Extreme and Rare Events: Modeling, Algorithms, and Applications
Poster Session Abstracts
September 26, 2026
2:30-4:00pm

Mechanisms and Pathways of Extreme Events in Partially-Observed Stochastic Dynamical Systems
Charlotte Moser, University of Wisconsin Madison

Extreme events occur across the natural, engineering, and socioeconomic sciences, where rare but high-impact episodes can lead to disproportionate consequences that pose major challenges for prediction and risk management. Existing studies have mainly focused on the statistics, sampling, forecasting, and attribution of extremes from observable variables. In this paper, we develop a mathematical framework for studying the mechanisms and pathways of extreme events in partially-observed stochastic dynamical systems with hidden variables. By integrating data assimilation with information-theoretic and trajectory-based diagnostics, we infer latent precursor dynamics from observations, quantify their uncertainty, and determine how their influence propagates toward observed extreme events. Conditional Gaussian models provide a tractable analytical setting for deriving closed-form diagnostics, while the framework extends through numerical methods. The analysis proceeds from two complementary perspectives. From a trajectory-wise viewpoint, we compare filtering and smoothing distributions to identify the onset of hidden precursors and quantify temporal influence. From a statistical viewpoint, we construct event-conditioned hidden-state distributions to identify sensitive triggering directions, latent pathways, and multiple classes of extreme-event mechanisms through clustering. Three numerical examples illustrate the methodology. In an intermittent stochastic model, hidden damping dynamics emerge before observed bursts, where discrepancies between the filter and smoother provide an onset diagnostic. In a stochastic model with damping and forcing, separate damping-induced, forcing-driven, and mixed pathways to extremes are identified. In a nonlinear topographic-flow model, distinct mechanisms and pathways for blocking and unblocking patterns associated with observed extreme events are revealed.

Efficient Blackbox Posterior Approximation via Gaussian process
Hwanwoo Kim, University of Texas at Dallas

Gaussian processes (GPs) provide a powerful framework for surrogate modeling in Bayesian inference problems where likelihood evaluations are computationally expensive. However, standard GP surrogates can exhibit poor approximation quality when design points are selected suboptimally. In this poster, I will present enhanced GP surrogate methods that strategically combine random exploration with upper confidence bound (UCB) acquisition strategies, specifically tailored for posterior approximation tasks. I will demonstrate how this hybrid sampling approach significantly improves posterior approximation accuracy while enabling efficient uncertainty quantification. Through applications in pharmacokinetic/pharmacodynamic (PK/PD) models, I will illustrate the practical advantages of this methodology for computationally intensive inference problems commonly encountered in scientific applications.

Efficient Importance Sampling of Brownian Motion Avoiding Poissonian Obstacles for Long Time
Hao Liu, Stanford University

We design and analyze an efficient importance sampling estimator of the probability that, over a long time horizon, a Brownian particle avoids obstacles of a fixed radius placed according to an independent Poisson random field. Donsker and Varadhan identify the exponential decay rate of this probability as the optimal balance between the volume cleared of obstacles and the spectral cost of confinement, and the confinement results of Sznitman and Povel show that the particle survives by localizing its path in a ball of (roughly) deterministic radius but with a random center. This strategy makes direct Monte Carlo exponentially ineffective over long horizons, and makes a practical yet asymptotically optimal sampler non-trivial to design, since the sampler must track the random region swept out by the path. For every fixed dimension \(d\ge2\) and every fixed obstacle radius, we construct a sequential importance sampler that is unbiased and asymptotically optimal: its coefficient of variation grows sub-exponentially in the time horizon. The proof-oriented allocation uses only a subexponential number of fine Gaussian increments in its candidate-augmented skeleton; individual candidates are generated and scored independently, which suits parallel implementation. In contrast to natural sequential importance sampling ideas, our method avoids solving eigenvalue problems repeatedly on the evolving explored regions. Numerical experiments illustrate the efficiency of the method.

Fine-Tuning Generative Models for Extreme Events via CVaR-Penalized Wasserstein Gradient Flows
Zhizhen Zhang, Umass Amherst

We propose CVaR-penalized Generative Particle Algorithm (CVaR-GPA), a robust, tail-agnostic algorithm for fine-tuning generative models to learn heavy-tailed distributions and capture extreme events, requiring no prior knowledge or estimation of the target's tail characteristics. The method is the Wasserstein gradient flow of the Lipschitz-regularized Kullback-Leibler (KL) divergence penalized by a Conditional Value-at-Risk (CVaR) discrepancy term: the Lipschitz-regularized KL divergence enables robust learning under minimal assumptions on the target distribution, while the CVaR penalty restores the velocity that otherwise vanishes prematurely in the under-sampled tails. The penalized flow admits a bounded but non-Lipschitz velocity field. This departs from the Lipschitz transport maps of standard generators, which preserve the tail behavior of a light-tailed source, and enables transport toward heavier-tailed targets. To define this flow on empirical measures, we derive the first-variation subgradients of CVaR from its Rockafellar-Uryasev representation, valid precisely where the classical density-based formula fails. The particle algorithm CVaR-GPA fine-tunes the output samples of any pre-trained model, without access to its architecture, and runs on an adaptive time horizon set by a kinetic-energy stopping criterion rather than a preset depth. On synthetic isotropic and anisotropic Student-t target distributions, Neal's funnel distribution, and the real-world high-dimensional Fama-French 25 portfolio dataset, CVaR-GPA dramatically improves global and tail accuracy on heavy-tailed targets over the pre-trained baseline.

Data-Efficient Anomaly Diagnosis via Koopman Operators
Sanchita Ghosh, Texas Tech University

Modern energy system infrastructures integrate smart connectivity and automation technologies to provide better value of services while paving the way for safer and sustainable urban society. However, extreme and rare anomalous events of physical faults and/or cyberattacks gravely undermine the reliability and safety of such systems. While the problem of isolating the anomaly source is generally based on system models and measurement is generally ill-posed, it is crucial for dispatching appropriate remedial measures to mitigate safety-critical threats. To address these challenges, my research proposes two Koopman operator-based diagnostic algorithms for large-format battery packs. The first algorithm utilizes the Kullback-Leibler divergence-based distributional deviations in Koopman mode statistics to detect short-circuit as well as to identify the shorted module/s in a battery pack. The second algorithm leverages the characteristics of Koopman eigenfunctions and modes to isolate between actuation and sensor cyberattacks. Both algorithms adopt an online small-data learning strategy and do not rely on sensor redundancy or restrictive system assumptions. The efficacy of the algorithms is validated with high-fidelity simulations under rate-limited measurements and realistic uncertainties.

Gibbs conditioning and atypical consensus on random regular graphs

I-Hsun Chen, Brown University

Consider a social network in which each individual holds a binary opinion. The ”consensus” of this system quantifies the overall agreement between neighboring individuals. When the underlying network is regular, we show that, conditioned on an atypical consensus, the joint distribution of the opinions behaves (asymptotically) like a certain Ising measure on random regular graphs. We also identify a phase transition as the consensus exceeds a critical threshold. In a more general setting, we surprisingly reveal an abnormal degeneracy and provide criteria that rule out such behavior. The proofs leverage a recently obtained tractable form of the large deviation rate function for component empirical measures of random regular graphs with i.i.d. marks, and entail the analysis of associated non-convex constrained optimization problems. This is joint work with Ivan Lee, Kavita Ramanan and Sarath Yasodharan.

Quantification of the dynamical subspace of extreme events in complex systems through targeted nudged ensembles

Isabella Thiel, MIT

Extreme events in chaotic flows are often associated with transient finite-time instabilities, yet these instability directions are hard to compute and analyze in high-dimensional systems because they require tangent or adjoint dynamics, which are often hard to implement or too expensive. Here, the starting point is a trajectory that contains extreme events, either measured or generated by a perfect model, combined with an imperfect yet qualitatively informative model that is used for the characterization of the local dynamics. Using these ingredients, we introduce a non-intrusive framework in which a \emph{nudged stochastic ensemble} of the imperfect model "around" the available trajectory of the perfect model acts as a Jacobian-free estimator of local finite-time deformation, allowing for the characterization of the dynamics related to extreme events. We prove that in the small-noise, finite-time regime, the centered ensemble covariance satisfies a perturbed Lyapunov equation whose leading-order terms are generated by the same deformation kernels as the Cauchy-Green tensor of the \emph{nudged flow}. Under a spectral-gap condition, the leading empirical covariance eigenspace therefore approximates the dominant finite-time instability subspace with finite-ensemble sampling error.
We exploit this result to generate extreme event statistics using a few ensemble summaries, by conditioning a correction operator for a coarsely-resolved topographic two-layer quasi-geostrophic flow. The conditioned model improves tail-sensitive, i.e. extreme-event dependent, and spatial-exceedance diagnostics in data-limited regimes, including cases where the training horizon is orders of magnitude shorter than the return times associated with extreme events probed in testing. The results show that local finite-time instability information can be extracted non-intrusively and used to improve long-horizon emulation of extreme-event statistics.

Extreme Event Aware Learning and Optimal Observable Discovery

Kai Chang, MIT

Quantifying and predicting rare and extreme events is challenging because such events are infrequent, severe, and expensive to simulate. Existing data-driven methods often require multiple extremes in the training data or sampling process, leading to accurate predictions in quiescent regimes but high epistemic uncertainty in extreme-event regions. To overcome this limitation, we introduce Extreme Event Aware (η-) Learning, which does not require extreme events in the available data. The method reduces uncertainty even in uncharted extreme regimes by enforcing during training the statistics of an observable indicative of extremeness, obtained from qualitative knowledge or unlabeled data. This statistical regularization results in models that fit observed data while remaining consistent with prescribed observable statistics, enabling the generation of unprecedented extreme events. In the second part, we discuss a recently developed bilevel variational principle for discovering tail-optimal observable when extreme events are absent in the training data. We demonstrate both methods on prototype and real-world numerical examples.

Why DDIM Hallucinates More Than DDPM: A Theoretical Analysis of Reverse Dynamics

Grigorios Chrysos, University of Wisconsin-Madison

We theoretically study the hallucination phenomena in two canonical diffusion samplers: the stochastic Denoising Diffusion Probabilistic Model (DDPM) and the deterministic Denoising Diffusion Implicit Model (DDIM). We analyze the reverse ODE (DDIM) and SDE (DDPM) for a Gaussian mixture target, proving that after a critical time τ, (a) DDIM can become stuck on the segment connecting the two nearest modes and (b) DDPM *stochasticity* helps it become unstuck from this region, thus avoiding hallucination. Our empirical validation verifies that DDPM has a significantly lower hallucination rate than DDIM when this region is entered. Building on our observations, we exhibit how using additional stochastic steps can help DDIM avoid hallucinations and offer new insights on how to design improved samplers.

Seasonal trend assessment of US extreme precipitation via changepoint segmentation

Jaechoul Lee, Northern Arizona University

This study examines seasonal trends in US seasonal maximum precipitation, recognizing that long-term trends vary by season and that seasonal extremes often exhibit nonstationarities such as periodicities and changepoints. Ignoring these features can bias trend estimates and obscure real climatic patterns. To improve accuracy, we identify changepoints in seasonal maximum precipitation arising from gauge changes, observer turnover, or station relocations. Multiple changepoints are estimated using a penalized likelihood framework with a generalized extreme value model that incorporates periodic behavior and location shifts in the data. A genetic algorithm search efficiently explores the large space of possible changepoint numbers and timings. We also compute seasonal return levels for extreme precipitation to characterize seasonal changes in risk. The approach is illustrated using two stations, and national-scale results are summarized through spatial maps. Accounting for changepoints reveals more variable, spatially coherent, and season-specific trends than studies that ignore them, with increasing trends more prevalent in fall across the South and along the East Coast.

Multi-component matching queues with abandonment and buffers: Fluid, diffusion limit and large deviations

Souvik Ray, Iowa State University

We consider a regulated matching queue system in an assembly-to-order (ATO) production line. Here each product is assembled from several distinct perishable components which arrive randomly over time at the ATO station and wait in their respective class-specific queues until they are either matched or their patience expires. A match occurs instantaneously whenever one component from each category is available, and the matched set of ingredients exits the system immediately. Due to resource constraints, each queue is subject to a (possibly infinite) buffer capacity constraint. We prove the law of large numbers result for the fluid scaling system, where the limiting dynamics is characterized as the unique solution to a muti-dimensional Skorokhod problem whose reflection is oblique and solution map is non-Lipschitz. Sample path large deviation principles are established and exponentially decaying rates of probabilities for certain interesting rare events are obtained. We also derive the diffusion limit for the system in the heavy traffic regime.