

Department of Mathematics and Statistics


Preprint MPS-2014-20

30 May 2014

Lazy ABC

by

Dennis Prangle



Lazy ABC

Dennis Prangle*

Abstract

Approximate Bayesian computation (ABC) performs statistical inference for otherwise intractable probability models by accepting parameter proposals when corresponding simulated datasets are sufficiently close to the observations. Producing the large quantity of simulations needed requires considerable computing time. However, it is often clear before a simulation ends that it is unpromising: it is likely to produce a poor match or require excessive time. This paper proposes lazy ABC, an ABC importance sampling algorithm which saves time by sometimes abandoning such simulations. This makes ABC more scalable to applications where simulation is expensive. By using a random stopping rule and appropriate reweighting step, the target distribution is unchanged from that of standard ABC. Theory and practical methods to tune lazy ABC are presented and demonstrated on the computationally demanding spatial extremes application of Erhardt and Smith (2012), producing efficiency gains, in terms of effective sample size per unit CPU time, of roughly 3 times for a 20 location dataset, and 8 times for 35 locations.

Keywords: importance sampling, ABC, unbiased likelihood estimators, spatial extremes

*University of Reading. Email d.b.prangle@reading.ac.uk

1 Introduction

Approximate Bayesian computation (ABC) algorithms are a popular method of inference for a wide class of otherwise intractable probability models in applications such as population genetics, ecology, and systems biology (Beaumont, 2010; Marin et al., 2012). They select parameter vectors θ for which datasets y simulated from the model of interest are sufficiently close to the observations. A bottleneck is the computational cost of producing the large quantity of model simulations needed, which becomes increasingly severe for more detailed models. However, it is often clear during a simulation that it is unpromising. For example it is likely to produce a poor match or to require excessive computation time. This paper presents *lazy ABC*, an importance sampling method which abandons some such simulations, a step referred to as *early stopping*, exploiting information from incomplete simulations to save time. The result is an ABC algorithm which is more scalable to applications where simulation is computationally demanding.

In more detail, standard ABC is based on a random likelihood estimator \hat{L}_{ABC} , which is 1 for a close match of simulated and observed data, and zero otherwise. The algorithm can be shown to target a distribution corresponding to the approximate likelihood $L_{\text{ABC}}(\theta) = \text{E}[\hat{L}_{\text{ABC}}|\theta]$. Lazy ABC is based on an alternative estimator \hat{L}_{lazy} . This equals zero with probability $1 - \alpha$ – if early stopping is performed – and otherwise equals the \hat{L}_{ABC} estimator multiplied by a weight. Letting the weight equal $1/\alpha$ makes \hat{L}_{lazy} an unbiased estimator of $L_{\text{ABC}}(\theta)$. Results on random likelihood estimates show that importance sampling (and Markov chain Monte Carlo) algorithms based on $\hat{L}_{\text{lazy}}(\theta)$ therefore target the same distribution as standard ABC. No further approximation has been introduced.

The lazy ABC estimator trades off an increase in variance for a reduction in computation time. It is shown that for this to be most advantageous α should be (1) larger when there is a high probability of the simulated dataset Y being a close match to the observations (2) smaller when the expected time to complete the simulation is large. To achieve this α is based on X , a random variable encapsulating some preliminary information about Y .

The final likelihood estimator is based on X and Y . However when early stopping occurs a realised value of zero is obtained without drawing a value of Y . In programming terminology this is an example of *lazy evaluation*, which is the basis for the method's name.

The paper presents theoretical results on the optimal tuning of α in lazy ABC, making precise the two properties just outlined. This choice is asymptotically optimal in terms of maximising efficiency, which is defined as effective sample size (ESS) per unit CPU time. Based on this, a framework for tuning in practice is also presented. The main requirements are the estimation of the probability of close matches and of expected remaining computation times. Both of these are conditional on θ and x (realised values of X). As x is typically high dimensional this estimation is not feasible, so instead it is recommended to base the choice of α on some low dimensional *decision statistics* $\phi(\theta, x)$. A computationally demanding example is presented based on the spatial extremes application of Erhardt and Smith (2012) where lazy ABC increases the efficiency by roughly 3 times for modestly sized data and 8 times for a larger example.

The focus of this paper is on importance sampling, which is widely used by ABC practitioners and particularly amenable to parallelisation. However the lazy ABC approach is also applicable to other algorithms, such as Markov chain Monte Carlo (MCMC) and sequential Monte Carlo (SMC), as discussed in the final section. Another extension is to situations where the exact likelihood, or a non-negative unbiased estimator, is available but expensive to compute. The paper discusses how a less expensive non-negative unbiased estimator can be produced. Essentially the same theoretical results apply as for lazy ABC, but it is argued that practical application is more challenging.

Several recent papers have proposed speeding up ABC by fitting a model to (θ, y) pairs, simulated either in a preliminary stage or in earlier ABC iterations (Buzbas and Rosenberg, 2013; Meeds and Welling, 2014; Moores et al., 2014; Wilkinson, 2014). This model is then sometimes or always used in place of the original model of interest in the inference algorithm. A potential application of lazy ABC is to make use of such approximate models (their

predictions given θ forming X in the notation above) to gain speed benefits without incurring additional approximation errors. More generally, there has been much interest over the past decade in Bayesian inference algorithms with random weights (e.g. Beaumont, 2003; Andrieu and Roberts, 2009; Fearnhead et al., 2010; Tran et al., 2014). A novelty of lazy ABC is that it introduces a random factor to the weights to reduce computation time, rather than to deal with intractability.

Rejection control in sequential Bayesian algorithms (Liu et al., 1998) uses a similar idea to lazy ABC. Here after the first t stages of the sequential analysis, a proposal (typically a sequence of latent states at times $1, 2, \dots, t$) with weight w is allowed to continue with probability $\alpha = \min(1, w/c)$ for some constant c . On continuation the weight is updated to w/α and otherwise a new proposal is generated. Novelty of the current work are finding an optimal form for α and allowing it depend on information other than w .

The remainder of the paper is structured as follows. Section 2 contains background material. To help later developments this presents ABC within the framework of random weight importance sampling. Section 3 gives the lazy ABC algorithm and proves it targets the correct distribution. Section 4 presents theory and practical methods for tuning the algorithm. A related result is also given on the optimal importance distribution for standard ABC importance sampling. Section 5 contains the application to spatial extremes and Section 6 is a discussion. Appendices contain proofs and material on lazy ABC with multiple stopping decisions.

2 Importance sampling

Consider analysing data y_{obs} under a probability model with density $\pi(y|\theta)$ and parameters θ . The likelihood is defined as $L(\theta) = \pi(y_{\text{obs}}|\theta)$. Bayesian inference introduces a prior distribution with density $\pi(\theta)$ and aims to find the posterior distribution $\pi(\theta|y_{\text{obs}}) = \pi(\theta)L(\theta)/\pi(y_{\text{obs}})$, where $\pi(y_{\text{obs}}) = \int \pi(\theta)L(\theta)d\theta$, or at least to estimate the posterior expect-

tation $E[h(\theta)|y_{\text{obs}}]$ of a generic function $h(\theta)$. Importance sampling is a simple method to do this. Parameter values $\theta_{1:N}$ are simulated independently from an importance density $g(\theta)$ and given weights $w_i = L(\theta_i)\pi(\theta_i)/g(\theta_i)$ (n.b. $\theta_{1:N}$ represents the sequence $(\theta_i)_{1 \leq i \leq N}$. Similar notation is used later.) It is assumed throughout that $g(\theta) > 0$ whenever $\pi(\theta) > 0$. Each of the (θ_i, w_i) pairs can be computed in parallel, allowing for efficient implementation.

A Monte Carlo estimate of $E[h(\theta)|y_{\text{obs}}]$ is $\mu_h = \frac{\sum_{i=1}^N h(\theta_i)w_i}{\sum_{i=1}^N w_i}$. Two properties of importance sampling estimates are

$$\mu_h \rightarrow E[h(\theta)|y_{\text{obs}}] \quad \text{almost surely as } N \rightarrow \infty, \quad (1)$$

$$E[N^{-1} \sum_{i=1}^N w_i] = \pi(y_{\text{obs}}). \quad (2)$$

See Geweke (1989) for proof that (1) holds under weak conditions. To prove (2) note that each w_i is an unbiased estimator of $\pi(y_{\text{obs}})$. Estimating this is of interest for model comparison.

2.1 Notation

The remainder of the paper is largely concerned with the distribution of random variables produced in an iteration of various importance sampling algorithms. Henceforth, expectations and probabilities that involve quantities produced by importance sampling should be read as being with respect to this distribution. In particular this means that below the marginal density of θ is taken to be $g(\theta)$. The preceding material in this section is the only time that a marginal density of $\pi(\theta)$ is used instead.

2.2 Random weights

Algorithm 1a describes random weight importance sampling (RW-IS), an importance sampling algorithm in which likelihood evaluations are replaced with random estimates of the likelihood. Under the condition that these estimates are non-negative and unbiased, the algorithm produces valid output, in the sense that (1) and (2) continue to hold. This can

be seen by noting that Algorithm 1a is equivalent to a deterministic weight importance sampling algorithm with augmented parameters (θ, ℓ) , prior density $\pi(\theta)\pi(\ell|\theta)$, importance density $g(\theta)\pi(\ell|\theta)$ and likelihood ℓ . Here ℓ is the realisation of the likelihood estimator and $\pi(\ell|\theta)$ is the conditional density of this estimator. This algorithm gives the correct marginal posterior for θ . See Tran et al. (2014) for a more detailed proof and Fearnhead et al. (2010) for further background, including the observation that non-negativity is not necessary. Non-negative unbiased likelihood estimators can also be used to produce valid MCMC algorithms (Beaumont, 2003; Andrieu and Roberts, 2009).

2.3 Approximate Bayesian computation

Many interesting models are sufficiently complicated that it is not feasible to calculate exact likelihoods or useful (i.e. reasonably low variance) unbiased estimators. ABC algorithms instead base inference on simulation from the model. Algorithm 1b uses this idea in a RW-IS framework to give ABC importance sampling (ABC-IS). The RW-IS likelihood estimator \hat{L} has been replaced with a Bernoulli estimator \hat{L}_{ABC} which equals 1 if $d(S(Y), S(y_{\text{obs}})) \leq \epsilon$ i.e. if the distance between summary statistics of the simulated and observed datasets is less than or equal to a threshold ϵ . This is typically a biased estimate of the likelihood and ABC-IS targets the posterior for an approximate likelihood $L_{\text{ABC}}(\theta) = E[\hat{L}_{\text{ABC}}|\theta]$.

A special case of ABC-IS is when $g(\theta) = \pi(\theta)$. The weights in this case equal zero or one, and it is often referred to as ABC rejection sampling. A generalisation of ABC-IS, considered in Section 6, is to use as a likelihood estimate $K(d(S(Y), S(y_{\text{obs}}))/\epsilon)$, where K is a density function known as the ABC kernel. Algorithm 1b uses a uniform kernel.

As $\epsilon \rightarrow 0$, the target distribution of ABC-IS converges to $\pi(\theta|S(y_{\text{obs}}))$. However, $\epsilon > 0$ is typically required to achieve a reasonable number of non-zero weights, so a trade-off must be made. The observed summary statistics $S(y_{\text{obs}})$ should ideally preserve most of the information on θ available from y_{obs} . However analysis of ABC algorithms shows that the quality of the approximation deteriorates with the dimension of $S(y)$. Therefore the

Input (general):

- Prior density $\pi(\theta)$ and importance density $g(\theta)$.
- Number of iterations to perform N .

Input (RW-IS):

- Likelihood estimator \hat{L} .

Input (ABC):

- Observed data y_{obs} .
- Summary statistics $S(\cdot)$, distance function $d(\cdot, \cdot)$ and threshold $\epsilon \geq 0$.

Algorithm:

Repeat the following steps N times.

- 1 Simulate θ^* from $g(\theta)$.
- 2 Simulate ℓ^* from $\hat{L}|\theta$.
In ABC this is done as follows:
 - 2a Simulate y^* from $Y|\theta^*$
 - 2b Set $\ell^* = \mathbb{1}[d(S(y^*), S(y_{\text{obs}})) \leq \epsilon]$.
- 3 Set $w^* = \ell^* \pi(\theta) / g(\theta)$.

Output:

A set of N pairs of (θ^*, w^*) values.

Algorithm 1: a) Random weight importance sampling (RW-IS) b) ABC importance sampling (ABC-IS); this is the special case where step 2 is implemented by 2a and 2b.

choice of $S(\cdot)$ involves a trade-off between low dimension and informativeness. For further background details on all aspects of ABC see the review articles of Beaumont (2010) and Marin et al. (2012).

3 Lazy ABC

This section defines lazy ABC and shows it produces valid results.

Definition 1. *Lazy ABC is Algorithm 1a, using a likelihood estimator of the form (3) under conditions C1-C3.*

$$\hat{L}_{\text{lazy}} = \begin{cases} \hat{L}_{ABC}/\alpha(\theta, X) & \text{with probability } \alpha(\theta, X) \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

C1 $\alpha(\theta, x)$ is a function with codomain $[0, 1]$.

C2 $\alpha(\theta, x) > 0$ whenever $\Pr(\hat{L}_{ABC} > 0|\theta, x) > 0$.

C3 The random variable X is such that both $X|\theta$ and $Y|\theta, x$ can be simulated from.

The following theorem shows that the estimator \hat{L}_{lazy} can be used in a RW-IS algorithm, or a pseudo-marginal MCMC algorithm, and give valid results.

Theorem 1. *Conditional on θ , \hat{L}_{lazy} is a non-negative unbiased estimator of $L_{ABC}(\theta)$.*

Proof. Non-negativity is immediate. For unbiasedness first observe that $E(\hat{L}_{\text{lazy}}|\theta, x)$ equals zero when $\alpha(\theta, x) = 0$ and $E(\hat{L}_{ABC}|\theta, x)$ otherwise. By C2 if $\alpha(\theta, x) = 0$ then $\Pr(\hat{L}_{ABC} > 0|\theta, x) = 0$ and so $E(\hat{L}_{ABC}|\theta, x) = 0$. Hence $E(\hat{L}_{\text{lazy}}|\theta, X) = E(\hat{L}_{ABC}|\theta, X)$. Taking expectations over X gives the required result. \square

The key feature of the estimator \hat{L}_{lazy} is that it can be evaluated by the following steps given parameter proposal θ^* .

1. Simulate x^* from $X|\theta^*$ and let $a^* = \alpha(\theta^*, x^*)$.
2. With probability a^* continue to step 3. Otherwise output $\ell^* = 0$.
3. Simulate y^* from $Y|\theta^*, x^*$. Calculate ℓ_{ABC}^* , the corresponding value of $\hat{L}_{\text{ABC}}|Y = y^*$.
Output $\ell^* = \ell_{\text{ABC}}^*/a^*$.

The desired behaviour is that simulating from $X|\theta$ is computationally cheap but can be used in steps 1 and 2 to reject many unpromising importance sampling iterations. The expensive part of the likelihood calculation, simulating from $\hat{L}_{\text{ABC}}|\theta, x$, is only carried out when necessary, exploiting lazy evaluation.

A special case of lazy ABC is to take $\alpha(\theta, x) = \mathbb{1}[\Pr(\hat{L}_{\text{ABC}} > 0|\theta, x) > 0]$, making step 2 deterministic. This is useful when X is a partial observation of $S(Y)$ allowing a lower bound to be placed on $d(S(Y), S(y_{\text{obs}}))$.

Several more general examples follow which form the main focus of the paper. Each assumes that Y is a deterministic function of a latent vector $X_{1:p}$ for some p , such that it is possible to simulate from $X_1|\theta$ and $X_i|\theta, x_{1:i-1}$ for all $1 < i \leq p$. Further examples are given in Appendix A which use the same framework to consider multiple stopping decisions.

Example 1: Partial simulation Let $X = X_{1:t}$ for some $t < p$.

Example 2: Partial calculation of S Assume that computing $S(Y)$ involves calculating variables $X'_{1:q}$ which are deterministic transformations of Y , and that this is the most expensive part of simulating \hat{L}_{ABC} . Let $X = (X_{1:p}, X'_{1:t})$ for some $t < q$. (This is applied in Section 5.2.)

Example 3: Random stopping times As for either previous example but with t replaced by a random stopping time variable T . This allows a stopping decision once a particular event has occurred.

Simulation from $X|\theta$ is referred to as the *initial simulation stage* and from $S(Y)|\theta, x$ as

the *continuation simulation stage*. It is often useful later to have $\alpha(\theta, x) = \alpha(\phi(\theta, x))$ where $\phi(\theta, x)$ is referred to as the *decision statistics*.

Notation is now introduced for expected CPU times: $\bar{T}_1(\theta)$ is for steps 1 and 2 above conditional on θ , $\bar{T}_2(\theta, \phi)$ is for step 3 conditional on (θ, ϕ) and $\bar{T}(\theta)$ is for simulation from \hat{L}_{ABC} conditional on θ . The first two are roughly the times of the initial simulation and continuation stages, but are defined to cover all steps involved in simulating from \hat{L}_{lazy} .

It is assumed that

$$\bar{T}(\theta) \leq \bar{T}_1(\theta) + \text{E}[\bar{T}_2(\theta, \phi)|\theta]. \quad (4)$$

Roughly speaking this states that drawing from \hat{L}_{lazy} conditional on no early stopping takes at least as long as drawing from \hat{L}_{ABC} . The difference is due to computational overheads of considering stopping. It is also convenient to define $\bar{T}_1 = \text{E}[\bar{T}_1(\theta)]$ and $\bar{T}_2(\phi) = \text{E}[\bar{T}_2(\theta, \phi)|\phi]$.

3.1 Lazy importance sampling

The above approach can be generalised to non-ABC situations to give *lazy importance sampling* (LIS). This is Algorithm 1a using a likelihood estimator of the form:

$$\hat{L}_{\text{lazy}} = \begin{cases} \hat{L}/\alpha(\theta, X) & \text{with probability } \alpha(\theta, X) \\ 0 & \text{otherwise} \end{cases}$$

In addition to conditions C1 and C2 above assume:

- C4 The distribution $(X, \hat{L})|\theta$ is such that $\hat{L}|\theta$ is a non-negative unbiased estimator of $L(\theta)$, and both $X|\theta$ and $\hat{L}|\theta, x$ can be simulated from.

This framework can be used when \hat{L} is an expensive unbiased estimator. It also allows cases where either or both of X and \hat{L} are non-random. For example, X may be a deterministic approximation of the likelihood and $\hat{L}|\theta$ may be a point mass at $L(\theta)$.

All theorems and proofs of this paper also hold for lazy importance sampling, replacing $L_{\text{ABC}}(\theta)$ and \hat{L}_{ABC} with $L(\theta)$ and \hat{L} , and making other small modifications noted in the text. In particular Theorem 1 shows that given conditions C1, C2 and C4, $\hat{L}_{\text{lazy}}|\theta$ is a non-negative unbiased estimator of $L(\theta)$. However the practical application of lazy importance sampling is challenging as discussed in Section 6.

4 Tuning

There is considerable freedom to tune lazy ABC through the choice of X (when to consider stopping) and α (the function assigning continuation probabilities). Section 4.1 proves a result on the most efficient choice of α . This theory is used in Section 4.2 to motivate practical tuning methods.

Note that the case where $\alpha(\phi)$ is based on $\phi \in A$ for discrete A does not require the theoretical results below. Here $\alpha(\phi)$ values can be selected by numerical optimisation of an estimate of the algorithm's efficiency based on pilot simulations. The methods that follow detail construction of such an estimate.

4.1 Theory

A commonly used tool for the analysis of importance sampling algorithms is the effective sample size (ESS). Liu (1996) argued that typically the variance of the importance sampling estimator is roughly equal to that of N_{eff} independent samples where

$$N_{\text{eff}} = NE(W)^2/E(W)^2,$$

and the random variable W is the weight generated in an iteration of importance sampling. The argument of Liu generalises immediately to RW-IS algorithms through the interpretation of them as importance sampling algorithms on an augmented parameter space given in Section 2.2.

Define efficiency as N_{eff}/T where T is the CPU time of the algorithm (i.e. ignoring any execution time savings due to parallelisation.) Assume that T follows a central limit theorem in N . Then the delta method gives that for large N efficiency asymptotically equals

$$\frac{\text{E}(W)^2/\text{E}(W^2)}{\text{E}(T)/N}.$$

Theorem 2. *Fix some decision statistics $\phi(\theta, x)$. Amongst continuation probability functions of the form $\alpha(\theta, x) = \alpha(\phi(\theta, x))$, asymptotic efficiency is maximised by the following expression for some $\lambda > 0$,*

$$\alpha(\phi) = \min \left\{ 1, \lambda \left[\frac{\text{E}[\hat{L}_{ABC} \frac{\pi(\theta)^2}{g(\theta)^2} | \phi]}{\bar{T}_2(\phi)} \right]^{1/2} \right\}. \quad (5)$$

Proof. See Appendix B.1. □

Remark 1. *Suppose $\pi(\theta)/g(\theta) = u(\phi)$ i.e. this fraction is completely determined by ϕ . For example this is the case in ABC rejection sampling where $g(\theta) \equiv \pi(\theta)$. Then (5) becomes*

$$\alpha(\phi) = \min \left\{ 1, \lambda u(\phi) \left[\frac{\gamma(\phi)}{\bar{T}_2(\phi)} \right]^{1/2} \right\}, \quad (6)$$

where $\gamma(\phi) = \text{E}[\hat{L}_{ABC} | \phi] = \text{Pr}(d(S(Y), S(y_{\text{obs}})) \leq \epsilon | \phi)$.

Remark 2. *Theorem 2 and Remark 1 hold for LIS with \hat{L}_{ABC} replaced by \hat{L}^2 .*

A simple closed form expression for λ does not appear possible. In the practical tuning methods below λ is found numerically, and the behaviour of this numerical estimate investigated by simulation study (see Figure 1B).

By viewing ABC-IS as a special case of lazy ABC, Theorem 2 can be applied to find the optimal choice of $g(\theta)$ for ABC-IS.

Corollary 1. *The asymptotic efficiency of ABC-IS is maximised by $g(\theta) \propto \pi(\theta) \left[\frac{\gamma(\theta)}{\bar{T}(\theta)} \right]^{1/2}$, where $\gamma(\theta) = \text{E}(\hat{L}_{ABC} | \theta)$.*

Proof. See Appendix B.2. □

Remark 3. *A corresponding result to Corollary 1 holds for RW-IS with $\gamma(\theta) = \mathbb{E}(\hat{L}^2|\theta)$.*

Remark 4. *The special case of Corollary 1 with $\bar{T}(\theta)$ constant matches the result of Appendix A in Fearnhead and Prangle (2012).*

Note that it is not clear what the optimal choice of $g(\theta)$ is for lazy ABC. The examples later use typical choices from the ABC literature, but a better choice may improve lazy ABC performance further.

4.2 Methods

Theorem 2 motivates choosing α by estimating (6). This section details a method to implement this approach. Its effectiveness is discussed in Section 6.

Tuning begins with a pilot run of N' iterations of ABC-IS. This is used to estimate $\gamma(\phi)$ and $\bar{T}_2(\phi)$ for various choices of X and ϕ , considering only ϕ is such that Remark 1 can be applied. Under each of these choices, λ is found by numerically maximising an estimate of efficiency. The optimal choice of X and ϕ is then made. Note that estimation of $\gamma(\phi)$ and $\bar{T}_2(\phi)$ is challenging if ϕ is high dimensional e.g. for $\phi = (\theta, x)$. Therefore a low dimensional ϕ is recommended. To ensure Remark 1 applies $\pi(\theta)/g(\theta)$ can form one component of ϕ if necessary. Following tuning, N iterations of lazy ABC are performed (unless the estimated efficiency gains are judged inadequate). Detailed comments on several aspects of this method follow.

4.2.1 Estimation of $\bar{T}_2(\phi)$

It may often suffice to treat $\bar{T}_2(\phi)$ as constant and estimate it as the mean CPU time of the continuation stage in the pilot run. This is the case if knowledge of the simulation process shows the number of computational operations required is unaffected by ϕ , or if the pilot run

shows $\bar{T}_2(\phi)$ varies little relative to $\gamma(\phi)$. Alternatively, statistical methods such as regression can be used for estimation, which is straightforward when ϕ is low dimensional.

4.2.2 Estimation of $\gamma(\phi)$

Estimation of $\gamma(\phi)$ is more difficult. Two approaches are suggested: the “standard” approach, producing $\hat{\gamma}^{(1)}$, attempts accurate estimation but involves strong assumptions; the “conservative” approach, producing $\hat{\gamma}^{(2)}$, sacrifices accuracy to improve robustness. They are based on two equivalent expressions for $\gamma(\phi)$: $\Pr(d(S(Y), S(y_{\text{obs}})) \leq \epsilon|\phi)$ and $E[\hat{L}_{\text{ABC}}|\phi]$. Examples of successful implementations of both approaches are given in Section 5.

The standard approach is to directly model the relationship between $\phi(\theta, X)$ and $d(S(Y), S(y_{\text{obs}}))$ and use this to estimate $\Pr(d(S(Y), S(y_{\text{obs}})) \leq \epsilon|\phi)$. However a difficulty is that for most ϕ values this involves extrapolating into the tails of the distribution of $d(S(Y), S(y_{\text{obs}}))|\phi$. See Figure 1A for example. This creates a danger of underestimating the optimal α values and potentially producing very large importance sampling weights.

The conservative approach is to select ϵ_1 following the pilot run such that a sufficiently large number of its simulations $y_{1:N'}$ satisfy $d(S(y_i), S(y_{\text{obs}})) \leq \epsilon_1$. Let z_i be indicator variables denoting meeting this condition and model the relationship between z_i and the simulated ϕ_i values. One method, used in the application later, is non-parametric logistic regression following Wood (2011). This approach is effectively tuning the method based on an ϵ value larger than that of interest. This is an inefficient way to sample from the target of interest. However, if tuning can be done well for the larger ϵ value then this method is safe from producing any dangerously large importance weights, as discussed further in Section 4.2.5. Nonetheless, for ϕ regions where there are no $z_i = 1$ values the conservative estimate is still based on extrapolation and unlikely to be accurate. Consequences of this are discussed in Section 6.

4.2.3 Estimating efficiency

The tuning method outlined above requires the use of N' pilot run iterations to estimate the efficiency of lazy ABC under various choices of tuning details (in particular X , ϕ and α). It is sufficient to estimate $[E(W^2)E(T)]^{-1}$, as this equals efficiency up to a constant of proportionality. This can be used to estimate efficiency relative to ABC-IS, which is a particularly interpretable form of the results as it shows the efficiency improvement of using lazy ABC.

Assume that for a particular choice of tuning details the following are available for $1 \leq i \leq N'$: $t_i^{(1)}$ - initial simulation stage time; $t_i^{(2)}$ - continuation simulation stage time; α_i - continuation probability; $\hat{\gamma}_i$ - estimate of $E(\hat{L}_{\text{ABC}}|\phi_i)$; u_i - ratio $\pi(\theta)/g(\theta)$. An estimate up to proportionality of efficiency is then $[\widehat{W^2\hat{T}}]^{-1}$ where $\widehat{W^2} = N'^{-1} \sum_{i=1}^{N'} u_i^2 \hat{\gamma}_i / \alpha_i$ and $\hat{T} = \sum_{i=1}^{N'} t_i^{(1)} + \sum_{i=1}^{N'} \alpha_i t_i^{(2)}$. An estimate of efficiency of ABC-IS is formed by taking $\alpha \equiv 1$. Note that this typically overestimates T due to the overheads of considering stopping (see (4)). A more precise estimate would be possible using further pilot simulations of standard ABC.

4.2.4 Combining pilot and main run output

To make efficient use of the pilot run, it can be used in the final output as well as for tuning. This is done by appending the pilot sequence of (θ, w) pairs to that from the main algorithm. Loosely speaking, since each individual sequence targets the same distribution, so does the combined sequence. More technically, it is straightforward to see that ABC versions of relations (1) and (2) are roughly true for the combined sequence when N and N' are large, and are exactly true as $N \rightarrow \infty$ regardless of N' . Also note that on appending the sequences, gains in efficiency are possible by multiplying the weights of one sequence by a constant, but this is not implemented here as little improvement was observed in the application later.

4.2.5 Choice of ϵ

In ABC-IS, an appropriate value ϵ is often unknown a priori and is instead chosen based on the simulated $d(S(Y), s(y_{\text{obs}}))$ values. For lazy ABC in this situation one can use the pilot run to select a preliminary conservative choice of ϵ_1 as in Section 4.2.2 and perform lazy ABC with $\epsilon = \epsilon_1$. Alternative values of ϵ can then be investigated by updating the realised \hat{L}_{ABC} values in the weight calculations. For $\epsilon < \epsilon_1$ this simply reduces the number of non-zero weights. However $\epsilon \gg \epsilon_1$ is not recommended as this may introduce large weights and destabilise the importance sampling approximation.

5 Example: spatial extremes

This section uses lazy ABC in a computationally demanding application of ABC to spatial extremes introduced by Erhardt and Smith (2012).

5.1 Background

The observation $y_{t,d}$ represents the maximum measurement (e.g. of rainfall) during year t at location $x_d \in \mathbb{R}^2$. There are D locations and T years. The data are treated as T independent replications of a spatial distribution. Several models based on extreme value theory have been proposed, and Erhardt and Smith concentrate on the *Schlather process* (Schlather, 2002). This is based on independent identically distributed mean zero stationary Gaussian processes $Y_i(x)$ where $i = 1, 2, \dots$ and $x \in \mathbb{R}^2$. The correlation between locations x and x' is given by the correlation function $\rho(h)$ where $h = \|x - x'\|_2$. Let s_i be draws from a Poisson process on $(0, \infty)$ with intensity $\mu^{-1}s^{-2}$, where $\mu = \text{E}[\max(0, Y(x))]$. Then the Schlather process is

$$Z(x) = \max_i s_i \max(0, Y_i(x)).$$

Erhardt and Smith focus on the Whittle-Matérn correlation function with zero nugget

$$\rho(h; c, \nu) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{h}{c}\right)^\nu K_\nu\left(\frac{h}{c}\right),$$

where Γ is the gamma function and K_ν is the modified Bessel function of the third kind with order ν . This has two parameters: range $c > 0$ and smoothness $\nu > 0$.

A density function for the Schlather process is not available for $D > 2$, making inference difficult. Schlather (2002) provides a near-exact algorithm to simulate from the process based on only a finite number of copies of Y_i , motivating the use of ABC by Erhardt and Smith. They applied ABC rejection and importance sampling with a uniform prior on $[0, 10]^2$ and investigated several choices of summary statistics. The analysis here focuses on the choice they find most successful, based on *tripletwise extremal coefficient estimators*. Given a triple of 3 locations, i, j, k , this estimator is

$$\hat{\theta}_{ijk} = \frac{T}{\sum_{t=1}^T 1/\max(y_{t,i}, y_{t,j}, y_{t,k})}.$$

There are $O(D^3)$ such summaries, so Erhardt and Smith calculate a vector m of mean values within 100 clusters of triples, and use these as summary statistics. Their clustering process finds triples of similar shapes, ignoring differences of location and rotation. The ABC distance function between two vectors m_1 and m_2 of cluster means is

$$d(m_1, m_2) = \sum_{i=1}^{100} |m_{1i} - m_{2i}|. \tag{7}$$

Although applying dimension reduction techniques to such high dimensional summaries has been shown to often improve ABC results (Fearhead and Prangle, 2012), this is not investigated here as the aim is to investigate the efficiency improvements of lazy ABC.

Implementing the analysis below used the R packages “SpatialExtremes” (Ribatet et al., 2013) to simulate from the Schlather process and “ABCExtremes” to implement some details

of the approach of Erhardt and Smith.

5.2 Methods

Exploratory investigation of ABC code with $D = 20$ and $T = 100$ showed that the majority of time was spent simulating the data (7.1ms/iteration) and calculating extremal coefficient estimates (17.9ms/iteration), with the remaining steps being brief (3.1ms/iteration). The time costs of the first two of these scaled with D as roughly proportional to D and D^3 respectively, so the latter is expected to dominate for large D . Furthermore, interrupting and then resuming operations during the calculation of extremal coefficients is much simpler to implement than during simulation of data. Therefore the initial simulation stage of the lazy ABC analysis was chosen to be simulating the data at all locations, and extremal coefficient estimates at a subset of locations L . The continuation simulation stage was to calculate the remaining extremal coefficient estimates.

The decision statistic \hat{d} was constructed as follows. Let m_{1i} be the i th cluster mean for the observed data. Let \hat{m}_{2i} be the i th cluster mean for the simulated data using only extremal coefficient estimates available at the initial simulation stage, and B be the set of clusters for which any such estimates are available. Then define $\hat{d} = \sum_{i \in B} |m_{1i} - \hat{m}_{2i}|$. This is an estimate of the ABC distance d (7). It could be improved by estimating typical \hat{m}_{2i} values for $i \notin B$ but including such constant terms has no effect on the analysis below.

It was assumed that $\bar{T}_2(\hat{d})$ is constant as, given D , L and T , the continuation stage always involves the same number of calculations. The value was estimated by the mean CPU time for this stage in the pilot run. Analyses were performed using both the standard and conservative γ estimators. To calculate $\hat{\gamma}_1(\hat{d})$ from pilot run output, the relationship between \hat{d} and d was modelled statistically. Exploratory analysis showed that there was a roughly linear relationship, but for some choices of L this was heteroskedastic (see Figure 1A). Furthermore, for small \hat{d} the distribution of $d|\hat{d}$ was skewed. So $\Pr(d \leq \epsilon|\hat{d})$ was estimated based on a linear regression of d on \hat{d} with a Box-Cox transformation, using only

simulations with nearby values of \hat{d} . This was done for several \hat{d} values and interpolated estimates elsewhere formed $\hat{\gamma}_1(\hat{d})$. For the importance sampling case, $\log u$ was also included in each regression so that a number of functions mapping u to estimates of $\Pr(d \leq \epsilon | \hat{d}, u)$ for various \hat{d} values were produced which were used for interpolation. Calculation of $\hat{\gamma}_2$ was as described in Section 4.2.2, taking ϵ_1 to give 100 acceptances in the pilot run. Given estimates of \bar{T}_2 and γ , tuning was performed as described in Section 4.2, with optimisation over possible choices of L by backwards selection.

Three simulation studies were performed. The first replicated the rejection sampling analysis of Erhardt and Smith on several simulated datasets. These used $D = 20, T = 100$ and true parameter values shown in Table 1. Each dataset used a different set of observation locations with integer coordinates sampled from $[0, 10]^2$. The first analysis was a replication of the standard ABC analysis, using ϵ values corresponding to 200 acceptances. Then lazy ABC was performed on the same datasets under each method of estimating γ . To compare the methods fairly, lazy ABC used the same ϵ value as standard ABC and reused its random seeds so that the sequence of (θ, X, Y) realisations is also the same.

The second simulation study investigated rejection sampling for a single larger simulated dataset with $D = 35, c = 0.5$ and $\nu = 1$. Locations were chosen as before. As in a real application ϵ was not assumed to be known in advance and the approach of Section 4.2.5 was used to select this post-hoc. A complication for this dataset was that the simulation of Gaussian processes was difficult when both parameters were large: the default “direct method”, based on Choleski decomposition, sometimes produced numerical errors. Simulation was possible via the turning bands method (TBM) but much slower (roughly 150 times the CPU time). A two stage simulation method was implemented. First the direct method was attempted and if this failed TBM was used. To save time lazy ABC was implemented with multiple stopping decisions, the first taking place after attempting the direct method. This has a binary decision statistic indicating success or failure. The second stopping decision is as described earlier. Tuning was performed as described in Appendix A.1.2, using $\hat{\gamma}_1$ fitted as described

above by either the standard or conservative tuning method. The standard method used ϵ_1 to give 30 acceptances in the pilot run. As before all analyses reused the same random seeds.

Finally an importance sampling analysis was performed on the larger dataset. A sample of 10^4 log parameter values was taken from simulations of the preceding standard ABC analysis with distances below the 0.3 quantile. A Gaussian mixture distribution was constructed with locations given by this sample and variances equal to twice the empirical variance of the sample. After truncation to the prior support, this was used to give $g(\theta)$, where θ now represents the log parameters. This choice follows the suggestions of Beaumont et al. (2009), noting that using the log scale produced a better fit to the sample and that the subsample was used to avoid slow density calculations. The preceding $D = 35$ analysis was then repeated. As discussed in Section 4.2, $u = \pi(\theta)/g(\theta)$ was included as a decision statistic. Estimation of γ and \bar{T}_2 was performed as before with u included in the γ estimate as described earlier.

All ABC analyses performed 10^6 total iterations. For lazy ABC 10^4 of these comprised the pilot run.

5.3 Results

Figure 1 illustrates some details of tuning for one case of the $D = 20$ study. The results are shown in Table 1. For all datasets lazy ABC is roughly 4 times more efficient under the standard tuning method and 3 times under the conservative method. Efficiency gains for conservative tuning are slightly less than estimated. This is because the estimate is made for a choice of ϵ_1 larger than the final ϵ . The mean weights were also investigated, as these are useful in model selection as an estimate of $\pi(y_{\text{obs}})$. All lazy ABC estimates differed from the standard ABC estimate by no more than 4%.

Table 2 shows results for the $D = 35$ dataset. In the initial rejection sampling analysis lazy ABC improved efficiency by roughly 8 times. For importance sampling the improvement factor is 2, showing that lazy ABC still improves efficiency, although this is harder when g concentrates on plausible choices of θ . For example, standard ABC now spends negligible

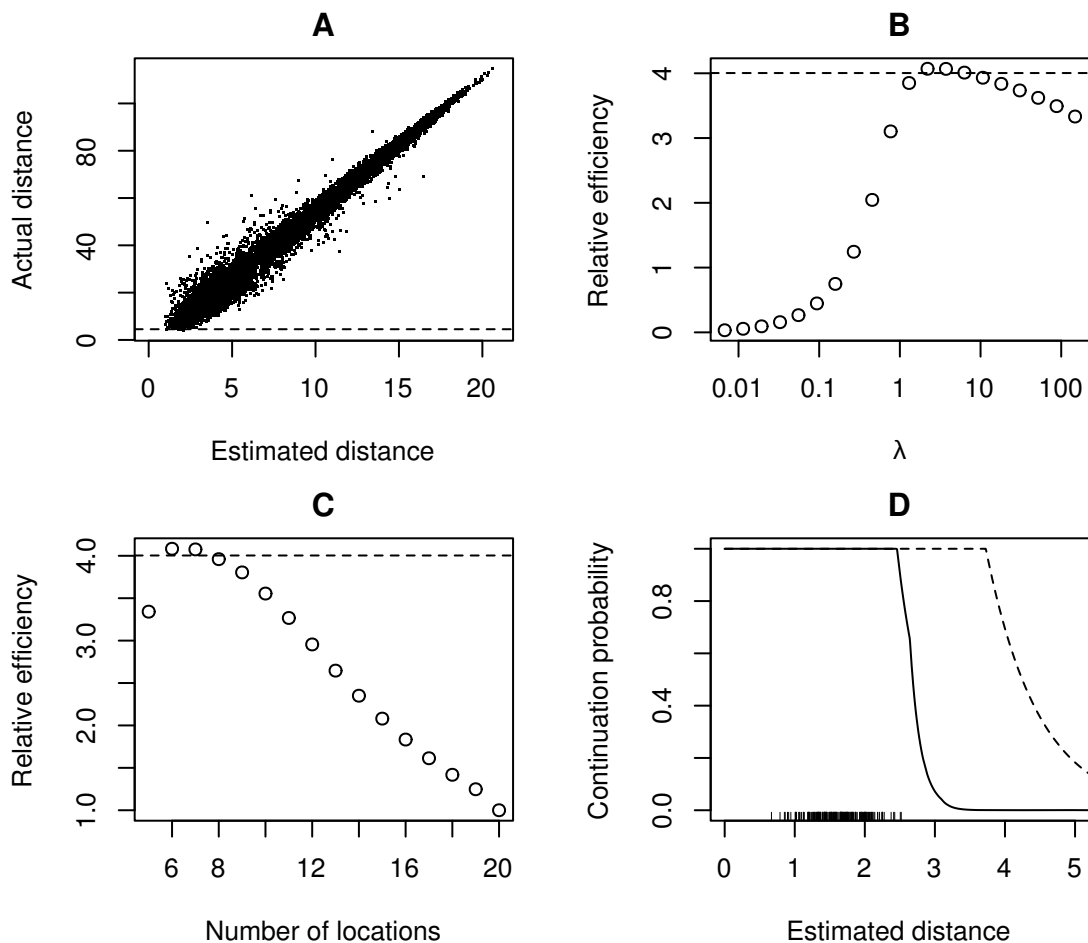


Figure 1: Details of a simulation study applying lazy ABC to spatial extremes corresponding to the first row of Table 1. Panels A-C concentrate on the standard tuning approach. *Panel A* Pilot run values of \hat{d} and d . The dashed line shows the value of ϵ . *Panel B* Estimated efficiency for different values of λ . The dashed line shows the realised efficiency. *Panel C* Estimated efficiency for the best choices of L of various lengths output by backwards selection. The dashed line shows the realised efficiency. *Panel D* Values of \hat{d} and α from non-pilot simulations under standard (solid line) and conservative (dashed line) tuning. The marks on the horizontal axis indicate the simulations which resulted in positive weights. (For this panel conservative tuning was performed using L as selected by standard tuning.)

Range	Smooth	Standard	Lazy			Relative efficiency	
		Time (10 ³ s)	Time (10 ³ s)	Sample size	ESS	Estimated	Actual
0.5	1	32.0	8.0 (11.6)	196 (199)	196.0 (198.7)	4.08 (3.28)	4.00 (2.79)
1	1	31.3	7.3 (9.8)	200 (200)	199.9 (200.0)	4.34 (4.31)	4.35 (3.25)
1	3	31.3	8.2 (11.2)	194 (198)	182.5 (196.5)	3.77 (3.43)	3.51 (2.79)
3	1	31.2	7.7 (11.1)	194 (200)	189.9 (200.0)	4.18 (3.56)	3.89 (2.86)
3	3	31.2	7.4 (11.0)	192 (199)	175.8 (199.0)	4.43 (3.65)	3.79 (2.87)
5	3	31.3	8.3 (11.1)	200 (200)	200.0 (200.0)	3.73 (3.49)	3.85 (2.87)

Table 1: Simulation study replicating Erhardt and Smith (2012). For each dataset a choice of ϵ was made under standard ABC so that the accepted sample size (and therefore ESS) was 200, and the same value was used for lazy ABC. Lazy ABC figures are shown for both the standard $\hat{\gamma}$ estimate and, in brackets, the conservative estimate. The lazy ABC output includes the pilot run as described in Section 4.2.4, and also includes the tuning time (roughly 120 seconds for the standard approach and 210 for the conservative). Iterations were run in parallel and computation times are summed over all cores used. For all datasets efficiency (ESS/time) to 1 significant figure was 0.006 for standard ABC and 0.02 or 0.03 under either approach to lazy ABC.

time on TBM simulations. As before lazy ABC estimates of $\pi(y_{\text{obs}})$ differed from those of standard ABC by no more than 4%. In both cases a post-hoc selection of ϵ has been used successfully.

Table 2 shows that under rejection sampling the standard tuning ESS is considerably smaller than the sample size. This is due to two simulations which are given importance weights of 10. These have $\hat{\gamma}$ values of roughly 10^{-4} which appears to be an underestimate. Conservative tuning avoids large importance weights to give an ESS of roughly 200 and improves the relative efficiency for the same final choice of ϵ . For importance sampling conservative tuning also performs better. The reason here is not obvious but may be a better final selection of L .

6 Discussion

This paper has introduced lazy ABC, a method to speed up inference by ABC importance sampling without introducing further approximations. The approach is to abandon some unpromising simulations before they are complete. By using a probabilistic stopping rule

	ϵ	Standard			Lazy			Relative efficiency
		Time (10 ³ s)	Sample size	ESS	Time (10 ³ s)	Sample size	ESS	
RS standard	2.61	241.4	210	210	22.8	200	139.6	7.0
RS conservative	2.61	241.4	207	207	25.5	200	197.7	9.0
IS standard	2.33	136.4	209	168	61.9	200	165	2.2
IS conservative	2.33	136.4	209	168	51.0	200	162	2.6

Table 2: Simulation study on a spatial extremes dataset with $D = 35$. Results are shown for rejection and importance sampling with standard and conservative tuning. The rejection sampling output was used to create the importance density. The final choice of ϵ is shown. For IS the two ϵ values are equal but there is a small difference for RS. The lazy ABC output includes the pilot run and the tuning time.

and weighting the accepted simulations accordingly, the algorithm targets exactly the same distribution as standard ABC, in the sense that Monte Carlo estimates of functions $h(\theta)$ and of the model evidence converge to unchanged values.

Results have been provided on the optimal tuning of the lazy ABC stopping rule and used to motivate a practical tuning method. This has been demonstrated for a computationally challenging application where it has produced improvements in efficiency (ESS/CPU time) over standard ABC of up to 8 times. One case of this application involved multiple stopping decisions. This illustrated two potential uses of lazy ABC: firstly to consider stopping every simulation based on whether it appears promising, secondly to consider stopping after particular events which are suspected a priori to indicate unpromising results.

The tuning method is based on estimating the optimal choice of $\alpha(\phi)$, (6). The most difficult part was estimating $\gamma(\phi) = \Pr(d(S(Y), S(y_{\text{obs}})) \leq \epsilon | \phi)$ from pilot run data. Two approaches to this were described, a standard approach of direct estimation and a conservative approach of estimation using a larger ϵ value than is of interest for ABC. The latter approach improves robustness and make estimation simpler at the cost of some inefficiency. Both approaches performed well in the simulation studies but some improvements are desirable. Firstly, estimation of $\gamma(\phi)$ involves extrapolation which may produce inaccurate results. Secondly, several choices by the user are required, especially for the standard approach. A more automated approach would be useful for lazy versions of ABC SMC algorithms, where a new choice of α would be needed for each ϵ value, or alternatively for lazy ABC algo-

rithms which adapt α as more simulations become available. It would be of interest to find suboptimal but robust choices of α addressing these issues.

Lazy ABC with multiple stopping decisions is an extension to the framework of the main paper and is described in Appendix A. A tuning method is given when the decision statistics for all stopping decisions are discrete, and also some cases where one decision statistic is continuous. For more complex cases tuning results are not available. For now it is recommended to discretise most decision statistics to avoid this difficulty.

Also, Section 3.1 showed that a generalisation to the non-ABC setting, lazy importance sampling, is possible, and the theoretical results of the paper carry over to this. However exploratory analysis suggests tuning this in practice is more challenging than lazy ABC. This is because it is necessary to estimate $\gamma(\phi) = E(\hat{L}^2|\phi)$ (see Remark 2), and this expectation can be strongly influenced by the upper tail of $\hat{L}|\phi$ which is hard to estimate from pilot run output. For lazy ABC, $\hat{L}_{\text{ABC}}|\phi$ is Bernoulli avoiding this difficulty. A related point is that lazy ABC can be generalised to allow a non-uniform ABC kernel. This gives \hat{L}_{ABC} with a known upper bound so that estimation of $\gamma(\phi)$ seems feasible.

This paper has concentrated on importance sampling, which is widely used by ABC practitioners, but the lazy ABC approach can be extended to ABC versions of MCMC and SMC, which are more efficient algorithms. The tuning results are applicable to SMC algorithms, but further practical methods are needed, as mentioned above. Further theory on optimal tuning is necessary for MCMC, although good performance may be possible with ad-hoc tuning. Examining the connections between lazy ABC and rejection control (Liu et al., 1998) may also be fruitful, especially to design algorithms in which partial X simulations are resampled and continued many times.

Acknowledgements: Thanks to Chris Sherlock and Richard Everitt for helpful conversations, and Robert Erhardt for advice on the ABCExtremes package.

A Multiple stopping decisions

The lazy ABC framework of Section 3 allows multiple stopping decisions, as follows. As in that section assume Y is a deterministic transformation of a latent vector $X_{1:p}$.

Example A1: Multiple stopping decisions Let $X = X_{1:p}$ and $\alpha(\theta, x) = \prod_{i=1}^s \alpha^{(i)}(\theta, x_{1:t_i})$. Thus, for each $1 \leq i \leq s$, once simulation of $X_{1:t_i}$ has been performed then \hat{L}_{lazy} is set to zero with a certain probability, in which case no further simulation is necessary. It is often be useful to let $\alpha^{(i)}(\theta, x_{1:t_i}) = \alpha^{(i)}(\phi_i(\theta, x_{1:t_i}))$. That is, each stopping decision has associated decision statistics ϕ_i .

Example A2: Multiple random stopping times As for Example A1 but with each t_i replaced with a random stopping time variable T_i . This permits stopping to be considered when various random events occur, without imposing a fixed order of occurrence.

The following alternative characterisation of these examples is useful below.

Lemma 1. *For any $1 \leq i \leq s$, Examples A1 and A2 can be represented as a lazy importance sampling algorithm with continuation probability $\alpha^{(i)}(\phi_i)$ and*

$$\hat{L} = \begin{cases} \hat{L}_{ABC}/\beta_i(\theta, X) & \text{with probability } \beta_i(\theta, X) \\ 0 & \text{otherwise,} \end{cases}$$

where $\beta_i(\theta, x) = \prod_{j \neq i} \alpha^{(j)}(\phi_j)$.

Proof. The likelihood estimator stated can easily be verified to have the same distribution as \hat{L}_{lazy} . □

It is also helpful to define $\bar{T}_{2i}(\theta, \phi_{1:s})$ as the expected time remaining from the calculation of ϕ_i until the likelihood estimate is computed conditional on θ and $\phi_{1:s}$, and $\bar{T}_{2i}(\phi_i) = \mathbb{E}[\bar{T}_{2i}(\theta, \phi_{1:s})|\phi_i]$.

A.1 Tuning

The efficiency estimate of Section 4.2.3 can be used in a multiple stopping decision setting given a choice of α . It is necessary to update the estimator \hat{T} given there which is usually a straightforward task. Sections A.1.1 and A.1.2 describe situations of practical interest where the optimal form of α can be derived. However in general the problem is challenging, as illustrated by Section A.1.3.

A.1.1 Discrete decision statistics

Suppose $\alpha(\theta, x) = \prod_{i=1}^s \alpha^{(i)}(\phi_i)$ where $\phi_i(\theta, x)$ takes values in $\{1, 2, \dots, d_i\}$ for d_i finite. Tuning requires selecting a finite number of $\alpha^{(i)}(\phi_i)$ values to optimise the efficiency estimate, which is possible by standard numerical optimisation methods. However note that producing an efficiency estimate as in Section 4.2.3 becomes difficult for large s .

A.1.2 One continuous decision statistic

Suppose $\alpha(\theta, x) = \prod_{i=1}^s \alpha^{(i)}(\phi_i)$ where $\phi_1(\theta, x)$ is continuous and $\phi_i(\theta, x)$ is as in Section A.1.1 for $i > 1$. Also suppose there exists $u_1(\phi_1) = \pi(\theta)/g(\theta)$, so that this fraction is fully determined by ϕ_1 . Applying Lemma 1, Theorem 2 and Remarks 1 and 2 gives that efficiency is optimised by

$$\alpha^{(1)}(\phi_1) = \min \left\{ 1, \lambda u_1(\phi_1) \left[\frac{\gamma_1(\phi_1)}{\bar{T}_{21}(\phi_1)} \right]^{1/2} \right\}, \quad (8)$$

where $\gamma_i(\phi_1) = \mathbb{E}[\zeta(\beta_i(\theta, X)) \mathbb{1}\{d(S(Y), S(y_{\text{obs}})) \leq \epsilon\} | \phi_i]$, $\zeta(0) = 0$ and $\zeta(x) = x^{-1}$ for $x > 0$.

In general $\gamma_1(\phi_1)$ and $\bar{T}_{21}(\phi_1)$ depend on $\alpha^{(i)}(\phi_i)$ for $i > 2$ and so must be estimated several times during the tuning process which is costly. A special case where this can be avoided is when $\phi_{2,p}$ is fully determined by ϕ_1 (and so typically the decision associated with α_1 is guaranteed to occur last). For example this is the situation in the second simulation study of Section 5.2.

A.1.3 Multiple continuous decision statistics

Consider the setting of A.1.2 with the modification that every $\phi_i(\theta, x)$ is continuous and there exists a corresponding function $u_i(\phi_i) = \pi(\theta)/g(\theta)$. The same approach as above gives equations of the form

$$\alpha^{(i)}(\phi_i) = \min \left\{ 1, \lambda_i u_i(\phi_i) \left[\frac{\gamma_i(\phi_i)}{\bar{T}_{2i}(\phi_i)} \right]^{1/2} \right\},$$

for $i = 1, \dots, s$. The definition of γ_i involves $\alpha^{(j)}$ for all $j \neq i$, and \bar{T}_{2i} will also involve many of these terms. Thus deriving the optimal $\alpha^{(i)}$ functions involves solving a complicated system of non-linear implicit equations.

B Tuning proofs

All results are proved for the general case of LIS as described in Section 3.1. For lazy ABC replace \hat{L} with \hat{L}_{ABC} .

B.1 Proof of Theorem 2

In LIS the importance sampling weight W equals $\frac{\hat{L}\pi(\theta)}{\alpha(\phi)g(\theta)}$ with probability $\alpha(\phi)$ and zero otherwise. Hence:

$$\mathbb{E}(W^2) = \int \frac{\mathbb{E}[\hat{L}^2|\theta, \phi, y]\pi(\theta)^2}{\alpha(\phi)g(\theta)^2} \pi(\phi, y|\theta)g(\theta)d\theta d\phi dy = \int \frac{\xi(\phi)}{\alpha(\phi)} g(\phi)d\phi, \quad (9)$$

where $\xi(\phi) = \mathbb{E} \left[\hat{L}^2 \left(\frac{\pi(\theta)}{g(\theta)} \right)^2 \middle| \phi \right]$ (which equals $\mathbb{E} \left[\hat{L}_{\text{ABC}} \left(\frac{\pi(\theta)}{g(\theta)} \right)^2 \middle| \phi \right]$ in the ABC case) and $g(\phi) = \int \pi(\phi|\theta)g(\theta)d\theta$.

The expected time of a single iteration of the LIS algorithm is

$$\mathbb{E}(T)/N = \bar{T}_1 + \int \alpha(\phi)\bar{T}_2(\theta, \phi)\pi(\phi|\theta)g(\theta)d\theta d\phi = \bar{T}_1 + \int \alpha(\phi)\bar{T}_2(\phi)g(\phi)d\phi. \quad (10)$$

Note that $E(W)$ is a constant, so choosing $\alpha(\phi)$ to maximise the expression for asymptotic efficiency in Section 4.1 is equivalent to minimising $E(W^2)E(T)/N$. Call this problem P . Consider also the problems $P(v)$, minimising $E(W^2)$ under the constraint $E(T)/N = v$ and $P(v, \mu)$, minimising $E(W^2) + \mu[E(T)/N - v]$, or equivalently

$$\int \left[\frac{\xi(\phi)}{\alpha(\phi)} + \mu\alpha(\phi)\bar{T}_2(\phi) \right] g(\phi)d\phi. \quad (11)$$

Note that $P(v, \mu)$ is a Lagrange multiplier form of $P(v)$. Consider only $\mu > 0$. Also, let Υ be the set of $E(T)/N$ values attainable by some choice of α .

First we consider minimising (11) subject to $0 \leq \alpha(\phi) \leq 1$. This can be done by pointwise optimisation of the integrand. With α unconstrained the solution is

$$\alpha^*(\phi) = \lambda \left[\frac{\xi(\phi)}{\bar{T}_2(\phi)} \right]^{1/2}, \quad (12)$$

where $\lambda = \mu^{-1/2}$. Also note that $\alpha^*(\phi)$ may sometimes be infinite. The derivative of the integrand with respect to α is negative for $\alpha < \alpha^*$. Hence if $\alpha^*(\phi) > 1$, the constrained solution is $\alpha(\phi) = 1$, giving the global solution (5) from the theorem statement.

Substituting (5) into (9) and (10) shows that the resulting values of $E(W^2)$ and $E(T)/N$ are continuous in λ . Furthermore all $E(T)/N$ values in Υ are attainable by (5) under some choice of λ . Hence given $v \in \Upsilon$ there is some μ^* for which the solution to $P(v, \mu^*)$ has $E(T)/N = v$. This must also be a solution to $P(v)$ since otherwise a superior choice of α for $P(v)$ is also superior for $P(v, \mu^*)$. Now choose v^* so that the solution to $P(v^*)$ minimises $E(W^2)E(T)/N$. This must be a solution to P since otherwise a superior choice of α for P is superior to the solution already found for some $P(v)$.

B.2 Proof of Corollary 1

RW-IS can be seen as a special case of LIS where $\phi = \theta$, $\bar{T}_1(\theta) = 0$ and $\bar{T}_2(\phi) = \bar{T}(\theta)$. Repeating the working above to optimise the choice of $\alpha(\theta)g(\theta)$ gives the unconstrained

solution:

$$\alpha(\theta)g(\theta) = \lambda\pi(\theta) \left[\frac{\gamma(\theta)}{\bar{T}(\theta)} \right]^{1/2}. \quad (13)$$

(Recall the LIS definition of $\gamma(\theta)$ from Remark 3 and note that the lazy ABC definition in Corollary 1 can be derived from this.) Various choices of α , such as $\alpha \equiv 1$, give a solution which also meets the constraint on α . These all give algorithms which are equivalent to RW-IS with $g(\theta) \propto \pi(\theta) \left[\frac{\gamma(\theta)}{\bar{T}(\theta)} \right]^{1/2}$ as claimed.

References

- Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, pages 697–725.
- Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164(3):1139–1160.
- Beaumont, M. A. (2010). Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution and Systematics*, 41:379–406.
- Beaumont, M. A., Cornuet, J. M., Marin, J.-M., and Robert, C. P. (2009). Adaptive approximate Bayesian computation. *Biometrika*, 96(4):983–990.
- Buzbas, E. O. and Rosenberg, N. A. (2013). AABC: approximate approximate Bayesian computation when simulating a large number of data sets is computationally infeasible. *Preprint*. Available at <http://www.arxiv.org/abs/1301.6282>.
- Erhardt, R. J. and Smith, R. L. (2012). Approximate Bayesian computing for spatial extremes. *Computational Statistics & Data Analysis*, 56(6):1468–1481.
- Fearnhead, P., Papaspiliopoulos, O., Roberts, G. O., and Stuart, A. (2010). Random-weight particle filtering of continuous time processes. *Journal of the Royal Statistical Society: Series B*, 72(4):497–512.

- Fearnhead, P. and Prangle, D. (2012). Constructing summary statistics for approximate Bayesian computation: Semi-automatic ABC. *Journal of the Royal Statistical Society, Series B*, 74:419–474.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, 57(6):1317–1339.
- Liu, J. S. (1996). Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing*, 6(2):113–119.
- Liu, J. S., Chen, R., and Wong, W. H. (1998). Rejection control and sequential importance sampling. *Journal of the American Statistical Association*, 93(443):1022–1031.
- Marin, J.-M., Pudlo, P., Robert, C. P., and Ryder, R. J. (2012). Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180.
- Meeds, E. and Welling, M. (2014). GPS-ABC: Gaussian process surrogate approximate Bayesian computation. *Preprint*. Available at <http://www.arxiv.org/abs/1401.2838>.
- Moore, M. T., Mengersen, K., and Robert, C. P. (2014). Pre-processing for approximate Bayesian computation in image analysis. *Preprint*. Available at <http://www.arxiv.org/abs/1403.4359>.
- Ribatet, M., Singleton, R., and R Core team (2013). *SpatialExtremes: Modelling Spatial Extremes*. R package version 2.0-0.
- Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes*, 5(1):33–44.
- Tran, M.-N., Scharth, M., Pitt, M. K., and Kohn, R. (2014). Importance sampling squared for Bayesian inference in latent variable models. *Preprint*. Available at <http://www.arxiv.org/abs/1309.3339>.
- Wilkinson, R. D. (2014). Accelerating ABC methods using Gaussian processes. *Preprint*. Available at <http://www.arxiv.org/abs/1401.1436>.

Wood, S. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B*, 73(1):3–36.