Objective Bayes Methodology Conference 2025
June 8-12, 2025

Stavros Niarchos Foundation Cultural Center (SNFCC), Athens-Greece

Abstracts of Tutorial and Talk Presentations

National Technical University of Athens

ISBA INTERNATIONAL SOCIETY FOR BAYESIAN ANALYSIS

ΟΠΑ AUEB

Contents

Sunday, 8 June

Tutorial 1

Dealing with model uncertainty through Bayesian Model Averaging by **Mark Steel**, University of Warwick

An often overlooked but critical aspect of statistical practice is model uncertainty. Any statistical analysis can only be as good as the model you use. This is a relevant concern in all areas of application, but particularly in social sciences where there is typically no universally accepted theory or simple relationships or "laws". One important and potentially dangerous consequence of neglecting model uncertainty is that we assign more precision to our inference than is warranted by the data, and this leads to overly confident (and often misguided) decisions and predictions. In this tutorial we will consider Bayesian Model Averaging (BMA) as a natural and principled approach to dealing with model uncertainty. We will introduce the basic ideas and implementation details of BMA in the context of variable selection in the normal linear regression framework, a commonly used setting in many practical applications. We briefly sketch how to expand the context to deal with potentially endogenous covariates (unobserved confounding) and discuss two applications in economics: identifying drivers of economic growth and estimating the returns to education.

References:

Fernández, C., E. Ley, and M. Steel (2001). Benchmark priors for Bayesian model averaging, Journal of Econometrics, 100, 381–427.

Lenkoski, A., T. Eicher, and A. Raftery (2014). Two-stage Bayesian model averaging in endogenous variable models, Econometric Reviews, 33, 122–51.

Liang, F., R. Paulo, G. Molina, M. Clyde, and J. Berger (2008). Mixtures of g-priors for Bayesian variable selection, Journal of the American Statistical Association, 103, 410–23.

Steel, M.F.J. (2020). Model Averaging and its Use in Economics, Journal of Economic Literature, 58, 644-719. https://arxiv.org/abs/1709.08221.

Steiner, G. and Steel, M.F.J. (2025). Bayesian Model Averaging in Causal Instrumental Variable Models, preprint. https://arxiv.org/abs/2504.13520.

Tutorial 2

Introduction to e-values by **Rianne de Heide**, University of Twente

E-values are emerging in many hypothesis testing situations, and they turn out to have interesting properties. In this tutorial I will introduce them and make the connection to the objective Bayesian paradigm. I will end with some exciting recent results in multiple testing.

Tutorial 3

Causal Reasoning with Observational Data Based on Graphical Models by **Guido Consonni**, Università Cattolica del Sacro Cuore, Milano

Understanding causality is essential for answering questions that go beyond statistical association—questions about interventions, changes, and underlying mechanisms. This tutorial offers an accessible introduction to causal reasoning from observational data, tailored for researchers with no prior background in causal modeling. Grounded in the framework of graphical models, the session presents a broad overview of key concepts and tools that allow us to move from associations to causal conclusions. We will explore the distinction between statistical and causal models, the role of interventions, and how causal graphical models (such as Directed Acyclic Graphs) enable formal reasoning through the do-operator and graph surgery. Topics will include structural causal models, the ladder of causation, independent causal mechanisms, and sparse mechanism shift. We will also touch on methods for causal discovery from data—including constraint- and score-based approaches—along with the core principles of identification and estimation. Finally, the tutorial will briefly discuss the relevance of causality in machine learning, covering causal representation learning and transferability across tasks.

Tutorial 4

Bayesian nonparametric modelling of dependent distributions by **Jim Griffin**, University College London

Bayesian nonparametric methods are concerned with making inference about infinite dimensional objects, such as distributions. One area which has received a lot of attention is modelling related distributions, such as student scores on a national test in difference schools. This tutorial will review some of the different models that have been developed, such as additive, hierarchical and nested process and discuss how these relate to commonly used parametric approaches to dependent random variables, such as hierarchical and factor models.

Monday, 9 June

Session 1

Chair: Dimitris Fouskakis

Counterexamples that help define the objective Bayesian philosophy by **James Berger**, Duke University

Discussant: Luis R. Pericchi, University of Puerto Rico

Learning statistics from counter-examples: ancillary statistics was a famous article by Debabrata Basu (Basu, 2011). For a recent volume in Sankyha honoring Basu, I wrote an article (Berger, 2024) with essentially the same title (minus ancillarity). In this talk, I present those counterexamples that support and help to define objective Bayesian statistics.

FAB Prediction by **Peter Hoff**, Duke University

Discussant: Gunnar Taraldsen, Norwegian University of Science and Technology

A standard statistical inference task is to construct a prediction region, that is, a set of plausible values for an unobserved random object. The standard Bayesian approach to prediction via posterior predictive distributions permits complex modeling, incorporation of prior information and sharing of information across groups. However, the frequentist coverage of such regions may be poor, particularly in hierarchical models. To remedy this situation, we propose a class of prediction procedures that maintain exact frequentist coverage, while also taking into account prior or indirect information. Examples of such regions include the Bayes-optimal conformal prediction region, and prediction procedures in hierarchical models of continuous and categorical data.

Bayesian restricted likelihood and model diagnostics by **Steven MacEachern**, Ohio State University

Discussant: Radu Craiu, University of Toronto

Box advocated prior predictive checks as a natural way for Bayesians to assess a model. These checks perform well for well-specified models but perform poorly when the Bayesian model makes use of a weakly informative prior distribution. Partial posterior predictive checks (Bayarri and Berger) reduce the impact of the prior distribution while allowing us to create predictive summaries that are exchangeable with the observed data, under the assumption that the model is correct. The typical implementation of these checks involves a random split of the data into a training set for the partial update and a test set for the predictive summary. This talk explores targeted splits of the data that, along with the choice of predictive summary, allow the analyst to check various aspects of the model.

Session 2

Chair: Merlise Clyde

Scoring rule-based priors for regression models by **Isadora Antoniano-Villalobos**, Università Ca' Foscari Venezia

Discussant: Marco Ferreira, Virginia Tech

Typically, objective prior distributions for regression models separate the coefficients from the intercept and the regression variance. In other words, a standard objective criterion is used to define the prior for the variance and intercept parameters of the model, while, conditional on the variance, the regression coefficients are assigned a multivariate prior with a dependence structure derived from the design matrix. Recently, the multivariate Lomax distribution was proposed as a multivariate objective prior based on a scoring rule approach. We explore this proposal as a prior distribution for the full set of regression parameters. In particular, the characterization of the multivariate Lomax as an exponential mixture of Laplace densities results in robust estimators associated with the proposed Bayesian model.

Distance-Based Mixture Models for Prior Specification in Spatial Bayesian Analysis by **Clara Grazian**, University of Sidney

Discussant: Janet Van Niekerk, King Abdullah University of Science and Technology

In spatial statistical modeling, prior distributions play a critical role in capturing underlying spatial dependencies and guiding inference, particularly in regions with sparse data. This work introduces a novel framework that leverages distance-based approaches using mixture models to define default prior distributions tailored to spatial problems. By incorporating spatial distances into the mixture model structure, the proposed approach flexibly accommodates spatial heterogeneity without imposing strong parametric assumptions. The method allows for the clustering of spatial locations based on similarity in their response profiles, enabling the derivation of spatially adaptive priors. We illustrate the effectiveness of the framework through simulation studies and applications to real-world spatial data, demonstrating improved predictive accuracy and uncertainty quantification compared to conventional spatial priors.

Bayesian Smoothing and Feature Selection via Variational Automatic Relevance Determination by **Feng Liang**, University of Illinois

Discussant: Xenia Miscouridou, University of Cyprus

This study introduces Variational Automatic Relevance Determination (VARD), a novel approach for fitting sparse additive regression models in high-dimensional settings. VARD stands out by independently assessing the smoothness of each feature while precisely determining whether its contribution to the response is zero, linear, or nonlinear. Additionally, we present an efficient coordinate descent algorithm for implementing VARD. Empirical evaluations on both simulated and real-world datasets demonstrate VARD's superior performance compared to alternative variable selection methods for additive models.

Tuesday, 10 June

Session 3

Chair: Cristiano Villa

Pushing the limits of variable selection with external data by **David Rossell**, Universitat Pompeu Fabra, Barcelona

Discussant: Maria Kalli, King's College London

There are well-known mathematical limits under which consistent variable selection is possible, in terms of sample size n, number of variables p, signal strength, and dependence between covariates. In many applications, one also has external data that may (or may not) indicate the promise of each variable. Abundant literature shows that incorporating such external data empirically leads to better inference, but there is little theoretical understanding why. We cover this gap by showing that, using external data, one may consistently recover the true covariate set under more general conditions than otherwise possible, and at faster rates. We study the Gaussian sequence model and linear regression and discuss the connections to empirical Bayes.

Proper Bayes Minimax Multiple Shrinkage Estimation by **Ed George**, University of Pennsylvania

Discussant: Christopher Hans, Ohio State University

For the canonical problem of estimating a multivariate normal mean under squared error loss, minimax multiple shrinkage estimators adaptively shrink estimates towards multiple points and subspaces, thereby enhancing the scope of potential risk reduction while maintaining the protection guarantee of minimaxity. Motivated from a Bayesian point of view, the construction of such minimax estimators has relied on using mixtures of improper priors yielding superharmonic marginals. Indeed, up to now, even just the existence of proper Bayes minimax multiple shrinkage estimators has remained a challenging open problem, one that Bill and I struggled with for over 30 years. Happily, Bill ultimately came up with a novel unbiased-estimate-of-risk argument to demonstrate, for the first time, the existence of such estimators, including the existence of proper Bayes minimax multiple shrinkage estimators based on mixtures of the Strawderman-type priors which he pioneered in 1971. Not only are such multiple shrinkage estimators automatically admissible, but they also allow for the interpretation of their adaptive mixture weights as valid posterior probabilities. (This work is joint with Pankaj Bhagwat and Bill Strawderman).

Model Uncertainty via Predictive Resampling by **Holmes Chris**, Ellison Institute of Technology

Discussant: Nikos Demiris, Athens University of Economics and Business

When presented with a finite dataset and a collection of candidate models, an important challenge is quantifying uncertainty about the model that best represents the underlying population. Traditional approaches use priors over models and their parameters, or rank models according to a model selection criterion. In this talk we discuss an alternative perspective: treating the unobserved portion of the population as the source of model uncertainty. Under this view, the "best" model would be evident if the full population were observed, and uncertainty arises solely from missing information.

We operationalize this using a predictive resampling framework. Missing observations are sequentially generated from one-step-ahead predictive distributions, guided by a consistent model selection rule. This generates a completed dataset, which yields a best model. Repeating the entire procedure, provides a sample of best models, offering a direct estimate of model uncertainty. The approach is 'Bayesian' in providing a conditional distribution on candidate models, and 'objective' in avoiding subjective priors and marginalization over parameter spaces, circumventing challenges faced by approaches using Bayes factors. This is joint work with Vik Shirvaikar and Stephen Walker.

Session 4

Chair: Anabel Forte

Bayesian calculus and predictive characterizations of extended feature allocation models by **Federico Camerlenghi**, Università degli Studi di Milano-Bicocca

Discussant: Michele Guindani, University of California, Los Angeles

Since the introduction of the Indian buffet process, there has been a growing interest in Bayesian nonparametric priors for feature allocation models, where each observation is a random finite set of features. These priors can also be used at a latent level, such as in image segmentation and in latent factor models. In the present talk, we first provide a unified framework for extended feature allocation models, with the aim of modeling features' interaction, such as attractiveness and repulsiveness. We develop a full Bayesian analysis of these models, deriving closed-form expressions and presenting several illustrative examples. Second, we characterize specific priors within this general class based on the induced predictive rules. In particular, we identify priors where the probability of discovering new features depends (i) solely on the sample size or (ii) on both the sample size and the distinct number of observed features. These predictive characterizations, referred to as Johnson ``sufficientness'' postulates, are general, and they provide guidance on the prior's specification. We finally demonstrate the advantages of our modeling approach through an application in spatial statistics.

Loss-based priors for BART models by **Fabrizio Leisen**, King's College London

Discussant: Luciana Dalla Valle, Università di Torino

The purpose is to present a novel prior for tree topology within Bayesian additive regression trees (BART) models. This approach quantifies the hypothetical loss in information and the loss due to the complexity associated with choosing the wrong tree structure. The resulting prior distribution is compellingly geared toward sparsity, a critical feature considering BART models' tendency to overfit. The method incorporates prior knowledge into the distribution via two parameters that govern the tree's depth and balance between its left and right branches. Additionally, a default calibration is proposed for these parameters, offering an objective version of the prior. The method's efficacy is demonstrated on both simulated and real datasets.

Session 5

Chair: Ed George

Safe anytime-valid e-value Bayes factors for one-factorial ANCOVA -- Labeling invariance and growth-rate optimality by **Alexander Ly**, Centrum Wiskunde & Informatica, Amsterdam

Discussant: Fulvio De Santis, Sapienza Università di Roma

We examine the K-sample means problem while accounting for the effects of covariates. Our inference focuses on (1) testing the null hypothesis that all K samples share an identical but unknown mean and (2) providing simultaneous uncertainty quantification for the K-1 average treatment effects. We show how a specific class of Bayes factor naturally emerges in this problem, when the objective is to develop a methodology that ensures explicit exact frequentist control over type I error and coverage rates regardless of how, or even if, data collection is stopped. (*Joint work with Udo Boehm, Wouter Koolen, and Peter Grunwald)***.**

On Bayesian inference about the smoothness parameter in Gaussian Matern random fields by **Victor De Oliveira**, University of Texas

Discussant: Mario Peruggia, Ohio State University

The Matern family of covariance functions is currently the most commonly used for the analysis of geostatistical data due to its ability to describe different smoothness behaviors. Yet, in many applications, the smoothness parameter is set at an arbitrary value. This practice is due partly to computational challenges faced when attempting to estimate all covariance parameters and partly to unqualified claims in the literature stating that geostatistical data have little or no information about the smoothness parameter. In this talk, I propose a new class of easy-to-compute default priors for the smoothness parameter. These priors approximate reference priors, but their analysis and computation are considerably simpler. It is shown that the posterior distribution of the parameters based on these priors is proper, and Bayesian inferences about the covariance parameters based on these priors have satisfactory frequentist properties, much better than those based on maximum likelihood. The methodology is illustrated with two real data sets.

Wednesday, 11 June

Session 6

Chair: Guido Consonni

On variable selection with non-linear effects by **Rui Paulo**, Universidade de Lisboa

Discussant: Veronika Rockova, University of Chicago

Traditional variable selection assumes that the effect on the mean response of each of the p potential explanatory variables is linear. From a Bayesian model selection point of view, this can be addressed by comparing the 2^p models via posterior model probabilities and computing the ensuing posterior inclusion probabilities of each variable. We enlarge this model space by allowing the possibility that the effect of each variable is non-linear, which we model via a Gaussian process. This results in a model space containing 3^p competing models: each predictor is either inactive, has a linear effect, or has a non-linear effect. In this talk, we discuss various aspects of implementing such a strategy, including computation, prior specification for the parameters specifying the covariance function of the Gaussian processes, and on the model space. Results are promising; in particular, when the effect is indeed linear, we obtain marginal inclusion probabilities that are in close agreement with conventional variable selection; when the effect is non-linear, we typically see an increase in the marginal inclusion probability.

O'Bayes reIMagined: probabilistic approximations of possibilistic inferential models by **Ryan Martin**, North Carolina State University

Discussant: Angelos Alexopoulos, Athens University of Economics and Business

When prior information is lacking, the go-to strategy for probabilistic inference is to combine a "default prior" and the likelihood via Bayes's theorem. Objective Bayes, (generalized) fiducial inference, etc. fall under this umbrella. This construction is natural, but the corresponding posterior distributions generally only offer limited, approximately valid uncertainty quantification. The present paper takes a reimagined approach, yielding posterior distributions with stronger reliability properties. Specifically, the proposed construction starts with an inferential model (IM), one that takes the mathematical form of a data-driven possibility measure and features exactly valid uncertainty quantification, and then returns a so-called inner probabilistic approximation thereof. This inner probabilistic approximation inherits many of the original IM's desirable properties, including credible sets with exact coverage and asymptotic efficiency. The approximation also agrees with the familiar Bayes/fiducial solution obtained in applications where the model has a group transformation structure. A Monte Carlo method for evaluating the probabilistic approximation is presented, along with numerical illustrations.

Bayesian estimation in high-dimensional Hawkes processes by **Judith Rousseau**, Université Paris Dauphine

Discussant: Stefan Franssen, Ceremade, Dauphine PSL

Multivariate Hawkes processes form a class of point processes describing self and inter exciting/inhibiting processes. There is now a renewed interest of such processes in applied domains and in machine learning, but there exists only limited theory about inference in such models, in particular in high dimensions. To be more precise, the intensity function of a linear Hawkes process has the following form: for each dimension $k \leq K$

$$\lambda^k(t) = \sum_{\ell \leq K} \int_0^{t^-} h_{\ell k}(t-s)dN_s^\ell + \nu_k, \quad t \in [0,T]$$

where $(N_\ell, \ell \leq K)$ is the Hawkes process and $\nu_k > 0$. There have been some recent theoretical results on Bayesian estimation in the context of linear and nonlinear multivariate Hawkes processes, but these results assumed that the dimension K was fixed. Convergence rates were studied assuming that the observation window T goes to infinity. In this work we consider the case where K is allowed to go to infinity with T. We consider generic conditions to obtain posterior convergence rates and we derive, under sparsity assumptions, convergence rates in L_1 norm and consistent estimation of the graph of interactions.

Session 7

Chair: Mark Steel

Combining confounded and unconfounded in heterogeneous treatment effect modelling by **Ioanna Manolopoulou**, University College London

Discussant: Stavros Nikolakopoulos, University of Ioannina

Building statistical models using non-randomly sampled data is a well-known challenge in statistics, and is especially challenging when any part of the statistical model is not fully identifiable. In causal inference, and in particular in the estimation of heterogeneous treatment effects, this arises when observational data are used which may be affected by unobserved confounding. One approach to correct for such confounding is to combine data with and without unobserved confounding. However, when the unconfounded data are not representative of the whole population, the effect of de-confounding will be poor for subsets of the population that fall outside the range of these data. Depending on the structure of the model and the nature of the prior distributions used within a Bayesian model, this will be addressed by borrowing information from other parts of the space. In this work, we highlight the importance of building models that can account for uncertainty due to unobserved confounding in regions where no de-confounding is possible. To this end, we embed a combination of data with and without unobserved confounding into Bayesian Causal Forests (BCF), and make use of a data-dependent tempered likelihood to harness as much reliable information from the unconfounded data as possible, without leading to over-confidence in regions of poor identifiability. We implement our methods on a set of simulated and real data examples. (*Joint work with Ilina Yozova and Nathan McJames).*

Posterior asymptotics of high-dimensional spiked covariance model by **Jaeyong Lee**, Seoul National University

Discussant: Luis Carvalho, Boston University

We consider Bayesian inference on the spiked eigenstructures of high-dimensional covariance matrices; in other words, we are interested in estimating the eigenvalues and corresponding eigenvectors of high-dimensional covariance matrices in which a few eigenvalues are significantly larger than the rest. We impose an inverse-Wishart prior distribution on the unknown covariance matrix and derive the posterior distributions of the eigenvalues and eigenvectors by transforming the posterior distribution of the covariance matrix. Moreover, we justify the proposed method by demonstrating that the posterior distribution of the spiked eigenvalues and corresponding eigenvectors converges to the true parameters under the spiked high-dimensional covariance assumption. We prove that the posterior distribution of the spiked eigen-vector attains the minimax optimality under the single spiked covariance model. Simulation studies and real data analysis demonstrate that our proposed method outperforms all existing methods in quantifying uncertainty.

Thursday, 12 June

Session 8

Chair: Brunero Liseo

Fast inference for Bayesian nonparametrics via moments of functional means by **Ramses Mena**, Universidad Nacional Autónoma de México

Discussant: Natalia Bochkina, University of Edinburgh

Using the symmetry inherent to exchangeable sequences driven by certain classes of Bayesian nonparametric priors, a system of partial difference equations characterising moments of random means is derived and solved. At the outset, the results can be used to compute such quantities in a fast manner, for instance to elucidate Bayesian non-parametric models for a given data set or to attain posterior inferences in a fast manner. However, they can also be used for other inference problems. Some properties and practical examples will be explored.

Selection of latent variables based on grouped Bayes factors by **Garcia Gonzalo Donato**, University of Castilla-La Mancha

Discussant: Marilena Barbieri, Università di Roma Tre

It is very common in applied studies the consideration of variables that cannot be measured directly. Their integration in a study is approached by means of several observable variables or proxies conceived to capture the idiosyncrasy of the underlying conceptual magnitudes. For example, ruralness of a county is a conceptual social aspect which is not directly observable but which is clearly related with, for instance, rate of households with access to public water or rate of households using wood to heat the home. These variables that literature has named ``latent'' are very popular in disciplines like economics, psychology, epidemiology, etc. In this research we tackle the problem of selecting latent variables, proposed as potential regressors for a response of interest. We investigate on the notion of grouped Bayes factors; a weighted average of Bayes factors built with the proxies generating the latent variable. We discuss about the choice of weights and show that standard default choices may lead to unsatisfactory results, leading us to argue in favor of informed possibilities that take into account the correlation structure of the proxies. The resulting method is fully Bayesian and objective and can be implemented in problems with large number of predictors as we illustrate in real datasets.

Session 9

Chair: Ioannis Ntzoufras

Bayesian model criticism using uniform parametrization checks by **Jeffrey Miller**, Harvard University

Discussant: Yichen Zhu, Università Bocconi, Milano

Models are often misspecified in practice, making model criticism a key part of Bayesian analysis. It is important to detect not only when a model is wrong, but which aspects are wrong, and to do so in a computationally convenient and statistically rigorous way. We introduce a novel method for model criticism based on the fact that if the parameters are drawn from the prior, and a dataset is generated according to the assumed likelihood, then a sample from the posterior will be distributed according to the prior. Building upon this idea, we propose to reparameterize all random elements of the likelihood and prior in terms of independent uniform random variables. This makes it possible to aggregate across arbitrary subsets of data points and parameters to test for model departures using classical hypothesis tests for dependence or non-uniformity. We demonstrate empirically how this method of uniform parameterization checks (UPCs) facilitates model criticism, and we develop supporting theoretical results.

Bayesian Model Criticism: From Holdout Checks to Model Comparison by **Gemma Moran,** Rutgers University

Discussant: Leonardo Egidi, Università degli Studi di Trieste

In this talk, I will cover two recent works on Bayesian model criticism. The first is Holdout Predictive Checks (HPCs). HPCs are built on posterior predictive checks (PPCs), which check a model by assessing the posterior predictive distribution on the observed data. However, PPCs use the data twice - both to calculate the posterior predictive and to evaluate it - which can lead to uncalibrated p-values. HPCs, in contrast, compare the posterior predictive distribution to a draw from the population distribution, a held-out dataset. This method blends Bayesian modeling with frequentist assessment. Unlike the PPC, we prove that the HPC is properly calibrated. Empirically, we study HPCs on classical regression, a hierarchical model of text data, and factor analysis. In the second work, we introduce the posterior predictive null check (PPN), a method for Bayesian model criticism that helps characterize the relationships between models. The idea behind the PPN is to check whether data from one model's predictive distribution can pass a predictive check designed for another model. This form of criticism complements the classical predictive check by providing a comparative tool. A collection of PPNs, which we call a PPN study, can help us understand which models are equivalent and which models provide different perspectives on the data. With mixture models, we demonstrate how a PPN study, along with traditional predictive checks, can help select the number of components by the principle of parsimony. With probabilistic factor models, we demonstrate how a PPN study can help understand relationships between different classes of models, such as linear models and models based on neural networks. Finally, we analyze data from the literature on predictive checks to show how a PPN study can improve the practice of Bayesian model criticism.

The conference is sponsored by the International Society for Bayesian Analysis (ISBA), OB section of ISBA, National Technical University of Athens (NTUA) and Athens University of Economics and Business (AUEB).

Contact: obayes2025@aueb.gr