



II International Workshop on Bayesian Statistics

CUNEF Universidad, Madrid • June 4–5, 2026

Book of Abstracts

Contents

Workshop Overview	3
1 Program Overview	3
2 Detailed Program	5
3 Keynote Speakers	7
On smoothing risk patterns in areal data	7
A semi-parametric Bayesian GLM and applications to outcome dependent sampling	8
Objective Bayesian Model Uncertainty Quantification under Missing Data	9
Sparse Bayesian Factor Analysis for Gaussian and non-Gaussian Data	10
Advances in Scalable Gaussian Processes	11
Joint entropy search for multi-objective Bayesian optimization with constraints and multiple fidelities	12
High-dimensional model selection: transfer learning and projected computation	13
4 Contributed Talks	14
4.1 Session 1: Time Series & Financial Econometrics	14
Nonlinear Autoregressive Models for Functional Time Series with Bayesian Additive Regression Trees	14
State-dependent realized stochastic covariance model	14
Bayesian Extreme Value Theory with Hawkes-AR-Gumbel Dependence for Extreme CVaR Estimation in Operational Risk	15
Bayesian Markov-Switching Partial Reduced-Rank Regression.	16
4.2 Session 2: Model Selection	17
Shrinkage through multiple identifiability	17
Bayesian Model Checks and Improvements for Discrete Data	17
Objective Model Prior Probabilities in Variable Selection	18
4.3 Session 3: Applied Bayesian Methods	19
Information-based Bayesian Optimization with Expert Human Feedback	19
Applied Bayesian nonparametric modeling under likelihood ratio order constraints	20
Latent advertising stock and saturation for Marketing Mix Modeling (MMM): a field for Bayesian methods' expansion	21
4.4 Session 4: Adversarial Learning & Robust ML	22
A Unified Bayesian Framework for Adversarial Robustness	22
Bayesian Online Test Time Adaptation: A General Framework	22
Adversarial observations in Probabilistic State-Space Models for Robust Reinforcement Learning	23
Tracking Latent Goals for Robust Reinforcement Learning	23
A Bayesian computational framework for general security games	24
4.5 Session 5: Spatial & Temporal Modeling	25
Fair spatial predictions in urban analytics: a Bayesian approach	25
A Bayesian Spatiotemporal Varying-Coefficient Model for Multi-Patient Intracranial Recordings	25
A Bayesian spatio-temporal model for precipitation with two sources of zeros	26

Workshop Overview

The II International Workshop on Bayesian Statistics will take place on 4–5 June 2026 at CUNEF Universidad, Madrid. The workshop is planned as a single-track event with plenary talks, contributed sessions, coffee breaks, lunches, and a closing session. Registration is free thanks to the support of SEIO and CUNEF Universidad.

This draft book of abstracts differentiates between keynote and contributed talks. Keynote abstracts are included when available in the local workshop material; missing keynote titles or abstracts are marked as to be confirmed.

Program Overview

Thursday, 4 June 2026

8:00 – 9:00 Registration

9:00 – 9:15 Opening remarks

9:15 – 10:15 **Keynote 1** *Peter Müller*

10:15 – 11:15 **Keynote 2** *David Rossell*

11:15 – 11:45 *Coffee Break*

11:45 – 13:05 Contributed Session I: Time Series & Financial Econometrics

13:05 – 14:35 *Lunch*

14:35 – 15:35 Contributed Session II: Model Selection

15:35 – 16:35 **Keynote 3** *Stefano Cabras*

16:35 – 17:05 *Coffee Break*

17:05 – 18:05 Contributed Session III: Applied Bayesian Methods

18:05 – 19:05 **Keynote 4** *Sylvia Frühwirth-Schnatter*

20:00 Group Photo at the entrance of the ALMANSA campus

20:30 Conference dinner (at Sal Gorda, calle Beatriz de Bobadilla 9)

Friday, 5 June 2026

9:00 – 10:00 **Keynote 5** *Lola Ugarte*

10:00 – 11:40 Contributed Session IV: Adversarial Learning & Robust ML

11:40 – 12:15 *Coffee Break*

12:15 – 13:15 **Keynote 6** *José Miguel Hernández-Lobato*

13:15 – 14:45	<i>Lunch</i>
14:45 – 15:45	Contributed Session V: Spatial & Temporal Modeling
15:45 – 16:45	Keynote 7 <i>Daniel Hernández-Lobato</i>
16:45 – 17:00	Closing remarks and announcement of the María Eugenia Castellanos Award winner

Workshop Venue: C. de Leonardo Prieto Castro, 2, Moncloa – Aravaca, 28040 Madrid

Room: Aula F1.2 – LPC Campus

Presentation lengths:

- Keynote talk: 60 minutes (50 min presentation + 10 min Q&A)
- Contributed talk: 20 minutes (15 min presentation + 5 min Q&A)

Detailed Program

Thursday, 4 June 2026

8:00 – 9:00 Registration

9:00 – 9:15 Opening remarks

9:15 – 10:15 **Keynote 1** *Peter Müller: A semi-parametric Bayesian GLM and applications to outcome dependent sampling*

10:15 – 11:15 **Keynote 2** *David Rossell: High-dimensional model selection: transfer learning and projected computation*

11:15 – 11:45 *Coffee Break*

Contributed Session I: Time Series & Financial Econometrics 11:45 – 13:05

11:45 – 12:05 *Eoghan O’Neill: Nonlinear Autoregressive Models for Functional Time Series with Bayesian Additive Regression Trees*

12:05 – 12:25 *Igor F. B. Martins: State-dependent realized stochastic covariance model*

12:25 – 12:45 *María F. Pintado: Bayesian Markov-Switching Partial Reduced-Rank Regression*

12:45 – 13:05 *Juan Ballesteros-Gómez: Bayesian Extreme Value Theory with Hawkes-AR-Gumbel Dependence for Extreme CVaR Estimation in Operational Risk*

13:05 – 14:35 *Lunch*

Contributed Session II: Model Selection 14:35 – 15:35

14:35 – 14:55 *Lorenzo Capello: Bayesian Model Checks and Improvements for Discrete Data*

14:55 – 15:15 *Gonzalo García-Donato: Objective Model Prior Probabilities in Variable Selection*

15:15 – 15:35 *Carlos García Meixide: Shrinkage through multiple identifiability*

15:35 – 16:35 **Keynote 3** *Stefano Cabras: Objective Bayesian Model Uncertainty Quantification under Missing Data*

16:35 – 17:05 *Coffee Break*

Contributed Session III: Applied Bayesian Methods 17:05 – 18:05

17:05 – 17:25 *Yago Aguado-Carrillo-de-Albornoz: Information-based Bayesian Optimization with Expert Human Feedback*

17:25 – 17:45 *Adam Olivares: Applied Bayesian nonparametric modeling under likelihood ratio order constraints*

17:45 – 18:05 *Miguel Franco Pérez: Latent advertising stock and saturation for Marketing Mix Modeling: a field for Bayesian methods’ expansion*

18:05 – 19:05 **Keynote 4** *Sylvia Frühwirth-Schnatter: Sparse Bayesian Factor Analysis for Gaussian and non-Gaussian Data*

20:00 Group Photo in the ALMANSA campus

20:30 Conference dinner (at Sal Gorda, calle Beatriz de Bobadilla 9)

Friday, 5 June 2026

9:00 – 10:00	Keynote 5 <i>Lola Ugarte: On smoothing risk patterns in areal data</i>	
Contributed Session IV: Adversarial Learning & Robust ML		10:00 – 11:40
10:00 – 10:20	<i>Pablo G. Arce: A Unified Bayesian Framework for Adversarial Robustness</i>	
10:20 – 10:40	<i>Daniel Corrales: Bayesian Online Test Time Adaptation: A General Framework</i>	
10:40 – 11:00	<i>Miguel Santos: Adversarial observations in Probabilistic State-Space Models for Robust Reinforcement Learning</i>	
11:00 – 11:20	<i>Mario Chacón-Falcón: Tracking Latent Goals for Robust Reinforcement Learning</i>	
11:20 – 11:40	<i>Jose Manuel Camacho: A Bayesian computational framework for general security games</i>	
11:40 – 12:15	<i>Coffee Break</i>	
12:15 – 13:15	Keynote 6 <i>José Miguel Hernández-Lobato: Advances in Scalable Gaussian Processes</i>	
13:15 – 14:45	<i>Lunch</i>	
Contributed Session V: Spatial & Temporal Modeling		14:45 – 15:45
14:45 – 15:05	<i>Pepa Ramírez-Cobo: Fair spatial predictions in urban analytics: a Bayesian approach</i>	
15:05 – 15:25	<i>Dominik Wielath: A Bayesian Spatiotemporal Varying-Coefficient Model for Multi-Patient Intracranial Recordings</i>	
15:25 – 15:45	<i>Juan Marcen-Gutierrez: A Bayesian spatio-temporal model for precipitation with two sources of zeros</i>	
15:45 – 16:45	Keynote 7 <i>Daniel Hernández-Lobato: Joint entropy search for multi-objective Bayesian optimization with constraints and multiple fidelities</i>	
16:45 – 17:00	Closing remarks	

Keynote Speakers

On smoothing risk patterns in areal data

Authors and affiliations. María Dolores (Lola) Ugarte, Universidad Pública de Navarra.

Presenter/status. María Dolores (Lola) Ugarte, Full Professor.

Abstract.

In this talk, we will explore one of the most compelling applications of areal data: disease mapping. Following a brief historical note, we will introduce the most widely used univariate space-time models in this domain. Subsequently, we will delve into recently developed multivariate models, showcasing an interesting application. In addition, we will present a simple idea for analyzing large datasets, encompassing both univariate and multivariate models.

A semi-parametric Bayesian GLM and applications to outcome dependent sampling

Authors and affiliations. Peter Müller, Entejar Alam and Paul Rathouz, The University of Texas at Austin.

Presenter/status. Peter Müller, Full Professor.

Abstract.

We introduce an instance of a modified varying weight dependent Dirichlet process model to implement a semi-parametric generalized linear model. The model extends recently developed semi-parametric generalized linear models by adding a nonparametric Bayesian prior on the centering distribution of the generalized linear model. Building on familiar posterior simulation methods for mixtures with respect to normalized random measures, we introduce a modification to implement posterior simulation in the resulting semi-parametric model.

The motivating application is data analysis for a study with outcome dependent sampling, that is, when participants are sampled into a study based on an outcome variable, as well as some auxiliary covariates. We discuss how the outcome dependent sampling design can be accommodated in the proposed semiparametric generalized linear model with only minor modification of the inference scheme.

Objective Bayesian Model Uncertainty Quantification under Missing Data

Authors and affiliations. Stefano Cabras, Universidad Carlos III de Madrid.

Presenter/status. Stefano Cabras, Full Professor.

Abstract.

This research addresses the dual challenge of statistical model uncertainty and missing data within a unified objective Bayesian framework. Traditional methods, such as Rubin's multiple imputation rules, rely on an imputation model that is presumed known; however, in contexts of model uncertainty, this model is precisely the object of inference. This work proposes a probabilistic methodology for quantifying uncertainty in problems like variable selection and model averaging when part of the data is not available and imputation is needed to avoid data loss with the usual listwise deletion method. By focusing on marginal likelihoods and model posterior probabilities, it is possible to identify critical conditions for the ignorability of the missingness mechanism, namely Missing At Random (MAR) and prior separability.

A central contribution is the derivation of the g' -prior for linear regression, which generalizes Zellner's g -prior to accommodate the presence of incomplete covariates. Unlike standard approaches, this prior remains well-defined even in high-dimensional settings where the number of predictors exceeds the sample size. The proposed computational framework utilizes a Monte Carlo approximation to estimate the marginal likelihood, drawing samples from the posterior distribution of missing values and imputation parameters.

The efficacy of the g' -Bayes factor is demonstrated through five distinct experiments involving both synthetic and real-world datasets, such as the Ozone and Boston Housing sets. The empirical results consistently show that this approach preserves "oracle" evidence—the inference obtained from the full, complete dataset—more effectively than listwise deletion.

Finally, because of the presence of the g' parameters, such prior maybe sensitive to this and preliminary results based on Intrinsic Bayes Factors will be shown to illustrate further investigation on this topic.

Sparse Bayesian Factor Analysis for Gaussian and non-Gaussian Data

Authors and affiliations. Sylvia Frühwirth-Schnatter, WU Wien.

Presenter/status. Sylvia Frühwirth-Schnatter, Full Professor.

Abstract.

Factor analysis is a popular method to obtain a sparse representation of the covariance matrix of multivariate observations and to uncover the unobserved driving factors behind observed correlation. A challenge for factor models is to estimate the unknown number of factors and to recover an interpretable factor loading matrix from the data. Research in the area of sparse Bayesian factor analysis successfully addresses these issues within a Bayesian framework through the help of variable selection and shrinkage priors, see Frühwirth-Schnatter, Hosszejni and Lopes, 2025 for a recent review. However, most approaches to sparse Bayesian factor analysis rely on the assumption that the idiosyncratic errors are Gaussian.

After a review of methods for Gaussian sparse Bayesian factor analysis, this restrictive assumption is alleviated in the present talk by assuming that the idiosyncratic errors are non-Gaussian, following Azzalini's multivariate skew-normal or skew-t distribution. Such a non-Gaussian factor model has a stochastic representation as a Gaussian factor model with an additional factor following a truncated standard normal distribution and the skewness parameters acting as factor loadings. This representation is exploited to perform sparse Bayesian inference and learn the number of factors and the sparse loading matrix also for non-Gaussian factor models.

Advances in Scalable Gaussian Processes

Authors and affiliations. José Miguel Hernández-Lobato, University of Cambridge.

Presenter/status. José Miguel Hernández-Lobato, Full Professor.

Abstract.

Gaussian processes are a powerful tool for uncertainty quantification and decision-making, but their scalability is challenged by cubic computational costs. This talk explores key advancements in scalable Gaussian processes that leverage stochastic optimization through iterative solvers. First, I will introduce a simple and effective stochastic optimization algorithm, demonstrating its superiority over preconditioned conjugate gradients and variational approximations in regression and Bayesian optimization tasks. Second, I will present innovative strategies for hyperparameter optimization in large datasets, including pathwise gradient estimators, warm-starting solvers, and early stopping, achieving significant speed-ups and improved efficiency. Together, these contributions push the boundaries of Gaussian process scalability, enabling their use in large-scale and ill-conditioned settings.

Joint entropy search for multi-objective Bayesian optimization with constraints and multiple fidelities

Authors and affiliations. Daniel Hernández Lobato, Universidad Autónoma de Madrid.

Presenter/status. Daniel Hernández Lobato, Associate Professor.

Abstract.

Bayesian optimization (BO) methods can be used to solve efficiently problems with several objectives and constraints. Each objective and constraint is considered a black-box function that is expensive to evaluate, lacking a closed-form expression. BO methods use a model of each black-box to guide the search for the problem's solution. Specifically, they make intelligent decisions about where each black-box function should be evaluated next with the goal of finding the solution using a few evaluations only. Sometimes, however, the black-boxes may be evaluated at different fidelity levels. A lower fidelity is simply a cheap proxy for the corresponding black-box. These lower fidelities correlate with the actual black-boxes to optimize and can, therefore, be used to reduce the overall cost of solving the optimization problem. Here, we propose Multi-fidelity Joint Entropy Search for Multi-objective Bayesian Optimization with Constraints (MF-JESMOC), a BO method for solving the aforementioned problems. MF-JESMOC chooses the next point, and fidelity level at which to evaluate the black-boxes, as the combination that is expected to reduce the most the joint entropy of the Pareto set and the Pareto front, normalized by the fidelity's evaluation cost. We use Deep Gaussian processes to model each black-box and the dependencies between fidelities. These are powerful probabilistic models that can learn the dependency structure among fidelity levels of each black-box. Several experiments show that MF-JESMOC outperforms other state-of-the-art methods for multi-objective BO with constraints and different fidelity levels in both synthetic and real-world problems.

High-dimensional model selection: transfer learning and projected computation

Authors and affiliations. David Rossell, Universitat Pompeu Fabra.

Presenter/status. David Rossell, Associate Professor.

Abstract.

We define as model selection the structural learning problem where the goal is to learn the subset of truly zero parameters in a model of interest, such as canonical generalized linear models, generalized additive models or graphical models. This is of course one of the most classical and fundamental problems in Statistics and Machine Learning. L_0 penalization methods and their related Bayesian model selection counterparts have optimal mathematical properties for this task, yet much mainstream literature considers such methods to be either unnecessary or impractical. This talk discusses these two objections, and how to ameliorate the issues: it is useful to do this, and can we do this computationally? We first discuss how these methods can be useful, not only in theory but also in practice, and how to improve their performance via data integration (also called transfer learning). For example, sparse model recovery methods enjoy excellent asymptotic properties when certain sparsity and signal strength (betamin) conditions hold, but these assumptions often don't hold in some application domains. We show that data integration pushes the mathematical conditions under which consistent model recovery is possible. Regarding the second objection of computational impracticality, we review recent optimization and MCMC literature showing that, under somewhat strict sparsity assumptions, the computational cost scales linearly with the problem dimension (with high probability, asymptotically). A key practical issue is that such results assume that one can quickly score each candidate model (at constant cost), but even in least-squares the cost is (at least) quadratic in the model dimension and grows also with the sample size n . We propose a new class of projected model selection criteria that score models at constant cost, after an initial pre-processing step, and which enjoy the same asymptotic and practical performance as the costlier exact model scores.

Contributed Talks

Session 1: Time Series & Financial Econometrics

Nonlinear Autoregressive Models for Functional Time Series with Bayesian Additive Regression Trees

Authors and affiliations. Jiahao Cao, University of Texas Health Science Center at Houston; Eoghan O'Neill, University College Dublin; Maria Grith, NEOMA Business School; Anastasija Teterova, Erasmus University Rotterdam; Guanyu Hu, Michigan State University.

Presenter/status. Eoghan O'Neill, Assistant Professor.

Abstract.

This paper introduces Bayesian Additive Regression Tree Models for Function-on-Function Regression. The outcome function is modelled as a linear combination of data-adaptive basis functions. The coefficients of basis functions are determined by sums of trees that can split on both scalar and functional variables, including the lag of the dependent variable. Splitting rules for functions are defined by inner products between functional covariates and linear combinations of fixed basis functions, distinct from the aforementioned data-adaptive basis. We consider a prior on functional splits that allows Markov chain Monte Carlo tree samples to adapt to the data by placing higher probability on selecting a subset of relevant basis functions for a splitting rule. The forecasting performance of the method is evaluated in an application to option pricing implied volatility surface data.

State-dependent realized stochastic covariance model

Authors and affiliations. Igor F. B. Martins, Örebro University Sweden.

Presenter/status. Igor F. B. Martins, PostDoctoral Researcher.

Abstract.

This paper proposes a state-dependent realized stochastic covariance model based on a mixture-of-experts framework. The model extends realized stochastic covariance specifications by allowing latent log-variances and correlations components to evolve according to a finite mixture of persistent processes with state membership determined by a multinomial logit index. In an empirical application, the model identifies economically meaningful states linked to macroeconomic conditions. Accounting for state dependence improves out-of-sample forecasts of realized covariance matrices relative to benchmark models. The proposed model achieves the lowest forecast errors under Frobenius norm and Stein loss functions with Diebold–Mariano tests supporting these improvements. In a global minimum-variance portfolio application, the proposed model delivers the lowest realized portfolio variance highlighting the importance of state-dependent covariance dynamics.

Bayesian Extreme Value Theory with Hawkes-AR-Gumbel Dependence for Extreme CVaR Estimation in Operational Risk

Authors and affiliations. Juan Ballesteros-Gómez, Universidad Pontificia Comillas and Banco Santander; Eduardo C. Garrido-Merchán, Universidad Pontificia Comillas and Instituto de Investigación Tecnológica; Pedro Pablo Pérez-Velasco, Universidad Pontificia Comillas and Banco Santander.

Presenter/status. Juan Ballesteros-Gómez, PhD Student.

Abstract.

Operational risk capital estimation under Basel II/III requires quantifying aggregate losses at extreme confidence levels (99.9%), yet the standard Loss Distribution Approach (LDA) assumes independence between loss frequency and severity –an assumption frequently violated during stress episodes when both the number and the size of losses increase simultaneously. Furthermore, maximum likelihood estimation of tail parameters ignores parameter uncertainty, leading to overconfident and potentially biased risk estimates at extreme quantiles.

We propose a Bayesian framework that combines Extreme Value Theory with a dynamic dependence architecture –the Hawkes-AR-Gumbel model –for operational risk CVaR estimation at confidence levels up to 99.995%. Our model integrates three mechanisms that capture empirically documented features of operational losses:

1. AR(1) latent stress process: $Z_t = \phi Z_{t-1} + \epsilon t$ captures the persistence of crisis regimes, where adverse conditions in one period propagate into the next.
2. Hawkes self-excitation for frequency: $\lambda(t) = \mu(Z_t) + \sum_{t_i < t} \eta e^{\kappa(t-t_i)}$ generates event clustering and overdispersion, reflecting the empirical observation that operational loss events tend to arrive in bursts rather than uniformly.
3. Gumbel copula for upper-tail dependence: $C_\theta(u, v) = \exp\left(-\left[(-\ln u)^\theta + (-\ln v)^\theta\right]^{1/\theta}\right)$ links frequency and severity innovations through an asymmetric copula that strengthens dependence specifically in the extreme tail –unlike the symmetric dependence imposed by shared Normal latent factors.

Inference is performed via Markov Chain Monte Carlo (MCMC) using PyMC, yielding full posterior distributions for all parameters (~ 10 structural parameters plus annual latent variables). CVaR at arbitrary confidence levels is estimated through posterior predictive Monte Carlo simulation, naturally integrating over all sources of uncertainty –parametric, latent, and structural.

We compare three models on simulated operational risk data generated from the Hawkes-AR-Gumbel process:

1. The independent model (standard LDA) ignores all dependence and underestimates CVaR at 99.995% by approximately 40%.
2. The shared latent factor model, which uses a Normal Z_t to create symmetric dependence, captures some of the dependence but fails to account for temporal persistence, event clustering, and upper-tail asymmetry.
3. The Hawkes-AR-Gumbel model recovers the true dependence structure, correctly estimating CVaR at extreme levels. The upper-tail dependence coefficient ($\lambda_U \approx 0.6$) ensures that extreme frequency and severity events are appropriately synchronized, while the AR persistence and Hawkes self-excitation capture temporal dynamics that static models miss.

Our approach is computationally tractable (under 15 minutes on a standard laptop), directly applicable to internal loss databases, and provides risk managers with a principled tool for honest capital quantification that accounts for the dynamic, clustered, and tail-dependent nature of operational losses.

Keywords. Operational risk, Extreme Value Theory, Generalized Pareto Distribution, Bayesian inference, CVaR, Expected Shortfall, Hawkes process, Gumbel copula, autoregressive latent stress, MCMC, frequency- severity dependence.

Bayesian Markov-Switching Partial Reduced-Rank Regression.

Authors and affiliations. María F. Pintado, CUNEF Universidad; Matteo Iacopini, Luiss University; Luca Rossini, University of Milan and Fondazione Eni Enrico Mattei; Alexander Y. Shestopaloff, Queen Mary University of London and University of Guelph.

Presenter/status. María F. Pintado, Assistant Professor.

Abstract.

Reduced-Rank (RR) regression is a powerful dimensionality reduction technique but it overlooks any possible group configuration among the responses by assuming a low-rank structure on the entire coefficient matrix. Moreover, the temporal change of the relations between predictors and responses in time series induce a possibly time-varying grouping structure in the responses. To address these limitations, a Bayesian Markov-switching partial RR (MS-PRR) model is proposed, where the response vector is partitioned in two groups to reflect different complexity of the relationship. A simple group assumes a low-rank linear regression, while a complex group exploits nonparametric regression via a Gaussian Process. Differently from traditional approaches, group assignments and rank are treated as unknown parameters to be estimated. Then temporal persistence in the regression function is accounted for by a Markov-switching process that drives the changes in the grouping structure and model parameters over time. Full Bayesian inference is performed via a partially collapsed Gibbs sampler, which allows uncertainty quantification without the need for trans-dimensional moves. Applications to two real-world macroeconomic and commodity data demonstrate the evidence of time-varying grouping and different degrees of complexity both across states and within each state.

Session 2: Model Selection

Shrinkage through multiple identifiability

Authors and affiliations. Carlos García Meixide and David Ríos Insua, ICMAT-CSIC.

Presenter/status. Carlos García Meixide, PhD Student.

Abstract.

We present a strategy for combining estimators arising from multiple identification functionals of the same (causal) parameter, along with uncertainty regions that satisfy different coverage guarantees. High dependence due to evaluating several functions of the data on the same realization is addressed through a working assumption that delivers a consistent estimator of the target estimand, even under the marginal likelihood misspecification this assumption entails. Our procedure offers a novel perspective on weak identifiability bias so that well-posed setups yield a degenerate prior in our setting. We illustrate this framework in the context of augmenting randomized controlled trials with observational data through identification results under hidden confounding based on instrumental variables. Additionally, we highlight connections with the low-bias regime in the minimax data fusion literature.

Bayesian Model Checks and Improvements for Discrete Data

Authors and affiliations. Lorenzo Cappello, Universitat Pompeu Fabra.

Presenter/status. Lorenzo Capello, Assistant Professor.

Abstract.

A central step in data analysis is model choice. The standard workflow consists of selecting a candidate model, assessing its fit, and, if deficiencies are found, identifying how to improve it and formulating a new candidate to repeat the process. Standard goodness-of-fit tests are well suited for the first step, but, if the null is rejected, they provide little guidance on which aspects of the model are inadequate. Posterior predictive checks address this issue, but introduce other ones, such as the selection of test statistics. In the presence of discrete data, we propose an exponential family-based approach in which the candidate model is augmented by exponential tilting with respect to a vector of statistics designed to capture potential deficiencies. Our key contribution is an automated procedure to define interpretable statistics. The procedure exploits the spectral decomposition of reversible Markov chains with the candidate model as the stationary distribution. We show why such eigenvectors provide meaningful diagnostics of model misfit. Finally, we present inference algorithms to select which statistics to include in this exponential family formulation and show how our proposal integrates naturally into the model criticism and refinement workflow.

Objective Model Prior Probabilities in Variable Selection

Authors and affiliations. James Berger, Duke University; Gonzalo García-Donato, Universidad de Castilla-La Mancha; Elías Moreno, Real Academia de Ciencias; Luis Pericchi, University of Puerto Rico.

Presenter/status. Gonzalo García-Donato, Full Professor.

Abstract.

For many years it was routine to use equal model prior probabilities in Bayesian model uncertainty analysis. At least twenty years ago it became clear that this was problematic, as it tended to favor models that were far too large within the increasingly vast model spaces being explored in genomics and other fields. A popular replacement was to adopt a suggestion of Harold Jeffreys for the variable selection problem in which a total of k possible variables are being considered for inclusion in the model: give the collection of all models containing d variables ($d = 0, \dots, k$) prior probability $1/(k + 1)$ and then divide this prior probability equally among the models in the collection. Many other choices of model prior probabilities that impose severe parsimony have also been introduced. We begin by reviewing the problems with using equal model prior probabilities and then discuss some serious problems with the Jeffreys choice. Finally, we introduce and study a number of objective alternative choices of model prior probabilities, from both numerical and theoretical perspectives.

Session 3: Applied Bayesian Methods

Information-based Bayesian Optimization with Expert Human Feedback

Authors and affiliations. Yago Aguado-Carrillo-de-Albornoz, Universidad Autónoma de Madrid; Daniel Fernández-Sánchez, ELLIS Institute Finland and Aalto University; Daniel Hernández-Lobato, Universidad Autónoma de Madrid.

Presenter/status. Yago Aguado-Carrillo-de-Albornoz, PhD Student.

Abstract.

Bayesian Optimization (BO) is a framework for the optimization of costly black-box functions. Among the various approaches that exist, information-based methods generally exhibit good performance. In some settings, however, we may have access to extra information beyond the evaluation of the objective: human expert knowledge. Specifically, we assume that, with some cost smaller than the evaluation of the actual objective, we may query the opinion of an expert on whether a point should be evaluated or not. This information can be useful to avoid evaluating points that are already known to provide sub-optimal solutions, hopefully speeding up the optimization process. Therefore, in this context, we propose a framework for information-based BO that is able to exploit the aforementioned expert feedback. Since the beliefs of the expert are expected to be correlated with reality, we model dependencies between the actual objective and the expert feedback via a multi-output Gaussian Process (GP). Then, we propose an acquisition function based on Entropy Search (ES) to solve the optimization problem efficiently. Specifically, such an acquisition is able to suggest a candidate point to evaluate next, incorporating to the candidate point a choice, via intelligent decisions, on whether to query the expert or perform a direct evaluation. More precisely, the acquisition takes into account the expected information gain about the solution of the optimization problem when querying or not the expert. The consequence is that it directly determines whether to make direct evaluations or ask for human feedback. The expert feedback is strictly binary (accept/reject), and there exists the possibility of human error, which motivates the use of a likelihood robust to errors. We do so by introducing a learnable parameter in our model, which determines the probability of the expert being wrong or acting in an adversarial manner. Unfortunately, the acquisition function described is intractable. Therefore, we approximate the required computations for entropy evaluation using Variational Inference (VI) and expectation propagation (EP). The method is still lacking empirical evaluation, but overall, we expect that it outperforms pure non-human-in-the-loop BO approaches by efficiently exploiting the less costly expert opinion. Moreover, we also expect it to attain better results than those of frameworks that use expert feedback, but do not consider information-based acquisition functions. Summing up, the proposed approach should speed-up the optimization process, while also making a more efficient use of resources.

Applied Bayesian nonparametric modeling under likelihood ratio order constraints

Authors and affiliations. Adam Olivares and Víctor Peña, Universitat Politècnica de Catalunya; Michael Jauch and Andrés F. Barrientos, Florida State University.

Presenter/status. Adam Olivares, PhD Student.

Abstract.

In many applications, domain knowledge often suggests that two distributions should satisfy a stochastic order. Such orderings arise naturally in fields such as economics and finance, as well as biology, medicine, and public health. Jauch et al. (2025) introduced mixture representations for univariate distribution functions F and G so that $F \leq G$ with respect to the likelihood ratio order. Distributions F and G are ordered in the likelihood ratio order when the ratio of their probability density functions, $f(x)/g(x)$, is a non-increasing function in x . To model this, Jauch et al. (2025) showed that the ratio of two probability density functions is monotone if and only if one can be expressed as a mixture of one-sided truncations of the other. They proposed a Bayesian nonparametric model for density estimation using Dirichlet process mixtures to enforce this constraint and applied it to medical data.

We extend this framework to K univariate stochastically ordered distributions $F_1 \leq \dots \leq F_K$ under the likelihood ratio order, yielding a flexible class of Bayesian nonparametric models for conditional densities with an ordinal covariate. This framework applies to density estimation, ordinal regression, and hidden Markov models. We develop slice-within-Gibbs MCMC algorithms for Dirichlet process mixtures that impose the ordering constraints through likelihood-ratio-ordered mixture representations. Through simulation studies, we show that when the likelihood ratio order holds, the proposed constrained models improve inference relative to independent mixtures, particularly in small or unbalanced samples. We illustrate the methodology with applications to animal movement data, intergenerational income mobility, and retinopathy, analyzed separately for smokers and non-smokers. In hidden Markov models, imposing likelihood ratio constraints on emission densities eliminates label-switching by construction.

Latent advertising stock and saturation for Marketing Mix Modeling (MMM): a field for Bayesian methods' expansion

Authors and affiliations. Miguel Franco Pérez and Xavier Puig Oriol, Universitat Politècnica de Catalunya.

Presenter/status. Miguel Franco Pérez, PhD Student.

Abstract.

Assessing advertising profitability is a cornerstone of data-driven marketing. Marketing Mix Modeling (MMM) quantifies the causal impact of advertising expenditure on Key Performance Indicators (KPIs)—typically sales—by addressing two fundamental non-linear transformations: Adstock and Saturation.

Adstock transformation accounts for the "carryover" effect, where sales at time depend on both current and historical advertising spend up to a window of size. This is modeled via a Rational Distributed Lag Model (RDLM) governed by decay functions (e.g., Geometric or Weibull PDF), where the resulting cumulative adstock, is treated as a latent variable estimated jointly with the decay parameters. This transformation tries to represent the temporal pattern of advertising campaign effects given by Marketing theory with pure decays or growing until a peak before the decaying. Subsequently, this latent adstock is transformed through a Saturation function (e.g., Hill, Logistic, or Michaelis-Menten) to reflect diminishing returns, after a certain point more accumulated advertising may lead not to more sales, as well as potential "warm-up" periods. To ensure correct attribution, these models incorporate temporal controls—such as trend and seasonality—and external covariates.

Bayesian methods are increasingly the standard for MMM due to the inherent complexity and non-linearity of the parameter space. While early frameworks like Meta's Robyn utilized meta-heuristic optimizations (that still included some Bayesian elements like Bayesian pre-modelling of time effects or non-gradient optimization for parameters grid search through Bayesian Estimation of the parameter space), contemporary libraries such as Google's Meridian and PyMC Marketing leverage purely Bayesian inference via Markov Chain Monte Carlo (MCMC) and No-U-Turn Samplers (NUTS).

For dealing with this models, Bayesian methods are gaining predominance. Even the very first open-source library Robyn (Meta®), that was more a Machine Learning metaheuristic rather than an inferential approach, relied on Prophet estimations (a previous also Meta® library) that apply Bayesian coefficients for dealing with time effects (trend with Laplacian null priors for the change-points and seasonality with Fourier terms). Furthermore, current developments and releases, namely Meridian (Google®) and PyMC Marketing (PyMC Labs®), are based on purely Bayesian inference through MCMC and NUTS algorithms.

The Bayesian framework is uniquely suited to this domain because it effectively manages the high-dimensional, non-linear interactions between adstock and saturation while allowing for the seamless integration of informative priors. This enables marketing practitioners to ground latent parameter estimation in historical intuition and experimental benchmarks.

This presentation aims to deconstruct the structure of Marketing Mix Models, with a specific focus on the joint estimation of the latent stock variable and its underlying decay parameters. By arguing the statistical adequacy of the Bayesian approach in dealing with these complex functional forms, the session will provide a deeper understanding of model stability. Finally, an on-going comparative study of Robyn, Meridian, and PyMC Marketing will be presented to demonstrate how their differing architectures behave under varied advertising strategies and the resulting non-linear interactions.

Keywords. Marketing Mix Modeling, sales attribution, adstock models, lag-stock models, non-linear saturation effect, Robyn, Meridian, PyMC Marketing.

Session 4: Adversarial Learning & Robust ML

A Unified Bayesian Framework for Adversarial Robustness

Authors and affiliations. Pablo G. Arce, ICMAT-CSIC; Roi Naveiro, CUNEF Universidad; David Ríos Insua, ICMAT-CSIC.

Presenter/status. Pablo G. Arce, PhD Student.

Abstract.

The vulnerability of machine learning models to adversarial attacks remains a critical security challenge. Traditional defenses, such as adversarial training, typically robustify models by minimizing a worst-case loss. However, these deterministic approaches do not account for uncertainty in the adversary's attack. While stochastic defenses placing a probability distribution on the adversary exist, they often lack statistical rigor and fail to make explicit their underlying assumptions. To resolve these issues, we introduce a formal Bayesian framework that models adversarial uncertainty through a stochastic channel, articulating all probabilistic assumptions. This yields two robustification strategies: a proactive defense enacted during training, aligned with adversarial training, and a reactive defense enacted during operations, aligned with adversarial purification. Several previous defenses can be recovered as limiting cases of our model. We empirically validate our methodology, showcasing the benefits of explicitly modeling adversarial uncertainty.

Bayesian Online Test Time Adaptation: A General Framework

Authors and affiliations. Daniel Corrales and David Ríos Insua, ICMAT-CSIC.

Presenter/status. Daniel Corrales, PhD Student.

Abstract.

Standard supervised machine learning algorithms typically assume identically independent distributed samples in both training and deployment. Further, these algorithms do not usually leverage learning from unlabelled data, effectively wasting task-relevant information. In the standard test-time adaptation literature, machine learning models are first trained using a labelled source dataset and then are adapted in deployment using only the pre-trained model and the incoming unlabelled data. This work addresses the issue of adaptation in scenarios of test-time distributional shifts by designing a fully probabilistic framework that generalises most common approaches in the literature, and, importantly, provides uncertainty estimates on predictions. We exemplify the applicability of this methodology with a set of use cases, growing in complexity in modelling and data dynamics, thus proving the usefulness and scalability of the framework.

Adversarial observations in Probabilistic State-Space Models for Robust Reinforcement Learning

Authors and affiliations. Miguel Santos and David Ríos Insua, ICMAT-CSIC.

Presenter/status. Miguel Santos, PhD Student.

Abstract.

Decision-making under partial or adversarial observability depends on accurate latent-state estimation and a coherent characterization of uncertainty. In this work, we study adversarial perturbations in linear probabilistic state-space models commonly used in reinforcement learning. We focus on settings where an attacker manipulates observations subject to likelihood-based constraints that preserve consistency with the underlying observation model. We examine how these plausible yet strategically distorted observations affect posterior state estimates, uncertainty propagation, and, ultimately, downstream policy decisions. This framework offers a principled approach to analyzing the vulnerability of inference-driven control systems and provides a foundation for the development of more robust reinforcement learning methods. The problem is especially relevant in safety-critical applications such as robotics, where reliable operation under sensor noise, partial failures, and adversarial perturbations is essential.

Tracking Latent Goals for Robust Reinforcement Learning

Authors and affiliations. Mario Chacón-Falcón, ICMAT-CSIC.

Presenter/status. Mario Falcón, PhD Student.

Abstract.

Multi-agent scenarios often face non-stationary environments, where agents' true objectives are hidden and constantly shifting. This is especially true in the presence of adversarial agents, where standard reinforcement learning struggles. To handle this, we propose an approach that explicitly infers an opponent's unobserved rewards by treating them as continuous latent variables. Rather than assuming access to the opponent's payoff, our method uses a Sequential Monte Carlo (SMC) filter within an Expectation-Maximization (EM) algorithm to continuously obtain these hidden intents from observed behavior. We integrate these Bayesian estimates with an information-seeking objective that encourages our agent to actively probe the adversary when their strategy is ambiguous. Furthermore, we apply this inference within a learned world model for safe offline training. We demonstrate that our framework identifies behavioral shifts and adapts, allowing robust performance in adversarial settings.

A Bayesian computational framework for general security games

Authors and affiliations. Jose Manuel Camacho, ICMAT-CSIC; Roi Naveiro, CUNEF Universidad; David Ríos Insua, ICMAT-CSIC.

Presenter/status. Jose Manuel Camacho, PostDoctoral Researcher.

Abstract.

Security games are a powerful and adaptable tool for analyzing strategic and operational problems in defense and homeland security (DHS). However, these games are often analyzed using game-theoretic models that assume common knowledge, even though that assumption rarely holds in DHS settings. Adversarial Risk Analysis (ARA) offers a Bayesian alternative that is better suited to these settings. This work presents a computational approach to studying general security games from an ARA perspective. The methodology can solve any security game formulated as a bi-agent influence diagram, handling both discrete and continuous decision spaces, as well as multiple decisions for each agent.

Session 5: Spatial & Temporal Modeling

Fair spatial predictions in urban analytics: a Bayesian approach

Authors and affiliations. Alan E. Gelfand, Duke University; J. Carlos Jiménez-Revuelta, Universidad de Sevilla; Pepa Ramírez-Cobo, Universidad de Cádiz.

Presenter/status. Pepa Ramírez-Cobo, Associate Professor.

Abstract.

This paper introduces a Bayesian spatial modeling framework that integrates fairness constraints into hierarchical inference. Using a Leroux CAR prior to model spatial dependence, we incorporate a regularization term that penalizes disparities in group-level predictive outcomes. The resulting formulation allows fairness to be controlled within the posterior distribution, providing a flexible way to balance predictive accuracy and equity. An empirical application to urban data illustrates how the method reduces disparities in access to healthcare services.

A Bayesian Spatiotemporal Varying-Coefficient Model for Multi-Patient Intracranial Recordings

Authors and affiliations. Dominik Wielath, Lorenzo Cappello and David Rossell, Universidad Pompeu Fabra.

Presenter/status. Dominik Wielath, PhD Student.

Abstract.

Recording human brain activity by placing electrodes directly on the brain surface enables capturing fast-paced neural computations with a higher signal-to-noise ratio and greater spatial specificity than commonly used extracranial methods. However, electrode placement guided by surgical needs results in limited and uneven spatial coverage, creating the need for integrating data across patients. We develop a hierarchical spatiotemporal Gaussian process varying-coefficient model to estimate the effects of task-specific covariates on brain activity, accounting for spatial dependence, misaligned electrode locations across patients, and between-patient heterogeneity. We apply the model to high-frequency brain activity recordings at 112 electrodes across 9 neurosurgical patients. Each patient performed up to 200 trials of a simple gambling task, with recordings provided as a 57-point time series over 2,850 ms for each trial. The model estimates spatiotemporal effect surfaces jointly across electrodes and time. Exploiting repeated measurements together with separable space-time kernels yields Woodbury- and Kronecker-based computations that make posterior inference computationally tractable. We use the model to study whether task-related effect surfaces exhibit transferable structure across patients via leave-one-patient-out prediction, comparing a patient-specific intercept-only baseline with a model that incorporates effects learned from the remaining patients. The framework provides a principled approach to inference for multi-patient intracranial data with irregular spatial coverage across patients.

A Bayesian spatio-temporal model for precipitation with two sources of zeros

Authors and affiliations. Juan Marcen-Gutierrez, Jorge Castillo-Mateo, Jesús Asín, and Ana C. Cebrián, Universidad de Zaragoza; Alan E. Gelfand, Duke University.

Presenter/status. Juan Marcen-Gutierrez, PhD Student.

Abstract.

Modeling precipitation presents statistical challenges due to its sparsity, positive skewness, and complex spatio-temporal dependence. A critical issue in observed precipitation data is the high frequency of zeros. This frequency of zeros is doubly inflated by a true absence of precipitation and by an observed left-censoring induced by the detection thresholds of measurement devices. To address this, we propose a Bayesian hierarchical double zero-inflated spatio-temporal model designed to recover the precipitation lost due to these detection limits. The model characterizes rainfall through two latent processes: a binary occurrence process with a probit link, and a strictly positive intensity process with a truncated Gamma distribution. The truncated part accounts for the precipitation occurrence and intensity lost due to measurement thresholds. Both components include Gaussian processes and autoregression to capture spatio-temporal dependence. Working within a Bayesian framework enables posterior inference with exact uncertainty quantification. This allows us to formally quantify the discrepancy between the latent true and observed processes. We provide full uncertainty quantification for both the unrecorded precipitation occurrences and the undetected precipitation amounts. The methodology is applied to a 15-year daily spring precipitation dataset from 70 tipping-bucket gauges in the Ebro Basin (Spain).
