Interacting Particle Systems and Ensemble Based Algorithms
Abstracts

5–8 October 2026

September 28, 2026

A Fast Spectral Particle Method for the Landau Equation

Giacomo Borghi

Heriot Watt University, UK

The talk will present a novel deterministic particle method for the homogeneous Landau equation. By exploiting the continuity formulation of the PDE, the spectral particle method reconstructs the velocity vector field with fast Fourier transform algorithms, reaching a computational cost which is linear in the number of particles. Theoretical analysis highlights the error decomposition while numerical experiments show how the spectral particle method is able to obtain highly accurate solutions with respect to baseline methods. Joint work with Lorenzo Pareschi.

Transformers for Operator Learning on Probability Measures: From Theory to Scientific Workflows

Edoardo Calvello

Lawrence Berkeley National Laboratory, USA

Attention mechanisms, the core of transformer architectures that power modern large language models, have been widely successful at modeling nonlocal correlations in data. Motivated by recent formulations of attention as a measure-to-measure mapping, in this talk we show how the mathematical properties of the attention mechanism enable the design of neural architectures approximating nonlinear operators on probability measures. We leverage this perspective to design novel methods for nonlinear data assimilation, surpassing the performance of ensemble Kalman method baselines. Starting from the theory developed, we discuss key considerations and open questions towards the development of scalable transformer-based Bayesian inference workflows.

Convergence of Kernelized Wasserstein Flows

Lénaïc Chizat

National University of Singapore, Singapore

Several machine-learning algorithms can be described, in the mean-field limit, by continuity equations whose velocity field is regularized by a smoothing kernel. Examples include Wasserstein gradient flows of kernel mean discrepancies (KMD), which arise in the large-width limit of shallow neural-network training, and Stein variational gradient flow (SVGF), the continuum limit of a deterministic interacting-particle method for sampling. Despite their popularity, even convergence of these dynamics has remained open in many settings.

I will present an approach that establishes quantitative local convergence for these flows, yielding sharp polynomial rates for Riesz-type kernels. In particular, it provides the first convergence guarantees for shallow networks with regular target parameter distributions. The analysis reveals how regularity of the target and of the kernel govern the speed of convergence. The key difficulty is that the energy dissipation controls a weaker norm than the energy; the proof overcomes this loss of coercivity through Sobolev interpolation and propagation of regularity.

This is joint work with Maria Colombo, Roberto Colombo and Xavier Fernández-Real.

Opinion Dynamics: Socio-economic Factors and Social Networks

Bertram Düring

University of Warwick, UK

In recent years the study of many-agent social systems have received growing attention. Mathematical methods from kinetic theory provide a powerful approach to study these phenomena. In the modelling of opinion dynamics this leads to (systems of) Boltzmann-type equations and approximate Fokker-Planck-type equations.

Voter demographics and socio-economic factors play an important role in opinion formation. We discuss recent attempts to include voter demographics and socio-economic factors in kinetic leader-follower models of opinion formation and show simulation results for recent general elections in the United Kingdom.

We also discuss how to introduce topological information about the social network in kinetic opinion formation models. We discuss a graphon-based kinetic model for selective media influence inspired by the theory of Hallin‘s spheres.

TBA

Franca Hoffmann

California Institute of Technology, USA

Pending

Continuity Equations with Non-local Mobility and Applications to Self-attention

Daniel Kelly

University of Oxford, UK

We will consider a modified aggregation equation in which the associated velocity field is normalized by a non-local mobility term. The motivation to study such an equation comes from applications in machine learning. In particular, recent advances in large language model capabilities have been predicated on the development of the "transformer" architecture and the underlying mechanism of "self-attention". Fundamentally, self-attention constitutes a system of interacting particles, for which our model may be viewed as a continuum limit. Previous works formally suggested a gradient flow interpretation of this equation with respect to a modified Wasserstein distance. This characterization has been made rigorous when the underlying spatial domain is compact. We will consider the case of an unbounded spatial domain.

Uniform-in-Time Propagation of Chaos and Long-Time Asymptotics for Consensus-Based Optimization

Dohyean Kim

California Institute of Technology, USA

Consensus-Based Optimization (CBO) is a class of derivative-free methods for global optimization based on interacting stochastic particle systems. Particles are attracted toward an objective-weighted consensus point that favors regions of low objective values. The resulting nonlinear and nonlocal interactions pose mathematical challenges for understanding both the mean-field approximation and the long-time behavior of these systems.

In this talk, we establish uniform-in-time propagation of chaos for first-order CBO and its kinetic (second-order) extension. Using stability estimates for the weighted consensus point and quantitative concentration inequalities, we show that the particle

system tracks its mean-field limit at the classical Monte Carlo rate uniformly over infinite time horizons. For CBO with anisotropic noise, the error prefactor is independent of the problem dimension. We also discuss almost sure exponential convergence of the particle system and the additional challenges in establishing long-time stability when velocity dynamics are introduced.

We then turn to the long-time asymptotics of a dynamically regularized first-order CBO model. By combining Wasserstein convergence with Fisher-information estimates, we establish convergence of the rescaled mean-field density in L^1, characterizing its asymptotic profile as the original distribution concentrates around a consensus point.

The propagation-of-chaos results are joint work with Nicolai Gerber, Franca Hoffmann, and Urbain Vaes (first order), and with Seung-Yeal Ha and Franca Hoffmann (second order). The long-time asymptotics results are joint work with Wuchen Li and Franca Hoffmann.

Inferring Interactive Kernels in the Mean-field Regime

Qin Li

University Wisconsin-Madison, USA

Inverse problem is to infer unknown parameters in a system using measurement. We study how to infer interactive kernels using macroscopic quantities such as density of an interactive particle system. Would the reconstruction agree if one views the system as finite number of agents or as represented by an absolutely continuous density? We show that these two perspectives do not give the same reconstruction, but are asymptotically equivalent.

Evolving Ensemble of Agents

Yang Liu

National University of Singapore, Singapore

We introduce the Evolving Ensemble of Agents (EvE), a decentralized framework that organizes existing, highly capable coding agents into a live, co-evolving system for algorithmic discovery. Rather than reinventing the wheel within the "LLMs as optimizers" paradigm, EvE fixes the base agent substrate and focuses entirely on evolving the cumulative guidance and skills that dictate agent behaviors. By maintaining two co-evolving populations, namely functional code solvers and agent guidance states, the system evaluates agents through a synchronous race, updating their empirical Elo ratings based on the marginal gains they contribute to the current solver state. When applied to a research bottleneck in In-Context Operator Networks (ICON), EvE autonomously discovered a robust rescale-then-interpolate mechanism that enables reliable example-count generalization. Crucially, controlled ablations reveal the absolute necessity of stage-dependent agent adaptation to navigate the shifting search landscapes of complex codebases. Compared to variants driven by a fixed initial agent or even a frozen "best-evolved" agent, EvE uniquely avoids phase mismatch, demonstrating that organizing agents into a self-revising ensemble is the fundamental driver for breaking through static performance ceilings.

TBA

Tan Minh Nguyen
National University of Singapore, Singapore

Pending

Long-Time Accuracy of Ensemble Kalman Filters: Chaotic Dynamics, Learned Models, and Localization

Daniel Sanz-Alonso
The University of Chicago, USA

Ensemble Kalman filters are widely used for state estimation in high-dimensional dynamical systems, but understanding their accuracy over long time horizons remains challenging. I will discuss two complementary results. First, for partially observed chaotic systems—including Lorenz models and the Navier–Stokes equations—we establish long-time accuracy of suitably inflated ensemble Kalman filters. The theory also accommodates surrogate dynamics, providing guarantees when the forecast model is learned from data. Second, for deterministic square-root ensemble Kalman filters in linear–Gaussian systems, we prove high-probability bounds, uniform in time, for mean and covariance errors relative to the Kalman filter. The required ensemble size depends on effective rank and unstable dimension rather than ambient dimension. For weakly coupled spatial systems, localization further reduces this dependence to local intrinsic dimension, at the cost of a controlled localization bias. Together, these results identify mechanisms that enable accurate long-time filtering with small ensembles. This talk is based on joint work with Xiaoou Cheng and Nathan Waniorek.

The Stein-log-Sobolev Inequality and the Exponential Rate of Convergence for the Continuous Stein Variational Gradient Descent Method

Jakub Skrzeczkowski

University of Oxford, UK

The Stein Variational Gradient Descent method is a variational inference method in statistics that has recently received a lot of attention. The method provides a deterministic approximation of the target distribution, by introducing a nonlocal interaction with a kernel. Despite the significant interest, the exponential rate of convergence for the continuous method has remained an open problem, due to the difficulty of establishing the related so-called Stein-log-Sobolev inequality. Here, we prove that the inequality is satisfied for each space dimension and every kernel whose Fourier transform has a quadratic decay at infinity and is locally bounded away from zero and infinity. Moreover, we construct weak solutions to the related PDE satisfying exponential rate of decay towards the equilibrium. The main novelty in our approach is to interpret the Stein-Fisher information as a duality pairing between H^{-1} and H^1, which allows us to employ the Fourier transform. We also provide several examples of kernels for which the Stein-log-Sobolev inequality fails, partially showing the necessity of our assumptions. This is a joint work with J. A. Carrillo and J. Warnett.

Wasserstein Gradient Flow of Maximum Mean Discrepancy with Energy Kernel

Lihan Wang

National University of Singapore, Singapore

Wasserstein gradient flows of Maximum Mean Discrepancy have recently found applications in statistics and machine learning. The standard theory of Wasserstein gradient flows in spaces of probability measures provides well-posedness for smooth kernels, however, the kernels that perform well in applications are not differentiable and the theory does not apply. Here we consider the "energy kernels" $k(x, y) = -|x - y|^q$ for $q \in (0, 2)$ which allow for long-range effects. We develop well-posedness theory in L^p spaces and show that these solutions can be approximated well by interacting particles, which justifies their use. We also provide results that show that at least in 1D the convergence can be exponential when the solution is close to the target and also prove that no quantitative rate of convergence holds for general initial data. Joint work with Matthew Rosenzweig and Dejan Slepčev (CMU).

Stochastic Interacting Particle Field Method in the Computation of Keller-Segel Systems

Zhongjian Wang

Nanyang Technological University, Singapore

In this talk, I will discuss the recent progress on our Stochastic Interacting Particle Field method in the computation of Keller-Segel systems. The main theme of the SIPF methods is to combine the particle approximation to the densities and the Fourier approximation to the chemo-attractant fields. Such framework is shown to be general applicable in various setup of Keller-Segel systems, including parabolic-elliptic, fully parabolic and Haptotaxis models. It is also applicable to Reaction-Diffusion-Advection equations related to the nonlinear wave front speed calculations. We will also discuss the convergence analysis of the SIPF where we will take the latest particle-in-cell derivative as an example.

Stein Variational Gradient Descent dynamics for Highly Concentrated Kernels

Jethro Warnett

University of Oxford, UK

Stein Variational Gradient Descent (SVGD) is a widely used in practice algorithm for scalable sampling with deterministic particle updates. We study its behavior in the singular limit where the kernel bandwidth tends to zero. In this regime, we show that the nonlocal SVGD dynamics converge to a local evolution equation that can be formally interpreted as a Wasserstein gradient flow with quadratic mobility. We analyze this singular limit in two settings: integrable kernels and weighted kernels. In the weighted case, the proof is supported by recently established Stein-log-Sobolev inequalities, which provide the necessary functional control. Overall, our results clarify how SVGD collapses from a nonlocal interacting particle system to a local gradient-flow dynamics as the kernel concentrates.

Score-Based Diffusion Models: Dynamics, Geometry, and Learning

Enrique Zuazua

Friedrich Alexander University of Erlangen-Nürnberg, Germany

Score-based diffusion models have become a central paradigm in generative machine learning, while their mathematical foundations remain only partially understood. In this talk, we develop a PDE and dynamical-systems perspective on generation, guidance, and score learning.

Using the heat flow as a canonical framework, we first exploit Li–Yau inequalities and entropy methods to establish stability and concentration of reverse-time dynamics toward the data support. We then study mixed score fields, including mixture-of-experts and classifier-free guidance, and show through Laplace–Varadhan asymptotics that their small-time dynamics is governed by explicit geometric potentials determined by the data supports.

Finally, we address learned scores as dynamic approximate representations, asking which analytical and geometric structures must be preserved to ensure reliable generation. This leads to Wasserstein and entropy stability estimates and to the hierarchy representation –structure – stability–generalization, connecting stochastic dynamics, PDE analysis, geometry, and machine learning.
