List of projects

FMS020F–NAMS002 Statistical inference for partially observed stochastic processes, 2016

Work in groups of two (if this is absolutely not possible for some reason, please let the lecturers know in advance).

However do not form groups spontaneously, instead email `umberto@maths.lth.se` with your **individual** preferences for the listed projects. That is, specify which project is your absolute favourite, then which one is your second favourite option and finally your least favourite one. We'll use your preferences to form groups, and though we will try to respect your wishes we also need to avoid the situation where everyone is choosing the same project. In addition, you are invited to suggest a research project (notice this is entirely optional): this can only happen if (1) you are going to use some of the methods illustrated during the course and (2) your supervisor approves your proposal, which should be challenging enough, i.e. not material/examples that you have already "solved" during your studies. If you choose this option, your project proposal should be detailed in the materials description and goals and should be *emailed simultaneously to the course lecturers and CC to your supervisor.*

You are not required to write a project report. Prepare instead slides for a presentation (Powerpoint or similar, LaTeX Beamer etc.) and email those to the relevant lecturer **at least 24 hrs before the date of your presentation**.

The titles below report the reference person for the specific project, EL: Erik Lindström, JL: Johan Lindström, UP: Umberto Picchini.

1 Particle marginal methods and ABC – reference person UP

Particle filters (sequential Monte Carlo, SMC) are certainly useful when dealing with state space models. The underlying assumption is that even when the transition density $p(x_t|x_{t-1};\theta)$ is unknown in closed form, when using a specific SMC algorithm (bootstrap filter) we only need the ability to forward simulate from the state-model for $\{X_t\}$. However, at the very least it is required knowledge of the distribution for the observations model so that $p(y_t|x_t;\theta)$ can be evaluated for any t.

Here we consider the case where $p(y_t|x_t;\theta)$ is unavailable in closed form but we can somehow sample from it. An α-stable stochastic volatility model is an example of "intractable model". The model is defined as

$$\begin{cases} y_t = e^{x_t/2}v_t \\ x_t = a + bx_{t-1} + \sigma u_t \end{cases} \tag{1}$$

with u_t and v_t mutually independent and independent across t. Notice y_t is known as "log return" in financial literature. We assume $u_t \sim \mathcal{N}(0,1)$ and v_t has symmetric alpha-stable distribution with $\alpha \in (0,2]$, unit scale and zero location parameter, i.e. $v_t \sim \mathcal{S}_\alpha(1,0,0)$. In this case $p(y_t|x_t)$ is unknown, however we can easily simulate a realization from it (conditionally on x_t and θ) as we know how to simulate the v_t. If U is a uniformly distributed random variable on $[-\pi/2, \pi/2]$ and W is an independent exponential random variable with mean 1, then

$$Z = \begin{cases} \sin(\alpha U)[\cos(U)]^{-1/\alpha}\{W^{-1}\cos[(\alpha-1)U]\}^{\frac{1-\alpha}{\alpha}} & (\alpha \neq 1) \\ \tan(U), \quad (\alpha = 1) \end{cases} \tag{2}$$

has $\mathcal{S}_\alpha(1,0,0)$ distribution, e.g. Calvet and Czellar [2014] (notice you can disregard the specific methodology illustrated in Calvet and Czellar).

Unknowns are $\theta = (a, b, \sigma, \alpha)$ which we would like to estimate from data (either simulated or from a data file). Although Calvet and Czellar [2014] provide an ABC filter for such problem, an option is to use the (perhaps less advanced but easier to implement) filter by Jasra et al. [2012] which we discussed at lecture. Such filter is also given as Algorithm 2 in Picchini and Samson [2015] (but you can disregard the rest). Notice Jasra et al. [2012] also provide a nice scheme for selecting the ABC threshold ϵ_t, by picking at each time t a percentile of the distances $|y_t - y_t^{i*}|$ produced by the cloud of particles, $i = 1, ..., N$.

Using such filtering approach, construct[1] a pseudo-marginal MCMC routine to estimate θ, using a fixed (known) starting state $x_0 = 0$. You can use whatever kernel function you consider appropriate to evaluate the discrepancies $|y_t - y_t^{i*}|$, e.g. an identity function (as in Jasra et al. [2012]) and/or a Gaussian kernel (as in Picchini and Samson [2015]).

Some possibilities (but feel free to explore):

1. Produce a dataset $y_1, ..., y_T$ of length $T = 3000$ using $a = -0.0445$, $b = 0.996$, $\sigma = 0.0615$, $\alpha = 1.9874$ (numbers are from Table 2 in Calvet and Czellar).

2. Calvet and Czellar use a specific dataset retrieved from `http://www.federalreserve.gov/releases/h10/hist/dat00_uk.htm` then they selected log-returns from 3 January 2006 to 30 December 2011. There are tools to retrieve data via R (for example `http://www.quantmod.com/documentation/getSymbols.FRED.html` and `https://github.com/jcizel/FredR`) and then you can import those in the software you prefer.

3. you can try some other possibility of your interest.

2 PMCMC for discretely observed diffusions – reference person EL

The Pseudo Marginal Metropolis Hastings algorithm, see Andrieu and Roberts [2009] allows for unbiased estimates of the likelihood function, like the one produced by the Pedersen technique. This idea was used in Stramer et al. [2011] who applied PMMH to several different discretely observed models.

[1] Notice Jasra et al. [2012], and the results provided therein, only pertain the construction of a filter. They do not perform parameter estimation.

Implement at least one of their models, preferably using some importance sampler like the Durham-Gallant sampler, Durham and Gallant [2002], and combine it with for example adaptive MCMC, see Andrieu and Thoms [2008] or more advanced MCMC ideas taking advantage of gradient and Hessian information, see Dahlin et al. [2015]. Optimal balance between the Monte Carlo error of the estimate and the Markov chain was discussed in e.g Doucet et al. [2015].

Extra/voluntary: The fact that we no longer needs to optimize over the likelihood function opens up for new ideas with regards to approximate Bridge samplers. Please contact Erik (`erikl@maths.lth.se`) if this would be of interest to you.

3 Laplace approximation for diffusions – reference person EL

The simplest approximation of the transition density for an SDE defined by

$$\mathrm{d}X_t = \mu(t, X_t)\mathrm{d}t + \sigma(t, X_t)\mathrm{d}W_t \tag{3}$$

is the explicit (or implicit) Euler Maruyama (E-M) approximations

$$\begin{aligned} X_{t+\Delta} &= X_t + \mu(t, X_t)\Delta + \sigma(t, X_t)\Delta W_t && \text{Explicit} \\ X_{t+\Delta} &= X_t + \mu(t, X_{t+\Delta})\Delta + \sigma(t, X_{t+\Delta})\Delta W_t && \text{Implicit} \end{aligned}$$

We can combine these expressions

$$\begin{aligned} p(x_t|x_0) &= \int p(x_t|x_s)p(x_s|x_0)\mathrm{d}x_s && (4) \\ &\approx \int p^{Implicit\ E-M}(x_t|x_s)p^{Explicit\ E-M}(x_s|x_0)\mathrm{d}x_s && (5) \\ &= \phi(x_t; m_t + x_s, P_t)\phi(x_s; m_0, P_0) && (6) \end{aligned}$$

where $\phi(x, m, P)$ is a Gaussian density, $m_t = \mu(t, x_t)(t - s)$, $m_0 = x_0 + \mu(0, x_0)s$, $P_t = \sigma(t, x_t)\sigma(t, x_t)^T(t-s) = \Sigma(t, x_t)(t-s)$ and $P_0 = \sigma(0, x_0)\sigma(0, x_0)^T s = \Sigma(0, x_0)s$. The point is that we can solve this integral in closed form, cf. [Petersen et al., 2008, Eq. (371)], arriving at

$$\begin{aligned} p(x_t|x_0) &= \phi(x_t; m_0 + m_t, P_0 + P_t) \\ &= \phi\left(x_t; x_0 + \mu(0, x_0)s + \mu(t, x_t)(t - s), \Sigma(0, x_0)s + \Sigma(t, x_t)(t - s)\right) \end{aligned}$$

which turns out to be a mid point scheme!

A more advanced decomposition may not be possible to solve in closed form for general models, but it would be interesting to see how good a well designed Laplace approximation of Eq. (4) could be.

4 Multiresolution for SPDE-based GMRFs – reference person JL

In Kou et al. [2012] a multi-resolution method for diffusion processes is described. The goal of this project is to evaluate the method for use on the SPDEs

[Lindgren et al., 2011, Lindgren and Rue, 2008] where the multi-resolution is represented through nested triangular basis functions, i.e. nested linear B-splines (see figure). Another option is to use Daubechies wavelets [Bolin and Lindgren, 2013].

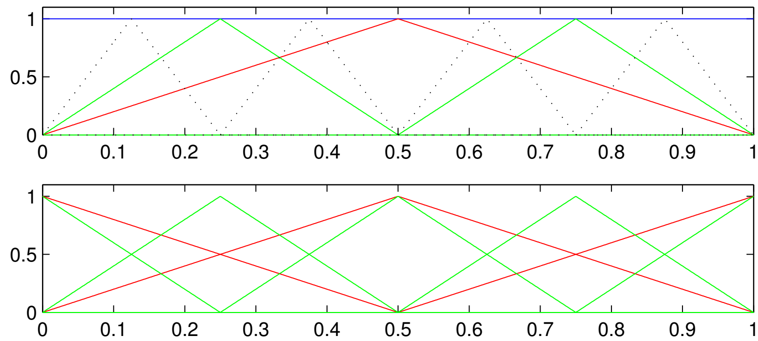

A suitable test cases is the time points of coal mine accidents which can be obtained as `data(coal)` from `library(boot)` in R. These data could be modelled as a log-Gaussian Cox process [Simpson et al., 2015] where the latent log-intensity is described using a multi-resolution solution to the spde

$$(\kappa^2 - \Delta)^{1/2} x(s) = \frac{1}{\sqrt{\tau}} W(s),$$

which approximates an AR(1)-process. To estimate the latent process and the unknown parameters (κ^2 and τ) the methods from [Kou et al., 2012] should be combined with adaptive MCMC [Andrieu and Thoms, 2008] for the parameters and MALA-like updates [Girolami and Calderhead, 2011] for the latent x-field.

References

Christophe Andrieu and Gareth O Roberts. The pseudo-marginal approach for efficient monte carlo computations. *The Annals of Statistics*, pages 697–725, 2009.

Christophe Andrieu and Johannes Thoms. A tutorial on adaptive mcmc. *Statistics and Computing*, 18(4):343–373, 2008.

David Bolin and Finn Lindgren. A comparison between markov approximations and other methods for large spatial data sets. *CSDA*, 61:7–21, 2013. doi: 10.1016/j.csda.2012.11.011.

L. Calvet and V. Czellar. Accurate methods for approximate Bayesian computation filtering. *Journal of Financial Econometrics*, 2014. doi: 10.1093/jjfinec/nbu019.

Johan Dahlin, Fredrik Lindsten, and Thomas B Schön. Particle metropolis–hastings using gradient and hessian information. *Statistics and computing*, 25 (1):81–92, 2015.

Arnaud Doucet, MK Pitt, George Deligiannidis, and Robert Kohn. Efficient implementation of markov chain monte carlo when using an unbiased likelihood estimator. *Biometrika*, page asu075, 2015.

Garland B Durham and A Ronald Gallant. Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business & Economic Statistics*, 20(3):297–338, 2002.

M. Girolami and B. Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *JRSSB*, 73:123–214, 2011. doi: 10.1111/j.1467-9868.2010.00765.x.

A. Jasra, S. Singh, J. Martin, and E. McCoy. Filtering via approximate bayesian computation. *Statistics and Computing*, 22(6):1223–1237, 2012.

S. Kou, P. Olding, M. Lysy, and J. Liu. A multiresolution method for parameter estimation of diffusion processes. *JASA*, 107:1558–1573, 2012.

F. Lindgren, H. Rue, and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *JRSSB*, 73(4):423–498, 2011. doi: 10.1111/j.1467-9868.2011.00777.x.

Finn Lindgren and Håvard Rue. On the second-order random walk model for irregular locations. *SJS*, 35:691–700, 2008.

Kaare Brandt Petersen, Michael Syskind Pedersen, et al. The matrix cookbook. *Technical University of Denmark*, 7:15, 2008.

U. Picchini and A. Samson. Coupling stochastic EM and approximate bayesian computation for parameter inference in state-space models. *arXiv:1512.04831*, 2015.

Daniel Simpson, Janine Illian, Finn Lindgren, Sigrunn Sorbye, and Håvard Rue. Going off grid: Computationally efficient inference for log-gaussian cox processes. *arXiv*, 1111.0641:n/a, 2015.

Osnat Stramer, Matthew Bognar, et al. Bayesian inference for irreducible diffusion processes using the pseudo-marginal approach. *Bayesian Analysis*, 6(2):231–258, 2011.