REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
21-09-2023	Final Report	15-Jun-2020 - 14-Jun-2023

4. TITLE AND SUBTITLE
Final Report: Geometric Approaches to Near-Optimization

5a. CONTRACT NUMBER
W911NF-20-1-0168

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

6. AUTHORS

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
Massachusetts Institute of Technology (MIT
77 Massachusetts Avenue
NE18-901
Cambridge, MA 02139 -4307

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
76946-NC.40

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Justin Solomon
UU	UU	UU	UU		19b. TELEPHONE NUMBER 617-324-6738

Standard Form 298 (Rev 8/98)
Prescribed by ANSI Std. Z39.18

RPPR Final Report
as of 26-Sep-2023

Agency Code: 21XD

Proposal Number: 76946NC
Agreement Number: W911NF-20-1-0168

INVESTIGATOR(S):

Name: Justin Solomon
Email: jsolomon@mit.edu
Phone Number: 6173246738
Principal: Y

Organization: **Massachusetts Institute of Technology (MIT)**
Address: 77 Massachusetts Avenue, Cambridge, MA 021394307
Country: USA
DUNS Number: 001425594
EIN: 042103594

Report Date: 14-Jun-2023
Date Received: 21-Sep-2023
Final Report for Period Beginning 15-Jun-2020 and Ending 14-Jun-2023
Title: Geometric Approaches to Near-Optimization
Begin Performance Period: 15-Jun-2020
End Performance Period: 14-Jun-2023
Report Term: 0-Other
Submitted By: Justin Solomon
Email: jsolomon@mit.edu
Phone: (617) 324-6738
Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 12
STEM Participants: 9

Major Goals: The high-level goal of this project is to develop geometrically-motivated algorithms for near-optimization, which involves identifying multiple nearly-optimal solutions for variational problems. The motivation here is that objective functions often have shallow regions where many candidate solutions provide reasonable levels of performance. Identifying near-optima helps engineers understand possible trade-offs when selecting a final solution to a problem, can reveal structure in the objective function, and can suggest secondary objective functions to tie-break between nearly indistinguishable points.

The algorithms studied in this project are built from geometric theory, understanding the near-optimal region as a shape embedded in a high-dimensional design space. The shape of this region captures the flexibility to adjust the solution to an optimization problem without affecting the objective value significantly.

Below, we list the goals articulated in our original proposal:

* Near-optimization in multi-objective problems (100% complete)

* Coreset-based near-optimization (100% complete)

* Near-optimization for infinite-dimensional problems (90% complete)

* Exploration of additional direction, including the following candidate topics:

    - Distributionally-robust near-optimization (50% complete)

    - Combinatorial near-optimization (25% complete)

    - Mixed stochastic-deterministic near-optimization (50% complete)

    - Near-optimization on curved spaces (25% complete)

In addition to these research-oriented tasks, a secondary goal is to support the development of the relatively new MIT Geometric Data Processing group led by Prof. Justin Solomon, and to open lines of communication and feedback between the group and various research/engineering organizations supported by the Army.

Supplementary grants tied to this work also were used to launch the Summer Geometry Initiative (SGI), a program designed to improve recruitment and retention of graduate research students in related research. Separate reports provide details about this extremely successful program. Thanks in large part to the ARO's consistent support, SGI has grown from a seedling program to a mainstay in the geometry processing community.

Overall, we are pleased to report that this ARO-funded research project was an extreme success. The effort funded by this grant led to a variety of high-impact publications related to near-optimization and related problems, and the team is well-poised to continue its work in this space over the long term. We have a second ARO grant on closely-related topics that will continue beyond the term of this proposal, and we are seeking support to perform this research over an even longer term.

Accomplishments: We are pleased to report that this ARO-funded research was an overwhelming success. Modulo the usual adjustments in mathematical research as experiments and proof verify or refute the conjectures laid out in our research plan, our final work package reflects successful delivery on all major work items outlined in our proposal.

Before providing detail on the latest (2022-2023) reporting cycle, we can summarize some accomplishments over the full term of the grant, organized by task in the original proposal:

(I) Near-Optimization in Multi-Objective Problems: A collaboration between PhD student Lingxiao Li and MEng student Erica Chiu showed how the Pareto frontier of multi-objective problems can be mapped out using both Markov chain Monte Carlo (MCMC) and variational inference (VI) algorithms. This work formed the core of Ms. Chiu's thesis and also was presented at venues including the Joint Statistical Meetings (JSM) and ACM SIGGRAPH, the latter a collaboration with manufacturing experts who applied our work to problems in planning of complex manufacturing procedures.

(II) Coreset-Based Near-Optimization: Our PhD student Sebastian Claici and postdoctoral associate Aude Genevay collaborated to carry out this work item, which became a key component of Dr. Claici's PhD thesis and led to a well-referenced technical report released to the broader research community. Moreover, this work led our team to launch a broader effort in optimization and machine learning using approximations of the objective functions and input data, built on mathematically well-founded principles developed in our coreset work. For example, Dr. Genevay and intern Gaspard Beugnot used insight from this project to derive coarsening/clustering methods for optimal transport, and we are now applying similar strategies to compress point clouds that are input to 3D computer vision models.

(III) Near-Optimization for Infinite-Dimensional Problems: Work on this item has been spearheaded by PhD students Lingxiao Li and Chris Scarvelis, with collaboration from postdoctoral associates Aude Genevay, Tal Shnitzer, and others. This research task has led to particularly wide-ranging consequences and future work within our group. We started by considering infinite-dimensional optimization problems in optimal transport, leading to a series of several publications tackling the Wasserstein barycenter problem, Wasserstein gradient flows, Bayesian inference, and inverse Riemannian metric learning from an optimal transport model. This work also inspired our team to study related problems in geometric measure theory and sampling, most recently incorporating insight from denoising diffusion probabilistic models.

(IV) Additional Directions: Many of the optional additional directions articulated in our proposal have been fruitful, leading to research products and ongoing collaborations. For instance, we incorporated distributional robustness into our optimal transport-style optimization problems, with a journal paper giving detailed mathematical analysis of the final model. Our collaborations with manufacturing experts also led us to introduce combinatorial search strategies to our Pareto frontier mapping methods.

The text above summarizes accomplishments during the full term of the grant, since this a final report. Specifically in the 2022-2023 term, we have continued the research in the proposal in several directions to wrap up our work on this project. Many of our latest projects have introduced neural networks and related nonlinear models to capture complexity in near-optimization problems. Some concrete examples include the following:

* Given the Hamiltonian of a dynamical physics problem, our SIGGRAPH 2023 publication data-free learning of reduced-order kinematics shows how to trace out the curved space of "near optima" (i.e., in physical language, low energy states) of a physical system. This practically minded technique yields nonlinear changes of coordinates for physical simulation that facilitate data generation and simulation with relatively few degrees of freedom.

* Building on our work in infinite-dimensional near-optimization, we have developed theoretically grounded methods for solving partial differential equations and gradient flows in probability spaces. These algorithms hold promise to tackle both classical inference problems in probability as well as more complex time-varying models that incorporate dynamics.

Training Opportunities: The Geometric Data Processing (GDP) group supports several programs to train future generations of researchers in geometry and related fields at levels spanning from introductory to advanced.

Locally, the group currently houses three postdoctoral research associates. In addition to mentorship from Prof. Solomon, our postdoctoral associates have participated in mentorship and advising programs hosted by the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL), including seminars/workshops on the job search and giving talks. Postdocs (and students) who have received support from this grant have moved onto top faculty and industrial research positions.

Prof. Solomon and his team mentor a large number of undergraduate students through MIT's UROP and SuperUROP programs; these undergraduate research programs support semester- and year-long research, respectively. Some UROP group members were supported through this grant.

Undergraduates hosted in the Geometric Data Processing group are afforded the same consideration as PhD students, including weekly meetings with Prof. Solomon and/or GDP student mentors. In total, the group has hosted over 40 undergraduate researchers since 2016, including 10+ year-long intensive SuperUROPs. These students also receive research advising from the Electrical Engineering and Computer Science (EECS) department and EECS Communication Lab.

Each summer, the group hosts a high school student researcher through the Research Science Institute program. The summer student in 2023 worked with ARO-supported PhD student Chris Scarvelis to refine his model for Riemannian metric learning from optimal transport.

The group also has taken on a number of long-term research visitors, who receive mentorship at MIT and also interact with the students in the group. In spring 2023, the team hosted a student from ETH Zurich who worked on optimizing for structured geometric embeddings of datasets.

Prof. Solomon has also incorporated results from the research supported by this grant into his professional education course titled "Modeling and Optimization for Machine Learning" (MIT Professional Education), co-taught annually until 2022 with Prof. Suvrit Sra, as well as a new "Applied Numerical Algorithms" offered at MIT for the first time fall 2023.

In summer 2021, the MIT Geometric Data Processing Group launched a program titled the Summer Geometry Initiative (SGI), thanks in part to a supplemental grant from the ARO. The program continued in summers 2022, 2023 with similar levels of success. This program was extremely successful in accomplishing its goals of broadening the diversity of the set of students exposed to geometry processing research; several SGI Fellows already have gone on to pursue graduate education in our discipline. We will provide additional details on the ARO report for this supplemental grant.

Finally, Prof. Solomon and staff regularly attend and contribute to a broad variety of research conferences and gatherings, including the International Conference on Machine Learning (ICML), NeurIPS, the Symposium on Geometry processing, ACM SIGGRAPH, and several events supported by the Society for Industrial and Applied Mathematics (SIAM). In the coming year, Prof. Solomon will be the local chair of the Symposium on Geometry Processing (SGP), a top conference in geometry processing.

RPPR Final Report
as of 26-Sep-2023

Results Dissemination: Research supported by this ARO grant has been disseminated to the technical and general-interest communities through a variety of avenues. ARO support is acknowledged in all printed research reports and related products.

Of course, the major means of dissemination in our discipline is through publication. Readers are invited to see the official list of research products accompanying this report for a full list.

A unique aspect of the research by Prof. Solomon and students/staff in the MIT Geometric Data Processing Group is that it appears in an interdisciplinary set of venues. Publications reported in this project appeared in several top-tier venues including:

* ACM Transactions on Graphics (computer graphics)
* ACM SIGGRAPH (computer graphics)
* International Conference on Learning Representations (machine learning)
* Computer Vision and Pattern Recognition (computer vision)

Beyond publications, Prof. Solomon and staff have given invited talks at a variety of institutions and venues. Some recent invited talks incorporating materials resulting from our ARO-funded research include: Geometry and Topology Meet Data Analysis and Machine Learning (GTDAML), the Paris Workshop on Interpolation of Measures, the UC Santa Barbara ECE Distinguished Lecture, the Simon Fraser University Applied and Computational Mathematics Seminar, and the Duke Applied Mathematics and Analysis Seminar.

The group disseminates its work to the broader community through several means. They founded and host the annual New England Symposium on Graphics, open to the broad area computer graphics research community; this program was on hold during the pandemic but is slated to continue soon. The Summer Geometry Initiative (SGI) annually brings summer research experiences in related research to undergraduate and MS students

Material from the ongoing research also is incorporated into Prof. Solomon's courses ("Shape Analysis" and "Applied Numerical Algorithms"); materials for these courses are available for free online, including lecture videos, exercises, and course notes under development.

Honors and Awards: Research and staff on our team have received several honors and awards during the reporting period. These include the following:

* MathWorks Fellowship (PhD student Leticia Mattos Da Silva)

* Best Paper Award, ICML 2023 Workshop on Neural Compression (postdoc Kimia Nadjahi)

Prof. Solomon has received broad recognition of his work, including the following:

* Harold E. Edgerton Faculty Achievement Award

* Tenure, MIT Department of Electrical Engineering and Computer Science

* ACM SIGGRAPH Technical Papers Test-of-Time Award

Please refer to past reports for other awards won by our team during the period of ARO support.

ARO support also has been instrumental in the development and growth of the MIT Geometric Data Processing Group. During the period of support, the team has graduated seven PhD students: Mazdak Abulnaga, Sebastian Claici, David Palmer, Dmitriy Smirnov, Yu Wang, Yue Wang, and Paul Zhang.

Protocol Activity Status:

Technology Transfer: Nothing to Report

PARTICIPANTS:

Participant Type: PD/PI
Participant: Justin Solomon
Person Months Worked: 6.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Postdoctoral (scholar, fellow or other postdoctoral position)
Participant: Tal Shnitzer Dery
Person Months Worked: 3.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Dmitriy Smirnov
Person Months Worked: 5.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Lingxiao Li
Person Months Worked: 8.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Yu Wang
Person Months Worked: 4.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Erica Chiu
Person Months Worked: 5.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Christopher Scarvelis
Person Months Worked: 5.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Undergraduate Student
Participant: Timothy Qian

Person Months Worked: 3.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Undergraduate Student
Participant: Hanna Yang
Person Months Worked: 5.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Undergraduate Student
Participant: Liane Xu
Person Months Worked: 2.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Undergraduate Student
Participant: Zoe Marschner
Person Months Worked: 2.00 **Funding Support:**
Project Contribution:
National Academy Member: N

ARTICLES:

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM SIGGRAPH
Publication Identifier Type: DOI Publication Identifier: 10.1145/3478513.3480551
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: Sum-of-Squares Geometry Processing
Authors: Zoë Marschner, Paul Zhang, David Palmer, Justin Solomon
Keywords: SOS relaxation, geometry
Abstract: eometry processing presents a variety of difficult numerical problems, each seeming to require its own tailored solution. This breadth is largely due to the expansive list of geometric primitives, e.g., splines, triangles, and hexahedra, joined with an ever-expanding variety of objectives one might want to achieve with them. With the recent increase in attention toward higher-order surfaces, we can expect a variety of challenges porting existing solutions that work on triangle meshes to work on these more complex geometry types. In this paper, we present a framework for solving many core geometry processing problems on higher-order surfaces. We achieve this goal through sum-of-squares optimization, which transforms nonlinear polynomial optimization problems into sequences of convex problems whose complexity is captured by a single degree parameter.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM SIGGRAPH Asia
Publication Identifier Type: DOI Publication Identifier: 10.1145/3478513.3480568
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: Interactive All-Hex Meshing via Cuboid Decomposition
Authors: Lingxiao Li, Paul Zhang, Dmitriy Smirnov, Mazdak Abulnaga, Justin Solomon
Keywords: Hex meshing, cuboids
Abstract: Standard PolyCube-based hexahedral (hex) meshing methods aim to deform the input domain into an axis-aligned PolyCube volume with integer corners; if this deformation is bijective, then applying the inverse map to the voxelized PolyCube yields a valid hex mesh. A key challenge in these methods is to maintain the bijectivity of the PolyCube deformation, thus reducing the robustness of these algorithms. In this work, we present an interactive pipeline for hex meshing that sidesteps this challenge by using a new representation of PolyCubes as unions of cuboids. We begin by deforming the input tetrahedral mesh into a near-PolyCube domain whose faces are close but not perfectly aligned to the major axis directions. We then build a PolyCube by optimizing the layout of a set of cuboids with user guidance to closely fit the deformed domain. Finally, we construct an inversion-free pullback map from the voxelized PolyCube to the input domain while optimizing for mesh quality metrics.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM Transactions on Graphics
Publication Identifier Type: DOI Publication Identifier: 10.1145/3450626.3459750
Volume: 40 Issue: 4 First Page #: 1
Date Submitted: 8/31/21 12:00AM Date Published: 8/1/21 4:00PM
Publication Location:
Article Title: Pareto gamuts: exploring optimal designs across varying contexts
Authors: Liane Makatura, Minghao Guo, Adriana Schulz, Justin Solomon, Wojciech Matusik
Keywords: Pareto front, multi-objective optimization
Abstract: Constrained by the limitations of learning toolkits engineered for other applications, such as those in image processing, many mesh-based learning algorithms employ data flows that would be atypical from the perspective of conventional geometry processing. As an alternative, we present a technique for learning from meshes built from standard geometry processing modules and operations. We show that low-order eigenvalue/eigenvector computation from operators parameterized using discrete exterior calculus is amenable to efficient approximate backpropagation, yielding spectral per-element or per-mesh features with similar formulas to classical descriptors like the heat/wave kernel signatures. Our model uses few parameters, generalizes to high-resolution meshes, and exhibits performance and time complexity on par with past work.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM Transactions on Graphics
Publication Identifier Type: DOI Publication Identifier: 10.1145/3450626.3459801
Volume: 40 Issue: 4 First Page #: 1
Date Submitted: 8/31/21 12:00AM Date Published: 8/1/21 8:00PM
Publication Location:
Article Title: Fast quasi-harmonic weights for geometric data interpolation
Authors: Yu Wang, Justin Solomon
Keywords: skinning, interpolation, optimization
Abstract: Wasserstein barycenters provide a geometric notion of the weighted average of probability measures based on optimal transport. In this paper, we present a scalable algorithm to compute Wasserstein-2 barycenters given sample access to the input measures, which are not restricted to being discrete. While past approaches rely on entropic or quadratic regularization, we employ input convex neural networks and cycle-consistency regularization to avoid introducing bias. As a result, our approach does not resort to minimax optimization. We provide theoretical analysis on error bounds as well as empirical evidence of the effectiveness of the proposed approach in low-dimensional qualitative scenarios and high-dimensional quantitative experiments.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM Transactions on Graphics
Publication Identifier Type: DOI Publication Identifier: 10.1145/3450626.3459797
Volume: 40 Issue: 4 First Page #: 1
Date Submitted: 8/31/21 12:00AM Date Published: 8/1/21 4:00AM
Publication Location:
Article Title: HodgeNet: learning spectral geometry on triangle meshes
Authors: Dmitriy Smirnov, Justin Solomon
Keywords: machine learning, Hodge Laplacian
Abstract: Constrained by the limitations of learning toolkits engineered for other applications, such as those in image processing, many mesh-based learning algorithms employ data flows that would be atypical from the perspective of conventional geometry processing. As an alternative, we present a technique for learning from meshes built from standard geometry processing modules and operations. We show that low-order eigenvalue/eigenvector computation from operators parameterized using discrete exterior calculus is amenable to efficient approximate backpropagation, yielding spectral per-element or per-mesh features with similar formulas to classical descriptors like the heat/wave kernel signatures. Our model uses few parameters, generalizes to high-resolution meshes, and exhibits performance and time complexity on par with past work.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Uncertainty in Artificial Intelligence
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: Improving Approximate Optimal Transport Distances using Quantization
Authors: Gaspard Beugnot, Aude Genevay, Kristjan Greenewald, Justin Solomon
Keywords: optimal transport, clustering
Abstract: Optimal transport (OT) is a popular tool in machine learning to compare probability measures geometrically, but it comes with substantial computational burden. Linear programming algorithms for computing OT distances scale cubically in the size of the input, making OT impractical in the large-sample regime. We introduce a practical algorithm, which relies on a quantization step, to estimate OT distances between measures given cheap sample access. We also provide a variant of our algorithm to improve the performance of approximate solvers, focusing on those for entropy regularized transport. We give theoretical guarantees on the benefits of this quantization step and display experiments showing that it behaves well in practice, providing a practical approximation algorithm that can be used as a drop-in replacement for existing OT estimators.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: NeurIPS
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: Large-Scale Wasserstein Gradient Flows
Authors: Petr Mokrov, Alexander Korotin, Lingxiao Li, Aude Genevay, Justin Solomon, Evgeny Burnaev
Keywords: Optimal transport, gradient flow
Abstract: Wasserstein gradient flows provide a powerful means of understanding and solving many diffusion equations. Specifically, Fokker-Planck equations, which model the diffusion of probability measures, can be understood as gradient descent over entropy functionals in Wasserstein space. This equivalence, introduced by Jordan, Kinderlehrer and Otto, inspired the so-called JKO scheme to approximate these diffusion processes via an implicit discretization of the gradient flow in Wasserstein space. Solving the optimization problem associated to each JKO step, however, presents serious computational challenges. We introduce a scalable method to approximate Wasserstein gradient flows, targeted to machine learning applications. Our approach relies on input-convex neural networks (ICNNs) to discretize the JKO steps, which can be optimized by stochastic gradient descent. Unlike previous work, our method does not require domain discretization or particle simulation.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: NeurIPS
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: Do Neural Optimal Transport Solvers Work? A Continuous Wasserstein-2 Benchmark
Authors: Alexander Korotin, Lingxiao Li, Aude Genevay, Justin Solomon, Alexander Filippov, Evgeny Burnaev
Keywords: Optimal transport, benchmark, learning
Abstract: Despite the recent popularity of neural network-based solvers for optimal transport (OT), there is no standard quantitative way to evaluate their performance. In this paper, we address this issue for quadratic-cost transport—specifically, computation of the Wasserstein-2 distance, a commonly-used formulation of optimal transport in machine learning. To overcome the challenge of computing ground truth transport maps between continuous measures needed to assess these solvers, we use inputconvex neural networks (ICNN) to construct pairs of measures whose ground truth OT maps can be obtained analytically. This strategy yields pairs of continuous benchmark measures in high-dimensional spaces such as spaces of images. We thoroughly evaluate existing optimal transport solvers using these benchmark measures.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Transactions on Machine Learning Research
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 9/21/23 12:00AM Date Published:
Publication Location:
Article Title: k-Mixup Regularization for Deep Learning via Optimal Transport
Authors: Kristjan Greenewald, Anming Gu, Mikhail Yurochkin, Justin Solomon, Edward Chien
Keywords: Optimal transport, regularization, learning
Abstract: Mixup is a popular regularization technique for training deep neural networks that improves generalization and increases robustness to certain distribution shifts. It perturbs input training data in the direction of other randomly-chosen instances in the training set. To better leverage the structure of the data, we extend mixup in a simple, broadly applicable way to k-mixup, which perturbs k-batches of training points in the direction of other k-batches. The perturbation is done with displacement interpolation, i.e. interpolation under the Wasserstein metric. We demonstrate theoretically and in simulations that k-mixup preserves cluster and manifold structures, and we extend theory studying the efficacy of standard mixup to the k-mixup case. Our empirical results show that training with k-mixup further improves generalization and robustness across several network architectures and benchmark datasets of differing modalities.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Computer Graphics Forum
Publication Identifier Type: DOI Publication Identifier: 10.1111/cgf.14370
Volume: 40 Issue: 5 First Page #: 231
Date Submitted: 8/31/21 12:00AM Date Published: 8/1/21 8:00AM
Publication Location:
Article Title: Frame Field Operators
Authors: D. Palmer, O. Stein, J. Solomon
Keywords: Differential operator, bilaplacian, frame field
Abstract: Differential operators are widely used in geometry processing for problem domains like spectral shape analysis, data interpolation, parametrization and mapping, and meshing. In addition to the ubiquitous cotangent Laplacian, anisotropic second-order operators, as well as higher-order operators such as the Bilaplacian, have been discretized for specialized applications. In this paper, we study a class of operators that generalizes the fourth-order Bilaplacian to support anisotropic behavior. The anisotropy is parametrized by a symmetric frame field, first studied in connection with quadrilateral and hexahedral meshing, which allows for fine-grained control of local directions of variation. We discretize these operators using a mixed finite element scheme, verify convergence of the discretization, study the behavior of the operator under pullback, and present potential applications.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: SIAM Journal on Imaging Sciences
Publication Identifier Type: DOI Publication Identifier: 10.1137/21m1433058
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: A Splitting Scheme for Flip-Free Distortion Energies
Authors: Oded Stein, Jiajin Li, Justin Solomon
Keywords: optimization, parameterization, meshes, geometry
Abstract: We introduce k-variance, a generalization of variance built on the machinery of random bipartite matchings. K-variance measures the expected cost of matching two sets of k samples from a distribution to each other, capturing local rather than global information about a measure as k increases; it is easily approximated stochastically using sampling and linear programming. In addition to defining k-variance and proving its basic properties, we provide in-depth analysis of this quantity in several key cases, including one-dimensional measures, clustered measures, and measures concentrated on low-dimensional subsets of R^n. We conclude with experiments and open problems motivated by this new way to summarize distributional shape.
Distribution Statement: 5-Distribution authorized to DoD Components only
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Submitted
Publication Identifier Type: DOI Publication Identifier: 10.1145/3508499
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: Co-Optimization of Design and Fabrication Plans for Carpentry
Authors: Haisen Zhao, Max Willsey, Amy Zhu, Chandrakana Nandi, Zachary Tatlock, Justin Solomon, Adriana Sc
Keywords: carpentry, multi-objective optimization, discrete
Abstract: Past work on optimizing fabrication plans given a carpentry design can provide Pareto-optimal plans trading off between material waste, fabrication time, precision, and other considerations. However, when developing fabrication plans, experts rarely restrict to a single design, instead considering families of design variations, sometimes adjusting designs to simplify fabrication. Jointly exploring the design and fabrication plan spaces for each design is intractable using current techniques. We present a new approach to jointly optimize design and fabrication plans for carpentered objects.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 0-Other
Journal: ArXiv
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/31/21 12:00AM Date Published:
Publication Location:
Article Title: Wasserstein Measure Coresets
Authors: Sebastian Claici, Aude Genevay, Justin Solomon
Keywords: optimal transport, coreset
Abstract: The proliferation of large data sets and Bayesian inference techniques motivates demand for better data sparsification. Coresets provide a principled way of summarizing a large dataset via a smaller one that is guaranteed to match the performance of the full data set on specific problems. Classical coresets, however, neglect the underlying data distribution, which is often continuous. We address this oversight by introducing Wasserstein measure coresets, an extension of coresets which by definition takes into account generalization. Our formulation of the problem, which essentially consists in minimizing the Wasserstein distance, is solvable via stochastic gradient descent. This yields an algorithm which simply requires sample access to the data distribution and is able to handle large data streams in an online manner. We validate our construction for inference and clustering.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Computer Graphics Forum
Publication Identifier Type: DOI Publication Identifier: 10.1111/cgf.14064
Volume: 39 Issue: 5 First Page #: 1
Date Submitted: 8/31/21 12:00AM Date Published: 8/1/20 4:00AM
Publication Location:
Article Title: Medial Axis Isoperimetric Profiles
Authors: Paul Zhang, Daryl DeFord, Justin Solomon
Keywords: Isoperimetric profile, optimization, medial axis
Abstract: Recently proposed as a stable means of evaluating geometric compactness, the isoperimetric profile of a planar domain measures the minimum perimeter needed to inscribe a shape with prescribed area varying from 0 to the area of the domain. While this profile has proven valuable for evaluating properties of geographic partitions, existing algorithms for its computation rely on aggressive approximations and are still computationally expensive. In this paper, we propose a practical means of approximating the isoperimetric profile and show that for domains satisfying a “thick neck” condition, our approximation is exact. For more general domains, we show that our bound is still exact within a conservative regime and is otherwise an upper bound. Our method is based on a traversal of the medial axis which produces efficient and robust results.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Computer Graphics Forum
Publication Identifier Type: DOI Publication Identifier: 10.1111/cgf.14074
Volume: 39 Issue: 5 First Page #: 133
Date Submitted: 8/31/21 12:00AM Date Published: 8/1/20 4:00AM
Publication Location:
Article Title: Hexahedral Mesh Repair via Sum-of-Squares Relaxation
Authors: Z. Marschner, D. Palmer, P. Zhang, J. Solomon
Keywords: hex mesh, repair, SOS relaxation
Abstract: The validity of trilinear hexahedral (hex) mesh elements is a prerequisite for many applications of hex meshes, such as finite element analysis. A commonly used check for hex mesh validity evaluates mesh quality on the corners of the parameter domain of each hex, an insufficient condition that neglects invalidity elsewhere in the element, but is straightforward to compute. Hex mesh quality optimizations using this validity criterion suffer by being unable to detect invalidities in a hex mesh reliably, let alone fix them. We rectify these challenges by leveraging sum-of-squares relaxations to pinpoint invalidities in a hex mesh efficiently and robustly. Furthermore, we design a hex mesh repair algorithm that can certify validity of the entire hex mesh. We demonstrate our hex mesh repair algorithm on a dataset of meshes that include hexes with both corner and face-interior invalidities.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: NeurIPS
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/31/21 12:00AM Date Published: 8/31/21 6:11PM
Publication Location:
Article Title: Continuous Regularized Wasserstein Barycenters
Authors: Lingxiao Li, Aude Genevay, Mikhail Yurochkin, Justin Solomon
Keywords: optimal transport, barycenter
Abstract: Wasserstein barycenters provide a geometrically meaningful way to aggregate probability distributions, built on the theory of optimal transport. They are difficult to compute in practice, however, leading previous work to restrict their supports to finite sets of points. Leveraging a new dual formulation for the regularized Wasserstein barycenter problem, we introduce a stochastic algorithm that constructs a continuous approximation of the barycenter. We establish strong duality and use the corresponding primal-dual relationship to parametrize the barycenter implicitly using the dual potentials of regularized transport problems. The resulting problem can be solved with stochastic gradient descent, which yields an efficient online algorithm to approximate the barycenter of continuous distributions given sample access. We demonstrate the effectiveness of our approach and compare against previous work on synthetic examples and real-world applications.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Journal of Machine Learning Research
Publication Identifier Type: Other Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/31/21 12:00AM Date Published: 12/2/20 8:13PM
Publication Location:
Article Title: Incorporating Unlabeled Data into Distributionally Robust Learning
Authors: Charlie Frogner, Sebastian Claici, Edward Chien, Justin Solomon
Keywords: distributionally-robust learning, optimal transport
Abstract: We study a robust alternative to empirical risk minimization called distributionally robust learning [DRL], in which one learns to perform against an adversary who can choose the data distribution from a specified set of distributions. We illustrate a problem with current DRL formulations, which rely on an overly broad definition of allowed distributions for the adversary, leading to learned classifiers that are unable to predict with any confidence. We propose a solution that incorporates unlabeled data into the DRL problem to further constrain the adversary. We show that this new formulation is tractable for stochastic gradient-based optimization and yields a computable guarantee on the future performance of the learned classifier, analogous to - but tighter than - guarantees from conventional DRL.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: SIAM Journal on Mathematics of Data Science
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: k-Variance: A Clustered Notion of Variance
Authors: Justin Solomon, Kristjan Greenewald, Haikady Nagaraja
Keywords: optimal transport, variance
Abstract: We introduce k-variance, a generalization of variance built on the machinery of random bipartite matchings. K-variance measures the expected cost of matching two sets of k samples from a distribution to each other, capturing local rather than global information about a measure as k increases; it is easily approximated stochastically using sampling and linear programming. In addition to defining k-variance and proving its basic properties, we provide in-depth analysis of this quantity in several key cases, including one-dimensional measures, clustered measures, and measures concentrated on low-dimensional subsets of R^n. We conclude with experiments and open problems motivated by this new way to summarize distributional shape.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: International Conference on Learning Representations
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/31/21 12:00AM Date Published: 3/13/21 7:17PM
Publication Location:
Article Title: Continuous Wasserstein-2 Barycenter Estimation without Minimax Optimization
Authors: Alexander Korotin, Lingxiao Li, Justin Solomon, Evgeny Burnaev
Keywords: optimal transport, barycenters
Abstract: Wasserstein barycenters provide a geometric notion of the weighted average of probability measures based on optimal transport. In this paper, we present a scalable algorithm to compute Wasserstein-2 barycenters given sample access to the input measures, which are not restricted to being discrete. While past approaches rely on entropic or quadratic regularization, we employ input convex neural networks and cycle-consistency regularization to avoid introducing bias. As a result, our approach does not resort to minimax optimization. We provide theoretical analysis on error bounds as well as empirical evidence of the effectiveness of the proposed approach in low-dimensional qualitative scenarios and high-dimensional quantitative experiments.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 5-Submitted
Journal: Submitted
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published:
Publication Location:
Article Title: Rewiring with Positional Encodings for Graph Neural Networks
Authors: Rickard Bruel-Gabrielsson, Mikhail Yurochkin, Justin Solomon
Keywords: graph, neural network
Abstract: Several recent works use positional encodings to extend the receptive fields of graph neural network (GNN) layers equipped with attention mechanisms. These techniques, however, extend receptive fields to the complete graph, at substantial computational cost and risking a change in the inductive biases of conventional GNNs, or require complex architecture adjustments. As a conservative alternative, we use positional encodings to expand receptive fields to any r-ring. Our method augments the input graph with additional nodes/edges and uses positional encodings as node and/or edge features. Thus, it is compatible with many existing GNN architectures. We also provide examples of positional encodings that are non-invasive, i.e., there is a one-to-one map between the original and the modified graphs.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: International Conference on Learning Representations
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 9/21/23 12:00AM Date Published:
Publication Location:
Article Title: Learning Proximal Operators to Discover Multiple Optima
Authors: Lingxiao Li, Noam Aigerman, Vladimir G. Kim, Jiajin Li, Kristjan Greenewald, Mikhail Yurochkin, Justin S
Keywords: proximal operators, learning
Abstract: Finding multiple solutions of non-convex optimization problems is a ubiquitous yet challenging task. Most past algorithms either apply single-solution optimization methods from multiple random initial guesses or search in the vicinity of found solutions using ad hoc heuristics. We present an end-to-end method to learn the proximal operator of a family of training problems so that multiple local minima can be quickly obtained from initial guesses by iterating the learned operator, emulating the proximal-point algorithm that has fast convergence. The learned proximal operator can be further generalized to recover multiple optima for unseen problems at test time, enabling applications such as object detection.
Distribution Statement: 1-Approved for public release; distribution is unlimited.
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Symposium on Computer Animation
Publication Identifier Type: DOI Publication Identifier: 10.1111/cgf.14621
Volume: Issue: First Page #:
Date Submitted: 9/21/23 12:00AM Date Published:
Publication Location:
Article Title: Wassersplines for Stylized Neural Animation
Authors: Paul Zhang, Dmitriy Smirnov, Justin Solomon
Keywords: optimal transport, splines, animation
Abstract: Much of computer-generated animation is created by manipulating meshes with rigs. While this approach works well for animating articulated objects like animals, it has limited flexibility for animating less structured free-form objects. We introduce Wassersplines, a novel trajectory inference method for animating unstructured densities based on recent advances in continuous normalizing flows and optimal transport. The key idea is to train a neurally-parameterized velocity field that represents the motion between keyframes. Trajectories are then computed by advecting keyframes through the velocity field. We solve an additional Wasserstein barycenter interpolation problem to guarantee strict adherence to keyframes. Our tool can stylize trajectories through a variety of PDE-based regularizers to create different visual effects. We demonstrate our tool on various keyframe interpolation problems to produce temporally-coherent animations without meshing or rigging.
Distribution Statement: 1-Approved for public release; distribution is unlimited.
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: International Conference on Learning Representations
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 9/21/23 12:00AM Date Published:
Publication Location:
Article Title: Riemannian Metric Learning via Optimal Transport
Authors: Christopher Scarvelis, Justin Solomon
Keywords: optimal transport, Riemannian metric
Abstract: We introduce an optimal transport-based model for learning a metric tensor from cross-sectional samples of evolving probability measures on a common Riemannian manifold. We neurally parametrize the metric as a spatially-varying matrix field and efficiently optimize our model's objective using a simple alternating scheme. Using this learned metric, we can non-linearly interpolate between probability measures and compute geodesics on the manifold. We show that metrics learned using our method improve the quality of trajectory inference on scRNA and bird migration data at the cost of little additional cross-sectional data.
Distribution Statement: 1-Approved for public release; distribution is unlimited.
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM Transactions on Graphics
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 9/21/23 12:00AM Date Published:
Publication Location:
Article Title: Symmetric Volume Maps
Authors: S. Mazdak Abulnaga, Oded Stein, Polina Golland, Justin Solomon
Keywords: volume map, tetrahedral mesh
Abstract: Although shape correspondence is a central problem in geometry processing, most methods for this task apply only to two-dimensional surfaces. The neglected task of volumetric correspondence—a natural extension relevant to shapes extracted from simulation, medical imaging, and volume rendering—presents unique challenges that do not appear in the two-dimensional case. In this work, we propose a method for mapping between volumes represented as tetrahedral meshes. Our formulation minimizes a distortion energy designed to extract maps symmetrically, i.e., without dependence on the ordering of the source and target domains. We accompany our method with theoretical discussion describing the consequences of this symmetry assumption, leading us to select a symmetrized ARAP energy that favors isometric correspondences. Our final formulation optimizes for near-isometry while matching the boundary.
Distribution Statement: 1-Approved for public release; distribution is unlimited.
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: International Conference on Machine Learning
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published: 7/1/22 8:00AM
Publication Location:
Article Title: Log-Euclidean Signatures for Intrinsic Distances Between Unaligned Datasets
Authors: Tal Shnitzer, Mikhail Yurochkin, Kristjan Greenewald, Justin Solomon
Keywords: dataset distance, diffusion
Abstract: The need for efficiently comparing and representing datasets with unknown alignment spans various fields, from model analysis and comparison in machine learning to trend discovery in collections of medical datasets. We use manifold learning to compare the intrinsic geometric structures of different datasets by comparing their diffusion operators, symmetric positive-definite (SPD) matrices that relate to approximations of the continuous Laplace-Beltrami operator from discrete samples. Existing methods typically assume known data alignment and compare such operators in a pointwise manner. Instead, we exploit the Riemannian geometry of SPD matrices to compare these operators and define a new theoretically-motivated distance based on a lower bound of the log-Euclidean metric. Our framework facilitates comparison of data manifolds expressed in datasets with different sizes, numbers of features, and measurement modalities.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Computer Vision and Pattern Recognition
Publication Identifier Type: Publication Identifier:
Volume: Issue: First Page #:
Date Submitted: 8/27/22 12:00AM Date Published: 6/1/22 8:00AM
Publication Location:
Article Title: DeepCurrents: Learning Implicit Representations of Shapes with Boundaries
Authors: David Palmer, Dmitriy Smirnov, Stephanie Wang, Albert Chern, Justin Solomon
Keywords: geometric measure theory, deep learning
Abstract: Recent techniques have been successful in reconstructing surfaces as level sets of learned functions (such as signed distance fields) parameterized by deep neural networks. Many of these methods, however, learn only closed surfaces and are unable to reconstruct shapes with boundary curves. We propose a hybrid shape representation that combines explicit boundary curves with implicit learned interiors. Using machinery from geometric measure theory, we parameterize currents using deep networks and use stochastic gradient descent to solve a minimal surface problem. By modifying the metric according to target geometry coming, e.g., from a mesh or point cloud, we can use this approach to represent arbitrary surfaces, learning implicitly defined shapes with explicitly defined boundary curves. We further demonstrate learning families of shapes jointly parameterized by boundary curves and latent codes.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM Transactions on Graphics
Publication Identifier Type: DOI Publication Identifier: 10.1145/3450626.3459797
Volume: 40 Issue: 4 First Page #:
Date Submitted: 9/21/23 12:00AM Date Published: 7/19/21 4:00AM
Publication Location:
Article Title: HodgeNet: learning spectral geometry on triangle meshes
Authors: Dmitriy Smirnov, Justin Solomon
Keywords: machine learning, triangle meshes
Abstract: Constrained by the limitations of learning toolkits engineered for other applications, such as those in image processing, many mesh-based learning algorithms employ data flows that would be atypical from the perspective of conventional geometry processing. As an alternative, we present a technique for learning from meshes built from standard geometry processing modules and operations. We show that low-order eigenvalue/eigenvector computation from operators parameterized using discrete exterior calculus is amenable to efficient approximate backpropagation, yielding spectral per-element or per-mesh features with similar formulas to classical descriptors like the heat/wave kernel signatures. Our model uses few parameters, generalizes to high-resolution meshes, and exhibits performance and time complexity on par with past work.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

RPPR Final Report
as of 26-Sep-2023

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: SIAM Journal on Mathematics of Data Science
Publication Identifier Type: DOI Publication Identifier: 10.1137/20m1385895
Volume: 4 Issue: 3 First Page #:
Date Submitted: 9/21/23 12:00AM Date Published: 7/7/22 4:00AM
Publication Location:
Article Title: k-Variance: A Clustered Notion of Variance
Authors: Justin Solomon, Kristjan Greenewald, Haikady Nagaraja,
Keywords: variance, optimal transport
Abstract: We introduce k-variance, a generalization of variance built on the machinery of random bipartite matchings. k-variance measures the expected cost of matching two sets of samples from a distribution to each other, capturing local rather than global information about a measure as k increases; it is easily approximated stochastically using sampling and linear programming. In addition to defining k-variance and proving its basic properties, we provide in-depth analysis of this quantity in several key cases, including one-dimensional measures, clustered measures, and measures concentrated on low-dimensional subsets of R^n. We conclude with experiments and open problems motivated by this new way to summarize distributional shape.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: IEEE Transactions on Medical Imaging
Publication Identifier Type: DOI Publication Identifier: 10.1109/tmi.2021.3128743
Volume: 41 Issue: 4 First Page #:
Date Submitted: 9/21/23 12:00AM Date Published: 4/1/22 4:00AM
Publication Location:
Article Title: Volumetric Parameterization of the Placenta to a Flattened Template
Authors: S. Mazdak Abulnaga, Esra Abaci Turk, Mikhail Bessmeltsev, P. Ellen Grant, Justin Solomon, Polina Gol
Keywords: geometric mapping, placental imaging
Abstract: We present a volumetric mesh-based algorithm for flattening the placenta to a canonical template to enable effective visualization of local anatomy and function. Monitoring placental function in vivo promises to support pregnancy assessment and to improve care outcomes. We aim to alleviate visualization and interpretation challenges presented by the shape of the placenta when it is attached to the curved uterine wall. To do so, we flatten the volumetric mesh that captures placental shape to resemble the well-studied ex vivo shape. We formulate our method as a map from the in vivo shape to a flattened template that minimizes the symmetric Dirichlet energy to control distortion throughout the volume. Local injectivity is enforced via constrained line search during gradient descent. We evaluate the proposed method on 28 placenta shapes extracted from MRI images in a clinical study of placental function. We achieve sub-voxel accuracy in mapping the boundary of the placenta to the template.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Physical Review E
Publication Identifier Type: DOI Publication Identifier: 10.1103/physreve.104.064130
Volume: 104 Issue: 6 First Page #:
Date Submitted: 9/21/23 12:00AM Date Published: 12/20/21 5:00AM
Publication Location:
Article Title: Empirical sampling of connected graph partitions for redistricting
Authors: Elle Najt, Daryl DeFord, Justin Solomon
Keywords: monte carlo, graph partition
Abstract: The space of connected graph partitions underlies statistical models used as evidence in court cases and reform efforts that analyze political districting plans. In response to the demands of redistricting applications, researchers have developed sampling methods that traverse this space, building on techniques developed for statistical physics. In this paper, we study connections between redistricting and statistical physics, and in particular with self-avoiding walks. We exploit knowledge of phase transitions and asymptotic behavior in self-avoiding walks to analyze two questions of crucial importance for Markov chain Monte Carlo analysis of districting plans. First, we examine mixing times of a popular Glauber dynamics-based Markov chain and show how the self-avoiding walk phase transitions interact with mixing time. We examine factors new to the redistricting context that complicate the picture, notably the population balance requirements, and connectivity requirements.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM Transactions on Graphics
Publication Identifier Type: DOI Publication Identifier: 10.1145/3476576.3476627
Volume: 40 Issue: 4 First Page #:
Date Submitted: 9/21/23 12:00AM Date Published: 8/1/21 4:00AM
Publication Location:
Article Title: Fast quasi-harmonic weights for geometric data interpolation
Authors: Yu Wang, Justin Solomon
Keywords: Computer Graphics and Computer-Aided Design
Abstract: We propose quasi-harmonic weights for interpolating geometric data, which are orders of magnitude faster to compute than state-of-the-art. Currently, interpolation (or, skinning) weights are obtained by solving large-scale constrained optimization problems with explicit constraints to suppress oscillative patterns, yielding smooth weights only after a substantial amount of computation time. As an alternative, our weights are obtained as minima of an unconstrained problem that can be optimized quickly using straightforward numerical techniques. We consider weights that can be obtained as solutions to a parameterized family of second-order elliptic partial differential equations. By leveraging the maximum principle and careful parameterization, we pose weight computation as an inverse problem of recovering optimal anisotropic diffusivity tensors. In addition, we provide a customized ADAM solver that significantly reduces the number of gradient steps.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Computer Graphics Forum
Publication Identifier Type: DOI Publication Identifier: 10.1111/cgf.14621
Volume: 41 Issue: 8 First Page #:
Date Submitted: 9/21/23 12:00AM Date Published: 12/1/22 5:00AM
Publication Location:
Article Title: Wassersplines for Neural Vector Field Controlled Animation
Authors: P. Zhang, D. Smirnov, J. Solomon
Keywords: Computer Graphics and Computer-Aided Design
Abstract: Much of computer-generated animation is created by manipulating meshes with rigs. While this approach works well for animating articulated objects like animals, it has limited flexibility for animating less structured free-form objects. We introduce Wassersplines, a novel trajectory inference method for animating unstructured densities based on recent advances in continuous normalizing flows and optimal transport. The key idea is to train a neurally-parameterized velocity field that represents the motion between keyframes. Trajectories are then computed by advecting keyframes through the velocity field. We solve an additional Wasserstein barycenter interpolation problem to guarantee strict adherence to keyframes. Our tool can stylize trajectories through a variety of PDE-based regularizers to create different visual effects. We demonstrate our tool on various keyframe interpolation problems to produce temporally-coherent animations without meshing or rigging.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: ACM Transactions on Graphics
Publication Identifier Type: DOI Publication Identifier: 10.1145/3592105
Volume: 42 Issue: 4 First Page #:
Date Submitted: 9/21/23 12:00AM Date Published: 7/26/23 8:00PM
Publication Location:
Article Title: Variational quasi-harmonic maps for computing diffeomorphisms
Authors: Yu Wang, Minghao Guo, Justin Solomon
Keywords: injective mapping, optimization
Abstract: Computation of injective (or inversion-free) maps is a key task in geometry processing, physical simulation, and shape optimization. Despite being a longstanding problem, it remains challenging due to its highly nonconvex and combinatoric nature. We propose computation of variational quasi-harmonic maps to obtain smooth inversion-free maps. Our work is built on a key observation about inversion-free maps: A planar map is a diffeomorphism if and only if it is quasi-harmonic and satisfies a special Cauchy boundary condition. We hence equate the inversion-free mapping problem to an optimal control problem derived from our theoretical result, in which we search in the space of parameters that define an elliptic PDE. We show that this problem can be solved by minimizing within a family of functionals. Similarly, our discretized functionals admit exactly injective maps as the minimizers, empirically producing inversion-free discrete maps of triangle meshes.
Distribution Statement: 1-Approved for public release; distribution is unlimited.
Acknowledged Federal Support: **Y**

CONFERENCE PAPERS:

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: SIGGRAPH '23: Special Interest Group on Computer Graphics and Interactive Techniques Conference
Date Received: 21-Sep-2023 Conference Date: 06-Aug-2023 Date Published: 23-Jul-2023
Conference Location: Los Angeles CA USA
Paper Title: Sum-of-Squares Collision Detection for Curved Shapes and Paths
Authors: Paul Zhang, Zoë Marschner, Justin Solomon, Rasmus Tamstorf
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: SIGGRAPH '23: Special Interest Group on Computer Graphics and Interactive Techniques Conference
Date Received: 21-Sep-2023 Conference Date: 06-Aug-2023 Date Published:
Conference Location: Los Angeles, CA, USA
Paper Title: Data-Free Learning of Reduced-Order Kinematics
Authors: Nicholas Sharp, Cristian Romero, Alec Jacobson, Etienne Vouga, Paul Kry, David I.W. Levin, Justin So
Acknowledged Federal Support: **Y**

DISSERTATIONS:

Publication Type: Thesis or Dissertation
Institution: Massachusetts Institute of Technology
Date Received: 21-Sep-2023 Completion Date: 5/1/22 4:00AM
Title: Deep Learning on Geometry Representations
Authors: Dmitriy Smirnov
Acknowledged Federal Support: **Y**

Publication Type: Thesis or Dissertation
Institution: Massachusetts Institute of Technology
Date Received: 21-Sep-2023 Completion Date: 5/1/22 8:00AM
Title: Uniform Sampling over Level Sets
Authors: Erica Chiu
Acknowledged Federal Support: **N**

Publication Type: Thesis or Dissertation
Institution: Massachusetts Institute of Technology
Date Received: 21-Sep-2023 Completion Date: 5/1/20 10:15PM
Title: Structure as simplification: transportation tools for understanding data
Authors: Sebastian Claici
Acknowledged Federal Support: **N**

Publication Type: Thesis or Dissertation
Institution: Massachusetts Institute of Technology
Date Received: 21-Sep-2023 Completion Date: 5/2/23 2:18AM
Title: Geometric Computing Beyond the Laplacian
Authors: Yu Wang
Acknowledged Federal Support: **N**

Partners

,

I certify that the information in the report is complete and accurate:
Signature: Justin Solomon
Signature Date: 9/21/23 6:23PM

No figures, charts, or formulas are required for this report; please reference other sections of the report for a summary of our work.