

Report No. 5/2025

DOI: 10.4171/OWR/2025/5

Mathematical Imaging and Surface Processing

Organized by
Mirela Ben-Chen, Haifa,
Antonin Chambolle, Paris,
Benedikt Wirth, Münster

2 February – 7 February 2025

ABSTRACT. The workshop “Mathematical Imaging and Surface Processing” is designed to foster exchange and collaboration between applied mathematicians and applied computer scientists on problems arising in computer graphics, geometry processing and image analysis. The participants are also interested in computational and theoretical issues of data science, and theory and methods of machine learning were featured in many of the talks. About 40 participants, with many younger post-docs or PhD students, could exchange and work on these themes, while coming from quite different backgrounds (academic and industrial). Several talks emphasized the strength of fundamental theoretical concepts, such as symmetries of Lie groups or curvature in Riemannian geometry, for important applications in graphics, vision, construction and architecture.

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

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

License: Unless otherwise noted, the content of this report is licensed under CC BY SA 4.0.

Introduction by the Organizers

The workshop *Mathematical Imaging and Surface Processing* followed two previous events on similar topics, in January 2016 (workshop # 1604) and August 2022 (workshop # 2234). It gathered more than 40 participants. Participants were applied computer scientists and applied mathematicians (from the US, Israel, Germany, France, etc.), interested in computational geometry and mathematical imaging (graphics, analysis, machine learning). There was a strong participation of younger researchers, with four talks by PhD students and several by young

post-docs on cutting edge research in geometry processing, geometric control and generative modeling in imaging. There was strong and interesting exchange between participants from quite different backgrounds, all sharing common concerns and a common mathematical culture.

During the first day, the talks addressed the geometry of diffeomorphisms, motion optimization, the design of inflatable surfaces, improvements to the “AlphaFold” protein generative model, optimal transportation and generative AIs. We also had a talk on new geometric measure theoretic tools for geometry processing and on analytical and computational tools for computing signed distances. On Tuesday, impressive methods for surface matching were presented, mixing non-linear analysis and machine learning techniques, talks about the use of harmonic functions and conformal maps in computational geometry, and about the discretization of geometric quantities. Before the walk to Sankt Roman on Wednesday, the presentations focussed on applied and theoretical machine learning while the last talk presented impressive applications of theoretical geometry to architectural design. Machine learning and neural network techniques were also an important topic on Thursday, together with an impressive computer graphics talk on the simulation of wrinkles and their fascinating motion. The main topic for Friday was the computation of Wasserstein barycenters, and we could admire the impressive Gromov–Wasserstein “average” of computer graphics shapes such as a bear and a cow. Two other talks addressed neural network based image processing, showing in particular how one could derive “frugal” and efficient learned denoisers.

The workshop was a great success. The atmosphere was both professional and informal, and all the talks were highly appreciated. The mixture of theoretically oriented talks and very applied ones was interesting to all. The strong presence of younger researchers allowed them to feel fully part of the group and led to fruitful exchange between the participants. It was very advantageous to mix participants from various fields, all interested in data science, image and geometry processing, some relying on machine learning based tools, other on more “classical” mathematical (geometric) analysis, and yet others on both, as they could benefit from exchange with colleagues of quite different backgrounds, yet sharing many common interests and a common language. We hope that it will still be possible to organise such a conference in the future, on similar topics and most importantly, mixing again applied mathematicians and computer scientists and featuring a high participation of younger colleagues.

Workshop: Mathematical Imaging and Surface Processing

Table of Contents

Alex Bronstein (joint with Advaith Maddipatla, Nadav Sellam, Meital Bojan, Sanketh Vedula, Paul Schanda, and Ailie Marx) <i>Inverse problems with experiment-guided AlphaFold</i>	215
Yingying Ren (Samara) <i>Computational Homogenization for Inverse Design of Surface-based Inflatables</i>	219
Johannes Hertrich (joint with Robert Gruhlke) <i>Importance Corrected Neural JKO Sampling</i>	220
Albert Chern (joint with Mohammad Sina Nabizadeh, Ritoban Roy-Chowdhury, Hang Yin, and Ravi Ramamoorthi) <i>Discrete Fluid Dynamics with Sub-Riemannian Geometry</i>	221
Florine Hartwig (joint with Oliver Gross, Martin Rumpf, Peter Schröder) <i>Optimal motion from shape change</i>	223
Nicole Feng (joint with Ioannis Gkioulekas and Keenan Crane) <i>Generalized Signed Distance</i>	226
Leticia Mattos Da Silva (joint with Silvia Sellán, Oded Stein, Justin Solomon) <i>Geometric Algorithms, Stochasticity and the Fokker-Planck Equation</i> ...	230
Xavier Pennec (joint with Tom Szwagier, Dimbihery Rabenoro) <i>Flag spaces for Geometric statistics</i>	231
Ron Kimmel (joint with Amir Bracha, Thomas Dages, Yaniv Wolf) <i>Learning geometry – a perspective on reconstruction, isometry, partiality and topology</i>	233
Josua Sassen (joint with Patrick Dondl, Martin Rumpf) <i>A Phase-field Approach to Optimal Reinforcing Networks</i>	234
Mark Gillespie (joint with Denise Yang, Mario Botsch, and Keenan Crane) <i>Ray Tracing Harmonic Functions</i>	236
Amir Vaxman (joint with Mirela Ben-Chen, Shir Rorberg, Kaibo Hu, Christian Müller, Ofir Weber, and Brad Wang) <i>Conformal Maps and Directional Fields: Can two walk together, except they be agreed?</i>	238
Michael Kazhdan (joint with Max Kohlbrenner, Sing-Chun Lee, Marc Alexa) <i>Poisson Manifold Reconstruction (beyond co-dimension one)</i>	241

Noam Rotstein	
<i>Consistent Image Editing – Manipulating Images Only Where Change is Needed</i>	242
Emile Pierret (joint with Bruno Galerne)	
<i>Diffusion Models for Gaussian Distributions: Exact Solutions and Wasserstein Errors</i>	243
Ruben Wiersma (joint with Ahmad Nasikun, Elmar Eisemann, Klaus Hildebrandt)	
<i>Intrinsic approaches to learning and computing on curved surfaces</i>	245
Agnès Desolneux (joint with Hermine Biermé, Valentin de Bortoli, Julie Delon, Antoine Salmona)	
<i>Isoperimetric inequalities: Gaussian, on Gaussian and applications</i>	247
Nicolas Papadakis (joint with Antonin Chambolle, Samuel Hurault, Arthur Leclaire, Jean Prost, Marien Renaud)	
<i>Convergent Plug and Play methods for image restoration</i>	248
Helmut Pottmann (joint with Yu-Chou Chiang, Martin Kilian, Daoming Liu, Christian Müller, Davide Pellis, Florian Rist, Eike Schling, Johannes Wallner, Bolun Wang, Hui Wang)	
<i>Geometry and computational design of strip structures</i>	252
Zorah Lähner	
<i>Geometric Optimization for Plasma Fusion</i>	252
Niloy J. Mitra (joint with Romy Williamson)	
<i>Rethinking Geometry Processing via Neural Surfaces</i>	253
Nicholas Sharp (joint with Tianchang Shen)	
<i>Representations of Surface Meshes for Optimization, Inverse Problems, and Learning</i>	254
Roi Poranne	
<i>Smoothed L0 Optimization</i>	256
Etienne Vouga (joint with Zhen Chen, Melina Skouras, Danny Kaufman)	
<i>Failures and Successes Simulating Wrinkles on Surfaces</i>	258
Simon Schwarz (joint with Anja Sturm and Max Wardetzky)	
<i>(In-)consistency of the cotan Laplacian and convergence of the heat method</i>	260
Romy Williamson (joint with Niloy Mitra)	
<i>Neural Geometry Processing via Spherical Neural Surfaces</i>	262
Robert Beinert (joint with Florian Beier, Gabriele Steidl)	
<i>A geometric optimal transport framework for 3d shape interpolation</i>	263
Sebastian Neumayer (joint with Mehrsa Pourya, Erich Kobler, Michael Unser)	
<i>DEALing with Image Reconstruction: Deep Attentive Least Squares</i>	266

Samuel Vaiter (joint with Alberto Bietti, Nicolas Keriven)	
<i>Graph Neural Networks on Large Random Graphs</i>	267
Julie Delon (joint with Eloi Tanguy, Nathaël Gozlan)	
<i>Beyond Wasserstein: a general fixed-point approach for computing</i> <i>barycentres of measures</i>	268

Abstracts

Inverse problems with experiment-guided AlphaFold

ALEX BRONSTEIN

(joint work with Advait Maddipatla, Nadav Sellam, Meital Bojan,
Sanketh Vedula, Paul Schanda, and Ailie Marx)

Proteins are inherently dynamic entities, sampling a continuum of conformational states to fulfill their biological roles. Experimental techniques such as X-ray crystallography, nuclear magnetic resonance (NMR) spectroscopy, and cryo-electron microscopy (cryo-EM) inherently report on ensemble-averaged data rather than singular static snapshots. In X-ray crystallography, the resolved electron density map represents a spatial and temporal average over all molecules in the crystal lattice, with regions of flexibility manifesting as diffuse or poorly resolved density. NMR spectroscopy measures the interaction between nuclear spins (e.g., magnetization transfer due to nuclear Overhauser effect, NOE) and spins and electrons (e.g., chemical shifts) arising from dynamic conformational ensembles in solution, with these experimental restraints used computationally to resolve compatible structural states. Cryo-EM similarly resolves multiple conformational states, as individual particles frozen in vitreous ice adopt distinct orientations and conformations, which are computationally classified into discrete or continuous flexibility ranges.

On the computational front, *ab initio* protein structure determination based on modeling the molecule’s free energy and its subsequent minimization (e.g., Rosetta and many of its variants [3, 4]) have been only partially successful and computationally expensive. A giant leap in protein structure prediction resulted from the fundamental discovery of the coevolution of contacting residues [7, 8], underlying deep learning-based models such as AlphaFold [1, 10], which had a groundbreaking impact on structural biology and was awarded the recent 2024 Nobel Prize in Chemistry.

Protein structure predictors are trained exclusively on X-ray crystallographic models, which are themselves *fitted* to electron density maps averaged over trillions of molecule instances. While it has been recognised several decades ago that the conformations of proteins in crystals are heterogeneous [6, 16], early crystallographic refinements prioritized single-conformer models. Advances in resolution, the more widespread application of room-temperature crystallographic experiments (as opposed to those performed at 100 K), and progress in refinement tools now permit explicit modeling of alternative conformations (“altlocs”) within overlapping density regions [6, 17, 18]. Recent studies analyzing the PDB reveal that such multi-conformer annotations are widespread, reflecting inherent structural variability captured in crystallography [14]. In NMR spectroscopy, the experimental observables, such as inter-atomic distances or bond-vector orientations reflect the time- and ensemble average, and NMR structures are always reported as bundles of conformations. However, AlphaFold’s training objective

– to predict a single “most probable” structure – biases its output toward static snapshots, effectively marginalizing conformational heterogeneity encoded in its training data. Over the past year, multiple sequence-conditioned protein structure generative models like AlphaFlow [9], and the recent AlphaFold3 [1] have been proposed to move beyond the one-sequence-one-structure paradigm. However, since these approaches remain *trained on unimodally-modeled PDB entries* derived predominantly from crystallographic data, the generated ensembles fail to capture the full heterogeneity implied by experimental measurements, thus limiting their practical utility [15]. This emphasizes the need for new models that can explicitly model protein ensembles that are faithful to experimental measurements. Developing such models is the focus of the present work.

In this work, we introduce *experiment-guided AlphaFold3*, a computational framework that integrates experimental data with deep learning priors to generate structural ensembles consistent with experimental observables. Our key insight is that AlphaFold3 can be viewed as a strong sequence-conditioned protein structure prior that may be further leveraged to solve inverse problems in the space of protein structures. By solving these inverse problems under the prior imposed by AlphaFold3, we bridge the gap between data-driven predictions and experimental evidence, yielding ensembles that are both physically plausible and experimentally consistent.

Experiment-guided AlphaFold3. Our primary technical contribution is a three-stage ensemble-fitting pipeline. First, we present Guided AlphaFold3, where we adapt the diffusion-based structure module of AlphaFold3 to incorporate experimental measurements during sampling. To properly handle ensemble measurements, we introduce a *non-i.i.d. sampling scheme that jointly samples the ensemble*, directing conformational exploration toward regions compatible with the experimental constraints. We show that this approach effectively captures multimodal ensemble measurements, where standard i.i.d. sampling methods fail. To our knowledge, this represents the first application of guided sampling within AlphaFold3 for experimental structural resolution. Second, we address artifacts introduced during guided sampling by using AlphaFold2’s computationally efficient force-field relaxation step, effectively projecting candidate structures onto physically realistic conformations. Finally, we develop a matching-pursuit ensemble selection algorithm to iteratively refine the ensemble by maximizing agreement with experimental data while preserving structural diversity. We validate our framework through case studies on two foundational challenges in structural biology: (1) X-ray crystallographic structure modeling, where we recover conformational heterogeneity obscured in static electron density maps, and (2) NMR structure determination, where we resolve ensembles that obey NOE-derived distance restraints.

Improved crystal density modeling. X-ray crystallography is one of the most accurate techniques for protein structure determination. A typical pipeline involves the crystallization of protein samples and the subsequent fitting of atomic structures to electron density maps generated from X-ray diffraction patterns.

However, this procedure is expensive, time-consuming, and often requires manual intervention by crystallographers [5]. As a result, several structures deposited in the PDB exhibit human-induced biases that can degrade structural accuracy. Another limitation of crystallographic pipelines is the misleading notion of “single crystal and single structure”. However, the PDB exhibits multimodality in the density that cannot be fully captured by models like AlphaFold3 that predict single structures. This limitation, recognized early on in protein crystallography [16], is particularly evident in *altloc regions* [15], where multiple conformations coexist in the same lattice [6, 17]. This inadequacy presents a compelling case for protein generative models to improve crystallographic structural modeling.

Hence, we introduce *Density-guided AlphaFold3*, which guides AlphaFold3 generated structural ensembles to be faithful to experimental electron density maps. Density-guided AlphaFold3 renders structures that are consistently more faithful to the observed electron density maps than unguided AlphaFold3. In some cases, the guided structure outperforms PDB-deposited structure’s faithfulness to the density. Additionally, guided structures capture structural heterogeneity more accurately than AlphaFold3. In some cases, guided structures capture the structural heterogeneity that PDB-deposited structures fail to model. Lastly, we are able to leverage the strong prior learned by AlphaFold3 to generate density-faithful ensembles in a fraction of the time required by conventional X-ray crystallography pipelines [2]. In our opinion, this advancement not only improves the accuracy of computational structural modeling but also has the potential to automate workflows for crystallographers.

Accelerated NMR ensemble structure determination. Solution-state NMR enables the study of proteins in near-physiological aqueous environments, capturing conformational heterogeneity through nuclear interaction restraints such as nuclear Overhauser effects (NOEs) and scalar couplings (J-couplings). NMR-based structure determination typically employs restrained molecular dynamics (MD) simulations, requiring hundreds of independent trajectories to adequately sample conformational spaces consistent with experimental data—a computationally intensive process that struggles to balance accuracy, efficiency, and ensemble diversity [11, 13]. Here, we propose *NOE-guided AlphaFold3*, which refines AlphaFold-generated structural ensembles to satisfy NOE-derived distance restraints. The resulting ensembles adhere to experimental NOE data more faithfully than AlphaFold3 predictions and, *in some cases, even surpass the accuracy of existing PDB-deposited NMR ensembles*. In particular, we demonstrate that the ensembles produced by NOE-guided AlphaFold3 on ubiquitin, a benchmark system for NMR structure and dynamics, accurately capture experimentally observed conformational flexibility, as independently validated against experimentally-measured N-H S^2 order parameters; [12]). In contrast, standard AlphaFold3 predictions generate overly rigid ensembles inconsistent with ubiquitin’s dynamic behavior. Finally, we note that our method dramatically *improves the NMR structure determination process from many hours to a few minutes, while retaining the accuracy obtained*

through MD. We believe this will enable new experimental workflows for NMR structural biologists.

REFERENCES

- [1] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, pages 1–3, 2024.
- [2] Paul D Adams, Pavel V Afonine, Gáabor Bunkóczi, Vincent B Chen, Ian W Davis, Nathaniel Echols, Jeffrey J Headd, L-W Hung, Gary J Kapral, Ralf W Grosse-Kunstleve, et al. Phenix: a comprehensive python-based system for macromolecular structure solution. *Acta Crystallographica Section D: Biological Crystallography*, 66(2):213–221, 2010.
- [3] Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N Kinch, R Dustin Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021.
- [4] Minkyung Baek, Ryan McHugh, Ivan Anishchenko, Hanlun Jiang, David Baker, and Frank DiMaio. Accurate prediction of protein-nucleic acid complexes using rosettafoldna. *Nature methods*, 21(1):117–121, 2024.
- [5] Jennifer A. Doudna. Structural genomics of rna. *Nature Structural Biology*, 7(11):954–956, Nov 2000.
- [6] Nicholas Furnham, Tom L Blundell, Mark a DePristo, and Thomas C Terwilliger. Is one solution good enough? *Nat. Struct. Mol. Biol.*, 13(3):184–185; discussion 185, 2006.
- [7] U. Gobel, C. Sander, R. Schneider, and A. Valencia. Correlated mutations and residue contacts in proteins. *Proteins: Structure, Function, and Bioinformatics*, 18(4):309–317, 1994.
- [8] Thomas A Hopf, Charlotta PI Schärfe, João PGLM Rodrigues, Anna G Green, Oliver Kohlbacher, Chris Sander, Alexandre MJJ Bonvin, and Debora S Marks. Sequence co-evolution gives 3d contacts and structures of protein complexes. *elife*, 3:e03430, 2014.
- [9] Bowen Jing, Bonnie Berger, and Tommi Jaakkola. Alphafold meets flow matching for generating protein ensembles. *arXiv preprint arXiv:2402.04845*, 2024.
- [10] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.
- [11] Oliver F Lange, Nils-Alexander Lakomek, Christophe Fares, Gunnar F Schroder, Korvin FA Walter, Stefan Becker, Jens Meiler, Helmut Grubmuller, Christian Griesinger, and Bert L De Groot. Recognition dynamics up to microseconds revealed from an rdc-derived ubiquitin ensemble in solution. *science*, 320(5882):1471–1475, 2008.
- [12] SF Lienin, T Bremi, B Brutscher, R Bruschweiler, and RR Ernst. Anisotropic intramolecular backbone dynamics of ubiquitin characterized by nmr relaxation and md computer simulation. *Journal of the American Chemical Society*, 120(38):9870–9879, 1998.
- [13] Kresten Lindorff-Larsen, Robert B Best, Mark A DePristo, Christopher M Dobson, and Michele Vendruscolo. Simultaneous determination of protein structure and dynamics. *Nature*, 433(7022):128–132, 2005.
- [14] viv A. Rosenberg, Ailie Marx, and Alexander M. Bronstein. A dataset of alternately located segments in protein crystal structures. *Scientific Data*, 11(1):783, Jul 2024.
- [15] Aviv A. Rosenberg, Sanketh Vedula, Alex M. Bronstein, and Ailie Marx. Seeing double: Molecular dynamics simulations reveal the stability of certain alternate protein conformations in crystal structures. *bioRxiv*, 2024.
- [16] Janet L Smith, Wayne A Hendrickson, Richard B Honzatko, and Steven Sheriff. Structural heterogeneity in protein crystals. *Biochemistry*, 25(18):5018–5027, 1986.
- [17] Henry van den Bedem and James S Fraser. Integrative, dynamic structural biology at atomic resolution—it’s about time. *Nat. Meth.*, 12(4):307–318, 2015.

- [18] Stephanie A Wankowicz, Ashraya Ravikumar, Shivani Sharma, Blake Riley, Akshay Raju, Daniel W Hogan, Jessica Flowers, Henry van den Bedem, Daniel A Keedy, and James S Fraser. Automated multiconformer model building for x-ray crystallography and cryo-em. *eLife*, 12:RP90606, 2024.

Computational Homogenization for Inverse Design of Surface-based Inflatables

YINGYING REN (SAMARA)



Surface-based inflatables are thin-shell structures whose inflated shape emerges from a spatially varying in-plane contraction induced by strategically placed fusing curves. We present a computational framework for the inverse design of such structures, formulating the problem within the broader context of geometric and topological optimization. We first develop numerical homogenization techniques to characterize and optimize geometric and mechanical behavior of periodic inflatable patches with arbitrary fusing patterns, enabling systematic exploration of the design space of surface-based inflatables.

Our inverse design pipeline consists of two stages. Given a target surface, we first perform a coarse optimization step that utilizes the data generated from homogenization to compute an optimal flattening of the target surface and rapidly explore fusing curve topologies. This step involves a local-global parametrization algorithm and a custom nonlinear parametrization scheme. We then further refine the fusing curves using a PDE-constrained nested optimization scheme that directly minimizes the distance of the inflated structure to the target surface. We validate our results by fabricating a set of physical prototypes and demonstrate that our approach is able to produce inflated surfaces that accurately approximate a wide range of 3D geometries and achieve superior structural performance compared to previous methods.

REFERENCES

- [1] Y. Ren, J. Panetta, S. Suzuki, U. Kusupati, F. Isvoranu, and M. Pauly, “Computational Homogenization for Inverse Design of Surface-based Inflatables,” *ACM Transactions on Graphics*, vol. 43, no. 4, pp. 1–18, July 2024. Available: <https://doi.org/10.1145/3658125>

Importance Corrected Neural JKO Sampling

JOHANNES HERTRICH

(joint work with Robert Gruhlke)

In order to sample from an unnormalized probability density function, we propose in [1] to combine continuous normalizing flows (CNFs) with rejection-resampling steps based on importance weights. We relate the iterative training of CNFs with regularized velocity fields to a JKO scheme and prove convergence of the involved velocity fields to the velocity field of the Wasserstein gradient flow (WGF). The alternation of local flow steps and non-local rejection-resampling steps allows us to overcome local minima or slow convergence of the WGF for multimodal distributions. Since the proposal of the rejection step is generated by the model itself, they do not suffer from common drawbacks of classical rejection schemes. The arising model can be trained iteratively, reduces the reverse Kulback-Leibler (KL) loss function in each step, allows to generate iid samples and moreover allows for evaluations of the generated underlying density. Numerical examples show that our method yields accurate results on various test distributions including high-dimensional multimodal targets and outperforms the state of the art in almost all cases significantly.

In an additional project [2], we use the proposed importance-based rejection steps for iterative fine-tuning of diffusion models. More precisely, we assume that we are given a diffusion model and aim to sample from a related posterior distribution. To this end, we add an additional guidance term to the drift of the backward SDE, which is called the h -transform from [3]. Then, we iteratively sample from our current estimation of the h -transform, adapt the dataset by the importance-based rejection steps and retrain the h -transform using the supervised loss from [3]. We prove that this procedure minimizes a certain loss function and demonstrate the effectiveness of this framework on class-conditional sampling and reward fine-tuning for text-to-image diffusion models.

REFERENCES

- [1] J. Hertrich and R. Gruhlke. *Importance Corrected Neural JKO Sampling*, arXiv preprint arXiv:2407.20444, 2024.
- [2] A. Denker, S. Padhy, F. Vargas and J. Hertrich. *Iterative Importance Fine-Tuning of Diffusion Models*, arXiv preprint arXiv:2502.04468, 2025.
- [3] A. Denker, F. Vargas, S. Padhy, K. Didi, S. Mathis, V. Dutordoir, R. Barbano, E. Mathieu, U. J. Komorowska, P. Lio. *DEFT: Efficient Fine-Tuning of Diffusion Models by Learning the Generalised h -transform*, Advances in Neural Information Processing Systems, vol. 37, 2024.

Discrete Fluid Dynamics with Sub-Riemannian Geometry

ALBERT CHERN

(joint work with Mohammad Sina Nabizadeh, Ritoban Roy-Chowdhury,
Hang Yin, and Ravi Ramamoorthi)

The incompressible Euler equations or Navier–Stokes equations exhibit rich phenomena such as intricate vortex interaction at different scales. We aim at reproducing these structures in a discrete framework to facilitate faithful numerical simulation. In order to design such framework we dive into geometric formulations for these fluid equations, which reveal the mathematical structures responsible for vortex conservation. Specifically, vortex conservation (circulation conservation) is equivalent to that the fluid state lies in the same coadjoint orbit arising from the Lie–Poisson structure in the dual space of the Lie algebra of velocity fields, a formulation known in the Hamiltonian fluid dynamics. Our goal is to recover such property in the discrete setting, in contrast to the classical weak formulations in finite elements. This leads to new finite-dimensional dynamical systems mimicking the Euler fluid with close relation to Carnot–Caratheodory sub-Riemannian geometry on a Lie group. The majority of the talk and background material is presented in [1].

1. PRELIMINARY

Let M be a Riemannian manifold representing the fluid domain. Let $G := \text{SDiff}(M)$ be the infinite dimensional Lie group of volume-preserving diffeomorphisms on M . The Lie algebra $\mathfrak{g} := \mathfrak{sdiff}(M)$ is the space of divergence-free tangent vector fields on M with the no-through boundary condition at ∂M . Each vector field $\mathbf{u} \in \mathfrak{g}$ represents a velocity field of the fluid. Each tangent vector $\dot{\phi} \in T_{\phi}G$ at any flow map $\phi \in G$ can be expressed by $\dot{\phi} = \mathbf{u} \circ \phi$ for some $\mathbf{u} \in \mathfrak{g} = T_{\text{id}}G$, which is the velocity field corresponding to an infinitesimal motion $\dot{\phi}$ of the flow map. Define a Riemannian metric for G by (twice) the kinetic energy $\|\dot{\phi}\|^2 := \int_M |\mathbf{u}|^2$ of the velocity field.

Definition 1 (Euler flow). *The time-dependent flow maps for Euler fluid flows are geodesics on G .*

Note that the Riemannian metric on G is (right-)translation-invariant since it depends only on $\dot{\phi}$'s translation \mathbf{u} to \mathfrak{g} . This allows for the Euler equation be described only by the variable $\mathbf{u} \in \mathfrak{g}$:

$$\dot{\mathbf{u}} + \nabla_{\mathbf{u}} \mathbf{u} = -\text{grad } p \quad \text{div } \mathbf{u} = 0.$$

Equivalently, in terms of circulation data $\eta = [\mathbf{u}^\flat] \in \mathfrak{g}^* = \Omega^1(M)/d\Omega^0(M)$ (i.e., η are 1-form only to be evaluated on closed curve), circulation η is governed by the Hamiltonian flow under the Hamiltonian $H(\eta) := \|\eta^\sharp\|^2$. This latter Hamiltonian formulation is given in the sense of the first two items of the following preliminaries:

- Hamiltonian system can be described only using a function H , called Hamiltonian, defined on a Poisson manifold $(P, \{\cdot, \cdot\})$, which is a manifold P together with a Lie algebra $\{\cdot, \cdot\}$ for functions $C^\infty(P)$ with an extra condition that $\{f, \cdot\}: C^\infty(P) \rightarrow C^\infty(P)$ is a derivation. The Hamiltonian flow is given by the vector field as the derivation $-\{H, \cdot\}$.
- The dual vector space of any Lie algebra is automatically a Poisson manifold. This Poisson manifold is foliated into coadjoint orbits. Any Hamiltonian flow stay in the coadjoint orbit containing the initial condition.
- A Poisson preserving action on a Poisson manifold P by a Lie group G induces (locally) a momentum map $J: P \rightarrow \mathfrak{g}^*$, where \mathfrak{g} is the Lie algebra of G . This momentum map is always a Poisson map.

The last item is a key property we will use later.

2. DISCRETE EULER EQUATION

Our discrete Euler equation starts with a finite dimensional subspace $I: \mathfrak{B} \hookrightarrow \mathfrak{g}$, where \mathfrak{B} can be understood as a velocity data sitting on a grid, and I can be thought of as a mimetic (div-free preserving) interpolation operator. The subspace \mathfrak{B} has an induced inner product structure from \mathfrak{g} . The discrete Euler flow is then formulated as an evolution of a velocity vector \mathbf{u} constrained to \mathfrak{B} . Right-translation of $\mathfrak{B} \subset \mathfrak{g}$ to everywhere on G defines a sub-Riemannian structure.

Definition 2 (*I-discrete Euler flow*). *Given a subspace $I: \mathfrak{B} \hookrightarrow \mathfrak{g}$ which defines a sub-Riemannian structure on G , the discrete Euler flow is a sub-Riemannian geodesic on G .*

Previous methods, pioneered by [2], derived this geodesic variational problem with the primary variable being the velocity $\mathbf{u} \in \mathfrak{B}$, which yields a weak formulation of Euler equation on \mathfrak{B} and \mathfrak{B}^* . However, the Poisson structure \mathfrak{g}^* is lost when it is replaced by the finite dimensional space \mathfrak{B}^* . Note that \mathfrak{B} is generally not a Lie subalgebra, and thus \mathfrak{B}^* is not a Poisson space.

We argue that the proper sub-Riemannian geodesic equation should be formulated for the circulation data $\eta \in \mathfrak{g}^*$ in the Hamiltonian formulation. Note that the adjoint $I^\top: \mathfrak{g}^* \rightarrow \mathfrak{B}^*$ allows us to pullback the kinetic energy $H_{\mathfrak{B}^*}$ defined on the finite dimensional space \mathfrak{B}^* to the Poisson space \mathfrak{g}^* .

Proposition 1. *The I-discrete Euler flow is governed by the Hamiltonian flow with Hamiltonian $H = H_{\mathfrak{B}^*} \circ I^\top$ on the Poisson space \mathfrak{g}^* . Any solution to I-discrete Euler flow stays on a coadjoint orbit.*

The property about coadjoint orbit conservation ensures that *I*-discrete Euler flows preserve circulation and vorticity like for the continuous Euler flows.

Finally, to numerically simulate *I*-discrete Euler flow, we consider an auxiliary Poisson space. In our work we take a symplectic space (Σ, σ) , which is a special case of Poisson space. We also let the volume-preserving diffeomorphism group G act on Σ symplectomorphically. Then there is a momentum map $J: \Sigma \rightarrow \mathfrak{g}^*$, which is in particular a Poisson map. By the Poisson structure preservation of J , the Hamiltonian system on \mathfrak{g}^* can be pulled back to Σ .

Theorem 1. *Let $s(t) \in \Sigma$ be the solution to the Hamiltonian system on the symplectic space (Σ, σ) with Hamiltonian $H_\Sigma = H \circ J$. Then $J(s(t)) \in \mathfrak{g}^*$ is the solution to the discrete Euler equation.*

In [1] we set Σ as the phase space for a particle system. The resulting dynamical system becomes an improved version of the Fluid Implicit Particles (FLIP) method. In addition to energy conservation, it has the non-trivial property of coadjoint orbit conservation. In contrast to previous particle-in-grid methods that involve multiple separately introduced interpolation methods and discretized operators, in our method, everything is derived from one single interpolation scheme I . In particular, a high order interpolation scheme I ensures the global order of accuracy for the method. Numerical comparisons demonstrate exact energy and Casimir conservation, and outstanding visuals of energetic vortex interactions.

REFERENCES

- [1] Mohammad Sina Nabizadeh, Ritoban Roy-Chowdhury, Hang Yin, Ravi Ramamoorthi, and Albert Chern. 2024. Fluid Implicit Particles on Coadjoint Orbits. *ACM Trans. Graph.* 43, 6, Article 270 (December 2024), 38 pages. <https://doi.org/10.1145/3687970>
- [2] D. Pavlov, P. Mullen, Y. Tong, E. Kanso, J.E. Marsden, M. Desbrun, Structure-preserving discretization of incompressible fluids, *Physica D: Nonlinear Phenomena*, Volume 240, Issue 6, 2011, Pages 443-458, ISSN 0167-2789, <https://doi.org/10.1016/j.physd.2010.10.012>

Optimal motion from shape change

FLORINE HARTWIG

(joint work with Oliver Gross, Martin Rumpf, Peter Schröder)

Exploiting geometric principles of shape change induced motion, we formulate a variational approach to compute optimal motion paths in world coordinates induced by shape changes in the shape space of deforming bodies. This applies in particular to the motion of snakes. Our continuous model consists of an inner dissipation energy measuring the energy necessary for the shape deformation of the body, e.g. due to metabolism, and an outer dissipation associated with friction measuring the energy dissipated through the motion in the environment. We illustrate how the computation of optimal motion paths can be formulated as the problem of finding sub-Riemannian geodesic paths between two positioned shapes in world coordinates or constraint to period deformations in shape space enabling us to find optimal deformation gaits.

The model is discretized in a consistent way allowing us to compute these geodesic paths numerically. Using this framework, we can simulate motion trajectories of snakes as well as low-dimensional systems such as the Purcell swimmer.

Riemannian metric on the space of positioned shapes. We consider a Riemannian metric on the space of positioned shapes \mathcal{M} . We regard \mathcal{M} as the total space of a principal $SE(3)$ -bundle $\pi: \mathcal{M} \rightarrow \mathcal{S}$ over the space of shapes \mathcal{S} . The projection π maps each element of \mathcal{M} to its underlying shape in \mathcal{S} , identifying configurations that differ only by a rigid body transformation.

For velocities $v \in T_\gamma \mathcal{M}$ the *inner dissipation* is measured as

$$\langle v, v \rangle_{\mathcal{A}} := \int_{\gamma} \frac{1}{2} \text{Hess } W_{\gamma}[\text{Id}](v, v) d\gamma,$$

where W_{γ} is an elastic deformation energy which measures bending and membrane distortion of the shape γ undergoing a deformation [2]. This energy is invariant under rigid body motions. Hence, it is invariant under repositioning of the shapes and constant on the $\text{SE}(3)$ -fibers. The Hessian is symmetric semi positive-definite with a six-dimensional kernel reflecting the infinitesimal rigid body motions.

To account for the *outer dissipation* due to friction we additionally introduce

$$\langle v, v \rangle_{\mathcal{B}} := \int_{\gamma} \frac{1}{2} \langle B_{\gamma} v, v \rangle d\gamma,$$

where $\langle \cdot, \cdot \rangle$ denotes the standard Euclidean metric and B_{γ} accounts for local frictional anisotropies as described in [1]. Specifically, we consider the case where movement in tangential direction causes less dissipation than movement in normal direction.

For $\beta \in \mathbb{R}_{>0}$ the *total dissipation metric* is then given by

$$\langle v, v \rangle_{\mathcal{D}} := \langle v, v \rangle_{\mathcal{A}} + \beta \langle v, v \rangle_{\mathcal{B}}.$$

This gives a well-defined Riemannian metric on \mathcal{M} . For a motion path $\gamma(t)_{t \in [0,1]}$ in \mathcal{M} we define the corresponding path energy

$$\mathcal{E}(\gamma(t)_{t \in [0,1]}) = \int_0^1 \langle \dot{\gamma}(t), \dot{\gamma}(t) \rangle_{\mathcal{D}} dt.$$

Optimal motion paths as sub-Riemannian geodesics. To compute optimal motion paths we have to account for the fact that physical motion is inherently constrained and is restricted to admissible directions dictated by the shape dynamics. The moving organism can only control its shape deformation.

We can decompose the tangent space $T_{\gamma} \mathcal{M} = V_{\gamma} \oplus H_{\gamma}$, where $V_{\gamma} := \ker(d\pi) \subset T_{\gamma} \mathcal{M}$ is the *vertical space* corresponding to infinitesimal rigid body motions and $H_{\gamma} := V_{\gamma}^{\perp}$ is the *horizontal space*.

A consequence of Helmholtz's principle of least dissipation [3] is that the motion induced by infinitesimal shape change should be orthogonal—with respect to the dissipation metric—to the shape-change orbits i.e. the $\text{SE}(3)$ -fibers V_{γ} . This condition introduces a horizontal constraint on the motion, restricting it to the sub-bundle H_{γ} of the tangent space. This formulation transforms the problem of computing optimal motion paths into the computation of sub-Riemannian geodesics. It ensures that the resulting motion path $\gamma(t)$ follows energy-minimizing trajectories within the admissible directions $H_{\gamma(t)}$ for all $t \in [0, 1]$. For a detailed overview of this concept, see for example [4].

Since the inner dissipation remains invariant under $\text{SE}(3)$ variations, the horizontal constraint is defined solely based on the outer dissipation.

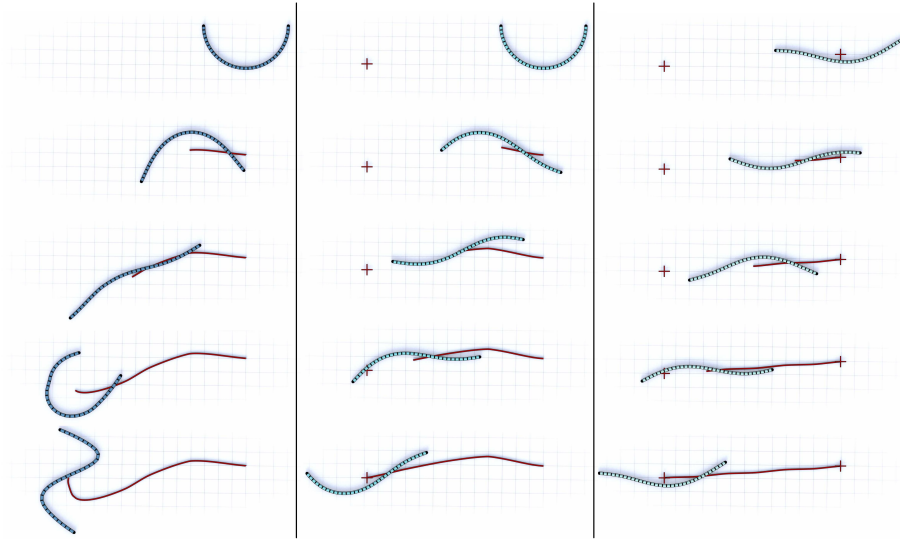


FIGURE 1. The columns show solutions of the three different inverse problems (1), (2), and (3) (from left to right). The red line shows a tracking of the center of mass along the computed path. The red crosses indicate the target displacement that should be achieved in case (2) and (3).

To compute an optimal motion path between two positioned shapes $\gamma_0, \gamma_1 \in \mathcal{M}$ we therefor have to solve the following minimization problem

$$(1) \quad \operatorname{argmin}_{\dot{\gamma} \in H_{\gamma}, \gamma(0)=\gamma_0, \gamma(1)=\gamma_1} \mathcal{E}(\gamma(t)_{t \in [0,1]}).$$

Alternatively, we consider a fixed start shape γ_0 and optimize over end shapes $\gamma_1 \in \mathcal{M}_g$, where \mathcal{M}_g is given by all shapes in \mathcal{M} with a fixed rigid body motion component g relative to γ_0 , i.e.

$$(2) \quad \operatorname{argmin}_{\dot{\gamma} \in H_{\gamma}, \gamma(0)=\gamma_0, \gamma(1) \in \mathcal{M}_g} \mathcal{E}(\gamma(t)_{t \in [0,1]}).$$

Moreover, we compute optimal periodic motions that result in a specified displacement $g \in \operatorname{SE}(3)$. For this we solve the following problem

$$(3) \quad \operatorname{argmin}_{\dot{\gamma} \in H_{\gamma}, \pi(\gamma(0))=\pi(\gamma(1)), \gamma(1)=g(\gamma(0))} \mathcal{E}(\gamma(t)_{t \in [0,1]}).$$

Those periodic paths project onto a closed loop in shape space \mathcal{S} , which may be traversed multiple times. Figure 1 shows exemplary results for all three inverse problems.

REFERENCES

- [1] O. Gross, Y. Soliman, M. Padilla, F. Knöppel, U. Pinkall, P. Schröder, *Motion from shape change*, ACM Trans. Graph. 42.4 (2023): 107-1.
- [2] B. Heeren, M. Rumpf, P. Schröder, M. Wardetzky, *Exploring the geometry of the space of shells*, Computer Graphics Forum. Vol. 33. No. 5. (2014).
- [3] H. Helmholtz, *Zur Theorie der stationären Ströme in reibenden Flüssigkeiten*, Wiss. Abh 1 (1868): 223-230.
- [4] R. Montgomery, *A tour of subriemannian geometries, their geodesics and applications*, No. 91. American Mathematical Soc., (2002).

Generalized Signed Distance

NICOLE FENG

(joint work with Ioannis Gkioulekas and Keenan Crane)

1. THE SIGNED HEAT METHOD FOR GENERALIZED SIGNED DISTANCE

Signed distance functions (SDFs) are essential for many problems. But while it is straightforward to compute the SDF of a closed shape free of defects, it is difficult to robustly compute the SDF of “broken” shapes that have been corrupted by holes, noise, or self-intersections. The core problem is that ordinary signed distance depends on a well-defined notion of inside and outside, which is not well-defined for broken shapes. Our goal is to directly compute signed distance to the true, uncorrupted shape underlying a corrupted observation of the shape.

The *signed heat method (SHM)* achieves this goal [6]. The algorithm takes as input some (possibly broken) oriented curves on a surface domain, or surfaces in \mathbb{R}^3 . The algorithm returns an approximation of the signed distance to the unknown uncorrupted geometry. A significant advantage of the method is that it computes distance to a reconstruction of the shape, without needing to explicitly perform a costly reconstruction and contouring procedure.

The SHM relies on the fact that on Riemannian manifolds, short-time vector diffusion yields parallel transport along shortest geodesics asymptotically as the diffusion time $t \rightarrow 0$ [3, Theorem 2.30]. We infer that diffusing the normal vectors to some source geometry Ω for a small time yields vectors obtained by parallel-transporting normals along a minimal geodesics, and hence conclude that diffused normals will be parallel with the gradient of signed distance. Thus, we simply normalize the diffused normal vectors to obtain a unit-magnitude vector field Y_t that approximates the gradient of a signed distance function, which we obtain by integrating Y_t . Concretely, the algorithm is:

- (1) Solve a vector diffusion equation $\frac{d}{dt}X_t = \Delta^\nabla X_t$, with initial conditions X_0 representing the normal vectors to the source geometry $\Omega \subset M$. In the discrete setting, we use a fixed diffusion time t determined by theoretical [16] and empirical [18] findings.
- (2) Evaluate the vector field $Y_t := X_t / \|X_t\|$.

- (3) Solve a Poisson equation $\Delta\phi = \nabla \cdot Y_t$ on M , with boundary conditions $\frac{\partial\phi}{\partial n} = n \cdot Y_t$ on ∂M .

The algorithm ends up amounting to solving two sparse linear systems. In addition to its simplicity, the SHM offers the following advantages:

- The method is based purely on intrinsic PDEs, and hence can be applied to virtually any discretization of geometry in both two and three dimensions. For example, we apply our method to triangle meshes, polygon meshes, point clouds, voxelized surfaces, tet meshes, and regular grids.
- Further, a careful discretization on triangle meshes using *Crouzeix-Raviart elements* enables a purely intrinsic discretization that enables further robustness to poor triangle quality using *intrinsic triangulations* [7].
- The method robust across a wide variety of errors in both the source and domain geometry.
- The method applies out-of-the-box on non-manifold and/or self-intersecting surface meshes.
- The method can compute unsigned distance simply by changing the direction of the normal vectors diffused in the first step. One can even combine both signed and unsigned distance in a single distance field.
- Because the method is variational, it easily admits additional constraints, such as level set constraints.
- The method yields good surface reconstruction compared to the classic reconstruction algorithms of *winding numbers* [10] and *Poisson Surface Reconstruction* [12].

2. CAN WE IMPROVE UPON THE SIGNED HEAT METHOD?

Given how well the SHM works, a natural question is how to extend the method even further. For example, the method has currently been formulated to solve a global system that amortizes performance over many queries dense in the domain. But we might also aim to develop an output-sensitive version that optimizes for individual pointwise queries — for instance, vision and robotics tasks like path planning and collision detection typically only require signed distance at a single query point at a time.

When the domain $M = \mathbb{R}^n$, one might try to turn the PDE formulation of the SHM into an equivalent integral formulation. Unfortunately, distance functions do not decay at infinity, preventing well-posed integral formulations. Using an integral equation on a finite domain also results in an intractable formulation (because of a recursive term, and an intractable source term.)

On the other hand, there exist “convolutional” methods for surface reconstruction, which estimate an underlying surface by convolving the source geometry against a particular kernel. It turns out that one can also derive convolutional methods for distance, but none can provide generalized signed distance like the SHM. In particular, all such convolutional methods are special cases of one of the

following formulas:

(1)

$$\tilde{d}(x) = -\frac{1}{\lambda} \log \left[\int_{\Omega} h_x(z) \exp(-\varphi_z(x)) dz \right], \text{ or } \tilde{d}(x) = \frac{\int_{\Omega} h_x(z) \exp(-\varphi_z(x)) dz}{\int_{\Omega} \exp(-\varphi_z(x)) dz}$$

with $\Omega \subset \mathbb{R}^n$, $\varphi_z(x) = \|x - z\|_2^2$, $\lambda \rightarrow \infty$. The formulas can be derived either:

- using *viscosity solutions* of the signed eikonal equation, and applying an exponential change of variables, namely a variant of the *Hopf-Cole transformation* (a.k.a. *Cole-Hopf*) [2, 4, 5, 9, 13];
- asymptotic analysis of Laplace-type integrals as $\lambda \rightarrow \infty$ using *Laplace's method* [1, 19], which generalizes the former.

Instances of these formulas are found widely across mathematics, computer science, and engineering, for example:

- Similar asymptotic analysis was used to obtain *Varadhan's formulas* for unsigned geodesic distance [20].
- The “LogSumExp” function (a.k.a. *Kreisselmeier-Steinhauser function*) used as a smooth maximum/minimum operator. Its gradient yields the “softmax” function (a.k.a. *Boltzmann distribution*) which is often used as an activation function in neural networks, is the core of the “attention” mechanism in machine learning, and underlies many kernel regression methods.

While these formulas are perfectly fine in theory, unfortunately they fail for discrete geometry, for example when Ω is sampled as a point cloud. In particular, we always obtain a distance function whose level sets are to the *sampled* geometry, rather than to the *underlying surface* from which the discrete geometry is sampled. Worse, any additional regularization scheme we might try proves futile.

The reason is that as $\lambda \rightarrow \infty$, the approximations in Equation (1) are completely determined by the behavior of either φ or h around the global minimizer of the exponential argument φ , and the global minimizer will be some point in the discrete sampling of Ω , rather than a point on a approximation of the underlying surface. This means that coming up with a robust convolutional distance algorithm is just as difficult as our original problem of computing pointwise distance to some nicely completed surface: if we knew a good φ , we would already have solved our problem (or vice versa). Any scheme we might try for regularizing either the point cloud or exponential kernel amounts to altering the h function which has no effect asymptotically, so we always get the same non-robust behavior. So while we need high λ to get good distance properties, the asymptotic properties of the exponential function rapidly negates the effect of any regularization we might choose. (Note that we can't simply change the function φ inside the exponential, because this would change the metric according to which distance is measured.)

There have been several variants of the general convolutional formulas presented here, and none of them are robust [8, 11, 15, 17]; at best, they sacrifice distance properties by using small λ , to only focus on reconstruction [10, 12, 14]. This is in contrast to the SHM, which because it includes an intermediate step of normalizing

gradient vectors, does obtain good distance accuracy to a good reconstruction of the zero level set.

Ultimately, the non-robust behavior of convolutional distance methods is essentially based on scalar diffusion, in contrast to the more robust vector diffusion of the SHM, which makes full use of the higher-order information available. The asymptotic behavior of the exponential function also underlies the behavior of many other regression methods, including those used for clustering (in which case they are good) and diffusion models (in which case they are bad). It is still unclear how one might extend the SHM to different computational paradigms, or improve upon the SHM, but I think there are still yet-to-be explored connections between distance computation and problems of regression, both “old” or “new”.

REFERENCES

- [1] Belyaev, A., and Fayolle, P.-A. Accuracy improvements for convolutional and differential distance function approximations. arXiv preprint arXiv:2412.09200 (2024).
- [2] Belyaev, A. G., and Fayolle, P.-A. On variational and pde-based distance function approximations. *Comput. Graph. Forum* 34, 8 (dec 2015), 104–118.
- [3] Berline, N., Getzler, E., and Vergne, M. *Heat Kernels and Dirac Operators*, 1 ed. Grundlehren Text Editions. Springer Berlin, Heidelberg, 1992.
- [4] Cole, J. D. On a quasi-linear parabolic equation occurring in aerodynamics. *Quarterly of applied mathematics* 9, 3 (1951), 225–236.
- [5] Crandall, M. G., and Lions, P.-L. Viscosity solutions of hamilton-jacobi equations. *Transactions of the American mathematical society* 277, 1 (1983), 1–42.
- [6] Feng, N., and Crane, K. A heat method for generalized signed distance. *ACM Trans. Graph.* 43, 4 (July 2024).
- [7] M. Gillespie, N. Sharp, K Crane, *Integer coordinates for intrinsic geometry processing*, ACM Transactions on Graphics (TOG) **40** (2021), 1–13.
- [8] Gurumoorthy, K. S., and Rangarajan, A. A schrödinger equation for the fast computation of approximate euclidean distance functions. In *Proceedings of the Second International Conference on Scale Space and Variational Methods in Computer Vision* (Berlin, Heidelberg, 2009), SSVM '09, Springer-Verlag, p. 100–111.
- [9] Hopf, E. The partial differential equation $u_t + uu_x = \mu_{xx}$. *Communications on Pure and Applied mathematics* 3, 3 (1950), 201–230.
- [10] Jacobson, A., Kavan, L., and Sorkine-Hornung, O. Robust inside-outside segmentation using generalized winding numbers. *ACM Trans. Graph.* 32, 4 (July 2013).
- [11] Karam, C., Sugimoto, K., and Hirakawa, K. Fast convolutional distance transform. *IEEE Signal Processing Letters* 26, 6 (2019), 853–857.
- [12] Kazhdan, M., Bolitho, M., and Hoppe, H. Poisson surface reconstruction. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing* (Goslar, DEU, 2006), SGP '06, Eurographics Association, p. 61–70.
- [13] Lipman, Y. Phase transitions, distance functions, and implicit neural representations, 2021.
- [14] Madan, A., and Levin, D. I. W. Fast evaluation of smooth distance constraints on co-dimensional geometry. *ACM Transactions on Graphics* 41, 4 (July 2022), 1–17.
- [15] Rangarajan, A., and Gurumoorthy, K. S. A schrödinger wave equation approach to the eikonal equation: Application to image analysis. In *Proceedings of the 7th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition* (Berlin, Heidelberg, 2009), EMMCVPR '09, Springer-Verlag, p. 140–153.
- [16] S. Schwarz, A. Sturm, M Wardetzky, *Heat kernel asymptotics for scaling limits of isoradial graphs*, *Potential Analysis* (2024), 1–20.

- [17] Sethi, M., Rangarajan, A., and Gurumoorthy, K. The schrödinger distance transform (sdt) for point-sets and curves. In 2012 IEEE Conference on Computer Vision and Pattern Recognition (2012), pp. 198–205
- [18] Sharp, N., Soliman, Y., and Crane, K. The vector heat method. *ACM Trans. Graph.* 38, 3 (June 2019).
- [19] Tibshirani, R. J., Fung, S. W., Heaton, H., and Osher, S. Laplace meets moreau: Smooth approximation to infimal convolutions using laplace’s method, 2024.
- [20] Varadhan, S. R. S. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics* 20, 2 (1967), 431–455.

Geometric Algorithms, Stochasticity and the Fokker-Planck Equation

LETICIA MATTOS DA SILVA

(joint work with Silvia Sellán, Oded Stein, Justin Solomon)

The Fokker-Planck equation is a second-order parabolic partial differential equation (PDE) that describes the time evolution of the probability density function of a process driven by a stochastic differential equation (SDE). My talk describes algorithms that solve different formulations of the Fokker-Planck equation appearing in low-dimensional geometry processing and high-dimensional probabilistic methods and machine learning. Our methods draw from different corners of mathematics, including the theory of viscosity solutions and information geometry, leading to tools with wide-ranging applications, from optimal transport tasks and geometric modeling to image resampling.

Let $\mathcal{M} \subset \mathbb{R}^3$ be a compact surface embedded in \mathbb{R}^3 , possibly with boundary $\partial\mathcal{M}$; take ∇ to be the gradient operator and Δ to be the Laplacian operator, with the convention that Δ is negative semidefinite. We first examine the Fokker-Planck equation as an example of second-order parabolic PDE of the following form:

$$(1) \quad \frac{\partial u}{\partial t} + H(x, \nabla u, u) = \varepsilon \Delta u,$$

where $\varepsilon \geq 0$ and $u(x, t): \mathcal{M} \times [0, \infty) \rightarrow \mathbb{R}$ is an unknown variable function. The function $H(x, \nabla u, u): T^*\mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}$ is the Hamiltonian on $T^*\mathcal{M}$ of

$$(2) \quad \frac{\partial u}{\partial t} + H(x, \nabla u, u) = 0,$$

which is known as the Hamilton-Jacobi equation.

We focus on the Cauchy problem for (1), that is, with a prescribed initial condition $u(x, 0) = u_0(x)$, where $u_0: \mathcal{M} \rightarrow \mathbb{R}$ is some scalar function, and boundary conditions determining the behavior of $u(x, t)$ at $x \in \partial\mathcal{M}$. Our strategy exploits the fact that, under mild assumptions imposed on H , nonlinear systems appearing in implicit time integration for (2) can be understood as the minimizer for a carefully-designed optimization problem. We leverage this fact with a splitting technique and appropriate spatial discretization to propose a convex relaxation for solving parabolic PDE on curved triangle meshes.

Next, we consider the class of self-maps obtained by entropy-regularized transport from a measure to itself. Formally, we define a mirror Schrödinger bridge to

be the minimizer of the KL divergence $D_{\text{KL}}(\mathbb{P} \parallel \mathbb{P}^0)$ over path measures \mathbb{P} with both initial and final marginal distributions equal to a probability measure π on \mathbb{R}^n , where \mathbb{P}^0 is an Ornstein-Uhlenbeck process with noise $\sigma > 0$. In this setting, we examine the Fokker-Planck equation as the following SDE:

$$d\mathbf{X}_t = f_t(\mathbf{X}_t)dt + \sigma d\mathbf{W}_t,$$

where f_t denotes the drift function and \mathbf{W}_t denotes the Wiener process.

We give an efficient algorithm leveraging time-symmetry to *estimate* the solution to mirror Schrödinger bridges with applications to resampling in generative modeling. In more detail, we propose the following iterative scheme:

$$\begin{aligned} \mathbb{P}^{2k+1} &= \arg \min_{\mathbb{P} \in \mathbb{D}(\pi, \cdot)} D_{\text{KL}}(\mathbb{P} \parallel \mathbb{P}^{2k}) && \text{(direct } D_{\text{KL}} \text{ projection)} \\ (3) \quad \mathbb{P}^{2k+2} &= \arg \min_{\mathbb{P} \in \mathbb{S}} D_{\text{KL}}(\mathbb{P}^{2k+1} \parallel \mathbb{P}), && \text{(reverse } D_{\text{KL}} \text{ projection)} \end{aligned}$$

where $\mathbb{D}(\pi, \cdot)$ denotes the space of path measures with initial marginal fixed to be π and \mathbb{S} is the set of time-symmetric path measures with no marginal constraints. We show that, up to a first-order approximation, the reverse D_{KL} projection in (3) admits an analytical solution, allowing for shorter runtime for convergence, and we propose a bound for the difference between the exact solution to the mirror Schrödinger bridge and our estimate.

Please refer to [1, 2] for more comprehensive discussions, further results and applications of the works presented in my talk.

REFERENCES

- [1] L. Mattos Da Silva, O. Stein, J. Solomon, *A framework for solving parabolic partial differential equations on discrete domains*, ACM Transactions on Graphics (2024).
- [2] L. Mattos Da Silva, S. Sellán, J. Solomon, *Through the looking glass: mirror Schrödinger bridges*, arXiv:2410.07003 [cs.LG] (2024).

Flag spaces for Geometric statistics

XAVIER PENNEC

(joint work with Tom Szwagier, Dimbihery Rabenoro)

Generalizing PCA to manifolds: Barycentric Subspaces Analysis. Statistically reduction of the dimension is a key issue in numerous problems. When data belong to a manifold, we first need to define families of parametric subspace in manifolds which could play the role of principal subspaces. Geodesic shooting along the main eigenvectors of the covariance matrix at the Fréchet mean point generates a Geodesic Subspace (GS) in tangent PCA. The point and tangent vectors defining that GS that can also be optimized to best fit the data such as proposed in Principal Geodesic Analysis (PGA) and Geodesic PCA (GPCA). To restore the full symmetry between the parameters, we proposed in [1] to use the *Exponential Barycentric subspace* (EBS) defined as the locus of weighted means of $k + 1$ reference points (with positive or negative weights). The EBS is locally

a stratified spaces of maximal dimension k comprised of critical points of the weighted variance satisfying the barycentric equation $\sum_i \lambda_i \log_x(x_i) = 0$. Its metric completion is called the *affine span* $\text{Aff}(x_0, \dots, x_k)$. Such spaces generalise the notion of Geodesic Subspaces which can be shown to be the limit when reference points coalesce towards a local 1-jet.

Barycentric subspaces and affine spans can naturally be nested by defining an ordering of the reference points. This allows the construction of forward or backward nested sequence of subspaces. However, these methods optimized for one subspace at a time and cannot optimize the explained (or unexplained) variance simultaneously for all the subspaces of the flag. In order to obtain a global criterion, PCA in Euclidean spaces was rephrased in [1] as an optimization on the flags of linear subspaces of the accumulated unexplained variance criterion. This generalizes nicely to flags of affine spans in Riemannian manifolds and gives a particularly appealing generalization of PCA on manifolds, called Barycentric Subspaces Analysis (BSA).

The curse of isotropy: From PCA to Principal Subspaces analysis. Considering PCA as an optimization of flags spaces gives an interesting geometric point of view, even in Euclidean spaces. Indeed, one usually consider the succession of unidimensional eigenmodes for the interpretation of the data in PCA, but more general flags with higher dimensional subspaces naturally arise with the geometric point of view. They correspond to covariance matrix with repeated eigenvalues, in which case eigenspaces are stable but eigenvectors may be freely rotated within each eigenspace. This raises an important issue about the interpretation of PCA modes, called the curse of isotropy [2]: principal components associated with equal eigenvalues show large intersample variability and are arbitrary combinations of potentially more interpretable components.

Most users overlook the problem because empirical eigenvalues are almost surely distinct in practice due to sampling errors with a finite number of samples. In [2], we propose to identify datasets that are likely to suffer from the curse of isotropy by introducing a generative Gaussian model with repeated eigenvalues and comparing it to traditional models via the principle of parsimony. This yields an explicit criterion to detect the curse of isotropy in practice. We notably argue that in a dataset with 1000 samples, all the eigenvalue pairs with a relative eigengap lower than 21% should be assumed equal. This demonstrates that the curse of isotropy cannot be overlooked. In this context, we propose to transition from fuzzy principal components to more interpretable principal subspaces. The final methodology, coined *principal subspace analysis* is extremely simple and shows promising results on a variety of datasets from different fields.

A geometric formulation of CLT for flags. Estimating principal subspaces rather than eigenvectors raises the question of the uncertainty of the estimated flag: with a statistical point of view, one thus looks for confidence regions for principal subspaces. The previous works of Anderson and Tyler were limited to confidence regions on individual eigenvectors or on one single eigenspace at a time. In [3],

we develop an asymptotic method to infer the collection of all principal subspaces together, i.e. the full flag of eigenspaces of this covariance matrix. Our approach is based on the Riemannian homogeneous geometry of the flag manifold. However, even if flags generalize Grassmann and Steifel manifolds, they are generally not symmetric, and the Riemannian logarithm is not known in closed form. To get around this problem, we develop an approach based on the embedding of the flag manifold in the product of Grassmannians, where we can show a central limit theorem and a χ^2 distribution of the Mahalanobis distance.

REFERENCES

- [1] X. Pennec, *Barycentric subspace analysis in manifolds*, **Annals of Statistics** **46(6A)** (2018), p.2711–2746.
- [2] T. Szwagier, X. Pennec, *The curse of isotropy: from principal components to principal subspaces*, **arXiv:2307.15348**, (2023).
- [3] D. Rabenoro, X. Pennec, *A geometric framework for asymptotic inference of principal subspaces in PCA*, **arXiv:2209.02025**, (2022).

Learning geometry – a perspective on reconstruction, isometry, partiality and topology

RON KIMMEL

(joint work with Amir Bracha, Thomas Dages, Yaniv Wolf)

Geometry reconstruction is a pivotal aspect of computer vision, focusing on the extraction of accurate geometric structures in 3D from various 2D visual data representations. Recent advancements have introduced innovative methodologies to enhance the precision and applicability of geometric reconstructions.

The discussion begins with an exploration of stereo vision, particularly addressing the challenge of computational shape reconstruction from random dot stereogram images. In this context, the significance of prior information is emphasized as crucial for achieving accurate reconstructions [1]. This foundational understanding sets the stage for more advanced techniques.

One such technique involves the utilization of Gaussian Splatting (3DGS) for surface reconstruction. This method optimizes a cloud of Gaussian elements to render realistic images from novel viewpoints. However, directly deriving geometries from these Gaussians presents challenges due to potential inconsistencies. To address this, a novel approach employs pre-trained stereo-matching models to extract depth information, leading to smoother and more accurate 3D mesh representations. This process entails rendering stereo-aligned image pairs from the 3DGS model, obtaining depth profiles through stereo analysis, and subsequently fusing these profiles into a cohesive mesh [2]. This strategy has demonstrated superior reconstruction capabilities, particularly in real-world scenarios captured via standard devices.

Complementing this, advancements in partial shape matching have been achieved through the development of the Wormhole Loss criterion [3]. This method addresses the complexities of matching partial and complete surfaces by treating

them as manifolds equipped with geodesic distances. The Wormhole Loss meticulously identifies consistent point correspondences by analyzing intrinsic geodesic distances, distances to surface boundaries, and extrinsic distances in the embedding space. This comprehensive approach enhances the accuracy of partial shape matching, facilitating improved alignment and comparison of geometric structures.

Collectively, these advancements signify substantial progress in the field of geometry reconstruction, offering refined tools and methodologies for both complete and partial shape analysis.

REFERENCES

- [1] Ron Kimmel. 3D Shape Reconstruction from Autostereograms and Stereo. *Journal of Visual Communication and Image Representation*, 13(2):324–333, 2002. <https://www.cs.technion.ac.il/~ron/PAPERS/autostereograms2002.pdf>.
- [2] Yaniv Wolf, Amit Bracha, and Ron Kimmel. GS2Mesh: Surface Reconstruction from Gaussian Splatting via Novel Stereo Views. In *Computer Vision – ECCV 2024*, Lecture Notes in Computer Science, vol. 15147, pp. 207–224. Springer, 2025. https://doi.org/10.1007/978-3-031-73024-5_13.
- [3] Amit Bracha, Thomas Dagès, and Ron Kimmel. Wormhole loss for partial shape matching. In *Advances in Neural Information Processing Systems*, 2024. URL: <https://arxiv.org/abs/2410.22899>.

A Phase-field Approach to Optimal Reinforcing Networks

JOSUA SASSEN

(joint work with Patrick Dondl, Martin Rumpf)

We introduce a phase-field approach to compute optimal reinforcing networks for elastic membranes based on a sharp-interface model introduced by Alberti et al. [1]. Formally, this sharp interface model reads as follows: The goal is to reinforce a two-dimensional membrane described as a bounded Lipschitz domain $\Omega \subset \mathbb{R}^2$ by adding a one-dimensional reinforcement $S \subset \Omega$ in the most efficient way. To this end, for a given exterior load $f \in L^2(\Omega)$, we consider the vertical displacement u of the elastic membrane. The corresponding effect of S on this displacement is described by the energy

$$(1) \quad \mathcal{E}_f(S) := \inf_{u \in C_c^\infty(\Omega)} \frac{1}{2} \int_{\Omega} |\nabla u|^2 \, dx + \frac{r}{2} \int_S |\nabla u|^2 \, d\mathcal{H}^1 - \int_{\Omega} f u \, dx.$$

It has to be maximized to find the optimal reinforcement as it is the negative compliance. Hence, the optimization problem of Alberti et al. [1] consists in finding the best reinforcement S among all networks with total length bounded by a prescribed L , i.e.

$$\underset{S \in \mathcal{A}_L}{\text{maximize}} \quad \mathcal{E}_f(S),$$

with

$$\mathcal{A}_L := \{S \text{ closed connected subset of } \Omega \text{ with } \mathcal{H}^1(S) \leq L\}.$$

Alberti et al. [1] provide additional considerations on a relaxation of this problem to prove the existence of solutions. Our goal is to introduce a phase-field approximation of this problem as well as its numerical implementation.

To this end, we employ an Ambrosio–Tortorelli-type phase-field model with the elastic membrane problem (1) entering as a PDE-constraint. For the diffuse perimeter, we consider the functional

$$\mathcal{P}_\varepsilon(v) := \int_{\Omega} \varepsilon |\nabla v|^2 + \frac{1}{4\varepsilon} v^2 \, dx.$$

In the membrane problem, the reinforcement is achieved via a diffuse indicator function leading to the energy

$$\mathcal{E}_f(v) := \inf_u \int_{\Omega} \frac{1}{2} \left(1 + \frac{r}{\varepsilon} \rho(v) \right) |\nabla u|^2 \, dx - \int_{\Omega} f u \, dx$$

with reinforcement factor $r \in \mathbb{R}_{\geq 0}$ and weighting function $\rho(v) = e^{-\frac{(1-v)^2}{\eta^2}}$ with bump width η . The weighting function ensures that the phase-field only has a reinforcing effect if it is close to one and, thus, prevents smearing out of the phase-field for larger ε (and lower resolutions in the numerical computations). Furthermore, to ensure connectedness of the computed structures and well-posedness of the problem, we use an adapted version of the connectedness constraint introduced by Dondl et al. [2]. This functional becomes

$$\mathcal{C}(v) := \int_{\Omega \times \Omega} H(v(x)) H(v(y)) \text{dist}^{F(v)}(x, y) \, dx \, dy$$

for our Ambrosio–Tortorelli-like phase-fields with weighted distance

$$\text{dist}^F(x, y) := \inf \left\{ \int_c F \, d\mathcal{H}^1 \mid c: [0, 1] \rightarrow \Omega, c(0) = x, c(1) = y \right\}$$

and appropriate helper functions such that $F \equiv 0$ inside the phase, $F > 0$ outside the phase, $H \equiv 1$ inside the phase, and $H \equiv 0$ outside the phase and smooth interpolations in between. With all these ingredients in place, our overall phase-field model becomes

$$\underset{v \in H^1(\Omega)}{\text{minimize}} \quad \lambda \mathcal{P}_\varepsilon(v) + \frac{1}{\varepsilon^\alpha} \mathcal{C}(v) - \theta \mathcal{E}_f(v)$$

with $\alpha = 2 + 2\gamma + \beta$ for some $\beta > 0$ and parameter γ stemming from the helper functions in the connectedness constraint.

We discretize the problem on triangle meshes in a straight-forward fashion by using affine finite elements. For the numerical optimization, we employ the L-BFGS method where we compute the derivative of $\mathcal{E}_f(v)$ using the adjoint method. We explored multiple ways to discretize the connectedness constraint and landed on using the so-called heat method [3] for distance computations, mitigating grid-dependent solutions, especially for lower resolution meshes. In Figure 1, we show an experiment reproducing a result from Alberti et al. [1] under refinement of the interface parameter ε and the spatial resolution.




FIGURE 1. Parameter study with respect to the interface width ε .

REFERENCES

- [1] G. Alberti, G. Buttazzo, S. Guarino Lo Bianco, and É. Oudet, Optimal reinforcing networks for elastic membranes, *Networks and Heterogeneous Media* 14.3 (2019), pp. 589-615. doi: 10.3934/nhm.2019023
- [2] P. W. Dondl, A. Lemenant, and S. Wojtowytsch, Phase field models for thin elastic structures with topological constraint, *Archive for Rational Mechanics and Analysis* 223.2 (2017), pp. 693-736. doi: 10.1007/s00205-016-1043-6
- [3] K. Crane, C. Weischedel, and M. Wardetzky, The heat method for distance computation, *Communications of the ACM* 60.11 (2017), pp. 90-99. doi: 10.1145/3131280

Ray Tracing Harmonic Functions

MARK GILLESPIE

(joint work with Denise Yang, Mario Botsch, and Keenan Crane)

Sphere tracing is a fast and effective algorithm for visualizing surfaces encoded by signed distance functions (SDFs), which have become a centerpiece in a wide range of visual computing algorithms. This talk discusses an analogous algorithm for a completely different class of functions, *harmonic functions*, opening up a whole new set of possibilities. For instance, our new algorithm can be used to directly visualize smooth surfaces reconstructed from point clouds (via Poisson surface reconstruction) or polygon soup (via generalized winding numbers) without performing linear solves or mesh extraction. It can also be used to render nonplanar polygons (including those with holes), and to visualize key mathematical objects, including knots, links, spherical harmonics, and Riemann surfaces (Figure 1).

Sphere Tracing. The sphere tracing algorithm [2] ray traces level sets of a function f with a known *Lipschitz bound*, i.e., a value $C > 0$ such that, for all points x, y we have $|f(x) - f(y)| \leq C\|x - y\|$. At any time t , no point of the isosurface of value f^* is contained within a ball of radius $R = (f(r(t)) - f^*)/C$ —providing a conservative step size for ray tracing. An important special case are signed distance functions, with Lipschitz constant $C = 1$. Similarly, we will truncate a ray by a “safe” sphere—the key difference is that this safe radius is determined by reasoning about harmonic functions, rather than Lipschitz functions.

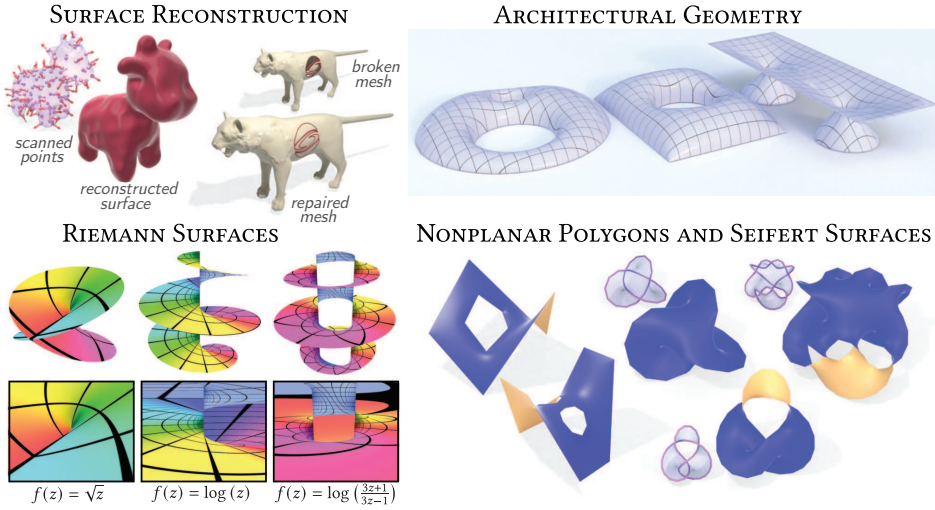


FIGURE 1. Level sets of harmonic functions arise in applications such as surface reconstruction or architectural geometry (*top*), as well as in the visualization of mathematical objects such as Riemann surfaces or Seifert surfaces (*bottom*).

Harnack Tracing. Let $f : B_R(x_0) \rightarrow \mathbb{R}_{\geq 0}$ be a positive harmonic function on the ball $B_R(x_0)$ of radius $R > 0$ centered at $x_0 \in \mathbb{R}^n$. The *Harnack inequality*, provides a conservative upper and lower bound on the value of f for every point $x \in B_R(x_0)$ in terms of the distance $\rho := \|x - x_0\| < R$ and the value $f(x_0)$:

$$\frac{(R - \rho)R^{n-2}}{(R + \rho)^{n-1}} f(x_0) \leq f(x) \leq \frac{(R + \rho)R^{n-2}}{(R - \rho)^{n-1}} f(x_0)$$

We first rewrite the Harnack inequality in a form that makes it more directly useful for our algorithms. Suppose we are standing at a point $x_0 \in \mathbb{R}^3$ and want to know the maximum step size ρ_{lower} we can take in any direction v such that $f(x_0 + \rho v)$ is no smaller than a given lower bound $f_- \in (0, f(x_0))$. In 3D, we have

$$\rho_{\text{lower}} := \frac{R}{2} \left(-(2 + a_-) + \sqrt{a_-^2 + 8a_-} \right), \quad \text{where } a_- := f(x_0)/f_-.$$

Likewise, if we want to avoid exceeding an upper bound $f^+ \in (f(x_0), \infty)$, then we have a maximum step size

$$\rho_{\text{upper}} := \frac{R}{2} \left(a_+ + 2 - \sqrt{a_+^2 + 8a_+} \right), \quad \text{where } a_+ := f(x_0)/f_+.$$

Hence, if our goal is to intersect a level set with value f^* , we take the step size

$$(1) \quad \rho := \frac{R}{2} \left| a + 2 - \sqrt{a^2 + 8a} \right|, \quad \text{where } a := f(x_0)/f^*.$$

This bound provides a conservative step size, whether $f(x_0)$ is above or below f^* .

Now, let $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ be harmonic, and let $r(t)$ be the ray $r(t) := r_0 + t v$ starting at $r_0 \in \mathbb{R}^3$ and moving in direction $v \in \mathbb{R}^3$. For a target value f^* , the Harnack tracing algorithm computes the smallest time $t^* > 0$ such that $f(r(t^*)) = f^*$, i.e., the time where the ray first pierces the surface

$$\mathcal{S} := f^{-1}(f^*) = \{x \in \mathbb{R}^3 \mid f(x) = f^*\}.$$

To do so, we use Equation (1) to find the largest step size ρ for which the ray is guaranteed not to pass through \mathcal{S} . Since the Harnack inequality applies only to *positive* functions, we first “shift” f within a local ball to get a safe step size.

Explicitly, suppose that for any point $x \in \mathbb{R}^3$, we have a radius $R > 0$ and value $c \in \mathbb{R}$ such that $f(y) > c$ for all points $y \in B_R(x)$. To find a step size ρ at some time t , we evaluate Equation (1), but using the shifted function value $f(r(t)) - c$ and shifted target value $f^* - c$. We then increment t by ρ and repeat this process until $f(r(t))$ is sufficiently close to f^* , or t exceeds some maximum time t_{\max} . This simple subroutine can then be incorporated into any rendering algorithm to visualize level sets of harmonic functions, as shown in Figure 1.

REFERENCES

- [1] M. Gillespie, D. Yang, M. Botsch, K. Crane, *Ray Tracing Harmonic Functions*, ACM Transactions on Graphics 43(4):1–18, 2024.
- [2] J. C. Hart. *Sphere Tracing: A Geometric Method for the Antialiased Ray Tracing of Implicit Surfaces*. The Visual Computer 12(10):527–545, 1996.

Conformal Maps and Directional Fields: Can two walk together, except they be agreed?

AMIR VAXMAN

(joint work with Mirela Ben-Chen, Shir Rorberg, Kaibo Hu, Christian Müller, Ofir Weber, and Brad Wang)

The purpose of this talk is to draw ties between notions of discrete conformality and directional fields on triangles meshes, and demonstrate how properties and structure from the continuum are preserved under the discretization. Rather than give concrete answers, this talk hinted at connections between such discretizations as interesting topics of research into a greater unified discrete conformal theory.

Discrete conformal equivalence. Two triangle meshes $\mathcal{M} = \{\mathcal{V}_{\mathcal{M}}, \mathcal{E}, \mathcal{F}\}$ and $\mathcal{N} = \{\mathcal{V}_{\mathcal{N}}, \mathcal{E}, \mathcal{F}\}$ of the same connectivity are discrete conformally-equivalent [7] when there existing a vertex-based (log)-conformal factor $u_{\mathcal{V}}$ such that

$$\forall e \in \mathcal{E}, \quad l_{ij}^{\mathcal{M}} = l_{ij}^{\mathcal{N}} e^{\frac{u_i + u_j}{2}}.$$

$l_{ij}^{\mathcal{M}}$ (resp. $l_{ij}^{\mathcal{N}}$) is the length of edge e_{ij} in \mathcal{M} (resp. \mathcal{N}). Discrete conformality preserves continuous conformal structure by being Möbius invariant, for instance, but does not define a continuous transformation of the triangle mesh.

Piecewise-Möbius maps (PCM). As an attempt to parameterize entire-triangle transformations, Vaxman *et al.* [8] defined per-face Möbius maps M_t , $\forall t \in \mathcal{T}$, that had to conform on mutual edges:

$$\forall e_{ij} \in f \cap g, f, g \in \mathcal{F}, \quad M_f(i) = M_g(i), \quad M_f(j) = M_g(j)$$

In this representation, discrete conformality was characterized by further having $|M'_f(i)| = |M'_g(i)| = \exp(u_i)$ (resp. j), for a local complex parameterization of the flap (f, g) with variable z . This setup allowed for a per-face definition, but still didn't define a continuous interpolator, since the agreement on the vertices of the edge did not extend to the entire edge (otherwise we would have had $M_f = M_g$).

Blended piecewise Möbius maps (BPM). Rorberg *et al.* [6] offered to blend the PCM maps inside every triangle, to obtain a continuous Möbius invariant transformation. Consider a single triangle $t = ijk$, with Möbius map M_t , and maps on the neighboring triangles neighbors M_1, M_2, M_3 on corresponding edges ij, jk, kl . For each edge we compute the *log Möbius error* $l\delta_{t1} = \log(M_t M_1^{-1})$, and similarly for $l\delta_{t2}$ and δ_{t3} . Then, we blend the log-error inside the triangle:

$$l\delta_t(z \in t) = \frac{B_{ij}(z)l\delta_{ij} + B_{jk}(z)l\delta_{jk} + B_{ki}(z)l\delta_{ki}}{B_{ij}(z) + B_{jk}(z) + B_{ki}(z)},$$

Where $B_{ij}(z) = d(z, e_{ij})^{-1}$, and similarly for jk and ki . We then compute:

$$M(z) = \exp\left(\frac{1}{2}l\delta_t(z)\right) M_t$$

The log-exp blending is very common in rotation interpolation [1], and here used to blend matrices in log space, where the underlying assumption is that the error is close enough to the identity to avoid double-cover issues. This method finally provided an interpolant. However, this interpolant did not create continuous conformal maps from discrete conformal maps, but rather just approximately; this might be an impossible task to achieve within one triangle in general, since it is akin to interpolating harmonic maps locally. BPMs are generalizable to curved triangle meshes by a local parameterization of each triangle and its four neighbors; by virtue of the Möbius equivariance of the scheme, the computed maps connect together seamlessly.

Discrete conformal maps and cross fields. In a different perspective on PCM and DC maps, we look at the (complex) derivative of a Möbius map, where $M(z) = \frac{az+b}{cz+d}$, $ad - bc = 1 \Rightarrow M'(z) = (cz + d)^{-1}$. Derivative of holomorphic functions, by way of the Cauchy-Riemann equations, define a field of crosses (a cross-field) on the surface of the triangle. In case a PCM transformation is discrete conformal, for two neighboring triangles f and g on an edge e_{ij} we have $|cz_{f,i|j} + d| = |cz_{g,i|j} + d|$. The means that $Re(\log(M_f(i|j))) = Re(\log(M_g(i|j))) = u_{i|j}$, which is analogous to tangent-continuous discrete 1-forms in $H(curl)$ in discrete exterior calculus [4].

Differential polar representations. It is possible to represent cross-fields (and vector fields in general) by looking at their rotation angles (or: phases), which are the imaginary parts of these logarithms. Such algorithms (e.g., [3, 5]) represented vector fields by single vectors per faces, with (dual) edge-based “connection angles” θ_{fg} , so that for two vectors v_f and v_g on adjacent edges, we have:

$$v_g = e^{i\theta_{fg}} v_f.$$

To counter the discrete Gaussian curvature (=angle defect), one must have:

$$d_0^T \theta = 2\pi I - \kappa_{\mathcal{M}},$$

$\kappa_{\mathcal{M}}$ is the discrete Gaussian curvature of the mesh \mathcal{M} , and I the (vertex-based) index of v . d_0 is the discrete differential operator, where d_0^T essentially sums dual-edge quantities into vertices as cycles around them. Using the polar Hodge decomposition, we have that:

$$\begin{aligned} d_0^T \theta &= d_0^T d_1^T \psi_f + d_0^T *_1 d_t \psi_v + d_0^T h \\ d_0 \theta &= \Delta \psi_v = \kappa_v - \kappa_{\mathcal{M}}. \end{aligned}$$

where ψ_f is the coexact (face-based) potential, ψ_v is the exact (vertex-based) potential, and h a possible harmonic component. This is directly analogous to the Yamabe equation, used in the discrete conformal literature [2].

The discrete Hilbert transform. As another relation between fields in the polar representation and discrete conformal maps, we consider ψ and $\theta = *_1 d_0 \psi_v$ in a flat domain in \mathbb{C} . In this case, ψ_v is just a vertex-based harmonic function: $\Delta_0 \psi_v = 0$. Next consider the resulting face-based cross-field, parameterized by a single phase η_f , where we have that:

$$\begin{aligned} d_1^T \eta_f &= \theta \Rightarrow \\ \Delta_2 \eta_f &= d_1 *_1^{-1} d_1^T \eta_f = d_1 d_0 \psi_v = 0. \end{aligned}$$

And thus, the face-based phase η_f is a face-based harmonic function. This is analogous to phase and log-conformal factor being conjugate harmonic functions for a holomorphic function. Thus, $\psi_v \rightarrow \eta_f$ (the computation of a cross-field from a curvature potential in polar form) is a discrete Hilbert transform.

REFERENCES

- [1] Alexa, M. Linear combination of transformations. ACM Trans. Graph. 21, 3 (July 2002), 380–387.
- [2] Ben-Chen, M., Gotsman, C., and Bunin, G. Conformal flattening by curvature prescription and metric scaling. Computer Graphics Forum 27, 2 (2008).
- [3] Crane, K., Desbrun, M., and Schröder, P. Trivial connections on discrete surfaces. Computer Graphics Forum 29, 5 (2010).
- [4] Desbrun, M., Hirani, A., Leok, M., and Marsden, J. Discrete exterior calculus. preprint, arXiv:math.DG/0508341, 2005.
- [5] Ray, N., Vallet, B., Li, W. C., and Lévy, B. N-symmetry direction field design. ACM Transactions on Graphics 27, 2 (2008).
- [6] Rorberg, S., Vaxman, A., and Ben-Chen, M. Bpm: Blended piecewise möbius maps. Computer Graphics Forum 42, 5 (2023), e14914.

- [7] Springborn, B., Schröder, P., and Pinkall, U. Conformal equivalence of triangle meshes. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 1–11.
- [8] Vaxman, A., and Ben-Chen, M. Dupin meshing: A parameterization approach to planar hex-dominant meshing. Tech. Rep. CS-2015-01, Department of Computer science,technion-IIT, 2015.

Poisson Manifold Reconstruction (beyond co-dimension one)

MICHAEL KAZHDAN

(joint work with Max Kohlbrenner, Sing-Chun Lee, Marc Alexa)

A central task in geometry processing is the reconstruction of a surface from point samples. Most of the work in this area considers two-dimensional surfaces embedded in 3D. Among different ways to define a surface from point samples, level-set techniques and, in particular, (screened) Poisson Surface Reconstruction [1,2] have proven to be particularly resilient to common sampling artifacts such as noise, outliers, and missing data. The reason for this robustness is that they treat the reconstruction problem globally.

The aim of this work is to generalize the class of Poisson Surface Reconstruction methods to oriented d -dimensional manifolds embedded in an n -dimensional space with *co-dimension* $\bar{d} = n - d > 1$. The input is a sampling \mathcal{S} of the manifold, consisting of locations of the samples and (oriented) frames spanning the \bar{d} -dimensional normal space. The output is a vector valued function $F : \mathbb{R}^n \rightarrow \mathbb{R}^{n-d}$ that represents the manifold as a level-set, i.e.

$$\mathcal{M}_{\mathcal{S}} = \{\mathbf{x} \in \mathbb{R}^n : F(\mathbf{x}) = \mathbf{0}\}.$$

As in screened Poisson Surface Reconstruction, we solve for the coordinate functions of F whose gradients span the normal frames prescribed at the sample points. This poses several challenges.

When $\bar{d} > 1$, there is a continuum of frames that can be used to represent the same oriented \bar{d} -dimensional normal space at a point. As the reconstructed manifold should only depend on the oriented normal space at the samples, we require a representation that is agnostic to the particular choice of frame.

For a similar reason, there can be a continuum of solutions F whose gradients span the normal space at the sample points. Often, this reflects the fact that the underlying optimization is non-convex making it hard to find solutions that avoid getting trapped in unwanted local minima.

Our main contributions are choices in representation and optimization that lead to a practical extension of screened Poisson Surface Reconstruction to higher co-dimension, overcoming the mentioned challenges:

- We argue that the right mathematical tool to deal with the multiplicity of frames is *exterior algebra*: we encode the local normal space of the manifold as an exterior product — making the representation independent of the particular choice of frame.

- We show that using the exterior product, the problem of solving for the coordinate functions of F reduces to minimizing a multi-quadratic energy. The continuous approach can be discretized using a finite basis, and we specifically develop a discretization for the case $\bar{d} = 2$.
- We propose a hierarchical approach for minimizing the energy that smooths out the energy landscape at coarser resolutions, keeping the optimization from getting trapped in local minima and consistently returning the same solution, regardless of initial guess.

We demonstrate the efficacy of our approach in reconstructing curves in 3D and surfaces in 4D, analyzing run-time performance, convergence, and stability in the presence of noise. We conclude by offering some discussion and proposing directions for future work.

REFERENCES

- [1] M. Kazhdan, M. Bolitho, and H. Hoppe *Poisson Surface Reconstruction*, Proceedings of the Fourth Eurographics Symposium on Geometry Processing (2006), 61–70.
- [2] M. Kazhdan, and H. Hoppe *Screened Poisson Surface Reconstruction*, ACM Transactions on Graphics **32** (2013), 29:1–29:13.

Consistent Image Editing – Manipulating Images Only Where Change is Needed

NOAM ROTSTEIN

In this work we explore a novel framework for high-fidelity, text-guided image editing that deliberately confines modifications to only those regions where change is necessary, thereby preserving the overall integrity of the original image. Conventional diffusion-based editing techniques often result in global alterations that inadvertently disrupt areas that should remain untouched, leading to inconsistencies and loss of context. To address this challenge, our approach introduces two complementary methods that together redefine image editing as a process of controlled, minimal intervention.

Our first method, *Paint by Inpaint*, is founded on the insight that object removal through inpainting is a relatively well-understood task compared to the more challenging process of object insertion. By leveraging extensive segmentation datasets and state-of-the-art inpainting models, we construct a large-scale dataset of natural images paired with their corresponding object-removed versions. This inversion of the inpainting process enables us to train a diffusion model to “paint” objects back into an image using solely textual instructions. As a result, the model learns to accurately reintroduce objects into their correct spatial contexts, ensuring that the inserted elements blend seamlessly with the existing background. This technique effectively mitigates issues such as residual artifacts and inconsistent texture transitions that typically arise in direct object addition, yielding edited images with remarkable fidelity and contextual coherence.

Complementing this is our second method, *Pathways on the Image Manifold*, which reconceptualizes image editing as a continuous transformation rather than

a one-shot modification. Recognizing that abrupt edits can undermine the natural appearance of an image, we employ a video-generation framework to produce a series of intermediate frames that trace a smooth trajectory from the source image to the desired edited state. In this paradigm, the desired change is first distilled into a temporal editing caption—a concise narrative describing how the image should evolve over time. A video diffusion model then synthesizes a sequence of plausible frames along this continuous path, and an automated frame selection mechanism, guided by a vision-language model, identifies the optimal frame that best captures the intended edit. This gradual approach not only ensures that modifications occur in small, controlled increments but also leverages the inherent temporal coherence of video data to produce a final output that faithfully retains the original image’s structural and aesthetic qualities.

The synergy between *Paint by Inpaint* and *Pathways on the Image Manifold* lies in their shared emphasis on targeted modification. By focusing on altering only the regions that require change, our combined approach achieves a level of editing consistency that is particularly well-suited for applications in augmented reality, interior design, social media, and data augmentation for computer vision. Our extensive evaluations, conducted on both natural and synthetic benchmarks, demonstrate that our methods not only improve objective metrics such as L1/L2 distances and CLIP-based similarity scores but also receive high marks in subjective human assessments. Furthermore, the continuous editing paradigm introduced by the manifold pathway method opens promising avenues for future research in video-based image transformation, enabling the generation of smooth animations and dynamic visualizations that preserve source fidelity.

REFERENCES

- [1] N. Wasserman, *Paint by Inpaint: Learning to Add Image Objects by Removing Them First*, arXiv preprint arXiv:2404.18212, 2024.
- [2] N. Rotstein, *Pathways on the Image Manifold: Image Editing via Video Generation*, arXiv preprint arXiv:2411.16819, 2024.

Diffusion Models for Gaussian Distributions: Exact Solutions and Wasserstein Errors

EMILE PIERRET

(joint work with Bruno Galerne)

Diffusion models have emerged as a powerful framework for generative modeling, particularly in high-dimensional data applications such as image and speech synthesis. These models are based on a progressive transformation of a data distribution into a Gaussian distribution using a forward stochastic differential equation (SDE). The generative process is then performed by reversing this transformation through a backward SDE or an equivalent probability flow ordinary differential equation (ODE) [1]. While these models have demonstrated impressive empirical

performance, their theoretical convergence properties remain an active area of research [4–6]. One of the key challenges in understanding diffusion models lies in analyzing the various sources of errors that arise during the generative process, including initialization, discretization, truncation, and score approximation errors. In this talk, we restrict our study to Gaussian data distributions, which allow for exact analytical solutions of the backward SDE and probability flow ODE. By leveraging the fact that the score function in this setting is a linear operator, we derive closed-form expressions for the evolution of the generated samples and establish that all solutions, including their numerical discretizations, remain Gaussian. This property enables us to compute exact Wasserstein errors induced by different sources of approximation and rigorously compare the performance of various sampling schemes.

Analysis of the error types in diffusion models. A major component of our study focuses on the quantification of errors that impact the generative process. We identify four primary sources of errors:

- The initialization error: approximating the terminal distribution of the forward process with a standard Gaussian introduces an initialization discrepancy that propagates through the backward integration.
- The discretization error: The numerical integration of the SDE or ODE requires discretization, leading to errors that depend on the choice of numerical scheme.
- The truncation error: Due to numerical instabilities, the backward process is often truncated at a small positive time, affecting the final generated distribution.
- The score approximation error: In practical applications, the score function is estimated using a neural network rather than being known exactly, introducing an additional source of discrepancy.

By studying these errors in the Gaussian setting, we derive explicit Wasserstein distances between the true distribution and its sampled approximation under various numerical schemes.

Observation of the exact Wasserstein errors without score approximation. To validate our theoretical findings, we conduct experiments comparing different discretization schemes, including Euler-Maruyama, Exponential Integrator for the backward SDE and Euler and Heun’s methods for the probability flow ODE in terms of exact Wasserstein errors, not considering the score approximation error. Our analysis confirms that Heun’s method offers the most accurate sampling, as empirically observed [2], closely following the theoretical probability flow ODE while maintaining numerical stability. A key contribution of our study is providing an analytical benchmark for evaluating the impact of numerical approximations in diffusion models. The impact of truncation error is minimal for sufficiently small truncation time step, reinforcing the viability of standard practice in diffusion models. This is particularly useful for improving the reliability of diffusion-based generative modeling.

Experimental study of the score approximation. We train a neural network to approximate the score function for a Gaussian distribution of microtextures, named the ADSN model [3]. We observe that stochastic methods, such as Euler-Maruyama, demonstrate resilience to score approximation errors, making them more robust in practical settings.

Implications and Future Work. Our study highlights the importance of precise numerical methods in diffusion models, providing a theoretical framework to guide the selection of sampling schemes. While our analysis is restricted to Gaussian distributions, an important next step is to extend this work to more complex data distributions, such as Gaussian mixture models or real-world datasets. Understanding how these numerical errors behave in non-Gaussian settings will be crucial for further advancements in the field. Additionally, future research should focus on refining neural network architectures for score function estimation to minimize approximation errors. A promising direction is to develop adaptive sampling techniques that dynamically adjust integration steps based on local error estimates, further improving generative accuracy. By bridging the gap between theoretical analysis and practical implementation, our work contributes to the development of more robust and efficient diffusion models. The insights provided in this study offer a strong foundation for optimizing hyperparameters and improving the fidelity of generated samples.

REFERENCES

- [1] 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)
- [2] T. Karras, M. Aittala, T. Aila, S. Laine, *Elucidating the Design Space of Diffusion-Based Generative Models*, Neural Information Processing Systems, (2022)
- [3] B. Galerne, Y. Gousseau, J.-M. Morel *Random Phase Textures: Theory and Synthesis*, IEEE Transactions on Image Processing, (2011)
- [4] V. De Bortoli, *Convergence of denoising diffusion models under the manifold hypothesis*, (2022)
- [5] M. Chen, K. Huang, T. Zhao, M. Wang *Score Approximation, Estimation and Distribution Recovery of Diffusion Models on Low-Dimensional Data*, Proceedings of Machine Learning Research, (2023)
- [6] H. Lee, J. Lu, Y. Tan, *Convergence for score-based generative modeling with polynomial complexity*, Neural Information Processing Systems Workshop on Score-Based Methods, (2022)

Intrinsic approaches to learning and computing on curved surfaces

RUBEN WIERSMA

(joint work with Ahmad Nasikun, Elmar Eisemann, Klaus Hildebrandt)

In this talk, we consider intrinsic approaches to learning and computing on curved surfaces for the following tasks: analyzing 3D shapes using convolutional neural networks (CNNs) and solving linear systems on curved surfaces. We propose more efficient and better performing algorithms for these tasks with intrinsic geometry.

The first two parts consider CNNs on curved surfaces. We would like to find patterns with meaningful directional information, such as edges or corners. On images, it is straightforward to define a convolution operator that encodes directional information, as the pixel grid provides a global reference for directions. Such a global coordinate system is not available for curved surfaces. With Harmonic Surface Networks [1], we apply a 2D kernel to the surface using local coordinate systems. The rotation of these coordinate systems around the normal is ambiguous, which is a problem for pattern recognition. We overcome this ambiguity by computing complex-valued, rotation-equivariant features and transport these features between coordinate systems with parallel transport along shortest geodesics.

Part two introduces DeltaConv [2]. DeltaConv is a convolution operator based on geometric operators from vector calculus, such as the Laplacian. A benefit of the Laplacian is that it is invariant to local coordinate systems. This solves the problem of a missing global coordinate system. However, the Laplacian operator is also isotropic. That means it cannot pick up on directional information. DeltaConv constructs anisotropic operators by splitting the Laplacian into gradient and divergence and applying a non-linearity in between. The resulting convolution operators are demonstrated on learning tasks for point clouds and achieve state-of-the-art results with a relatively simple architecture.

Part three involves solving linear systems on curved surfaces. This is relevant for many applications in geometry processing: smoothing data, simulating or animating 3D shapes, or machine learning on surfaces. A common way to solve large systems on grid-based data is a multigrid method. Multigrid methods require a hierarchy of grids and the operators that map between the levels in the hierarchy. We show that these components can be defined for curved surfaces with irregularly spaced samples using a hierarchy of graph Voronoi diagrams. The resulting approach, Gravo Multigrid [3], achieves solving times comparable to the state-of-the-art, while taking an order of magnitude less time for pre-processing: from minutes to seconds for meshes with over a million vertices.

REFERENCES

- [1] R. Wiersma, E. Eisemann, K. Hildebrandt, *CNNs on Surfaces using Rotation-Equivariant Features*, ACM Trans. Graph. **39** (2020).
- [2] R. Wiersma, A. Nasikun, E. Eisemann, K. Hildebrandt, *DeltaConv: Anisotropic Operators for Geometric Deep Learning on Point Clouds*, ACM Trans. Graph. **41** (2022).
- [3] R. Wiersma, A. Nasikun, E. Eisemann, K. Hildebrandt, *A Fast Geometric Multigrid Method for Curved Surfaces*, ACM SIGGRAPH 2023 Conference Proceedings (2023).

Isoperimetric inequalities: Gaussian, on Gaussian and applications

AGNÈS DESOLNEUX

(joint work with Hermine Biermé, Valentin de Bortoli, Julie Delon,
Antoine Salmona)

In the first part of the talk I presented a joint work with Antoine Salmona, Valentin de Bortoli and Julie Delon, in which we use the Gaussian isoperimetric formula ([1]) to quantify the expressivity of push-forward generative models of the form $X = g(Z)$ where Z follows μ_p the Gaussian standard distribution on \mathbb{R}^p and $g : \mathbb{R}^p \rightarrow \mathbb{R}^d$ is a Lipschitz function (typically a neural network). The law of X is then the push-forward measure $g_{\#}\mu_p$.

To be more precise, the Gaussian isoperimetric formula states that if $A \subset \mathbb{R}^p$ is a Borel set, then

$$\mu_p^+(A) := \liminf_{\varepsilon \rightarrow 0^+} \frac{\mu_p(A_\varepsilon) - \mu_p(A)}{\varepsilon} \geq \varphi(\Phi^{-1}(\mu_p(A))),$$

where φ is the density of the 1d standard Gaussian distribution, Φ is its cumulative distribution and $A_\varepsilon = \{x \in \mathbb{R}^p; d(x, A) \leq \varepsilon\}$. A consequence of this formula is that if g is a Lipschitz function, then for all Borel sets $A \subset \mathbb{R}^d$ we have

$$\text{Lip}(g) (g_{\#}\mu_p)^+(\partial A) \geq \varphi(\Phi^{-1}(g_{\#}\mu_p(A))).$$

This formula can then be used to give bounds on the minimal Lipschitz constant that the function g needs to have in order to generate a Gaussian mixture. The formula also shows that the push-forward distribution $g_{\#}\mu_p$ cannot have a disconnected support. More precisely, we can quantify this thanks to the Gaussian isoperimetric inequality and get the following bound: Let ν be a measure on \mathbb{R}^d on two disconnected manifolds M_1 and M_2 such that $\nu(M_1) = \lambda$ and $\nu(M_2) = 1 - \lambda$, with $\lambda \in [1/2, 1)$, and let $g : \mathbb{R}^p \rightarrow \mathbb{R}^d$ be a Lipschitz function. Then,

$$d_{\text{TV}}(g_{\#}\mu_p, \nu) \geq \int_{\Phi^{-1}(\lambda)}^{d(M_1, M_2)/2\text{Lip}(g) + \Phi^{-1}(\lambda)} \varphi(t) dt$$

where $d(M_1, M_2) = \inf\{\|m_1 - m_2\|; m_1 \in M_1, m_2 \in M_2\}$.

In the second part of the talk, I described how to compute the Lipschitz-Killing (LK) curvatures of the excursion sets of a stationary Gaussian random field. The expected values of these LK curvatures have explicit formulas that are functions what we called geometrical spectral moments (that are explicit functions of the eigenvalues of the Hessian of the covariance matrix). These geometrical spectral moments are ordered and this ordering induces inequalities between the mean LK curvatures, which can be seen as generalized isoperimetric inequalities on Gaussian excursion sets. Indeed in 2D for instance, the formula becomes the following: Let X be a C^2 stationary Gaussian standard random field with covariance function ρ .

Then for a.e. t , denoting the excursion set $E_X(t) = \{x \in \mathbb{R}^2; X(x) \geq t\}$, we have

$$\begin{aligned} C_0^*(X, t) &:= \frac{1}{2\pi} \mathbb{E} \left(\int_{\partial E_X(t) \cap (0,1)^2} \kappa_X(x) \mathcal{H}^1(dx) \right) = \gamma_{\text{TC}} \frac{1}{(2\pi)^{3/2}} t e^{-\frac{t^2}{2}} \\ C_1^*(X, t) &:= \frac{1}{2} \mathbb{E} \left(\int_{\partial E_X(t) \cap (0,1)^2} \mathcal{H}^1(dx) \right) = \sqrt{\gamma_{\text{Per}}} \frac{1}{4} e^{-\frac{t^2}{2}} \\ C_2^*(X, t) &:= \mathbb{E} (|E_X(t) \cap (0,1)^2|) = 1 - \Phi(t), \end{aligned}$$

where Φ is the cumulative of the 1d standard Gaussian distribution, and where denoting γ_1 and γ_2 the two eigenvalues of $\Gamma = -D^2\rho(0)$, we have

$$\gamma_{\text{TC}} = \sqrt{\gamma_1 \gamma_2} \leq \left(\frac{1}{2\pi} \int_0^{2\pi} \sqrt{\gamma_1 \cos^2(\theta) + \gamma_2 \sin^2(\theta)} d\theta \right)^2 = \gamma_{\text{Per}}.$$

Notice that this implies the following isoperimetric inequality:

$$C_0^* \times \varphi(\Phi^{-1}(1 - C_2^*)) \leq \frac{4}{\pi^2} (C_1^*)^2 \times \Phi^{-1}(1 - C_2^*).$$

In 2D, the interpretation of the LK curvatures is the following: C_2^* is the volume of the excursion set in $U = (0,1)^2$, C_1^* is half its perimeter and C_0^* is its total curvature (divided by 2π), that is related to the Euler characteristic thanks to the Gauss-Bonnet Theorem.

We used these results to estimate the anisotropy of a Gaussian random field being only given one observed excursion set. We also extended these results to 3D Gaussian field. This is part of a joint work with Hermine Biermé.

REFERENCES

- [1] Sudakov, V. N. and Tsirelson, B. S. Extremal properties of half-spaces for spherically invariant measures. *Journal of Soviet Mathematics*, 1978.
- [2] A. Salmona, V. de Bortoli, J. Delon and A. Desolneux, *Can Push-forward Generative Models Fit Multimodal Distributions?*, Conference NeurIPS, 2022.
- [3] H. Biermé and A. Desolneux, *The anisotropy of 2D or 3D Gaussian random fields through their Lipschitz-Killing curvature densities*, to appear in *Annals of Applied Probability*, 2025.

Convergent Plug and Play methods for image restoration

NICOLAS PAPADAKIS

(joint work with Antonin Chambolle, Samuel Hurault, Arthur Leclaire,
Jean Prost, Marien Renaud)

In this talk, we present a class of convergent algorithms to solve inverse problems in imaging using explicit and non-convex regularizations learned through deep neural networks. We consider a general observation model of the form $y = Ax^* + \xi$ where $y \in \mathbb{R}^m$ is the observed degraded image, $x^* \in \mathbb{R}^n$ is the true image we want to approximate, $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear degradation operator, and $\xi \sim \mathcal{N}(0, \nu^2 \text{id})$ is

an additive Gaussian noise representing measurement uncertainty. This problem is classically addressed by solving an optimization problem of the following form:

$$(1) \quad x^* \in \arg \min_{x \in \mathbb{R}^n} \lambda f(x) + g(x),$$

where f is a term that attaches to the degraded observation, g is a regularization term, and $\lambda > 0$ is a regularization parameter that weights the importance of the two terms. The regularization term allows for the inclusion of prior information about the regularity of the desired solution. A long-standing problem is to design functions g that reflect relevant regularity on x . Recent deep neural network models have produced remarkable results in various imaging applications. It remains a challenge to determine how to best incorporate deep learning models into a regularization function g that has the necessary properties to allow the use of established optimization techniques.

Plug-and-Play Regularization. The minimization of (1) is generally performed using first-order optimization algorithms, known as proximal algorithms. These algorithms operate individually on the two terms via two operators: the gradient descent operator $\text{id} - \tau \nabla f$ and the proximal operator $\text{Prox}_{\tau f}$. For example, the proximal gradient descent (PGD) algorithm alternates between a proximal operation and a gradient descent step on either of the functions f and g :

$$(\text{PGD}) \quad x_{k+1} = \text{Prox}_{\tau g} \circ (\text{id} - \tau \lambda \nabla f)(x_k) \quad \text{or} \quad x_{k+1} = \text{Prox}_{\tau \lambda f} \circ (\text{id} - \tau \nabla g)(x_k).$$

Plug-and-Play (PnP) methods [8, 10] establish a connection between these optimization algorithms and deep neural network models through the introduction of an image denoiser D_σ . Their derivation is based on the theoretical analysis of optimal MAP and MMSE denoisers. Given a generic denoiser D_σ , for example, a deep neural network pre-trained to denoise a Gaussian noise-degraded image with standard deviation σ , these methods use D_σ in place of a descent operator on the regularization term g . In RED algorithms [8], the denoiser replaces the gradient descent operator: $D_\sigma = \text{id} - \nabla g$. In PnP algorithms [10], the denoiser replaces the proximal operator: $D_\sigma = \text{Prox}_{\tau g}$. The PGD algorithm thus becomes

$$(\text{PnP-PGD}) \quad x_{k+1} = D_\sigma \circ (\text{id} - \tau \lambda \nabla f)(x_k)$$

$$(\text{RED-PGD}) \quad x_{k+1} = \text{Prox}_{\tau \lambda f} \circ (\text{id} - \tau (\text{id} - D_\sigma))(x_k)$$

State-of-the-art results for various inverse problems have been obtained with these processes, that consider a simple Gaussian denoiser to restore a variety of inverse problems. However, in PnP and RED frameworks, the denoiser is used with few theoretical guarantees of convergence. Indeed, since a generic denoiser does not generally express itself exactly as a proximal operator or a gradient, convergence results do not easily follow. Most convergence results limit their analysis to non-expansive denoisers, which implies a significant drop in performance. Moreover, as regularization is only implicitly included via the denoising operation, (PnP-PGD) and (RED-PGD) algorithms do not minimize an explicit functional. This strongly limits the interpretability of the result and the numerical control of the algorithm.

Gradient step denoiser. We introduce a Gradient-Step denoiser [2] that is expressed exactly as a gradient descent step on a differentiable potential $g_\sigma : \mathbb{R}^n \rightarrow \mathbb{R}$:

$$(2) \quad D_\sigma = \text{id} - \nabla g_\sigma$$

We choose g_σ in the form $g_\sigma(x) = \frac{1}{2}\|x - N_\sigma(x)\|^2$, with $N_\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^n$ being a C^1 class neural network. With this formulation, we can parameterize N_σ with any differentiable neural network architecture that has proven effective for image denoising, such as a UNet. Despite being constrained to be a gradient field, D_σ performs as well as the same unconstrained UNet denoising network.

With the gradient step denoiser (2), the (RED-PGD) algorithm takes the form of an optimization algorithms to optimize $\lambda f + g_\sigma$. Since the function g_σ is non-convex, the proof of convergence of the algorithm relies on known results of in non-convex optimization [1] showing convergence of iterates to critical points of (1). Experiments in deblurring and super-resolution confirm the theoretical convergence results and show that the algorithm achieve state-of-the-art performance.

Proximal Denoiser for Convergence. To prove the convergence of PnP methods, we propose a second denoiser trained to take exactly the form of a proximal operator [3]. To do this, we demonstrate that, for g_σ with L -Lipschitz gradient and $L < 1$, the *Gradient-Step* denoiser (2) is a proximal operator, that is,

$$D_\sigma = \text{Id} - \nabla g_\sigma = \text{Prox}_{\phi_\sigma}$$

for an explicit $M = \frac{L}{L+1}$ -weakly convex function $\phi_\sigma : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$.

With this proximal denoiser, the (PnP-PGD) algorithm takes the form of a true proximal gradient descent algorithm to optimize $\lambda f + \phi_\sigma$. Note that D_σ is expressed as $\text{Prox}_{\phi_\sigma}$ and not as $\text{Prox}_{\tau\phi_\sigma}$ (with a time step $\tau > 0$) as used in the (PGD) algorithm. This involves a fixed time step $\tau = 1$ and constrains the choice of the regularization parameter λ to ensure the convergence of the algorithm. We then propose an α relaxation of the PGD algorithm [4]. When the convergence constraint of the PGD algorithm restricts the sum of the Lipschitz constant of $\lambda \nabla f$ and the weak convexity constant M of ϕ_σ , the α PGD algorithm, on the other hand, limits the product of these two terms. This gives freedom for the choice of λ when M approaches 0. Experiments demonstrate the effectiveness of these PnP algorithms. In particular, PnP- α PGD remedies the notable performance loss of PnP-PGD attributable to the restriction on the value of the parameter λ .

Stochastic Denoising regularization. PnP and RED approaches rely on a non-standard use of a denoiser D_σ , trained on noisy images but applied on images that are less and less noisy along the iterations. This contrasts with recent algorithms based on diffusion models [5], where the denoiser is applied only on re-noised images. Generalizing the gradient descent version of the RED method [8], we introduce the Stochastic deNOising REgularization (SNORE) algorithm [6], which applies the denoiser only on images with noise of the adequate level:

$$(\text{SNORE}) \quad x_{k+1} = x_k - \tau_k \left(\lambda \nabla f(x_k) + \frac{x_k - D_\sigma(x_k + \sigma \xi_k)}{\sigma^2} \right),$$

for a random realization $\xi_k = \mathcal{N}(0, \text{Id})$. This is a stochastic gradient descent algorithm targeting a critical point of the non-convex function $\lambda f + \tilde{g}_\sigma$ for the explicit stochastic regularization defined as

$$\tilde{g}_\sigma(x) = \frac{1}{\sigma^2} \mathbb{E}_{\xi \sim \mathcal{N}(0, \text{Id})} g_\sigma(x + \sigma \xi).$$

A convergence analysis of this algorithm and its annealing extension are presented. The method is competitive with respect to state-of-the-art methods, while presenting improved performance in terms of visual quality thanks to the noise injection inspired by diffusion models.

We finally introduce an equivariant regularization [7] encompassing stochastic versions of RED algorithms such as SNORE and the equivariant scheme of [9].

REFERENCES

- [1] H. Attouch, J. Bolte, B. F. Svaiter, *Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized gauss-seidel methods*. Math. Programming, **137** (2013), 91–129.
- [2] S. Hurault, A. Leclaire, N. Papadakis, *Gradient step denoiser for convergent plug-and-play*, International Conference on Learning Representations (2022).
- [3] S. Hurault, A. Leclaire, N. Papadakis, *Proximal denoiser for convergent plug-and-play optimization with nonconvex regularization*, International Conference on Machine Learning (2022).
- [4] S. Hurault, A. Chambolle, A. Leclaire, N. Papadakis, *Convergent plug-and-play with proximal denoiser and unconstrained regularization parameter*, Journal of Mathematical Imaging and Vision **66** (2024), 616–638.
- [5] J., Ho, I. Jain, P. Abbeel, *Denoising diffusion probabilistic models*. Advances in Neural Information Processing Systems, **33** (2020), 6840–6851.
- [6] M. Renaud, J. Prost, A. Leclaire, N. Papadakis, *Plug-and-Play image restoration with Stochastic deNoising REgularization*, International Conference on Machine Learning (2024).
- [7] M. Renaud, A. Leclaire, N. Papadakis, *Equivariant Denoisers for Image Restoration. International Conference on Scale Space and Variational Methods* (2025).
- [8] Y. Romano, M. Elad, P. Milanfar, *The little engine that could: Regularization by denoising (RED)*. SIAM Journal on Imaging Sciences, **10** (2017), 1804–1844.
- [9] M. Terris, T. Moreau, N. Pustelnik, J. Tachella, *Equivariant plug-and-play image reconstruction*, Conference on Computer Vision and Pattern Recognition (2024).
- [10] S. V. Venkatakrishnan, C. A. Bouman, B. Wohlberg, *Plug-and-play priors for model based reconstruction*. In IEEE Global Conference on Signal and Information Processing (2013).
- [11] P. Vincent, *A connection between score matching and denoising autoencoders*. Neural computation, **23** (2011), 1661–1674.

Geometry and computational design of strip structures

HELMUT POTTMANN

(joint work with Yu-Chou Chiang, Martin Kilian, Daoming Liu, Christian Müller, Davide Pellis, Florian Rist, Eike Schling, Johannes Wallner, Bolun Wang, Hui Wang)

The fabrication of gridshells from straight or circular flat strips of bendable material leads to remarkable geometric structures. With further practical restrictions, such as structures whose strips are tangential or normal to an underlying reference surface, we obtain multiple relations to classical and discrete differential geometry. The design space includes surfaces with a constant ratio of principal curvatures, principal symmetric nets and linear Weingarten surfaces [3]. The latter occur as kinetic structures that can be deployed from a collapsed straight or circular state [1]. The arrangement of strips in form of 3-webs or 4-webs leads to various challenging and largely unsolved geometric problems concerning webs of asymptotic and geodesic curves on surfaces [2, 4, 5].

REFERENCES

- [1] D. Liu, D. Pellis, Y.-C. Chiang, F. Rist, J. Wallner, H. Pottmann, *Deployable strip structures*, ACM Trans. Graphics **42** (2023), 103:1–103:16.
- [2] C. Müller, H. Pottmann, *The Geometry of Discrete Asymptotic-Geodesic 4-Webs in Isotropic 3-Space*, Monatsh. Math. **203** (2024), 223–246.
- [3] D. Pellis, H. Wang, M. Kilian, F. Rist, H. Pottmann, C. Müller, *Principal symmetric meshes*, ACM Trans. Graphics **39** (2020), 127:1–127:17.
- [4] E. Schling, H. Wang, S. Hoyer, H. Pottmann, *Designing asymptotic geodesic hybrid gridshells*, Comp. Aided Design **152** (2022), 103378.
- [5] B. Wang, H. Wang, E. Schling, H. Pottmann, *Rectifying strip structures*, ACM Trans. Graphics **42** (2023), 256:1–256:19.


Geometric Optimization for Plasma Fusion

ZORAH LÄHNER

Due to reach of scientific break-even of a plasma fusion reactor recently, which means more energy was produced during the fusion than was necessary to start it and keep running, the hope for efficient, clean energy through plasma fusion is again rising. There exist several conceptually different designs for plasma fusion reactors, one of which is the **stellarator**. One of the defining properties of stellarators, in comparison to the more advanced tokamaks, is their complex plasma flow. The complexity allows for configuration that are very efficient and stable but at the same makes the simulation of its properties and the optimization harder.

The plasma flow is guided and kept in a vacuum by a collection of magnetic coils $\{C_i\}_i$ which define a magnetic field M along which the plasma flows in a toric shape. While the plasma flow happens on the entire inside of this shape, all its properties and whether the flow leads to an equilibrium is defined by its boundary T . In theory the dependency of the flow is completely on C_i , however, the existing

simulations can only compute the plasma flow and magnetic field M from T and through a second simulation the magnetic coils C_i can be optimized from M . The exact shape of T is normally described in toroidal coordinates (ζ, η, ϕ) which describes the deformation along each angle along the primary and secondary torus axes with the following formula (see inset figure):



$$R = \frac{\sinh \zeta}{\cosh \zeta - \cos \eta}, \quad Z = \frac{\sin \eta}{\cosh \zeta - \cos \eta}.$$

The advantage of this representation is its separability under the toroidal harmonics which means the deformation can be easily written in a set of coefficients $a \in \mathbb{R}^k$ using the k lowest frequency toroidal harmonics. A low number for k is required for simulations to converge systematically, however, this introduces several symmetries in the solution and it is unknown whether these are beneficial.

There are several torus configurations known which lead to high reaction efficiency and have a very stable plasma flow, however, another crucial aspect comes from the engineering constraints of the coils and the rest of the reactor, like the complexity of the coils themselves, the sensitivity of the plasma flow stability to inaccuracies from manufacturing, and the manufacturing costs in general. Only very few configurations which do well on all evaluation metrics exists and none of them are extraordinary. The question whether a highly efficient and stable but at the same buildable plasma flow for stellarators exist remains completely open.

Apart from several open questions on the engineering side, there are aspects in which machine learning might make important contributions in the exploration of the huge plasma flow configuration space:

- Make simulations more stable through machine learning prediction which can detect problems in convergence.
- Learn a meaningful space of deformations for T which excludes unnatural solutions, for example with self-intersections or very uneven slice areas which are known to not lead to an equilibrium.
- Learn to make meaningful predictions of new deformations T which are likely to show interesting behaviour.
- Generate a differentiable surrogate model for the flow simulation which would allow a more direct exploration of the configuration space.

Rethinking Geometry Processing via Neural Surfaces

NILOY J. MITRA

(joint work with Romy Williamson)

Neural surfaces (e.g., neural map encoding [1], deep implicits, and neural radiance fields) have recently gained popularity because of their generic structure (e.g., multi-layer perceptron) and easy integration with modern learning-based setups. Traditionally, we have a rich toolbox of geometry processing algorithms designed

for polygonal meshes to analyze and operate on surface geometry. However, neural representations are typically discretized and converted into a mesh before applying any geometry processing algorithm. This is unsatisfactory and, as we show, unnecessary. This talk will describe our recent work [2] on developing geometry processing tools that act directly on neural representations without unnecessary discretization or meshing. I will also list open questions and discuss how we may approach them.

REFERENCES

- [1] L. Morreale, N. Aigeman, V. Kim, N. J. Mitra, *Neural Surfaces Maps*, IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021.
- [2] R. Williamson, N. J. Mitra, *Neural Geometry Processing via Spherical Neural Surfaces*, Computer Graphics Forum (Proc. Eurographics), 2025.

Representations of Surface Meshes for Optimization, Inverse Problems, and Learning

NICHOLAS SHARP

(joint work with Tianchang Shen)

Surface meshes are widely used in geometric computing to represent surfaces embedded in 3 dimensional space (or equivalently objects, as the boundary of a sub-volume of space). Many formulations and algorithms are available for the forward analysis of such surfaces, but it is much less clear how to optimize for a surface, or equivalently solve inverse problems where the surface is unknown—the same setting arises in geometric machine learning. Put differently, it is straightforward to describe problems where a surface is the *input*, but much less straightforward to describe problems where a surface is the *output*. This work considers two classes of strategies for such problems.

The first class is isosurfacing-based methods. The main idea of isosurfacing is to represent a surface mesh as the 0 levelset of a scalar function over space: by evolving this scalar function, we implicitly evolve the corresponding surface. Importantly, we are interested in not just the continuous surface, but its discretization, and the particular choice of discretization has significant impact for numerics and modeling in downstream tasks. Most commonly, the Marching Cubes algorithm implies a particular discretization, by evaluating the scalar function at a regular lattice, and extracting a triangle mesh whose vertices lie along the lattice edges. In [1], this approach is generalized to allow the lattice grid and extracted vertex locations to warp and deform via a carefully constructed parameterization; this additionally flexibility allows for significant improvements in the quality of the extracted discrete mesh. However, the image of potential meshes that can be extracted by this method is still greatly limited, to a narrow family of lattice-like meshes; it certainly does not parameterize the full space of surface meshes as possible outputs.




FIGURE 1. A halfedge mesh represents discrete mesh connectivity via a set of halfedges, or equivalently face-sides. Each halfedge is related to two others: a *next* and a *twin*; these relationships can be viewed as permutations following a few simple properties.

The second class of approaches are graph-based, treating mesh representation akin to graph representation. The work of [2] begins with a standard graph-representation approach, embedding the vertices of a mesh in high-dimensional space, and connecting with an edge only those vertices which are sufficiently close according to some distance metric. To represent face connectivity, an approach inspired by so-called halfedge mesh representation is developed, observing that after some manipulations, representing the connectivity of a halfedge mesh is equivalent to representing an ordering of the neighbors of each vertex. This ordering is parameterized, and used as a continuous representation for surface mesh connectivity, particularly in the context of machine learning. A strength of this approach is that it can represent meshes of general connectivity, without restriction to a narrow family of structures, and it even guarantees manifold connectivity. However conversely, there is the potential to generate nonsensical geometrically tangled meshes, and the representation is not differentiable in the continuous geometric sense.

It remains to see what approaches for representing surfaces will prove most effective in applications such as inverse problems and machine learning, which have surprisingly similar needs.

REFERENCES

- [1] Shen, Tianchang and Munkberg, Jacob and Hasselgren, Jon and Yin, Kangxue and Wang, Zian and Chen, Wenzheng and Gojcic, Zan and Fidler, Sanja and Sharp, Nicholas and Gao, Jun, *Flexible Isosurface Extraction for Gradient-Based Mesh Optimization*, ACM Trans. Graph., 42 (2023), 0730-0301.
- [2] Tianchang Shen and Zhaoshuo Li and Marc Law and Matan Atzmon and Sanja Fidler and James Lucas and Jun Gao and Nicholas Sharp, *SpaceMesh: A Continuous Representation for Learning Manifold Surface Meshes*, SIGGRAPH Asia 2024 Conference Papers, 10.1145/3680528.3687634.

Smoothed L_0 Optimization

ROI PORANNE

The L_0 norm counts the number of non-zero elements in a vector and naturally produces sparse solutions. However, the inherent discontinuous nature of the L_0 norm leads to combinatorial optimization problems that are notoriously difficult. The key insight across our works is that many problems in geometry can be naturally formulated as an L_0 optimization problem. We further rely on a common pattern, which is to replace the L_0 norm with a smooth approximation that is gradually sharpened during the optimization process, thereby transforming a highly non-convex problem into a sequence of more tractable ones. It has been applied for multiple applications in geometry processing for finding optimal cuts for UV mapping [4], generating panel layouts for architectural design [2], and for robotic motion planning [3].

The approach, also known more generally as homotopy optimization or the continuation method [1], enables employing traditional smooth optimization techniques. To demonstrate, consider the problem of UV mapping. Simplistically, the problem is to find a mapping of a triangle mesh to the plane that minimizes distortion. The mapping does not have to be continuous; edges can be *cut* to introduce a discontinuity, which enables finding lower distortion mappings. However, the number of cut edges should be minimized as well. The problem then boils down to striking a balance between minimizing distortion and minimizing the number of cuts. While the former is a continuous function, the latter is discrete. In [4] we address this by treating the mesh as a *polygon soup*, and attempting to find a mapping that minimizes distortion, but also attempts to minimize the difference between cut edge pairs *in the L_0 sense*. We define this edge *separation* measure with a smooth approximation of the L_0 penalty:

$$\|t\|_\delta = \sum \frac{t_i^2}{t_i^2 + \delta}$$

The optimization is performed using Newton's method, and is visualized in a user interface. During optimization, the user can observe the progress of the optimization, and adapt several optimization parameters as they see fit. Most importantly is the δ parameter, which dictates the shape of the approximation, and how similar it is to the true L_0 norm. This formulation relaxes the classical binary notion of seams: instead of tagging each edge as either a seam or a regular edge, we allow it to be something in-between during the optimization. As the algorithm converges, the edges come closer to either becoming seams or regular edges, until they are close enough to be "snapped" to one of the two binary choices.

In another work on architectural panelization [2], we approximate surfaces with planar, spherical, or cylindrical panels through continuous deformation. The challenge is to encourage the formation of panels while minimizing distortion. We formulate this through minimizing the difference between the normals of adjacent faces in the L_0 sense, and rely on a similar approximation. Again, as δ decreases,

this function approaches the L_0 norm, encouraging sparsity in the differences between normals of adjacent faces, resulting in well-defined panels.

Finally, in [3], we developed a toolpath generation technique that for hot-wire cutting material with a bendable wire. There, we quantify the efficiency of a cut using a metric based on robust partial surface matching. Parts of a cut that are not sufficiently close to the target shape receive zero reward, which is naturally expressed through an L_0 -like norm again.

A critical aspect of these approaches is determining when and how to decrease the smoothing parameter. In our interactive applications, we allow users to control this parameter, enabling them to explore the solution space and make design decisions. Alternatively, we can employ adaptive strategies that decrease the parameter based on optimization progress, tightening the approximation as the solution converges. For example, in our work on robotic hot-wire cutting, we gradually decrease δ whenever progress—as measured by the difference in objective values between subsequent iterations—falls below a given threshold.

The interactive nature of our approach allows users to guide the optimization through additional constraints or objectives. For UV mapping, users can interactively move vertices, cut mesh parts, join seams, or control the placement of the parameterization patches. In panelization, users can pin vertices and encourage the merging and splitting of panels. This interactivity addresses a common challenge in design workflows, where not all objectives can be explicitly formulated, and user judgment plays a crucial role.

The methodology presented here represents a powerful framework for addressing discrete optimization problems in geometry processing and robotics, particularly when interactivity and user control are essential. Future directions include exploring acceleration techniques for larger problems, incorporating additional constraints such as global overlap prevention for UV mapping, and extending the approach to new application domains.

REFERENCES

- [1] E.L. Allgower, K. Georg, *Introduction to Numerical Continuation Methods*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2003).
- [2] E. Jadon, B. Thomaszewski, A.A. Apolinarska, R. Poranne, *Continuous deformation based panelization for design rationalization*, SIGGRAPH Asia 2022 Conference Papers **41** (2022), 1–8.
- [3] S. Duenser, R. Poranne, B. Thomaszewski, S. Coros, *Robocut: Hot-wire cutting with robot-controlled flexible rods*, ACM Transactions on Graphics (TOG) **39** (2020), 98:1–98:15.
- [4] R. Poranne, M. Tarini, S. Huber, D. Panozzo, O. Sorkine-Hornung, *Autocuts: simultaneous distortion and cut optimization for UV mapping*, ACM Transactions on Graphics (TOG) **36** (2017), 1–11.

Failures and Successes Simulating Wrinkles on Surfaces

ETIENNE VOUGA

(joint work with Zhen Chen, Melina Skouras, Danny Kaufman)

Fine wrinkles on surfaces are ubiquitous in nature, with everyday examples including the wrinkling of cloth, skin, and mylar balloons, or patterns of dunes in the desert. These wrinkle patterns change over time, but notice that the coupling between the dynamics of the wrinkles and the substrate is largely one-way: changes to the substrate geometry change the wrinkle patterns, but the *macroscale* behavior of thin objects only depends on averaged statistical properties of the wrinkle pattern (the wrinkle orientation, frequency, and amplitude; adding corrugations to a thin shell increases its bending resistance, for instance) but does not depend on the precise wrinkle geometry or location.

This observation motivates a two-scale representation of a wrinkled surface as a wrinkle-free *base surface* $\mathbf{r}_b : \Omega \rightarrow \mathbb{R}^3$, for $\Omega \subset \mathbb{R}^2$, on top of which wrinkles are added as a displacement field parameterized by amplitude $a : \Omega \rightarrow \mathbb{R}$ and phase $\phi : \Omega \rightarrow S^1$. In this talk we consider simple, sinusoidal wrinkle profiles, with the final wrinkled surface given by

$$\mathbf{r}_m = \mathbf{r}_b + a \cos(\phi) \hat{\mathbf{n}}_b,$$

where $\hat{\mathbf{n}}_b$ is the surface normal to \mathbf{r}_b ; but note that in practice wrinkles also induce a tangential displacement that depends on phase. The base surface \mathbf{r}_b can be discretized as a coarse triangle mesh \mathcal{T} , i.e. a piecewise-linear function interpolating the mesh vertex positions. That leaves the question of how to represent the wrinkle field variables. Some desiderata include:

- it must be possible to represent fine wrinkles on the surface, with multiple wrinkle periods per triangle. This requirement precludes a direct piecewise-linear discretization of ϕ as a phase prescribed on the vertices of \mathcal{T} ;
- it must be possible to represent singularities in the wrinkle field: points at which $a \rightarrow 0$, $\|\mathrm{d}\phi\| \rightarrow \infty$, and phase is undefined.
- since our goal is to simulate dynamics, there must exist a well-defined and differentiable map from the discrete degrees of freedom \mathbf{q} to continuous functions a, ϕ .

1. FREQUENCY REPRESENTATION

Wrinkle variables can be loosely taxonomized as *fast* or *slow*. Fast variables change significantly within one wrinkle period, whereas slow variables can be considered constant within one wrinkle period (away from singularities). Example of fast variables include: the wrinkle displacement $a \cos \phi$ itself and the wrinkle phase ϕ . Slow variables include the wrinkle amplitude a as well as the frequency $\omega = \mathrm{d}\phi$.

The above observations motivate a discretization of the wrinkle field as an amplitude and frequency one-form ω , both of which are slow variables and can thus be discretized as piecewise-linear functions on \mathcal{T} without aliasing. The main

challenge of this frequency representation is that phase must be integrated from frequency, by solving a Poisson-like problem of the form $\min_{\phi} \|\mathrm{d}\phi - \omega\|^2$. Sophisticated algorithms for recovering phase from frequency exist which automatically place appropriate singularities [1].

Unfortunately, the frequency representation possesses several downsides that make it unsuitable for simulating dynamics. Most frequency one-forms are not locally or globally integrable, and the closest-possible phase given by the above Poisson problem is not a continuous function of ω . Of course, the recovered phase is only determined up to a global constant phase shift; but even removing this symmetry (by prescribing the phase at a single point, for instance) does not help. Consider, as a thought experiment, a cylinder parameterized on the square $[0, 1] \times [0, 1]$ with the left and right boundaries identified; and consider the time-varying frequency $\omega(x, y; t) = (t, 0)$. The corresponding wrinkle field must have an integer number of waves, with that integer jumping discontinuously at half-integer values of t .

2. COMPLEX WRINKLE FIELDS

In recent work [2], we proposed a complex wrinkle field (CWF) representation of the wrinkle pattern as:

- a complex number z_i per vertex i of \mathcal{T} , which represents the wrinkle amplitude $|z_i|$ and phase $\arg z_i$ at the vertex;
- a real frequency ω_{ij} per oriented edge ij of the mesh (with $\omega_{ij} = -\omega_{ji}$), representing the frequency of the wrinkles in the edge direction.

We do not impose a hard compatibility constraint between the z and ω ; these are *independent* variables.

For a point P inside a triangle with vertices 1, 2, 3, the following interpolant gives the amplitude and phase at P :

$$z(P) = \sum_{j=1}^3 \alpha_j z_j \exp \left(i \left[\alpha_{j+1} \omega_{j(j+1)} + \alpha_{j+2} \omega_{j(j+2)} \right] \right)$$

where $\alpha_1, \alpha_2, \alpha_3$ are the barycentric coordinates of P . Note that this interpolant is only C^0 at the edges of \mathcal{T} ; though in the paper we describe a subdivision scheme for CWFs that empirically yields smooth wrinkle fields.

3. SIMULATING DYNAMICS WITH COMPLEX WRINKLE FIELDS

In theory, simulating dynamics with CWFs is simply a matter of applying Lagrangian mechanics to a formulation of kinetic and potential energy expressed in the CWF variables. Some early experiments in this direction yielded partial successes on simple test cases such as a mylar cylinder twisting over time. However, several unsolved challenges remain:

- since the curvature of the base surface \mathbf{r}_b is a slow variable, we would ideally like \mathbf{r}_b to *remain* wrinkle-free during dynamics, with the wrinkle field variables absorbing normal displacement of the surface. We do not

see this behavior in practice: instead, over time \mathbf{r}_b displaces to yield an aliased, noisy approximation of the wrinkled surface. Perhaps there is a better discrete surface representation (such as by a tensor-product spline or subdivision surface?) that forces the base mesh to remain wrinkle-free;

- all of the above problem setup assumed that locally the surface exhibits only one wrinkle frequency. In practice, *wrinkles-on-wrinkles* do occur, especially on boundary layers near clamped boundaries (such as around the “equator” of an inflated mylar balloon, near where the two disks of mylar have been sewn together);
- a gauge symmetry remains in the CWF parameterization of wrinkle fields: mapping $\omega \mapsto -\omega$ and $\phi \mapsto 2\pi - \phi$ yields a second CWF that represents the same displaced surface \mathbf{r}_m . This remaining symmetry is an obstacle to effective simulation: we observe in experiments that when wrinkles form on a surface, several wrinkled regions nucleate on the surface with opposite frequency orientation, separated by wrinkle-free regions where $a = 0$. As the wrinkled regions grow and eventually meet, they become *glassy* domains of locally-consistent frequency orientation, separated by a thin stripe of zero amplitude that does not go away due to the frequency discontinuity across the stripe. A reduced representation of CWFs without the gauge symmetry (where direction fields are used instead of a one-form for frequency?) is needed.

REFERENCES

- [1] F. Knöppel, K. Crane, U. Pinkall, and P. Schröder, *Stripe Patterns on Surfaces*, ACM Transactions on Graphics (TOG) **32**(4) (2015), 1–11.
- [2] Z. Chen, D. Kaufman, M. Skouras, and E. Vouga, *Complex Wrinkle Field Evolution*, ACM Transactions on Graphics (TOG) **42**(4) (2023), 1–19.

(In-)consistency of the cotan Laplacian and convergence of the heat method

SIMON SCHWARZ

(joint work with Anja Sturm and Max Wardetzky)

The *heat method* introduced in [3] is a popular algorithm in graphics for computing the geodesic distance on polyhedral surfaces. Originally motivated by Varadhan’s asymptotic

$$(1) \quad -2t \log p_t(x, y) \xrightarrow{t \rightarrow 0} d_g^2(x, y)$$

for the heat kernel p_t on a Riemannian manifold (M, g) with geodesic distance d_g (see [8]), the heat method consists of three steps: In a first step, the heat kernel is approximated by solving a single backward Euler step in time:

$$(2) \quad u_h - t \Delta_h u_h = \delta_\gamma,$$

where $t > 0$ is some small time, γ a subset of the domain and Δ_h some discrete Laplace operator on a given triangulation T_h . Instead of using the function u_h

directly for computing the geodesic distance as in Equation (1), numerical experiments show that it is more advantageous to first consider its normalized gradients

$$X_h = \frac{\nabla u_h}{|\nabla u_h|}$$

and recover an approximation d_h of the geodesic distance $d_g(\gamma, \cdot) := \min_{y \in \gamma} d_g(y, \cdot)$ by solving the Poisson equation

$$\Delta_h d_h = \nabla \cdot X_h \quad \text{with} \quad d_h|_\gamma = 0.$$

This three-step approach can also be applied to other related problems, e.g., efficient computation of parallel transport, see [7], and the signed distance function, see [4]. Although the heat method is – along with fast marching, see in particular [6] – the standard tool for distance computation in graphics, a proof of convergence under mesh refinement has been lacking so far. We can prove the following:

Theorem 1. *Let $(T_h)_{h>0}$ be a family of shape regular Delaunay triangulations of a two-dimensional, compact manifold $M \subset \mathbb{R}^3$. Suppose that the maximum edge length of T_h is bounded from above by h and let Δ_h be the cotan Laplace operator. Consider the solution u_h of the linear system in Equation (2) with $t = h^{2\beta}$ for any $0 < \beta < 1$. Then,*

$$-h^\beta \log u_h(x) \xrightarrow{h \rightarrow 0} d_g(\gamma, x) := \min_{y \in \gamma} d_g(y, x)$$

uniformly in x , where d_g denotes the geodesic distance on M .

The main problem in the convergence proof is that the cotan Laplacian is not consistent, i.e., that for a function $f : M \rightarrow \mathbb{R}$

$$\lim_{h \rightarrow 0} \|\Delta_h f - \Delta_g f\|_\infty \neq 0$$

in general. We show, however, that a local *averaging property* holds for the cotan Laplacian on any shape regular Delaunay triangulation:

Theorem 2. *Let $f : M \rightarrow \mathbb{R}$ be a smooth function and $(T_h)_{h>0}$ a family of shape regular Delaunay triangulations of (M, g) with edge lengths bounded from above by h . For any $r_h > 0$ such that $r_h \rightarrow 0$ but $r_h/h \rightarrow \infty$ as $h \rightarrow 0$, we have*

$$(3) \quad \sup_{u \in V_h} \left| \Delta_g f(u) - \frac{1}{|B_{r_h}(u)|} \sum_{v \in B_{r_h}(u) \cap V_h} A_v \Delta_h f(v) \right| \xrightarrow{h \rightarrow 0} 0,$$

where A_v is the dual area of a vertex $v \in V_h$, $B_r(u)$ the r -ball on M with respect to the geodesic distance d_g , and

$$|B_r(u)| := \sum_{v \in B_r(u) \cap V_h} A_v$$

for any $u \in V_h$ and $r > 0$.

This averaging property is closely related to the construction of non-local, but consistent discrete Laplace operators on triangulations and point clouds in [1, 5].

Using this averaging property and the homogenization methods from [2], we can prove Theorem 1. We expect that the combination of (3) and homogenization methods can also be applied to the convergence analysis of other algorithms in geometry processing relying on the cotan Laplacian.

REFERENCES

- [1] M. Belkin, J. Sun, and Y. Wang. *Constructing Laplace operator from point clouds in \mathbb{R}^d* , Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (2009), 1031–1040.
- [2] F. Chaplais. *Averaging and Deterministic Optimal Control*, SIAM J. Control Optim. **25.3** (1987), 767–780.
- [3] 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), 1–11.
- [4] N. Feng and K. Crane. *A Heat Method for Generalized Signed Distance*, ACM Trans. Graph. **43.4** (2024), 1–19.
- [5] K. Hildebrandt and K. Polthier. *On approximation of the Laplace–Beltrami operator and the Willmore energy of surfaces*, Comput. Graph. Forum **30.5** (2011), 1513–1520.
- [6] R. Kimmel and J. A. Sethian. *Computing geodesic paths on manifolds*, Proc. Natl. Acad. Sci. USA **95.15** (1998), 8431–8435.
- [7] N. Sharp, Y. Soliman, and K. Crane. *The Vector Heat Method*, ACM Trans. Graph. **38.3** (2019), 1–19.
- [8] S. R. S. Varadhan. *On the behavior of the fundamental solution of the heat equation with variable coefficients*, Comm. Pure Appl. Math. **20.2** (1967), 431–455.

Neural Geometry Processing via Spherical Neural Surfaces

ROMY WILLIAMSON

(joint work with Niloy Mitra)

Neural surfaces (e.g., neural map encoding, deep implicit and neural radiance fields) have recently gained popularity because of their generic structure (e.g., multi-layer perceptron) and easy integration with modern learning-based setups. Traditionally, we have a rich toolbox of geometry processing algorithms designed for polygonal meshes to analyze and operate on surface geometry. In the absence of an analogous toolbox, neural representations are typically discretized and converted into a mesh, before applying any geometry processing algorithm. This is unsatisfactory and, as we demonstrate, unnecessary. In this work, we propose a *spherical neural surface* representation for genus-0 surfaces and demonstrate how to compute core geometric operators directly on this representation. Namely, we estimate surface normals and first and second fundamental forms of the surface, as well as compute surface gradient, surface divergence and Laplace Beltrami operator on scalar/vector fields defined on the surface. Our representation is fully seamless, overcoming a key limitation of similar explicit representations such as

Neural Surface Maps [1]. These operators, in turn, enable geometry processing directly on the neural representations *without* any unnecessary meshing. We demonstrate illustrative applications in (neural) spectral analysis, heat flow and mean curvature flow, and evaluate robustness to isometric shape variations. We propose theoretical formulations and validate their numerical estimates, against analytical estimates, mesh-based baselines, and neural alternatives, where available. By systematically linking neural surface representations with classical geometry processing algorithms, we believe that this work can become a key ingredient in enabling *neural geometry processing*. Code will be released on the project webpage. If interested, please read the full paper: Neural Geometry Processing via Spherical Neural Surfaces [2].

REFERENCES

- [1] L. Morreale et al., *Neural Surface Maps*, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2021), 4639–4648, <https://doi.org/10.48550/arXiv.2103.16942>
- [2] R. Williamson and N. Mitra, *Neural Geometry Processing via Spherical Neural Surfaces*, (2025), <https://doi.org/10.48550/arXiv.2407.07755>

A geometric optimal transport framework for 3d shape interpolation

ROBERT BEINERT

(joint work with Florian Beier, Gabriele Steidl)

The Gromov-Wasserstein (GW) framework [9, 12] is a non-linear generalization of the classical, linear optimal transport problem, which seeks a relaxed correspondence between two gauged measure spaces while preserving their internal geometry. From a mathematical perspective, a gauged measure space (gm-space) is a triple $\mathbb{X} := (X, g, \xi)$ consisting of a Polish space X , a Borel probability measure $\xi \in \mathcal{P}(X)$, and a gauge function $g \in L_{\text{sym}}(X \times X, \xi \otimes \xi)$, i.e. a symmetric, square-integrable function with respect to the product measure $\xi \otimes \xi$. For two gm-spaces $\mathbb{X}_i := (X_i, g_i, \xi_i)$, the quadratic GW transport [9] itself is based on the so-called GW distance:

$$\text{GW}(\mathbb{X}_1, \mathbb{X}_2) := \inf_{\pi \in \Pi(\mathbb{X}_1, \mathbb{X}_2)} \left(\iint_{(X_1 \times X_2)^2} (g_1(x_1, x'_1) - g_2(x_2, x'_2))^2 d\pi(x_1, x_2) d\pi(x'_1, x'_2) \right)^{\frac{1}{2}}.$$

Here, $\Pi(\mathbb{X}_1, \mathbb{X}_2)$ denotes the set of transport plans $\pi \in \mathcal{P}(X_1 \times X_2)$ with marginals $(P_{X_i})_{\#} \gamma = \pi \circ P_{X_i}^{-1} = \xi_i$, where P_{X_i} is the projection onto the component X_i . Since the measures can live on different Polish spaces, the GW distance serves as a valuable tool for the analysis of objects that do not possess a natural Euclidean embedding or should be studied independently of it. Crucial applications can be found in shape and coherent structure matching [1, 6, 9], classification [2, 4], and interpolation [3, 5, 7, 11].

To tackle the latter, we consider GW barycenters, which are generalized Fréchet means with respect to the GW distance. Since GW is not a metric, we combine




FIGURE 1. One iteration of the Fréchet mean method to approximate the GW barycenter: 1. Lifting of the input spaces \mathfrak{X}_1 and \mathfrak{X}_2 to the tangent space $\mathfrak{T}_{\mathfrak{Y}}$ at \mathfrak{Y} via the logarithmic map ($\mathbb{L}\text{og}_{\mathfrak{Y}}$). 2. Computation of the barycenter in the tangent space. 3. Projection of the tangential barycenter back to \mathfrak{M} via the exponential map ($\mathbb{E}\text{xp}_{\mathfrak{Y}}$).

the indefinite gm-spaces into equivalence classes $\mathfrak{X} := \llbracket \mathbb{X} \rrbracket := \{\mathbb{Y} \mid \text{GW}(\mathbb{X}, \mathbb{Y}) = 0\}$. Equipping the space \mathfrak{M} of these equivalence classes with the GW distance, we obtain a complete, geodesic metric space [12]. For gm-classes \mathfrak{X}_i , $i = 1, \dots, N$, and for weights $\rho_i \geq 0$ with $\sum_{i=1}^N \rho_i = 1$, any GW barycenter is a solution of

$$(1) \quad \arg \min_{\mathfrak{Y} \in \mathfrak{M}} \sum_{i=1}^N \rho_i \text{GW}^2(\mathfrak{X}_i, \mathfrak{Y}).$$

Notice that there are no restriction on the gm-space \mathfrak{Y} ; so we have to determine the underlying Polish space, the Borel probability measure, and the gauge function. Problems of this kind are also called free barycenter problems. To show the existence and to characterize the solutions of (1), we rely on the multi-marginal transport problem

$$(2) \quad \arg \min_{\pi \in \Pi(\mathbb{X}_1, \dots, \mathbb{X}_N)} \frac{1}{2} \iint_{X_{\times}^2} \sum_{i,j=1}^N \rho_i \rho_j |g_i(x_i, x'_i) - g_j(x_j, x'_j)|^2 d\pi(x_{\times}) d\pi(x'_{\times})$$

with the Cartesian product $X_{\times} := \times_{i=1}^N X_i$ as well as $x_{\times} := (x_1, \dots, x_N) \in X_{\times}$ and the multi-marginal plans $\Pi(\mathbb{X}_1, \dots, \mathbb{X}_N) := \{\pi \in \mathcal{P}(X_{\times}) \mid (P_{X_i})_{\#} \pi = \xi_i\}$. The barycenter problem (1) and the multi-marginal problem (2) are equivalent, and any barycenter has the form

$$\left\{ \left[\left(X_{\times}, \sum_{i=1}^N \rho_i, \hat{\pi} \right) \right] \mid \hat{\pi} \text{ solves (2), } \mathbb{X}_i = (X_i, g_i, \xi_i) \in \mathfrak{X}_i \right\} \neq \emptyset,$$

see [3, 5]. On the basis of the GW framework for Gaussians [8], the GW barycenter can here be computed analytically [3]. In general, solving (2) is, however, numerical intractable for $N > 2$.



FIGURE 2. Interpolation between the left and right shape using the proposed relaxed fixpoint iteration. Note that any GW barycenter is only given up to isometries. In this example, the intermediate shapes are mirrored in comparison with the input shapes.

Motivated by obtaining a tractable algorithm, we study the geometry and tangential structure [12] of the induced GW space \mathfrak{GM} . To approximate GW barycenters, we lift a known fixpoint iteration for the computation of Fréchet means in Riemannian manifolds [10] to the GW setting. The basic idea behind the procedure is outlined in Figure 1. The core step is hereby the computation of the tangential barycenter, whose existence is guaranteed, and which can be partially characterized, see [3]. Relaxing the calculation of the tangential barycenter, we propose the following fixpoint iteration: $\mathfrak{Y}^{(0)} \in \mathfrak{GM}$ and $\mathfrak{Y}^{(n+1)} \in \text{TB}_\rho(\mathfrak{Y}^{(n)})$, where

$$\text{TB}_\rho(\mathfrak{Y}) := \left\{ \left[\left(X_\times, \sum_{i=1}^N \rho_i g_i, (P_{X_\times})_\# \gamma \right) \right] \mid \gamma \in \mathcal{P}(Y \times X_\times), (P_{Y \times X_i})_\# \gamma \in \Pi_o(\mathbb{Y}, \mathbb{X}_i) \right\}$$

and where Π_o denotes the set of optimal GW plans. This iteration is simple to implement in practice, since it merely requires the gluing of bi-marginal GW plans, and monotonically improves the quality of the barycenter [3]:

$$\sum_{i=1}^N \rho_i \text{GW}^2(\mathfrak{x}_i, \mathfrak{Y}^{(n)}) \geq \sum_{i=1}^N \rho_i \text{GW}^2(\mathfrak{x}_i, \mathfrak{Y}^{(n+1)}).$$

Moreover, every subsequence of the computed iteration contains a further subsequence converging to a fixpoint, i.e. $\mathfrak{Y} \in \text{TB}(\mathfrak{Y})$. Finally, we provide numerical evidence of the potential of this method, including multi-3d-shape interpolations as outlined in Figure 2.

REFERENCES

- [1] F. Beier, *Gromov–Wasserstein transfer operators*, International Conference on Scale Space and Variational Methods in Computer Vision (2023), 614–626.
- [2] F. Beier, R. Beinert, *Multi-marginal approximation of the linear Gromov–Wasserstein distance*, Proceedings in Applied Mathematics and Mechanics **22** (2022), e202200072.
- [3] F. Beiert, R. Beinert, *Tangential fixpoint iterations for Gromov–Wasserstein barycenters*, arXiv (2024), arXiv:2403.08612.
- [4] F. Beier, R. Beinert, G. Steidl, *On a linear Gromov–Wasserstein distance*, IEEE Transactions on Image Processing **31** (2022), 7292–7305.

- [5] F. Beier, R. Beinert, G. Steidl, *Multi-marginal Gromov–Wasserstein transport and barycenters*, Information and Inference **12** (2023), 2753–2781.
- [6] F. Beier, M. Piening, R. Beinert, G. Steidl, *Joint metric space embedding by unbalanced OT with Gromov–Wasserstein marginal penalization*, arXiv (2025), arXiv:2502.07510.
- [7] S. Chowdhury, Samir and T. Needham, *Gromov–Wasserstein averaging in a Riemannian framework*, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (2020), 842–843.
- [8] J. Delon, A. Desolneux, A. Salmona, *Gromov–Wasserstein distances between Gaussian distributions*, Journal of Applied Probability **59**(4) (2022), 1178–1198.
- [9] F. Mémoli, *Gromov–Wasserstein distances and the metric approach to object matching*, Foundations of Computational Mathematics **11**(4) (2011), 417–487.
- [10] X. Pennec, *Intrinsic statistics on Riemannian manifolds: basic tools for geometric measurements*, Journal of Mathematical Imaging and Vision **25**(1) (2006), 127–154.
- [11] G. Peyré, M. Cuturi, J. Solomon, *Gromov–Wasserstein averaging of kernel and distance matrices*, International Conference on Machine Learning (2016), 2664–2672.
- [12] K.-T. Sturm, *The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces*, Memoirs of the American Mathematical Society **290** (2023).

DEALing with Image Reconstruction: Deep Attentive Least Squares

SEBASTIAN NEUMAYER

(joint work with Mehrsa Pourya, Erich Kobler, Michael Unser)

State-of-the-art image reconstruction methods frequently depend on deep learning models with highly complex architectures and extensive parameterization. While these methods achieve impressive results, they often lack interpretability, robustness, and theoretical convergence guarantees. In this talk, we discuss an alternative: a data-driven reconstruction approach inspired by classical Tikhonov regularization. Our method iteratively refines intermediate reconstructions by solving a sequence of quadratic problems, offering a principled framework that bridges traditional regularization techniques with deep learning.

Each iteration of the proposed method consists of two key components: (i) learned filters that extract salient image features and (ii) an attention mechanism that locally adjusts the penalty applied to these features. This allows the regularization process to adapt dynamically to the structure of the image, leading to improved reconstructions with enhanced interpretability and stability. During the process, each quadratic update can be efficiently computed using the conjugate gradient method. To this end, we can use the previous results for a warm start.

On the theory side, we establish the uniqueness of each update, the existence of fixed points, and conditions for convergence of the process. Furthermore, we prove a stability result for the resulting reconstruction operator, demonstrating robustness against variations in the input data. This makes our approach particularly well-suited for ill-posed inverse problems, where classical methods struggle to provide competitive performance and deep learning models often lack stability and insight into the reconstruction process.

We validate our approach through extensive experimental evaluations on various inverse problems, including denoising, deblurring, and inpainting. Our method

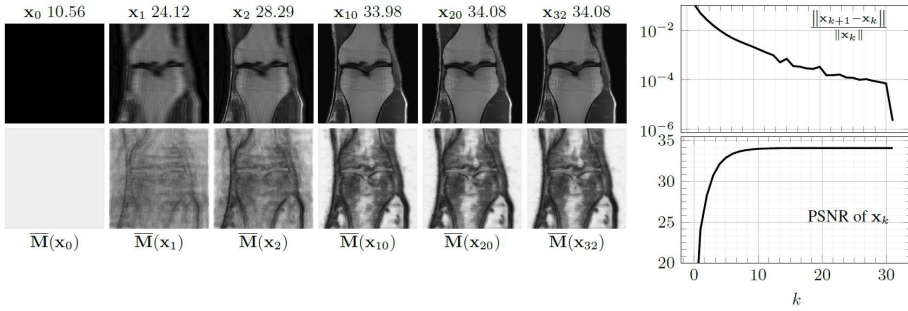


FIGURE 1. Reconstruction path and convergence plots.

achieves performance on par with leading plug-and-play and learned regularizer approaches, while offering a significantly more interpretable and theoretically grounded reconstruction process. Additionally, we visualize the learned attention mechanism to highlight its role in adaptively refining the reconstruction, see Figure 1 for an example path. More visualizations can be found in the associated preprint [1].

Overall, our work presents a novel synthesis of classical inverse problem regularization and modern data-driven techniques. By combining the strengths of both paradigms, we provide a robust, efficient, and interpretable alternative to purely deep learning-based reconstruction methods.

REFERENCES

- [1] M. Pourya, E. Kobler, M. Unser, S. Neumayer, *DEALing with image reconstruction: Deep Attentive Least Squares*, arXiv preprint #2502.04079 (2025)

Graph Neural Networks on Large Random Graphs

SAMUEL VAITER

(joint work with Alberto Bietti, Nicolas Keriven)

We consider graphs where each node represents a latent variable in a compact space X , drawn i.i.d. from a distribution P . Bernoulli edges are formed according to a similarity kernel $W(\cdot, \cdot)$, scaled by a factor α_n that controls the graph sparsity. As n (the number of nodes) grows, such *latent position random graphs* approximate continuous structures on X . We proved the following results [1, 3] on the convergence, stability, and expressivity of Graph Neural Networks (GNNs) on such random graphs. These results heavily rely on chaining arguments and concentration properties of graph shift operators [2, 4].

Convergence of GNNs. A first result is that a Graph Neural Network (GNN) can be regarded as an operator (a “continuous GNN”, cGNN) in the latent space. Under mild conditions, the discrete GNN converges to its continuous counterpart

in the relatively sparse regime ($\alpha_n \sim \log n/n$): as $n \rightarrow \infty$, the outputs of the (discrete) GNN become close to sampling the latent-space function of the cGNN.

Stability of cGNNs. If two random graphs are drawn from different latent distributions or kernels, we proved that the corresponding GNN outputs remain close. This mirrors the idea of spatial stability in convolutional neural networks on Euclidean domains [5], here interpreted on random graph topologies.

Expressivity of GNNs. For node-level tasks, one can think of GNNs as approximating functions over X . A basic set of “base functions” can be expanded by GNN operations (convolution with the graph operator, linear combinations, and non-linear activations). Under sufficient hypothesis on the features, these expanded sets can approximate broader classes of functions over X .

Improving expressivity with Positional Encodings. When nodes lack natural features, *positional encodings* help break undesirable symmetries of the GNNs. For instance, *Eigenvector-based encodings* feed the first few eigenvectors of the adjacency or Laplacian (with sign invariance), approximating the principal spectral modes of the underlying continuous operator. *Distance-based encodings* use powers of the graph operator to summarize node “distances.” Under suitable assumptions, such PEs converge to well-defined latent-space functions, improving GNNs’ expressive range.

REFERENCES

- [1] N. Keriven and S. Vaiter. What functions can Graph Neural Networks compute on random graphs? The role of Positional Encoding, In *NeurIPS*, 2023.
- [2] N. Keriven and S. Vaiter. Sparse and Smooth: improved guarantees for Spectral Clustering in the Dynamic Stochastic Block Model. *Electron. J. Stat.*, 16(1):1330–1366, 2022.
- [3] N. Keriven, A. Bietti, and S. Vaiter. Convergence and Stability of Graph Convolutional Networks on Large Random Graphs. In *NeurIPS*, 2020.
- [4] J. Lei and A. Rinaldo. Consistency of spectral clustering in stochastic block models. *Ann. Stat.*, 43(1):215–237, 2015.
- [5] S. Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.

Beyond Wasserstein: a general fixed-point approach for computing barycentres of measures

JULIE DELON

(joint work with Eloi Tanguy, Nathaël Gozlan)

1. INTRODUCTION

Wasserstein barycentres are a fundamental concept in Optimal Transport theory, providing a way to compute average distributions between multiple probability measures. These barycentres preserve the geometric structure of the underlying distributions, making them particularly well-suited for machine learning applications.

Consider two probability measures μ and ν defined on metric spaces $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$, along with a cost function $c : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$. The optimal transport cost between μ and ν with respect to the ground cost c is given by

$$\mathcal{T}_c(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c d\pi,$$

where $\Pi(\mu, \nu)$ denotes the set of probability measures on $\mathcal{X} \times \mathcal{Y}$ whose marginals are μ and ν . Given K different cost functions c_k and K probability measures ν_k defined on different spaces \mathcal{Y}_k , the barycentre problem in this context can be formulated as

$$(1) \quad \bar{\mu} \in \operatorname{argmin}_{\mu} \sum_{k=1}^K \lambda_k \mathcal{T}_{c_k}(\mu, \nu_k).$$

When the metric space $(\mathcal{X}, d_{\mathcal{X}})$ coincides with $(\mathcal{Y}, d_{\mathcal{Y}})$ and is Polish, and when the cost function is given by $c = d_{\mathcal{X}}^p$ with $p \geq 1$, the quantity $W_p(\mu, \nu) := (\mathcal{T}_{d^p}(\mu, \nu))^{\frac{1}{p}}$ defines a distance between probability measures with finite moments of order p , known as the p -Wasserstein distance. In this case, the barycentre $\bar{\mu}$ defined above is referred to as a Wasserstein barycentre.

The theoretical study of Wasserstein barycentres originates from the seminal work of Carlier and Ekeland [2], who investigated the existence, uniqueness, and dual formulations of barycentre problems with general continuous cost functions. Subsequent research refined and extended these results, particularly for the W_2 distance in Euclidean spaces [1], as well as for W_1 [4] and W_p [3].

2. COMPUTING BARYCENTRES FOR GENERIC TRANSPORT COSTS

From a computational standpoint, computing Wasserstein barycentres is known to be a particularly difficult problem, classified as NP-hard [5]. This complexity arises because exact algorithms have an exponential computational cost, either in terms of the space's dimension or the number of marginal distributions.

A practical approach, introduced in [7], is based on a fixed-point algorithm inspired by the computation of Fréchet means on manifolds. At each iteration, the current barycentre μ is replaced by its push-forward $\left(\sum_{k=1}^K \lambda_k T_k\right) \# \mu$, where the T_k are optimal transport maps between μ and ν_k , provided these maps exist. The authors of [8] were the first to rigorously prove the convergence of this fixed-point method in the case where the measures ν_k are absolutely continuous. Specifically, they demonstrated that a subsequence of iterates converges to a fixed point and that, if this fixed point is unique, it corresponds to a barycentre. Their study focuses on W_2 barycentres. Although their proof is restricted to absolutely continuous measures, this fixed-point approach is widely used in practice for discrete measures and serves as the foundation for free-support methods in numerical optimal transport libraries [6].

In this work, we extend this fixed-point approach to compute barycentres of probability measures for general transport costs, i.e., solutions to the optimisation problem (1). Our assumptions are minimal: we consider compact spaces

and assume that the ground costs c_k are continuous and satisfy the property that $\operatorname{argmin}_x \sum_{k=1}^K \lambda_k c_k(x, x_k)$ is uniquely defined. We do not require the existence of optimal transport maps between μ and the ν_k , nor do we impose specific conditions on the probability measures μ and ν_k . We propose an iterative fixed-point algorithm that generalises the approach of [8] in this more general setting. We establish that the sequences generated by our algorithm have convergent subsequences, that their limits must be fixed points of a certain mapping G , and that any barycentre of (1) is necessarily a fixed point of G . Furthermore, we prove that these results hold even when using entropic regularised transport costs.

From a numerical perspective, our approach makes it possible to extend the recent definition of generalised Wasserstein barycentres introduced in [9], particularly by allowing for nonlinear mappings between the ambient space and the subspaces associated with the measures ν_k . Additionally, our method facilitates the efficient computation of barycentres for the mixture Wasserstein metric [10].

REFERENCES

- [1] Agueh, Martial and Carlier, Guillaume, “Barycenters in the Wasserstein Space,” SIAM Journal on Mathematical Analysis, 2011.
- [2] Carlier, Guillaume and Ekeland, Ivar, “Matching for teams”, Economic theory, 2010.
- [3] Brizzi, Camilla and Friesecke, Gero and Ried, Tobias, “ p -Wasserstein barycenters,” arXiv preprint arXiv:2405.09381, 2024.
- [4] Carlier, Guillaume and Chenchene, Enis and Eichinger, Katharina, “Wasserstein Medians: Robustness, PDE Characterization, and Numerics,” SIAM Journal on Mathematical Analysis, 2024.
- [5] Altschuler, Jason and Boix-Adsera, Enric, “Wasserstein barycenters are NP-hard to compute”, SIAM Journal on Mathematics of Data Science, 2022.
- [6] Flamary, Rémi and Courty, Nicolas and Ehrhardt, Matthias and Nielsen, Frank, “POT: Python Optimal Transport library,” 2021.
- [7] Cuturi, Marco and Doucet, Arnaud, “Fast Computation of Wasserstein Barycenters,” Proceedings of the 31st International Conference on Machine Learning, 2014.
- [8] Álvarez-Esteban, Pedro C and Del Barrio, E and Cuesta-Albertos, JA and Matrán, C, “A fixed-point approach to barycenters in Wasserstein space”, Journal of Mathematical Analysis and Applications, 2016.
- [9] Delon, Julie and Gozlan, Nathaël and Saint-Dizier, Alexandre, “Generalized Wasserstein barycenters between probability measures living on different subspaces”, Annals of Applied Probability, 2022.
- [10] Delon, Julie and Desolneux, Agnès, “A Wasserstein-type distance in the space of Gaussian mixture models,” SIAM Journal on Imaging Sciences, 2020.

Participants**Dr. Robert Beinert**

Sekr. MA 4-3
Technische Universität Berlin
Institut für Mathematik
Straße des 17. Juni 136
10623 Berlin
GERMANY

Prof. Dr. Mirela Ben-Chen

Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 3200003
ISRAEL

Sascha Beutler

Institut für Analysis und Numerik
Fachbereich Mathematik
Universität Münster
Einsteinstraße 62
48149 Münster
GERMANY

Prof. Dr. Alexander Bronstein

Computer Science Department
TECHNION
Israel Institute of Technology
32000 Haifa
ISRAEL

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. Dr. Julie Delon

UFR
Mathématiques-Informatique Université
Paris Cité
45, rue des Saints Pères
75270 Paris Cedex 06
FRANCE

Prof. Dr. Agnes Desolneux

Centre Borelli
ENS Paris-Saclay
4, avenue des sciences
91190 Gif-sur-Yvette Cedex
FRANCE

Nicole Feng

Computer Science Department
Carnegie Mellon University
Pittsburgh PA 15213-3890
UNITED STATES

Dr. Mark Gillespie

Department of Computer Science
École Polytechnique
1 rue Honoré d'Estienne d'Orves
91120 Palaiseau
FRANCE

Florine Hartwig

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

Dr. Johannes Hertrich

University College London
169 Euston Road
NW12AE London
UNITED KINGDOM

Prof. Dr. Michael Kazhdan

Department of Computer Science
The John Hopkins University
229 Malone Hall
Baltimore, MD 21218
UNITED STATES

Prof. Dr. Ron Kimmel

Department of Computer Sciences
TECHNION
Israel Institute of Technology
32000 Haifa
ISRAEL

Prof. Dr. Leif Kobbelt

Lehrstuhl für Informatik VIII
RWTH Aachen
Ahornstraße 55
52074 Aachen
GERMANY

Jun.-Prof. Dr. Zorah Lähner

Institut für Informatik II
Universität Bonn
Friedrich-Hirzebruch Allee 5
53115 Bonn
GERMANY

Leticia Mattos Da Silva

Computer Science and Artificial
Intelligence Laboratory
Massachusetts Institute of Technology
32 Vassar St, Bldg 32-D475A
Cambridge, MA 02139
UNITED STATES

Prof. Dr. Niloy Mitra

Department of Computer Science
University College London
169 Euston Road
London NW1 2AE
UNITED KINGDOM

Prof. Dr. Sebastian Neumayer

Technische Universität Chemnitz
Fakultät für Mathematik
Straße der Nationen 62
09111 Chemnitz
GERMANY

Natalia Pacheco-Tallaj

Department of Mathematics
Massachusetts Institute of
Technology
77 Massachusetts Avenue
Cambridge, MA 02139-4307
UNITED STATES

Prof. Dr. Nicolas Papadakis

Institut de Mathématiques de Bordeaux
351, cours de la Liberation
33405 Talence Cedex
FRANCE

Dr. Xavier Pennec

Inria Sophia-Antipolis, Epione team
Côte d'Azur University and INRIA
2004 Route des Lucioles
P.O. Box 93
06902 Sophia-Antipolis Cedex
FRANCE

Emile Pierret

Département de Mathématiques
U.F.R. Faculté des Sciences
Université d'Orléans
B.P. 6759
45067 Orléans Cedex 02
FRANCE

Dr. Roi Poranne

Dept. of Mathematics & Computer
Sciences
University of Haifa
Mount Carmel
Haifa 31905
ISRAEL

Prof. Dr. Helmut Pottmann

Institut für Diskrete Mathematik
und Geometrie
Technische Universität Wien
Wiedner Hauptstraße 8-10
1040 Wien
AUSTRIA

Dr. Samara Ren

Institute of Science and Technology
Austria
(IST Austria)
Am Campus 1
3400 Klosterneuburg
AUSTRIA

Noam Rotstein

Department of Computer Science
Technion – Israel Institute of
Technology
Haifa 3486114
ISRAEL

Prof. Dr. Martin Rumpf

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

Dr. Josua Sassen

Centre Borelli
ENS Paris-Saclay
4, avenue des sciences
91190 Gif-sur-Yvette
FRANCE

Prof. Dr. Peter Schröder

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

Simon Schwarz

Institut für Mathematische Stochastik
Georg-August-Universität Göttingen
Goldschmidtstraße 7
37077 Göttingen
GERMANY

Dr. Nicholas Sharp

NVIDIA Toronto AI Lab
431 King St W
Toronto ON M5V 3M4
CANADA

Nikolas Uesseler

Mathematisches Institut
Universität Münster
Einsteinstr. 62
48149 Münster
GERMANY

Samuel Vaiter

CNRS, Laboratoire J. A. Dieudonné
Parc Valrose
06000 Nice
FRANCE

Dr. Amir Vaxman

School of Informatics, University of
Edinburgh
10 Crichton St., Newington
EH8 9AB Edinburgh
UNITED KINGDOM

Romy Williamson

Department of Computer Science
University College London
169 Euston Road
NW1 2AE London
UNITED KINGDOM

Prof. Dr. Etienne Vouga

Department of Computer Science
The University of Texas at Austin
2317 Speedway, Stop D 9500
78712 Austin
UNITED STATES

Prof. Dr. Benedikt Wirth

Fachbereich Mathematik und Informatik
Universität Münster
Einsteinstraße 62
48149 Münster
GERMANY

Dr. Ruben Wiersma

Department of Computer Science
ETH Zürich
Universitätstr. 6
8092 Zürich
SWITZERLAND

Prof. Dr. Denis Zorin

Computer Science Department
Courant Institute of Mathematical
Sciences
New York University
60 Fifth Ave
New York, NY 10011
UNITED STATES