
Foreword

ON SEQUENTIAL MONTE CARLO: AN OVERVIEW

The sequential Monte Carlo (SMC) methodology is a family of powerful Monte Carlo methods for high dimensional static or dynamic problems. It is built upon sequential importance sampling and incorporates various other Monte Carlo techniques such as resampling, rejection sampling, auxiliary variable method, Markov chain Monte Carlo (MCMC), etc. As an alternative to Markov Chain Monte Carlo (MCMC) method, SMC can be very useful in cases where MCMC methods are inefficient or are inappropriate (such as on-line updating). One of the earliest forms of SMC can be dated to Hammersley and Morton (1954) and Rosenbluth and Rosenbluth (1955), who noticed that sequentially simulating a self-avoiding random walk (SAW) with one-step look-ahead is a good strategy but is biased. They demonstrated that this bias can be corrected by sample reweighting. Since the SAW model serves as a prototype for biopolymer modeling, this idea opens up new frontiers for cross-fertilization between computational science and many application areas.

In the early 1990s, researchers in the signal processing and statistics communities independently rediscovered the power of sequential sampling and updating strategies for coping with nonlinear state-space models and general missing data problems (Gordon, Salmond and Smith (1993); Kong, Liu and Wong (1994); Liu and Chen (1995)). These strategies were later coined the general name “sequential Monte Carlo” (Liu and Chen (1998)), and they reignited a new wave of exciting developments that have revolutionized the original sequential importance sampling method and greatly broaden the applications (Doucet, de Freitas and Gordon (2001); Andrieu, Doucet and Holenstein (2010)). A particularly new twist in this new wave is the explicit formulation of the resampling strategy (Liu and Chen (1995, 1998)), which unifies the improvement ideas existing in the molecular simulation literature.

This special issue celebrates the 70th anniversary of the publication of the first SMC method used on a computer. The nine featured articles from authors all over the world represent the rich spectrum of the SMC methodology and its applications. These articles cover new and improved SMC-type algorithms, new algorithms based on SMC, new methods that use SMC as an indispensable ingredient, and new ways of adapting existing SMC methods for new tasks.

Zhang, Song and Liang introduce the Langevinized EnKF (LEnKF), which combines the ensemble Kalman filter (EnKF) with mini-batch data updates for stochastic gradient Langevin dynamics, as an enhanced sequential preconditioned SGLD sampler (Li et al. (2016)). They show both theoretically and empirically that this novel approach is superior to traditional methods such as EnKF or particle filters in both accuracy and scalability, particularly for large-scale dynamic data. Crucinio and Johansen introduce a divide-and-conquer approach to filtering, which partitions the state variable into lower-dimensional components for SMC updates. The components are then recursively merged to reconstruct the full filtering distribution. This method is less reliant on factorizing transition densities and observation likelihoods compared to other methods, making it applicable to a wider range of models. Cardoso, Moulines and Olsson develop a Parisian Particle Gibbs (PPG) sampler for estimating smoothed expectations of additive functionals in general state-space models, which uses the conditional version of the PARIS algorithm (Olsson and Westerborn (2017)) as an inner loop in a Gibbs sampling method. The new algorithm has the same computational complexity as PARIS but reduces the bias significantly. The authors also provide theoretical results on the bounds on bias and variance.

Iodies, Ning and Wheeler discuss challenges of statistical inference for partially observed nonlinear dynamic systems, particularly meta-populations with multiple units. They propose the Iterated Block Particle Filter (IBPF), which combined an iterated filtering likelihood maximization technique with a block particle filter, to alleviate the curse of dimensionality. This method allows for the estimation and inference of both unit-specific and shared parameters. Rimella, Jewell and Fearnhead address challenges in Individual-Based Epidemic Models, where standard particle filters suffer from inefficiency due to the mismatch of proposal and categorical observations. The authors propose a way to approximate the optimal proposals efficiently, in terms of both estimation accuracy and computational complexity, by exploiting the conditional independence structure of the model and avoiding marginalization among the whole state space. Cai, Chen and Lin explore the use of SMC to efficiently generate sample paths from dynamic systems with strong constraints. The authors propose a strategy based on time-varying look-ahead timescale to identify optimal resampling priority scores using an ensemble of forward or backward pilots.

For continuous-time hidden Markov models with a Cox process observation, Jin, Singh and Chopin propose to start with the Poisson unbiased estimates of the path integrals, retain the positive parts for nonnegative likelihood estimation, and introduce a clever trick to limit the possible truncation bias when enforcing nonnegativity. Li, Wang, Deng and Liu discuss challenges of estimating the gradient of the log-likelihood function in statistical models with hidden variables

under the SMC framework, of which the main one arises from the discreteness of the particles and the complexity introduced by resampling. They propose to replace the traditional discrete resampling mechanism with a kernel-smoothed resampling procedure and demonstrate its improved performance in various model settings. Hou and Wong propose the UDSMC scheme for protein structural sampling. Each step of the scheme involves a multi-descendants upsampling stage, followed by a downsampling (resampling) stage. The downsampling stage employs either the “optimal resampling” method of Fearnhead and Clifford (2003) or the simple multinomial resampling, depending on the scenarios.

The articles featured in this special issue have been instrumental in shaping our recent research activities and thoughts in the area of SMC. These articles have not only provided us with valuable insights and knowledge, but also challenged our preconceived notions about SMC. We hope that these articles will inspire you to explore new ideas and perspectives in this exciting area of research, and that they will help you discover new applications and advancements in the field.

Acknowledgments

I am most grateful to my co-editor Arnaud Doucet, the editorial team of *Statistica Sinica*, and all participating authors. Without their wholehearted dedication it is impossible to have this special issue ready in such a timely manner. I am also grateful to my students Yichao Li and Da Wang, who have helped greatly with reviewing and summarizing the submitted articles.

Jun S. Liu
Guest Editors for this special issue
Harvard University

References

- Andrieu, C., Doucet, A. and Holenstein, R. (2010). Particle Markov Chain Monte Carlo methods. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **72**, 269–342.
- Doucet, A., de Freitas, N. and Gordon, N. J. (2001). *Sequential Monte Carlo Methods in Practice*. Springer.
- Fearnhead, P. and Clifford, P. (2003). On-line inference for hidden Markov models via particle filters. *Journal of Royal Statistical Society. Series B (Statistical Methodology)* **65**, 887–899.
- Gordon, N. J., Salmond, D. J. and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings on Radar and Signal Processing* **140**, 107–113.
- Hammersley, J. M. and Morton, K. W. (1954). Poor man’s Monte Carlo. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **16**, 23– 38.

- Kong, A., Liu, J. and Wong, W. (1994). Sequential imputations and Bayesian missing data problems. *Journal of the American Statistical Association* **89**, 278–288.
- Li, C., Chen, C., Carlson, D. and Carin, L. (2016). Preconditioned Stochastic Gradient Langevin Dynamics for deep neural networks. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, 1788–1794.
- Liu, J. and Chen, R. (1995). Blind deconvolution via sequential imputations. *Journal of the American Statistical Association* **90**, 567–576.
- Liu, J. and Chen, R. (1998). Sequential monte carlo methods for dynamic systems. *Journal of the American Statistical Association* **93**, 1032–1044.
- Olsson, J. and Westerborn, J. (2017). Efficient particle-based online smoothing in general hidden Markov models: The PaRIS algorithm. *Bernoulli* **23**, 1951–1996.
- Rosenbluth, M. N. and Rosenbluth, A. W. (1955). Monte Carlo calculation of the average extension of molecular chains. *The Journal of Chemical Physics* **23**, 356–359.