Recent Advances in Sequential Monte Carlo
19th – 21st September 2012

Contents

THE UNIVERSITY OF WARWICK

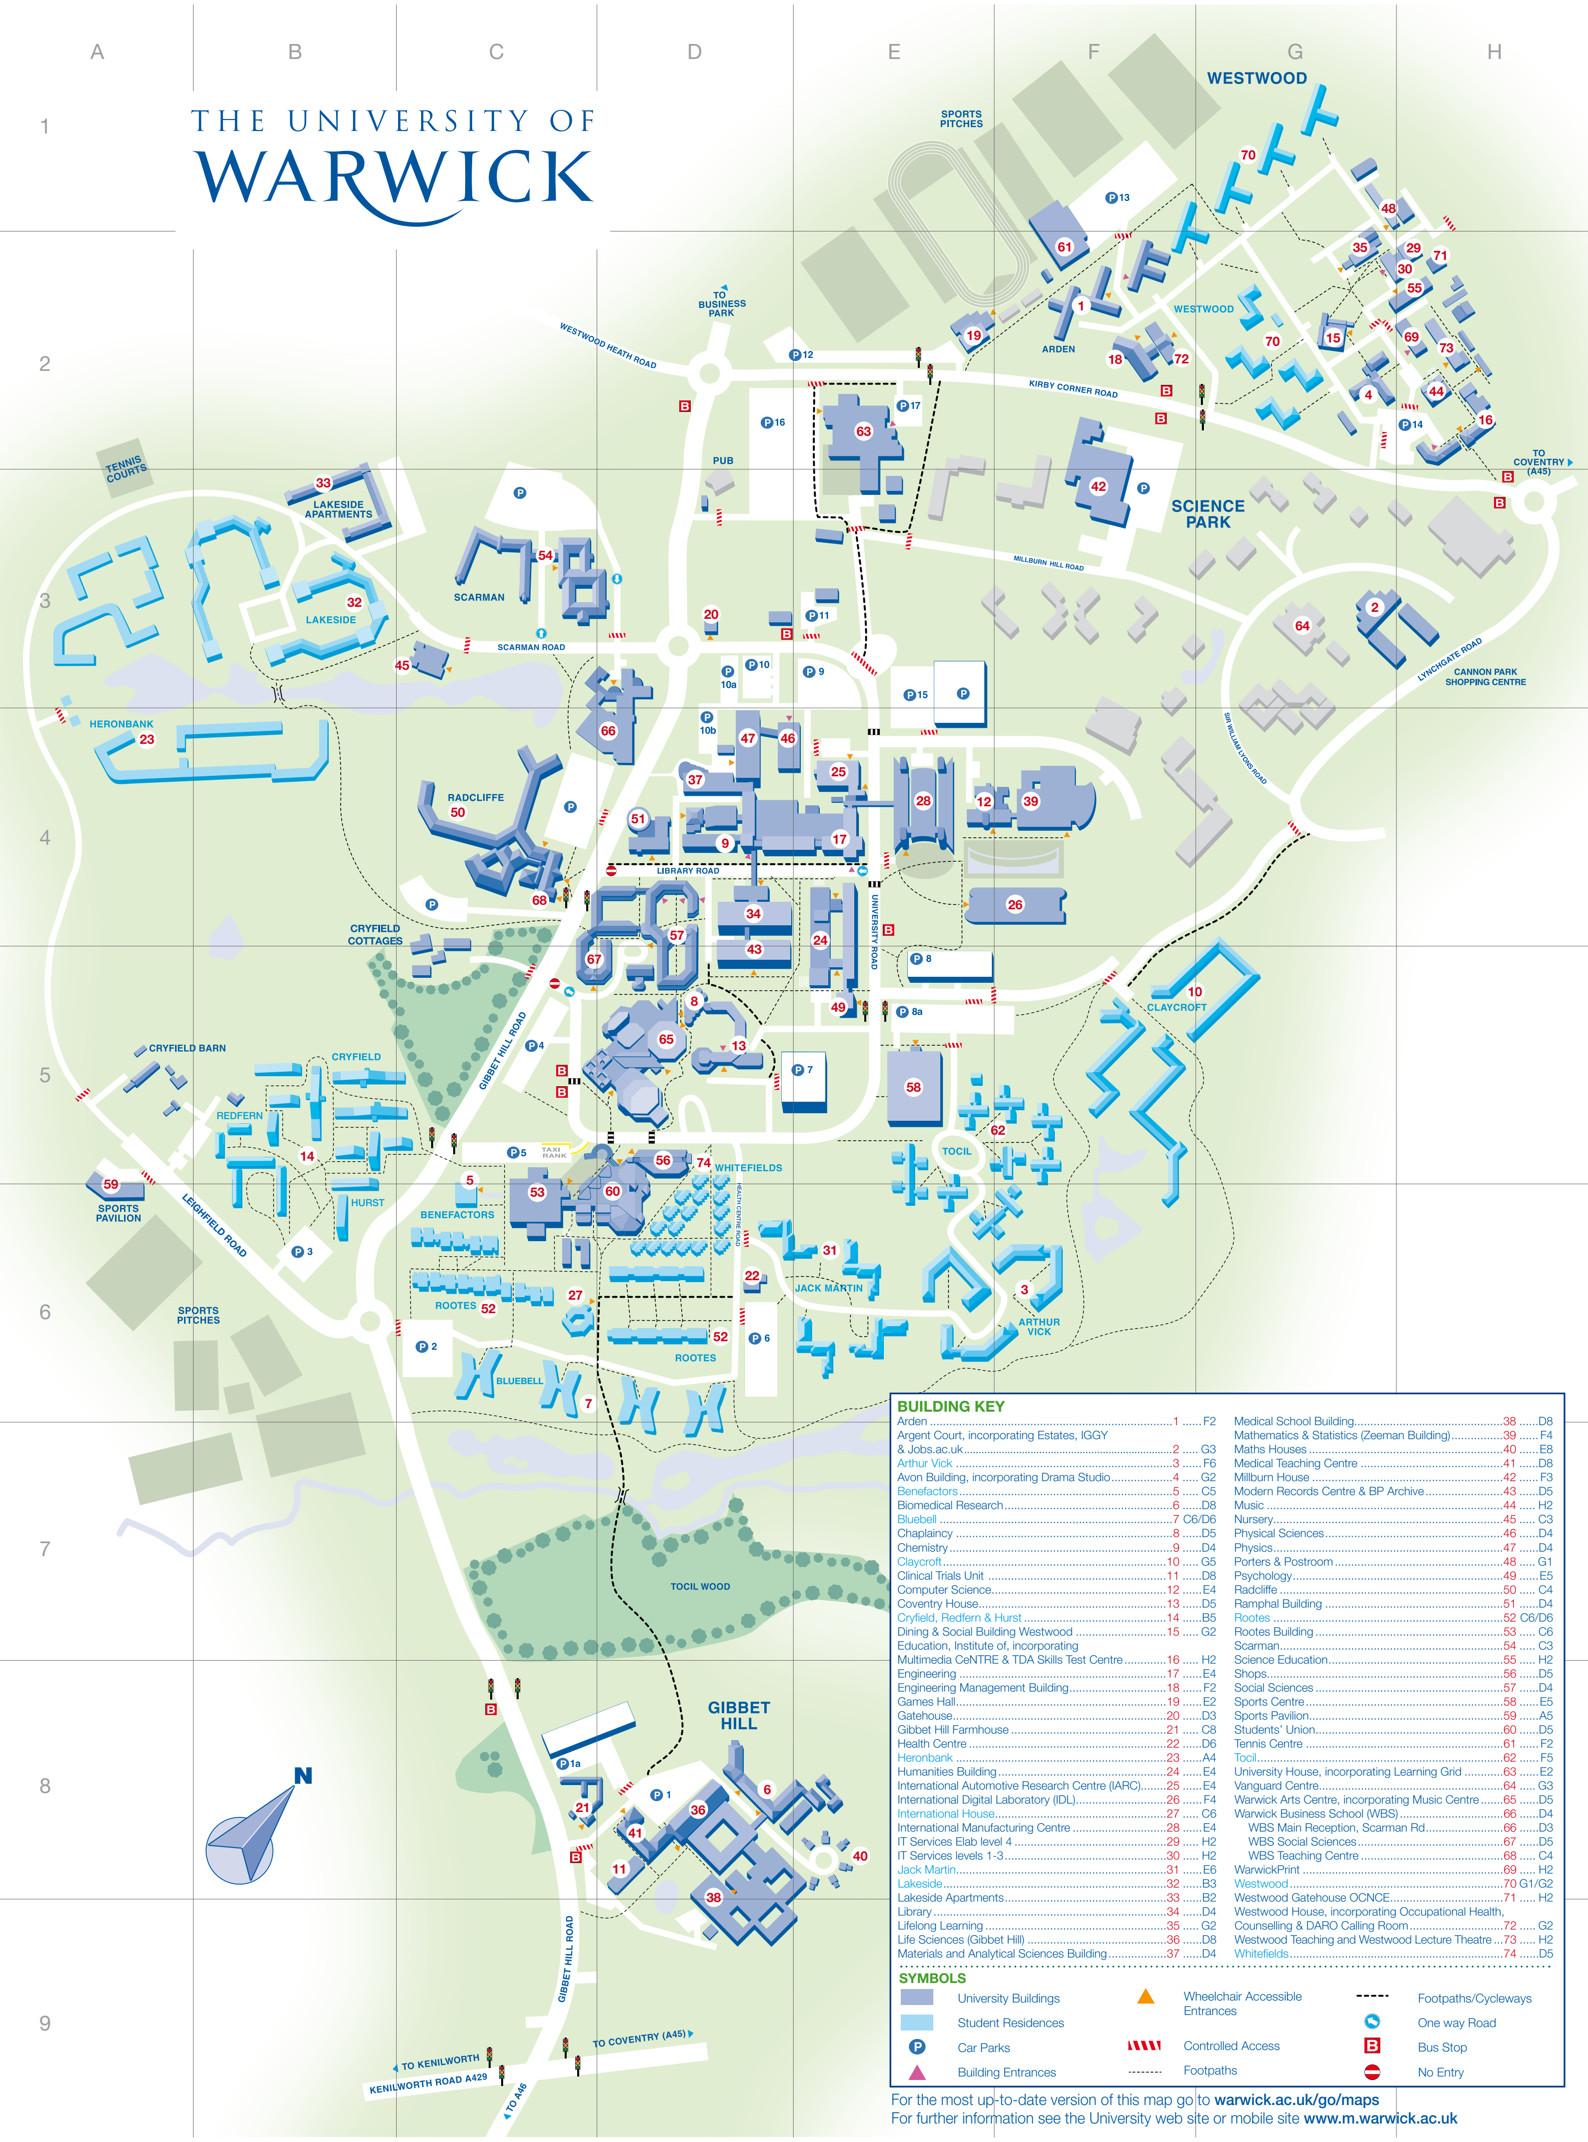

BUILDING KEY

Arden	1	F2
Argent Court, incorporating Estates, IGGY & Jobs.ac.uk	2	G3
Arthur Vick	3	F6
Avon Building, incorporating Drama Studio	4	G2
Benefactors	5	C5
Biomedical Research	6	D8
Bluebell	7	C6/D6
Chaplaincy	8	D5
Chemistry	9	D4
Claycroft	10	G5
Clinical Trials Unit	11	D8
Computer Science	12	E4
Coventry House	13	D5
Cryfield, Redfern & Hurst	14	B5
Dining & Social Building Westwood	15	G2
Education, Institute of, incorporating Multimedia CeNTRE & TDA Skills Test Centre	16	H2
Engineering	17	E4
Engineering Management Building	18	F2
Games Hall	19	E2
Gatehouse	20	D3
Gibbet Hill Farmhouse	21	C8
Health Centre	22	D6
Heronbank	23	A4
Humanities Building	24	E4
International Automotive Research Centre (IARC)	25	E4
International Digital Laboratory (IDL)	26	F4
International House	27	C6
International Manufacturing Centre	28	E4
IT Services Elab level 4	29	H2
IT Services levels 1-3	30	H2
Jack Martin	31	E6
Lakeside	32	B3
Lakeside Apartments	33	B2
Library	34	D4
Lifelong Learning	35	G2
Life Sciences (Gibbet Hill)	36	D8
Materials and Analytical Sciences Building	37	D4
Medical School Building	38	D8
Mathematics & Statistics (Zeeman Building)	39	F4
Maths Houses	40	E8
Medical Teaching Centre	41	D8
Millburn House	42	F3
Modern Records Centre & BP Archive	43	D5
Music	44	H2
Nursery	45	C3
Physical Sciences	46	D4
Physics	47	D4
Porters & Postroom	48	G1
Psychology	49	E5
Radcliffe	50	C4
Ramphal Building	51	D4
Rootes	52	C6/D6
Rootes Building	53	C6
Scarman	54	C3
Science Education	55	H2
Shops	56	D5
Social Sciences	57	D4
Sports Centre	58	E5
Sports Pavilion	59	A5
Students' Union	60	D5
Tennis Centre	61	F2
Tocil	62	F5
University House, incorporating Learning Grid	63	E2
Vanguard Centre	64	G3
Warwick Arts Centre, incorporating Music Centre	65	D5
Warwick Business School (WBS)	66	D4
WBS Main Reception, Scarman Rd	66	D3
WBS Social Sciences	67	D5
WBS Teaching Centre	68	C4
WarwickPrint	69	H2
Westwood	70	G1/G2
Westwood Gatehouse OCNCE	71	H2
Westwood House, incorporating Occupational Health, Counselling & DARO Calling Room	72	G2
Westwood Teaching and Westwood Lecture Theatre	73	H2
Whitefields	74	D5

SYMBOLS

- University Buildings
- Student Residences
- Car Parks
- Building Entrances
- Wheelchair Accessible Entrances
- Controlled Access
- Footpaths
- Footpaths/Cycleways
- One way Road
- Bus Stop
- No Entry

For the most up-to-date version of this map go to **warwick.ac.uk/go/maps**
For further information see the University web site or mobile site **www.m.warwick.ac.uk**

Administrative Details

Workshop Registration

Registration on Wednesday will take place in the lobby of the Mathematics & Statistics Building (Map Key 39), between 10 am and 12 noon.

Getting Here

Information on getting to the University of Warwick from Coventry, as well as from other directions locally and further afield, can be found at `http://www.warwick.ac.uk/about/visiting/`

Accommodation

Accommodation is in en-suite rooms on campus in either the Arthur Vick or Jack Martin buildings. Keys and directions can be collected from the reception of the Rootes Social Building (Map Key 53). All rooms have linen and toiletries. Kitchen facilities are available although meals are provided throughout the workshop. Rooms will be available after 15:00 for check in. All bedrooms must be vacated by 9:30am on Friday 21st.

Car Parking

Workshop participants are invited to use conference car park 7, 8a or 15. Car parking is free of charge during the workshop in the conference car parks. If you park in one of these car parks, you will need a token to exit — these should be collected from Rootes reception.

Internet Access

Campus: Wireless access is most easily available via eduroam — `http://www.eduroam.org/` — which is supported across most of the Warwick campus (but not at the time of writing in the accommodation buildings). Speak to one of the organisers for details of other options.

Accommodation: *Wired* internet access is available in all bedrooms. Details of how to log onto the system will be displayed in each individual bedroom, but participants will need to bring their own Ethernet cable. Ethernet cables can be purchased from Costcutter on campus (Map Key 56).

Start.Warwick

The Start.Warwick app, available for iPads, iPhones and Android devices from `http://www2.warwick.ac.uk/insite/news/intnews2/start_warwick_app`, provides useful information on travel and an interactive map of the campus amongst other things.

Facilities

Supermarkets:	Costcutter (Map Key 56)
	Tesco (Map Grid Reference H4)
Food & Drink:	Dirty Duck (Map Key 60)

	Terrace Bar (Map Key 60)
	Varsity (Map Grid Reference D3)
Coffee Shops:	Curiositea (Map Key 60)
	Costa (Map Key 53)
	Arts Centre (Map Key 65)
Cinema:	Map Key 65
Theatre:	Map Key 65
Sports Centre:	Map Key 58
Health Centre:	Map Key 22
Pharmacy:	Map Key 60

Help, Information & Telephone Numbers

Department

Department of Statistics
University of Warwick
Gibbet Hill Road
Coventry
CV4 7AL

Telephone:	024 7657 4812
Fax:	024 7652 4532
Map Key:	**39**

Emergency Numbers

Emergency:	Internal - 22222; External - 024 7652 2222
Security:	Internal - 22083; External - 024 7652 2083
Organiser:	Internal - 50919; External - 024 7615 0919 (Adam Johansen)

Transport

Swift Taxis (Coventry):	024 7676 7676
Trinity Street Taxis:	024 7699 9999
National Rail Enquiries:	08457 484 950

Timetable

All activities will take place in the Mathematics & Statistics Building (Map Key 39), with talks in room MS.03 (signposted from lobby), unless otherwise stated.

Wednesday 19th September

Time	Speaker	Title	Pg
10:00	**Registration**	**-**	**-**
11:45	**Welcome**	**Welcome & Administrative Details**	**-**
12:00	**Nicolas Chopin**	**Particles as auxiliary variables: PMCMC, SMC^2, SMC-ABC, and all that jazz**	**9**
13:00	**Lunch**	**-**	**-**
14:00	Peter Bunch	Rao-Blackwellised Smoothers	8
14:30	Fredrik Lindsten	Ancestral Sampling for Particle Gibbs	14
15:00	Liangliang Wang	Phylogenetic tree construction using sequential Monte Carlo algorithms on posets	18
15:30	Nikolaus Schweizer	Non-asymptotic Error Bounds for Sequential MCMC Methods in Multimodal Settings	17
16:00	**Coffee Break**	**-**	**-**
16:30	Andreas Eberle	Quantitative bounds for Markov chain based Monte Carlo methods in high dimensions	10
17:00	Nick Whiteley	Efficiency and optimal sampling of particle filters	19
17:30	Éric Moulines	On the convergence of Island particle models	14
18:00	**Finish**	**-**	**-**
19:30	**Dinner**	**Rootes Social Building (Map Key 53)**	**-**

Thursday 20th September

Time	Speaker	Title	Pg
7:30	**Breakfast**	**Rootes Social Building (Map Key 53)**	**-**
9:00	**Pierre Del Moral**	**On the exponential concentration properties of sequential Monte Carlo methods**	**10**
10:00	François Le Gland	On the asymptotic variance of SMC algorithms with adaptive resampling	13
10:30	Dan Crisan	Particle Filters with Random Resampling Times	9
11:00	**Coffee Break**	**-**	**-**
11:30	Sinan Yildirim	Approximate Bayesian Computation for Maximum Likelihood Estimation in Hidden Markov Models	20
12:00	Sylvain Le Corff	Convergence of a particle based online estimation procedure in hidden Markov models	13
12:30	Murray Pollock	Filtering for discretely observed jump diffusions	15
13:00	**Lunch**	**Rootes Social Building (Map Key 53)**	**-**
14:30	Paul Fearnhead	Continuous-time Importance Sampling for Multivariate Diffusions	11
15:00	Omiros Papaspiliopoulos	On computable filters	14
15:30	Alexandre Bouchard-Côté	Applications of SMC to the analysis of partially observed jump processes	8
16:00	**Coffee Break**	**-**	**-**
16:30	Sylvain Rubenthaler	Perfect simulation algorithm of a trajectory under a Feynman-Kac law	16
17:00	Sumeet Singh	Parameter Estimation for Hidden Markov Models with Intractable Likelihoods	18
17:30	Pierre Jacob	Resampling on parallel architectures	12
18:00	**Finish**	**-**	**-**
19:30	**Conference Dinner**	**Sutherland Suite, Rootes Social Building (Map Key 53)**	**-**

Friday 21st September

Time	Speaker	Title	Pg
7:30	**Breakfast**	**Rootes Social Building (Map Key 53)**	**-**
9:00	**Peter Jan van Leeuwen**	**Efficient Particle Filters for high-dimensional systems**	**12**
10:00	Gareth Peters	Incorporating Importance Sampling to the Evaluation of Subexponential Compound Process Tails, Closed Form Asymptotic Tail Expansions and Risk Measures	15
10:30	Yan Zhou	Bayesian Model Comparison via Sequential Monte Carlo Samplers	20
11:00	**Coffee Break**	**-**	**-**
11:30	Krzysztof Łatuszyński	Stability of Continuous-time Importance Sampling for Multivariate Diffusions	13
12:00	Gareth Ridall	Online model selection for Bayesian Motor Unit Number Estimation (MUNE)	16
12:30	Nikolas Kantas	A particle method for approximating principal eigen-functions and related quantities	12
13:00	**Lunch**	**Rootes Social Building (Map Key 53)**	**-**
14:30	Darren Wilkinson	Particle MCMC for partially observed Markov processes	19
15:00	Andrew Golightly	Efficient particle MCMC for exact inference in stochastic biochemical network models through approximation of expensive likelihoods	11
15:30	Thomas Schön	Learning Wiener models	17
16:00	Christophe Andrieu	Properties of pseudo-marginal MCMC algorithms	8
16:30	**Poster Session**	**Wine and light refreshments will be provided**	**-**
18:30	**Finish**	**-**	**-**

Talk Abstracts

Properties of pseudo-marginal MCMC algorithms

Christophe Andrieu
University of Bristol

In this presentation we will explore various theoretical properties of pseudo-marginal algorithms and related sampling schemes which have been recently established.

Applications of SMC to the analysis of partially observed jump processes

Alexandre Bouchard-Côté[1]. Seong Hwan-Jun[1], Liangliang Wang[2]
University of British Columbia[1], *University of Western Ontario*[2]

Many problems in modern computational biology and related domains involve partially observed continuous time jump processes defined over a large or countably infinite space. One example is a biological sequence evolving over a generational continuum and undergoing mutations affecting its length and contents. A second example is the 3D configuration of a molecule following a folding pathway. In all of these cases, an important and challenging task is to simulate paths given observations at a finite set of time points. In this talk, I will describe how we used SMC and Particle MCMC (PMCMC) to attack some these problems. Thanks to SMC, the global simulation problem is efficiently broken into smaller problems equivalent to end-point conditioned simulations. These sub-problems are in turn approached using methods reminiscent of a sampling equivalent of informed search techniques. If time permits, I will also describe Entangled Monte Carlo (EMC), a randomized algorithm we have developed to efficiently parallelize expensive SMC calculations such as those described in the first half of the talk. Inspired by ideas from perfect simulation, EMC avoids the transmission of particles between nodes, and instead reconstructs them from the particle genealogy. We found that the computational gain from parallelization using EMC significantly outweighs the cost of particle reconstruction.

Rao-Blackwellised Smoothers

Peter Bunch[1], Simo Särkkä[2], Simon Godsill[1]
University of Cambridge[1], *Aalto University*[2]

We will discuss some recent developments in Rao-Blackwellised particle smoothers in which we are able to simulate efficiently from the joint marginal smoothing distribution of a partially non-Gaussian state space model. This is in contrast with other approaches which operate either in the complete state-space, or using only a partial marginalisation of the linear part of the state.

Particles as auxiliary variables: PMCMC, SMC2, SMC-ABC, and all that jazz

Nicolas Chopin[1], Pierre Jacob[2], Omiros Papaspiliopoulos[3], Sumeet Singh[4]
ENSAE ParisTech[1], *Université Paris-Dauphine*[2], *Universitat Pompeu Fabra*[3], *University of Cambridge*[4]

This talk is an introduction to both particle filtering and sampling algorithms based on auxiliary particle variables, such as PMCMC (Particle Markov chain Monte Carlo), SMC^2 (Sequential Monte Carlo Square), and related algorithms. I will start with a gentle description of particle filtering, a generic class of Monte Carlo methods that make it possible to analyse sequentially (i.e. to filter the states and so on) state-space models. Then I will turn to a probabilistic description of these particle filters; or more precisely a probabilistic description of the joint distribution of all the random variables generated in the course of the algorithm. I will emphasise how essential is this probabilistic description in order to derive exact auxiliary variables-based, through naive examples. Then I will describe a general framework to include particle auxiliary variables in either MCMC or importance sampling; the former leading to PMCMC, while the latter leading to SMC^2.

Particle Filters with Random Resampling Times

Dan Crisan, Olasunkanmi Obanubi
Imperial College London

Particle filters are numerical methods for approximating the solution of the filtering problem which use systems of weighted particles that (typically) evolve according to the law of the signal process. These methods involve a corrective/resampling procedure which eliminates the particles that become redundant and multiplies the ones that contribute most to the resulting approximation. The correction is applied at instances in time called resampling/correction times. Practitioners normally use certain overall characteristics of the approximating system of particles (such as the effective sample size of the system) to determine when to correct the system. As a result, the resampling times are random. However, in the continuous time framework, all existing convergence results apply only to particle filters with deterministic correction times. In this paper, we analyse (continuous time) particle filters where resampling takes place at times that form a sequence of (predictable) stopping times. We prove that, under very general conditions imposed on the sequence of resampling times, the corresponding particle filters converge. The conditions are verified when the resampling times are chosen in accordance to effective sample size of the system of particles, the coefficient of variation of the particles weights and, respectively, the (soft) maximum of the particles weights. We also deduce central-limit theorem type results for the approximating particle system with random resampling times.

On the exponential concentration properties of sequential Monte Carlo methods

Pierre Del Moral
Université Bordeaux I

We present some new concentration inequalities for sequential Monte Carlo methods.

We analyze different types of stochastic particle models, including particle profile occupation measures, genealogical tree based evolution models, particle free energies, as well as backward Markov chain particle models.

We illustrate these results with a series of topics related to computational physics and biology, stochastic optimization, signal processing and Bayesian statistics, and many other probabilistic machine learning algorithms.

Quantitative bounds for Markov chain based Monte Carlo methods in high dimensions

Andreas Eberle
Universität Bonn

The constants in most of the known rigorous error bounds for MCMC and SMC methods depend exponentially on the dimension of the underlying state space. This is sometimes due to a failure of the MC methods in high dimensions, whereas in other cases it is due to estimates that are highly dimension dependent. In this talk we consider two different mathematical approaches for deriving non-asymptotic upper bounds in high dimensions that in simple models behave well as the dimension goes to infinity.

For a Metropolis-adjusted Langevin algorithm with semi-implicit Euler proposals, we apply a coupling method to bound the Wasserstein distance from equilibrium after a finite number of steps. In the case of log-concave perturbations of Gaussian measures, the obtained estimates are fairly explicit and dimension independent in a sense to be specified. For sequential Monte Carlo samplers, functional inequalities can be applied to derive L^p error bounds for a given finite number of particles. If appropriate functional inequalities hold uniformly in the dimension (e.g. this is the case for product models), a number of particles and steps that are each of order $O(d)$ are sufficient to attain a given error bound.

Continuous-time Importance Sampling for Multivariate Diffusions

Paul Fearnhead[1], Gareth Roberts[2], Giorgos Sermaidis[1], Krzysztof Łatuszyński[2]
Lancaster University[1], *University of Warwick*[2]

Inference for multivariate diffusion processes is challenging due to the intractability of the dynamics of the process. Most methods rely on high frequency imputation and discrete-time approximations of the continuous-time model, leading to biased inference. Recently, methods that are able to perform inference for univariate diffusions which avoid time-discretisation errors have been developed. However these approaches cannot be applied to general multivariate diffusions.

Here we present a novel, continuous-time Importance Sampling method that enables inference for general continuous-time Markov processes whilst avoiding time-discretisation errors. The method can be derived as a limiting case of a discrete-time sequential importance sampler, and uses ideas from random-weight particle filters, retrospective sampling and Rao-Blackwellisation.

Efficient particle MCMC for exact inference in stochastic biochemical network models through approximation of expensive likelihoods

Andrew Golightly
Newcastle University

Recently proposed particle MCMC methods provide a flexible way of performing Bayesian inference for parameters governing stochastic kinetic models defined as Markov (jump) processes (MJPs). Each iteration of the scheme requires an estimate of marginal likelihood calculated from the output of a sequential Monte Carlo scheme (also known as a particle filter). Consequently, the method can be extremely computationally intensive. We therefore aim to negate many instances of the expensive likelihood calculation through use of a fast approximation. We consider two approximations: the chemical Langevin equation (CLE) or diffusion approximation and a linear noise approximation (LNA). Either an estimate of marginal likelihood under the CLE, or the tractable marginal likelihood under the LNA can be used to calculate a first step acceptance probability. Only if a proposal is accepted under the approximation do we then run a sequential Monte Carlo scheme to compute an estimate of the marginal likelihood under the true MJP and construct a second stage acceptance probability that permits exact (simulation based) inference for the MJP. We therefore avoid expensive calculations for proposals that are likely to be rejected. We illustrate the method by considering inference for parameters governing a Lotka-Volterra system and a simple model of gene expression.

Resampling on parallel architectures

Pierre Jacob[1], Pierre Del Moral[2], Lawrence Murray[3], Gareth Peters[4]
Université Paris-Dauphine[1], Université Bordeaux I[2], CSIRO[3], University College London[4]

In SMC algorithms the resampling steps typically constitute the bottleneck towards full parallelization, since they require synchronization of all the cores and collecting all the weights. I will discuss a resampling step which relies on multiple pair-wise interactions instead of full synchronous interactions, and hence makes it more suitable for implementation on parallel computing architecture or possibly cloud computing. The idea is to embed the discrete time state space model under study into a continuous time model where the potential function is piece-wise constant, and then to approximate a resampling scheme that is classical in the continuous time filtering literature and more suitable to parallel computing.

Efficient Particle Filters for high-dimensional systems

Peter Jan van Leeuwen
University of Reading

Particle filters have long been considered unsuitable in high-dimensional systems due to the so-called 'curse of dimensionality'. In this talk I will show that it is not so much the dimension of the system, but the number of independent observations that is the problem. Furthermore, a solution is presented in which the deterministic part of the proposal density is used to position the majority of particles such that they have equal weights. The stochastic part of the proposal density is taken of small amplitude to allow only minimal changes to these weights. Finally this new scheme is applied to a 65,000 dimensional highly nonlinear fluid dynamical system using only 32 particles. It will be shown that the scheme is robust using several performance measures.

A particle method for approximating principal eigen-functions and related quantities

Nikolas Kantas[1], Nick Whiteley[2]
University College London[1], University of Bristol[2]

Perron-Frobenius theory treats the existence of a positive eigen-vector associated with the principal eigen-value $\lambda_\star$ of a non-negative matrix, say Q . A simple method for approximating this eigen-vector involves computing the iterate $\lambda_\star^{-n}Q^{(n)}$, for large n. In the more general case that Q is a non-negative integral kernel, an extended Perron-Frobenius theory applies, but it is typical that neither the principal eigen-function nor the iterate $\lambda_\star^{-n}Q^{(n)}$ can be computed exactly. In this setting we propose and study an interacting particle algorithm which yields a numerical approximation of the principal eigen-function and the associated twisted Markov kernel. We study a collection of random integral operators underlying the algorithm, address some of their mean and path-wise properties, and obtain L_r error estimates. Examples are provided in the context of a classical neutron model studied by Harris, a Bellman optimality equation and a rare event estimation problem. For the rare event problem we show how the proposed algorithm allows unbiased approximation of a Markov importance sampling method by conditional simulation.

Stability of Continuous-time Importance Sampling for Multivariate Diffusions

Krzysztof Łatuszyński[1], Paul Fearnhead[2], Gareth Roberts[1], Giorgos Sermaidis[2]
University of Warwick[1], *Lancaster University*[2]

I will present stability results of the continuous-time importance sampling (CIS) scheme for irreducible diffusions (presented in the talk by Paul Fearnhead) that includes unbiasedness and L^p stability of weights. Uniform ellipticity of the Kolmogorov diffusion operator and polynomial growth of coefficients is assumed and the results present an interplay between smoothness of drift and diffusion coefficients and probabilistic implementation parameters of the CIS algorithm.

Convergence of a particle based online estimation procedure in hidden Markov models

Sylvain Le Corff
Télécom ParisTech

Online variants of the Expectation Maximization (EM) algorithm have recently been proposed to perform parameter inference with large data sets or data streams in independent latent models and in hidden Markov models. Nevertheless, the convergence properties of these algorithms remain an open problem in the hidden Markov case. This contribution deals with a new online EM algorithm which updates the parameter at some deterministic times. Some convergence results are derived even in the case of hidden Markov models. These properties rely on the assumption that some intermediate quantities are available in closed form or can be approximated by Monte Carlo methods when the Monte Carlo error vanishes rapidly enough. We propose an algorithm which approximates these quantities using Sequential Monte Carlo methods.

On the asymptotic variance of SMC algorithms with adaptive resampling

François Le Gland
INRIA Rennes

Popular and efficient SMC algorithms use adaptive strategies based on heuristic considerations, where resampling is performed only when some criterion (effective number of particles, entropy of the sample, etc.) reaches a prescribed threshold. Del Moral, Doucet and Jasra (2012) have recently studied such adaptive resampling strategies, and have obtained a CLT as the sample size goes to infinity. The contribution of this talk is to provide some alternate, and hopefully practical, expression for the asymptotic variance. Special attention is given to two simple toy examples: rare event simulation, where importance weights can take only 0/1 values, and filtering of linear Gaussian systems.

Ancestral Sampling for Particle Gibbs

Fredrik Lindsten
Linköpings Universitet

We present a novel method in the family of particle MCMC methods that we refer to as particle Gibbs with ancestor sampling (PG-AS). Similarly to the existing PG with backward simulation (PG-BS) procedure, we use backward sampling to (considerably) improve the mixing of the PG kernel. Instead of using separate forward and backward sweeps as in PG-BS, however, we achieve the same effect in a single forward sweep. We apply the PG-AS framework to the challenging class of non-Markovian state-space models. We develop a truncation strategy of these models that is applicable in principle to any backward-simulation-based method, but which is particularly well suited to the PG-AS framework. In particular, as we show in a simulation study, PG-AS can yield an order-of-magnitude improved accuracy relative to PG-BS due to its robustness to the truncation error. Several application examples are discussed, including Rao-Blackwellized particle smoothing and inference in degenerate state-space models.

On the convergence of Island particle models

Éric Moulines[1], Cyrille Dubarry[1], Randal Douc[1], Pierre Del Moral[2]
Télécom ParisTech[1], *Université Bordeaux I*[2]

Recently, the development of parallel computing methods lead to the study of parallel implementation of this particle approximation. In this paper we focus on the so-called *island particle models*, consisting in running N_2 interacting particle systems each of size N_1. This problem gives rise to several questions among which the optimization of the size of the islands N_1 compared to the number of islands N_2 for a given total computing cost $N := N_1N_2$ and the opportunity of letting the islands interact. When the N_2 islands are run independently, the bias induced in each of them only depends on their population size N_1; thus it can be interesting to introduce an interaction between the islands even though this interaction increases the final estimator variance through the sampling steps. We propose here to study the asymptotic bias and variance in each case so that when N_1 and N_2 are large and fixed, we can determine which choice is the best in terms of asymptotic mean squared error.

On computable filters

Omiros Papaspiliopoulos[1], Matteo Ruggiero[2]
Universitat Pompeu Fabra[1], *University of Torino*[2]

I will review previous work and some current of my own (with Matteo Ruggiero, Turin) on computable filters, in the sense that the update and prediction operations are dealt with by means of a finite (albeit increasing along the iterations) computation.

Incorporating Importance Sampling to the Evaluation of Subexponential Compound Process Tails, Closed Form Asymptotic Tail Expansions and Risk Measures

Gareth Peters
University College London

The framework of Loss Distributional Approach modelling for loss process involving compound processes comprised of individual models for the severity of the loss and the frequency of loss occurrences in a period of time for single and multiple risks is the most widely utilised approach to risk and insurance modelling. This talk gives an overview of a class of closed form LDA models that are developed for modelling under Basel III in the context of low frequency, high severity loss processes. Typically, in such settings one my restrict to the class of severity models in the sub-exponential family. Within this context, several families of closed form LDA annual loss distributions, densities, survival functions are proposed. Then under these models, closed form approximations to risk measures such as VaR, Expected Shortfall and Insurance mitigated variants via haircuts or stochastic stop-loss premiums are developed. These results are obtained via a novel combination of the properties of the convolutional semi-group formed for such models and knowledge of the Fractional Lower Order Moments. The closed form solutions will be compared to alternative stochastic quantile and tail expectation results recently developed and known as Single Loss Approximations as well as results obtained for VaR estimation under assumptions suitable for Extreme Value Theory large deviation results to hold. In terms of Monte Carlo estimates to compare results, standard Monte Carlo and numerical solutions via Importance Sampling for Panjer Recursions (Volterra Integral Equations) are developed and compared. The aim of this is to carefully detail when each approximation is suitable for utilisation under the Basel III regulatory framework.

Filtering for discretely observed jump diffusions

Murray Pollock, Adam Johansen, Gareth Roberts
University of Warwick

In this talk we will introduce a novel algorithm for filtering discretely observed jump diffusions. Typically a jump diffusion transition density would be intractable, necessitating some form of approximation or discretisation.

Instead, we extend recent methodology introduced by Casella and Roberts [Casella, B., Roberts, G.O., 2011. Exact Simulation of Jump-Diffusion Processes with Monte Carlo Applications. Methodology and Computing in Applied Probability, 13-3, 449-473] which enables finite dimensional representations of jump diffusion sample paths to be simulated exactly.

In addition to a brief outline of the methodology we will also discuss a number of interesting, novel, Rao-Blackwellisations which arise and we will consider some illustrative examples.

Online model selection for Bayesian Motor Unit Number Estimation (MUNE)

P.G. Ridall[1], S. Taylor[1], C. Sherlock[1], A.N. Pettitt[2], D. Henderson[3], P.A. McCombe[3,4], C. Thomas[5]

Lancaster University[1], *Queensland University of Technology*[2], *Royal Brisbane and Women's Hospital*[3], *University of Queensland*[4], *University of Miami*[5]

Ridall et al. (2007) showed how estimation of the number of functioning motor units can be carried out by trans-dimensional simulation using reversible jump Markov Chain Monte Carlo (Green 1995) However this approach relies on convergence of a complex Markov chain, which is questionable when the dimension of the model is large. We propose a state space model that relies on sequential marginalisation of states and parameters over a particle approximation to the system. The state process is driven by an increasing stimulus activating an increasing number of units. We use a discrete approximation and forward recursions to marginalise the state process. Conditional on these states, we model the measurements of action potential by a Normal-Gamma observation equation, represented by four sufficient statistics for each particle. These are updated deterministically as new firing states are realized through the reampling process. Furthermore, the predictive distribution is available in terms of the statistics in closed form. Our simulations have shown that the success of this algorithm depends on the ability to retain these firing histories for as long as possible. The number of permutations of firing patterns is at its greatest in the middle of the experiment, when the greatest number of units are firing stochastically. Like Fearnhead and Liu (2007) we use a two stage stratified resampling process. Here the number of particles is varies over time according to the current complexity of the model. We validate our methods by carrying out estimation on simulated data and demonstrate its use by applying it to clinical data taken from a patient with spinal injury.

Perfect simulation algorithm of a trajectory under a Feynman-Kac law

Sylvain Rubenthaler[1], Christophe Andrieu[2], Arnaud Doucet[3]

Université Nice Sophia Antipolis[1], *University of Bristol*[2], *University of Oxford*[3]

I will present an algorithm for perfect simulation of trajectories under a Feynman-Kac law. This algorithm is a combination of simulation from the past and Metropolis. One interesting application is the law of directed polymers.

Learning Wiener models

Thomas Schön
Linköpings Universitet

The Wiener model consists of a linear dynamical system followed by a static nonlinearity. Learning this model structure in its most general case presents some interesting challenges that we consider in this talk. The maximum likelihood solution (via expectation maximisation) requires nontrivial integrals to be solved, where the particle smoother presents a good way forward. The Bayesian solution is obtained using a particle Gibbs (one of the members in the PMCMC family) algorithm with backward simulation. In exploiting a nonparametric representation (Gaussian process) of the nonlinearity and using an automatic order determination of the linear dynamical system we obtain a truly data driven method.

Non-asymptotic Error Bounds for Sequential MCMC Methods in Multimodal Settings

Nikolaus Schweizer
Universität Bonn

We consider Sequential MCMC methods, i.e., Sequential Monte Carlo applied in the context of stochastic numerical integration for target measures μ which cannot feasibly be attacked directly with standard MCMC methods due to the presence of multiple well-separated modes. The basic idea is to approximate the target distribution μ with a sequence of distributions $\mu_0, \ldots, \mu_n$ such that $\mu_n = \mu$ is the actual target distribution and such that μ_0 is easy to sample from. This sequence of distribution is approximated using Sequential MCMC. Our results center around two main questions: 1) Under which conditions can the smoothing effect of the MCMC steps balance the additional variance introduced into the system through the importance sampling resampling step? 2) Under which conditions does the particle dynamics work well in multimodal settings where conventional MCMC methods are trapped in local modes? We address both questions by proving suitable non-asymptotic error bounds which depend explicitly on a) an upper bound on relative densities, b) constants associated with global or local mixing properties of the MCMC dynamics (hyper-boundedness and spectral gap), and c) the amount of probability mass shifted between effectively disconnected components of the state space as we move from μ_0 to μ_n.

Preprint: arXiv *1205.6733v1*

Parameter Estimation for Hidden Markov Models with Intractable Likelihoods

Sumeet Singh
University of Cambridge

Approximate Bayesian computation (ABC) is a popular technique for approximating likelihoods and is often used in parameter estimation when the likelihood functions are analytically intractable. Although the use of ABC is widespread in many fields, there has been little investigation of the theoretical properties of the resulting estimators. In this paper we give a theoretical analysis of the asymptotic properties of ABC based maximum likelihood parameter estimation for hidden Markov models. In particular, we derive results analogous to those of consistency and asymptotic normality for standard maximum likelihood estimation. We also discuss how Sequential Monte Carlo methods provide a natural method for implementing likelihood based ABC procedures.

Preprint: http://arxiv.org/abs/1103.5399v1

Phylogenetic tree construction using sequential Monte Carlo algorithms on posets

Liangliang Wang[1], Alexandre Bouchard-Côté[2], Arnaud Doucet[3]
University of Western Ontario[1], *University of British Columbia*[2], *University of Oxford*[3]

A main task in evolutionary biology is phylogenetic tree reconstruction which determines the ancestral relationships among different species based on observed molecular sequences, e.g. DNA data. When a stochastic model, typically Continuous Time Markov Chain (CTMC), is used to describe the evolution, the phylogenetic inference depends on unknown evolutionary parameters. Bayesian inference provides a general framework for phylogenetic analysis, able to implement complex models of sequence evolution and to provide a coherent treatment of uncertainty for the groups on the tree. However, the application of Bayesian methods to large scale phylogenetics problems is increasingly limited by computational issues, motivating the development of methods that can complement existing Markov Chain Monte Carlo (MCMC) schemes. Sequential Monte Carlo (SMC) methods have been recently developed to address phylogenetic inference problems but currently available techniques are only applicable to a very restricted class of phylogenetic tree models compared to MCMC. We present a novel SMC method to approximate posterior phylogenetic tree distributions which is applicable to a general class of models and can be easily combined with MCMC to infer evolutionary parameters. Our method only relies on the existence of a flexible partially ordered set (poset) structure and is more generally applicable to sampling problems on combinatorial spaces. We demonstrate that the proposed SMC algorithm provides consistent estimates under weak assumptions, is computationally fast and is additionally easily parallelizable.

Efficiency and optimal sampling of particle filters

Nick Whiteley[1] and Anthony Lee[2]
University of Bristol[1] *and University of Warwick*[2]

This talk reports on an investigation of alternative sampling laws for particle algorithms and the influence of these laws on the efficiency of particle approximations of marginal likelihoods in hidden Markov models.
Amongst a very broad class of particle system transition operators, we characterise an essentially unique transition which is optimal relative to an asymptotic-in-time variance growth rate criterion. The sampling structure of the algorithm defined by this optimal transition turns out to be only quite subtly different from standard algorithms and yet the fluctuation properties of the estimates it provides are, in some ways, dramatically different.
The structure of the optimal transition suggests a new class of algorithms, which we term "twisted" particle filters, and which we validate with asymptotic analysis of a more traditional nature, in the regime where the number of particles tends to infinity.

Particle MCMC for partially observed Markov processes

Darren Wilkinson
Newcastle University

A number of interesting statistical applications require the estimation of parameters underlying a nonlinear multivariate continuous time Markov process model, using partial and noisy discrete time observations of the system state. Bayesian inference for this problem is difficult due to the fact that the discrete time transition density of the Markov process is typically intractable and computationally intensive to approximate. Nevertheless, it is possible to develop particle MCMC algorithms which are exact, provided that one can simulate exact realisations of the process forwards in time. Such algorithms, often termed "likelihood free" or "plug-and-play" are very attractive, as they allow separation of the problem of model development and simulation implementation from the development of inferential algorithms. Such techniques break down in the case of perfect observation or high-dimensional data, but more efficient algorithms can be developed if one is prepared to deviate from the likelihood free paradigm, at least in the case of diffusion processes. The methods will be illustrated using examples from population dynamics and stochastic biochemical network dynamics.

Approximate Bayesian Computation for Maximum Likelihood Estimation in Hidden Markov Models

Sinan Yildirim[1], Tom Dean[1], Sumeetpal S. Singh[1], Ajay Jasra[2]
University of Cambridge[1], National University of Singapore[1]

In this work, we present methodology for implementing maximum likelihood estimation (MLE) in hidden Markov models (HMMs) with intractable likelihoods in the context of approximate Bayesian computation (ABC). We show how both batch and online versions of gradient ascent MLE and expectation maximisation (EM) algorithms can be used for those HMMs using the ABC approach to confront the intractability. We demonstrate the performance of our methods first with examples on estimating the parameters of two intractable distributions, which are the -stable and g-and-k distributions, and then with an example on estimating the parameters of the stochastic volatility model with α-stable returns. Key words: approximate Bayesian computation, hidden Markov models, maxi- mum likelihood estimation, gradient ascent MLE, expectation maximisation.

Bayesian Model Comparison via Sequential Monte Carlo Samplers

Yan Zhou
University of Warwick

Bayesian model comparison for the purpose of selection, averaging and validation has attracted considerable attentions in the past few decades. Various Monte Carlo methods have been widely used in this context for calculation of the posterior probabilities. Sequential Monte Carlo (SMC) samplers are a class of Monte Carlo algorithms well suited for this purpose. In this work, we outline strategies of using SMC for Bayesian model comparison. Further the SMC algorithm can be used together with the path sampling estimator in a straightforward way. In addition, we propose adaptive strategies for both the specification of SMC sequence of distributions and MCMC kernels. They can improve the results considerably with little computational cost. The robustness and efficiency of the SMC algorithms for Bayesian model comparison, and the improvement of adaptive strategies are studied empirically. It is also shown that SMC algorithms is much more scalable compared to some other contemporary methods on modern parallel hardware.

Poster Abstracts

Metropolis-Hastings Implementations of the Forward-Filtering-Backward-Sampling Algorithm for Particle Smoothing

Peter Bunch
University of Cambridge

The conventional forward-filtering-backward-sampling algorithm for particle smoothing generates samples from the joint smoothing distribution by sequentially sampling backwards from the forward filter approximations. This algorithm has a computational complexity quadratic in the number of particles, which is often restrictive. In addition, the support of the smoothing distribution is restricted to the values which appear in the filtering approximation. Here a Metropolis-Hastings sampling procedure is used to improve the efficiency of the forward-filtering-backward-sampling algorithm, achieving comparable error performance but with a lower execution time. In addition, an algorithm for approximating the joint smoothing distribution without limited support is presented, which achieves simultaneous improvements in both execution time and error. These algorithms also provide a greater degree of flexibility over existing methods, allowing a trade-off between execution time and estimation error, controlled by the length of the Markov chains.

Multiple Auxiliary Particle Filter (MAPF)

Xi Chen
University of Cambridge

Electromagnetic source localization is a technique that enables the studies of neural dynamical activities on a millisecond timescale using MEG/EEG data. It aims to reveal neural activities in cortical area which can not be seen with imaging methods such as fMRI. We model the problem under a Bayesian multi-targets tracking framework, a Multiple Auxiliary Particle Filter (MAPF) is developed to estimate the dipolar source dynamics on a millisecond timescale. Simulation results and analysis of the results for both single dipole and two dipoles scenarios are shown.

Sequential Monte Carlo for dynamic image-based lighting

Johan Dahlin
Linköpings Universitet

Image-based lighting is an approach for illuminating computer generated objects inserted in real-world environments. The lighting of the object is governed by the light transport equation, which is an integral taking into account the light from the environment hitting the object as well as the optical and geometrical properties of the object. Traditionally, this integral is solved by importance sampling for each frame independently. In this work, we consider the use of particle filters/smoothers to solve the light transport equation when the environment map changes with time. This approach is often a more robust and computationally cheaper alternative to the traditional methods.

Inference for Piecewise Deterministic Processes via Sequential Monte Carlo Methods

Axel Finke
University of Warwick

We discuss filtering and fixed-parameter estimation for discretely-observed piecewise-deterministic processes (PDPs) via sequential Monte Carlo methods. PDPs are stochastic processes that jump randomly at a countable number of stopping times but otherwise evolve deterministically in continuous time.

For instance, PDPs can be used to model the trajectories of fighter aircraft. Here, a pilot accelerates at random times which correspond to jumps in the acceleration dynamics. Between such jumps, the trajectory of the aircraft is just a deterministic function of the aircraft's position, speed, and acceleration at the most recent jump.

We apply the sequential Monte Carlo filter for PDPs, introduced in Whiteley, Johansen and Godsill (2011), to this model. In addition, we develop particle MCMC methods for fixed-parameter estimation in this context.

Inference in CAR(1) processes driven by alpha-stable Levy processes

Tatjana Lemke
University of Cambridge

We present Poisson sum series representations (PSSR) for alpha-stable random variables and alpha-stable processes. In particular we focus on continuous-time autoregressive (CAR) models driven by alpha-stable Levy processes. Our representations aim to provide a conditionally Gaussian framework, which allows parameter estimation using Rao-Blackwellised versions of state of the art Bayesian computational methods such as particle filters and Markov chain Monte Carlo (MCMC).

Our modified Poisson sum series representation for stable stochastic integrals, which provides a Gaussian framework conditioned on two sets $\{\Gamma_i\}_{i=1}^{\infty}$ and $\{V_i\}_{i=1}^{\infty}$ makes it possible to deal with the CAR(1) process. Then, the variables which need to be sampled are the stability index α, factors μ_W and σ_W^2 in mean and variance of the conditionally Gaussian framework and the above mentioned latent variables $\{\Gamma_i\}_{i=1}^{\infty}$ and $\{V_i\}_{i=1}^{\infty}$.

In practice the series representations need to be truncated, however to overcome the issues due to truncation of the series, novel residual approximations are developed. Simulations using a Rao-Blackwellised SMC, which treats μ_W as part of the state vector. demonstrate the potential of these Poisson sum series representations for inference in CAR(1) processes.

Online Inference for dynamical system: move-resample approach

Reinaldo Marques, Geir Storvik
University of Oslo

In particle filtering algorithms, the posterior $\pi(x_t|y_{1:t})$ (for simplicity we here assume the static parameters to be known) is approximated by a set of weighted samples $\{(x_t^i, w_t^i), i = 1, \ldots, M\}$. A problem with such methods is degeneracy problem in that either the weights have huge variability (typically with one or a few samples dominating in weights) or high correlations between the samples. Updating the samples by a few MCMC steps have been suggested as an improvement in this case. The general setup is to first resample the particles such that all particles are given equal weight (resample-move algorithm). Thereafter the MCMC steps are applied in order to make the identical samples diverge. In this work we consider an alternative strategy where the order of MCMC updates and the resampling steps are switched, i.e. MCMC updates are performed first. The main advantage with such an approach is that by performing MCMC updates, the weights can be updated simultaneously, making them less variable. Such an approach has not been considered earlier, supposedly due to the difficulty in updating the corresponding weights. We show some results using synthetic data.

Exact Changepoint Distributions and Sequential Monte Carlo samplers

Christopher Nam, John Aston, Adam Johansen
University of Warwick

Quantifying the uncertainty of changepoints is an important aspect of both theoretical and applied statistics. Several existing approaches provide different estimates, such as the number and locations of changepoints, and thus there is a need to assess their plausibility. This poster will review a methodology in quantifying the uncertainty of changepoints in light of parameter uncertainty. The methodology combines recent work on exact changepoint distributions via Hidden Markov Models (HMM) and Finite Markov Chain Imbedding (FMCI), and accounts for parameter uncertainty via Sequential Monte Carlo (SMC) samplers. In addition, SMC also allows us to account for the number of states in a HMM being ultimately unknown. The methodology will be applied to econometric data.

Filtering for discretely observed jump diffusions

Murray Pollock, Adam Johansen, Gareth Roberts
University of Warwick

In this poster we will outline some extensions to the work presented earlier in the workshop whereby we introduced a novel algorithm for filtering discretely observed jump diffusions.

Particle filtering for estimation of static dipoles in magnetoencephalogaphy

Alberto Sorrentino
University of Warwick

In MagnetoEncephaloGraphy (MEG) one looks for dynamic estimates of brain activity from non-invasive, millisecond-by-millisecond recordings of the magnetic fields produced by large neural populations inside the brain and measured outside the scalp. Under the dipole approximation, a single neural population is approximated as a point source (the dipole), characterized by a location, an orientation and a strength; in this case brain activity amounts to a time-varying set of point sources. The non-linear dependence of the data on the dipole location and the dynamic nature of the problem encourage application of particle filtering techniques. Particle filters are sequential Monte Carlo algorithms that produce samples approximately distributed according to the posterior distribution of an evolving unobserved system. In previous work we proposed a Sequential Importance Resampling (SIR) particle filter to estimate dynamically the number of sources and the source parameters from MEG data, achieving performances comparable to state-of-the-art methods, but with a notably higher level of automation. However, current dipoles were modelled as moving randomly within the brain, a key assumption for the SIR particle filter to work, but in conflict with the physiological interpretation of the current dipole as the activity of a particular neural population. Here we explicitly assume that the current dipoles do not move, and construct both a more realistic model and a more effective algorithm based on the Resample Move idea of Gilks and Berzuini - that approximates reasonably the posterior distribution of a time-varying dipole configuration under the assumption of stationary dipole locations.

Preprint: arXiv *1205.6310v1*

Bayesian Model Comparison via Sequential Monte Carlo Samplers and Path Sampling

Yan Zhou
University of Warwick

Model comparison for the purposes of selection, averaging and validation is a problem which is found throughout statistics. Within the Bayesian paradigm, these problems all require the calculation of the posterior probabilities of models within a particular class. Substantial progress has been made in recent years, but there are numerous difficulties in the practical implementation of existing schemes. The current work develops sequential Monte Carlo (SMC) sampling strategies to characterize the posterior model probabilities using either the standard estimator or path sampling estimators. It is demonstrated that these methods can be efficient in terms of Monte Carlo variance. In addition, some adaptive schemes for both the SMC distribution sequences and MCMC kernels are developed. This leads to an automatic algorithm with minimal manual tuning efforts while being robust and efficient. Also, the SMC algorithms are demonstrated to be very scalable and easy to implement when being parallelized. This is demonstrated with GPU and multi-core CPU implementations.

Participant List

Name	Affiliation	Email
Aggelis Alexopoulos	Athens University	alex-agel@hotmail.com
Javier Amezcua	University of Reading	j.amezcuaespinosa@reading.ac.uk
Christophe Andrieu	University of Bristol	c.andrieu@bristol.ac.uk
John Aston	University of Warwick	j.a.d.aston@warwick.ac.uk
Alex Beskos	University College London	a.beskos@ucl.ac.uk
Peter Bunch	University of Cambridge	pb404@cam.ac.uk
Alexandre Bouchard-Côté	University of British Columbia	bouchard@stat.ubc.ca
Xi Chen	University of Cambridge	xc253@cam.ac.uk
Nicolas Chopin	ENSAE, ParisTech	nicolas.chopin@ensae.fr
Dan Crisan	Imperial College London	d.crisan@imperial.ac.uk
Johan Dahlin	Linköpings Universitet	johan.dahlin@isy.liu.se
Pierre Del Moral	Université Bordeaux I	pierre.del-moral@inria.fr
Arnaud Doucet	University of Oxford	doucet@stats.ox.ac.uk
Andreas Eberle	Universität Bonn	eberle@iam.uni-bonn.de
Paul Fearnhead	Lancaster University	p.fearnhead@lancaster.ac.uk
Axel Finke	University of Warwick	A.Finke@warwick.ac.uk
Xiaoxing Fu	University of Leicester	xxf1@le.ac.uk
Simon Godsill	University of Cambridge	sjg@eng.cam.ac.uk
Andrew Golightly	Newcastle University	andrew.golightly@newcastle.ac.uk
Pierre Jacob	Université Paris-Dauphine	pierre.jacob@ensae.fr
Peter Jan van Leeuwen	University of Reading	p.j.vanleeuwen@reading.ac.uk
Lan Jiang	University of Cambridge	lj281@cam.ac.uk
Adam Johansen	University of Warwick	a.m.johansen@warwick.ac.uk
Nikolas Kantas	University College London	n.kantas@ucl.ac.uk
Robert Kohn	University of New South Wales	r.kohn@unsw.edu.au
Krzysztof Łatuszyński	University of Warwick	k.g.latuszynski@warwick.ac.uk
Anthony Lee	University of Warwick	Anthony.Lee@warwick.ac.uk
Tatjana Lemke	University of Cambridge	tl340@cam.ac.uk
Sylvain Le Corff	Télécom ParisTech	sylvain.lecorff@gmail.com
François Le Gland	INRIA Rennes	francois.legland@inria.fr
Fredrik Lindsten	Linköpings Universitet	lindsten@isy.liu.se
Reinaldo Marques	University of Oslo	ramarque@math.uio.no
Éric Moulines	Télécom ParisTech	eric.moulines@telecom-paristech.fr
Christopher Nam	University of Warwick	c.f.h.nam@warwick.ac.uk
Christopher Nemeth	Lancaster University	c.nemeth@lancaster.ac.uk
Emre Ozkan	Linköpings Universitet	emre@isy.liu.se
Omiros Papaspiliopoulos	Universitat Pompeu Fabra	omiros.papaspiliopoulos@upf.edu
Gareth Peters	University College London	garethpeters78@gmail.com
Murray Pollock	University of Warwick	murray.pollock@warwick.ac.uk
Peter Gareth Ridall	Lancaster University	g.ridall@lancaster.ac.uk
Sylvain Rubenthaler	Université Nice Sophia Antipolis	Sylvain.Rubenthaler@unice.fr
Thomas Schön	Linköpings Universitet	schon@isy.liu.se
Nikolaus Schweizer	Universität Bonn	nschweizer@iam.uni-bonn.de
François Septier	Institut Mines-Telecom / Telecom Lille1 / LAGIS UMR CNRS 8219	francois.septier@telecom-lille1.eu
Sumeetpal Singh	University of Cambridge	sss40@cam.ac.uk
Alberto Sorrentino	Università di Genova	sorrentino@dima.unige.it
Simon Taylor	Lancaster University	s.taylor2@lancaster.ac.uk
Liangliang Wang	University of Western Ontario	wang@stats.uwo.ca
Nick Whiteley	University of Bristol	nick.whiteley@bristol.ac.uk
Darren Wilkinson	Newcastle University	darren.wilkinson@ncl.ac.uk
Kasia Wolny	University of Warwick	k.b.wolny@warwick.ac.uk
Sinan Yildirim	University of Cambridge	sy276@cam.ac.uk
Yan Zhou	University of Warwick	yan.zhou@warwick.ac.uk