Frontiers in Applied Analysis

June 3–6, 2025

Hall of Arts (HOA) 160

The workshop brings together researchers at the forefront of a broad spectrum of applied analysis for discussing and exchange of ideas. Funding for junior participants and an opportunity for them to present their work have been available. The workshop is made possible by the NSF Grant DMS 2342349, RTG *Frontiers in Applied Analysis* and the Center of Nonlinear Analysis of CMU.

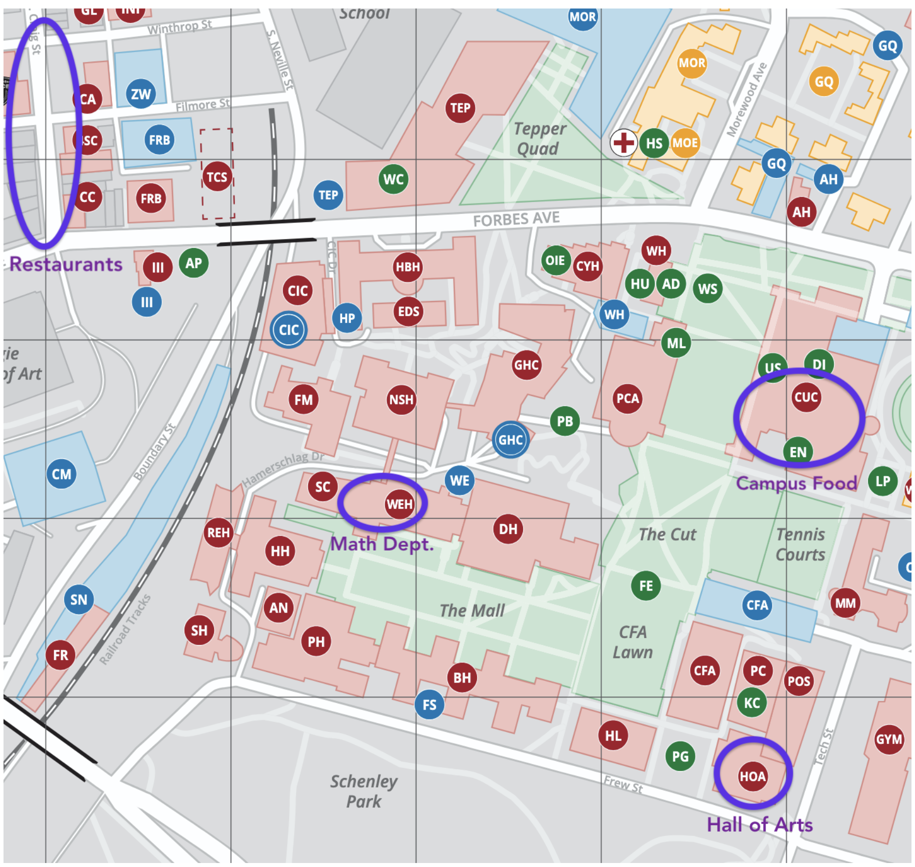

Figure 1: Talks are in Hall of Arts (HOA 160). Main food locations for lunch are also indicated. Math Department is in Wean Hall, where one can also find good espresso.

Tuesday, June 3

Time	Session Details
08:20–09:00	**Breakfast**
09:00–10:00	**John Ball** *Nonlinear elasticity and computer vision*
10:00–11:00	**Gianni Dal Maso** *A matrix-valued measure associated to the derivatives of a generalised function of bounded deformation*
11:00–11:25	**Coffee Break**
11:30–12:00	**Kerrek Stinson** *A linearization result for crack growth*
12:00–12:30	**Gokul Nair** *Scaling law analysis for crystals on curved surfaces*
12:30–2:20	**Lunch**
2:30–3:30	**Gautam Iyer** *Rebalanced Annealing: Efficiently sampling from multimodal distributions*
3:30–4:30	**Lightning Talks**
4:30–6:00	**Poster Reception**

Wednesday, June 4

Time	Session Details
08:20–09:00	**Breakfast**
09:00–10:00	**Lenya Ryzhik** *Diffusion methods and sampling with PDE: an amateur perspective*
10:00–11:00	**Ulisse Stefanelli** *A free boundary problem in accretive growth*
11:00–11:25	**Coffee Break**
11:30–12:00	**Sebastian Munoz** *Free boundary regularity and support propagation in mean field games and optimal transport*
12:00–12:30	**Francesco Nobili** *Isoperimetric planar tilings with unequal cells*
12:30–2:20	**Lunch**
2:30–3:30	**Inwon Kim** *Supercooled Stefan problem and stochastic optimization*
3:30–4:00	**Raghav Venkatraman** *Minimax rates for the estimation of eigenpairs of weighted Laplace-Beltrami operators on manifolds*
4:00–4:30	**Coffee Break**
4:30–5:30	**Charles Smart** *Recent progress in elliptic homogenization*

Thursday, June 5

Time	Session Details
08:20–09:00	**Breakfast**
09:00–10:00	**Carola Schönlieb** *Some topics in structure preserving deep learning*
10:00–11:00	**Nick Boffi** *Stochastic interpolants: from generative modeling to generative science and engineering*
11:00–11:25	**Coffee Break**
11:30–12:00	**Xiaochuan Tian** *Sparse Neural Network Solutions to PDEs: Variational and RKBS Perspectives*
12:00–12:30	**Giovanni Brigati** *Kinetic Optimal Transport*
12:30–2:20	**Lunch**
2:30–3:30	**Martin Rumpf** *Regularizing autoencoder for manifold embedding and Riemannian calculus on latent spaces*
3:30–4:00	**Jan-Eric Sulzback** *Slow manifolds in multiscale PDEs and their application*
4:00–4:30	**Coffee Break**
4:30–5:20	**Jeff Calder** *On the continuum limit of t-SNE visualization*

Friday, June 6

Time	Session Details
08:20–09:00	**Breakfast**
09:00–10:00	**Wilfrid Gangbo** *Hamilton-Jacobi equations in non-commutative variables*
10:00–10:30	**Elias Hess-Childs** *A universal total anomalous dissipator*
11:00–11:30	**Andrew Warren** *Estimation of one-dimensional structures in data*
11:00–11:25	**Coffee Break**
11:30–12:30	**Selim Esedoglu** *Vectorial median filters and curvature motion of networks*

Abstracts

John Ball *Nonlinear elasticity and computer vision*
Abstract: A nonlinear elasticity model for comparing images is formulated and analyzed, in which optimal transformations between images are sought as minimizers of an integral functional. The existence of minimizers in a suitable class of homeomorphisms between image domains is established under natural hypotheses, and the question of whether for affinely related images the minimization algorithm delivers the linear transformation as the unique minimizer is discussed. This is joint work with Chris Horner.

Gianni Dal Maso *A matrix-valued measure associated to the derivatives of a generalised function of bounded deformation*
Abstract: We associate to every function $u \in GBD(\Omega)$ a measure μ_u with values in the space of symmetric matrices, which generalises the distributional symmetric gradient Eu defined for functions of bounded deformation. We show that this measure μ_u admits a decomposition as the sum of three mutually singular matrix-valued measures μ_u^a, μ_u^c, and μ_u^j, the absolutely continuous part, the Cantor part, and the jump part, as in the case of $BD(\Omega)$ functions. We then characterise the space $GSBD(\Omega)$, originally defined only by slicing, as the space of functions $u \in GBD(\Omega)$ such that $\mu_u^c = 0$.

Kerrek Stinson *A linearization result for crack growth*
Abstract: The Griffith theory of fracture is based on the idea that elastic energy can be dissipated by the formation of cracks. Taking the variational viewpoint (a la Francfort and Marigo), we will discuss a recent result proving crack growth in a nonlinearly elastic material converges to the linear analogue as the stiffness of the material tends to infinity. This is the first linearization result for crack growth, which makes no a priori assumptions on the geometry of the evolving crack. Based on joint work with Manuel Friedrich and Pascal Steinke.

Gokul Nair *Scaling law analysis for crystals on curved surfaces*
Abstract: Curved crystals are different from their flat counterparts. Experiments where colloidal particles are put onto liquid droplets have shown that, with enough particles, the resulting crystals are highly anisotropic and fractal-like, even though they form slowly enough that energy minimization should apply. Motivated by the surrounding literature, we study a continuum model of curved crystals involving minimization of a bulk, non-Euclidean elastic energy plus a surface energy (perimeter) term. We prove upper and lower bounds on the scaling law of the minimum energy with respect to its parameters, and find two laws distinguishing anisotropic and isotropic crystal growth. Mathematically, the key ingredient is a new "thin isoperimetric inequality" whose optimizers are not spheres. This interpolation-style inequality yields lower bounds on the energy of a curved crystal, which are optimal in the expected anisotropic regime. This is joint work with Ian Tobasco.

Gautam Iyer *Rebalanced Annealing: Efficiently sampling from multimodal distributions*
Abstract: A computationally challenging problem that arises in many situations is to sample from the Gibbs distribution when the configuration space is large. When the energy function

is convex, there are several algorithms that perform well even in high dimensions. When the energy landscape is complicated, the cost of existing algorithms is exponential in the inverse temperature, and are impractical when the temperature is small. We present an algorithm (Rebalanced Annealing) whose cost is (a low degree) polynomial in the inverse temperature. This is joint work with R. Han and D. Slepčev.

Lenya Ryzhik *Diffusion methods and sampling with PDE: an amateur perspective*
Abstract: The Ermon-Song diffusion algorithm introduced in their papers 2019 and 2021 is a basic approach to producing additional samples of a ”complicated” probability distribution starting with samples from another, known distribution from which sampling is easier. Without claiming any expertise in the generative or training aspects of the model, we will discuss this algorithm in the context of PDE methods of speeding up sampling, as well as the convergence of a discrete version of the Ermon-Song algorithm when the distribution from which one needs to sample has singular support. The ingredients in the proof are all completely elementary but to the best of our knowledge some of them may be new. This is a joint work with Ayya Alieva and Gautam Iyer.

Ulisse Stefanelli *A Free Boundary Problem in Accretive Growth*
Abstract: Growth is a fundamental process in many biological, natural, and technological systems. In the case of accretive growth, material is added at the boundary of the growing system. Mathematically, this process can be modeled by a Hamilton-Jacobi equation coupled with a PDE for an activation field (e.g., nutrient, temperature, or stress) defined on the evolving domain. The resulting coupled system is of free boundary type.

In this talk, I will present an analysis of such a free boundary problem. At first, I will discuss some basic geometry and regularity for the sublevels of the stationary Hamilton-Jacobi equation, which model the growing domain. In particular, I will check that such domains are regular enough to admit a uniform Poincare‘ inequality. This allows us to prove an existence result for the coupled free boundary problem.

The talk is based on collaborations with Elisa Davoli (TU Wien), Katerina Nik (KAUST), and Giuseppe Tomassetti (Rome 3).

Sebastian Munoz *Free boundary regularity and support propagation in mean-field games and optimal transport*
Abstract: In this talk, we present recent results on the regularity of first-order mean-field games systems. We focus on systems where the initial density is a compactly supported function on the real line. Our results show that the solution is smooth in regions where the density is strictly positive and that the density itself is globally continuous. Additionally, the speed of propagation is determined by the behavior of the cost function for small densities. When the coupling is entropic, we demonstrate that the support of the density propagates with infinite speed. On the other hand, when $f(m) = m^\theta$ with $\theta > 0$, we prove that the speed of propagation is finite. In this case, we establish that under a natural non-degeneracy assumption, the free boundary is strictly convex and enjoys $C^{1,1}$ regularity. We also establish sharp estimates on the speed of sup- port propagation and the rate of long time decay for the density. Our methods are based on the analysis of an elliptic equation

satisfied by the flow of optimal trajectories. The results also apply to mean-field planning problems, characterizing the structure of minimizers of a class of optimal transport problems with congestion.

Francesco Nobili *Isoperimetric planar tilings with unequal cells*
Abstract: In this talk, we consider an isoperimetric problem for planar Tilings allowing for unequal repeating cells. We discuss general existence and regularity results and we study classification results for double Tilings, i.e. Tilings with two repeating cells. In this case, we explicitly compute the associated energy profile and give a complete description of the phase transitions. Based on joint works with M. Novaga and E. Paolini.

Inwon Kim *Supercooled Stefan problem and stochastic optimization*
Abstract: The supercooled Stefan problem describes freezing of supercooled water into ice, by diffusion as well as the change of an enthalpy function. It is known to be unstable with finite-time singularities. We will discuss construction of global solutions to this problem via its connection with a stochastic optimization problem, as well as some qualitative properties of the solutions via its connection to the potential theory.

Raghav Venkatraman *Minimax rates for the estimation of eigenpairs of weighted Laplace-Beltrami operators on manifolds*
Abstract: This talk has two parts. In the first, we study the problem of estimating eigenpairs of elliptic differential operators from samples of a distribution ρ supported on a manifold $\mathcal{M}$. The operators discussed in the paper are relevant in unsupervised learning and in particular are obtained by taking suitable scaling limits of widely used graph Laplacians over data clouds. We study the minimax risk for this eigenpair estimation problem and explore the rates of approximation that can be achieved by commonly used graph Laplacians built from random data. More concretely, assuming that ρ belongs to a certain family of distributions with controlled second derivatives, and assuming that the d-dimensional manifold $\mathcal{M}$ where ρ is supported has bounded geometry, we prove that the statistical minimax rate for approximating eigenvalues and eigenvectors in the $H^1(\mathcal{M})$-sense is $n^{-2/(d+4)}$, a rate that matches the minimax rate for a closely related density estimation problem. To the best of our knowledge, our results are the first statistical lower bounds for this type of eigenpair estimation problem.

In the second part, we then revisit the literature studying Laplacians over proximity graphs in the large data limit and prove that eigenpairs of these graph-based operators can induce manifold agnostic estimators with an error of approximation that, up to logarithmic corrections, matches our minimax lower bounds, providing in this way a concrete statistical basis for the claim that graph Laplacian based estimators are, essentially, optimal for this estimation problem.

Time permitting, we indicate how these results draw inspiration from quantitative stochastic homogenization, and compare the rates of convergence in the case of sparse graphs (high contrast) and dense graphs (low contrast), and close with providing a flavor of the proofs of these results.

Joint work with Nicolas Garcia Trillos and Chenghui Li (Wisconsin), and Scott Armstrong (Courant and Sorbonne).

Charles Smart *Recent progress in elliptic homogenization*
Abstract: I will survey some recent results in elliptic homogenization, both in the periodic and random case. This will include some discussion of the recent renormalization theory by Armstrong, Bou-Rabee, and Kuusi.

Carola-Bibiane Schönlieb. *Some topics in structure preserving deep learning*
I will discuss some of our recent works on structure preserving deep learning for the design of neural networks with specific properties - such as non-expansiveness or mass conservation - and their application to imaging and to the solution of partial differential equations.

Nick Boffi. *Stochastic interpolants: from generative modeling to generative science and engineering*
While diffusion-based generative models have achieved state-of-the-art performance across diverse data modalities, their design remains largely empirical, lacking a systematic framework for use in domain-specific applications across science and engineering. In this talk, I will introduce a mathematical framework that unifies flows and diffusions, substantially expanding the design space of flow-based generative models. To this end, I will define a process called a stochastic interpolant that establishes an exact connection between two arbitrary probability measures in finite time. I will then show how this construction enables efficient learning of generative models described by ordinary or stochastic differential equations. Empirically, I will demonstrate that our approach outperforms diffusion models in high-resolution image synthesis at no additional computational cost, and that it enjoys the flexibility to trade computational budget for sample quality at inference time. I will further illustrate how to apply the framework to design tailored generative models for problems in inverse imaging and probabilistic forecasting of turbulent fluids. Building on this foundation, I will conclude by describing our recent extension of the framework to learning the flow map of an ordinary differential equation, which avoids solving a differential equation at inference time and can generate samples in a single neural network evaluation. This dramatically improves the efficiency of modern generative models and makes feasible real-time engineering applications such as robotic planning and control.

Xiaochuan Tian. *Sparse Neural Network Solutions to PDEs: Variational and RKBS Perspectives*
Abstract: We propose a general framework for solving nonlinear PDEs using neural networks. To avoid over-parameterization and eliminate redundant features, a regularization approach that encourages sparsity is explored within the framework of shallow radial basis function (RBF) networks. An adaptive training process iteratively adds neurons to maintain a compact network structure. Existence theory is established via the calculus of variations, and a representer theorem is derived for reproducing kernel Banach spaces (RKBS) associated with one-hidden-layer neural networks of possibly infinite width. An error estimate for the neural network approximation to the PDE is also derived. Training is performed using a second-order semismooth Newton method with gradient boosting. The approach is compared with the reproducing kernel Hilbert space (RKHS) framework and Gaussian process

methods. This is a joint work with Konstantin Pieper (Oak Ridge National Lab) and Zihan Shao (UC San Diego).

Giovanni Brigati. *Kinetic Optimal Transport*
Abstract: In the recent preprint https://arxiv.org/pdf/2502.15665, in collaboration with Jan Maas and Filippo Quattrocchi, we study three possible notion of “minimal global acceleration” between probability measures, allowing a fixed time T in order to transfer mass from one measure to another. Equivalence among the three definitions is shown, together with existence of “optimal kinetic transport plans/maps/dynamical transport plans” and even of a force field steering particles along superpositions of Newton’s equations in an optimal way.

Then, a non-parametric notion of minimal acceleration is introduced, where the time span T is further resource to be optimised. This way, we achieve a discrepancy functional on the space of probability measures, inducing a geometry which is adapted to kinetic equations. Actually, symplectic and degenerate Riemannian effects combine in a way which is reminiscent of Hörmander’s hypoellipticity. Probability measure-valued curves, which are absolutely continuous in the new geometry, are fully characterised as solutions to Vlasov’s PDEs, up to time reparametrisations. This fact reflects into an explicit structure of the associated metric tangent cone, and of the consequent computation of the metric derivative of kinetic-a.c. curves.

Applications are discussed in relation with interpolation between images, optimal steering of vehicles, and trajectory inference in biology.

Martin Rumpf. *Regularizing autoencoder for manifold embedding and Riemannian calculus on latent spaces*
Abstract: Autoencoders are widely used in machine learning for dimension reduction of high-dimensional data. The encoder embeds the input data manifold into a lower-dimensional latent space, while the decoder represents a parametrization of the data manifold by the manifold in latent space.

The talk presents and analyzes a regularization of the encoder based on a loss functional that prefers isometric, extrinsically flat embeddings. The loss functional is computed via Monte Carlo integration with different sampling strategies for pairs of points on the input manifold. A geometric loss functional of the embedding map is shown to be the Γ-limit of the sampling-dependent loss functionals.

Furthermore, it will be discuss how to implement a Riemannian calculus directly on the lower-dimensional latent manifold. To this end, the latent manifold is defined as an implicit surface in latent space based on a learned projection. Then, geodesic boundary and initial value problems are discretized with the implicit surface constraint.

Numerical experiments show that smooth manifold embeddings into latent space are obtained and Riemannian interpolation and extrapolation can be efficiently computed on the latent manifold.

This is joint work with Juliane Braunsmann, Florine Hartwig, Marko Rajković, Josua Sassen, and Benedikt Wirth.

Jan-Eric Sulzbach. *Slow manifolds in multiscale PDEs and their application*
Abstract: In this talk I will introduce the concept of slow manifolds as an approach to reduce

the complexity of a multiple time-scale PDE system. The idea is that in a fast-slow system the fast variable approaches a state in which its dynamics only depend on the slow variable. Therefore, one can reduce the system to a self-contained and lower dimensional equation, posed on the slow manifold, which only depends on the slow variable.

In the first part of this talk I will present the main ideas of this approach for a finite dimensional dynamical system. This builds the foundation for the second part, where the goal is to extend the results to an infinite-dimensional setting in general Banach spaces. The main result is a generalization of the so-called Fenichel-Tikhonov theory. In the last part of the talk I will present applications to biological and ecological systems and present a way to approximate the slow manifold numerically.

This talk is based on joint work with Christian Kuehn (TUM).

Jeff Calder. *On the continuum limit of t-SNE visualization*
Abstract: The t-stochastic neighbor embedding (t-SNE) method, and its variants such as UMAP, have become the de facto methods for visualizing high dimensional data in a wide variety of fields. Nevertheless, their behavior and properties are very poorly understood, and they often generate artifacts in visualizations that do not represent features in the original data. In this talk, we will give an overview of t-SNE, and discuss some of our recent work towards understanding t-SNE through a continuum limit variational analysis. t-SNE is based on attraction/repulsion dynamics, and the continuum limit is a gradient regularized variational problem with strong nonconvexities in the gradient. Surprisingly, the one dimensional problem is well-posed. We will give an overview of the proof of this, and give some conjectures for the higher dimensional setting.

Joint work with Ryan Murray, Adam Pickarski, Jingcheng Lu.

Wilfrid Gangbo. *Hamilton-Jacobi equations in non-commutative variables*
Abstract: We study Hamilton-Jacobi equations for games where the order of strategic choices matters and the number of players can be infinite. For instance, we study optimal control problems on non commutative state spaces such as $M_n(\mathbb{C})^d_{sa}$, where $d \in \mathbb{N}$ and $M_n(\mathbb{C})_{sa}$ is the set of self-adjoint matrices with complex entries. The theory of free probability provides us with quotient spaces which ensures that the sub-level sets of our actions are compact. This allows us to describe control problems involving the large-n limit of Brownian motion on $M_n(\mathbb{C})^d_{sa}$.(The essence of this talk is based on a collaborative work with D. Jekel, K. Nam and A. Palmer).

Elias Hess-Childs. *A universal total anomalous dissipator*
Abstract: Anomalous dissipation describes the tendency for a turbulent fluid to dissipate energy at a constant rate independent of the molecular viscosity, despite viscosity being the ultimate mechanism of dissipation. It is a cornerstone of phenomenological turbulence theory and is taken as a basic axiom in Kolmogorov's highly successful K41 theory. Despite this, a rigorous mathematical demonstration of these effects in fluid models remains elusive.

To study anomalous dissipation in a more tractable setting, recent work has focused on constructing incompressible vector fields that induce persistent energy loss in scalar advection-diffusion equations in the vanishing noise limit. In this talk, I will provide an overview of anomalous dissipation and discuss my recent work with Keefer Rowan, where we construct a universal total anomalous dissipator—a vector field that completely dissipates any initial data in unit time in the vanishing noise limit.

Specifically, we construct a vector field such that the laws of the associated SDEs remain diffuse even as the noise vanishes. In fact, as the noise approaches zero, the laws converge to the uniform distribution on the torus.

Andrew Warren. *Estimation of one-dimensional structures in data*
Abstract: Given a data distribution which is concentrated around a one-dimensional structure, can we infer that structure? We consider versions of this problem where the distribution resides in a metric space and the 1d structure is assumed to either be the range of an absolutely continuous curve, or a connected set of finite 1d Hausdorff measure. In each of these cases, we relate the inference task to solving a variational problem where there is a tradeoff between data fidelity and simplicity of the inferred structure; the variational problems we consider are closely related to the so-called "principal curve" problem of Hastie and Stuetzle as well as the "average-distance problem" of Buttazzo, Oudet, and Stepanov. For each of the variational problems under consideration, we establish existence of minimizers, stability with respect to the data distribution, and consistency of a discretization scheme which is amenable to Lloyd-type numerical methods. Lastly, we consider applications to estimation of stochastic processes from partial observation, as well as the lineage tracing problem from mathematical biology.

Selim Esedoglu *Vectorial median filters and curvature motion of networks*
Abstract: The median filter is a standard tool in image processing (a typical application is removing salt and pepper noise from digital pictures). It is also well known as a monotone discretization of the level set formulation of two-phase motion by mean curvature.

In fact, median filters turn out to be closely connected, in a precise sense, to another class of algorithms for curvature motion: threshold dynamics. This precise connection between two disparate families of algorithms – level set methods and threshold dynamics – allows porting over to one (the level set) world what we have learned in recent years in the context of the other (threshold dynamics).

In particular, we give a variational (minimizing movements) formulation of the median filter (and thus exhibit a Lyapunov function, implying unconditional energy stability). Furthermore, we discuss extending median filters to multiphase, weighted motion by mean curvature of networks. The resulting level set methods are new: they are vectorial median filters. They avoid the issues that plagued previous attempts in the level set literature for this challenging evolution, and can accommodate a much wider range of surface tensions.