# ABC MCMC: a survey of theoretical results 

Anthony Lee<br>University of Warwick \& Alan Turing Institute

Validating and Expanding ABC Workshop
Banff International Research Station
February 20th, 2017

## Outline

# ABC pseudo-marginal Markov chains 

Comparison and order

## Outline

## ABC pseudo-marginal Markov chains

## Comparison and order

## Intractable likelihood functions

- Let $y_{\text {obs }} \sim f_{\text {obs }}\left(\cdot \mid \theta_{0}\right)$ be the observed data.
- Assume $f_{\text {obs }}\left(y_{\text {obs }} \mid \cdot\right)$ is intractable.
- Assume can draw $x \sim f_{\text {obs }}(\cdot \mid \theta)$ for any $\theta \in \Theta$.
- Approximation I: replace $y_{\text {obs }}$ with $y:=s\left(y_{\mathrm{obs}}\right)$.
- $f(y \mid \cdot)$ is typically also intractable.
- We can draw $x \sim f(\cdot \mid \theta)$ for any $\theta \in \Theta$.


## Intractable likelihood functions

- Approximation II: $\tilde{f}(y \mid \theta):=\int K(x, y) f(x \mid \theta) \mathrm{d} x$.
- $\tilde{f}$ is in some sense "even less" tractable than $f$.
- Standard choices include:

$$
K(x, y) \propto \mathbb{I}\{d(x, y) \leq \epsilon\}, \quad K(x, y)=\mathcal{N}(y ; x, \epsilon I)
$$

- Alternatives exist, e.g.

$$
\bar{f}(y \mid \theta)=\int f_{A}\left(y \mid \phi_{N}\left(x_{1: N}, \theta\right)\right) \prod_{i=1}^{N} f\left(x_{i} \mid \theta\right) \mathrm{d} x_{1: N}
$$

- $f_{A}$ is multivariate normal $\Rightarrow$ synthetic likelihood [Wood, 2010].


## Why is it useful?

- Denote by $p$ the prior density for $\theta$.
- An auxiliary target can be defined:

$$
\pi(\theta, w) \propto p(\theta) \tilde{f}(y \mid \theta) w Q_{\theta}(w)
$$

where $W \sim Q_{\theta}$ is non-negative and $\mathbb{E}_{Q_{\theta}}[W]=1$.

1. $\tilde{f}(y \mid \theta) W$ is a non-negative, r.v. with expectation $\tilde{f}(y \mid \theta)$
2. $\pi(\theta)=\int \pi(\theta, w) \mathrm{d} w \propto p(\theta) \tilde{f}(y \mid \theta)$.

- Rejection/importance sampling algorithms then follow.
- We can simulate a $\pi(\theta, w)$-invariant Markov chain.


## ABC-MCMC pseudo-marginal kernels

- To sample from $P(\theta, w ; \cdot)$ :

1. Draw $\theta^{\prime} \sim q(\theta, \cdot)$ and $w^{\prime} \sim Q_{\theta^{\prime}}$.
2. Output $\left(\theta^{\prime}, w^{\prime}\right)$ w.p.

$$
1 \wedge \frac{p\left(\theta^{\prime}\right) \tilde{f}\left(y \mid \theta^{\prime}\right) w^{\prime} q\left(\theta^{\prime}, \theta\right)}{p(\theta) \tilde{f}(y \mid \theta) w q\left(\theta, \theta^{\prime}\right)}
$$

otherwise output $(\theta, w)$.

- Drawing $w^{\prime} \sim Q_{\theta^{\prime}}$ is equivalent to producing an unbiased estimate $\tilde{f}\left(y \mid \theta^{\prime}\right) w^{\prime}$ of $\tilde{f}\left(y \mid \theta^{\prime}\right)$.


## ABC examples of unbiased estimators

- Pseudo-marginal methods [Beaumont, 2003, Andrieu and Roberts, 2009] are generally applicable.
- Marjoram et al. [2003]:

$$
w^{\prime}=K\left(x^{\prime}, y\right) / \tilde{f}\left(y \mid \theta^{\prime}\right), \quad x^{\prime} \sim f\left(\cdot \mid \theta^{\prime}\right)
$$

- Becquet and Przeworski [2007]:

$$
w^{\prime}=\frac{1}{N} \sum_{i=1}^{N} K\left(x_{i}^{\prime}, y\right) / \tilde{f}\left(y \mid \theta^{\prime}\right), \quad x_{i}^{\prime} \stackrel{i i d}{\sim} f\left(\cdot \mid \theta^{\prime}\right)
$$

- We denote the corresponding kernel by $P_{N}$.
- There are other possibilities, e.g. $r$-hit estimators [Lee, 2012]


## Exact/marginal kernel $P_{\star}$

- We can compare this kind of chain with an "exact" variant.
- To sample from $P_{\star}(\theta ; \cdot)$ :

1. Draw $\theta^{\prime} \sim q(\theta, \cdot)$
2. Output $\theta^{\prime}$ w.p.

$$
1 \wedge \frac{p\left(\theta^{\prime}\right) \tilde{f}\left(y \mid \theta^{\prime}\right) q\left(\theta^{\prime}, \theta\right)}{p(\theta) \tilde{f}(y \mid \theta) q\left(\theta, \theta^{\prime}\right)}
$$

otherwise output $\theta$.

- Can think of this as $P_{\infty}$, or the case where $w=1$.


## Outline

## ABC pseudo-marginal Markov chains

Comparison and order

## Performance measures

- To keep things simple, we will consider only

1. Asymptotic variance of ergodic averages:

$$
\operatorname{var}(f, P):=\lim _{n \rightarrow \infty} n \operatorname{var}\left(\frac{1}{n} \sum_{i=1}^{n} f\left(\theta_{i}, w_{i}\right)\right)
$$

where $\left(\theta_{0}, w_{0}\right) \sim \pi$.
2. Geometric ergodicity (GE):

$$
\left\|P^{n}\left(\theta_{0}, w_{0} ; \cdot\right)-\pi(\cdot)\right\|_{T V} \leq C(x) \rho^{n} .
$$

- Reversible $P: P$ is geometrically ergodic $\Rightarrow$ finite asymptotic variance for all $f \in L^{2}(\pi)$.
- Almost necessary, for $\Longleftrightarrow$ variance bounding instead of GE.


## Comparisons with $P_{\star} 1 / 2$

1. (B) [Andrieu and Vihola, 2015] For any $f \in L^{2}(\pi)$ with $f: \Theta \rightarrow \mathbb{R}, \operatorname{var}(f, P) \geq \operatorname{var}\left(f, P_{\star}\right)$.
2. (G) [Andrieu and Roberts, 2009, Andrieu and Vihola, 2015] If $W_{\theta} \sim Q_{\theta}$ is uniformly bounded in $\theta$, then $P_{\star} \mathrm{GE} \Rightarrow P \mathrm{GE}$ (at least for positive $P$ ).
3. (G) [Andrieu and Vihola, 2015] Under technical conditions on $f \in L^{2}(\pi)$ with $f: \Theta \rightarrow \mathbb{R}$,

$$
\lim _{N \rightarrow \infty} \operatorname{var}\left(f, P_{N}\right)=\operatorname{var}\left(f, P_{\star}\right)
$$

4. (B) [Andrieu and Roberts, 2009, Andrieu and Vihola, 2015] If $W_{\theta} \sim Q_{\theta}$ is unbounded for "enough" $\theta$ then $P$ cannot be GE (not typically a problem in ABC).

## Comparisons with $P_{\star} \quad 2 / 2$

4 (B) [Lee and Łatuszyński, 2014, Andrieu and Vihola, 2015] If $W_{\theta} \sim Q_{\theta}$ is bounded but not uniformly so, then $P$ might not inherit GE from $P_{\star}$.

- For $K(x, y)=\mathbb{I}(d(x, y) \leq \epsilon), \tilde{f}(y \mid \theta)>0$ for all $\theta$ with $\tilde{f}(y \mid \theta) \rightarrow 0$ as $\|\theta\| \rightarrow \infty$ and $q$ "local" then $P_{N}$ cannot be GE for any $N$.

5 (G) [Deligiannidis and Lee, 2016] If $\sup _{\theta} \operatorname{var}\left(W_{\theta}\right)<\infty$ and $P$ GE then $\operatorname{var}(f, P)<\infty$ for any $f \in L^{2}(\pi)$ with $f: \Theta \rightarrow \mathbb{R}$.

## Ordering $P$ 's

- [Andrieu and Vihola, 2016] If $\left\{W_{\theta} ; \theta \in \Theta\right\} \leq_{c x}\left\{W_{\theta}^{\prime} ; \theta \in \Theta\right\}$ then $\operatorname{var}(f, P) \leq \operatorname{var}\left(f, P^{\prime}\right)$.
- Implies that $\operatorname{var}\left(f, P_{N}\right) \leq \operatorname{var}\left(f, P_{N+1}\right)$ for $N \in \mathbb{N}$.
- Also motivates stratification in ABC and dependent estimators.
- But how much better is $P_{N+1}$ compared to $P_{N}$ ?
- Improvement diminishes eventually as $\operatorname{var}\left(f, P_{\star}\right) \leq \operatorname{var}\left(f, P_{N}\right)$.


## Computational considerations

- [Bornn et al., 2017, Sherlock et al., 2016] Let $M \leq N$. Then

$$
M\left[\operatorname{var}\left(f, P_{M}\right)+\operatorname{var}_{\pi}(f)\right] \leq N\left[\operatorname{var}\left(f, P_{N}\right)+\operatorname{var}_{\pi}(f)\right]
$$

which implies

$$
\operatorname{var}\left(f, P_{1}\right) \leq N\left[\operatorname{var}\left(f, P_{N}\right)+\operatorname{var}_{\pi}(f)\right]-\operatorname{var}_{\pi}(f)
$$

i.e. simple averaging cannot bring "too much" benefit.

- $P_{N}$ positive implies $\operatorname{var}\left(f, P_{1}\right) \leq(2 N-1) \operatorname{var}\left(f, P_{N}\right)$.
- Also shows that $\operatorname{var}\left(f, P_{N}\right)<\infty \Longleftrightarrow \operatorname{var}\left(f, P_{1}\right)<\infty$.
- If comp. cost is proportional to $N$, often best to use $N=1$.


## Discussion $1 / 2$

- There exist provably more robust Markov chains, e.g. 1-hit ABC [Lee et al., 2012, Lee and Łatuszyński, 2014], r-hit variants [Lee, 2012], correlated pseudo-marginal methods [Deligiannidis et al., 2015].
- Understanding is still incomplete.
- Other Monte Carlo methods, e.g. SMC samplers / PMC.
- Choice of summary statistics.
- How to exploit mappings $F^{-}(U)=X \sim f(\cdot \mid \theta)$ where $U \sim \mathcal{U}\left([0,1]^{d}\right)$.


## Discussion $2 / 2$

- There are potential benefits to alternative approximate likelihoods. E.g., in a very simple scenario [Price et al., 2017],

$$
\bar{f}_{N}(y \mid \theta)=\int f_{A}\left(y \mid \phi_{N}\left(x_{1: N}, \theta\right)\right) \prod_{i=1}^{N} f\left(x_{i} \mid \theta\right) \mathrm{d} x_{1: N}
$$

is comp. more robust than $\tilde{f}(y \mid \theta)=\int K_{\epsilon}(x, y) f(x \mid \theta) \mathrm{d} x$.

- $N$ acts like $1 / \epsilon$, controls some approximation error.
- Natural estimator of $\bar{f}_{N}(y \mid \theta)$ converges in prob. as $N \rightarrow \infty$ with cost $\mathcal{O}(N)$, but for a given dimension $d$ one needs $\mathcal{O}\left(N^{d / 2}\right)$ samples to stabilize the natural estimator of $\tilde{f}(y \mid \theta)$.
- Of course, in general $\bar{f}_{N}(y \mid \theta) \nrightarrow f(y \mid \theta)$ as $N \rightarrow \infty$.


## References I

C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. Ann. Statist., 37(2):697-725, 2009.
C. Andrieu and M. Vihola. Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. Ann. Appl. Probab., 25(2):1030-1077, 2015.
C. Andrieu and M. Vihola. Establishing some order amongst exact approximations of MCMCs. Ann. Appl. Probab., 26(5):2661-2696, 2016.
M. A. Beaumont. Estimation of population growth or decline in genetically monitored populations. Genetics, 164:1139-1160, 2003.
C. Becquet and M. Przeworski. A new approach to estimate parameters of speciation models with application to apes. Genome research, 17(10): 1505-1519, 2007.
L. Bornn, N. S. Pillai, A. Smith, and D. Woodard. The use of a single pseudo-sample in approximate Bayesian computation. Statistics and Computing, 2017. To appear.
G. Deligiannidis and A. Lee. Which ergodic averages have finite asymptotic variance? arXiv preprint arXiv:1606.08373, 2016.
G. Deligiannidis, A. Doucet, and M. K. Pitt. The correlated pseudo-marginal method. arXiv preprint arXiv:1511.04992, 2015.

## References II

A. Lee. On the choice of MCMC kernels for approximate Bayesian computation with SMC samplers. In Proceedings of the Winter Simulation Conference, 2012.
A. Lee and K. Łatuszyński. Variance bounding and geometric ergodicity of Markov chain Monte Carlo kernels for approximate Bayesian computation. Biometrika, 101(3):655-671, 2014.
A. Lee, C. Andrieu, and A. Doucet. Discussion of "Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation" by Fearnhead and Prangle. J. R. Stat. Soc. Ser. B Stat. Methodol., 74(3):449-450, 2012.
P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain Monte Carlo without likelihoods. Proceedings of the National Academy of Sciences, 100 (26):15324-15328, 2003.
L. F. Price, C. C. Drovandi, A. Lee, and D. J. Nott. Bayesian synthetic likelihood. Journal of Computational and Graphical Statistics, 2017. To appear.
C. Sherlock, A. Thiery, and A. Lee. Pseudo-marginal Metropolis-Hastings using averages of unbiased estimators. arXiv preprint arXiv:1610.09788, 2016.
S. N. Wood. Statistical inference for noisy nonlinear ecological dynamic systems. Nature, 466(7310):1102-1104, 2010.

