UQDM 2025

Workshop on Uncertainty Quantification for Dynamical Modelling

University of Edinburgh July 9–11

Engineering and Physical Sciences Research Council

THE UNIVERSITY *of* EDINBURGH

LONDON MATHEMATICAL SOCIETY EST. 1865

Useful Information

Talks will be held in Room **6206** of the James Clerk Maxwell Building (JCMB). Take the lifts from the JCMB foyer to floor 6, then follow the corridor on your left (the “62_ _” corridor) until you reach 6206.

Refreshments and **lunches** will be served in the Magnet Café area on the **third floor** of JCMB. The Magnet is right in front of the stairs, or to your left as you exit the lifts.

The **workshop dinner** will be held on Thursday 11th at 18:30 at Blonde Restaurant.

Programme

IS: Invited Speaker, CT: Contributed Talk, ST: Short Talk.

Wednesday 9 July

12:00–13:20	**Registriation and welcome lunch**		
13:20–13:30	**Opening remarks**		
13:30–14:15	IS	**Eviatar Bach** University of Reading	Learning probabilistic filters for data assimilation
14:15–15:00	IS	**Svetlana Dubinkina** Vrije Universiteit Amsterdam	Data assimilation with randomized observations
15:00–15:30	**Refreshments**		
15:30–16:00	CT	**Jiaao Wang** University of Edinburgh	Boltzmann Machine, Contrastive Divergence Method and Information Geometry
16:00–16:45	IS	**Hossein Mohammadi** University of Edinburgh	Adaptive sampling for Gaussian process models

Thursday 10 July

09:00–09:45	IS	**Chris Oates** Newcastle University	Prediction-Centric Uncertainty Quantification via MMD
09:45–10:30	IS	**Abdul-Lateef Haji-Ali** Heriot-Watt University	Multilevel Path Branching for Digital Options
10:30–11:00	**Refreshments**		
11:00–11:45	IS	**Laura Mansfield** University of Oxford	Uncertainty Quantification of Machine Learning Parameterisations in Climate Models
11:45–12:30		**Martin Brolly** University of Edinburgh	Representing model error in Earth system predictions
12:30–14:00	**Lunch**		
14:00–14:30	CT	**Nathan Kirk** Illinois Institute of Technology	Generating Representative Samples
14:30–15:15	IS	**James Maddison** University of Edinburgh	High dimensional uncertainty quantification in glaciological inverse problems
15:15–15:45	**Refreshments**		
15:45–16:00	ST	**Simon Urbainczyk** Heriot-Watt University	STRIDE: Sparse Techniques for Regression in Deep Gaussian Processes

16:00–16:15	ST	**Anastasia Istratuca** University of Edinburgh	When and how to smooth a random field
16:15–17:00	IS	**Sebastian Krumscheid** Karlsruhe Institute of Technology	Identifying Effective Stochastic Surrogate Models from Multiscale Data
18:30	**Workshop dinner**		

Friday 11 July

09:00–09:45	IS	**Bojana Rosic** University of Twente	Sub-structured Neural Approaches for Uncertainty-Aware Modeling of Stochastic Dynamics
09:45–10:30	IS	**Peter Challenor** University of Exeter	Uncertainty quantification for dynamical systems: the dynamical emulator
10:30–11:00	**Refreshments**		
11:00–11:30	CT	**Naratip Santitissadeekorn** University of Surrey	Influence network reconstruction from discrete time-series of count data modelled by multidimensional Hawkes processes
11:30–12:15	IS	**Andrew Curtis** University of Edinburgh	Variational Bayesian Full Waveform Inversion and Efficient Assessment of Prior Hypotheses in a 3D World
12:15–13:00	IS	**Katy Tant** University of Glasgow	Approximating the Local Elastic Tensor in Complex Media using a Variational Bayesian Approach to Ultrasonic Travel-Time Tomography
13:00–14:00	**Lunch**		

Learning probabilistic filters for data assimilation

Eviatar Bach

IS

University of Reading

Filtering – the task of estimating the conditional distribution for states of a dynamical system given partial and noisy observations – is important in many areas of science and engineering, including weather and climate prediction. However, the filtering distribution is generally intractable to obtain for high-dimensional, nonlinear systems. Filters used in practice, such as the ensemble Kalman filter (EnKF), provide biased probabilistic estimates for nonlinear systems and have numerous tuning parameters.

I will present a framework for learning a parameterized analysis map – the transformation that takes samples from a forecast distribution, and combines with an observation, to update the approximate filtering distribution – using variational inference. In principle this can lead to a better approximation of the filtering distribution, and hence smaller bias. We show that this methodology can be used to learn the gain matrix, in an affine analysis map, for filtering linear and nonlinear dynamical systems; we also study the learning of inflation and localization parameters for an EnKF. The framework developed here can also be used to learn new filtering algorithms with more general forms for the analysis map.

I will also present some recent work on learning corrections to the EnKF using permutation-invariant neural architectures, leading to superior performance compared to leading methods in filtering chaotic systems. Lastly, I will present some ideas for learning filters using other probabilistic cost functions.

Representing model error in Earth system predictions

Martin Brolly

University of Edinburgh

Predictions of the Earth system are subject to several sources of uncertainty. Most well known is uncertainty in initial conditions. Decades of research have produced a range of data assimilation algorithms to deal with this part of the problem. Less well known but also critically important is uncertainty due to model error. Earth system models routinely neglect, or represent only crudely, important components of the Earth system. Moreover, finite resources mean that even the most important components (the ocean and atmosphere) are not simulated with sufficient resolution. Parameterisations aim to compensate for these errors. In particular stochastic parameterisations aim to quantify the uncertainty that such errors introduce.

I will discuss the challenges in constructing accurate stochastic parameterisations, both in general and in the context of an idealised model of geophysical turbulence. I will present approaches to learning stochastic parameterisations from data, emphasising the need to break traditional assumptions and highlighting a promising route to extending the range of predictability in forecasting.

Uncertainty quantification for dynamical systems: the dynamical emulator

Peter Challenor IS

University of Exeter

Traditionally a lot of uncertainty quantification has been done using fast surrogate models (or emulators). These are often based on Gaussian processes trained on carefully designed experiments. The classic emulator approximates $y = f(x)$ where f is some complex set of equations (usually ODEs or PDEs) that is very expensive to solve numerically. In the case of dynamical systems we are interested in a different problem, $x_{t+1} = f(x_t)$, the dynamical properties of the system (or the flow map). We show that it is possible to build a Gaussian process based emulator for this problem. We discuss the computational issues and give examples of simple dynamical systems we can emulate. Finally, we show how we can scale the methods up to emulate a digital twin of an environmental system (the Tamar Estuary).

Data assimilation with randomized observations

Svetlana Dubinkina IS

Vrije Universiteit Amsterdam

In collaboration with Nazanin Abedini (VU Amsterdam) and Jana de Wiljes (U. Of Ilmenau)

Ensemble Kalman filtering is widely used in many applications, and it can be analyzed via the continuous ensemble Kalman-Bucy framework. The corresponding ensemble Kalman-Bucy filter exhibits long-time stability and accuracy with fully observed state, as has been shown in Wiljes and Tong (2020). In this work, under some condition we show similar results but with partially observed state. Furthermore, we claim that this condition needs to be satisfied with high probability in order to provide filter stability, consequently leading to randomized observations.

Multilevel Path Branching for Digital Options

Abdul-Lateef Haji-Ali IS

Heriot-Watt University

I'll present a new Monte Carlo estimator for pricing digital options when the underlying assets follow stochastic differential equations solved via standard time-stepping schemes like Euler–Maruyama or Milstein. The key idea is to use repeated, hierarchical path splitting to improve the estimator's strong convergence rate with respect to the time step. This results in a Multilevel Monte Carlo method with computational complexity comparable to classical MLMC for smooth payoffs, and dramatically lower than that of traditional Monte Carlo for digital options.

When and how to smooth a random field

Anastasia Istratuca ST

University of Edinburgh, Maxwell Institute Graduate School

We consider the computational efficiency of Multilevel Monte Carlo (MLMC) methods applied to elliptic partial differential equations with random coefficients, where the aforementioned coefficient is extremely oscillatory. In this case, MLMC is limited in the choice of coarsest mesh on the first level, which can render the estimator more costly than standard Monte Carlo methods. To alleviate this, we smooth the random field on the coarse levels, so that these can be chosen independently of the correlation length of its covariance function. We consider two different approaches for doing so and highlight advantages and disadvantages for both.

Generating Representative Samples

Nathan Kirk **CT**

Illinois Institute of Technology

Approximating a probability distribution using a discrete set of points is a fundamental task in modern scientific computation, with applications in uncertainty quantification among other things. We discuss recent advances in this area, including the use of Stein discrepancies and various optimization techniques. In particular, we introduce Stein-Message-Passing Monte Carlo (Stein-MPMC), a graph neural network model and an extension of the original Message-Passing Monte Carlo mode; the first machine-learning algorithm for generating low-discrepancy (space-filling) point sets. Additionally, we present a generalized Subset Selection algorithm, a simpler yet highly effective optimization method.

High dimensional uncertainty quantification in glaciological inverse problems

James Maddison **IS**

University of Edinburgh

Established variational assimilation techniques are widely used for glaciological state estimates, combining satellite observations with a dynamical model to estimate otherwise unseen parameters such as basal sliding or rheology information. Since these techniques have a Bayesian interpretation, with the state estimate defining only a posterior maximizer, it is in principle possible to define and explore the posterior distribution and use this to quantify uncertainty. However, since this is a high-dimensional problem, and since evaluating the dynamical model is expensive, in practice it is highly challenging to gain more detailed posterior information for a real glaciological problem.

Here we apply the local Gaussian approximation, constructing local covariance information using a low rank update approximation for the posterior Hessian, and with the necessary second derivative information computed using a high-level autodiff approach. This method is applied to complicated glaciological problems, including a numerical model for the Amundsen basin, allowing estimates of uncertainties in important quantities such as the Volume Above Floatation.

Uncertainty Quantification of Machine Learning Parameterisations in Climate Models

Laura Mansfield IS

University of Oxford

Climate models simulate the atmospheric circulation using the governing equations, but are limited by model resolution, typically around 100km for simulations on climate timescales. Important processes occuring on lengthscales smaller than this, such as clouds, convection and atmospheric gravity waves, are not directly resolved and must instead be included through parameterisations. This involves simplifying assumptions and can add a significant source of uncertainty into climate models. Machine learning (ML) is emerging as a promising approach for learning parameterisations. ML parameterisations are typically trained on datasets generated by high resolution climate models or existing parameterisations (“offline”), but evaluated based on their performance when coupled into an existing climate model (“online”). Quantifying uncertainties associated with ML parameterisations is crucial for gaining insights into the reliability of hybrid ML-climate models.

I will discuss how we can estimate uncertainties associated with ML parameterisations, considering “epistemic” uncertainty from the ML model and “aleatoric” uncertainty originating from the training dataset. For this, we use the Lorenz 1996 system to explore parameterisation uncertainty by source. I will also present how in more realistic setting, offline evaluation of an ML parameterisation may show small uncertainties in predicted tendencies. However, these can propagate once coupled online, potentially leading to significant uncertainty in climate model circulation that we should consider carefully when building ML parameterisations.

Adaptive sampling for Gaussian process models

Hossein Mohammadi IS

University of Exeter

Gaussian processes (GPs) are generally regarded as the gold standard surrogate model for emulating computationally expensive computer-based simulators. However, the problem of training GPs as accurately as possible with a minimum number of model evaluations remains challenging. We address this problem by suggesting a novel adaptive sampling criterion called VIGF (variance of improvement for global fit). The improvement function at any point is a measure of the deviation of the GP emulator from the nearest observed model output. At each iteration of the proposed algorithm, a new run is performed where VIGF is the largest. Then, the new sample is added to the design and the emulator is updated accordingly. A batch version of VIGF is also proposed which can save the user time when parallel computing is available. The applicability of our method is assessed on a bunch of test functions and its performance is compared with several sequential sampling strategies. The results suggest that our method has a superior performance in predicting the benchmark functions in most cases.

Prediction-Centric Uncertainty Quantification via MMD

Chris Oates IS

Newcastle University

Deterministic mathematical models, such as those specified via differential equations, are a powerful tool to communicate scientific insight. However, such models are necessarily simplified descriptions of the real world. Generalised Bayesian methodologies have been proposed for inference with misspecified models, but these are typically associated with vanishing parameter uncertainty as more data are observed. In the context of a misspecified deterministic mathematical model, this has the undesirable consequence that posterior predictions become deterministic and certain, while being incorrect. Taking this observation as a starting point, we propose Prediction-Centric Uncertainty Quantification, where a mixture distribution based on the deterministic model confers improved uncertainty quantification in the predictive context. Computation of the mixing distribution is cast as a (regularised) gradient flow of the maximum mean discrepancy (MMD), enabling consistent numerical approximations to be obtained. Results are reported on both a toy model from population ecology and a real model of protein signalling in cell biology.

Sub-structured Neural Approaches for Uncertainty-Aware Modeling of Stochastic Dynamics

Bojana Rosic IS

University of Twente

In this talk the surrogate modeling of stochastic nonlinear structural dynamics governed by uncertain, anisotropic material parameters will be discussed. These dynamics arise in complex physical systems where uncertainty is inherent and spatially variable. A surrogate model in the form of a feedforward/LSTM neural network is trained probabilistically, leveraging sparsity-promoting priors to regularize the parameter space and enhance generalization. To improve scalability and interpretability, we propose substructuring the network using various domain decomposition techniques. This allows for localized learning in partitioned input spaces, enabling efficient and accurate modeling of high-dimensional problems.

Influence network reconstruction from discrete time-series of count data modelled by multidimensional Hawkes processes

Naratip Santitissadeekorn CT

University of Surrey

Identifying key influencers from time series data without a known prior network structure is a challenging problem in various applications, from crime analysis to social media. While much work has focused on event-based time series (timestamp) data, fewer methods address count data, where event counts are recorded in fixed intervals. We develop network inference methods for both batched and sequential count data. Here the strong network connection represents the key influences among the nodes. We introduce an ensemble-based algorithm, rooted in the expectation-maximization (EM) framework, and demonstrate its utility to identify node dynamics and connections through a discrete-time Cox or Hawkes process. For the linear multidimensional Hawkes model, we employ a minimization-majorization (MM) approach, allowing for parallelized inference of networks. For sequential inference, we use a second-order approximation of the Bayesian inference problem. Under certain assumptions, a rank-1 update for the covariance matrix reduces computational costs.

Approximating the Local Elastic Tensor in Complex Media using a Variational Bayesian Approach to Ultrasonic Travel-Time Tomography

Katy Tant IS

University of Glasgow

In ultrasonic imaging, to correctly focus scattered wave energy in the imaging domain, we require good knowledge of the underlying spatial distribution of material properties in the object of interest, as these can impact the speed and direction of the propagating waves. Travel-time tomography methods can be used to this end, inverting the fastest time of arrivals between pairs of transmitters and sensors to construct a map of some material property which varies in space. Given the usually high-dimensional and non-linear nature of these problems, much of the related literature has focused on driving these tomography approaches with Markov Chain Monte Carlo methods, which are of course computationally expensive. However, variational Bayesian Inversion approaches offer an efficient alternative.

In this work we apply the stochastic Stein Variational Gradient Descent to invert for some parameterisation of the spatially varying elastic tensor in complex media from ultrasonic travel time data. We show that the travel time data itself is not enough to fully constrain this problem in spatial domains which require high dimensional parameterisations, and sufficient prior knowledge on one of the three anisotropy parameters we introduce (scale, anisotropy strength and orientation) is required.

STRIDE: Sparse Techniques for Regression in Deep Gaussian Processes

Simon Urbainczyk **ST**

Heriot-Watt University

Gaussian processes (GPs) have gained popularity as flexible machine learning models for regression and function approximation with an in-built method for uncertainty quantification. However, GPs suffer when the amount of training data is large or when the underlying function contains multi-scale features that are difficult to represent by a stationary kernel. To address the former, training of GPs with large-scale data is often performed through inducing point approximations (also known as sparse GP regression (GPR)). To aid the latter, deep GPs have gained traction as hierarchical models that resolve multi-scale features by combining multiple GPs. In this work, we combine variational learning with MCMC to develop a particle-based expectation-maximisation method to simultaneously find inducing points within the large-scale data (variationally) and accurately train the GPs (sampling-based). The result is a highly efficient and accurate methodology for deep GP training on large-scale data. We test our method on standard benchmark problems.