Institution: Imperial College London

Unit of Assessment: 10 Mathematical Sciences

Title of case study: C3 - Applications of Bayesian methods in finance, credit scoring and target tracking

1. Summary of the impact (indicative maximum 100 words)

In recent years there has been an explosion of real data from areas as diverse as bioinformatics, genetics, engineering and finance. Coupled with this has been the development of complex and realistic Bayesian statistical models to represent these data. In order to use these models to perform (Bayesian) statistical inference, one is required to calculate integrals, which are unknown analytically. Most of the numerical methods used to approximate these integrals are based upon Monte Carlo methods of which some of the seminal work has been done at Imperial College London, for instance the `particle-filter' developed in 1993 [4]. These methods are now very widely used in finance for automated trading, calculating the probability of default for economies, and for target tracking in the defence sector and we give explicit exemplars of each. The numerical methods developed at Imperial have been important in applying realistic models to these varied application areas and have impacted companies and organisations as diverse as Maple-Leaf Capital LLC, QinetiQ and the Credit Research Initiative.

2. Underpinning research (indicative maximum 500 words)

Bayesian computation can be split, roughly, into two main components: Markov chain Monte Carlo (MCMC) and particle filtering/sequential Monte Carlo (SMC) methods. These are the two primary numerical techniques which are the current state-of-the art to perform Bayesian statistical inference (i.e. to draw conclusions that are interpretable to non-statisticians) from complex and realistic statistical models; that is, those models which reflect real-world phenomena. Due to a variety of demands in real applications (for example high frequency trading in finance or the human genome project in genetics), the need to be able to perform such inference has greatly increased within the past 20 years.

The first appearance of an implementation of the particle filter was in the seminal paper by Gordon, Salmond and Smith [1]; further important developments were made by the group of Imperial researchers led by Adrian Smith whilst he was at Imperial. The bootstrap particle filter is the basis of almost every exact computational algorithm used for `filtering' a problem appearing in finance and the defence industry. Some additional important theoretical and methodological contributions were made, individually, by Dan Crisan, Mike Pitt and Ajay Jasra (e.g. [2]) respectively. These Imperial-led projects (for instance under grant [G1]) include the widely used `auxiliary particle filter' [3] and `sequential Monte Carlo samplers' [2] methods; also the foundations of the theoretical understanding of these methods is in [4]. The SMC sampler method allows for a different filtering problem (to that mentioned above) to be addressed and this is something which is now used, for example, in the calculation of probability of defaults on a daily basis. The auxiliary particle filter is a substantial adaptation and improvement of the particle filter.

MCMC techniques were also put to use and developed by Adrian Smith's group. In particular, new MCMC methods and applications to important statistical models were at the forefront of the work that was done. These were some of the first researchers in the world to nurture this methodology for real Bayesian models and included Jon Wakefield, Chris Holmes, Dave Denison, David Stephens and Bani Mallick (culminating in the book [5]). These works were funded by grants [G2, G3]. As an example, of the resulting academic research that was done and applied, David Stephens and Matthew Gander used MCMC for financial models in 2001-2004 (see [6]).

All of this Imperial research effort can be thought of as the foundations of all of the current (and substantial) research being done in Bayesian computation.

Key contributors:

- The key Imperial staff and students involved in the above research include Adrian Smith (Professor, 1990-1998), David Stephens (RA then lecturer, 1990-2006), Bani Mallick (lecturer, 1994-1998), Dan Crisan (Professor, 2000-Present, RA 1995-1998), Jon Wakefield (lecturer, 1990-1996), Ajay Jasra (RF & lecturer, 2006-2011, previously PhD student, 2002-2005), Mike Pitt (RA, 1995-1999), Christopher C. Holmes (PhD, RA & lecturer, 1996-2004), Neil Gordon (PhD, 1990-1993, then Defense Research Agency), Matthew Gander (PhD, 2001-2004).
- Key external collaborators include A Doucet (University of Cambridge, University of Melbourne, University of British Columbia) and D Salmond (Defence Research Agency)

3. References to the research (* References that best indicate quality of underpinning research)

[1] * Gordon, N., Salmond, D. & Smith, A.F.M., '*Novel approach to nonlinear/non-Gaussian Bayesian state estimation*', Radar and Signal Processing, IET Proceedings F, 140 (2), p107-113 (1993). DOI.

[2] Del Moral, P., Doucet, A. & Jasra, A., '*Sequential Monte Carlo samplers*', J. R. Statist. Soc. B, 68, 411-436 (2006). DOI. [The work by A Jasra was conducted at Imperial but his author affiliation on the paper is University of Cambridge.]

[3] * Pitt, M.K., Shephard, N., '*Filtering via simulation: auxiliary particle filters',* Journal of the American Statistical Association, 94:446, 590-599 (1999). DOI.

[4] * Crisan, D., & Doucet, A., '*A survey of convergence results on particle filtering methods for practitioners'*, IEEE Trans. Sig. Proc., 50, 736-746 (2002). DOI.

[5] Denison, D.G.T., Holmes, C.C., Mallick, B.K. & Smith, A.F.M., '*Bayesian Methods for Nonlinear Classification and Regression*', Publ. Wiley: New York, ISBN-13: 978-0471490364 (2002).

[6] Gander M.P.S., Stephens D.A., *'Simulation and inference for stochastic volatility models driven by Levy processes',* Biometrika, 94 (3), 627-646 (2007). DOI.

Grants:

[G1] EPSRC Grant: EP/H000550/1, 'Increasing the efficiency of numerical methods for estimating the state of a partially observed system', PI: Dr D. Crisan, £314,974, 01/10/09-30/03/13.

[G2] EPSRC Grant: GR/G62103/01, 'Studies in Bayesian Computation and Display methodology', PI: Prof A. F.M. Smith, £100,960, 13/03/92-12/04/95.

[G3] EPSRC Grant: GR/L10437/01, 'Bayesian population pharmacokinetic & pharmacodynamics modelling: implementation and model selection', PI: Dr J. Wakefield, Co-I: N Best, D Spiegelhalter, £123,811, 01/10/96-31/03/99.

4. Details of the impact (indicative maximum 750 words)

The impact from the above described work on Bayesian computation has been varied and wide, impacting a number of sectors.

Finance:

Particle filters and MCMC are used in many financial institutions. One particular case study is the development of algorithmic trading strategies at Maple-Leaf Capital LLC. Algorithmic trading strategies use mathematical models and rules to decide upon buying and selling financial instruments, such as equities and futures. Often these models and rules are combined to generate a trading position. Prior to the work to be described, understanding the statistical interaction between strategies was challenging as estimation of these properties was often very difficult.

The head of quantitative analysis at Maple–Leaf Capital LLP in the period 2004-2010 confirms the use of the MCMC techniques in finance. The methods developed at Imperial College were used by the quantitative analysis team "*to infer a probability distribution associated to the statistical interaction between strategies…Once this information [was] available, it was used to calibrate the position of the overall combination of trading strategies*" which was then used "*to assist trades*" [A]. The exact impact of the algorithm (e.g. in terms of profits) is not publicly available, but the improvement in estimation (of statistical interaction of the strategies) was adopted due to apparent failures of existing methodology. The quantitative analysis team "*did not know of any method at that time, other than MCMC, which could have fitted this type of model*". The algorithms were thus "*critical in this part of the work*" [A].

This example of financial impact is presented in detail, but the methodology is used generically across the field: MCMC and particle filters are further utilized as methods to analyse and predict financial positions and complement and enhance existing methods. There have been numerous papers devoted to this application of particle filtering/SMC methods (e.g. [B]).

Credit research initiative (CRI), National University of Singapore (NUS):
This is a non-profit undertaking by the Risk Management Institute (RMI) at NUS that uses MCMC and particle filters for assistance in predicting credit risk rate in a "public good" approach to credit rating. NUS launched the CRI in 2009 to output predictions of probability of default (PDs) using advanced statistical models and intending to "*give the big rating agencies like Standard & Poor's, Moody's Investor Service and Fitch Ratings a run for their money*" [C] (see also [D, E]). As is well known in the popular literature, the credit prediction system melted down in 2007/2008, leading to the credit crunch and resulting financial crisis. There are, of course, a wide range of reasons for this, but one must be attributed to the problems of existing models and methods for prediction of PDs. One solution is to "*leverage open source models with fully transparent inputs and outputs*" with "*software and data…open to a worldwide peer review process…[to] facilitate their rapid improvement*" [F]. Such "*open source, transparent credit models and methodologies would eliminate conflicts of interest and bring the benefits of mass collaboration to the world of credit ratings*" [G].

In order to accurately fit the models, the SMC sampler technique of [2], co-developed at Imperial College London, is utilized by the CRI to provide online predictions. The CRI's methodology for parameter estimation is described in Duan, J.C. & A. Fulop (2013) which references [2] in describing the parameter estimation by SMC [H, I]. The SMC method is used to "*deal with the problem of high dimension of the parameter space*", allowing uncertainty to be properly assigned to the parameters [J]. The online predictions provided by the CRI would not be possible without such methods. The predictions are publicly available to anyone (subject to the decision of NUS).

The CRI website [K] offers daily predictions from a probability of default (PD) model for defaults of about 60,400 listed firms in 106 economies in Asia Pacific, North America, Europe, Latin America, the Middle East and Africa [L]. This web portal presents the outputs from this model, including daily updated PDs for individual firms in the aforementioned regions and aggregate PDs for different economies and sectors. The CRI has "*agreed to provide [the] Probability of Default to a number of financial institutions for their internal risk management and analysis*" and its website has over 2000 registered users [J]. The CRI initiative shows the "*potential for open source credit models to take their place next to proprietary software and agency ratings*" in the credit ratings industry [G].

Defence:
Target tracking is the notion of trying to estimate or predict the position and/or velocity of targets simultaneously given noisy sensor measurements. This has particular applications in the defence industry, where the `targets' could be enemy tanks/aircraft/submarines and the measurements are 'noisy' measurements recorded by sensors. These phenomena are often modelled by a state-space model.

Until the development of particle filtering methodology, one could only apply the most basic of state-space models, which are often unrealistic representations of the real data phenomena encountered in target tracking. The development of the bootstrap particle filter by Adrian Smith and co-workers in [1] was one of the most fundamental methods to allow one to fit, online, realistic state-space models. This work has become integral in the target tracking work of the UK defence industry, such as QinetiQ. The bootstrap particle filter has been routinely applied at QinetiQ and BAE systems since 1993 and plays a fundamental role in the defence of the United Kingdom and the ability to predict or estimate the position of the enemy.

Confirmation of the impact of the work in the defence sector comes from the current Principal Consultant (National Security) at BAE Systems Detica who is able to confirm the use of particle filters based around paper [1] in the period 2008-2012 whilst employed at QinetiQ [M]. Particle

filters allowed QinetiQ to "*tackle problems that typically had weak or no existing solution*" [M]. As an example, in the context of multi-target tracking, "*particle filters enabled [QinetiQ] to constrain objects to be on the road, improving localisation accuracy, use interacting models to constrain objects motion by other objects, and to perform inference in bearings only tracking problems*" [M]. These problems are "*routinely found in the defence industry*" and "*the particle filter played an important role in [QinetiQ's] work*" [M]. Unfortunately it is not possible to receive confirmation about precisely what was implemented in real systems however it can be confirmed that "*particle filters had a massive impact on the breath of problems that could be solved [in defense], allowing tracking systems to be deployed in scenarios that were previously impossible (or unreliable)*" [M].

Navigation and wireless networks:
Similar to the application for target tracking in the defence industry, particle filters/SMC are also used in navigation (GPS) and for tracking in, the now standard problem of, wireless sensor networks. Frameworks for positioning, navigation, and tracking problems have been developed and particle filters can be used for positioning based on cellular phone measurements, for integrated navigation in aircraft, and for target tracking in aircraft and cars. The particle filter enables a promising solution to the combined task of navigation and tracking, with possible application to collision avoidance systems in cars [e.g. N].

5. Sources to corroborate the impact (indicative maximum of 10 references)

[A] Letter from Quantitative Analysis, Tudor Capital LLP, formerly Head of Quantitative Analysis, Maple-Leaf Capital LLP (available from Imperial on request)
[B] Examples of the financial applications of SMC/particle filtering: DOI-1, DOI-2
[C] Today, Singapore article, 'Ratings Systems: New Asian Kid on the Block?', 17/7/09 (archived here).
[D] Reuters article, 'Singapore university seeks to break hold of credit-rating goliaths', 14/10/11 (archived here).
[E] Business Times article, 'NUS offers free global credit ratings of firms', 16/7/10 (archived here).
[F] http://www.guardian.co.uk/commentisfree/2013/feb/25/moodys-sp-credit-rating-agencies-need-reform (archived at https://www.imperial.ac.uk/ref/webarchive/rmf on 19/6/13)
[G] http://tabbforum.com/opinions/can-open-source-models-fix-the-credit-ratings-industry (archived at https://www.imperial.ac.uk/ref/webarchive/smf on 19/6/13)
[H] The methodology for the parameter estimation used in the CRI models, which references [2], is described in Duan, J.-C., A. Fulop, 2013, 'Multiperiod Corporate Default Prediction with Partially-Conditioned Forward Intensity' (archived here)
[I] Description of background documents for the CRI models, http://www.rmicri.org/about/backgrounddocs.php (archived at https://www.imperial.ac.uk/ref/webarchive/tmf on 17/5/13)
[J] Letter from Deputy Director of Education and Industry Relations, RMI, NUS, 9/6/13 (available from Imperial on request)
[K] http://www.rmicri.org/home/ (archived at https://www.imperial.ac.uk/ref/webarchive/wmf on 19/6/13)
[L] http://www.rmicri.org/about/aboutcri.php (archived at https://www.imperial.ac.uk/ref/webarchive/vmf on 19/6/13)
[M] Letter from Principal Consultant, National Security, BAE Systems (formerly at QinetiQ), 10/6/13 (available from Imperial on request)
[N] Examples of the use of SMC/particle filtering in navigation and wireless networks: DOI-1, DOI-2, DOI-3, DOI-4, DOI-5