

Three discussions of the paper ‘sequential quasi-Monte Carlo sampling’, by M. Gerber and N. Chopin, Read before The Royal Statistical Society at a meeting organised by the Research Section on Wednesday, December 10th, 2014

Julyan Arbel[†], Collegio Carlo Alberto and University of Torino

Igor Prünster[‡], Collegio Carlo Alberto and University of Torino

Christian P. Robert[§], Université Paris-Dauphine and University of Warwick

Robin J. Ryder[¶], Université Paris-Dauphine

Summary. This is a collection of three written discussions of the paper “sequential quasi-Monte Carlo sampling” by M. Gerber and N. Chopin, following the presentation given before the Royal Statistical Society in London on December 10th, 2014.

1. Consequences on Bayes nonparametrics (J. Arbel and I. Prünster)

We congratulate the authors for a stimulating paper, accessible to laypersons, and thank them for sharing their code. We would like to comment on the potential implication of SQMC in Bayesian nonparametrics.

Nonparametric mixtures are commonly used for density estimation and can be thought of as an extension of finite mixture models when the number of clusters is unknown. The setting is as follows: observations $\mathbf{y}_{1:t}$ are spread out into k_t clusters; the cluster labels, or allocation variables, are denoted by $\mathbf{x}_{1:t}$, and are interpreted as the *states* of $\mathbf{y}_{1:t}$ in the context of SMC. Note that the states are discrete and elements of $\{1, \dots, k_t\}$, where k_t denotes the number of clusters. The $n_{t,j}$ observations allocated to the same cluster j are iid from a given parametric kernel $K(\cdot; \theta_j)$. This model is static in the sense that data are usually not recorded sequentially. However, Pólya urn type schemes (Blackwell and MacQueen, 1973) essentially formulate this model sequentially by artificially thinking of observation \mathbf{y}_t to occur at time t . So the nonparametric mixture model can be cast as an SMC sampler as follows (see Liu, 1996; Fearnhead, 2004)

$$\begin{aligned} \mathbf{y}_t | \mathbf{x}_t, \theta &\sim K(\mathbf{y}_t; \theta_{\mathbf{x}_t}), \\ \mathbf{x}_t | \mathbf{x}_{1:t-1} &\sim f^X(\mathbf{x}_t | \mathbf{x}_{1:t-1}). \end{aligned}$$

The predictive distribution f^X is particularly simple for the Dirichlet process (Ferguson, 1973; Blackwell and MacQueen, 1973). Denote by α and G_0 the precision parameter and the base measure of the Dirichlet process. Then the kernel for the states is given by the following probability mass function on $\{1, \dots, k_{t-1} + 1\}$

$$m_t(\mathbf{x}_{1:t-1}, \mathbf{x}_t = j) = p_{t,j} \propto \begin{cases} n_{t-1,j} K_j(\mathbf{y}_t | \mathbf{x}_{1:t-1}) & \text{if } j \leq k_{t-1} \\ \alpha K_0(\mathbf{y}_t) & \text{if } j = k_{t-1} + 1 \end{cases} \quad (1)$$

where K_j and K_0 involve integrals of K and G_0 which can be calculated in conjugate cases and approximated otherwise. Note that $m_t(\mathbf{x}_{1:t-1}, \mathbf{x}_t)$ is also available in closed-form expression for broader classes of random probability measures including the two-parameter Poisson–Dirichlet process (Pitman and Yor, 1997) and normalised random measures with independent increments (Regazzini et al., 2003; James et al., 2009).

As stressed by Gerber and Chopin, the key ingredient of SQMC is the deterministic transform Γ_t . A possible choice for (1) is

$$\Gamma_t(\mathbf{x}_{1:t-1}, \mathbf{u}_t) = \min \left\{ j \in \{1, \dots, k_{t-1} + 1\} : \sum_{i=1}^j p_{t,i} > \mathbf{u}_t \right\} \quad (2)$$

where $\mathbf{u}_t \sim \mathcal{U}([0, 1])$. The case of discrete proposal $m_t(\mathbf{x}_{1:t-1}, \mathbf{x}_t)$ does not seem to be discussed in the paper. Would the authors expect a gain in efficiency by considering a quasi-Monte Carlo alternative as (2) instead of standard SMC?

[†]julyan.arbel@carloalberto.org
[‡]igor.pruenster@unito.it
[§]xian@ceremade.dauphine.fr
[¶]ryder@ceremade.dauphine.fr

2. Spectrum of the sqMC method (C. Robert)

Quasi-Monte Carlo methods are ignored by the statistical community, despite regular attempts by well-respected researchers to induce a broader use of those methods. The authors are to be congratulated and will hopefully contribute to a wider diffusion of qMC methods.

At a purely mathematical level, that randomised low discrepancy sequences produce both unbiased estimators and error rates of order $O(N^{-1} \log(N)^{d-1})$ at cost $O(N \log(N))$ implies that sqMC methods should *always* be replace regular Monte Carlo methods. Since this is not the case, we may wonder at the reasons. The major difficulty with qMC stands in its requirement to expressing Monte Carlo estimators as deterministic transforms of a fixed number of uniforms. This is a strong deterrent to the dissemination of qMC alternatives, as moving to a random number of baseline uniforms does not seem achievable in full generality. It is thus no surprise that the proposal appears in connection with particle filters. Indeed, this field centres on dynamic importance sampling and hence enjoys a local iid-ness that relates better to qMC integrators than single-chain MCMC algorithms. For instance, each particular resampling step consists in multinomial simulation, hence could be turned into qMC. Actually, lower variance solutions that amount to qMC, like systematic and residual sampling (Gundersen and Jensen, 1987; Fearnhead, 1998; Chopin, 2004), have been proposed in the particle literature.

Here, the authors apply qMC to the particles themselves, using a *global* low discrepancy sequence for the resampling and the move steps: as the cost of using the Hilbert curve grows quickly with dimension, two sequences in lieu of one would bring a significant cost reduction. Moreover, they still assume the deterministic transform representation

$$\mathbf{x}_t^n = \Gamma_t(\mathbf{x}_{t-1}^n, \mathbf{u}_t^n),$$

which is a stumbling block for those contemplating switching to qMC, as illustrated in the paper using only Normal baseline distributions. In a sequential setting with unknown parameters θ , the transform Γ changes each time θ is modified, which impacts computing cost when the inverse cdf is not available. Since simulating θ is unlikely to benefit from qMC improvements, one may question the relevance of the overall scheme: the most variable item in the Gibbs sampling steps expands its inefficiency to the joint kernel.

In connection with Lemma 1 and the sqMC approximation of the evidence, Rao-Blackwellisation using all *proposed* moves could be considered: is this easier and significant within qMC since the gain may be of a lower order? Similarly, the trick at the end of §4.2.1, using sqMC on a single instead of $(t+1)$ vectors, is unclear, but I wonder at a connection with the particle learning literature (Lopes et al., 2010; Carvalho et al., 2010).

In conclusion, I am looking forward the aftermath of this Read Paper that will expose whether qMC is bound to become the reference in simulation computational statistics or to remain a niche activity away from mainstream simulation.

3. ABC with sqMC (R. Ryder)

One of the reasons why the authors are so successful in their new methodology is the ability to know in advance the number of uniform variables which are needed for one step of the algorithm. This allows greater efficiency and, to use the terminology at the end of section 1.2 of the paper, the use of QMC point sets (fixed length) instead of QMC sequences (unbounded length). Nonetheless, as noted by the authors, the use of QMC sequences is at least theoretically possible.

When reading the paper, I was immediately curious about the possibility of adapting this methodology in an ABC (Approximate Bayesian Computation) setting. In an attempt to develop a naive ABC population Quasi Monte Carlo algorithm, it is natural to use QMC sequences instead of QMC point sets, as shown in Algorithm 1, which is adapted from Marin et al. (2012).

I tried this algorithm on a toy example: $y_k \sim N(x_k, \sigma^2)$ where (x_k) is a Markov chains taking values in $\{-2, 2\}$, with probability of switching at any iteration equal to θ , and attempted to estimate θ . For the generation of QMC sequences, I used the `randtoolbox` R package Dutang and Savicky (2013).

There are two points to note from this simple experiment:

- (a) The QMC version systematically outperforms the basic population ABC sampler, with a variance typically about 30 times smaller for an equal number of particles.
- (b) The computational time of the QMC version explodes in a non-linear fashion when the number of particles increases, as shown in Figure 1. This is intriguing, and might be due to the cost of generating a QMC sequence, at least in the implementation I used.

There is certainly much more work to be done to adapt these new ideas to an ABC setting. In particular, there is one situation where QMC point sets could certainly be useful: in ABC, one typically accepts a proposed value θ^* it that value leads to simulated data \mathbf{z} which is close to the observed data \mathbf{y} , in the sense that $d(\mathbf{y}, \mathbf{z}) < \epsilon$ for some pseudometric d and threshold ϵ . The threshold ϵ is note always chosen in advance, but is sometimes a quantile from a fixed number of attempts: there must surely be situation where QMC could be used to improve the estimator for such flavours of ABC. I was surprised not to find any investigation of this idea in the literature.

Algorithm 1 Likelihood-free population Quasi Monte Carlo sampler

Input: observations \mathbf{y} , decreasing sequence $(\epsilon_1, \dots, \epsilon_N)$, prior π with cdf F_π .

At iteration $t = 1$,

Set $j = 1$

for $i = 1$ to N **do**

repeat

 Generate u_j as the j th element in an 1 dimensional RQMC sequence.

 Set $\boldsymbol{\theta}_i^{(1)} = F_\pi^{-1}(u_j)$ and $\mathbf{z} \sim f(\mathbf{z} | \boldsymbol{\theta}_i^{(1)})$.

 Increment j .

until $\rho(\eta(\mathbf{z}), \eta(\mathbf{y})) \leq \epsilon_1$

 Set $\omega_i^{(1)} = 1/N$

end for

Take Σ_1 as twice the empirical variance of the $\boldsymbol{\theta}_i^{(1)}$'s

for $t = 2$ to T **do**

 Set $j = 1$

for $i = 1$ to N **do**

repeat

 Generate (u_j, v_j) as the j th element in a 2 dimensional RQMC sequence.

 Select an ancestor $\boldsymbol{\theta}_i^*$ from the $\boldsymbol{\theta}^{t-1}$ using v_j as for SQMC

 Generate $\boldsymbol{\theta}_i^{(t)}$ using $\boldsymbol{\theta}_i^*$ and u_j as for SQMC

 Increment j .

until $\rho(\eta(\mathbf{z}), \eta(\mathbf{y})) \leq \epsilon_t$

 Set $\omega_i^{(t)} \propto \pi(\boldsymbol{\theta}_i^{(t)}) / \sum_{j=1}^N \omega_j^{(t-1)} \varphi \left\{ \Sigma_{t-1}^{-1/2} \left(\boldsymbol{\theta}_i^{(t)} - \boldsymbol{\theta}_j^{(t-1)} \right) \right\}$

end for

 Take Σ_t as twice the weighted variance of the $\boldsymbol{\theta}_i^{(t)}$'s

end for

References

- Blackwell, D. and MacQueen, J. B. (1973). Ferguson distributions via Pólya urn schemes. *The Annals of Statistics*, pages 353–355.
- Carvalho, C., Johannes, M., Lopes, H., and Polson, N. (2010). Particle learning and smoothing. *Stat. Science*, 25(1):88–106.
- Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann. Statist.*, 32(6):2385–2411.
- Dutang, C., and Savicky, P. (2013) `randtoolbox`: Generating and Testing Random Numbers. R package version, 1.
- Fearnhead, P. (1998). *Sequential Monte Carlo methods in filter theory*. PhD thesis, University of Oxford.
- Fearnhead, P. (2004). Particle filters for mixture models with an unknown number of components. *Statistics and Computing*, 14(1):11–21.
- Ferguson, T. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1(2):209–230.
- Gundersen, H. J. G. and Jensen, E. B. (1987). The efficiency of systematic sampling in sterology and its prediction. *Journal of Microscopy*, 147(3):229–63.
- James, L. F., Lijoi, A., and Prünster, I. (2009). Posterior analysis for normalized random measures with independent increments. *Scandinavian Journal of Statistics*, 36(1):76–97.
- Liu, J. S. (1996). Nonparametric hierarchical Bayes via sequential imputation. *Annals of Statistics*, 24:910–930.
- Lopes, H., Carvalho, C., Johannes, M., and Polson, N. (2010). Particle learning for sequential Bayesian computation (with discussion and rejoinder). In Bernardo, J., M.J. Bayarri, J. B., Dawid, A., Heckerman, D., Smith, A., and West, M., editors, *Bayesian Statistics 9*, volume 9, pages 175–96. Oxford University Press.
- Pitman, J. and Yor, M. (1997). The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, 25(2):855–900.




Fig. 1. Computational time (in minutes) of the standard population ABC of Marin et al. (2012) (triangles) and QMC population ABC (circles) samplers, for increasing number of particles. The former increases linearly, but the latter does not.

Regazzini, E., Lijoi, A., and Prünster, I. (2003). Distributional results for means of normalized random measures with independent increments. *The Annals of Statistics*, 31:560–585.

Marin, J.-M., Pudlo, P., Robert, C. P., and Ryder, R.J. (2012). Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167-1180.