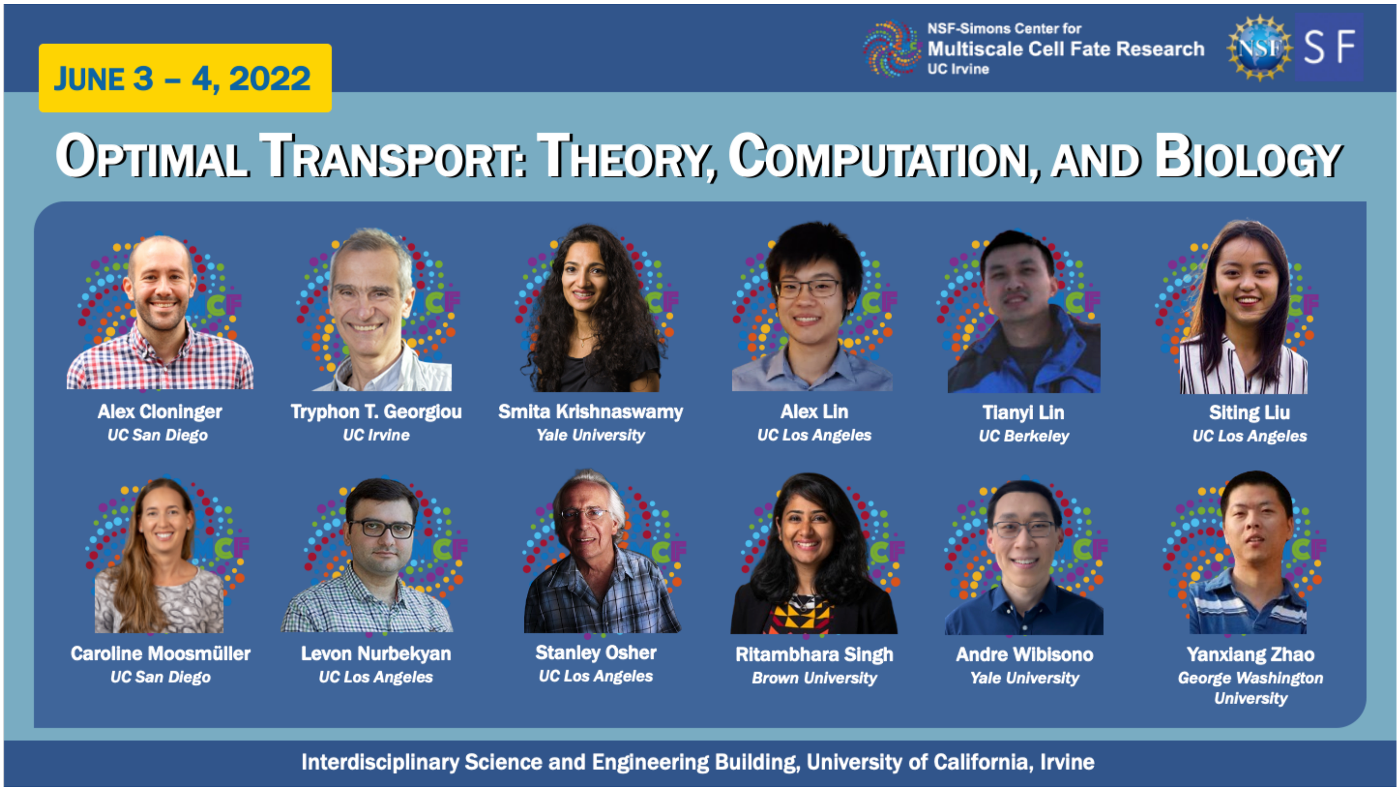

Day 1: Friday, June 3, 2022

07:50 – 08:00 Opening
Location: ISEB 1010

08:00 – 08:30 Stanley Osher
UC Los Angeles
Conservation Laws and Generalized Optimal Transport

08:35 – 09:05 Alex Cloninger
UC San Diego
Efficient Distribution Classification via Linearized Optimal Transport Embeddings

09:10 – 09:40 Ritambhara Singh
Brown University
Single-Cell Data Integration Using Gromov-Wasserstein Optimal Transport

09:45 – 10:10 Coffee Break

10:10 – 10:40 Smita Krishnaswamy
Yale University
Diffusion Earth Mover's Distance and Distribution Embeddings

10:45 – 11:15 Andre Wibisono
Yale University
Rapid Convergence of the Langevin Algorithm under Isoperimetry

11:20 – 11:50 Levon Nurbekyan
UC Los Angeles
Efficient Natural Gradient Method for Large-Scale Optimization Problems

12:00 – 01:00 Lunch

01:00 – 03:00 Discussion Session I
Location: ISEB 3010, 3020, 4010, 4044

03:00 – 03:30 Coffee Break

03:30 – 05:30 Discussion Session II
Location: ISEB 3010, 3020, 4010, 4044

Day 2: Saturday, June 4, 2022

08:00 – 09:45 Discussion Session I
Location: ISEB 3010, 3020, 4010, 4044

09:45 – 10:10 Coffee Break
Location: ISEB 4th Floor Terrace

10:15 – 12:00 Discussion Session II
Location: ISEB 3010, 3020, 4010, 4044

12:00 – 01:00 Lunch
Location: ISEB 4th Floor Terrace

All talks location: ISEB 1010

01:00 – 01:30 Tryphon T. Georgiou
UC Irvine
Optimal Mass Transport Meets Thermodynamics: On Power and Efficiency of Finite-Time Thermodynamic Engines

01:35 – 02:05 Siting Liu
UC Los Angeles
Controlling the Propagation of Epidemics via Mean-Field Control

02:10 – 02:40 Yanxiang Zhao
George Washington University
Supervised Optimal Transport

02:45 – 03:10 Coffee Break
Location: ISEB 1010

03:15 – 03:40 Caroline Moosmüller
UC San Diego
Analysis of Biological Data Sets via Optimal Transport

03:45 – 04:15 Tianyi Lin
UC Berkeley
Computational Wasserstein Barycenters: Combinatorial Structures and Accelerated Algorithms

04:20 – 04:50 Alex Lin
UC Los Angeles
APAC-Net: Alternating the Population and Agent Control via Two Neural Networks to Solve High-Dimensional Stochastic Mean Field Games

04:55 – 05:00 Closing Remarks

ABSTRACTS

Day 1: Friday, June 3, 2022

Conservation Laws and Generalized Optimal Transport

Stanley Osher, UCLA

In this talk, we connect Lax's entropy-entropy flux in conservation laws with optimal transport type metric spaces. Following this connection, we further design variational discretizations for conservation laws and mean field control of conservation laws. In particular, we design unconditionally stable time discretization methods that are easy to implement.

Efficient Distribution Classification via Linearized Optimal Transport Embeddings

Alex Cloninger, UCSD

Detecting differences and building classifiers between distributions, given only finite samples, are important tasks in a number of scientific fields. Optimal transport (OT) has evolved as the most natural concept to measure the distance between distributions, and has gained significant importance in machine learning in recent years. The largest drawback of OT in machine learning is that it only provides an unsupervised distance between distributions. In this talk, we discuss how optimal transport embeddings can be used to deal with this issue, both on a theoretical and a computational level. In particular, we'll show how to embed the space of distributions into an L2-space via OT, and how linear techniques can be used to classify families of distributions generated by simple group actions in any dimension. The proposed framework significantly reduces both the computational effort and the required training data in supervised settings. We demonstrate the benefits in pattern recognition tasks in imaging and provide some medical applications.

Single-Cell Data Integration Using Gromov-Wasserstein Optimal Transport

Ritambhara Singh, Brown University

Recent advances in sequencing technologies have allowed us to capture various aspects of the genome at single-cell resolution. However, with the exception of a few co-assaying technologies, it is not possible to simultaneously apply different sequencing assays on the same single cell. In this scenario, computational integration of multi-omic measurements is crucial to enable joint analyses. This integration task is particularly challenging due to the lack of sample-wise or feature-wise correspondences. In this talk, I will present our method - Single-Cell alignment with Optimal Transport (SCOT) - an unsupervised algorithm that uses Gromov-Wasserstein optimal transport to align single-cell multi-omics datasets. I will discuss features of SCOT - like automatic hyperparameter selection, its performance, and its extension to the integration setting where datasets can have disproportionate cell type representations.

Diffusion Earth Mover's Distance and Distribution Embeddings

Smita Krishnaswamy, Yale University

We propose a new fast method of measuring distances between large numbers of related high dimensional datasets called the Diffusion Earth Mover's Distance (EMD). We model the datasets as distributions supported on a common data graph that is derived from the affinity matrix computed on the combined data. In such cases where the graph is a discretization of an underlying Riemannian closed manifold, we prove that Diffusion EMD is topologically equivalent to the standard EMD with a geodesic ground distance. Diffusion EMD can be computed in {Õ}(n) time and is more accurate than similarly fast algorithms such as tree-based EMDs. We also show Diffusion EMD is fully differentiable, making it amenable to future uses in gradient-descent frameworks such as deep neural networks. Finally, we demonstrate an application of Diffusion EMD to single cell data collected from 210 COVID-19 patient samples at Yale New Haven Hospital. Here, Diffusion EMD can derive distances between patients on the manifold of cells at least two orders of magnitude faster than equally accurate methods. This distance matrix between patients can be embedded into a higher level patient manifold which uncovers structure and heterogeneity in patients. More generally, Diffusion EMD is applicable to all datasets that are massively collected in parallel in many medical and biological systems.

Rapid Convergence of the Langevin Algorithm under Isoperimetry

Andre Wibisono, Yale University

Sampling is a fundamental algorithmic task. Many modern applications require sampling from complicated probability distributions in high-dimensional spaces. While the setting of log concave target distribution is well-studied, it is important to understand the guarantees of sampling beyond the log concavity assumption. We study the Unadjusted Langevin Algorithm (ULA) for sampling from a probability distribution under isoperimetry conditions. ULA is a discretization of the Langevin dynamics in continuous time, which has a natural optimization interpretation as the gradient flow for minimizing Kullback-Leibler (KL) divergence under Wasserstein metric; moreover, isoperimetry also has a natural optimization interpretation as a sufficient condition for exponential convergence rate. Using this optimization perspective, we establish a convergence guarantee for ULA in KL divergence assuming the target distribution satisfies log-Sobolev inequality and the log density has bounded Hessian. Notably, we do not assume convexity or bounds on higher derivatives. We also show convergence guarantees in Rényi divergence assuming the limit of ULA satisfies either log-Sobolev or Poincaré inequality.

Efficient Natural Gradient Method for Large-Scale Optimization Problems

Levon Nurbekyan, UCLA

We propose an efficient numerical method for computing natural gradient descent directions with respect to a generic metric in the state space. Our technique relies on representing the natural gradient direction as a solution to a standard least-squares problem. Hence, instead of calculating, storing, or inverting the information matrix directly, we apply efficient methods from numerical linear algebra to solve this least-squares problem. We treat both scenarios where the derivative of the state variable with respect to the parameter is either explicitly known or implicitly given through constraints. We apply the QR decomposition to solve the least-squares problem in the former case and utilize the adjoint-state method to compute the natural gradient descent direction in the latter case.

Day 2: Saturday, June 4, 2022

Optimal Mass Transport Meets Thermodynamics: On Power and Efficiency of Finite-Time Thermodynamic Engines

Tryphon Georgiou, UCI

We discuss principles that underlie power transduction in thermodynamic systems that are subject to thermal fluctuations. Specifically, we will discuss how anisotropy in thermal fluctuations (chemical, electrical potential, and so on) allows generation of power in engineered and physical processes, and we will present specific embodiments of pertinent controlled power-generating mechanisms. Our theme follows a long line of developments, from classical and non-equilibrium thermodynamics, onto recent formalisms that aim to quantify precisely such finite-time thermodynamic transitions and power generation. The talk is based on joint works with Rui Fu (UCI), Olga Movilla (UCI), Amir Taghvaei (UCI) and Yongxin Chen (GaTech). Research funding by NSF and AFOSR is gratefully acknowledged.

Controlling the Propagation of Epidemics via Mean-Field Control

Siting Liu, UCLA

The coronavirus disease 2019 (COVID-19) pandemic is changing and impacting lives on a global scale. In this talk, I will introduce a mean-field game model for controlling the propagation of epidemics in a spatial domain. The control variable, the spatial velocity, is first introduced for the classical disease models, such as the Suspected-infected-recovered model. We provide fast numerical algorithms based on proximal primal-dual methods for this proposed model. Numerical experiments demonstrate that the proposed model illustrates how to separate infected patients in a spatial domain effectively. This talk is based on joint work with Wonjun Lee, Wuchen Li, Hamidou Tembine, and Stanley Osher.

Supervised Optimal Transport

Yanxiang Zhao, George Washington University

Optimal Transport, a theory for optimal allocation of resources, is widely used in various fields such as astrophysics, machine learning, and imaging science. However, many applications impose element-wise constraints on the transport plan which traditional optimal transport cannot enforce. Here we introduce Supervised Optimal Transport (sOT) that formulates a constrained optimal transport problem where couplings between certain elements are prohibited according to specific applications. sOT is proved to be equivalent to an l1 penalized optimization problem, from which efficient algorithms are designed to solve its entropy regularized formulation. We demonstrate the capability of sOT by comparing it to other variants and extensions of traditional OT in color transfer problems. We also study the barycenter problem in sOT formulation, where we discover and prove a unique reverse and portion selection (control) mechanism. Supervised optimal transport is broadly applicable to applications in which a constrained transport plan is involved and the original unit should be preserved by avoiding normalization.

Analysis of Biological Data Sets via Optimal Transport

Caroline Moosmüller, UCSD

Biological data sets, such as gene expressions or protein levels, are often high-dimensional, and thus difficult to interpret. Finding important structural features and identifying clusters in an unbiased fashion is a core issue for understanding biological phenomena. In this talk, we will describe unsupervised data analysis methodologies based on optimal transport to analysis gene expression data and time-series that describe biological processes such as the cell cycle.

Computational Wasserstein Barycenters: Combinatorial Structures and Accelerated Algorithms

Tianyi Lin, UCB

We study the fixed-support Wasserstein barycenter problem (FS-WBP), which consists in computing the Wasserstein barycenter of m discrete probability measures supported on a finite metric space of size n. We show first that the constraint matrix arising from the standard linear programming (LP) representation of the FS-WBP is not totally unimodular when m >= 3 and n >= 3. This result resolves an open question pertaining to the relationship between the FS-WBP and the minimum-cost flow (MCF) problem since it proves that the FS-WBP in the standard LP form is not an MCF problem when m >= 3 and n >= 3. We also develop a provably fast deterministic variant of the celebrated iterative Bregman projection (IBP) algorithm with a complexity bound of O(mn^{7/3}\eps^{−4/3}), where \eps \in (0, 1) is the desired tolerance. We also conduct extensive experiments with both synthetic data and real images and demonstrate the favorable performance of the proposed algorithm in practice.

APAC-Net: Alternating the Population and Agent Control via Two Neural Networks to Solve High-Dimensional Stochastic Mean Field Games

Alex Lin, UCLA

We present APAC-Net, an alternating population and agent control neural network for solving stochastic mean field games (MFGs). Our algorithm is geared toward high-dimensional instances of MFGs that are beyond reach with existing solution methods. We achieve this in two steps. First, we take advantage of the underlying variational primal-dual structure that MFGs exhibit and phrase it as a convex-concave saddle point problem. Second, we parameterize the value and density functions by two neural networks, respectively. By phrasing the problem in this manner, solving the MFG can be interpreted as similar to training a generative adversarial network (GAN). We show the potential of our method on up to 100-dimensional MFG problems.