Bayesian Inverse Problems and UQ
Poster Session Abstracts
March 3, 2026
4:00-5:00pm

Total HSIC: A New Index for Global Sensitivity Analysis
Troy Larsen, North Carolina State University

Global sensitivity analysis (GSA) is a fundamental component of uncertainty quantification used to identify which input parameters most influence a model's output. Traditional methods, such as variance-based Sobol' indices, are limited because they only capture influence that manifests through variance and assume statistical independence of inputs. This can lead to inaccurate sensitivity scores, particularly when dealing with complex, non-linear dependencies. We propose an alternative approach using the Hilbert-Schmidt independence criterion (HSIC). Unlike variance-based methods, HSIC is a moment-independent statistic that captures statistical dependence by embedding probability measures into reproducing kernel Hilbert spaces (RKHSs). Our findings demonstrate that this approach offers a compelling and flexible alternative to traditional GSA methods, capable of detecting a wider range of parameter influence.

Decentralized Proximal Stochastic Gradient Langevin Dynamics
Mohammad Rafiqul Islam, Florida State University

We propose Decentralized Proximal Stochastic Gradient Langevin Dynamics (\emph{DPSGLD}), a decentralized MCMC algorithm for sampling from a log-concave probability distribution constrained to a convex domain $\mk$. Constraints are enforced through a shared proximal regularization based on the Moreau–Yosida envelope, enabling unconstrained updates while preserving consistency with the target constrained posterior. We establish non-asymptotic convergence guarantees in the 2-Wasserstein distance for both individual agent iterates and their network averages. Our analysis shows that DPSGLD converges to a regularized Gibbs distribution and quantifies the bias introduced by the proximal approximation.

Computational Methods for Hyperparameter Estimation with applications to Separable Nonlinear Inverse Problems
Elle Buser, Emory University

Good estimates of hyperparameters can be critical for solving inverse problems accurately, but estimating these hyperparameters often requires solving a challenging and expensive optimization problem involving log determinants. In this work, we consider two optimization approaches, both of which involve stochastic average approximations (SAA) of the objective function. The first approach uses a Monte Carlo approach to approximate the objective function and gradient. The second approach is a majorization-minimization approach that replaces a challenging optimization problem with a sequence of simpler ones. Specifically, we use a majorization function for the log-determinant term and combine it with a randomized trace estimator to form a majorant. Both of these methods can be used within a general framework for hyperparameter estimation that includes both estimating standard hyperparameters (e.g., noise and prior variances) and estimating model parameters (e.g., unknown parameters that define the forward process). That is, we consider separable nonlinear inverse problems, where the forward model assumes a parametric form, and the goal is to improve the model parameters to more accurately solve the resulting linear inverse problem. By treating the model parameters as hyperparameters and using the proposed approaches for hyperparameter estimation, we provide an efficient approach to estimate model parameters in separable nonlinear inverse problems. A variety of numerical examples are provided from seismic tomography, image deblurring, super-resolution imaging, and velocity estimation from Navier-Stokes flow.

Transport vs MCMC for Conditional Sampling

Alex Johnson, University of Washington

A common task in uncertainty quantification is modeling a conditional distribution for parameters of interest given data. To compute descriptive statistics, one often requires independent samples from this distribution. The standard tools are Markov Chain Monte Carlo (MCMC) methods, which provide strong theoretical guarantees but often suffer from slow convergence and require re-initialization for each new conditioning variable. Measure transport methods, by contrast, construct deterministic or stochastic couplings between reference and target measures, offering amortized sampling and scalability. In this talk, we present a systematic comparison of measure transport and MCMC across a suite of benchmark UQ problems. Specifically, we construct triangular maps via flow and diffusion-based transport methods to demonstrate the efficacy of such a conditional sampling approach. Through our experiments, we highlight measure transport as a practical and powerful alternative to MCMC.

Error Analysis of Triangular Optimal Transport Maps for Filtering

Michele Martino, University of Washington

We present a systematic analysis of estimation errors for a class of optimal transport based algorithms for filtering and data assimilation. Along the way, we extend previous error analyses of Brenier maps to the case of conditional Brenier maps that arise in the context of simulation based inference. We then apply these results in a filtering scenario to analyze the optimal transport filtering (OTF) algorithm. An extension of that algorithm along with numerical benchmarks on various non-Gaussian and high-dimensional examples are provided to demonstrate its effectiveness and practical potential.

Latent Space Inference via Paired Autoencoders

Emma Hart, Emory University

This work describes a novel data-driven latent space inference framework built on paired autoencoders to handle observational inconsistencies when solving inverse problems. Our approach uses two autoencoders, one for the parameter space and one for the observation space, connected by learned mappings between the autoencoders' latent spaces. These mappings enable a surrogate for regularized inversion and optimization in low-dimensional, informative latent spaces. Our flexible framework can work with partial, noisy, or out-of-distribution data, all while maintaining consistency with the underlying physical models. The paired autoencoders enable reconstruction of corrupted data, and then use the reconstructed data for parameter estimation, which produces more accurate reconstructions compared to paired autoencoders alone and end-to-end encoder-decoders of the same architecture, especially in scenarios with data inconsistencies. We demonstrate our approaches on two imaging examples in medical tomography and geophysical seismic-waveform inversion, but the described approaches are broadly applicable to a variety of inverse problems in scientific and engineering applications.

Bayesian experimental design in the presence of scientific uncertainty

Sabina Sloman, University of Manchester

Bayesian (optimal) experimental design (BOED) is a principled framework for the design of experiments to estimate the latent parameter(s) of a statistical model. In the context of BOED, the optimal design maximizes a measure of the "expected information gained" about the parameter(s) of interest. In many cases, the scientific knowledge that informs the statistical model is partial or incomplete, and aspects of the model may be misspecified or express substantial uncertainty. This poster summarizes work on the problems misspecification can introduce in the context of BOED, focusing on two distinct forms of misspecification: (1) misspecification of nuisance parameters (factors influencing a target phenomenon that aren't of scientific interest), and (2) model misspecification (that arises due to a deliberate simplification of or incorrect belief about the data-generating process). In both cases, achieving robustness requires reconceptualizing the experimental design problem as the optimization of multiple objectives (as opposed to of the single objective of "information gained").

Bayesian Inference for Latent Gaussian Models Governed by PDEs

Sonia Reilly, Courant Institute, NYU

Latent Gaussian Models describe a wide range of problems in which an observed variable depends on an unknown parameter with a Gaussian prior and on some additional hyperparameters. Bayesian inference methods have been developed that marginalize out the dependence on the hyperparameters in a fast and accurate way when the dependence of the data on the parameter is governed by a simple operator. PDE governed inverse problems pose additional challenges, since existing methods do not account for problems for which evaluating the forward model is the most computationally expensive step. In this poster we will present a method for marginalizing hyperparameters efficiently for linear PDE governed problems using low-rank approximation of the prior-to-posterior precision update.

Sparse Dictionary-based Solution of Dynamic Inverse Problems

Aidan Mason-Mackay, University of Eastern Finland

In ill-posed dynamic inverse problems expected spatial features and temporal correlation between frames can be leveraged to improve the quality of the computed solution, in particular when the available data are limited and the dimensionality of the unknown is large. One way to take advantage of the spatial and temporal traits believed to characterize the solution is to encode them into the entries of a dictionary, and to seek the solution as a sparse linear combination of the dictionary atoms. To promote a vector of coefficients with mostly vanishing entries, we consider a stochastic extension of the dictionary coding problem model with a random hierarchical sparsity promoting prior. We compute the Maximum A Posteriori (MAP) estimate of the coefficient vector using the Iterative Alternating Sequential Algorithm (IAS), which has been demonstrated to efficiently solve inverse problems with minimal need for parameter tuning. The proposed methodology is tested on real-world dynamic Computed Tomography and MRI datasets, where it is compared to the popular Alternating Direction Method of Minimizers (ADMM). The computed examples show the that proposed methodology is competitive with the ADMM for compressed sensing, with a significantly lower sensitivity to hyper-parameter selection.

Zero-Shot Function Encoder-Based Differentiable Predictive Control

Hassan Iqbal, The University of Texas at Austin

We introduce a differentiable framework for zero-shot adaptive control over parametric families of nonlinear dynamical systems. Our approach integrates a function encoder-based neural ODE (FE-NODE) for system identification with differentiable predictive control (DPC) for offline self-supervised learning of explicit control policies. The FE-NODE captures nonlinear behaviors in state transitions and enables zero-shot adaptation to new systems without retraining. We formulate the online adaptation as a linear inverse problem, efficiently solving for latent basis coefficients via regularized least squares using limited system measurements. DPC efficiently learns control policies across system parameterizations, eliminating costly online optimization common in classical model predictive control. We demonstrate the efficiency, accuracy, and online adaptability of the proposed method across a range of nonlinear systems with varying parametric scenarios, highlighting its potential as a general-purpose tool for fast zero-shot adaptive control.

Efficient Dynamic Image Reconstruction with Motion Estimation

Toluwani Okunola, Tufts University

Dynamic inverse problems are challenging to solve due to the need to identify and incorporate appropriate regularization in both space and time. Moreover, the very large scale nature of such problems in practice presents an enormous computational challenge. In this work, in addition to the use of edge-enhancing regularization of spatial features, we propose a new regularization method that incorporates a temporal model that estimates the motion of objects in time. In particular, we consider the optical flow model that simultaneously estimates the motion and provides an approximation for the desired image, and we incorporate this information into the cost functional as an additional form of temporal regularization. We propose a computationally efficient algorithm to solve the jointly regularized problem that leverages a generalized Krylov subspace method. We illustrate the effectiveness of the prescribed approach on a wide range of numerical experiments, including limited angle and single-shot computerized tomography.

Priorconditioned Sparsity-Promoting Projection Methods for Deterministic and Bayesian Linear Inverse Problems

Jonathan Lindbloom, Dartmouth College

High-quality reconstructions of signals and images with sharp edges are needed in a wide range of applications. To overcome the large dimensionality of the parameter space and the complexity of the regularization functional, sparsity-promoting techniques for both deterministic and hierarchical Bayesian regularization rely on solving a sequence of high-dimensional iteratively reweighted least squares (IRLS) problems on a lower-dimensional subspace. Generalized Krylov subspace (GKS) methods are a particularly potent class of hybrid Krylov schemes that efficiently solve sequences of IRLS problems by projecting large-scale problems into a relatively small subspace and successively enlarging it. We refer to methods that promote sparsity and use GKS as S-GKS. A disadvantage of S-GKS methods is their slow convergence. In this work, we propose techniques that improve the convergence of S-GKS methods by combining them with priorconditioning, which we refer to as PS-GKS. Specifically, integrating the PS-GKS method into the IAS algorithm allows us to automatically select the shape/rate parameter of the involved generalized gamma hyper-prior, which is often fine-tuned otherwise. Furthermore, we proposed and investigated variations of the proposed PS-GKS method, including restarting and recycling (resPS-GKS and recPS-GKS). These respectively leverage restarted and recycled subspaces to overcome situations when memory limitations of storing the basis vectors are a concern. We provide a thorough theoretical analysis showing the benefits of priorconditioning for sparsity-promoting inverse problems. Numerical experiment are used to illustrate that the proposed PS-GKS method and its variants are competitive with or outperform other existing hybrid Krylov methods.

Estimating High-Dimensional Covariance Matrices with Hierarchical Rank Structure

Robin Armstrong, Cornell University

Many algorithms in data assimilation and model order reduction rely on sample-based estimates for a covariance matrix associated with the trajectory of a high-dimensional dynamical system. Due to computational constraints, the number of available samples is often far less than the dimension of the underlying state space. Under these circumstances, extracting meaningful covariance information requires that the noisy statistics of the sample are regularized with a structural assumption. We put forth a regularization technique that exploits the rank structure of submatrices representing cross-covariances between well-separated domains of space. We establish that low-rank truncations of these submatrices can be estimated from fewer samples than the submatrices themselves. We then show how this fact can be used to encode physics-informed regularizing assumptions onto the sample statistics, resulting in a hierarchically rank-structured covariance estimator. Through numerical experiments with a variety of dynamical systems, we demonstrate that these techniques are effective at reducing sampling errors in the covariance, with especially strong performance gains over traditional methods in the presence of multiscale behavior.

Enforcing detailed balance in inverse uncertainty quantification of biochemical reaction networks

Federica Milinanni, ICERM

Subcellular signaling pathways regulating, e.g., cell division or synapse plasticity, are often modeled as dynamic networks of reactions satisfying the law of mass-action. To remain physically grounded, these models must satisfy the Wegscheider conditions, a form of thermodynamic constraint requiring that the product of reaction rate constants around each cycle of reactions is equal in both directions of the cycle. These constraints lead directly to the law of detailed balance at equilibrium. To impose Wegscheider conditions, we need to identify all cycles in the reaction network, which can be computationally difficult, especially in large, complex systems. Furthermore, even when these constraints are explicitly available, they often cause inverse uncertainty quantification methods like MCMC to struggle. Instead of imposing Wegscheider conditions directly in the parameter space, our approach enforces physical consistency by penalizing parameters that violate detailed balance. By integrating this requirement into the likelihood, we promote inferred dynamics that are thermodynamically consistent, even when limited measurements lead to significant model unidentifiability.

A Scalable Sequential Framework for Dynamic Inverse Problems via Model Parameter Estimation

Aryeh Keating, Virginia Tech

Large-scale dynamic inverse problems, such as those arising in time-dependent imaging, are often severely ill-posed as a result of model complexity and the high dimensionality of the unknown parameters. Regularization is commonly employed to mitigate ill-posedness by incorporating prior information and structural constraints. However, classical regularization formulations are frequently infeasible in this setting due to prohibitive memory requirements, necessitating sequential methods that process data and state information online, eliminating the need to form the full space–time problem. In this work, we propose a memory-efficient framework for reconstructing dynamic sequences of undersampled images from computerized tomography data that requires minimal hyperparameter tuning. The approach is based on a prior-informed, dimension-reduced Kalman filter with smoothing. While well suited for dynamic image reconstruction, practical deployment is challenging when the state transition model and covariance parameters must be initialized without prior knowledge and estimated in a single pass. To address these limitations, we integrate regularized motion models with expectation–maximization strategies for the estimation of state transition dynamics and error covariances within the Kalman filtering framework. We demonstrate the effectiveness of the proposed method through numerical experiments on limited-angle and single-shot computerized tomography problems, highlighting improvements in reconstruction accuracy, memory efficiency, and computational cost.

Bayesian Inverse Problems with Sparse Priors and Randomization

Diego Arenas Mata, Virginia Tech

Sparse priors, such as the Laplace prior, are of particular interest in Bayesian inverse problems because they promote sparsity and preserve edges in the solution, which are often more appropriate than the smooth reconstructions obtained with Gaussian priors. Sampling from the resulting non-Gaussian posteriors can be challenging, however, particularly in high-dimensional settings. To address this, we build on the randomize-then-optimize (RTO) framework and its extension to l1 priors via variable transformation. We propose a method for posterior sampling in Bayesian inverse problems with Laplace priors that converts the prior into a standard Gaussian in the transformed space. The resulting nonlinear optimization problem is solved with the Levenberg–Marquardt algorithm, using a randomized preconditioner for the linearized system at each iteration. Performance is evaluated through numerical experiments on image deconvolution and computed tomography problems.

Towards Convergence Guarantees for Particle Swarm Optimization

Ritvik Teegavarapu, California Institute of Technology

Interacting particle systems have become a powerful tool for tackling global optimization problems, leveraging ensembles of particles governed by coupled stochastic differential equations. Among these methods, particle swarm optimization (PSO) is a widely used algorithm featured in many state-of-the-art optimization toolkits. However, the standard PSO dynamics are not readily suited for mean-field analysis due to their complex interaction structure. By modifying certain elements of the PSO formulation, one can recover as a consensus-based optimization (CBO) algorithm with memory, which admits a tractable mean-field limit and facilitates rigorous convergence analysis. In this paper, we develop some preliminary results on how the PSO parameter coupling effects existing convergence guarantees for CBO and its memory variant.

Bayesian Inference on SPD Manifolds: Geometry-Aware Learning of Posterior Covariances for PDE Inverse Problems

Pegah Amiri, University of Houston

We present work on Bayesian inference. We train a neural networks to learn an approximation of the covariance matrix of the posterior distribution. These matrices are symmetric positive definite (SPD) and live on a non-Euclidean Riemannian manifold. We consider a log-Euclidean framework that uses Riemannian geometry to map SPD matrices from their original manifold to tangent spaces through a logarithmic mapping. This approach allows for effective neural network training in tangent spaces while keeping geometric properties intact. We study the performance of NN predictions on noisy observations of varying complexity.

Kronecker Approximations of Covariance Matrices for Solving Inverse Problems

Srijon Sarkar, Emory University

Inverse problems arise in many applications where the task is to recover unknown parameters or data from indirect or noisy observations. In contrast to forward problems, where outputs are directly computed from known inputs, inverse problems are designed to infer the underlying variables for the observed effect. In many settings, large prior covariance matrices pose computational bottlenecks for solving inverse problems. To address this, we exploit Kronecker product approximations and demonstrate how these approximations can be used to solve inverse problems.