ABS26

The Applied Bayesian Statistics summer school

INTERPRETABLE BAYESIAN LEARNING FOR PHYSICAL AND ENGINEERING SCIENCES

Lecturer: Simon Mak
Assistant lecturer: Yen-Chun Liu

Villa del Grumello, Como, Italy, 6-10 July 2026

Abstracts

Participants' presentations

Organised by CNR-IMATI Milano in cooperation with Fondazione "Alessandro Volta", under the patronage of SIMAI - Società Italiana di Matematica Applicata e Industriale.

CONTENTS

Bayesian Optimization in Inverse Problems

Mike Baumann [(1)], *Simone Göttlich* [(1)], *and Simon Weissmann* [(1)]

(1) University of Mannheim, School of Business Informatics and Mathematics, Germany

Abstract: A central challenge in inverse problems is the principled selection of regularization parameters. In this talk, we address this problem through a data-driven bilevel optimization framework, in which the regularization parameters are learned directly from training data using Bayesian Optimization.
The lower-level problem is a variational reconstruction model, formulated as a convex and possibly non-smooth optimization problem depending on the observed data sample. The upper-level problem measures the quality of the reconstruction and defines the learning objective for the regularization parameters. Together, these two levels give rise to a stochastic, non-smooth, and generally non-convex optimization problem. This structure makes standard gradient-based methods difficult to apply, particularly because the lower-level solution depends implicitly on the parameters and may not be available in closed form.
To address these challenges, we model the upper-level loss as a black-box function using Gaussian process surrogate models and optimize it using Bayesian optimization techniques. A particular focus of our analysis is the practical setting in which the lower-level problems are solved only approximately, assuming the accuracy levels can be controlled. We investigate how such inexact lower-level solutions affect the performance and cost of the applied approach and suggest a bias-aware algorithm, trading off performance and cost.

Bayesian Hierarchical Modeling for Reliable Large-Scale Sensors Deployments and Applications

Anna Ballario [(1)]

(1) Politecnico di Torino, Dipartimento di Scienze Matematiche, Italy

Abstract: Reliability of large populations of sensors is a major challenge in modern industrial production and applied functional monitoring systems. The massive deployment of low-cost Micro-Electro-Mechanical Systems (MEMS) sensors across several technological domains requires calibration strategies that ensure metrological reliability while remaining feasible at industrial scale. However, traditional laboratory calibration procedures become impractical when sensors are produced in very large quantities, motivating the development of statistical approaches for large-scale calibration.

This work develops a Bayesian statistical framework for the virtual calibration and actual applications of large sensor batches. An earlier published model exploits information from a reference, laboratory-calibrated batch to infer the calibration properties of new lots of the same production process. By calibrating only a small subset of sensors from each batch, the model estimates key parameters that characterize the entire lot, providing robust reliability assessments while drastically reducing direct calibration effort.

Building upon this formulation, a hierarchical extension of the model is developed by introducing a Beta hyperprior distribution on the probability of detecting out-of-tolerance sensors, enabling the integration of prior industrial knowledge with explicit control of parameter variability. This approach softens prior deterministic assumptions about batch quality and enhances the flexibility and robustness of reliability estimates. The choice of hyperprior parameters therefore plays a crucial role in ensuring that the model coherently reflects prior industrial knowledge.

The study evaluates sensor batch reliability metrics under varying model parameters and introduces alternative metrics to address identified limitations. Model validation was conducted on a case study of 100 digital MEMS accelerometers calibrated at INRiM. Models based on weakly informative hyperpriors rely more strongly on the information conveyed by the likelihood. Consequently, when the observed data indicate a smaller percentage of defective sensors than that expected according to the prior, higher batch reliability and lower uncertainty associated with the statistically calibrated sensors are obtained. Conversely, when the data suggest a higher percentage of defective sensors, the same models produce a substantial increase in the associated uncertainty. This result highlights how the agreement between the prior assumption and the data-driven evidence plays a fundamental role in the assessment of the batch reliability and uncertainty.

Future research will extend the model by including utility and cost functions to formally represent the producer's risk in rejecting a batch and the consumer's risk in accepting it, integrating probabilistic calibration results with decision-making criteria. These studies will be crucial for balancing the benefits of using less informative hyperpriors, which may lead to higher estimated reliability of lots, against the drawback of a larger posterior variance, that is, a reduced confidence in the estimation of the true number of defective sensors.

Surrogate Multi-Fidelity Modeling for Financial Options Pricing Using Gaussian Processes

Mirko Piazzalunga [(1)] *and Tommaso Moretti* [(1)]

(1) Università degli Studi di Milano-Bicocca, Dipartimento di Statistica e Metodi Quantitativi, Italy

Abstract: Traditional financial option pricing faces a critical trade-off between computational efficiency and analytical accuracy. While the classic Black-Scholes model provides instantaneous, closed-form solutions, its restrictive assumptions, such as constant volatility, fail to capture the dynamics of volatile, real-world financial markets, leading to systematic pricing errors. In contrast, advanced high-fidelity approaches, such as Monte Carlo pricing under the Heston stochastic volatility model, offer superior precision but are computationally expensive, limiting their applicability in rapid pricing and risk analysis.

This research proposes a probabilistic multi-fidelity surrogate framework based on Gaussian Processes to bridge this gap. The objective is to combine inexpensive low-fidelity pricing evaluations with a limited number of computationally intensive high-fidelity simulations, thereby approximating high-fidelity option prices at significantly reduced computational cost. The framework is designed to learn and correct the structural bias of simplified pricing models while preserving predictive accuracy.

Before financial application, the framework is validated on the Branin benchmark function, a standard testbed for surrogate modeling to test the algorithm's convergence, its generic ability to correct an induced bias, and its reliability for broader surrogate modeling applications beyond finance. Once verified, the architecture is transferred to the financial domain through a controlled simulation protocol. A dense low-fidelity dataset is generated and evaluated through the Black-Scholes formula to capture the global curvature of the pricing function. Concurrently, a sparse dataset of strategic points is evaluated using a computationally heavy Monte Carlo simulations model to serve as the high-fidelity benchmark. These datasets are fused using a Co-Kriging architecture, enabling the surrogate to model both the global pricing structure and the residual discrepancy between fidelity levels.

The expected outcome is a surrogate model capable of achieving near high-fidelity pricing accuracy with substantially lower computational cost, while also quantifying epistemic uncertainty through Gaussian Process predictive variance. Future extensions will investigate more sophisticated low-fidelity baselines and validate the framework against real market option prices.

Sampling-based Batch Sequential Design by Stein Variational Gradient Descent

Xiaoxian Ding [(1)], *Penghui Fu* [(1)], *and C. F. Jeff Wu* [(1)]

(1) The Chinese University of Hong Kong, Shenzhen, School of Data Science, China

Abstract: Many real-world experimental design problems require a batch of experimental runs at each stage, where multiple points are selected and evaluated simultaneously. However, much of the computer experiments literature focuses on fully sequential, point-by-point methods. This work proposes a sampling-based framework for systematically converting a fully sequential design method into a batch sequential method. In particular, we use and adapt Stein variational gradient descent (SVGD) to efficiently sample a batch of points from a suitably constructed target distribution, balancing individual utility with batch diversity. We also address challenges that arise when applying SVGD to experimental design, including constrained design regions and pathologically uniform target distributions. The proposed framework is used to obtain batch versions of state-of-the-art fully sequential methods, and its performance is demonstrated through extensive numerical studies.

Data Efficient Optimization of Catalytic Reactions Using Bayesian Optimization

Ertugrul Furkan Düzenli [(1)]

(1) Max Planck Institut für Chemische Energiekonversion, Molecular Catalysis, Germany

Abstract: Bayesian optimization is increasingly used to accelerate chemical reaction optimization, but its performance relative to established design of experiments methods remains important to assess. In this study, sequential and parallel Bayesian optimization approaches were evaluated against a structured design of experiments reference for a constrained catalytic reaction system. Reaction performance was assessed using a combined multi response criterion that accounted for target product yield and selectivity. Across three case studies, the adaptive optimization approaches reached the reference performance level defined by the design of experiments reference with substantially fewer experiments. Sequential Bayesian optimization was most effective in reducing the total number of experiments, whereas the parallel approach reduced the number of optimization cycles by proposing multiple candidates before each model update. These results show that Bayesian optimization can provide a data efficient alternative to conventional experimental planning when reaction development is limited by experimental cost, time, or available experimental resources.

Bayesian machine learning for analog computation

Matteo Spaziani [1]

(1) Physics Department, La Sapienza University, Italy

Abstract: Analog computation has been slowly getting more and more attention for machine learning applications due to its higher energy efficiency and higher throughput. Its main disadvantage lies in the introduction of noise in the process of converting analog signals into digital ones and vice versa. The introduction of Bayesian inference can help machine learning models deal with these additional noise sources and achieve better accuracy when working in conditions of low signal-to-noise ratio.