Kody J law

Title: Bayesian static parameter estimation for partially observed diffusions using multilevel Monte Carlo
By Jasra, Kamatani, LAW, Zhou

Abstract: This talk will consider estimating parameters for partially observed SDE, which is known to be a challenging problem. A popular class of methods for solving this problem are particle MCMC (pMCMC) methods, such as particle marginal Metropolis-Hastings. Such methods leverage a non-negative unbiased estimator of the un-normalized target arising from a particle filter within a pseudo-marginal MCMC algorithm in order to obtain an asymptotically exact algorithm without ever evaluating the un-normalized target exactly. Here we assume furthermore that the SDE giving rise to the hidden process cannot be solved exactly, and must be approximated at finite resolution. It is well-known that in such contexts the multilevel Monte Carlo (MLMC) method can be used to substantially reduce the cost to achieve a given level of error. The idea is to represent the target expectation as a telescopic sum of increments of increasing cost, and estimate the increments using targets which are coupled in such a way that the increments have decreasing variance. A schedule of decreasing sample numbers can then be carefully constructed based upon the relationship between the variance and the cost, resulting in a substantial speedup. In the context of interest here it is not clear how to construct an exact coupling, and we instead appeal to a carefully constructed approximate coupling of the pairs of particle filters. It will be shown how to construct a consistent estimator with optimal speedup via the approximate coupling.
.

Adam Johansen

Title: On The Asymptotic Genealogy of Sequential Monte Carlo

Abstract: We consider a class of interacting particle system which covers many of the sequential Monte Carlo (SMC) methods, which are widely-used in applied statistics and cognate disciplines across a range of domains.

We consider the genealogical tree embedded into such particle systems by the resampling operation, and identify conditions, as well as an appropriate time-scaling, under which it converges to the Kingman coalescent in the infinite system size limit in the sense of finite-dimensional distributions. This makes the plethora of distributional results known for the Kingman coalescent available for analysis of SMC algorithms, which we illustrate by characterising the limiting mean and variance of the tree height, as well as of the total branch length of the tree. It also greatly improves the tractability of genealogies of SMC methods, which are known to be closely connected to the performance of these algorithms. The conditions which we require to prove convergence are strong, but we demonstrate by simulation that they do not appear to be necessary.

This is joint work with Jere Koskela, Paul Jenkins and Dario Spano.

Freddy Bouchet

Title: Computing return times using rare event algorithms.

Authors: Freddy Bouchet, Charles-Edouard Bréhier, Corentin Herbert, Thibault Lestang, and Francesco Ragone

Abstract: The average time between two occurrences of the same event, referred to as its return time (or return period), is a useful statistical concept for practical applications. However, due to their scarcity, reliably estimating return times for rare events is very difficult using either observational data or direct numerical simulations. For rare events, an estimator for return times can be built from the extrema of the observable on trajectory blocks. Here, we show that this estimator can be improved to remain accurate for return times of the order of the block size. More importantly, we show that this approach can be generalised to estimate return times from algorithms specifically designed to sample rare events. So far those algorithms often compute probabilities, rather than return times. The approach we propose provides a computationally extremely efficient way to estimate numerically the return times of rare events for a dynamical system, gaining several orders of magnitude of computational costs. We illustrate the method on two kinds of observables, instantaneous and time-averaged, using two different rare event algorithms, for a simple stochastic process, the Ornstein–Uhlenbeck process. As an example of realistic applications to complex systems, we finally discuss extreme heat waves and extreme values of the drag on an object in a turbulent flow.

Nick Whiteley

Title: The Viterbi process and parallelized estimation in high-dimensions

Abstract: The Viterbi process is the limiting maximum a-posteriori estimate of the unobserved path in a hidden Markov model as the length of the time-horizon grows. The existence of such a process suggests that approximate inference algorithms which process data segments in parallel may be accurate. It is shown that for models on state-space R^d satisfying a type of dissipative condition related to convexity, such approximations are indeed accurate and moreover scaleable to high dimensional problems because the rate of convergence to the Viterbi process does not necessarily depend on d.

Andrew Stuart

Title: Reconciling Bayesian Regularization And Total Variation Regularization
Joint work with Matt Dunlop (Caltech), Charlie Elliott (Warwick), and Viet Ha Haong (NTU).

Abstract: A central theme in classical algorithms for the reconstruction of discontinuous functions from observational data is perimeter regularization, for example via total variation penalization. On the other hand, sparse or noisy data often demands a probabilistic approach to the reconstruction of images, to enable uncertainty quantification; the Bayesian approach to inversion is a natural framework in which to carry this out. The link between Bayesian inversion methods and perimeter regularization, however, is not fully understood. In this talk two links are studied: (i) the MAP objective function of a suitably chosen phase-field Bayesian approach is shown to be closely related to a least squares plus perimeter regularization objective; (ii) sample paths of a suitably chosen Bayesian level set formulation are shown to possess finite perimeter and to have the ability to learn about the true perimeter. Furthermore, the level set approach is shown to lead to faster algorithms for uncertainty quantification than the phase field approach.

Anthony Lee

Title: Variance Estimation in the Particle Filter

Abstract: This talk concerns numerical assessment of Monte Carlo error in particle filters. We show that by keeping track of certain key features of the genealogical structure arising from resampling operations, it is possible to estimate variances of a number of Monte Carlo approximations that particle filters deliver. All our estimators can be computed from a single run of a particle filter. We establish that as the number of particles grows, our estimators are weakly consistent for asymptotic variances of the Monte Carlo approximations and some of them are also non-asymptotically unbiased. The asymptotic variances can be decomposed into terms corresponding to each time step of the algorithm, and we show how to estimate each of these terms consistently.

Mathias Rousset

Title: Central Limit Theorem (CLT) for Adaptive Multilevel Splitting and Fleming-Viot Particle Models.

Abstract: The Adaptive Multilevel Splitting (AMS) is an importance splitting particle method used to estimate rares events and to simulate conditionally on rare events. In this talk, we will recast AMS as a Fleming-Viot particle model, and present some recent techniques we have obtained to analyze its asymptotic fluctuations (CLT) of Feming-Viot particle systems, including in the case where particles may be killed by hard obstacles.

Etienne Mémin

Title: Stochastic transport for large-scale stochastic dynamics and tracking

Abstract: In this talk I will describe a formalism to derive in a systematic way a large-scale stochastic representation of fluid flows dynamics. This formalism enables to take into account the inherent uncertainty attached to the flow evolution. The uncertainty introduced here aims at representing, through a random field, the small-scale effects that are neglected in the large-scale evolution model.
The resulting dynamics is obtained from a stochastic representation of the Reynolds transport theorem and enables us to get a physically relevant derivation (i.e. from the usual conservation law) of the sought evolution laws. We will review and discuss such a representation for some classical geophysical models and we will show its application to the tracking of closed curves with an image-based data driven dynamics.

Sebastian Reich

Title: Data assimilation as a control problem

Abstract: Continuous-time data assimilation can be rephrased as an interacting particle system, mean field equation, respectively. The feedback particle filter (FPP) formulation is particularly attractive as it can be viewed as controlled version of the free dynamics in the spirit of Kalman's original work. The FPP also naturally leads to the popular ensemble Kalman-Bucy filter. The assimilation of discrete-in-time data is more delicate. However, application of a coupling approach, originally considered by Schroedinger, to the discrete-time data assimilation problem leads again to a controlled version of the underlying stochastic process. Alternatively, optimal transportation and ensemble transform methods can also be used to enhance available particle filters. I will survey these recent developments.

Peter Jan van Leeuwen

Title: Constructing Dimension-independent Particle Filters for High-dimensional Geophysical Problems
Peter Jan van Leeuwen, Javier Amezcua, Mengbin Zhu, Jacob Skauvold

In particle filtering three random variables appear: the index of the parent particle, the position of the particle, and the weight of the particle. In a standard setting one first draws the index from a pre-determined distribution, then the position of a particle from a proposal density, and the weight is then a function of these two. One can show that if the distribution of the indices is uniform there is an optimal choice for the proposal density for the positions of the particles, called the optimal proposal density. This density is optimal in the sense that it minimises the variance of the weights. However, one can easily do better than this, e.g. by choosing the index distribution such that all particles have equal weight. Even though the weights are equal and the particles are all different, it is easy to show that the particles will typically not represent the posterior density well. In this talk we change the order of the drawing, by first drawing the index, then the weights from a uniform distribution, and then determine the position of the particles from these two. There are an infinite number of ways to draw these particles, and I will discuss a simple straightforward choice. I is hard (for me) to determine the mathematical properties of particle filters like this, so I will refer to numerical experiments in high-dimensional settings where the posterior density is known, and show that a small number of particles can still provide essential features of that posterior density.

Alexandros Beskos

Title: Geometric MCMC for infinite-dimensional inverse problems

Abstract: Bayesian inverse problems often involve sampling posterior distributions on infinite-dimensional function spaces. Traditional Markov chain Monte Carlo (MCMC) algorithms are characterized by deteriorating mixing times upon mesh-refinement, when the finite-dimensional approximations become more accurate. Such methods are typically forced to reduce step-sizes as the discretization gets finer, and thus are expensive as a function of dimension. Recently, a new class of MCMC methods with mesh-independent convergence times has emerged. However, few of them take into account the geometry of the posterior informed by the data. At the same time, recently developed geometric MCMC algorithms have been found to be powerful in exploring complicated distributions that deviate significantly from elliptic Gaussian laws, but are in general computationally intractable for models defined in infinite dimensions. In this work, we combine geometric methods on a finite-dimensional subspace with mesh-independent infinite-dimensional approaches. Our objective is to speed up MCMC mixing times, without significantly increasing the computational cost per step (for instance, in comparison with the vanilla preconditioned Crank–Nicolson (pCN) method). This is achieved by using ideas from geometric MCMC to probe the complex structure of an intrinsic finite-dimensional subspace where most data information concentrates, while retaining robust mixing times as the dimension grows by using pCN-like methods in the complementary subspace. The resulting algorithms are demonstrated in the context of three challenging inverse problems arising in subsurface flow, heat conduction and incompressible flow control. The algorithms exhibit up to two orders of magnitude improvement in sampling efficiency when compared with the pCN method.

Igor Shevchenko

Title: Modelling uncertainty using stochastic transport noise in a multilayer quasi-geostrophic model
Colin Cotter, Dan Crisan, Darryl Holm, Wei Pan, Igor Shevchenko

Abstract: In this talk we present a stochastic parameterisation (derived from the stochastic variational approach for geophysical fluid dynamics introduced by Holm (Proc Roy Soc A, 2015)) for unresolved eddy motions in a two-layer quasi-geostrophic (QG) model in a horizontally periodic channel with forcing and dissipation. The parameterisation is based upon the idea of “transport noise”, which models the modifications to the velocity field due to unresolved flow dynamics. This model assumes that the transport of large scale components is accurate, but that the velocity field used to transport these components is missing contributions from unresolved scales. We present a time-integration scheme for the stochastic QG model and describe a procedure for estimating stochastic forcing to approximate unresolved components using data from high resolution simulations. Finally, we compare the high resolution model output with the stochastic model at lower resolution and demonstrate the results of uncertainty quantification experiments.

Christophe Andrieu

Title: When SMC meet SMC and SMC’

Wei Pan

Title: Numerically modelling stochastic Lie transport in fluid dynamics

Abstract: We present a numerical investigation of stochastic advection by Lie transport (SALT) for the damped and driven incompressible 2D Euler fluid flows. According to Holm (Proc Roy Soc, 2015) and Cotter et al. (2017), the principles of transformation theory and multi-time homogenisation, respectively, imply a physically meaningful, data-driven approach for decomposing the fluid transport velocity into its drift and stochastic parts, for a certain class of fluid flows. This results in a stochastic parameterisation for the fluid transport and gives us stochastic partial differential equation models for the fluid flows.

We develop a new methodology to implement this velocity decomposition. The resulting SPDE system is prescribed on a coarse resolution and is used in our data assimilation setup where the observations come from the fine scale PDE system.

Numerically, we perform the following analyses on this velocity decomposition. We first perform uncertainty quantification tests on the Lagrangian trajectories by comparing an ensemble of realisations of Lagrangian trajectories driven by the stochastic differential equation, and the Lagrangian trajectory driven by the ordinary differential equation. We then perform uncertainty quantification tests on the resulting stochastic partial differential equation by comparing the coarse-grid realisations of solutions of the stochastic partial differential equation with the ``true solutions" of the deterministic fluid partial differential equation, computed on a refined grid. In these experiments, we also investigate the effect of varying

the ensemble size and the number of prescribed stochastic terms. Further experiments are done to show the uncertainty quantification results "converge" to the truth, as the spatial resolution of the coarse grid is refined, implying our methodology is consistent. The uncertainty quantification tests are supplemented by analysing the L2 distance between the SPDE solution ensemble and the PDE solution. Statistical tests are also done on the distribution of the solutions of the stochastic partial differential equation. The numerical results confirm the suitability of the new methodology for decomposing the fluid transport velocity into its drift and stochastic parts, in the case of damped and driven incompressible 2D Euler fluid flows. Finally, we describe our data assimilation algorithm and show some encouraging preliminary results.

Michael Pitt

Title: The Correlated Pseudo-Marginal Method (Michael Pitt – presenter - with George Deligiannidis, Arnaud Doucet)

Abstract: The pseudo-marginal algorithm is a popular variant of the Metropolis--Hastings scheme which allows us to sample asymptotically from a target probability density π, when we are only able to estimate an unnormalized version of π pointwise unbiasedly. It has found numerous applications in Bayesian statistics as there are many scenarios where the likelihood function is intractable but can be estimated unbiasedly using Monte Carlo samples. Using many samples will typically result in averages computed under this chain with lower asymptotic variances than the corresponding averages that use fewer samples. For a fixed computing time, it has been shown in several recent contributions that an efficient implementation of the pseudo-marginal method requires the variance of the log-likelihood ratio estimator appearing in the acceptance probability of the algorithm to be of order 1, which in turn usually requires scaling the number N of Monte Carlo samples linearly with the number T of data points. We propose a modification of the pseudo-marginal algorithm, termed the correlated pseudo-marginal algorithm, which is based on a novel log-likelihood ratio estimator computed using the difference of two positively correlated log-likelihood estimators. We show that the parameters of this scheme can be selected such that the variance of this estimator is order 1 as $N,T\rightarrow\infty$ whenever $N/T\rightarrow 0$. By combining these results with the Bernstein-von Mises theorem, we provide an analysis of the performance of the correlated pseudo-marginal algorithm in the large T regime. In our numerical examples, the efficiency of computations is increased relative to the standard pseudo-marginal algorithm by more than 20 fold for values of T of a few hundreds to more than 100 fold for values of T of around 10,000-20,000.

Roland Potthast

Title: Ensemble and Particle Filters for Operational Global Weather Forecasting

Abstract: National weather services and research institution around the globe are reacting to the increasing need to estimate risk and distributions of both standard variables of weather and climate such as temperatures or humidity and high-impact phenomena such as strong precipitation, wind gusts and storm, tornados, hurricanes or fog on all scales of the forecasting process from minutes to days, from days to month, from month to years and decades.
The agenda of weather and climate forecasting and projection today includes the development and operation of ensemble forecasting system (EPS) on all scales, which have

the ability to model and describe the distribution of possible events and as such the variability of extreme weather, its variables and phenomena.
We describe the setup of the ensemble data assimilation (EDA) and ensemble forecasting systems (EPS), including hybrid ensemble variational data assimilation (LETKF+EnVAR), with new results on large-scale non-Gaussian/non-linear data assimilation.

In particular, we discuss a Localized Adaptive Particle Filter (LAPF) and a Local Markov Chain Particle Filter (LMCPF) and its realization on the operational scale of the global ICON weather forecasting model of DWD. We provide theoretical derivations and show results of experiments in the operational framework over one month of assimilation, where observation data include the full scale of operationally used global observations of DWD, including e.g. synop stations, radio sondes, airplanes, infrared and microwave satellite radiances, atmospheric motion vectors and radio occultations.

Jochen Broecker

Title: Statistical evaluation of data assimilation systems under serial dependence

Abstract: The performance of data assimilation systems needs to be evaluated in an objective manner, either with regard to certain specific applications or in a general sense.

In as much as the model for the signal is part of the data assimilation system, assessment of data assimilation systems might also serve as a means to assess model quality.

Evaluation of data assimilation systems can only be performed in a statistical sense, and in this contribution, we focus on the evaluation through performance indices or scores averaged over time, where the scores compare each individual forecast with the corresponding verification.

As many authors have noted, the fact that the correlation structure of the score time series might be very complicated presents a serious difficulty.

In order to apply a Law of Large Numbers or a Central Limit Theorem though, the correlation structure has to be brought at least partly under control.

As we will show, assuming that the forecasts provided by the data assimilation system are in a sense optimal, it becomes possible to analyse the correlation structure of the score time series and thus devise various tests regarding the performance of data assimilation systems.

Several evaluation strategies will be investigated in this light.

N.K. Nichols

Title: "Diagnosis, conditioning and regularization of observation error covariance matrices in ensemble data assimilation".

Joaquin Miguez

Title: Online convergence assessment of particle filters

Speaker: Joaquin Miguez (Univ. Carlos III of Madrid), joint work with Petar M. Djuric (Stony Brook Univ.) and Victor Elvira (IMT Lille Douai).

Abstract: Particle filters are broadly used to approximate posterior distributions in state-space models using random grids with weighted nodes (aka particles). While the convergence of the filter can be guaranteed, under mild assumptions, when the number of particles is increased, the quality of the approximation is usually unknown but strongly dependent on the number of particles and the underlying state space model. In this talk we discuss a novel method for assessing the convergence of particle filters in an online manner, as well as a simple scheme for the online adaptation of the number of particles based on this assessment. The new technique relies on model-invariant statistics that result from a sequential comparison between the actual observations and their predictive probability distributions as approximated by the particle filter. We provide a rigorous theoretical analysis of the proposed methodology, together with illustrative numerical examples.