pi}\in\Gamma(\mu,\nu)} \int_{\mathbb{R}^d\times\mathbb{R}^d} \|x-y\|_2^2 \,\mathrm{d}\boldsymbol{\pi}(x,y), \qquad \mu,\nu\in\mathcal{P}_2(\mathbb{R}^d),$$

where $\Gamma(\mu,\nu)$ denotes the probability measures on $\mathbb{R}^d\times\mathbb{R}^d$ with marginals μ and ν. A curve $\gamma : I \to \mathcal{P}_2(\mathbb{R}^d)$ on the open interval $I\subset\mathbb{R}$ is called absolutely continuous if there exists a function $m\in L^1(I)$ such that

$$W_2(\gamma(s),\gamma(t)) \le \int_s^t m(s)\,\mathrm{d}s, \qquad s,t\in I.$$

Absolutely continuous curves are closely related to the continuity equation [1, Thm 8.3.1]. More precisely, for an absolutely continuous curve $\gamma : I\to\mathcal{P}_2(\mathbb{R}^d)$ with metric derivative $|\gamma'|\in L_1(I)$ there exists a Borel velocity field $v:(t,x)\mapsto v_t(x)$ such that $v_t\in L_2(\gamma(t),\mathbb{R}^d)$ and

$$\|v_t\|_{L_2(\gamma(t),\mathbb{R}^d)} = |\gamma'|(t) \quad \text{a.e. } t\in I$$

and the continuity equation

$$\partial_t\gamma(t) + \nabla_x\cdot(\gamma(t)v_t) = 0 \quad \text{in } I\times\mathbb{R}^d$$

holds in the distributional sense and conversely. We are interested in *forward schemes* for the computation of Wasserstein gradient flows.

For a given a function $\mathcal{F}:\mathcal{P}_2(\mathbb{R}^d)\to(-\infty,+\infty]$, an absolutely continuous curve $\gamma:(0,+\infty)\to\mathcal{P}_2(\mathbb{R}^d)$ is called the *Wasserstein gradient flow* if its Borel velocity field $v_t\in \mathrm{T}_{\gamma(t)}\mathcal{P}_2(\mathbb{R}^d)$ belongs to the Fréchet subdifferential of $\mathcal{F}$ at $\gamma(t)$ for a.e. $t>0$. In other words, $\gamma:(0,+\infty)\to\mathcal{P}_2(\mathbb{R}^d)$ fulfills

$$\begin{aligned} (\mathrm{Id}, v_t)_{\#}\gamma(t) &\in -\boldsymbol{\partial}_-\mathcal{F}(\gamma(t)) \quad \text{a.e. } t>0,\\ \partial_t\gamma(t) + \nabla_x\cdot(\gamma(t)v_t) &= 0 \quad \text{in } I\times\mathbb{R}^d \end{aligned}$$

in the distributional sense, see [1, Def. 11.1.1]. Here $\mathrm{T}_{\gamma(t)}\mathcal{P}_2(\mathbb{R}^d)$ denotes the regular tangent space of $\mathcal{P}_2(\mathbb{R}^d)$ at $\gamma(t)$.

Interesting functionals are

$$\mathcal{F}(\mu) := \begin{cases} \int_{\mathbb{R}^d} V(x)\,d\mu(x) + \int_{\mathbb{R}^d} f(\rho(x))\,dx + \int_{\mathbb{R}^d} K(x,y)\,d\mu(x)d\mu(x) & \mu=\rho\lambda \\ +\infty & \text{o/w} \end{cases}$$

where the first and second summand lead to the Fokker-Planck equation (Langevin dynamics) when considering

$$\mathrm{KL}(\mu,\nu) = \int_{\mathbb{R}^d}\log\rho(x)\rho(x)\,\mathrm{d}x - \int_{\mathbb{R}^d}\log q(x)\rho(x)\,\mathrm{d}x, \quad \mu=\rho\lambda,\ \nu=q\lambda$$

the first and third summand appear in the discrepancy

$$\mathcal{F}(\mu) = D_K^2(\nu,\nu) := \frac{1}{2}\int_{\mathbb{R}^d}\int_{\mathbb{R}^d} K(x,y)\,\mathrm{d}\sigma(x)\mathrm{d}\sigma(y), \quad \sigma := \mu - \nu.$$

We are interested in the later one. Discrepancy gradient flows or its interaction energy part (third summand E_K) were recently considered e.g. in [2, 3] for various kernels. We want to deal with the kernel $K(x,y) = -\|x-y\|_2$ not contained in the mentioned papers, for which $\mathcal{F}$ is not λ-convex along generalized geodesics. Our motivation came from halftoning problems in image processing [4, 6]. To analyze the flow we have to work in the geometric tangent spaces of $\mathbf{T}_\mu\mathcal{P}_2(\mathbb{R}^d)$ of $\big(\mathcal{P}_2(\mathbb{R}^d), W_2\big)$ at μ. We define the set of directions of steepest descent at μ by

$$\mathrm{H}_-\mathcal{F}(\mu) := \Big\{ \big(\mathrm{H}_{\boldsymbol{v}}^-\mathcal{F}(\mu)\big)^- \cdot \boldsymbol{v} : \boldsymbol{v} \in \operatorname*{arg\,min}_{\substack{\boldsymbol{w}\in\mathbf{T}_\mu\mathcal{P}_2(\mathbb{R}^d),\\ \|\boldsymbol{w}\|_\mu=1}} \mathrm{H}_{\boldsymbol{w}}^-\mathcal{F}(\mu)\Big\}$$

with Hadamard derivative

$$\mathrm{H}_{\boldsymbol{v}}^-\mathcal{F}(\mu) = \liminf_{\substack{t\searrow 0\\ \tilde{\boldsymbol{v}}\to\boldsymbol{v}}} \frac{\mathcal{F}(\gamma_{\tilde{\boldsymbol{v}}}(t)) - \mathcal{F}(\mu)}{t}.$$

Then we call an absolutely continuous curve $\gamma : I \to \mathcal{P}_2(\mathbb{R}^d)$ a *Wasserstein steepest descent flow with respect to* $\mathcal{F}$ if the tangent vectors $\boldsymbol{v}_t = \dot{\boldsymbol{\gamma}}(t)$ exist for all $t \in I$ and are directions of steepest descent, i.e.,

$$\dot{\boldsymbol{\gamma}}(t) \in \mathrm{H}_-\mathcal{F}(\gamma(t)), \qquad t \in I.$$

For functions $\mathcal{F}$ which are λ-convex along generalized geodesics (or more general regular) we establish several relations with the Wasserstein gradient flow which can be characterized by a certain limit of the backward scheme of Jordan, Kinderlehrer and Otto [5].

Unfortunately the discrepancy functional with distance kernel is not convex along geodesics. For the interaction energy $E_{-\|\cdot\|_2}$ part of our special discrepancy function, we obtain the following result.

Proposition (Unique direction of steepest descent at δ_0).
The direction of steepest descent of $E_{-\|\cdot\|_2}$ at δ_0 is given by

$$\dot{\boldsymbol{\gamma}}(0) = \delta_0 \otimes \mu^*$$

with

i) $d = 1$: $\mu_* = \frac{1}{2}\lambda_{[-1,1]} \in \mathcal{P}_2(\mathbb{R})$
ii) $d = 2$: $\mu_* = \rho_R\lambda_{B_R(0)} \in \mathcal{P}_2(\mathbb{R}^2)$, $R = \pi/(2\sqrt{3})$ with

$$\rho_R(x) := \begin{cases} \frac{1}{2\pi R}\left(R^2 - \|x\|_2^2\right)^{-\frac{1}{2}}, & \|x\|_2^2 < R^2,\\ 0, & \text{else} \end{cases}$$

iii) $d = 3$: $\mu_* := \sigma_{\mathbb{S}^2} \in \mathcal{P}_2(\mathbb{R}^3)$, $R = \frac{2}{3}$

A numerical example for the discrepancy flow in 2D from $\delta_{(-1,0)}$ to $\delta_{(1,0)}$ is given in Figure 1.

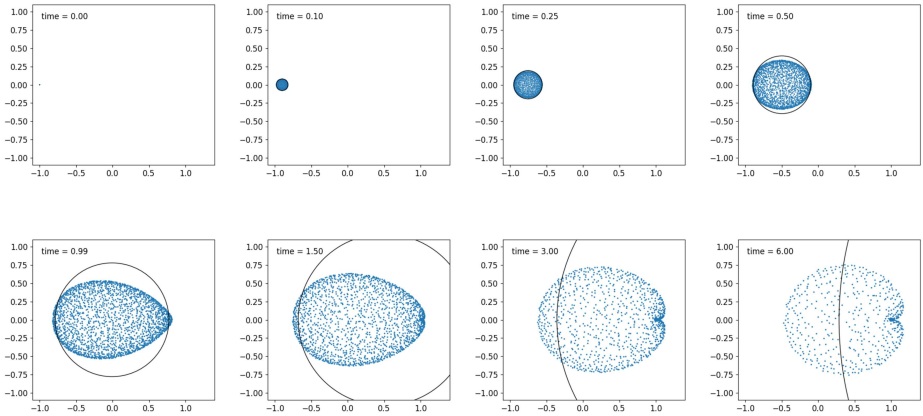

FIGURE 1. 2D particle gradient flow with respect to the discrepancy $\mathrm{D}_{-\|\cdot\|_2}(\delta_{(1,0)}, \cdot)$ starting in a small square around the point $(-1, 0)$. The black circles depict the border of the support of the measures of the geodesic $\gamma_{\boldsymbol{v}}(t)$ starting at $\gamma_{\boldsymbol{v}}(0) = \delta_{(-1,0)}$ in the direction of steepest descent $\boldsymbol{v}$. For small $0 < t < 0.25$ we observe good accordance with our theoretical findings.

REFERENCES

[1] L. Ambrosio, N. Gigli, and G. Savare. *Gradient Flows*. Lectures in Mathematics ETH Zürich. Birkhäuser, Basel, 2005. Topology **32**, 100–120 (1990).
[2] M. Arbel, A. Korba, A. Salim, and A. Gretton. *Maximum Mean Discrepancy Gradient Flow*. Curran Associates Inc., Red Hook, NY, USA, 2019.
[3] J. Carrillo, S. Lisini, and E. Mainini. Gradient flows for non-smooth interaction potentials. *Nonlinear Analysis: Theory, Methods & Applications*, **100**, 122–147 (2014).
[4] M. Ehler, M. Gräf, S. Neumayer, and G. Steidl. Curve based approximation of measures on manifolds by discrepancy minimization. Found. Comput. Math. **21**(6), 1595–1642 (2021).
[5] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of Fokker-Planck equation. SIAM J. Math. Anal. **29**(1), 1–17 (1996).
[6] Teuber, T., Steidl, G., Gwosdek, P., Schmaltz, C., Weickert, J.: Dithering by differences of convex functions. SIAM J. Imaging Sci. **4**(1), 79–108 (2011).

Computing with isometries and developable surfaces

JOHANNES WALLNER

(joint work with Felix Dellinger, Caigui Jiang, Helmut Pottmann, Florian Rist, Cheng Wang)

Developable C^2 surfaces constitute a prominent class of surfaces, besides being important for applications – they represent the shapes of thin sheet material as it bends from a flat state into space without stretching or tearing. It is interesting

that even now new results on developable surfaces can be achieved, and we would like to point here to the proof of existence of surfaces isometric to an annulus which has been folded along concentric circles; see figure below, left and [1].

Geometric modeling with developables is notoriously difficult, and there has been a great number of individual contributions. Nearly all of the well-known geometric properties of developables have been pressed into service for characterizing developability for different kinds of surface representations. These include global ones like the existence of an orthogonal network of geodesic curves [5], local ones like vanishing Gauss curvature, or the special geometry of tangent planes and rulings which developables are known to possess. Quite a few of these properties have led to effective computational treatments of developables, often by means of optimization.

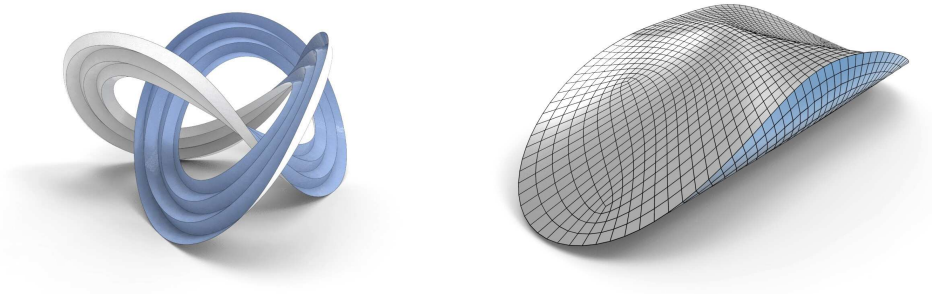

This presentation reports on progress made in recent years. We have used both splines and meshes to model developables with curved creases [2, 6]. We have also been investigating so-called checkerboard patterns associated to general quad meshes with regular combinatorics. This checkerboard pattern is constructed by edge-midpoint subdivision and allows for a useful notion of developability. It is based on the pattern of lines that emerge if discrete tangent planes associated with faces are intersected with the respective tangent planes associated with neighbours. On a general checkerboard pattern approximating a smooth parametric surface these lines represent samples of two distinct sections of the tangent line bundle. In the developable case however these two sections coalesce in the rulings which cover the the surface. As it turns out, this property lends itself to optimization, does not required a special layout of the edges of the mesh we are computing with, and enables the computational design of developable surfaces.

Another way of modeling developables is via local isometry to planar domain [3]. See the figure above, right, for a numerical experiment concerning a piecewise-developable surface which is determined by its unfolding, and which is not covered by the convex case treated by A. D. Alexandrov's 1942 theorem and its continuous analogues. Isometries are highly useful in their own right, e.g. for nonrigid isometric paneling [4]. Work on the topics mentioned above is in fact ongoing.

References

[1] L. Alese. *Propagation of curved folding: The folded annulus with multiple creases exists.* Beitr. Algebra Geom. **63** (2022), 19–43.

[2] C. Jiang, K. Mundilova, F. Rist, J. Wallner, and H. Pottmann. *Curve-pleated structures.* ACM Trans. Graphics **38**/6 (2019), article 169, 13 pp.

[3] C. Jiang, C. Wang, F. Rist, J. Wallner, and H. Pottmann. *Quad-mesh based isometric mappings and developable surfaces.* ACM Trans. Graphics **39**/4 (2020), article 128, 13pp.

[4] C. Jiang, H. Wang, V. Ceballos Inza, F. Dellinger, F. Rist, J. Wallner, and H. Pottmann. *Using isometries for computational design and fabrication.* ACM Trans. Graphics **40**/4 (2021), article 42, 12pp.

[5] M. Rabinovich, T. Hoffmann, O. Sorkine-Hornung. *Discrete geodesic nets for modeling developable surfaces.* ACM Trans. Graphics **37**/2, article 16, 17pp.

[6] C. Tang, P. Bo, J. Wallner, and H. Pottmann. *Interactive design of developable surfaces.* ACM Trans. Graphics **35**/2 (2016), article 12, 12pp.

Repulsive Curves and Surfaces

Henrik Schumacher

(joint work with Caleb Brakensiek, Keenan Crane, Chris Yu, and Philipp Reiter)

Functionals that penalize bending or stretching of a surface play a key role in geometric and scientific computing, but to date have ignored a very basic requirement: in many situations, surfaces must not pass through themselves or each other. In [1], [2], and [3] we develop a numerical framework for optimization of curve and surface geometry while avoiding (self-)collision.

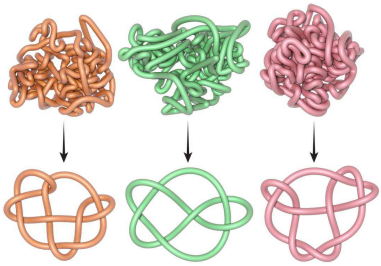

Figure 1. Minimization of the tangent-point energy of curves (subject to constant curve length) can be used to simplify knots and links.

Our starting point is the tangent-point energy, which effectively pushes apart pairs of points that are close in space but distant along the curve/surface. We develop a discretization of this energy for polygonal lines and triangle meshes. To further the optimization process we employ a novel optimization scheme based on a fractional Sobolev inner product developed in [1]. By using the gradient with respect to this inner product as descent direction, we can significantly reduce the number of gradient iterations to arrive at a local minimum (compared to methods based on the frequently employed L^2-gradient). In particular, the number

of iterations is largely independent of the mesh resolution. However, there is an extra cost in computing the Sobolev gradient: We have to solve a linear equation involving the Gram matrix of the employed inner product. Due to the nonlocal nature of this inner product, this matrix is dense. Hence direct linear solvers would be prohibitively expensive. Fortunately, we can solve the linear equation efficiently by employing a geometric multigrid method (cf. [3]) or by employing a CG method preconditioned by two sparse Poisson solves interleaved with the forward application of a further fractional differential operator (cf. [2]). We further accelerate the computation of the nonlocal energy, its derivative, and the forward action of the linear operators by hierarchical approximation techniques in spirit of the Barnes-Hut and fast multipole methods. This allows us to solve optimization problems with various constraints (e.g., area, volume) and in nearly linear time with respect to the number of degrees of freedom.

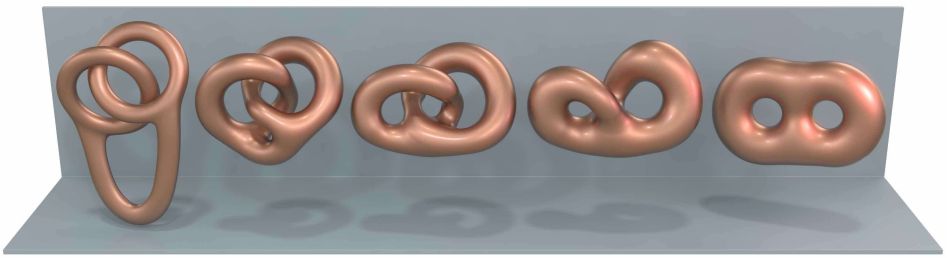

FIGURE 2. The downward gradient flow of the tangent-point energy is an isotopy. Here, it reveals that the seemingly nontrivially interlinked initial surface is actually isotopy-equivalent to the unknotted genus-2 handle body.

Finally, we explore how this machinery can be applied to problems in mathematical visualization, geometric modeling, and geometry processing.

Acknowledgements. This work is supported by the DFG project 282535003: *Geometric curvature functionals: energy landscape and discrete methods*, and by a postdoctoral fellowship of the German Academic Exchange Service (DAAD).

REFERENCES

[1] Ph. Reiter, and H. Schumacher *Sobolev gradients for the Möbius energy. ARMA*, 242(2):701–746, 2021.

[2] C. Yu, C. Brakensiek, H. Schumacher, and C. Keenan *Repulsive Surfaces. ACM Trans. Graph.*, 40(6):1–19, 2021.

[3] C. Yu, H. Schumacher, and C. Keenan *Repulsive Curves. ACM Trans. Graph.*, 40(2):1–21, 2021.

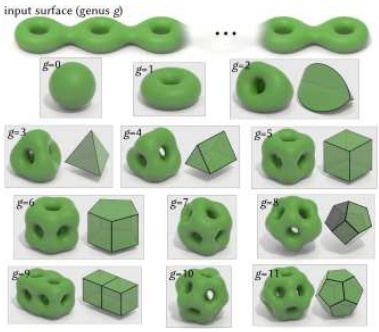

FIGURE 3. The conjectured minimizers of the tangent-point energy in various classes of unknotted handle bodies; some of them seem to enjoy surprising symmetries.

Shape Correspondence by Aligning Scale-Invariant LBO Eigenfunctions

AMIT BRACHA
(joint work with Oshri Halimi, Ron Kimmel)

When matching non-rigid shapes, the regular or scale-invariant Laplace-Beltrami Operator (LBO) eigenfunctions could potentially serve as intrinsic descriptors which are invariant to isometric transformations. However, the computed eigenfunctions of two quasi-isometric surfaces could be substantially different. Such discrepancies include sign ambiguities and possible rotations and reflections within subspaces spanned by eigenfunctions that correspond to similar eigenvalues. Thus, without aligning the corresponding eigenspaces it is difficult to use the eigenfunctions as descriptors. In this talk, I will present the possibility to model the relative transformation between the eigenspaces of two quasi-isometric shapes using a band orthogonal matrix, as well as present a framework that aims to estimate this matrix. Estimating this transformation allows us to align the eigenfunctions of one shape with those of the other, that could then be used as intrinsic, consistent, and robust descriptors. To estimate the transformation, we propose an unsupervised spectral-framework that uses descriptors given by the eigenfunctions of the scale-invariant version of the LBO. Then, using a spectral training mechanism, we find a band limited orthogonal matrix that aligns the two sets of eigenfunctions.

PH-CPF: Planar Hexagonal Meshing using Coordinate Power Fields

MICHAL EDELSTEIN
(joint work with Kacper Pluta, Amir Vaxman, Mirela Ben-Chen)

Trianglular meshes are the most prevalent surface representation in geometry processing. However, for many applications, such as architectural design, animation and numerical simulation, a more structured representation is required, such as

a quadrangular or hexagonal mesh. Specifically in architectural design, meshes with planar faces are amenable to realization using wood or glass panels, and are thus highly sought after. While planar quadrangular meshes have been intensively researched in recent years, planar hexagonal meshes remain an open challenge. Due to considerable technical difficulties, such as a strict restriction on allowable element orientations, and the necessary existence of non-convex faces, existing approaches are either limited to specific types of surfaces, or are not fully automatic, or both.

From the perspective of the geometry processing framework available for meshing, much progress has been made in recent years. Still, the "standard" approach relies on computing vector fields on the surface, which are the *candidate gradients* of the parameterization functions [1]. However, with the exception of a few recent works, there is no guarantee that these candidates are in fact *integrable*, i.e. that there exist parameterization functions whose gradients are these vector fields. Furthermore, while using the candidate gradients as the primary variables is convenient in some applications, in many cases, such as the planar hexagonal meshing application, and especially when the gradients are not expected to be orthogonal, this is not the most natural choice.

In our paper [3], we propose a novel optimization framework for parameterization-based meshing, whose primary variables are a pair of tangent vector fields, which are candidate *Coordinate Power Fields* (CPF). CPFs are a generalization of *coordinate vector fields*, which are the pushforward of the $2D$ coordinate grid under a given local patch parameterization. Unlike coordinate vector fields, coordinate power fields allow for rotational jumps in the parameterization, and thus (up to global holonomy constraints) lead to a seamless global paramterization. Specifically, at the heart of our framework lies a new *continuity* constraint, which two tangent vector fields should fulfill in order to be CPFs, and which provably guarantees the local existence of a seamless parameterization with quantized rotational jumps.

We leverage our framework to compute planar hexagonal meshes. Specifically, we define additional constraints on the CPFs, using a novel *Dupin metric* approach, such that a pushed forward hexagonal grid would be planarizable. We demonstrate that this holds if the CPFs are *conjugate*, and are additionally aligned and sized according to the *Dupin Indicatrix* [2]. By combining these constraints with our optimization framework we automatically generate planar hexagonal meshes that approximate a wide variety of shapes, including complex shapes which, to the best of our knowledge, are not achievable by existing automatic methods. Furthermore, on meshes where existing approaches are applicable, we achieve a comparable or better planarity error as existing approaches, while using a considerably smaller number of elements. Finally, we demonstrate that our framework is additionally applicable to planar quad meshing.

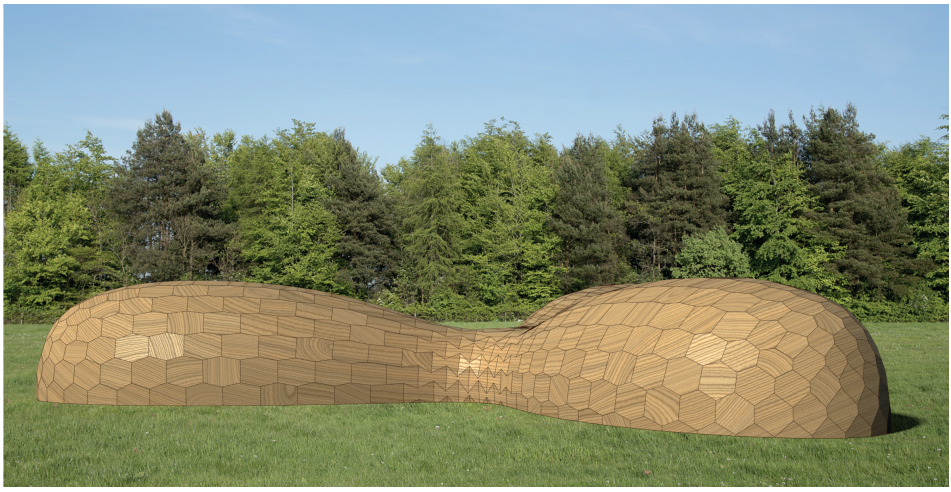

FIGURE 1. A computer rendered image of one of the planar hexagonal meshes computed with our approach.

REFERENCES

[1] D. Bommes, B. Lévy, N. Pietroni and E. Puppo, C. Silva, M. Tarini and D. Zorin, *Quad-mesh generation and processing: A survey*, Computer Graphics Forum, vol. 32, no. 6, pp. 51–76. 2013.
[2] M.P. Do Carmo, *Differential geometry of curves and surfaces: revised and updated second edition*, Courier Dover Publications, 2016.
[3] K. Pluta, M. Edelstein, A. Vaxman and M. Ben-Chen, *PH-CPF: planar hexagonal meshing using coordinate power fields*, ACM Transactions on Graphics (TOG), 40.4 (2021): 1-19.

Learning Low Bending and Low Distortion Manifold Embeddings

JULIANE BRAUNSMANN
(joint work with Marko Rajković, Martin Rumpf, Benedikt Wirth)

Autoencoders, which consist of an encoder and a decoder, are widely used in machine learning for dimension reduction of high-dimensional data. The encoder $\phi : M \to \mathbb{R}^l$ embeds the input data manifold $M \subset \mathbb{R}^n$ into a lower-dimensional latent space $\mathbb{R}^l$, while the decoder $\psi : \mathbb{R}^l \to \mathbb{R}^n$ represents the inverse map. A good regularity and structure of the embedded manifold may substantially simplify further data processing tasks such as cluster analysis or data interpolation. In [1] we propose a novel regularization for learning the encoder map ϕ by considering a loss functional that prefers isometric, extrinsically flat embeddings. We consider M to be a compact, m-dimensional Riemannian manifold (without boundary) with metric g and denote by $d_M(x, y)$ the Riemannian distance between any two points $x, y \in M$ and by $\mathrm{av}_M(x, y) = \exp_x \frac{v}{2}$ their geodesic midpoint, where $v \in T_x M$ is such that $\exp_x v = y$. For a fixed, admissible *sampling radius* $\epsilon > 0$ we consider

a finite training data set $\mathcal{S}_\epsilon \subset \{(x,y) \in M \times M : d_M(x,y) \leq \epsilon\}$ and define the *discrete sampling loss functional*

$$E^{\mathcal{S}_\epsilon}[\phi] \coloneqq \frac{1}{|\mathcal{S}_\epsilon|} \sum_{(x,y)\in\mathcal{S}_\epsilon} \left(\gamma(|\partial_{(x,y)}\phi|) + \lambda\, |\partial^2_{(x,y)}\phi|^2\right),$$

with $\lambda > 0$ and first and second order difference quotients defined as

$$\partial_{(x,y)}\phi \coloneqq \tfrac{\phi(y)-\phi(x)}{d_M(x,y)},\ \partial^2_{(x,y)}\phi \coloneqq 8\tfrac{\mathrm{av}_{\mathbb{R}^l}(\phi(x),\phi(y))-\phi(\mathrm{av}_M(x,y))}{d_M(x,y)^2},$$

where $\mathrm{av}_{\mathbb{R}^l}(a,b) = (a+b)/2$ denotes the linear average in $\mathbb{R}^l$ and $\gamma : [0,\infty) \to [0,\infty)$ a strictly convex function with minimum $\gamma(1) = 0$. In order to evaluate the functional, input triples $(x, y, \mathrm{av}_M(x,y))$ and $d_M(x,y)$ are needed.

The first term in $E^{\mathcal{S}_\epsilon}$ has a strict minimum for $|\partial_{(x,y)}\phi| = 1$ and thus promotes $|\phi(x)-\phi(y)| \approx d_M(x,y)$ and therefore low distortion and an approximate isometry. The second term in $E^{\mathcal{S}_\epsilon}$ penalizes the deviation of intrinsic averages on $\phi(M)$ from extrinsic ones in $\mathbb{R}^l$ and thus low bending.

When sampling first $x \in M$ uniformly and then y uniformly in $B^M_\epsilon(x)$, the geodesic ball around x, we have, as the number of samples goes to infinity, almost sure convergence to the *continuous sampling loss functional*

$$\mathcal{E}^\epsilon[\phi] \coloneqq \fint_M \fint_{B^M_\epsilon(x)} \gamma(|\partial_{(x,y)}\phi|) + \lambda\, |\partial^2_{(x,y)}\phi|^2 \,\mathrm{d}V_g(y)\,\mathrm{d}V_g(x),$$

where V_g denotes the Riemann–Lebesgue volume measure on M.

In [1] we show the above functional is consistent with the local limit loss functional defined on $H^2(M,\mathbb{R}^l)$ given by

$$\mathcal{E}[\phi] \coloneqq \int_M \Gamma(\operatorname{grad}\phi(x)) + \lambda \|\operatorname{Hess}\phi(x)\|^2 \,\mathrm{d}V_g(x) \quad \text{with}$$

$$\Gamma(W) \coloneqq \fint_{B^m_1(0)} \gamma(|g_x(W,\bar{w})|)\,\mathrm{d}w, \quad \|A\|^2 \coloneqq \fint_{B^m_1(0)} |g_x(A[x\bar{w}],x\bar{w})|^2\,\mathrm{d}w$$

for all $W \in T_xM$ and $A \in L(T_xM, T_xM)$ and $\bar{w} = \frac{w}{|w|}$. The first term penalizes deviation from an isometric embedding. The second term defines a squared norm on the space of symmetric endomorphisms on T_xM and thus penalizes a non-vanishing Riemannian Hessian and thus extrinsic bending.

Analogously to approaches in numerical analysis, in which an infinite dimensional optimization problem is approximated by introducing an auxiliary length scale parameter ϵ and restricting to a space of functions parametrized with finitely many parameters, in [2] we introduce function spaces $\mathcal{F}^\epsilon(M,\mathbb{R}^l) \subset H^2(M)$ and the modified functional

$$\mathcal{E}^\epsilon_{\mathcal{F}}[\phi] \coloneqq \begin{cases} \mathcal{E}^\epsilon[\phi] & \text{if } \phi \in \mathcal{F}^\epsilon(M,\mathbb{R}^l), \\ \infty & \text{else.} \end{cases}$$

In our application, the discrete function space $\mathcal{F}^\epsilon(M,\mathbb{R}^l)$ consists of deep neural networks. To relate the above discrete problem to the original one, the expressivity of the discrete function space has to increase in a way compatible with the discrete functional as $\epsilon \to 0$. In our case this corresponds to a regularity measure of the

discretized functions decreasing more slowly than ϵ asymptotically, which can be implemented as a weight constraint on the neural networks.

More formally, we ask for the following conditions:

- For every $\epsilon > 0$ $\mathcal{F}^\epsilon(M, \mathbb{R}^l)$ is closed as a subset of $L^2(M, \mathbb{R}^l)$,
- $\mathcal{F}^\epsilon(M, \mathbb{R}^l) \subset C^{2,1}(M, \mathbb{R}^l)$ with $\sup_{\phi^\epsilon \in \mathcal{F}^\epsilon(M,\mathbb{R}^l)} L_{\text{grad}}(\phi^\epsilon) + L_{\text{Hess}}(\phi^\epsilon) = o(\frac{1}{\epsilon})$,
- For every $\phi \in H^2(M, \mathbb{R}^l)$ there exists a sequence $\{\phi^\epsilon\}_{\epsilon>0}$ with $\phi^\epsilon \in \mathcal{F}^\epsilon(M, \mathbb{R}^l)$ such that $\lim_{\epsilon\to 0} \|\phi^\epsilon - \phi\|_{H^2(M,\mathbb{R}^l)} = 0$,

where $L_{\text{grad}}(\phi^\epsilon)$ and $L_{\text{Hess}}(\phi^\epsilon)$ denote the Lipschitz constants of the Riemannian gradient and Hessian of ϕ^ϵ, respectively, as defined in [3].

Theorem 1. *Under the above conditions, the nonlocal regularization energies $\{\mathcal{E}^\epsilon_{\mathcal{F}}\}_{\epsilon>0}$ converge to the continuous regularization energy $\mathcal{E}$ as $\epsilon \to 0$ in the sense of Mosco* [4] *in $H^2(M, \mathbb{R}^l)$. Further, for every small enough $\epsilon > 0$ there exists a minimizer ϕ^ϵ of $\mathcal{E}^\epsilon_{\mathcal{F}}$ and $\|\phi^\epsilon\|_{H^2(M,\mathbb{R}^l)} \leq C$ independently of ϵ.*

The proof relies on Taylor expansions, where the Lipschitz bounds ensure convergence of the remainder terms.

To test our regularization functional, we use artificial image manifolds where the manifold is explicitly known. In general imaging applications there exist methods to define a manifold structure and thus compute distances and averages, such as LDDMM [5], metamorphosis [7] or optimal transport [6]. Our method allows to either train the encoder ϕ separately by minimizing $E^{\mathcal{S}_\epsilon}[\phi]$ and to train the decoder map ψ subsequently by minimizing, for fixed ϕ, the *reconstruction loss*

$$R[\phi, \psi] = \frac{1}{2|\mathcal{S}_\epsilon|} \sum_{(x,y)\in\mathcal{S}_\epsilon} \|\psi(\phi(x)) - x\|^2_{L^2} + \|\psi(\phi(y)) - y\|^2_{L^2}$$

with $\|\cdot\|_{L^2}$ the discrete L^2-norm on pixel images. Alternatively, encoder and decoder can be trained simultaneously by minimizing $E^{\mathcal{S}_\epsilon}[\phi] + \kappa R[\phi, \psi]$ for some $\kappa > 0$. In our experiments, the latent space dimension l was chosen larger than the minimum value required for a smooth embedding. In this way, the encoder might achieve flatter embeddings. We observe in practice that the number of used dimensions, determined by PCA on a point cloud in latent space computed from a test set of images, indeed tends to be higher than the minimum required value for any choice of λ. However, removing the bending term, *i.e.*, setting $\lambda = 0$, leads to even more dimensions being used, since isometric embeddings are possible using arbitrarily many Euclidean dimensions and no preference on the amount of used dimensions is given. For x and y moderately far apart, we further observe that inclusion of the bending term leads to lower errors between averages obtained via $\psi(\frac{\phi(x)+\phi(y)}{2})$ and the ground truth averages $\text{av}_M(x, y)$, even if the encoder was only trained on point pairs with small distance.

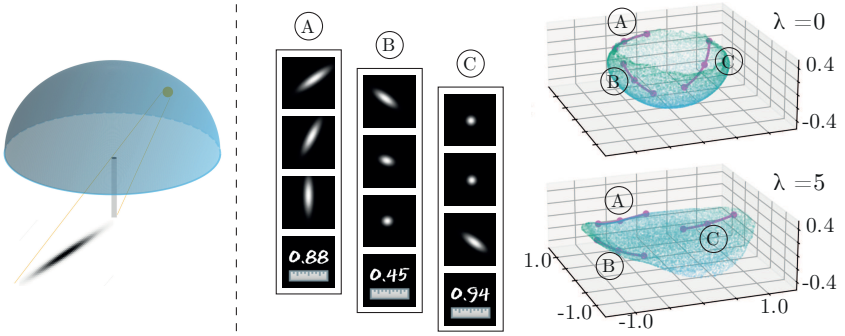

FIGURE 1. Example of an image manifold. Images represent shadows of a pole cast by a lightsource moving on the hemisphere, which induces the geometry. Three input triples are shown in the middle. On the right, the first three PCA components are used to visualize the obtained embeddings for $\lambda = 0$ and $\lambda = 5$.

REFERENCES

[1] J. Braunsmann, M. Rajković, M. Rumpf, B. Wirth, *Learning Low Bending and Low Distortion Manifold Embeddings*, Proceedings of the 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2021), 4411–4419.

[2] J. Braunsmann, M. Rajković, M. Rumpf, B. Wirth, *Learning Low Bending and Low Distortion Manifold Embeddings: Theory and Applications*, arXiv:2208.10193 (2022).

[3] N. Boumal, *An introduction to optimization on smooth manifolds*, To appear with Cambridge University Press (2022).

[4] U. Mosco, *Convergence of convex sets and of solutions of variational inequalities*, Advances in Mathematics, 3, Issue 4 (1969), 510–585.

[5] L. Younes, *Shapes and diffeomorphisms*, Applied Mathematical Sciences, 171 (2010).

[6] M. Cuturi, G. Peyré, *Computational Optimal Transport: With Applications to Data Science*, Foundations and Trends in Machine Learning, 11 (2019), 355–607

[7] A. Trouvé, L. Younes, *Local geometry of deformable templates*, SIAM J. Math. Anal., 37 (2006), 17–59

Towards convergence of the heat method via large deviations

SIMON SCHWARZ

(joint work with Anja Sturm and Max Wardetzky)

Consider a polyhedral surface Γ equipped with a discrete Laplace operator Δ. The heat method proposed in [1] is frequently used to compute the geodesic distance on Γ to a given subset γ.

The algorithm is motivated by the short-time asymptotics of the heat kernel p^M on a Riemannian manifold (M, g): In this case, a classical result by Varadhan [3] shows that

$$\lim_{t \to 0} -2t \log p_t^M(x, y) = d_g^2(x, y) \tag{1}$$

Algorithm 1 The Heat Method for distance computing

1: Approximate the heat kernel by $(\Delta - t\,\mathrm{Id})u = \delta_\gamma$ for some fixed time t.
2: Evaluate the vector field $X = -\nabla u/|\nabla u|$.
3: Solve the Poisson equation $\Delta\phi = \nabla \cdot X$.

holds uniformly for all $x, y \in M$. However, the heat method uses a backward Euler step approximation of the *discrete heat kernel*. Hence, Varadhan's formula does not apply in our setting – it is even known that the short-time asymptotics of the discrete heat kernel p^G on a graph $G = (V, E)$ equipped with any graph Laplacian differs substantially from the smooth case and satisfies (see [2])

$$\lim_{t\to 0} \frac{\log p_t^G(x,y)}{\log t} = d^c(x,y), \tag{2}$$

where d^c denotes the combinatorial graph distance. This difference in the convergence behavior implies that we have to decrease the edge length h of the polyhedral surface and time t simultaneously, to possibly obtain convergence to the geodesic distance. This is consistent to the suggestion in [1] to choose $t = h^2$ in step 1.

Our main result – proven by establishing a large deviation principle – makes this relation precise for sequences of planar infinite graphs with decreasing edge lengths and shows that there is a phase transition in β for $t = h^\beta$: If $\beta \leq 1$ we can recover the Euclidean distance corresponding to the behaviour in (1), while we obtain the combinatorial graph distance for $\beta > 1$ similar to (2).

Theorem 1. *Let $(\Gamma_h)_{h>0}$ be a sequence of planar infinite graphs embedded in $\mathbb{C}$ with decreasing edge length of order h satisfying some additional assumptions and let p_h, $h > 0$ denote the discrete heat kernels obtained by using a geometric graph Laplacian. Then, for any $\beta \in (0, 1]$, $x \in \mathbb{C}$ and fixed $t > 0$,*

$$\lim_{h\to 0} h^\beta \log p^h_{h^\beta t}(x,y) = -\frac{1}{2t}|x-y|^2$$

uniformly in $y \in \mathbb{C}$ over compact sets. If $\beta > 1$ and the limit

$$d_J(x,y) = \lim_{h\to 0} h d^c_h(x,y)$$

exists, where d^c_h is the combinatorial distance on Γ_h, then

$$\lim_{h\to 0} \frac{h}{\log h^{\beta-1}} \log p^h_{h^\beta}(x,y) = d_J(x,y)$$

for any $x, y \in \mathbb{C}$.

The theorem yields convergence of the heat method on the plane, since we find

$$h \log u_{h^2}(x,y) \sim \max_{c>0} \left(-c + h \log p^h_{ch}(x,y)\right) \sim -\sqrt{2}\,|x-y| \quad \text{as } h \to 0$$

by using a Laplace integral asymptotic. Notice that our theoretical insights suggests that it is possible to replace steps 2 and 3 in the heat method by directly computing $\frac{h}{\sqrt{2}} \log u$.

References

[1] K. Crane, C. Weischedel, and M. Wardetzky, Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow, *ACM Trans. Graph.* 32(5), 2013.

[2] M. Keller, D. Lenz, F. Münch, M. Schmidt, and A. Telcs, Note on short-time behavior of semigroups associated to self-adjoint operators, *Bull. Lond. Math. Soc.* 48(6):935–944, 2016.

[3] S. R. S. Varadhan, On the behavior of the fundamental solution of the heat equation with variable coefficients, *Comm. Pure Appl. Math.* 20(2):431–455. 1967.

Discrete Riemannian calculus on implicit latent manifolds

Florine Hartwig

(joint work with Juliane Braunsmann, Marko Rajković, Martin Rumpf, Josua Sassen, Benedikt Wirth)

Implicit representations of manifolds are a flexible and effective tool for many applications in geometry processing. However, computing geometric operators on them is challenging. We develop a scheme for discrete Riemannian calculus on implicit manifold representations, allowing us to efficiently compute discrete geodesics, exponential shooting, and parallel transport. This approach fits nicely into a general framework [1], which shows that it is a consistent discretization of Riemannian calculus. We show how this scheme can be applied to compute geometric operators on learned latent manifolds. This allows us to compute visually appealing shape interpolations of high-dimensional data on a learned low-dimensional manifold representation derived by an autoencoder.

Discrete geodesic calculus on implicit manifolds. Let $\mathcal{M} = \{X \in \mathbb{R}^n \mid \Phi(X) = 0\}$ be a d-dimensional embedded manifold with $\Phi : \mathbb{R}^n \to \mathbb{R}^{n-d}$ and $\operatorname{rank} D\Phi(X) \neq 0$ for all $X \in \mathbb{R}^n$. The Riemannian distance on $\mathcal{M}$ can be locally approximated by the Euclidean distance in ambient space. For the discretization of the geometric operators, we follow [1] and compute discrete geodesics on $\mathcal{M}$ by minimizing the discrete path energy

$$\mathbf{E}^K[(X_0 = X_A, \mathbf{X}, X_K = X_B)] = \sum_{i=1}^{K} |X_k - X_{k-1}|^2 .$$

To this end, we use an augmented Lagrangian approach to incorporate the constraint $\Phi(X_k) = 0$ for every point X_k of a discrete path $\mathbf{X}$ on $\mathcal{M}$. For this, we consider the augmented Lagrangian

$$L[\mathbf{X}, \Lambda, \mu] = \mathbf{E}^K[(X_0 = X_A, \mathbf{X}, X_K = X_B)] + \Lambda : \Phi(\mathbf{X}) + \frac{\mu}{2}\|\Phi(\mathbf{X})\|_2^2 ,$$

with Lagrangian multiplier $\Lambda \in \mathbb{R}^{n-d,K-1}$, penalty factor $\mu > 0$ and $\Phi(\mathbf{X}) = (\phi_i(X_k))_{k=0,\dots,K}^{i=1,\dots,n-d}$. This combines a pure penalty method which often leads to an ill-conditioned problem with a classical Lagrangian approach. Exponential shooting can be computed by discretizing the continuous geodesic equation $\ddot{X}(t) \subset T_{X(t)}\mathcal{M}$ for an arc-length parameterized curve $X(t) \subset \mathcal{M}$. See Figure 1 for exemplary results on analytic surfaces.

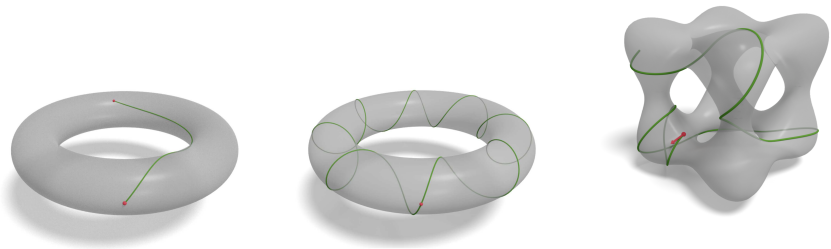

FIGURE 1. Geodesic on a torus (left), closed geodesic on a torus (middle) and exponential shooting on a quartic surface (right) .

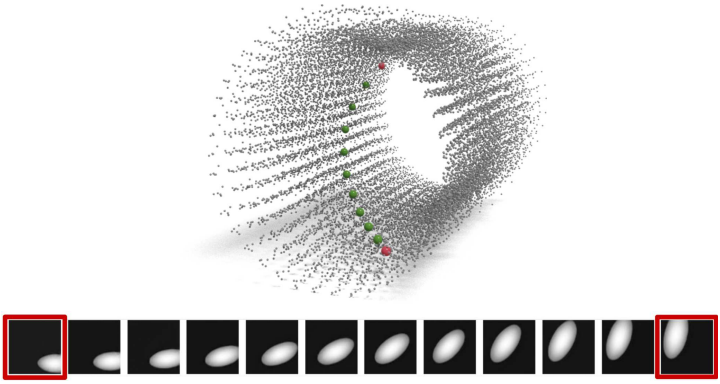

FIGURE 2. Results for a toy model with images of ellipses as input manifold. Top: Computed geodesic on the latent manifold with $K = 12$. The 4-dimensional latent manifold embeddedd in $\mathbb{R}^{16}$ is visualized by encoding 10k points of the input dataset and choosing 3 dimensions based on PCA to display (point cloud). Bottom: The geodesic is decoded to the image space.

Discrete geodesic calculus on latent manifolds. We consider the autoencoder derived in [2], that autoencoder promotes an isometric and intrinsically flat embedding of a high-dimensional input manifold. Hence, distances on the derived latent manifold $\mathcal{M}$ correspond to geodesic distances on the original data manifold. We follow a denoising approach [3] to learn a projection $p : \mathbb{R}^n \to \mathcal{M} \subset \mathbb{R}^n$ on the latent manifold $\mathcal{M}$. Then $\Phi = \mathrm{id} - p$ gives an approximate implicit representation of $\mathcal{M}$. As a numerical experiment, we consider images of anisotropic Gaussians, which are rotated, scaled and translated (experiment (G) of [2]). The learned manifold is a 4-dimensional cylinder $S^1 \times [0.5, 1.5] \times [-1, 1]^2$. See Figure 2 for a result.

References

[1] M. Rumpf, B. Wirth, *Variational time discretization of geodesic calculus*, IMA Journal of Numerical Analysis **35**(3) (2015), 1011–46.

[2] J. Braunsmann, M. Rajković, M. Rumpf, B. Wirth, *Learning Low Bending and Low Distortion Manifold Embeddings*, Proceedings of the 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2021), 4411–4419.

[3] G., Alain, Y. Bengio, *What regularized auto-encoders learn from the data-generating distribution*, The Journal of Machine Learning Research (2014) **15**(1), 3563–93.

Repulsive Shells

Josua Sassen

(joint work with Keenan Crane, Martin Rumpf, Henrik Schumacher)

Shape interpolation and extrapolation are among the core tasks in geometry processing. However, to date most approaches to these problems are not able to prevent self-intersections. We propose a method that is able to do so by leveraging recent work on repulsive energies. To this end, we build on previous work that studies the space of triangular surfaces as a Riemannian manifold with a metric based on discrete shell elasticity. We show how to incorporate the repulsive tangent-point energy into this metric, such that shapes with self-intersection are pushed infinitely far away. This guarantees that geodesics on this *shape space of repulsive shells* consist of intersection-free surfaces and thus form a useful tool for shape interpolation. We apply variational time-discretization and develop a trust-region approach with appropriate preconditioners for the resulting minimization problem. Finally, we explore further potential applications for the combination of elastic and repulsive energies.

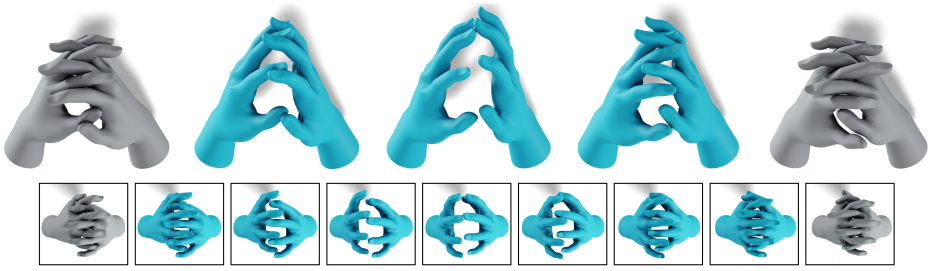

Figure 1. Intersection-free interpolation between clasping hands. Our proposed method computes interpolations between surfaces as geodesics on a Riemannian shape space, where the metric guarantees the avoidance of self-intersections.

The Space of Discrete Shells. We begin by recalling the underlying Riemannian structure on the space of immersions of a discrete surface introduced in [1]. To this end, we consider a triangular surface S_h with vertices $\mathbf{V}$, edges $\mathbf{E} \subset \mathbf{V} \times \mathbf{V}$, and faces $\mathbf{F} \subset \mathbf{V} \times \mathbf{V} \times \mathbf{V}$. Furthermore, we consider piecewise affine immersions $x\colon S_h \to \mathbb{R}^3$ describing the midsurface of a discrete thin shell and study piece-wise affine deformations of S_h. Then we consider the space $\mathcal{M}$ of all such immersion modulo rigid body motions and aim to equip it with a Riemannian metric. To this end, we consider a discrete elastic shell energy $\mathcal{W}[x, \tilde{x}] = \mathcal{W}_{\text{mem}}[x, \tilde{x}] + \mathcal{W}_{\text{bend}}[x, \tilde{x}]$ as introduced in [1, 2] consisting of a membrane contribution $\mathcal{W}_{\text{mem}}$ and a bending contribution $\mathcal{W}_{\text{bend}}$. To obtain a metric, we need to pass to a viscous model, which we do via Rayleigh's paradigm and can finally construct the metric $g^s_x(u, v) := \frac{1}{2}\ \mathrm{d}^2_y \mathcal{W}(x, y)\big|_{y=x}(u, v)$. With this metric at hand, we can define geodesics $c\colon [0, 1] \to \mathcal{M}$, variationally as critical points of the path energy

$$\mathcal{E}^s[c] := \int_0^1 g^s_{c(t)}(\dot{c}(t), \dot{c}(t))\,\mathrm{d}t,$$

for fixed endpoints $c(0)$ and $c(1)$. To compute geodesics numerically, we define for $(x_0, \dots, x_K) \in \mathcal{M}^{K+1}$ the discrete path energy

$$E^K_s[x_0, \dots, x_K] := K \sum_{k=1}^{K} \mathcal{W}(x_{k-1}, x_k),$$

which has been shown to be a consistent discretization in [3].

The Space of Repulsive Shells. To construct our modified metric that avoids self-intersections, we consider the tangent-point energy first introduced for curves in [4]. We first recall its formulation for a continuous surface $\mathcal{S}$ with immersion ϕ. For two points $s, t \in \phi(\mathcal{S})$ the tangent-point radius $R^\phi_{\text{tp}}(s, t)$ is defined as the radius of the smallest sphere through s and t that is also tangent to $\phi(\mathcal{S})$ at s. This radius can be expressed as $R^\phi_{\text{tp}}(s, t) := \frac{|s-t|^2}{2\,|P_\phi(s)\,(s-t)|}$, where $P_\phi(s)$ denotes the orthogonal projection onto the normal space at s. The tangent-point energy for integrability parameter $p \geq 1$ is then defined as the integral of the inverse tangent-point radius for all pairs of points, i.e.

$$\mathcal{V}(\phi) := \int_{\phi(\mathcal{S})} \int_{\phi(\mathcal{S})} \frac{|P_\phi(s)\,(s-t)|^p}{|s-t|^{2p}}\,\mathrm{d}s\,\mathrm{d}t.$$

In [5], it has been proven that for $p > 4$ this indeed provides a barrier against self-intersections, i.e. every C^1-immersion with finite energy $\mathcal{V}$ is already a $C^{1,\alpha}$-*embedding* with $\alpha = 1 - \frac{4}{p} > 0$. Due to its nonlocality, this energy is challenging to discretize and we follow the approach from [6] to obtain a discrete energy, which we will refer to as $\mathcal{V}_h$ going forward.

In our approach, we combine the measurement of infinitesimal membrane and bending distortion with infinitesimal changes of the tangent-point energy. The

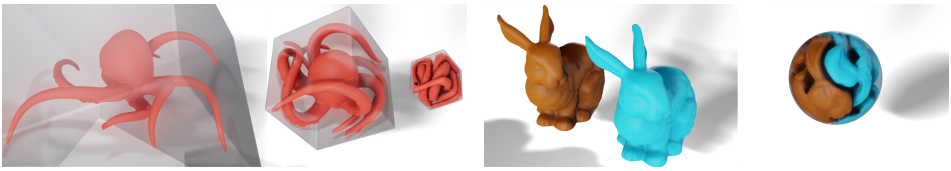

FIGURE 2. Packing objects by minimizing the sum of elastic, tangent-point energy and a barrier term for a given, shrinking geometry.

combination is realized by the graph of $\mathcal{V}_h$ over the shape space $\mathcal{M}$ of viscous shells. That is, we define

$$\mathcal{M}^{\mathcal{V}} := \{(x, \mathcal{V}_h(x)) \mid x \in \mathcal{M}\} \subset \mathcal{M} \times \mathbb{R}^+ \cup \{\infty\}.$$

We can picture $\mathcal{M}^{\mathcal{V}}$ as a mountain range over the underlying shape space, where the height of mountains corresponds to the amount of tangent-point energy a given shape carries.

Then, the metric $\tilde{g}$ obtained by pulling back the product metric from $\mathcal{M}^{\mathcal{V}}$ to $\mathcal{M}$ is given by $g_x(u,v) := g^s_x(u,v) + (\,\mathrm{d}_x\mathcal{V}_h\, u)\,(\,\mathrm{d}_x\mathcal{V}_h\, v)$, for $x \in \mathcal{M}$ and $u, v \in T_x\mathcal{M}$. Indeed, one can show that with this metric immersions with self-intersections are infinitely far away from ones without. The corresponding path energy is given by

$$\mathcal{E}[c] := \int_0^1 g^s_{c(t)}(\dot{c}(t), \dot{c}(t)) + (\,\mathrm{d}_{c(t)}\mathcal{V}_h\, \dot{c}(t))^2 \,\mathrm{d}t.$$

We can again apply the same time-discretization as before and obtain

$$E^K[x_0, \ldots, x_K] := K \sum_{k=1}^{K} \mathcal{W}(x_{k-1}, x_k) + |\mathcal{V}_h(x_{k-1}) - \mathcal{V}_h(x_k)|^2.$$

To minimize this energy numerically, we use a trust-region Newton method, where we approximate the Hessian of the tangent-point part of the discrete path energy via a Gauß-Newton approach and solve the resulting trust-region subproblem via Steihaug's CG method. In Figure 1, we show an exemplary result for a discrete geodesic between two clasping hands.

Elastic Deformations. The combination of elastic and tangent-point energy is also useful for applications beyond the shape space framework discussed so far. It models deformations of a shell repelling itself, which especially means that it avoids intersections. That means, we consider the variational problem

$$\min_x \; \mathcal{W}(\tilde{x}, x) + \mathcal{V}(x), \tag{1}$$

where $\tilde{x}$ is some given reference configuration and add further terms based on the specific application.

One application we consider is packing objects into tight spaces. To this end, we add a barrier term to (1) that forces the surface to stay inside a given geometry.

Then, in Figure 2, we reduce the size of the geometry step-by-step and thus get a tight packing of the shapes.

References

[1] B. Heeren, M. Rumpf, M. Wardetzky, B. Wirth, *Time-Discrete Geodesics in the Space of Shells*, Computer Graphics Forum **31** (2012), 1755–1764.

[2] E. Grinspun, A. Hirani, M. Desbrun, P. Schröder, *Discrete shells*, Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (2003), 62–67.

[3] M. Rumpf, B. Wirth, *Variational time discretization of geodesic calculus*, IMA Journal of Numerical Analysis **35** (2015), 1011–1046.

[4] G. Buck, J. Orloff, *A simple energy function for knots*, Topology and its Applications **61** (1995), 205–214.

[5] P. Strzelecki, H. von der Mosel, *Tangent-Point Repulsive Potentials for a Class of Non-smooth m-dimensional Sets in $\mathbb{R}^n$. Part I: Smoothing and Self-avoidance Effects*, Journal of Geometric Analysis **23** (2013), 1085–1139.

[6] C. Yu, C. Brakensiek, H. Schumacher, K. Crane, *Repulsive Surfaces*, ACM Transactions on Graphics **40** (2021), 1–19.

A finite element approach for minimizing line and surface energies arising in the study of singularities in liquid crystals

Dominik Stantejsky

(joint work with François Alouges, Antonin Chambolle)

Motivated by a problem originating in the study of defect structures in nematic liquid crystals [1, 2], we consider the numerical minimization of an energy consisting of the following energy

$$\mathcal{E}_0(T) = \mathbb{M}(T|_\Omega) + \int_{\mathcal{M}} |\nu_3| \, \mathrm{d}\mu_{T|_{\mathcal{M}}} + \beta \, \mathbb{M}(\partial T + \Gamma) \, .$$

The central object is a two dimensional flat chain T in $\mathbb{R}^3$ with coefficients in $\mathbb{Z}_2$. The first contribution to the energy is given by the mass of T outside an obstacle E with the notation $\Omega = \mathbb{R}^3 \setminus \overline{E}$. The second part is the integral of a density integrated over the part of T on the obstacle surface $\mathcal{M} := \partial E$. In our case the density is given by $|\nu_3|$, the absolute value of the $3-$component of the normal vector field on $\mathcal{M}$. The last term is the mass of the boundary ∂T reduced by a given prescribed curve Γ and weighted by a parameter $\beta \in (0, \infty)$. The latter is essential since for $\Gamma = 0$ the minimizer is trivial.

While the problem in [2] is stated for finite mass rectifiable flat chains with coefficients in $\mathbb{Z}_2$, we relax this assumption and consider T to be a normal current in order to obtain a convex problem. The minimization of the energy can been seen as generalization of the obstacle problem and Plateau problem. If $\beta \gg 1$ and $E = \emptyset$ this reduces to the classical Plateau problem which has recently been treated with similar methods in [5].

The algorithm we present in [3] is based on the Alternating Direction Method of Multipliers used to minimize a finite element representation of our energy. We represent T by Nédélec finite elements and ∂T by the $\mathbb{P}^0-$field given by its curl.

For vanishing curl away from ∂T, the Nédélec elements are a gradient and our algorithm reduces to the minimization of the total variation as e.g. in [4]. The obstacle surface $\mathcal{M}$ is extended to a boundary layer. This allows us to take into account the obstacle, the density $|\nu_3|$ and the parameter β in a unified way via penalization. A priori information about the minimizers allow us to use structured meshes to improve the quality of our solutions.

We are able to obtain minimizing configurations of the problem studied in [2] and validate our algorithm in the case of a spherical obstacle.

References

[1] F. Alouges, A. Chambolle, and D. Stantejsky. The saturn ring effect in nematic liquid crystals with external field: Effective energy and hysteresis. *Arch. Ration. Mech. Anal.*, 2021.
[2] F. Alouges, A. Chambolle, and D. Stantejsky. Convergence to line and surface energies in nematic liquid crystal colloids with external magnetic field. *Preprint*, 2022.
[3] D. Stantejsky. A finite element approach for minimizing line and surface energies arising in the study of singularities in liquid crystals. *In preparation*, 2022.
[4] A. Chambolle, and T. Pock. Approximating the total variation with finite differences or finite elements. *Geometric Partial Differential Equations - Part II, Elsevier*, 2021.
[5] S. Wang, and A. Chern. Computing Minimal Surfaces with Differential Forms. *ACM Trans. Graph.,* 40(4), 2021.

Learning from Synthetic 3D Priors for Real-World 3D Perception

Angela Dai
(joint work with Alexey Bokhovkin, Pablo Palafox)

Understanding the 3D structure of real-world environments is a fundamental challenge in machine perception, with many applications towards robotic navigation and interaction, content creation, and mixed reality scenarios. In this talk, we propose to develop learned neural parametric models to capture structural and object priors from large-scale synthetic shape datasets. We show that such neural parametric models can be used to model a space of shape parts, which can be used to fit to real-world observations from commodity sensors at test time (Figure 1). In particular, neural coordinate field representations can be leveraged to form parametric models of shape and part geometry as signed distance fields; given an input RGB-D scan with detected objects, rather than making independent predictions for each object's structure, the learned parametric model can be used to fit against both individual object observations and to multiple observations across a scene.

We further demonstrate the effectiveness of learned neural parametric models for deformable shapes to reconstruct and track depth sequence observations. Here, a disentanglement of deformable object representation into shape and pose, along with a structural decomposition into parts, enables robust fitting to new depth observations even under random initialization (Figure 2). These learned parametric models will enable the construction of intuitive, semantic primitives for future virtual or real-world interaction or manipulation of real environments.

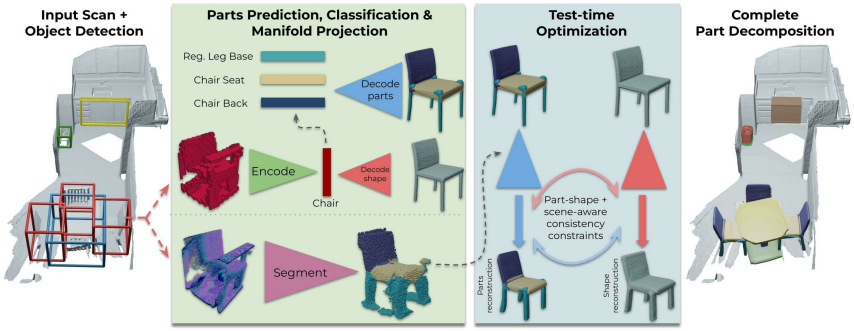

FIGURE 1. Learned parametric models of shape parts from synthetic shape data enable test-time fitting and holistic scene optimization for new real-world observations [1].

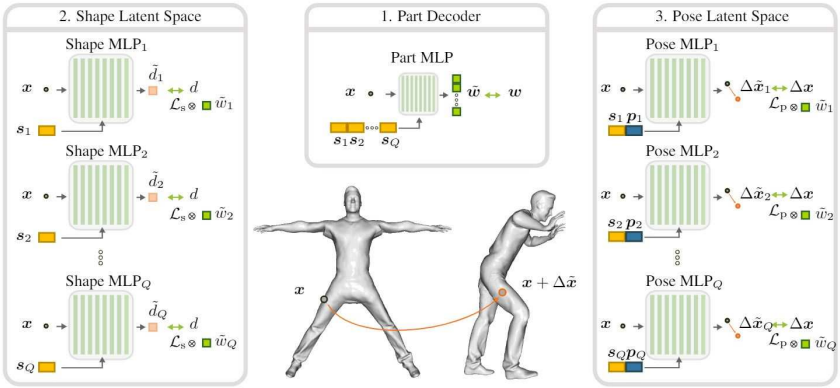

FIGURE 2. Learned parametric models can additionally model deformable shapes, disentangling shape and pose while leveraging structured part decompositions to enable robust fitting to new depth sequence observations [2].

REFERENCES

[1] A. Bokhovkin and A. Dai. Neural part priors: Learning to optimize part-based object completion in rgb-d scans. *arXiv preprint arXiv:2203.09375*, 2022.

[2] P. Palafox, N. Sarafianos, T. Tung, and A. Dai. Spams: Structured implicit parametric models. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2022.

Curvature effects in geometric statistics: empirical Fréchet mean and parallel transport accuracy

Xavier Pennec

Two fundamental tools for statistics on objects living in non-linear manifolds are the Fréchet mean and the parallel transport. We present in this talk new results to quantify the accuracy of these two fundamental tools that put forward the impact of the manifold curvature.

In the spirit of the Baker-Campbell-Hausdorff (BCH) formula for Lie groups, Gavrilov developed a coordinate free expansion of the logarithm of the composition of two Riemannian exponential maps [1]. The double exponential $\exp_x(v,u) = \exp_{\exp_x(v)}(\Pi_x^{\exp_x(v)} u)$ corresponds to a first geodesic shooting from the point x along the vector v, followed by a second geodesic shooting from $y = \exp_x(v)$ along the parallel transport $\Pi_x^y u$ of the vector u along the first geodesic. Its Riemannian logarithm at the point x has a surprisingly simple coordinate free Taylor expansion that only depends on the Riemannian curvature tensor and its covariant derivatives:

$$\begin{aligned} h_x(v,u) = \quad & \log_x(\exp_x(v,u)) = v + u + \tfrac{1}{6}R(u,v)(v+2u) \\ & + \tfrac{1}{24}(\nabla_v R)(u,v)(2v+5u) + \tfrac{1}{24}(\nabla_u R)(u,v)(v+2u) + O(5), \end{aligned}$$

where $O(5)$ is a polynomial of order 5 and more in the two variables u and v. We introduced in [2] in a companion Taylor expansion on manifolds called the neighbouring log which is measuring how the logarithm $u = log_x(y)$ changes when the foot-point x is geodesically moved in the direction of v:

$$\begin{aligned} l_x(v,w) = \quad & \Pi_{x_v}^x \log_{\exp_x(v)}(\exp_x(w)) = w - v + \tfrac{1}{6}R(w,v)(v-2w) \\ & + \tfrac{1}{24}(\nabla_v R)(w,v)(2v-3w) + \tfrac{1}{24}(\nabla_w R)(w,v)(v-2w) + O(5). \end{aligned}$$

These two tensorial and coordinate free Taylor expansions constitute a complete toolbox for the polynomial approximations of problems related to infinitesimal geodesic triangles. Moreover they are valid in general affine connection manifolds (with additional terms including torsion and its covariant derivative if they are not vanishing) and computations can be pushed to higher orders if needed.

Because we have in practice a limited number of samples, a problem in geometric statistics is to determine the properties of the empirical Fréchet mean of n IID samples in a Riemannian manifold. In sufficiently concentrated conditions, the empirical Fréchet mean exists and is unique for each sample, so that we can define its expected moments for a fixed number of samples. Using the above Taylor expansion of the Riemannian logarithm, we computed the Taylor expansion of these moments which were used in turn to compute the first and second order moments of empirical means of an IID n-sample [2]. The expected empirical mean turns out to have an unexpected non vanishing term (a bias) of order 4 in the distribution extension and in 1/n with respect to the number of samples. This bias term is a double contraction of the covariant derivative of the Riemanian

curvature with the covariance matrix, and vanishes for symmetric spaces:

$$\mathrm{Bias}(\bar{x}_n)^a = \mathbf{E}\left[\log_{\bar{x}}(\bar{x}_n)^a\right] = \tfrac{1}{6n}\left(1-\tfrac{1}{n}\right)\nabla_b R^a_{cde}\mathfrak{M}_2^{ce}\mathfrak{M}_2^{bd} + O(\epsilon^5).$$

In this formula, ϵ is the diameter of the support of the distribution. Likewise, the covariance of the empirical mean has a correction term in 1/n contracting twice the Riemannian curvature with the covariance:

$$\mathrm{Cov}(\bar{x}_n)^{ab} = \tfrac{1}{n}\big(\mathfrak{M}_2^{ab} - \tfrac{1}{3}\left(1-\tfrac{1}{n}\right)\mathfrak{M}_2^{cd}(\mathfrak{M}_2^{ae}R^b_{cde} + R^a_{cde}\mathfrak{M}_2^{be})\big) + O(\epsilon^5).$$

This term can be intrepreted as an extended Ricci curvature: in positively curved spaces, the convergence with the number of samples is slower than in Euclidean spaces while it is accelerated in negatively curved spaces. These curvature effects become important with large curvature and can drastically modify the estimation of the mean. They could partly explain the phenomenon of sticky means recently put into evidence in stratified spaces with negative curvature, and smeary means in positive curvature.

Parallel transport is a key geometric algorithm to compare local statistical models at different locations. This is the geometric equivalent of domain adaptation in machine learning. In previous works, we have build on the Schild's ladder principle to engineer a more symmetric discrete parallel transport scheme based on iterated geodesic parallelograms, called pole ladder. This scheme is surprisingly exact in only one step on symmetric spaces, which makes it quite interesting for many applications involving simple symmetric manifolds. For general manifolds, iterated Schild's and pole ladders were thought to be of first order with respect to the number of steps, similarly to other schemes based on Jacobi fields. However, the literature was lacking a real convergence performance analysis when the scheme is iterated. Using the previous Taylor expansions of geodesic triangles, we showed in [3] that pole ladder naturally converges with quadratic speed (one step being of order 4), and that Schild's ladder can be modified to perform identically even when geodesics are approximated by numerical schemes. This contrasts with Jacobi fields approximations that are bound to linear convergence. The extra computational cost of ladder methods is thus easily compensated by a drastic reduction of the number of steps needed to achieve the requested accuracy. Experiments showed that these theoretical errors are measured in prcatice with a high accuracy. This work closes several years of attempts to establish the numerical accuracy of parallel transport with discrete ladders methods.

References

[1] A. V. Gavrilov, *The double exponential map and covariant derivation.* Siberian Mathematical Journal **48(1)** (2007), 56–61.

[2] X.Pennec. *Curvature effects on the empirical mean in Riemannian and affine Manifolds: a non-asymptotic high concentration expansion in the small-sample regime* . AXIV: 1906.07418, June 2019.

[3] N. Guigui and X. Pennec. *Numerical Accuracy of Ladder Schemes for Parallel Transport on Manifolds.* Foundations of Computational Mathematics, **22** (2022), 757–790. ARXIV : 2007.07585

A Koopman Approach to Analyzing Sequence Neural Models

Omri Azencot
(joint work with Ilan Naiman)

Understanding the inner workings of predictive models is an essential requirement in many fields across science and engineering. This need is even more important nowadays with the emergence of neural networks whose visualization and interpretation is inherently challenging. Indeed, modern computational neural models often lack a commonly accepted knowledge regarding their governing mathematical principles. Consequently, while deep neural networks may achieve remarkable results on various complex tasks, explaining their underlying decision mechanisms remains a challenge. The goal of this talk is to help bridge this gap by proposing a new framework for the approximation, reasoning, and understanding of sequence neural models.

Sequence models are designed to handle time series data originating from images, text, audio, and other sources of information. One approach to analyzing sequence neural networks is through the theory and practice of dynamical systems [1, 2]. For instance, the temporal asymptotic behavior of a dynamical system can be described using the local analysis of its attractor states [3]. Similarly, recurrent models have been investigated in the neighborhood of their fixed points [4], leading to work that interprets trained RNNs for tasks such as sentiment analysis [5]. However, the local nature of these methods is a limiting factor which may lead to inconsistent results. Specifically, their approach is based on fixed-point analysis which allows to study the dynamical system in the neighborhood of a fixed-point. In contrast, our approach is global—it does not depend on a set of fixed-points, and it facilitates the exploration of the dynamics near and further away from fixed points.

Over the past few years, a family of data-driven methods was developed, allowing to analyze complex dynamical systems based on Koopman theory [6]. These methods exploit a novel observation by which nonlinear systems may be globally encoded using infinite-dimensional but *linear* Koopman operators. In practice, Koopman-based approaches are lossy as they compute a finite-dimensional approximation of the full operator. Nevertheless, it has been shown in the fluid dynamics [7, 8] and geometry processing [9, 10] communities that the dominant features of general nonlinear dynamical systems can be captured via a single matrix per system, allowing e.g., to align time series data [11]. Thus, we pose the following research question: can we design and employ a Koopman-based approach to analyze and develop a fundamental understanding of deep neural models?

Given a trained sequence neural network and a procedure to extract its hidden states, our Koopman-based method generates a moderate size matrix which faithfully describes the dynamics in the latent space. Unlike existing work, our approach is global and independent of a particular latent sample, and thus it can be virtually applied to any hidden state. A key advantage of our framework is that we can directly employ linear analysis tools on the approximate Koopman

operator to reason about the associated neural network. In particular, we show that the eigenvectors and eigenvalues of the Koopman matrix are instrumental for understanding the decision mechanisms of the model. For instance, we show in our results that the dominant eigenvectors carry crucial *semantic knowledge* related to the problem at hand. Moreover, the eigenvalues represent the *memory* of the network as they provide a timestamp for the temporal span of the respective eigenvectors. Finally and most importantly, Koopman-based methods such as ours are backed by rich theory and practice, allowing us to exploit the recent advances in Koopman inspired techniques for the purpose of developing a comprehensive understanding of sequence neural networks. Thus, the **key contribution** in this work is the novel application of Koopman-based methods for understanding sequential models, and the extraction of high-level interpretable and insightful understandings on the trained networks.

We focus our investigation on two learning tasks: sentiment analysis and electrocardiogram (ECG) classification. We will identify four eigenvectors in the sentiment analysis model whose roles are to highlight: positive words, negative words, positive pairs (e.g., "not bad"), and negative pairs. In addition, we demonstrate that the eigenvectors in the ECG classification task naturally identify dominant features in normal beat signals and encode them. Specifically, we show that four Koopman eigenvectors accurately capture the local extrema points of normal beat signals. These extrema points are fundamental in deciding whether a signal is normal or anomalous. Our results reinforce that the network indeed learns a robust representation of normal beat signals. Then, we will verify that the main components of the nonlinear network dynamics can be described using our Koopman matrices by measuring the difference in accuracy results, and the relative error in predicted states. Given the versatility of our framework and its ease of use, we advocate its utility in the analysis and understanding of neural networks, and we believe it may also affect the design and training of deep models in the future.

References

[1] K. Doya *Bifurcations of recurrent neural networks in gradient descent learning*, IEEE Transactions on neural networks (1993), Vol. 1, Issue. 75, 218.

[2] R. Pascanu, T. Mikolov, and Y. Bengio, *On the difficulty of training recurrent neural networks*, International conference on machine learning (2013), 1310–1318.

[3] Strogatz, S.H. *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering*, CRC press, (2018)

[4] D. Sussillo, O. Barak, *Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks*, Neural computation (2013), Vol.25, Issue 3, 626–649.

[5] N. Maheswaranathan, A. Williams, M. Golub, S. Ganguli, D. Sussillo *Reverse engineering recurrent networks for sentiment classification reveals line attractor dynamics*, Advances in Neural Information Processing Systems (2019), 15696–15705.

[6] B.O. Koopman, *Hamiltonian systems and transformation in Hilbert space*, Proceedings of the National Academy of Sciences (1931), May;17(5):315-8.

[7] O. Azencot, N.B. Erichson, V. Lin, M. Mahoney, *Forecasting sequential data using consistent Koopman autoencoders*, International Conference on Machine Learning (2020), Nov 21, 475–485, PMLR.

[8] I. Mezić, *Spectral properties of dynamical systems, model reduction and decompositions*, Nonlinear Dynamics (2005), Aug;41(1):309–25.
[9] A. Sharma, M. Ovsjanikov, *Weakly supervised deep functional maps for shape matching*, Advances in Neural Information Processing Systems (2020).
[10] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, L. Guibas, *Functional maps: a flexible representation of maps between shapes*, ACM Transactions on Graphics (ToG) (2012) Jul 1;31(4):1-1.
[11] O. Rahamim, R.Talmon, *Aligning sets of temporal signals with Riemannian geometry and Koopman operator*, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (2021) Jun 6, 5310–5314.

Metric Optimization in Penner Coordinates

Denis Zorin
(joint work with Ryan Capouellez)

Many parametrization and mapping-related problems in geometry processing can be viewed as metric optimization problems, i.e., computing a metric minimizing a functional and satisfying a set of constraints, such as flatness. Penner coordinates [1] are global coordinates on the space of metrics on meshes with a fixed vertex set and topology, but varying connectivity, making it homeomorphic to the Euclidean space of dimension equal to the number of edges in the mesh, without any additional constraints imposed, and reducing to logarithms of edge lengths when restricted to a fixed connectivity. These coordinates play an important role in the theory of discrete conformal maps [2, 3] enabling recent development of highly robust algorithms with convergence and solution existence guarantees for computing such maps [5, 6]. We demonstrate how Penner coordinates can be used to solve a general class of problems involving metrics, including optimization and interpolation, while retaining the key guarantees available for conformal maps [4].

References

[1] Robert C. Penner, *The decorated Teichmüller space of punctured surfaces*, Communications in Mathematical Physics, **113** (1987), 299–339.
[2] Xianfeng David Gu, Feng Luo, Jian Sun, Tianqi Wu, *A discrete uniformization theorem for polyhedral surfaces*, Journal of Differential Geometry, **109** (2018), 223–256.
[3] Boris Springborn, *Ideal hyperbolic polyhedra and discrete uniformization*, Discrete & Computational Geometry, **64**(1) (2020), 63–108.
[4] Ryan Capouellez, Denis Zorin, *Metric Optimization in Penner Coordinates*, arxiv (2022) https://arxiv.org/abs/2206.11456
[5] Mark Gillespie, Boris Springborn, Keenan Crane, *Discrete conformal equivalence of polyhedral surfaces*, ACM Transactions on Graphics, **40**(4) (2021).
[6] Marcel Campen, Ryan Capouellez, Hanxiao Shen, Leyi Zhu, Daniele Panozzo, Denis Zorin, *Efficient and robust discrete conformal equivalence with boundary*, ACM Transactions on Graphics, **40**(6) (2021), 1–16.

Shape Spaces: From geometry to biological plausibility

Laurent Younes
(joint work with Nicolas Charon)

This talk reviews several Riemannian metrics and evolution equations in the context of diffeomorphic shape analysis. After a short review of of various approaches at building Riemannian metrics on shape spaces (with a special focus on the foundations of the large deformation diffeomorphic metric mapping algorithm), the attention is turned to elastic metrics and to growth models that can be derived from them. In the latter context, a new class of metrics, involving the optimization of a growth tensor, is introduced and some of its properties are studied.

The talk is based on [1].

References

[1] Charon, N and Younes, L *Shape spaces: From geometry to biological plausibility*, arXiv:2205.01237 [math.DG]. To appear in Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging (2022).

Fast Nonlinear Vector Quantile Regression

Alex Bronstein

Quantile regression (QR) [1] is a well-known method which estimates a *conditional quantile* of a target variable Y, given covariates $\mathbf{X}$. A major limitation of QR is that it deals with a scalar-valued target variable, while many important applications require estimation of vector-valued responses. A trivial approach is to estimate conditional quantiles separately for each component of the vector-valued target. However this assumes statistical independence between targets, a very strong assumption rarely held in practice. Extending QR to high dimensional responses is not straightforward because (i) the notion of quantiles is not trivial to define for high dimensional variables, and indeed multiple definitions of multivariate quantiles exist [2]; (ii) quantile regression is performed by minimizing the *pinball loss* function [1], which is not defined for high dimensional responses.

Seminal works of [2] and [3] introduced a notion of quantiles for vector-valued variables, termed *vector quantiles*. Key to their approach is extending the notions of monotonicity and strong representation of scalar quantile functions to high dimensions, i.e.

(1) *Co-monotonicity*: $(Q_{\mathbf{Y}}(\boldsymbol{u}) - Q_{\mathbf{Y}}(\boldsymbol{u'}))^{\top}(\boldsymbol{u} - \boldsymbol{u'}) \geq 0,\ \forall\ \boldsymbol{u}, \boldsymbol{u'} \in [0,1]^d$

(2) *Strong representation*: $\mathbf{Y} = Q_{\mathbf{Y}}(\mathbf{U}),\ \mathbf{U} \sim \mathbb{U}[0,1]^d$

where $\mathbf{Y}$ is a d-dimensional random variable, and $Q_{\mathbf{Y}} : [0,1]^d \mapsto \mathbb{R}^d$ is its vector quantile function.

Moreover, [2] extended QR to vector-valued targets, which leads to *vector quantile regression* (VQR). This enables estimation of conditional vector quantile functions $Q_{\mathbf{Y}|\mathbf{X}}$ from samples drawn from $P_{(\mathbf{X},\mathbf{Y})}$, where $\mathbf{Y}$ is a d-dimensional target variable and $\mathbf{X}$ are k-dimensional covariates. They show that a conditional quantile function $Q_{\mathbf{Y}|\mathbf{X}}$ which obeys co-monotonicity (1) and strong representation (2) exists and is unique. Under the assumption of a linear specification $Q_{\mathbf{Y}|\mathbf{X}}(\boldsymbol{u};\boldsymbol{x}) = \boldsymbol{B}(\boldsymbol{u})^\top \boldsymbol{x} + \boldsymbol{a}(\boldsymbol{u})$, they formulate VQR as an optimal transport problem between the measures of $\mathbf{Y}|\mathbf{X}$ and $\mathbf{U}$, with the additional mean-independence constraint $\mathbb{E}[\mathbf{U}|\mathbf{X}] = \mathbb{E}[\mathbf{X}]$. 1 provides a visualization of these notions for a two-dimensional target variable. The primal formulation of this problem is large-scale linear program and is thus intractable for modestly-sized problems. A relaxed dual formulation which is amenable to gradient-based solvers exists [4], but results in violations of co-monotonicity.

The first goal of our work is to address the following limitations of [2, 4]: (i) the linear specification assumption on the conditional quantile function, and (ii) the violation of co-monotonicity when solving the inexact formulation of the VQR problem. The second goal of this work is to make VQR an accessible tool for off-the-shelf usage on large-scale high-dimensional datasets. Currently there is no available solver for the relaxed dual formulation of VQR, which is necessary in order to scale reasonably with problem size. Below we list our contributions.

Nonlinear VQR. To address the limitation of linear specification, we propose *nonlinear vector quantile regression*. The key idea is fit a nonlinear embedding function of the input features jointly with the regression coefficients. We demonstrate that nonlinear VQR can model complex conditional quantile functions significantly better than linear VQR.

Vector monotone rearrangement (VMR). We propose VMR, which resolves the co-monotonicity violations in estimated conditional vector quantile functions. We solve an optimal transport problem to *rearrange* the vector quantiles such that they satisfy co-monotonicity. It can be viewed as a vector extension to the monotone rearrangement originally proposed by [5] for scalar quantiles.

Scalable VQR. We introduce highly-scalable solvers for linear and nonlinear VQR. Our approach, inspired by [6] and [4], relies on solving the relaxed dual formulation of the VQR problem. We propose stochastic-gradient-based solvers which maintain a constant memory footprint regardless of problem size. We demonstrate that our solvers can scale to millions of samples and thousands of quantile levels and allow for GPU-acceleration.

Open-source software package. We release a feature-rich, well-tested software implementing estimation of vector quantiles, vector ranks, vector quantile contours, linear and nonlinear VQR, and VMR. To the best of our knowledge, this would be the first publicly available tool for estimating conditional vector quantile functions at scale.

The full paper can be found at https://arxiv.org/abs/2205.14977. Our open-source package `vqr` is available at https://github.com/vistalab-technion/vqr.

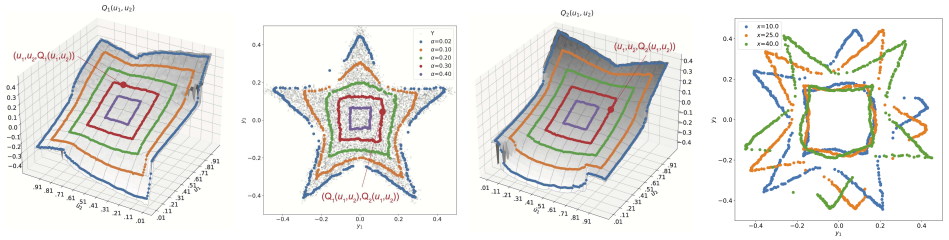

(A) (B)

FIGURE 1. Illustration of vector quantiles of a 2-dimensional star-shaped distribution, where $T = 50$ quantile levels were estimated in each dimension. (A) Vector quantiles (colored dots) are overlaid on their data distribution (middle). Different colors correspond to α-contours, each containing $100 \cdot (1 - 2\alpha)^2$ percent of the data, a generalization of confidence intervals for vector-valued variables. The vector quantile function (VQF) $Q_{\mathbf{Y}}(\boldsymbol{u}) = [Q_1(\boldsymbol{u}), Q_2(\boldsymbol{u})]^\top$ is co-monotonic with $\boldsymbol{u} = (u_1, u_2)$. The components Q_1, Q_2 of the VQF are shown as surfaces (left, right) with the corresponding vector quantiles overlaid. On Q_1, increasing u_1 for a fixed u_2 produces a monotonically increasing curve, and vice versa for Q_2. (B) Conditional vector quantile functions (CVQFs) for a joint distribution of $(\mathbf{X}, \mathbf{Y})$ where $\mathbf{Y}|\mathbf{X} = \boldsymbol{x}$ has a star-shaped distribution rotated by $\boldsymbol{x}$ degrees. The CVQF of $\mathbf{Y}|\mathbf{X} = \boldsymbol{x}$ changes nonlinearly with the covariates $\boldsymbol{x}$, while e.g. $\mathbb{E}[\mathbf{Y}|\mathbf{X}]$ remains the same. This demonstrates the challenge of estimating CVQFs from samples of the joint distribution. Two α-contours are depicted for each $\boldsymbol{x}$. Digital zoom-in recommended.

REFERENCES

[1] R. Koenker and G. Bassett. *Regression Quantiles*, Econometrica, 46(1):33, (1978).

[2] G. Carlier, V. Chernozhukov, and A. Galichon. *Vector quantile regression: An optimal transport approach*, Annals of Statistics, 44(3):1165–1192, (2016).

[3] V. Chernozhukov, A. Galichon, M. Hallin, and M. Henry. *Monge-Kantorovich depth, quantiles, ranks and signs*, The Annals of Statistics, 45(1):223–256, (2017).

[4] G. Carlier, V. Chernozhukov, G. De Bie, and A. Galichon. *Vector quantile regression and optimal transport, from theory to numerics*, Empirical Economics, (2020).

[5] V. Chernozhukov, I. Fernández-Val, and A. Galichon. *Quantile and probability curves without crossing*, Econometrica, 78(3):1093–1125, (2010).

[6] A. Genevay, M. Cuturi, G. Peyré, and F. Bach. *Stochastic optimization for large-scale optimal transport*, Advances in neural information processing systems, 29, (2016).

Geometric Measure Theory and Kelvin Geometry for Convexifying and Compactifying Computational Problems

Albert Chern

(joint work with Stephanie Wang, Sina Nabizadeh and Ravi Ramamoorthi)

We use geometric techniques to approach two classical computational challenges: optimizing non-convex variational problems and solving partial differential equations (PDEs) on unbounded domains. These problems are fundamental in shape synthesis and physical simulations. In our work, we use the ideas from Geometric Measure Theory and Felix Klein's geometric philosophy to re-examine these computational challenges. Surprisingly, many problems associated to non-convexity or non-compactness are removed just by suitable changes of variables.

1. Minimal Surface Problem

A classical computational differential geometry problem is the *Plateau problem*:

> **Problem I** *In a space M ($\mathbb{R}^3$ or a Riemannian manifold), extend a given boundary curve $\Gamma \subset M$ into a surface $\Sigma \subset M$ with minimal area.*

The traditional approach is based on (regularized) mean curvature flows on triangle meshes, which are (Sobolev) gradient flows of the area functional. However, this is a non-convex problem. Directly minimizing the area can lead to being stuck at local minima. Mean curvature flows can also develop finite time blow up. Initializing a correct surface is topologically and combinatorially challenging.

Geometric measure theory provides new ways of approaching this problem. In the theory, curves and surfaces are represented as differential forms. In an n-dimensional manifold, for each oriented k-dimensional surface Σ we construct an $(n-k)$-form $\delta_\Sigma \in \Omega^{n-k}(M)$, called a Dirac δ-form, with the defining property

$$\int_M \omega \wedge \delta_\Sigma = \int_\Sigma \omega \quad \text{for all continuous } \omega \in \Omega^k(M).$$

Geometric operations become familiar linear algebraic operations:

$$\delta_{A\cap B} = \delta_A \wedge \delta_B, \quad \delta_{\partial A} = (-1)^{n-k+1} d\delta_A$$

where $\dim(A) = k$. The area of a surface Σ becomes the L^1 norm (mass norm) for the δ-form:

$$\text{Area}(\Sigma) = \int_M |\delta_\Sigma|.$$

In terms of differential forms, the oriented Plateau problem becomes

$$\textbf{Problem II} \quad \underset{\Sigma}{\text{minimize}} \int_M |\delta_\Sigma| \quad \text{subject to} \quad d\delta_\Sigma = \delta_\Gamma$$

which we relax to a convex L^1 optimization problem with a linear constraint:

$$\textbf{Problem III} \quad \underset{\eta\in\Omega^1(M)}{\text{minimize}} \int_M |\eta| \quad \text{subject to} \quad d\eta = \delta_\Gamma.$$

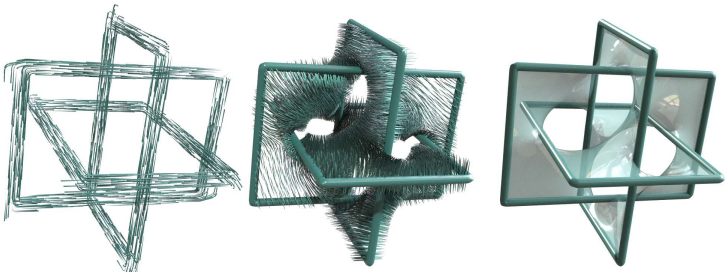

FIGURE 1. Plateau problem computed with differential forms. Given a boundary curve represented by a Dirac δ 2-form (left), we find the minimal surface as the L^1 minimizer subject to the exterior derivative constraint (middle) followed by a Poisson surface reconstruction (right).

Problem II and Problem III are equivalent due to the following estimate: Given any $\eta \in \Omega^1(M)$, $d\eta = \delta_\Gamma$, there exists a surface Σ, $d\delta_\Sigma = \delta_\Gamma$, such that $\int_M |\delta_\Sigma| \leq \int_M |\eta|$. To see this estimate, observe that $d\eta = \delta_\Gamma$ implies $\eta = d\theta$ for some $\mathbb{S}^1$-valued function θ branched at Γ, and use the Fubini theorem (co-area formula) $\int_M |\eta| = \frac{1}{2\pi} \int_{\mathbb{S}^1} \int_M |\delta_{\theta^{-1}(a)}| \, da \geq \int_M |\delta_{\theta^{-1}(a^*)}|$ for $a^* \in \mathbb{S}^1$ that yields the minimal level set of θ.

As such, the seemingly non-convex minimal surface problem is revealed as a convex one in disguise. We solve the convex Problem III numerically by standard L^1 optimization techniques known by ADMM, proximal gradient descent, Bregman iterations, etc [1]. A typical result is shown in Figure 1. We also solve Problem III by representing functions and forms by deep neural fields [2].

While the least squares problem for differential forms under linear constraints on their d yields the Hodge decomposition as the optimality condition, we obtain a generalization of the Hodge decomposition for our L^1 problem.

Every $\eta \in \Omega^k$ can be decomposed into

$$\eta = d\phi + \xi + \zeta$$

where $(\xi + \zeta)$ is a homologically constrained minimal surface, and ξ is a globally minimal surface.

2. Infinite Domain Problem

Many physical simulation problems take place on an open unbounded space. The non-compactness of the domain poses additional challenges, since we can no longer directly discretize the entire domain using a regular grid. Standard numerical approaches rely on coordinate mapping or domain truncation, yielding coordinate singularity or artifacts on the truncation boundary. We describe a general Kelvin transformation technique, which maps the infinite domain to a bounded one without creating singularities. The method is made possible by factoring out

FIGURE 2. A 2D slice of a numerical solution to the 3D Helmholtz equation on the infinite exterior domain of the Bunny. The solution is computed by a generalized Kelvin transform, consisting of inverting the domain and factoring out the singularity at infinity.

an asymptotic of the singularity induced by the coordinate stretching. The resulting transformation of functions can be understood as the natural transformation for fractional densities in geometric measure theory. In the viewpoint of Klein's Erlangen Program, the analysis reveals a "Kelvin Geometry," where objects are functions subject to Kelvin transforms, leaving the PDE of interest invariant. The key to solving the infinite domain problem is to recognize that the boundedness quality of the domain is not a geometrically invariant notion under Kelvin Geometry. Therefore, we can transform the infinite domain problem into a compact one without sacrificing numerical accuracy.

Our Kelvin geometric technique has been applied to the Poisson equation and Helmholtz equation [3] (Figure 2). On going work includes the Kelvin transform of the wave equation and the eikonal equation.

REFERENCES

[1] S. Wang and A. Chern, *Computing Minimal Surfaces with Differential Forms*, ACM Transactions on Graphics (2021), Vol. 40 Issue 4, August 2021, 113:1–113:14.

[2] D. Palmer, D. Smirnov, S. Wang, A. Chern and J. Solomon, *Deep Currents: Learning Implicit Representations of Shapes with Boundaries*, Conference on Computer Vision and Pattern Recognition (CVPR), 2022.

[3] M. S. Nabizadeh, R. Ramamoorthi and A. Chern, *Kelvin transformations for simulations on infinite domains*, ACM Transactions on Graphics (2021), Vol. 40 Issue 4, August 2021, 97:1–97:15.

Guaranteed Queries on General Neural Implicit Surfaces via Range Analysis

Nicholas Sharp
(joint work with Alec Jacobson)

Implicit Neural Surface Representations. Implicit surface representations encode an embedded 3D surface as the level set of some function $f : \mathbb{R}^3 \to \mathbb{R}$. By convention, the locations x where $f(x) < 0$ are inside the shape, and $f(x) > 0$ are outside of the shape. In *neural* implicit surfaces, the function f is a neural network f_θ with parameters θ, which can be fit to encode a surface of interest. The networks have proven recently useful for a variety of tasks in computer graphics and vision [3], in part because many tasks involving generating a potentially-complex surface can be approached via simple gradient-based optimization on f_θ. This abstract considers the computational geometry of such a representation: when we use neural networks as a function space to encode 3D shapes, how do we computationally evaluate queries against the shape such as finding intersections between a ray and the surface, or testing whether two such shapes intersect. An extended treatment of this work can be found in [1].

The primary past approach for such queries is to attempt to fit the neural network f_θ such that it has a signed distance function (SDF) property, meaning that it satisfies the Eikonal equation $|\nabla f_\theta| = 1$ almost everywhere and the magnitude $|f_\theta|$ gives the distance to the implicit surface. If f_θ is an SDF, then it is not too difficult to evaluate geometric queries. However, there are many disadvantages to the SDF representation: a neural network only ever approximately encodes an SDF, and the SDF property holds only at the conclusion of training, precluding the use of queries in loss functions. Instead, we will seek an approach which applies to *general* neural implicit surfaces, without any particular assumptions on the function encoded f_θ.

Range Analysis. Our key tool to evaluate geometric queries on general neural implicit surfaces is a numerical technique called *range analysis* [2]. Given some range of inputs to a function, range analysis automatically computes bounds on the output of the function over that range. These bounds are necessarily valid, in the sense that they contain the true output range of the function, but they are not necessarily tight—various range analysis schemes have been developed with different tradeoffs off tight bounds vs. computational cost. The most common form of range analysis is *interval arithmetic*, but we find that when applied to typical neural implicit surface networks, the bounds from interval arithmetic are so loose that they are not useful in practice. Instead, we leverage a 1^{st} order generalization called *affine arithmetic*, which gives much tighter bounds for neural networks (Figure 1).

Geometric Queries. Equipped with range analysis of neural networks, we can analyze regions of space, and potentially determine that the value of the implicit function is strictly positive or negative on that region. Whenever a region cannot

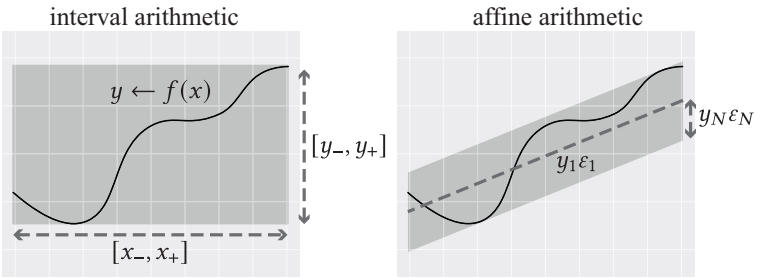

FIGURE 1. We empirically find affine arithmetic to be much more effective than interval arithmetic for analyzing neural implicit surfaces, likely because it is exact within the affine layers of neural networks.

be classified as such, either because the range analysis bounds are too loose, or because it actually intersects the level set, we subdivide the region into two smaller regions and recurse. This basic recipe becomes our building block to construct higher-level geometric queries, including ray casting, constructing bounding k-d trees, and intersection testing. These queries allow efficient computational-geometry-style queries on existing neural implicit surface representations for the first time, opening the door to a wide range of applications in surface processing.

REFERENCES

[1] N. Sharp & A. Jacobson, *Spelunking the Deep: Guaranteed Queries on General Neural Implicit Surfaces via Range Analysis*, ACM Trans. Graph. **41** (2022)
[2] J. Stolfi & L. De Figueiredo, *Self-validated numerical methods and applications*, 21st Brazilian Mathematics Colloquium **5** (1997)
[3] T. Davies & D. Nowrouzezahrai & A. Jacobson, *On the Effectiveness of Weight-Encoded Neural Implicit 3D Shapes*, (2020)

Computational optimal transport: mature tools and open problems

JEAN FEYDY

(joint work with Minh-Hieu Do, Olga Mula-Hernandez, Marc Niethammer, Gabriel Peyré, Bernhard Schmitzer, Thibault Séjourné, Zhengyang Shen, Anna Song, Alain Trouvé, François-Xavier Vialard)

Optimal transport is a fundamental tool to deal with discrete and continuous distributions of points [1, 2]. We can understand it either as a generalization of sorting to spaces of dimension $D > 1$, or as a nearest neighbor projection under a mass preservation constraint. Over the last decade, a sustained research effort on numerical foundations has led to a x1,000 speed-up for most transport-related computations. This has opened up a wide range of research directions in geometric data analysis, machine learning and computer graphics.

This talk will discuss the consequences of these game-changing numerical advances from a user's perspective. We will focus on:

(1) Mature libraries and software tools that can be used as of 2022, with a clear picture of the current state-of-the-art [3, 4, 5, 6, 7, 8].
(2) New ranges of applications in 3D shape analysis, with a focus on population analysis [9] and point cloud registration [10].
(3) Open problems that remain to be solved by experts in the field.

References

[1] G. Peyre and M. Cuturi. *Computational optimal transport*, https://optimaltransport.github.io/book/ (2018).
[2] J. Feydy. *Geometric data analysis, beyond convolutions*, PhD thesis, Université Paris-Saclay (2020), https://www.jeanfeydy.com/geometric data analysis.pdf.
[3] R. Flamary et al, *Python Optimal Transport*, https://pythonot.github.io/.
[4] M. Cuturi et al. *Optimal Transport Tools*, https://ott-jax.readthedocs.io/.
[5] Alice INRIA team. *Geogram software*, http://alice.loria.fr/software/geogram.
[6] B. Schmitzer, *MultiScale-OT toolbox*, https://bernhard-schmitzer.github.io/MultiScaleOT.
[7] Q. Mérigot et al., *PyMongeAmpere*, https://github.com/mrgt/PyMongeAmpere.
[8] J. Feydy, *GeomLoss*, https://www.kernel-operations.io/geomloss/.
[9] A. Song. *Generation of tubular and membranous shape textures with curvature functionals*, SIAM Journal of Mathematical Imaging and Vision (2022).
[10] Z. Shen et al. *Accurate point cloud registration with robust optimal transport*, NeurIPS 2021.

Adjoint mismatch

Dirk A. Lorenz

Optimization methods in mathematical imaging frequently need to evaluate linear operators and their adjoint: The minimization of an objective of the form

$$\tfrac{1}{2}\|Ax - b\|^2 + \alpha R(x)$$

in $x \in X$ over a Hilbert space X for some given linear operator A, some well behaved (e.g. convex, lower semi-continuous and coercive) function R, and some given data b can be done, for example, by the proximal gradient iteration $x^{k+1} = \operatorname{prox}_{\tau\alpha R}(x^k - \tau A^*(Ax^k - b))$, while for objectives of the form $F(Ax) + G(x)$ for equally well behaved F and G one can, for example, use the Chambolle-Pock iteration which iterates $x^{k+1} = \operatorname{prox}_{\tau G}(x^k - \tau A^* y^k)$, $y^{k+1} = \operatorname{prox}_{\sigma F^*}(y^k + \sigma A(2x^{k+1} - x^k))$ (where F^* is the convex conjugate of F). The former method converges if $0 < \tau < 2/\|A\|^2$ while the latter needs $0 < \sigma\tau < 1/\|A\|^2$ [1].

Examples for linear maps are derivative operators, interpolation operators, convolutions or more complicated "forward operators" like the Radon transform or solution operators for differential equations. For many operators, the adjoint is readily available, but in other cases, the forward operator is only available through a numerical method and the adjoint operator has to implemented seperately. A discretization of the adjoint is not necessarily the adjoint of the discretization. Sometimes people even use a mismatched adjoint on purpose, e.g. if the replacement of any application of A^* by some V^* is much faster to compute or produces

output which seems to be more plausible. The former happens, for example, for the Radon transform where there are faster and slower methods for the adjoint (which is known as backprojection), and the latter can occur, for example, in superresolution where the forward operator does sampling and the adjoint does some interpolation.

Up to now there are only a few results on convergence under mismatch available. In [3] and [4] the authors investigate the Landweber method under mismatch and [6] treats the randomized Kaczmarz iteration. The work [2] analyzes the effect of mismatch and [7] investigates the Chambolle-Pock method. For the latter, we could show that the method still converges if both G and F^* are strongly convex and the product of the constants of strong convexity is larger than a constant that depends on the quantity $\|A - V\|$.

The fact that sometimes linear operators A are available only via black-box implementations, i.e. one can give some input x to a program and gets back a vector Ax, and that the respective adjoint are available via other black-box implementations which give V^*y for vectors y raises some other questions: How can one compute $\|A\|$ (which is needed in the bounds for the step-sizes) if only black boxes are available and there is no access to an adjoint? Moreover, we assume that probabilistic sketching methods based on randomized numerical linear algebra are not applicable due to the hardware constraints that do not permit to store enough vectors to form a large enough sketched matrix. An even more difficult question is, how to compute $\|A - V\|$ if only a black-box for A and another one for V^* is available.

For the first question we propose to use *projected stochastic gradient ascent* for the problem

$$\max_{\|v\|=1} \|Av\|^2. \tag{1}$$

The gradient of the objective is A^TAv, but under our assumption, we can not compute it. However, we can get an unbiased estimate by using a trick: Let x be a random vector of appropriate size such that $\mathbb{E}(xx^T) = I$. Then it holds that

$$\mathbb{E}\left[(Av)^T(Ax)x\right] = \mathbb{E}\left[xx^TA^TAv\right] = A^TAv.$$

Thus, the iteration

$$v^{k+\frac{1}{2}} = v^k + \tau_k(Av^k)^T(Ax)x$$

$$v^{k+1} = \frac{v^{k+\frac{1}{2}}}{\|v^{k+\frac{1}{2}}\|}$$

has a good chance to converge to a solution of (1) for appropriate step-sizes τ_k. While it can be shown that there are simple methods to choose step-sizes that lead to a guaranteed increase of $\|Av^k\|$ in every iteration, a full convergence prove is still missing, but the method can be observed to converge in practice. The idea using random vectors x with covariance $\mathbb{E}(xx^T) = I$ is also used in the well-know trace estimator due to Hutchinson [5]. For the computation $\|A - V\|$ one can use

the characterization

$$\|A - V\| = \max_{\|u\|=\|v\|=1} \left[u^T(A - V)v\right] = \max_{\|u\|=\|v\|=1} \left[u^T(Av) - (V^T u)^T v\right]$$

and perform projected stochastic gradient ascent in a similar fashion.

References

[1] H. H. Bauschke and P. L. Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*, CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC, Springer, Cham, second ed., 2017.

[2] E. Chouzenoux, J.-C. Pesquet, C. Riddell, M. Savanier, and Y. Trousset. Convergence of proximal gradient algorithm in the presence of adjoint mismatch. *Inverse Problems*, 37(6):Paper No. 065009, 29, 2021.

[3] Y. Dong, P. Hansen, M. Hochstenbach, and N. A. B. Riis. Fixing nonconvergence of algebraic iterative reconstruction with an unmatched backprojector. *SIAM J. Sci. Comput.*, 41:A1822–A1839, 2019.

[4] T. Elfving and P. Hansen. Unmatched projector/backprojector pairs: Perturbation and convergence analysis. *SIAM J. Sci. Comput.*, 40, 2018.

[5] M.F. Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics - Simulation and Computation*, 19(2):433–450, jan 1990.

[6] D. A. Lorenz, S. D. Rose, and F. Schöpfer. The randomized Kaczmarz method with mismatched adjoint. *BIT Numerical Mathematics*, 58:1079–1098, 2018.

[7] D. A. Lorenz and Felix Schneppe. Chambolle-Pock's Primal-Dual Method with Mismatched Adjoint. Submitted for publication, arxiv.org/abs/2201.04928, 2022

Linear and Quadratic Shape Functions for Polygons and Polyhedra

Mario Botsch

(joint work with Astrid Bunge, Philipp Herholz, Olga Sorkine-Hornung, Marc Alexa, Misha Kazhdan)

Solving PDEs on surface meshes or volume meshes is a central goal in computer graphics and geometry processing. The continuous surfaces/volumes are typically discretized with either triangles/tetrahedra or quadrangles/hexahedra, and then the finite element method (FEM) is employed using linear P1 shape functions on the former or bilinear/trilinear Q1 shape functions on the latter. In several application scenarios, however, the meshes consist of *arbitrary* polygons or polyhedra, and we want to solve PDEs directly on these general meshes without having to remesh to standard elements.

One possibility is to use generalized barycentric coordinates [6] as shape functions, as done, e.g., by Wicke et al. [9] using mean value coordinates or by Martin et al. [7] using harmonic coordinates. The drawback of these approaches is that evaluation, derivation, and numerical integration is rather complex for the generalized barycentric shape functions. Alexa and Wardetzky [1] and de Goes et al. [5] proposed discrete differential operators for general polygon meshes based on discrete exterior calculus. While being much simpler to use than the generalized barycentric coordinates, these methods include a hyper-parameter that has to be

tuned to achieve optimal results. We propose a simple method for deriving linear and quadratic shape functions on polygon surface meshes and polyhedral volume meshes, which inherit the beneficial numerical properties of linear/quadratic shape functions for P1/P2 elements, reproduce them on triangles/tetrahedra, and thus generalize them to arbitrary polygons/polyhedra.

Linear Shape Functions for Polygon Meshes

The main idea is to (virtually) insert a center vertex into each polygon, thereby splitting it into a fan of triangles. This operation turns the general polygon mesh into a pure triangle mesh, on which we can use the standard linear P1 shape functions. For a polygon with n vertices $\mathbf{x}_1, \dots, \mathbf{x}_n$, the center vertex $\mathbf{x}_0$ is represented as an affine combination using the weights $\mathbf{w} = (w_1, \dots, w_n)^\mathsf{T}$:

$$\mathbf{x}_0(\mathbf{w}) = \sum_{j=1}^{n} w_j \mathbf{x}_j \quad \text{with} \quad \sum_{j=1}^{n} w_j = 1. \tag{1}$$

Its position is chosen as the minimizer of squared triangle areas summed over the triangles incident to the center vertex, being optimized directly in terms of the affine weights:

$$\min_{\mathbf{w}} \sum_{j} \operatorname{area}\left(\mathbf{x}_j, \mathbf{x}_{j+1}, \mathbf{x}_0(\mathbf{w})\right)^2 \quad \text{such that} \quad \sum_{j} w_j = 1. \tag{2}$$

As (2) is underdetermined, we choose the solution with minimal L2 norm $\|\mathbf{w}\|$, which can be computed through a simple linear system solve.

The resulting weights define the shape functions ϕ_i on the original polygon mesh in terms of the shape functions ψ_j on the refined triangle mesh:

$$\phi_i = \psi_i + w_i \psi_0. \tag{3}$$

Putting the weights of each polygon into a global prolongation matrix $\mathbf{P}$ (with $\mathbf{P}^\mathsf{T}$ begin the corresponding restriction matrix) allows to define the discrete Laplacian $\mathbf{L}$ and mass matrix $\mathbf{M}$ of the *polygon mesh* through the standard cotangent Laplacian $\mathbf{L}_{\text{tri}}$ and mass matrix $\mathbf{M}_{\text{tri}}$ of the refined *triangle mesh* as

$$\mathbf{L} = \mathbf{P}^\mathsf{T} \mathbf{L}_{\text{tri}} \mathbf{P} \quad \text{and} \quad \mathbf{M} = \mathbf{P}^\mathsf{T} \mathbf{M}_{\text{tri}} \mathbf{P}. \tag{4}$$

Through the prolongation and restriction matrices the simplicial refinement can be completely hidden from the user and can be efficiently implemented within the matrix assembly stage. As shown in Bunge et al. [2], the resulting differential operators on polygon meshes inherit the properties of standard P1 elements and provide the expected quadratic convergence rate under element refinement.

Linear Shape Functions for Polyhedral Meshes

The virtual refinement can straightforwardly be extended to polyhedral meshes by using a two-step prolongation process: First, each polygon is split into a triangle fan by inserting the center vertex that minimizes the sum of squared triangle areas, leading to a prolongation matrix $\mathbf{P}_{\text{face}}$. Second, each polyhedron (now having triangle faces) is split into a fan of tetrahedra by inserting the center vertex that

minimizes the sum of squared tetrahedron volumes, leading to a prolongation matrix $\mathbf{P}_{\text{cell}}$. Both prolongation weights can be determined through simple linear systems. The combined prolongation matrix $\mathbf{P} = \mathbf{P}_{\text{cell}} \cdot \mathbf{P}_{\text{face}}$ can then be used to "sandwhich" discrete differential operators computed on the refined tetrahedral mesh, analogously to (4). This idea was first proposed in [3] and leads to differential operators that inherit the properties of linear P1 shape functions on tetrahedral meshes, in particular the quadratic convergence rate under element refinement.

Quadratic Shape Functions for Polygon Meshes

Compared to their linear counterparts, higher-order shape functions feature faster convergence rate and decrease the well-known locking artifacts when simulating elastic materials with high Poisson ratio. Quadratic shape functions turned out to be a good choice in the large scale analysis of Schneider et al. [8].

When generalizing the virtual refinement idea to quadratic shape functions on polygon meshes, the virtually refined triangle meshes will have degrees of freedom both at the vertices and at the edge midpoints. Inserting a *virtual center vertex* into a polygon with n vertices therefore also leads to additional n *virtual edge vertices* on the edges connecting the polygon vertices to the virtual center vertex. The prolongation weights are now responsible for distributing the values at the $n+1$ virtual vertices to the n vertex nodes and n edge nodes of the original polygon. When we denote by $\mathcal{C}$ the nodes of the polygon mesh and by $\mathcal{K}$ the nodes of the refined triangle mesh, the polygon shape functions become

$$\phi_i = \psi_i + \sum_{j \in \mathcal{K}} w_{ij} \psi_j \quad \text{with} \sum_{j \in \mathcal{K}} w_{ij} = 1 \quad \text{for } i \in \mathcal{C}. \tag{5}$$

We optimize w_{ij} to yield as-smooth-as-possible piecewise quadratic shape functions, by minimizing the gradient jump across internal virtual edges incident to the center vertex. The resulting shape functions inherit the properties of P2 shape functions and can be shown to reproduce quadratic P2 elements on tetrahedra, to have quadratic precision, and to yield the desired cubic convergence rate [4].

Quadratic Shape Functions for Polyhedral Meshes

Combining the two-level prolongation of linear polyhedral shape functions with the smoothness optimization of quadratic polygon shape functions allows to derive quadratic shape functions for arbitrary polyhedra. The first prolongation matrix $\mathbf{P}_{\text{face}}$ is derived by minimizing cross-edge gradient jumps as in (5). For the second prolongation matrix $\mathbf{P}_{\text{cell}}$ one optimizes smoothness by penalizing the gradient jump across the virtual triangles incident to the cell's center vertex. The quadratic polyhedral shape functions provide the same beneficial properties as the in the surface case. See [4] for more details and comparisons to other approaches.

References

[1] Marc Alexa, Max Wardetzky, *Discrete Laplacians on General Polygonal Meshes*, ACM Transactions on Graphics 30:4 (2011), 102:1–102:10.

[2] Astrid Bunge, Philipp Herholz, Misha Kazhdan, Mario Botsch, *Polygon Laplacian Made Simple*, Computer Graphics Forum 39:2 (2020), 303–313.

[3] Astrid Bunge, Mario Botsch, Marc Alexa, *The Diamond Laplace for Polygonal and Polyhedral Meshes*, Computer Graphics Forum 40:5 (2021), 217–230.

[4] Astrid Bunge, Philipp Herholz, Olga Sorkine-Hornung, Mario Botsch, Michael Kazhdan, *Variational Quadratic Shape Functions for Polygons and Polyhedra*, ACM Transaction on Graphics 41:4 (2022), 54:1–54:14.

[5] Fernando de Goes, Andrew Butts, Mathieu Desbrun, *Discrete Differential Operators on Polygonal Meshes*, ACM Transactions on Graphics 39:4 (2020), 110:1–110:14.

[6] Kai Hormann, N. Sukumar, *Generalized Barycentric Coordinates in Computer Graphics and Computational Mechanics*, Taylor & Francis, 2017.

[7] Sebastian Martin, Peter Kaufmann, Mario Botsch, Martin Wicke, Markus Gross, *Polyhedral Finite Elements Using Harmonic Basis Functions*, Computer Graphics Forum 27:5 (2008), 1521–1529.

[8] Teseo Schneider, Yixin Hu, Xifeng Gao, Jeremie Dumas, Denis Zorin, Daniele Panozzo, Daniele, *A Large Scale Comparison of Tetrahedral and Hexahedral Elements for Solving Elliptic PDEs with the Finite Element Method*, ACM Transactions on Graphics 41:3 (2022), 23:1–23:14.

[9] Martin Wicke, Mario Botsch, Markus Gross, *A Finite Element Method on Convex Polyhedra*, Computer Graphics Forum 26:3 (2007), 355–364.

Near optimal statistical estimation of smooth optimal transport potentials

François-Xavier Vialard

(joint work with Adrien Vacher, Boris Muzellec, Alessandro Rudi, and Francis Bach)

It was recently shown in [5] that under smoothness conditions, the squared Wasserstein distance between two distributions can be efficiently computed with appealing statistical error upper bounds, avoiding the curse of dimension in the rate of convergence. However, rather than the distance itself, the object of interest for applications such as generative modelling or other downstream applications is the underlying optimal transport map. The optimal transport map being the gradient of the optimal potential in the classical dual formulation of optimal transport, computational and statistical guarantees need to be obtained on such a quantity. In this paper, we propose the first tractable algorithm for which the statistical L^2 error on the maps nearly matches the existing minimax lower-bounds for smooth map estimation. Our method is based on solving the semi-dual formulation of optimal transport with an infinite-dimensional sum-of-squares reformulation, and leads to an algorithm which has dimension-free polynomial rates in the number of samples, with potentially exponentially dimension-dependent constants. Therefore, for small dimensions but greater than 2,3 or 4, the theoretical constant is still tractable. This work settles the question of whether the smoothness of optimal

solutions can be taken advantage of from a computational and statistical point of view in order to estimate the optimal transport potentials.

Optimal transport is a linear optimization problem on a space of functions under an inequality constraint. A particular case of optimizing on non-negative functions can be found in polynomial optimization [3]. A typical problem of interest is the optimization of a polynomial function on a set constrained by polynomial inequalities. Leveraging, when available, a representation theorem such as Putinar's Positivestellensatz, the optimization problem can be reduced to a hierarchy of SDP problems, see [3]. Recent works develop this idea in Reproducing Kernel Hilbert Spaces [4], for non-convex optimization in [1] where the authors recast, as is standard, the problem of minimization of a function $f : D \subset \mathbb{R}^d \to \mathbb{R}$ defined on a domain D as a convex optimization problem, $\max c$ under the inequality constraint $c \leq f(x)$ for every $x \in D$. Obviously, this problem is computationally intractable in general and they propose to solve it under structural assumptions on $f(x) - c = \frac{1}{2}\langle \phi(x), A\phi(x)\rangle$ for a positive self-adjoint operator $A : \mathcal{H} \mapsto \mathcal{H}$ where $\mathcal{H}$ is a RKHS. The value of this new optimization problem is a priori less than the minimum value of f but it does coincide under the assumption that

$$f = \text{cste} + \frac{1}{2}\langle \phi(x), A_* \phi(x)\rangle \,, \tag{1}$$

for some A_*. The key point here is a representation result stating that a fairly large space of smooth functions (to be considered for optimization) can be represented by a sum of squares in RKHS, as in Equation (1). It turns out that optimal transport enjoys such a similar structure that can be used in a similar way as shown in [5].

In [5], an estimator that is computationally feasible using SDP programming is proposed. In order to get near-optimal estimation of the optimal transport potentials, this method needs to be refined in order to take advantage of the gain in strong convexity of the semi-dual functional associated with optimal transport. Indeed, the usual dual formulation of optimal transport presents no strong convexity since it is a linear optimization problem. However, optimizing on one of the potentials turns it into a more strongly convex functional, at least on smooth potentials. This strong convexity is exploited, together with the smoothness assumptions to refine a statistical estimator of the potentials. For further details, we refer to [6].

References

[1] Alessandro Rudi, Ulysse Marteau-Ferey, and Francis Bach. Finding Global Minima via Kernel Approximations. *arXiv e-prints*, page arXiv:2012.11978, December 2020.

[2] J. Andrew (Drew) Bagnell and Amir massoud Farahmand. Learning positive functions in a hilbert space, December 2015.

[3] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.

[4] Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Non-parametric Models for Non-negative Functions. *arXiv e-prints*, page arXiv:2007.03926, July 2020.

[5] Adrien Vacher and Boris Muzellec and Alessandro Rudi, and Francis Bach and François-Xavier Vialard. A Dimension-free Computational Upper-bound for Smooth Optimal Transport Estimation. *arXiv e-prints*, page arXiv:2101.05380, January 2021.
[6] Boris Muzellec and Adrien Vacher and Francis Bach and François-Xavier Vialard and Alessandro Rudi. A Dimension-free Computational Upper-bound for Smooth Optimal Transport Estimation. *arXiv e-prints*, page arXiv:2101.05380, December 2021.
[7] Holger Wendland and Christian Rieger. Approximate interpolation with applications to selecting smoothing parameters. *Numer. Math.*, 101(4):729–748, October 2005.

Towards Molecular Computational Anatomy?

Alain Trouvé
(joint work with Michael Miller, Daniel Tward, Laurent Younes)

Current descriptions of brain diseases usually need to put together several orders of magnitude ranging from the millimeter scale for tissues in standard imaging devices to the micron or even nano scale for neural cells and molecules.

Organizing these representations within a given patient or between a population for statistical modelling and understanding could be quite helpful but is still very challenging from a mathematical and computational perspective. In this talk, I will present our recent attempt to make a step in that direction in the context of computational anatomy based on two key assets: 1- Layered coarse-to-fine diffeomorphic transport based on idea coming from optimal control and riemannian geometry 2- Varifold based representations of information and reduction.

References

[1] M.I. Miller, D. Tward, A. Trouvé *Molecular Computational Anatomy: Unifying the Particle to Tissue Continuum via Measure Representations of the Brain*, BME Frontiers (To appear).
[2] M.I. Miller, A. Trouvé, L. Younes, *Image Varifolds on Meshes for Mapping Spatial Transcriptomics*, arXiv:2208.08376 (2022)

Analysis of 1-Lipschitz Neural Networks

Sebastian Neumayer
(joint work with Pakshal Bohra, Stanislas Ducotterd, Alexis Goujon, Dimitris Perdios, and Michael Unser)

The topics covered in this talk are related to the recent preprint [1]. Lipschitz constrained neural networks have several advantages compared to unconstrained ones and can be applied to various different problems. Consequently, they have recently attracted considerable attention in the deep learning community. Since designing and training expressive Lipschitz-constrained networks is very challenging, there is a need for improved methods and a better theoretical understanding. As the general case is very demanding, we restrict our attention to feed-forward neural networks with 1-Lipschitz component-wise activation functions and weight matrices with p-norm less or equal than one. This indeed leads to 1-Lipschitz

neural networks, for which naturally the question of expressiveness arises. Unfortunately, it turns out that networks with ReLU activation functions have provable disadvantages in this setting. Firstly, they cannot represent even simple piece-wise linear functions such as the hat function. Secondly, there exists a whole class of relatively simple functions that cannot be approximated in terms of the uniform norm on bounded boxes. To show this fact, we can make use of the second-order total variation and the fact that ReLU networks can only produce functions with bounded second-order total variation.

Due to these observations, we propose to use learnable spline activation functions with at least 3 linear regions instead. Clearly, this more complicated architecture should be motivated by theoretical findings. To this end, we prove that our architecture is optimal among all component-wise 1-Lipschitz activation functions in the sense that no other weight constrained architecture can approximate a larger class of functions. However, it remains an open question whether such NNs are universal approximators of 1-Lipschitz functions and our result can be seen as a first step towards its solution. Further, we prove that our proposed networks are in principal able to reproduce functions with arbitrary high second-order variation. Note that our architecture is also at least as expressive as the recently introduced non component-wise Groupsort activation function [2] for 2-norm-constrained weights. A more thorough comparison of linear splines to non component-wise activation functions is subtle, and it is so far unclear which choice leads to more expressive NNs in the remaining settings. Concerning the question of universality, the talk focused mainly on the approximation of scalar-valued functions $f\colon \mathbb{R}^d \to \mathbb{R}$. This also reflects the current state of research, where most results are only formulated for scalar-valued NNs. Extending these results to vector-valued functions appears highly non-trivial and should be addressed in future research. Finally, I would like to mention that little is known about the optimal structure for deep spline and Groupsort NNs, i.e., if it is more preferable to go deep or wide in architecture design.

On the numerical side, we are currently preparing a preprint with extensive experiments and details for an efficient implementation. For the implementation, we basically rely on a B-spline representation, which was already used before in [3]. However, we also need to take care of the additional Lipschitz constrained. This can be done in several ways, but naive approaches can lead to inferior training performance. To circumvent this issue, we instead propose to directly optimize over the set of 1-Lipschitz linear splines based on a method called *Differentiable Slope Clipping*. Our preliminary numerical results for one-dimensional function fitting, Wasserstein distance estimation and image reconstruction within the Plug-and-Play framework confirm that our architecture is at least competitive (often even better) with other recently proposed activation functions such as GroupSort, Householder activations and parametric ReLU, which were also all designed with the goal of increasing expressivity in mind. One additional advantage of our implementation over the other methods is that it can be applied to any already trained network by just initializing the linear splines accordingly. This possibly avoids

an expensive retraining. A Github repository with the implementation will be available soon.

References

[1] S. Neumayer, A. Goujon, P. Bohra, and M. Unser, *Approximation of Lipschitz Functions using Deep Spline Neural Networks*, arXiv:2204.06233 (2022).

[2] Q. Li, S. Haque, C. Anil, J. Lucas, R. Grosse, and J. Jacobsen, *Preventing Gradient Attenuation in Lipschitz Constrained Convolutional Networks*, Advances in Neural Information Processing Systems **32**, (2019), 15364–15376.

[3] P. Bohra, J. Campos, H. Gupta, S. Aziznejad, and M. Unser, *Learning activation functions in deep (spline) neural networks*, IEEE Open Journal of Signal Processing **1** (2020), 295–309.

Discrete geodesic calculus in the manifold of Sobolev curves

Benedikt Wirth

(joint work with Martin Rumpf)

The manifold of (closed) Sobolev curves is a well-known example of an infinite-dimensional shape space. It consists of immersions of the circle $\mathbb{S}^1$ into $\mathbb{R}^d$ with Sobolev regularity,

$$\mathrm{Imm}^m = \{c \in W^{m,2}(\mathbb{S}^1;\mathbb{R}^d) \,|\, c'(\theta) \neq 0 \text{ for all } \theta \in \mathbb{S}^1\},$$

where $m \geq 2$ and c' denotes the derivative with respect to the parametrization variable θ (which for $m \geq 2$ is everywhere defined). This manifold can be equipped with a Riemannian metric of Sobolev type,

$$g_c(\xi,\zeta) = \int_{\mathbb{S}^1} \sum_{i=1}^{m} \partial_s^i \xi \cdot \partial_s^i \zeta \,\mathrm{d}s$$

for any curve $c \in \mathrm{Imm}^m$ and tangent vectors $\xi, \zeta : \mathbb{S}^1 \to \mathbb{R}^d$. Above, $s = \int_0^\theta |c'(\tilde\theta)| \,\mathrm{d}\tilde\theta$ denotes arclength along the curve c so that

$$\mathrm{d}s = |c'(\theta)|\mathrm{d}\theta, \qquad \partial_s = \tfrac{\partial_\theta}{|c'(\theta)|}.$$

The induced Riemannian distance between two curves $c_0, c_1 \in \mathrm{Imm}^m$ can then be computed by minimizing the path energy $\mathcal{E}$ among all paths $(c_t)_{t\in[0,1]}$ in Imm^m with fixed endpoints c_0, c_1,

$$d^2(c_0,c_1) = \inf \mathcal{E}[(c_t)_{t\in[0,1]}] \qquad \text{with } \mathcal{E}[(c_t)_{t\in[0,1]}] = \int_0^1 g_{c_t}(\dot c_t, \dot c_t)\,\mathrm{d}t$$

(where $\dot c_t$ denotes the derivative of the path with respect to the time variable t). Furthermore, geodesics in the manifold of Sobolev curves can be defined as minimizers of this path energy for fixed endpoints.

Bruveris, Michor and Mumford have shown [1, 2] that this manifold of Sobolev curves is metrically and geodesically complete and that shortest geodesics between any two curves exist (as long as they lie in the same connected component, thus, if they have the same winding number when $d = 2$). To show this one exploits that the path energy $\mathcal{E}$ actually just behaves like the squared Sobolev norm of

$W^{1,2}([0,1];W^{m,2}(\mathbb{S}^1;\mathbb{R}^d))$, to which the direct method of the calculus of variations can readily be applied. This behaviour is not obvious since $\mathcal{E}$ contains arclength derivatives and arclength integration, while $W^{m,2}(\mathbb{S}^1;\mathbb{R}^d)$ only contains differentiation and integration with respect to the parametrization variable θ. The central trick is to show that $\log|c'|$ is uniformly bounded on any metric ball of Imm^m, which more or less (up to bounded terms) turns differentiation and integration with respect to s into differentiation and integration with respect to θ.

We aim to devise discrete approximations to geodesics in Imm^m which at the same time are numerically feasible and provably convergent. To this end we employ the framework of discrete geodesic calculus [3], that is, we define a discrete K-geodesic between $c_0, c_K \in \mathrm{Imm}^m$ as a minimizer of a discrete path energy

$$E[(c_0, c_1, \dots, c_K)] = \frac{1}{K}\sum_{j=1}^{K} W[c_{j-1}, c_j],$$

where $W[c_A, c_B]$ is a suitable approximation of the squared Riemannian distance $d^2(c_A, c_B)$. The existence and convergence proofs for discrete geodesics would naturally have to follow the compactness and lower semi-continuity type arguments from the continuous setting. To this end it is essential that the uniform bound on $\log|c'|$ and similar estimates carry over from the continuous to the discrete setting. Therefore we consider as initial approximation

$$W[c_A, c_B] = \mathcal{E}[(c_t^{AB})_{t\in[0,1]}] \quad \text{for } c_t^{AB} = tc_B + (1-t)c_A.$$

As a consequence, the discrete path energy $E[(c_0, \dots, c_K)]$ equals the continuous path energy $\mathcal{E}[(c_t)_{t\in[0,1]}]$ of the piecewise linear interpolation c_t of the points $c_0, \dots, c_K$, and boundedness of the discrete path energy therefore implies the same estimates as boundedness of the continuous path energy, as desired. However, the above choice of W is not yet computationally feasible since the time integral involved in

$$E[(c_t^{AB})_{t\in[0,1]}] = \int_0^1 \int_{\mathbb{S}^1} \sum_{i=1}^{m} |\partial_s^i \dot{c}_t^{AB}|^2 \,\mathrm{d}s\,\mathrm{d}t$$

cannot easily be evaluated (the spatial integration over $\mathbb{S}^1$ on the other hand is less critical and will later be approximated via quadrature). As a remedy we replace the integration $\mathrm{d}s = |(c_t^{AB})'|\mathrm{d}\theta$ by the larger $[t|c_B'| + (1-t)|c_A'|]\mathrm{d}\theta$,

$$W[c_A, c_B] = \int_0^1 \int_{\mathbb{S}^1} [t|c_B'| + (1-t)|c_A'|] \sum_{i=1}^{m} |\partial_s^i \dot{c}_t^{AB}|^2 \,\mathrm{d}\theta\,\mathrm{d}t.$$

This way the integrand of $W[c_A, c_B]$ becomes (after tedious calculations) a rational function in t, which can be explicitly integrated and thereby makes the discrete path energy numerically feasible. At the same time, the discrete path energy still bounds the continuous one from above, $E[(c_0, \dots, c_K)] \geq \mathcal{E}[(c_t)_{t\in[0,1]}]$, so that the compactness properties are still maintained and we obtain the following.

Theorem 1 (Existence and convergence of discrete geodesics)**.** *Let $c_A, c_B \in \mathrm{Imm}^m$. For K large enough there exists a discrete K-geodesic between c_A and*

c_B. Furthermore, as $K \to \infty$, the piecewise linear interpolations of these discrete geodesics converge weakly in $W^{1,2}([0,1]; W^{m,2}(\mathbb{S}^1;\mathbb{R}^d))$ (along a subsequence) to a continuous geodesic between c_A and c_B.

For a numerical implementation we additionally propose a spectral space discretization, that is, we represent a curve c by $2N+1$ Fourier components,

$$c(\theta) = \sum_{j=0}^{N} \begin{pmatrix} a_j^x \\ a_j^y \end{pmatrix} \cos(j\theta) + \begin{pmatrix} b_j^x \\ b_j^y \end{pmatrix} \sin(j\theta),$$

and we approximate any integral over $\mathbb{S}^1$ using trapezium rule quadrature with M equispaced points (the motivation is the exponential convergence of trapezium rule quadrature for smooth periodic functions).

Theorem 2 (Convergence of discrete geodesics)**.** *If $N = K^\alpha$ and $M = K^\beta$ for some arbitrary $\beta > \alpha > 0$, then spatially and temporally discretized geodesics converge to continuous ones in the same sense as before.*

The combination of a spectral discretization with the trapezium rule actually is essential here – had one instead used spline discretizations, the number M of quadrature points would have to be chosen as a substantially larger power of the time step number K. All above convergence results are obtained by actually showing Mosco convergence of the (time- or space-time-)discrete path energy to the continuous one. An example of a numerically computed discrete geodesic between a frog and a turtle shape from the MPEG-7 Core Experiment CE-Shape-1 is shown below. The last five shapes are obtained by discrete geodesic extrapolation of which one can show convergence as well.

References

[1] M. Bruveris, P. Michor, D. Mumford, *Geodesic completeness for Sobolev metrics on the space of immersed plane curves*, Forum Math. Sigma **2** (2014).

[2] M. Bruveris, *Completeness properties of Sobolev metrics on the space of curves*, J. Geom. Mech. **7** (2015), 125–150.

[3] M. Rumpf, B. Wirth, *Variational time discretization of geodesic calculus*, IMA J. Numer. Anal. **35** (2015), 1011–1046.

Reporter: Florine Hartwig

Participants

Mariem Abaach
Université de Paris
45 rue des Saints-Péres
75006 Paris
FRANCE

Prof. Dr. Nina Amenta
Department of Computer Science
University of California, Davis
One Shields Ave.
Davis, CA 95616-8633
UNITED STATES

Dr. Omri Azencot
Computer Science Department
Ben-Gurion University of the Negev
Beer-Sheva 84 105
ISRAEL

Prof. Dr. Mirela Ben-Chen
Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 3200003
ISRAEL

Prof. Dr. Mario Botsch
Lehrstuhl Informatik 7
TU Dortmund
Otto-Hahn-Straße 16
44227 Dortmund
GERMANY

Amit Bracha
Department of Mathematics
TECHNION
Israel Institute of Technology
Haifa 32000
ISRAEL

Juliane Braunsmann
Fachbereich Mathematik und Informatik
Universität Münster
Einsteinstraße 62
48149 Münster
GERMANY

Prof. Dr. Alexander Bronstein
Computer Science Department
TECHNION
Israel Institute of Technology
32000 Haifa
ISRAEL

Prof. Dr. Dorin Bucur
Laboratoire de Mathématiques
Université Savoie Mont Blanc
73376 Le Bourget-du-Lac
FRANCE

Dr. Blanche Buet
Laboratoire de Mathématiques d'Orsay
Université Paris Saclay
Bat. 307
91405 Orsay
FRANCE

Prof. Dr. Antonin Chambolle
CEREMADE, CNRS, Université Paris Dauphine
PSL Research University
Place de Lattre de Tassigny
75775 Paris Cedex 16
FRANCE

Dr. Albert R. Chern
Department of Computer Science and Engineering,
University of California, San Diego
9500 Gilman Dr MC 0404
La Jolla CA 92093
UNITED STATES

Prof. Dr. Keenan Crane
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh 15213
UNITED STATES

Prof. Angela Dai
Technische Universität München
Boltzmannstraße 3
85748 Garching bei München
GERMANY

Prof. Dr. Mathieu Desbrun
Department of Computer Science
California Institute of Technology
1200 East California Boulevard
Pasadena CA 91125
UNITED STATES

Dr. Olga Diamanti
Institut für Mathematik
Technische Universität Graz
Kopernikusgasse 24
8010 Graz
AUSTRIA

Michal Edelstein
Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 32000
ISRAEL

Jean Feydy
HeKA team, INRIA Paris
1 rue Jeanne d'Arc
92130 Issy-les-Moulineaux
FRANCE

Dr. Peter Gladbach
Institut für Angewandte Mathematik
Universität Bonn
Endenicher Allee 60
53115 Bonn
GERMANY

Prof. Dr. Eitan Grinspun
Departments of Computer Science and
(by cross-appointment) Mathematics
University of Toronto
100 St. George Street
Toronto ON M5S 1A1
CANADA

Prof. Dr. Philipp Grohs
Fakultät für Mathematik
Universität Wien
Oskar Morgenstern Platz 1
1090 Wien
AUSTRIA

Prof. Dr. Leonidas J. Guibas
Computer Science Department
Stanford University
353 Jane Stanford Way
Stanford CA 94305-9035
UNITED STATES

Florine Hartwig
Institut für Numerische Simulation
Universität Bonn
Endenicher Allee 60
53115 Bonn
GERMANY

Dr. Klaus Hildebrandt
Delft University of Technology, EEMCS,
Dept. Intelligent Systems
Van Mourik Broekmanweg 6
2628 XE Delft
NETHERLANDS

Prof. Dr. Ron Kimmel
Department of Computer Sciences
TECHNION
Israel Institute of Technology
32000 Haifa
ISRAEL

Felix J. Knöppel
Institut für Mathematik
MA 8 - 3
Technische Universität Berlin
Straße des 17. Juni 136
10623 Berlin
GERMANY

Dr. Dirk Lorenz
Institute for Analysis and Algebra
TU Braunschweig
Universitätsplatz 2
38106 Braunschweig
GERMANY

Prof. Dr. Niloy Mitra
Department of Computer Science
University College London
Gower Street
London WC1E 6BT
UNITED KINGDOM

Dr. Sebastian Neumayer
EPFL/STI/IMT/LIB
BM 4.134 (Bâtiment BM)
P.O. Box Station 17
1015 Lausanne
SWITZERLAND

Prof. Dr. Matteo Novaga
Dipartimento di Matematica
Università di Pisa
Largo Bruno Pontecorvo 5
56127 Pisa
ITALY

Dr. Xavier Pennec
Inria Sophia-Antipolis, Epione team
Côte d'Azur University and INRIA
2004 Route des Lucioles
P.O. Box BP93
06902 Sophia-Antipolis Cedex
FRANCE

Prof. Dr. Ulrich Pinkall
Institut für Mathematik
Technische Universität Berlin
Sekr. MA 8-1
Straße des 17. Juni 136
10623 Berlin
GERMANY

Prof. Dr. Helmut Pottmann
King Abdullah University of Science and Technology
23955-6900 Thuwal
SAUDI ARABIA

Shir Rorberg
Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 32000
ISRAEL

Noam Rotstein
Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 32000
ISRAEL

Prof. Dr. Martin Rumpf
Institut für Numerische Simulation
Universität Bonn
Endenicher Allee 60
53115 Bonn
GERMANY

Josua Sassen
Institut für Numerische Simulation
Universität Bonn
Endenicher Allee 60
53115 Bonn
GERMANY

Prof. Dr. Otmar Scherzer
Fakultät für Mathematik
Universität Wien
Oskar-Morgenstern-Platz 1
1090 Wien
AUSTRIA

Prof. Dr. Carola-Bibiane Schönlieb
Department of Applied Mathematics and Theoretical Physics (DAMTP)
Centre for Mathematical Sciences
Wilberforce Road
Cambridge CB3 0WA
UNITED KINGDOM

Prof. Dr. Peter Schröder
Annenberg Center, CMS
MC 305-16
California Institute of Technology
1200 E. California Blvd.
Pasadena CA 91125
UNITED STATES

Dr. Henrik Schumacher
Fakultät für Mathematik
TU Chemnitz
09107 Chemnitz
GERMANY

Simon Schwarz
Institut für Numerische und Angewandte Mathematik
Georg-August-Universität Göttingen
Lotzestraße 16-18
37073 Göttingen
GERMANY

Dr. Nicholas Sharp
Department of Mathematics
University of Toronto
100 St. George Street
Toronto ON M5S 1A1
CANADA

Yousuf Soliman
Department of Computing+Mathematical Sciences
MC 305-16
California Institute of Technology
1200 E California Blvd
Pasadena, CA 91125
UNITED STATES

Prof. Dr. Justin Solomon
EECS Department
M.I.T.
32 Vassar Street, office 32-D460
02139 Cambridge
UNITED STATES

Prof. Dr. Olga Sorkine-Hornung
Department of Computer Science
CNB G 106.2
ETH Zurich
Universitätstr. 6
8092 Zürich
SWITZERLAND

Dominik Stantejsky
Centre de Mathématiques Appliquées
École Polytechnique
Route de Saclay
91128 Palaiseau Cedex
FRANCE

Prof. Dr. Gabriele Steidl
TU Berlin
Institute of Mathematics
Straße des 17. Juni
10623 Berlin 10623
GERMANY

Dr. Oded Stein
MIT CSAIL
The Stata Center
Office 32-D475A
32 Vassar Street
Cambridge 02139 MA
UNITED STATES

Prof. Dr. Alain Trouvé
Centre Borelli
ENS Paris-Saclay
4, avenue des sciences
91190 Gif-sur-Yvette Cedex
FRANCE

Prof. Dr. Francois-Xavier Vialard
Laboratoire d'Informatique Gaspard Monge
Université Gustave Eiffel
UMR 8049
77454 Marne-la-Vallée
FRANCE

Prof. Dr. Etienne Vouga
Department of Computer Science
The University of Texas at Austin
2317 Speedway, Stop D 9500
78701 Austin
UNITED STATES

Prof. Dr. Johannes Wallner
Institut für Geometrie
Technische Universität Graz
Kopernikusgasse 24
8010 Graz
AUSTRIA

Prof. Dr. Max Wardetzky
Institut für Numerische und Angewandte Mathematik
Universität Göttingen
Lotzestraße 16-18
37083 Göttingen
GERMANY

Prof. Dr. Benedikt Wirth
Fachbereich Mathematik und Informatik
Universität Münster
Einsteinstraße 62
48149 Münster
GERMANY

Prof. Dr. Laurent Younes
John Hopkins University
Whiting School of Engineering
3400 North Charles Street
Baltimore, MD 21218-2608
UNITED STATES

Prof. Dr. Denis Zorin
Computer Science Department
Courant Institute of Mathematical Sciences
New York University
719 Broadway
New York, NY 10003
UNITED STATES