Sequential Monte Carlo Workshop

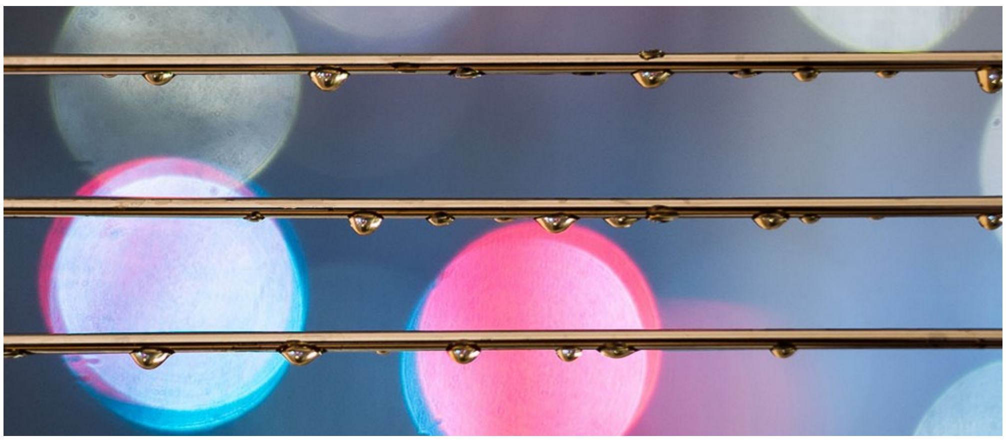

August 30 – September 1, 2017, Uppsala University

1 Acknowledgments

The SMC2017 workshop is financially supported by:

- The Centre for Interdisciplinary Mathematics (CIM) at Uppsala University
- The Swedish Research Council (VR)
- The Swedish Foundation for Strategic Research (SSF)
- The Royal Swedish Academy of Sciences (KVA)
- The Department of Information Technology at Uppsala University

Furthermore, we gratefully acknowledge LYNX Asset Management for sponsoring the conference dinner.

The front page photo is taken by Lawrence Murray.

Contents

2 Program

All events are held at Norrlands Nation, Västra Ågatan 14 (including welcome reception and conference dinner).

Clicking on a talk title jumps to the corresponding abstract.

2.1 Tuesday 29 August

18:00 Welcome reception

2.2 Wednesday, 30 August

8:00-8:50 Registration

8:50-9:00 Welcome

9:00-10:00 Session W1

Session chair: Fredrik Lindsten

09:00	Convergence of resampling algorithms *Nicolas Chopin, ENSAE*
09:30	Terrain Navigation using the Ambient Magnetic Field as a Map *Arno Solin, IndoorAtlas and Aalto University*

10:00-10:30 Coffee Break

10:30-12:00 Session W2: Computational Inference for Probabilistic Programs

Session chair: Thomas Schön

10:30	Inference in higher-order probabilistic programming languages *Brooks Paige, University of Cambridge*
11:00	Making sequential Monte Carlo easier to use via probabilistic programming *Marco Cusumano-Towner, Massachusetts Institute of Technology*
11:30	Delayed Sampling and Automatic Rao-Blackwellization of Probabilistic Programs *Lawrence Murray, Uppsala University*

12:00-14:00 Lunch

14:00-15:00 Session W3

Session chair: Adam Johansen

14:00	Log-concavity and importance sampling in high dimensions *Nick Whiteley, University of Bristol*
14:30	Kernel Algorithm for Gain Function Approximation in the Feedback Particle Filter *Prashant Mehta, University of Illinois*

15:00-15:30 Coffee Break

15:30-16:30 Session W4

Session chair: Alexandre Bouchard-Côté

15:30	Bayesian Inference for Big Data: Single- and Multi-Core Approaches *Murray Pollock, University of Warwick*
16:00	Online adaptation of the number of particles in sequential Monte Carlo methods *Víctor Elvira, Institute Mines Télécom*

2.3 Thursday, 31 August

9:00-10:00 Session Th1

Session chair: Nick Whiteley

9:00 Iterated Auxiliary Particle Filters and Their Application to Diffusion Bridge
Adam Johansen, University of Warwick

9:30 Latent state estimation using control theory
Bert Kappen, Radboud University Nijmegen, Gatsby Computational Neuroscience Unit, UCL

10:00-10:30 Coffee Break

10:30-12:00 Session Th2: High-Dimensional SMC and Data Assimilation

Session chair: Nicolas Chopin

10:30 Second order accurate ensemble transform particle filters
Jana de Wiljes, University of Potsdam

11:00 Approximate Smoothing and Parameter Estimation in High-Dimensional State-Space Models
Axel Finke, University College London

11:30 Multilevel Sequential Monte Carlo methods
Kody Law, Oak Ridge National Laboratory

12:00-14:00 Lunch

14:00-15:00 Session Th3

Session chair: Christian Naesseth

14:00 Controlled sequential Monte Carlo samplers
Jeremy Heng, Harvard University

14:30 On the stability and the uniform propagation of chaos properties of Ensemble Kalman-Bucy filters
Pierre Del Moral, INRIA Bordeaux

15:00-15:30 Coffee break and group photo

15:30-16:00 Poster Spotlights

16:00-00:00 Poster Session and Conference Dinner

2.4 Friday, 1 September

9:00-10:00 Session F1

Session chair: Lawrence Murray

9:00 SMC samplers for finite and infinite mixture models
Maria Lomeli Garcia, University of Cambridge

9:30 Numerically stable online estimation of variance in particle filters
Jimmy Olsson, Royal Institute of Technology

10:00-10:30 Coffee Break

10:30-12:00 Session F2

Session chair: Fredrik Lindsten

10:30 Variable dimension SMC with transformations and sequential Bayesian inference for the coalescent
Richard Everitt, University of Reading

11:00 Map Aided Position for Vehicles in GPS-denied Environments
Rickard Karlsson, NIRA Dynamics and Linköping University

11:30 Communication Efficient Sequential Monte Carlo
Deborshee Sen, National University of Singapore

12:00 Farewell

3 Talk abstracts

Convergence of resampling algorithms

Nicolas Chopin
ENSAE

This talk covers some new results on the asymptotic behaviour of the error introduced by the resampling step of a particle filter. General conditions are given for the consistency of the resampling scheme. It is also shown that, standard resampling schemes may be modified in a certain way so that the mean square error is $o(N^{-1})$. Joint work with Mathieu Gerber, and Nick Whiteley, both from Bristol University.

Terrain Navigation using the Ambient Magnetic Field as a Map

Arno Solin
IndoorAtlas and Aalto University

Variations in the ambient magnetic field can be used as features in indoor positioning and navigation. Here a probabilistic Gaussian process regression based framework is used for composing the vector field magnetic map. The map matching (positioning) is based on a particle filter, and follows the classical terrain matching framework used in aircraft positioning and navigation. The feasibility of this terrain matching approach is demonstrated in real-life indoor positioning examples, where both the mapping and positioning are done using a smartphone.

Inference in higher-order probabilistic programming languages

Brooks Paige
University of Cambridge

The emerging field of probabilistic programming aims to reduce the technical and cognitive overhead for writing and designing novel probabilistic models, by introducing a specialized programming language as an abstraction barrier between modeling and inference. While we would ideally be able to provide "automatic" inference for any probabilistic model, this proves severely challenging for models written in sufficiently expressive languages. In this talk I will discuss some of these difficulties, and show how sequential Monte Carlo methods provide a good general-purpose solution for inference in higher-order probabilistic programs. I will also present some recent work which aims to improve sampler efficiency in this setting by automatically learning proposal distributions.

Making sequential Monte Carlo easier to use via probabilistic programming

Marco Cusumano-Towner
Massachusetts Institute of Technology

Although sequential Monte Carlo techniques are broadly applicable in principle, they can be difficult to apply in practice. Two challenges faced by practitioners are (i) correctly implementing state-of-the-art SMC strategies, and (ii) measuring the approximation error of any given SMC scheme. This talk will present research aimed at addressing these two problems, building on ideas from probabilistic programming. First, this talk will introduce Gen, a new probabilistic programming platform that makes it easier to write complex inference algorithms that include use of SMC. Gen programs can specify generative models and custom inference strategies for those models, including proposals based on variational inference, deep learning, and proposals that are themselves programs written in Gen. Gen will be illustrated using inverse planning problems drawn from probabilistic robotics. Second, this talk will introduce a new algorithm for measuring the accuracy of approximate inference algorithms, based on a new estimator for the symmetric KL divergence between a general SMC sampler and its target distribution. This estimator, called AIDE, is based on the idea that Monte Carlo inference algorithms can be viewed as probabilistic models whose internal random choices correspond to auxiliary variables. AIDE will be illustrated on synthetic hidden Markov models and Dirichlet process mixture model examples.

Delayed Sampling and Automatic Rao-Blackwellization of Probabilistic Programs

Lawrence Murray
Uppsala University

We introduce a dynamic mechanism for the solution of analytically-tractable substructure in probabilistic programs. For inference with Sequential Monte Carlo, it automatically yields improvements such as locally-optimal proposals and Rao-Blackwellization. It works by maintaining a directed graph alongside the running program, evolving dynamically as the program triggers operations upon it. Nodes of the graph represent random variables, and edges the analytically-tractable relationships between them (e.g. conjugate priors and affine transformations). Each random variable is held in the graph for as long as possible, and is sampled only when used by the program in a context that cannot be resolved analytically. This allows it to be conditioned on as many observations as possible before sampling. The approach has been implemented in a new probabilistic programming language called Birch, in which several examples will be given.

Log-concavity and importance sampling in high dimensions

Nick Whiteley
University of Bristol

Direct application of importance sampling can perform badly in high dimensions, in the sense that for even very simple examples involving densities on R^d, computational effort much be increased exponentially with d in order to prevent the Monte Carlo variance of importance sampling approximations from exploding. The purpose of this talk is to illustrate that when the basic importance sampling idea is applied in a more sophisticated manner, which involves an importance sampling correction between the laws of certain diffusion processes constructed in terms of the target distribution known as Jarzynski's identity, the exponential growth of computational cost with dimension can be avoided. The assumption of log-concavity of the target distribution furnishes the diffusions in question with some appealing "dimension-free" ergodic properties.

This is joint work with Christophe Andrieu and James Ridgeway.

Kernel Algorithm for Gain Function Approximation in the Feedback Particle Filter

Prashant Mehta
University of Illinois

I will describe algorithms for numerical approximation of the gain function in the feedback particle filter. The exact gain function is the solution of a Poisson equation involving a probability-weighted Laplacian. The numerical problem is to approximate this solution using only particles sampled from the distribution. For the solution of this problem, a kernel algorithm will be presented along with certain asymptotic estimates for bias and variance. Comparisons with the ensemble Kalman filter will be described and also illustrated with results from some numerical experiments.

Bayesian Inference for Big Data: Single- and Multi-Core Approaches

Murray Pollock
University of Warwick

This talk will introduce novel methodologies for exploring posterior distributions by modifying methodology for exactly (without error) simulating diffusion sample paths. The methodologies discussed have found particular applicability to "Big Data" problems. We begin by presenting the Scalable Langevin Exact Algorithm (ScaLE) and recent methodological extensions (including Re-ScaLE, which avoids the need for particle approximation in ScaLE), which has remarkably good scalability properties as the size of the data set increases (it has sub-linear cost, and potentially no cost as a function of data size). ScaLE has particular applicability in the "single-core" big data setting - in which inference is conducted on a single computer. In the second half of the talk we will present methodology to exactly recombine inferences on separate data sets computed on separate cores - an exact version of "divide and conquer". As such this approach has particular applicability in the "multi-core" big data setting. We conclude by commenting on future work on the confluence of these approaches. Joint work with Hongsheng Dai, Paul Fearnhead, Adam Johansen, Divakar Kumar, Gareth Roberts.

Online adaptation of the number of particles in sequential Monte Carlo methods

Víctor Elvira
Institute Mines Télécom

Sequential Monte Carlo (SMC) methods are broadly used to approximate posterior distributions of hidden states in state-space models by means of sets of weighted particles. While the convergence is guaranteed when the number of particles tends to infinity, the quality of the approximation is usually unknown but strongly dependent on the number of particles. In this talk, we propose a novel methodology for assessing the convergence of the SMC method in an online manner, as well as a simple scheme for the online adaptation of the number of particles based on the convergence assessment. The methodology is based on a sequential comparison between the actual observations and their predictive probability distributions approximated by the SMC method.

Iterated Auxiliary Particle Filters and Their Application to Diffusion Bridge

Adam Johansen
University of Warwick

We present an ofine, iterated particle filter to facilitate statistical inference in general state space hidden Markov models. Given a model and a sequence of observations, the associated marginal likelihood L is central to likelihood-based inference for unknown statistical parameters. We define a class of "twisted" models: each member is specified by a sequence of positive functions ψ and has an associated ψ-auxiliary particle filter that provides unbiased estimates of L. We identify a sequence $\psi^\star$ that is optimal in the sense that the $\psi^\star$-auxiliary particle filter's estimate of L has zero variance. In practical applications, $\psi^\star$ is unknown so the $\psi^\star$-auxiliary particle filter cannot straightforwardly be implemented. We use an iterative scheme to approximate $\psi^\star$, and demonstrate empirically that the resulting iterated auxiliary particle filter significantly outperforms the bootstrap particle filter in challenging settings. This method is particularly well suited to approximate simulation of diffusion bridges as is shown by the empirical study.
Ref: P. Guarniero, A. M. Johansen and A. Lee. The Iterated Auxiliary Particle Filter. To appear in /Journal of the American Statistical Association/
[http://dx.doi.org/10.1080/01621459.2016.1222291]

Latent state estimation using control theory

Bert Kappen
Radboud University Nijmegen, Gatsby Computational Neuroscience Unit, UCL

In this talk, I will give an intuitive introduction to path integral control theory and illustrate some of the applications so far using Monte Carlo sampling. I will focus on the use of control theory for latent state estimation in time series models and discuss an application in fMRI.

Second order accurate ensemble transform particle filters

Jana de Wiljes
University of Potsdam

Sequential Monte Carlo methods (also called particle filters) provide a natural computational approach for state and parameter estimation problems in the context of nonlinear evolution equations. Yet the curse of dimensionality prevents particle filters to be employed for systems with a high-dimensional state space [2]. One option to overcome this limitation is to use localization techniques. However, due to the random resampling steps of standard particle filters, localization cannot be applied in a straightforward manner. The recently proposed ensemble transform particle filter (ETPF) [4] provides a solution to this problem and replaces the resampling step of a standard particle filter by a linear transformation. Further, this linear transformation of the particles allows to design a hybrid approach that couples particle filters with ensemble Kalman filters [3]. However, the transformation step is computationally expensive and leads to an underestimation of the ensemble spread for small and moderate ensemble sizes. These shortcomings have recently been addressed in [1] by developing second-order accurate extensions of the ETPF. These improvements allow one in particular to replace the exact solution of a linear transport problem by its Sinkhorn approximation. It is also demonstrated that the nonlinear ensemble transform filter (NETF) [6, 5] arises as a special case of our general framework. Numerical results for the Lorenz-63 and Lorenz-96 models demonstrate the effectiveness of the proposed modified particle filters.

Joint work with Sebastian Reich and Walter Acevedo.

References

[1] W. Acevedo, J. de Wiljes, and S. Reich, A second-order accurate ensemble transform particle filter, accepted at SIAM J. Sci Comp, https://arxiv.org/abs/1608.08179 (2017).

[2] T. Bengtsson, P. Bickel, and B. Li, Curse of dimensionality revisited: Collapse of the particle filter in very large scale systems, in IMS Lecture Notes - Monograph Series in Probability and Statistics: Essays in Honor of David F. Freedman, vol. 2, Institute of Mathematical Sciences, 2008, pp. 316–334.

[3] N. Chustagulprom, S. Reich, and M. Reinhardt, A hybrid ensemble transform filter for nonlinear and spatially extended dynamical systems, SIAM/ASA J. Uncertainty Quantification, 4 (2016), pp. 592–608.

[4] S. Reich, A nonparametric ensemble transform method for Bayesian inference, SIAM J. Sci. Comput., 35 (2013), pp. A2013– A2024.

[5] J. Tödter and B. Ahrens, A second-order exact ensemble square root filter for nonlinear data assimilation, Mon. Wea. Rev., 143 (2015), pp. 1347–1367.

[6] X. Xiong, I. Navon, and B. Uzungoglu, A note on the particle filter with posterior Gaussian resampling, Tellus, 85A (2006), pp. 456–460.

Approximate Smoothing and Parameter Estimation in High-Dimensional State-Space Models

Axel Finke
University College London

We present approximate algorithms for performing smoothing in a class of high-dimensional state-space models via sequential Monte Carlo methods ("particle filters"). In high dimensions, a prohibitively large number of Monte Carlo samples ("particles") – growing exponentially in the dimension of the state space – is usually required to obtain a useful smoother. Using blocking strategies as in Rebeschini and Van Handel (2015) (and earlier pioneering work on blocking), we exploit the spatial ergodicity properties of the model to circumvent this curse of dimensionality. We thus obtain approximate smoothers that can be computed recursively in time and in parallel in space. First, we show that the bias of our blocked smoother is bounded uniformly in the time horizon and in the model dimension. We then approximate the blocked smoother with particles and derive the asymptotic variance of idealised versions of our blocked particle smoother to show that variance is no longer adversely effected by the dimension of the model. Finally, we employ our method to successfully perform maximum-likelihood estimation via stochastic gradient-ascent and stochastic expectation–maximisation algorithms in a 100-dimensional state-space model.

This is joint work with Sumeetpal S. Singh (University of Cambridge).

Multilevel Sequential Monte Carlo methods

Kody Law
Oak Ridge National Laboratory

Bayesian inference provides a principled and well-defined approach to the integration of data into an a priori known distribution. The posterior distribution, however, is known only point-wise (possibly with an intractable likelihood) and up to a normalizing constant. Monte Carlo methods such as sequential Monte Carlo (SMC) samplers have been designed to sample such distributions. In the case that a complex and intractable model appears inside the statistical model, an additional approximation error is incurred from its necessarily finite resolution approximation. Recently, the multilevel Monte Carlo (MLMC) framework has been extended to such algorithms, so that approximation error can be optimally balanced with statistical sampling error, and ultimately the Bayesian inverse problem can be solved for the same asymptotic cost as solving the deterministic forward problem. This talk will concern the recent development of multilevel SMC (MLSMC) samplers and the resulting estimators for standard quantities of interest as well as normalizing constants. MLMC data assimilation methods, which combine dynamical systems with data in an online fashion, will also be presented, including ML particle filters and ensemble Kalman filters.

Controlled sequential Monte Carlo samplers

Jeremy Heng
Harvard University

In this talk, we introduce a class of controlled sequential Monte Carlo (SMC) samplers, that can be applied to both static models and state space models, and show connections with existing work. We provide theoretical analysis of our proposed methodology and demonstrate significant gains over state-of-the-art methods at a fixed computational complexity.

On the stability and the uniform propagation of chaos properties of Ensemble Kalman-Bucy filters

Pierre Del Moral
INRIA Bordeaux

The Ensemble Kalman filter is a sophisticated and powerful data assimilation method for filtering high dimensional problems arising in fluid mechanics and geophysical sciences. This Monte Carlo method can be interpreted as a mean-field McKean-Vlasov type particle interpretation of the Kalman-Bucy difusions. In contrast to more conventional particle filters and nonlinear Markov processes these models are designed in terms of a difusion process with a difusion matrix that depends on particle covariance matrices. Besides some recent advances on the stability of nonlinear Langevin type difusions with drift interactions, the long-time behaviour of models with interacting difusion matrices and conditional distribution interaction functions has never been discussed in the literature. One of the main contributions of the talk is to initiate the study of this new class of models. The talk presents a series of new functional inequalities to quantify the stability of these nonlinear difusion processes. In the same vein, despite some recent contributions on the convergence of the Ensemble Kalman filter when the number of sample tends to infinity very little is known on stability and the long-time behaviour of these mean-field interacting type particle filters. The second contribution of this talk is to provide uniform propagation of chaos properties as well as Lp-mean error estimates w.r.t. to the time horizon. Our regularity condition is also shown to be sufcient and necessary for the uniform convergence of the Ensemble Kalman filter. The stochastic analysis developed in this talk is based on an original combination of functional inequalities and Foster-Lyapunov techniques with coupling, martingale techniques, random matrices and spectral analysis theory.

SMC samplers for finite and infinite mixture models

Maria Lomeli Garcia
University of Cambridge

In this talk, a variety of novel SMC samplers are introduced for inference in two model classes: an infinite mixture model with Poisson-Kingman priors and a finite mixture model where a prior over the total number of components is used. Successively, I discuss how a sequential Bayesian hypothesis test can be performed, based on Bayes factor estimates for each model class, to compare the two competing hypotheses. This is possible because of the availability of model evidence estimates from the SMC output. I also present some preliminary experimental results of computational performance and of the Bayes factor test.

Numerically stable online estimation of variance in particle filters

Jimmy Olsson
Royal Institute of Technology

This talk discusses variance estimation in sequential Monte Carlo methods, alternatively termed particle filters. The variance estimator that we propose is a natural modification of that suggested by H.P. Chan and T.L. Lai [A general theory of particle filters in hidden Markov models and some applications. Ann. Statist., 41(6): 2877–2904, 2013], which allows the variance to be estimated in a single run of the particle filter by tracing the genealogical history of the particles. However, due particle lineage degeneracy, the estimator of the mentioned work becomes numerically unstable as the number of sequential particle updates increases. Thus, by tracing only a part of the particles' genealogy rather than the full one, our estimator gains long-term numerical stability at the cost of a bias. The scope of the genealogical tracing is regulated by a lag, and under mild, easily checked model assumptions, we prove that the bias tends to zero geometrically fast as the lag increases. As confirmed by our numerical results, this allows the bias to be tightly controlled also for moderate particle sample sizes. This is a joint work with Randal Douc, TELECOM SudParis.

Variable dimension SMC with transformations and sequential Bayesian inference for the coalescent

Richard Everitt
University of Reading

Whole genome sequencing has had a big impact in studying the evolutionary history of pathogens. At the core of many studies is the need to infer the clonal ancestry of a sample from sequence data. This is usually performed using a Bayesian approach, with a coalescent prior on ancestries, and using Markov chain Monte Carlo (MCMC) for inference. However, MCMC can be computationally expensive, and is inflexible in that it needs to be run from scratch whenever new data is received. This talk describes work on a sequential Monte Carlo approach to inference in models with coalescent priors. The proposed approach is also applicable to Bayesian model comparison. An application to mixtures of Gaussians is presented, in which a deterministic transformation is used to improve the efficiency of the SMC.

This is joint work with Richard Culliford, Felipe Medina Aguayo and Daniel Wilson and a draft is available at arxiv.org/abs/1612.06468

Map Aided Position for Vehicles in GPS-denied Environments

Rickard Karlsson
NIRA Dynamics and Linköping University

Most vehicle navigation systems rely on GPS positioning. With support from odometry and inertial sensors, this is a sufficiently accurate and robust solution, but there are future demands. Autonomous cars require higher accuracy and integrity. In GPS-denied environments (parking houses, tunnels, urban canyons) other solutions must be applied.

We present a localization application based on a prior that vehicles spend most time on the road, with odometer as the primary input. The map information is included in a Bayesian setting using the particle filter, rather than standard map matching. In extensive experiments the performance without GPS is shown to have basically the same quality as utilizing a GPS sensor. An indoor vehicle multi-level positioning algorithm is proposed that makes use of an indoor map, as well as dead-reckoning sensor information that is available in every car.

Communication Efficient Sequential Monte Carlo

Deborshee Sen
National University of Singapore

Distributed algorithms have become increasingly significant in recent years propelled by fast technological developments in parallel computing. For sequential Monte Carlo methods, the re-sampling step remains the main difficulty in attempting to parallelize them. We consider a recent algorithm, the so-called αSMC [2], which is an attempt at this. Interactions between particles in this algorithm are controlled by a sequence of "α" matrices. Our goal is to minimize interactions while still leading to stable algorithms. We prove that under standard assumptions the stability properties of the algorithm can be ensured by choosing well-connected, yet sparse, graphs. In particular, choosing Ramanujan graphs [1] lead to stable-in-time algorithms; and more generally, so do expander graphs. We next prove a central limit theorem when interactions are randomly chosen and we also prove that the asymptotic normalized variance of the filtering estimates produced by the αSMC with random interactions is stable as long as there is a certain minimum level of interaction. An offshoot of this is that the αSMC algorithm with random interaction is asymptotically equivalent to the bootstrap particle filter as long as the level of interaction increases to infinity with the number of particles, even if it is at a very slow rate.

Joint work with Alexandre Thiery.

[1] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. Combinatorica, 8(3):261–277, 1988.
[2] Nick Whiteley, Anthony Lee, and Kari Heine. On the role of interaction in sequential Monte Carlo algorithms. Bernoulli, 22(1): 494–529, 2016.

4 Poster abstracts

Simple Nudging Schemes for Particle Filtering

Ömer Deniz Akyıldız⋆ and Joaquín Míguez
Department of Signal Theory and Communications
Universidad Carlos III de Madrid
oakyildi@pa.uc3m.es,jmiguez@ing.uc3m.es

We investigate a new sampling scheme to improve the performance of particle filters in scenarios where either (a) there is a significant mismatch between the assumed model dynamics and the actual system producing the available observations, or (b) the system of interest is high dimensional and the posterior probability tends to concentrate in relatively small regions of the state space. The proposed scheme generates nudged particles, i.e., subsets of particles which are deterministically pushed towards specific areas of the state space where the likelihood is expected to be high, an operation known as *nudging* in the geophysics literature. This is a device that can be plugged into any particle filtering scheme, as it does not involve modifications in the classical algorithmic steps of sampling, computation of weights, and resampling. Since the particles are modified, but the importance weights do not account for this modification, the use of nudging leads to additional bias in the resulting estimators. However, we prove analytically that particle filters equipped with the proposed device still attain asymptotic convergence (with the same error rates as conventional particle methods) as long as the nudged particles are generated according to simple and easy-to-implement rules. Finally, we show numerical results that illustrate the improvement in performance and robustness that can be attained using the proposed scheme. In particular, we show the results of computer experiments involving a misspecified chaotic model and a large dimensional chaotic model, both of them borrowed from the geophysics literature.

Time-varying Combinations of Bayesian Dynamic Models and Equity Momentum Strategies

Nalan Baştürk[1], Stefano Grassi[2], Lennart Hoogerheide[3,4], Herman K. van Dijk[3,5,6], and Agnieszka Borowska[3,4,⋆]
[1]Maastricht University, [2]University of Rome, Tor Vergata, [3]Tinbergen Institute, [4]VU University Amsterdam, [5]Erasmus University Rotterdam, [6]Norges Bank
a.borowska@vu.nl

A novel dynamic asset-allocation approach is proposed where portfolios as well as portfolio strategies are updated at every decision period based on their past performance. A general class of models is specified, combining a dynamic factor model and a vector autoregressive model, where we also allow for stochastic volatility. A Bayesian strategy combination is introduced to combine two model-based strategies, a model momentum strategy based on fitted returns and a residual momentum strategy. This extends the mixture of experts analysis by allowing the strategy weights to be interdependent, time-dependent and incomplete. The presented estimation approach, originating from the forecast combination literature, relies on the implied state space structure of the joint model for the time series and strategy weights. Given the complexity of the resulting non-linear and non-Gaussian structure a novel and efficient particle filter is introduced, based on mixtures of Student's t distributions. Using US industry portfolio returns over almost a century of monthly data, our empirical results indicate that time-varying combinations of flexible models from our developed class with two momentum strategies outperform competing models in terms of mean returns and Sharpe ratios, as well as reduced volatility and the largest loss. The latter result demonstrates the usefulness of the proposed methodology from risk management perspective.

Reconstruction and estimation for non-parametric state-space models

Thi Tuyet Trang CHAU ⋆*, Pierre AILLIOT and Valerié MONBET*
IRMAR, University of Rennes 1
Email: thi-tuyet-trang.chau@univ-rennes1.fr

Missing data are present in many environmental data-sets and this work aims at developing a general method for imputing them. State-space models (SSM) have already extensively been used in this framework. The basic idea consists in introducing the true environmental process, which we aim at reconstructing, as a latent process and model the data available at neighboring sites in space and/or time conditionally to this latent process. A key input of SSMs is a stochastic model which describes the temporal evolution of the environmental process of interest. In many applications, the dynamic is complex and can hardly be described using a tractable parametric model. Here we investigate a data-driven method where the dynamical model is learned using a non-parametric approach and historical observations of the environmental process of interest. From a statistical point of view, we will address various aspects related to SSMs in a non-parametric framework. First we will discuss the estimation of the filtering and smoothing distributions, that is the distribution of the latent space given the observations, using sequential Monte Carlo approaches in conjunction with local linear regression. Then, a more difficult and original question consists in building a non-parametric estimate of the dynamics which takes into account the measurement errors which are present in historical data. We will propose an EM-like algorithm where the historical data are corrected recursively. The methodology will be illustrated and validated on an univariate toy example.

Upper and lower bounds on kernel density estimates in particle filtering

David Coufal⋆
Charles University in Prague, Faculty of Mathematics and Physics
Department of Probability and Mathematical Statistics

Sokolovská 83, 186 75 Praha 8, Czech Republic
e-mail: coufal@karlin.mff.cuni.cz

The upper bounds on kernel density estimates in particle filtering are derived for the Sobolev class of filtering densities. The bounds assure the convergence of kernel density estimates to the filtering density at each filtering time, provided the number of generated particles goes to infinity. The result is achieved by means of Fourier analysis. The lower bounds are discussed as well. They are delivered using the standard approach from non-parameteric estimation that employs tools from information theory. The lower bounds meet upper ones, which indicates optimality of the kernel density estimates in particle filtering.

References

[1] D. Coufal On convergence of kernel density estimates in particle filtering. *Kybernetika*, 5(52):735–756, 2016.
`http://www.kybernetika.cz/content/2016/5/735`

Delayed acceptance ABC-SMC

Richard Everitt

Approximate Bayesian computation (ABC) is now an established technique for statistical inference in the form of a simulator, and approximates the likelihood at a parameter θ by simulating auxiliary data sets x and evaluating the distance of x from the true data y. However, ABC is not computationally feasible in cases where using the simulator for each θ is very expensive. This paper investigates this situation in cases where a cheap, but approximate, simulator is available. The approach is to employ delayed acceptance Markov chain Monte Carlo (MCMC) within an ABC sequential Monte Carlo (SMC) sampler in order to, in a first stage of the kernel, use the cheap simulator to rule out parts of the parameter space that are not worth exploring, so that the "true" simulator is only run (in the second stage of the kernel) where there is a high chance of accepting proposed values of θ. We show that this approach can be used quite automatically, with the only tuning parameter choice additional to ABC-SMC being the number of particles we wish to carry through to the second stage of the kernel. Applications to stochastic differential equation models and latent doubly intractable distributions are presented.

This is joint work with Paulina Rowinska, and is available at arxiv.org/abs/1708.02230

Marginal sequential Monte Carlo for doubly intractable models

Richard Everitt

Bayesian inference for models that have an intractable partition function is known as a doubly intractable problem, where standard Monte Carlo methods are not applicable. The past decade has seen the development of auxiliary variable Monte Carlo techniques for tackling this problem; these approaches being members of the more general class of pseudo-marginal, or exact-approximate, Monte Carlo algorithms, which make use of unbiased estimates of intractable posteriors. Everitt (2017) (https://arxiv.org/abs/1504.00298) investigated the use of exact-approximate importance sampling (IS) and sequential Monte Carlo (SMC) in doubly intractable problems, but focussed only on SMC algorithms that used data-point tempering. This paper describes SMC samplers that may use alternative sequences of distributions, and describes ways in which likelihood estimates may be improved adaptively as the algorithm progresses, building on ideas from Moores (2015) (arxiv.org/abs/1403.4359) and Stoehr et al. (2017) (arxiv.org/abs/1706.10096).

This is joint work with Dennis Prangle, Philip Maybank and Mark Bell.

A rare event approach to high dimensional Approximate Bayesian computation

Richard Everitt

Approximate Bayesian computation (ABC) methods permit approximate inference for intractable likelihoods when it is possible to simulate from the model. However they perform poorly for high dimensional data, and in practice must usually be used in conjunction with dimension reduction methods, resulting in a loss of accuracy which is hard to quantify or control. We propose a new ABC method for high dimensional data based on rare event methods which we refer to as RE-ABC. This uses a latent variable representation of the model. For a given parameter value, we estimate the probability of the rare event that the latent variables correspond to data roughly consistent with the observations. This is performed using sequential Monte Carlo and slice sampling to systematically search the space of latent variables. In contrast standard ABC can be viewed as using a more naive Monte Carlo estimate. We use our rare event probability estimator as a likelihood estimate within the pseudo-marginal Metropolis-Hastings algorithm for parameter inference. We provide asymptotics showing that RE-ABC has a lower computational cost for high dimensional data than standard ABC methods. We also illustrate our approach empirically, on a Gaussian distribution and an application in infectious disease modelling.

This is joint work with Dennis Prangle and Theodore Kypraios and is available at arxiv.org/abs/1611.02492

Online Bayesian clustering of spatial temporal data

Shufei Ge⋆*, Shijia Wang and Liangliang Wang*
Department of Statistics and Actuarial Science, Simon Fraser University
shufeig@sfu.ca

In this paper, we propose an online Bayesian clustering method for the stream type of spatial temporal data based on mixtures of spatial spline regressions. MCMC is derived to make inference for the parameters in our proposed model. However, MCMC has to be rerun when new data arrive. To address this issue we propose an online SMC algorithm for our model. The data augmentation and mixture proposal improve the efficiency of the proposed algorithm. Our method is demonstrated by an application to tumor image data.

Accelerating MCMC with an approximation - Importance sampling versus delayed acceptance

Jouni Helske⋆*, Matti Vihola and Jordan Franks*
Department of Mathematics and Statistics, University of Jyväskylä
jouni.helske@iki.fi

We consider importance sampling (IS) type weighted estimators, based on Markov chain Monte Carlo (MCMC) which targets an approximate marginal of the target distribution [2]. The IS approach, based on unbiased estimators, is consistent, and provides a natural alternative to delayed acceptance (DA) pseudo-marginal/particle MCMC [1]. The IS approach enjoys many benefits against DA, including a straightforward parallelisation and additional flexibility in MCMC implementation. We compare the computational efficiency of IS and DA approaches in a geometric Brownian motion setting where the IS approach provides substantial efficiency improvements over DA.

References

[1] J. Andrés Christen and Colin Fox. Markov chain Monte Carlo using an approximation. *Journal of Computational and Graphical Statistics.*, 14(4), 2005.

[2] Matti Vihola, Jouni Helske, and Jordan Franks. Importance sampling type correction of Markov chain Monte Carlo and exact approximations. Preprint arXiv:1609.02541, 2016.

Online Bayesian Inference for High Dimensional Dynamic Spatio-Temporal Models

Sofia Karadimitriou Sofia

Spatio-temporal processes are geographically represented, i.e. in space, either by being a point, a field or a map and also they vary in time. We would like to make inference on the spatial and temporal change of certain phenomena which for instance could be the air pollution increasing or decreasing in time and ranging in a city or a country. In this era high dimensional datasets can be available, where measurements are observed daily or even hourly at more than one hundred weather stations or locations along with many predictors. Therefore, what we would like to infer is high dimensional and the analysis is difficult to come through due to high complexity of calculations or efficiency from a computational aspect. The first reduced dimension Dynamic Spatio Temporal Model (DSTM) was introduced by Wikle and Cressie (1999) to jointly describe the spatial and temporal evolution of a function observed subject to noise. A basic state space model is adopted for the discrete temporal variation, while a continuous autoregressive structure describes the continuous spatial evolution. Application of Wikle and Cressie's DTSM relies upon the pre-selection of a suit- able reduced set of basis functions and this can present a challenge in practice. In this poster we propose an online estimation method for high dimensional spatio-temporal data based upon DTSM on a spatio temporal count process which attempts to resolve this issue allowing the basis to adapt to the observed data. Specifically, we present a wavelet decomposition for the spatial evolution but where one would typically expect parsimony. This believed parsimony can be achieved by placing a Spike and Slab prior distribution on the wavelet coefficients. The aim of using the Spike and Slab prior, is to filter wavelet coefficients with low contribution, and thus achieve the dimension reduction with significant computation savings. We then propose a Hierarchical Bayesian State Space model, for the estimation of which we offer an appropriate Forward Filtering Backward Sampling algorithm under particle filtering for general state space models which includes static parameter estimation and Gibbs sampling steps for the Spike and Slab wavelet coefficients.

Backward simulation smoothing for Gaussian process state-space models

Roland Hostettler, Toni Karvonen, Filip Tronarp, and Simo Särkkä*
Department of Electrical Engineering and Automation
Aalto University, Espoo, Finland
toni.karvonen@aalto.fi

Gaussian process state-space models (GP-SSMs) are a class of nonlinear, nonparametric Bayesian state-space models where the state transition and measurement functions are modeled as Gaussian processes. This type of models has successfully been employed for modeling dynamic systems in different problems, for example in modeling for control or robotics [1, 2]. However, an inherent drawback is that GP-SSMs are non-Markovian due to the dependency of both the state transition density and the likelihood on the complete past trajectory. This can be alleviated by using approximation techniques such as sparse GPs [3] or Hilbert space approximations [4].

Nevertheless, an approach suitable for inference that does not require approximations to the system per se are sequential Monte Carlo methods which can readily handle non-Markovian state-space models [2, 5]. Hence, in this work, we develop a (naïve) backward simulation smoother for GP-SSMs. The proposed smoothing solution is capable of generating non-degenerate state trajectories and generally lowering the estimation error (as opposed to the more straight-forward joint filtering solution). Furthermore, it is also an important building block in maximum likelihood system identification [6, 7]. We demonstrate the proposed smoother's applicability in an indoor localization application and discuss some practical issues and possible future research directions.

References

[1] M. P. Deisenroth, R. D. Turner, M. F. Huber, U. D. Hanebeck, and C. E. Rasmussen, "Robust filtering and smoothing with Gaussian processes," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1865–1871, July 2012.

[2] R. Frigola, F. Lindsten, T. B. Schön, and C. E. Rasmussen, "Bayesian inference and learning in Gaussian process state-space models with particle MCMC," in *Advances in Neural Information Processing Systems 26*, 2013, pp. 3156–3164.

[3] J. Quiñonero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate Gaussian process regression," *Journal of Machine Learning Research*, vol. 6, pp. 1939–1959, 2005.

[4] A. Svensson, A. Solin, S. Särkkä, and T. Schön, "Computationally efficient Bayesian learning of Gaussian process state space models," in *19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 51, Cadiz, Spain, May 2016, pp. 213–221.

[5] F. Lindsten, T. B. Schön, and M. I. Jordan, "Bayesian semiparametric Wiener system identification," *Automatica*, vol. 49, no. 7, pp. 2053–2063, 2013.

[6] R. Frigola, Y. Chen, and C. E. Rasmussen, "Variational Gaussian process state-space models," in *Advances in Neural Information Processing Systems 27*, 2014, pp. 3680–3688.

[7] M. P. Deisenroth and S. Mohamed, "Expectation propagation in Gaussian process dynamical systems," in *Advances in Neural Information Processing Systems 25*, 2012, pp. 2609–2617.

Bayesian modelling and computation for surface-enhanced Raman spectroscopy

Matthew Moores**, Kirsten Gracie, Jake Carson, Karen Faulds, Duncan Graham, and Mark Girolami*

Department of Statistics, University of Warwick

http://warwick.ac.uk/mmoores

Raman spectroscopy can be used to identify molecules by the characteristic scattering of light from a laser. Each Raman-active dye label has a unique spectral signature, comprised by the locations and amplitudes of the peaks. The presence of a large, nonuniform background presents a major challenge to analysis of these spectra. We introduce a sequential Monte Carlo (SMC) algorithm to separate the observed spectrum into a series of peaks plus a smoothly-varying baseline, corrupted by additive white noise. The peaks are modelled as Lorentzian, Gaussian or Voigt functions, while the baseline is estimated using a penalised cubic spline. Our model-based approach accounts for differences in resolution and experimental conditions. We incorporate prior information to improve identifiability and regularise the solution. By utilising this representation in a Bayesian functional regression, we can quantify the relationship between molecular concentration and peak intensity, resulting in an improved estimate of the limit of detection. The posterior distribution can be incrementally updated as more data becomes available, resulting in a scalable algorithm that is robust to local maxima. These methods have been implemented as an R package, using RcppEigen and OpenMP.

References

[1] M. T. Moores, K. Gracie, J. Carson, K. Faulds, D. Graham & M. Girolami (2016). Bayesian modelling and quantification of Raman spectroscopy. *arXiv preprint*, `arXiv:1604.07299` [stat.AP]

Variational Sequential Monte Carlo

Christian A. Naesseth, Scott W. Linderman, Rajesh Ranganath, and David M. Blei*
Linköping University, Columbia University, Princeton University
`christian.a.naesseth@liu.se`

Variational inference underlies many recent advances in large scale probabilistic modeling. The success of variational approaches depends on (i) formulating a flexible parametric family of distributions; and (ii) optimizing the parameters to find the member of this family that most closely approximates the exact posterior. In this paper we present a new approximating family of distributions, variational sequential Monte Carlo (VSMC), and show how to optimize it in variational inference. VSMC melds variational inference (VI) and sequential Monte Carlo (SMC), providing practitioners with flexible, accurate, and powerful Bayesian inference. VSMC is a variational family that can approximate the posterior arbitrarily well, while still allowing for efficient optimization of its parameters.

A guided intermediate resampling particle filter for inference on high dimensional systems

Joonha Park and Edward L. Ionides*
Department of Statistics, University of Michigan
`joonhap@umich.edu`

Sequential Monte Carlo (SMC) methods, also known as the particle filter methods, have formed the basis of many inference methods on nonlinear partially observed Markov process (POMP) models. However, the performance of the particle filter deteriorates as the dimension of the state space increases [Bengtsson et al.(2008)]. This curse of dimensionality has been a limiting factor in inference using multidimensional nonlinear models. We present a novel particle filter, which we call a guided intermediate resampling filter (GIRF), that scales well to moderately high dimensions. This method is readily applicable to a broad range of models thanks to its plug-and-play property, which means that the user only needs to be able to simulate the process, while its transition density need not be evaluated. Our theoretical and experimental results indicate that this method scales much better than standard methods. We analyzed spatiotemporal epidemic data with a complex mechanistic model using our GIRF method, combining it with the iterated filtering approach for maximum likelihood estimation [Ionides et al.(2015)]. The new particle filter enabled likelihood based inference on the coupling between the measles epidemics in twenty cities in England and Wales in the mid-twentieth century.

References

[Bengtsson et al.(2008)] Thomas Bengtsson, Peter Bickel, and Bo Li. Curse-of-dimensionality revisited: collapse of the particle filter in very large scale systems. In *Probability and Statistics: Essays in Honor of David A. Freedman*, pages 316–334. Institute of Mathematical Statistics, 2008.

[Ionides et al.(2015)] Edward L Ionides, Dao Nguyen, Yves Atchadé, Stilian Stoev, and Aaron A King. Inference for dynamic and latent variable models via iterated, perturbed Bayes maps. *Proceedings of the National Academy of Sciences*, 112(3):719–724, 2015.

A probabilistic scheme for joint parameter estimation and state prediction in complex dynamical systems

Sara Pérez-Vieites**, Inés P. Mariño and Joaquín Míguez*
Department of Signal Theory & Communications, Universidad Carlos III de Madrid
`spvieites@tsc.uc3m.es`

Many problems in the geophysical sciences demand the ability to calibrate the parameters and predict the time evolution of complex dynamical models using sequentially-collected data. Here we introduce a general methodology for the joint estimation of the static parameters and the forecasting of the state variables of nonlinear, and possibly chaotic, dynamical models. The proposed scheme is essentially probabilistic. It aims at recursively computing the sequence of joint posterior probability distributions of the unknown model parameters and its (time varying) state variables conditional on the available observations. The latter are possibly partial and contaminated by noise. The new framework combines a Monte Carlo scheme to approximate the posterior distribution of the fixed parameters with filtering (or data assimilation) techniques to track and predict the distribution of the state variables. For this reason, we refer to the proposed methodology as nested filtering. In this paper we specifically explore the use of Gaussian filtering methods, but other approaches fit naturally within the new framework. As an illustrative example, we apply three different implementations of the methodology to the tracking of the state, and the estimation of the fixed parameters, of a stochastic two-scale Lorenz 96 system. This model is commonly used to assess data assimilation procedures in meteorology. For this example, we compare different nested filters and show estimation and forecasting results for a 4,000-dimensional system.

Global consensus Monte Carlo

Lewis Rendell⋆*, Adam Johansen, Anthony Lee and Nick Whiteley*
University of Warwick and University of Bristol
http://www.warwick.ac.uk/ljrendell

For problems involving large data sets, it is often practical or necessary to distribute the data across multiple processors. Consider a target probability density function expressed as a product of terms, each being the contribution of one such subset of the data. We introduce an instrumental hierarchical model, associating a parameter vector with each subset; these parameter vectors are conditionally independent when conditioned on a top-level parameter. This permits the construction of a Gibbs sampler that yields an approximation of the target density as a marginal of its invariant distribution. Specifying the conditional distribution of the artificial parameter vectors allows the trade-off between computational tractability and fidelity to the original model to be controlled, which may be achieved using a sequential Monte Carlo approach. In contrast to similar such algorithms, this approach requires no assumptions of approximate Gaussianity; we present promising empirical results in such cases, and initial investigations into the algorithm's theoretical properties.

A central limit theorem with application to inference in α-stable regression models

Marina Riabiz⋆ *and Tohid Ardeshiri and Simon J. Godsill*
Department of Engineering, University of Cambridge
mr622@cam.ac.uk

It is well known that the α-stable distribution, while having no closed form density function in the general case, admits a Poisson series representation (PSR) in which the terms of the series are a function of the arrival times of a unit rate Poisson process. In our previous work we have shown how to carry out inference for regression models using this series representation, which leads to a very convenient conditionally Gaussian framework, amenable to straightforward Gaussian inference procedures. The PSR has to be truncated to a finite number of terms for practical purposes. The residual terms have been approximated in our previous work by a Gaussian distribution with fully characterised moments. In this paper we present a new Central Limit Theorem (CLT) for the residual terms which serves to justify our previous approximation of the residual as Gaussian. Furthermore, we provide an analysis of the asymptotic convergence rate expressed in the CLT.

Volatility Spillovers with Multivariate Stochastic Volatility Models

Yuliya Shapovalova⋆ *and Michael Eichler*
Department of Quantitative Economics, Maastricht University
y.shapovalova@maastrichtuniversity.nl

Co-movements in financial time series suggest the presence of volatility spillover effects among financial markets. Understanding fundamentals behind this phenomena is important for portfolio managers and policy makers. Currently in the literature GARCH-type models is the dominating approach for detecting volatility spillovers. The inference is often based on notions of causality in mean and variance. In this paper, we aim to analyze volatility spillovers using a more natural approach for volatility modeling: multivariate stochastic volatility models (MSVM). The structure of MSVM allows to test for causality in volatility processes directly, and in contrast to GARCH models, causality in variance and causality in volatility do not coincide in this framework. However, due to the presence of latent volatility processes estimation of this class of models is a difficult task. We start with adopting an off-shelf solution – bootstrap and auxiliary particle filters in Particle Markov Chain Monte Carlo setting, and discuss limitations that arise in the multivariate case. We further discuss possibilities to improve the off-shelf methodology, by using different particle filtering schemes (particle efficient importance sampling or iterated auxiliary particle filter), adopting Hamiltonian Monte Carlo methods and combining pMCMC with Bayesian approximate methods to improve computational efficiency and speed up the convergence.

Terrain Navigation using the Ambient Magnetic Field as a Map

Arno Solin⋆
Department of Computer Science, Aalto University
http://arno.solin.fi

Variations in the ambient magnetic field can be used as features in indoor positioning and navigation. Here a probabilistic Gaussian process regression based framework [1] is used for composing the vector field magnetic map. The map matching (positioning) is based on a particle filter [2], and follows the classical terrain matching framework used in aircraft positioning and navigation. The feasibility of this terrain matching approach is demonstrated in real-life indoor positioning examples, where both the mapping and positioning are done using a smartphone.

References

[1] A. Solin, M. Kok, N. Wahlström, T. B. Schön, and S. Särkkä. Modeling and interpolation of the ambient magnetic field by Gaussian processes. Accepted for publication in IEEE Transactions on Robotics.

[2] A. Solin, J. Kannala, S. Särkkä, and E. Rahtu. Terrain navigation in the magnetic landscape: Particle filtering for indoor positioning. In *European Navigation Conference (ENC)*, pages 1–9. IEEE, 2016.

Multi-dipole estimation from (simultaneous) MEG and EEG data: a semi-analytic Sequential Monte Carlo approach.

Sara Sommariva[⋆,1] *and Gianvittorio Luria*[2] *and Alberto Sorrentino*[2]

1.Department of Neuroscience and Biomedical Engineering, Aalto University

2.Department of Mathematics, Università di Genova

sara.sommariva@aalto.fi

From a macroscopic point of view, any brain function is characterized by the activation and interaction of specific brain areas.
Magnetoencephalography (MEG) and Electroencephalography (EEG) are two modern neuroimaging techniques featuring an outstanding temporal resolution which makes them particularly suitable to study such functional networks. To this aim, MEG/EEG data are typically analyzed in two steps. First, the neural currents that have generated the recorded data are estimated (*source reconstruction*); then, the statistical relationship between the activity of the involved brain areas is computed (*functional connectivity analysis*).

In [1] and [2] we have shown how the Sequential Monte Carlo (SMC) approach introduced in [3] can be used to perform the first step of the process, namely source reconstruction. We modelled neural sources as the superimposition of an unknown number of point-like sources, termed dipoles, each one representing the activity of a small brain area. Both the number and the locations of the dipoles are assumed to remain fixed in time, while their moments, i.e. their intensities and orientations, are allowed to change. Under these hypothesis both the MEG and EEG forward problems can be written in terms of a conditionally linear model. Indeed, data depend linearly on the dipole moments, while depend non-linearly on their number and locations. Assuming a Gaussian prior for the linear variables and for the noise probability distribution, we used the SMC approach to approximate the posterior distribution of the number of sources and their locations while we computed analytically the conditional posterior of the time-varying dipole moments. Moreover, in [4] we have shown how, setting properly the noise covariance matrix, the presented algorithm can be used for the joint analysis of MEG and EEG data, resulting in a lower variance of the posterior distribution and thus in a lower uncertainty on the solution. Our next goal is to study the application of this algorithm for connectivity analysis, i.e. for recovering the dynamic interaction between the different active dipoles together with their time-courses.

References

[1] S. Sommariva and A. Sorrentino Sequential Monte Carlo samplers for semi-linear inverse problems and application to Magnetoencephalography *Inverse Problems*, 30 114020, 2014.

[2] A. Sorrentino, G. Luria and R. Aramini Bayesian Multi-Dipole Modeling of a Single Topography in MEG by Adaptive Sequential Monte Carlo Samplers *Inverse Problems*, 30 045010, 2014

[3] P. Del Moral, A. Doucet and A. Jasra Sequential Monte Carlo samplers *Journal Of The Royal Statistical Society Series B*, 68 41–36, 2006

[4] F. Rossi, G. Luria, S. Sommariva and A. Sorrentino Bayesian multi-dipole localization and uncertainty quantification from simultaneous EEG and MEG recordings *EMBEC 2017, IFMBE Proceedings*, 65, 2017

Sequential Monte Carlo for Static Bayesian Models with Independent MCMC Proposals

L. F. South, A. N. Pettitt & C. C. Drovandi*
School of Mathematical Sciences, Queensland University of Technology, Australia
ARC Centre of Excellence for Mathematical & Statistical Frontiers (ACEMS)
leah.south@hdr.qut.edu.au

Sequential Monte Carlo (SMC) represents a powerful alternative to Markov chain Monte Carlo (MCMC) for sampling from the posterior distribution of static Bayesian models. Under likelihood annealing SMC, an easy to sample from distribution is connected with the target through a sequence of reweighting, resample and move steps. The move step is typically the most computationally expensive step and can have an impact on the accuracy of evidence estimates and posterior expectation estimates. A common choice in the literature is to use several iterations of an MCMC kernel formed using the population of particles. However, the proposals made throughout the sequence of distributions are essentially wasted because only the final accepted particles target the posterior distribution. We propose to take further advantage of the population of particles by forming an independent proposal based on a copula model. As a by-product of the independent proposal choice, we are able to perform posterior inference using all proposals generated in the SMC process and we are able to consider a novel importance sampling based estimator of the marginal likelihood. We demonstrate that our approach can lead to more efficient posterior approximations and more precise estimates of the evidence compared with the multivariate normal random walk proposal.

References

[1] South LF, Pettitt AN, Drovandi CC (2017) Sequential Monte Carlo for static Bayesian models with independent MCMC proposals. http://eprints.qut.edu.au/101729/

[2] Gramacy R, Samworth R, King R (2010) Importance tempering. Statistics and Computing 20(1):1–7

[3] Nguyen TLT, Septier F, Peters GW, Delignon Y (2014) Improving SMC sampler estimate by recycling all past simulated particles. In: 2014 IEEE Workshop. In Statistical Signal Processing (SSP), pp p117–120

[4] Chopin N (2002) A sequential particle filter method for static models. Biometrika 89(3):539–552

[5] Tran MN, Giordani P, Mun X, Kohn R, Pitt MK (2014) Copula-type estimators for flexible multivariate density modeling using mixtures. Journal of Computational and Graphical Statistics 23(4):1163–1178

[6] Cornuet JM, Marin JM, Mira A, Robert CP (2012) Adaptive multiple importance sampling. Scandinavian Journal of Statistics 39(4):798–812

Maximum likelihood estimation by re-using the particles

Andreas Svensson, Fredrik Lindsten and Thomas B. Schön*
Department of Information Technology, Uppsala University
http://www.it.uu.se/katalog/andsv164

In the interest of maximum likelihood parameter estimation in nonlinear state-space models, we make the following observations concerning the particle filter:

- The propagation of the states from time t-1 to time t does not have to depend on the model parameters, but can be done using an arbitrary proposal.
- The resampling does not have to be done with respect to the importance weights, but can be made with respect to any weights.

These two points suggest that we first can run the particle filter with an initial parameter value to estimate the likelihood, and later also compute a likelihood estimate for another parameter value based on the very same particles. We demonstrate in simulated examples that this approach can provide competitive results in certain cases. A similar idea was proposed by [1], but has, to the best of our knowledge, not previously been fully explored for the use of parameter estimation.

References

[1] François Le Gland. Combined use of importance weights and resampling weights in sequential Monte Carlo methods *ESAIM: Proc.*, 19:85–100, 2007.

Phenotypic modelling of the polyfunctional T cell response

Nicola Trendel

T cells are integral parts of the adaptive immune system. Binding of the T cell receptor (TCR) to antigenic ligands can result in multiple different functional T cell responses but how these responses depend on the antigen affinity and dose is incompletely understood. Here, we derive a phenomenological model that is consistent with experimental phenotypes. We systematically measured TCR levels and cytokine release in response to antigens of varying affinity and dose at different time points after stimulation. Ligand-induced TCR internalisation is fast and irreversible on the experimental time scale. Consequently, cells stimulated with high doses of antigen shut down cytokine secretion earlier than cells stimulated with lower doses of antigen. A receptor internalisation model with readout-specific activation thresholds explains experimentally observed features. Next, we aim to systematically search models of equal or lesser complexity to identify a minimal, coarse-grained model encompassing only those assumptions necessary to explain the data, thus avoiding complexity bias. Our approach will likely include a combination of automated adaptive inference1 and ABC-SMC.

Rao-Blackwellized Particle Implementation Of Stochastic Expectation Maximization In MEG/EEG for Joint Estimation of Neural Sources and Connectivity

Narayan Subramaniyam[1], Sara Sommariva[1], Filip Tronarp[2⋆], Xi Chen[3], Simo Särkkä[2] and Lauri Parkkonen[1]

[1] Aalto University, Department of Neuroscience and Biomedical Engineering

[2] Aalto University, Department of Eleictrical Engineering and Automation

[3] Cavendish Laboratory, University of Cambridge

filip.tronarp@aalto.fi

Functional connectivity is an important topic in neuroscience which can serve as a tool for basic research as well as a clinical tool for diagnosis by, for instance, aiding in localisation of brain regions generating epilleptic seizures. Estimating the functional connectivity involves four components (i) determining the number of electromagnetic sources (ii) determining the location of the sources, (iii) determining the strength (dipole moment) of the sources and

(iv) determining their statistical interdependence. The current standard in the neuroscience community is to use a two-step approach where (ii) and (iii) are solved by, e.g, minimum norm estimation after which (iv) can be solved, however this induces a bias [1]. Here the problem is tackled by assuming the solution to (i) is known a priori while a state-space model is formulated for (ii) and (iii) given (iv), where the latter is assumed to impose linear dependence between sources. The problem is thus reduced to identification in a non-linear state-space model with a linear sub-structure which is solved by a Rao-Blackwellized particle filter combined with an expectation-maximisation algorithm. The proposed approach allows for reconstruction of neural sources and connectivity in a realistic and continuous head model [2].

References

[1] NP Subramaniyam and F Tronarp and S Särkkä and L Parkkonen, *Expectation—maximization algorithm with a nonlinear Kalman smoother for MEG/EEG connectivity estimation*, EMBEC & NBC 2017.

[2] Xi Chen and Simo Särkkä and Simon Godsill, *A Bayesian particle filtering method for brain source localisation*, Digital Signal Processing, 47:192-204, 2015.

Sequential Monte Carlo for Dynamic Latent Space Networks

Kathryn Turnbull, Christopher Nemeth, Matthew Nunes and Tyler McCormick*
STOR-i Centre for Doctoral Training, Lancaster University
k.turnbull@lancaster.ac.uk

The ubiquity of high-dimensional network data motivates the study of efficient inference schemes. When considering network data, we note that high-dimensional may refer to the number of nodes, the number of observations in time, or both. Our focus is on temporally evolving networks where we may utilise a state space modelling approach, hence providing a natural framework for the application of Sequential Monte Carlo methods. In particular, we consider a latent space network model whereby nodes are assumed to lie in a latent space and that the probability of a connection between two nodes is proportional to their similarity in this space. However, to avoid applying SMC to a high-dimensional and highly-dependent state space, we propose an inference scheme on the pairwise connection probabilities. Furthermore, we note that these methods can be extended to weighted edges and are amenable to online updating.

References

[1] Peter Hoff, Adrian Raftery, and Mark Handcock, *Latent Space Approaches to Social Network Analysis*, JASA, Volume 97, Issue 460, 2002.

[2] Daniel Sewell and Yuguo Chen, *Latent Space Models for Dynamic Networks*, JASA, Volume 110, Issue 512, 2015.

[3] Christophe Andrieu, Arnaud Doucet and Roman Holenstein, *Particle Markov chain Monte Carlo methods*, JRSS: Series B, Volume 72, Issue 3, 2010.

[4] Daniele Durante and David Dunson, *Nonparametric Bayes dynamic modelling of relational data*, Biometrika, Volume 101, Issue 4, 2014.

Multivariate Output Analysis for Markov Chain Monte Carlo

Dootika Vats**, James Flegal, Galin Jones*
Department of Statistics, University of Warwick
`http://warwick.ac.uk/dvats`

Markov chain Monte Carlo (MCMC) is a method of producing a correlated sample in order to estimate expectations with respect to a target distribution. A fundamental question is when should sampling stop so that we have good estimates of the desired quantities? The key to answering these questions lies in assessing the Monte Carlo error through a multivariate Markov chain central limit theorem. However, the multivariate nature of this Monte Carlo error has been ignored in the MCMC literature. I will give conditions for consistently estimating the asymptotic covariance matrix. Based on these theoretical results I present a relative standard deviation fixed volume sequential stopping rule for deciding when to terminate the simulation. This stopping rule is then connected to the notion of effective sample size, giving an intuitive, and theoretically justified approach to implementing the proposed method. The finite sample properties of the proposed method are then demonstrated in examples. The results presented in this talk are based on joint work with James Flegal (UC, Riverside) and Galin Jones (U of Minnesota).

A probabilistic scheme for joint parameter estimation and state prediction in complex dynamical systems

Sara Vieites

Many problems in the geophysical sciences demand the ability to calibrate the parameters and predict the time evolution of complex dynamical models using sequentially-collected data. Here we introduce a general methodology for the joint estimation of the static parameters and the forecasting of the state variables of nonlinear, and possibly chaotic, dynamical models. The proposed scheme is essentially probabilistic. It aims at recursively computing the sequence of joint posterior probability distributions of the unknown model parameters

and its (time varying) state variables conditional on the available observations. The latter are possibly partial and contaminated by noise. The new framework combines a Monte Carlo scheme to approximate the posterior distribution of the fixed parameters with fi ltering (or data assimilation) techniques to track and predict the distribution of the state variables. For this reason, we refer to the proposed methodology as nested filtering. In this paper we speci cally explore the use of Gaussian filtering methods, but other approaches fi t naturally within the new framework. As an illustrative example, we apply three different implementations of the methodology to the tracking of the state, and the estimation of the fi xed parameters, of a stochastic two-scale Lorenz 96 system. This model is commonly used to assess data assimilation procedures in meteorology. For this example, we compare different nested fi lters and show estimation and forecasting results for a 4,000-dimensional system.

Bayesian Phylogenetics using Particle Gibbs

Liangliang Wang⋆ *and Shijia Wang*
Department of Statistics and Actuarial Science, Simon Fraser University, Canada
`http://people.stat.sfu.ca/~lwa68`

Bayesian phylogenetics, which approximates a posterior distribution of phylogenetic trees, has become more and more popular with the development of Monte Carlo methods. Standard Bayesian estimation of phylogenetic trees can handle rich evolutionary models but requires expensive Markov chain Monte Carlo (MCMC) simulations, which may suffer from two difficulties, the curse of dimensionality and the local-trap problem. Our previous work [1] has shown that the combinatorial sequential Monte Carlo (CSMC) method can serve as a good alternative to MCMC in posterior inference over phylogenetic trees. However, the simple proposal distribution used in CSMC is inefficient to combine with MCMC in the framework of the particle Gibbs sampler. Moreover, CSMC is inapplicable to the particle Gibbs with ancestor sampling [2]. In this poster, we will present a more efficient CSMC method, called CSMC-BF, with a more flexible proposal. The proposed CSMC-BF can improve the performance of the particle Gibbs sampler compared with the original CSMC, and can be used in the particle Gibbs with ancestor sampling. We will demonstrate the advantages of the proposed CSMC-BF using simulation studies and real data analysis.

References

[1] Liangliang Wang, Alexandre Bouchard-Côté, and Arnaud Doucet. Bayesian phylogenetic inference using a combinatorial sequential Monte Carlo method. *Journal of the American Statistical Association*, 110(512):1362D1374, 2015.

[2] Fredrik Lindsten, Michael I. Jordan, and Thomas B. Schön Particle Gibbs with ancestor sampling. *Journal of Machine Learning Research*, 15(1):2145D2184, 2014.

An online sequential Monte Carlo EM algorithm for recommender system

Shijia Wang⋆ *and Liangliang Wang*
Department of Statistics and Actuarial Science
http://www.sfu.ca/~shijiaw/

Our work is motivated by a personalized recommender system with dynamic preference changes. The existing method is to use the batch EM algorithm for estimating the parameters in the hidden Markov model (HMM) model. A challenge with the model estimation is complex real data stream which leads to computational issues. We address this issue by proposing a particle based online EM algorithm for static parameters in an HMM. The proposed online SMC-EM combines the ideas of EM algorithm using complete sufficient statistics to reparameterize the model and the recursive property of sequential Monte Carlo (SMC). We demonstrate the effectiveness of our proposed method to recommend products using a simulation study and two real data sets. Both the simulation study and the real data analysis show the proposed online SMC-EM algorithm outperforms the batch EM for the HMM based dynamic model in terms of making recommendations.

Bias-variance trade-off for high-dimensional particle filters using artificial process noise

Anna Wigren⋆*, Fredrik Lindsten and Lawrence Murray*
Department of Information Technology, Uppsala University
anna.wigren@it.uu.se

For high-dimensional systems the particle filter suffers from degenerated particle weights both when using the standard and the optimal proposal density. However, in many practical applications the optimal proposal can significantly outperform the standard proposal and provide competitive results for fairly high-dimensional systems. Furthermore, the benefit of using the optimal proposal has been shown to increase when the magnitude of the process noise increases [1]. The optimal proposal is, however, difficult to use in practice since it is only possible to sample from it for a few specific types of models, such as the linear Gaussian model, but the system model can be approximated to enable the use of the optimal proposal. Here the aim is to approximate by first propagating according to the standard proposal and then adding artificial process noise. This introduces a bias but will also reduce the variance since it is possible to use the optimal proposal for the additional artificial noise step. The magnitude of the variance of the artificial noise can be varied to find the trade-off between bias and variance which gives the best overall performance. Simulation results show that a clear improvement in performance is possible.

References

[1] Chris Snyder, Thomas Bengtsson, and Mathias Morzfeld. Performance bounds for particle filters using the optimal proposal. *Monthly Weather Review*, 143(11):4750–4761, 2015.