Multi-level Approximate Bayesian Computation

Christopher Lester*

20 November 2018

1 Introduction

Well-designed mechanistic models can provide insights into biological networks that are imperceptible to machine learning techniques. For example, where suitable experimental data exists, mechanistic models can often provide strong evidence for the causative relationships that underpin a given biological model [1]. Mechanistic, multi-scale models have been able to assimilate physical features and behaviours that span multiple time- and spatial-scales [2].

To draw reliable conclusions from a mechanistic model, experimental data are often used to inform the planning, implementation and usage of the model. In particular, inference methods are frequently used to choose numerical parameters for a given model, or perhaps to select the most appropriate from a class of possible models. It is thus essential that uncertainties in inferred parameters, and hence, in any conclusions, are accurately quantified.

In this work, we focus on discrete state-space stochastic models that are governed by the Chemical Master Equation (the CME). Whilst originally designed to model biochemical reactions, CME-based models are now frequently used to describe a wide range of biological phenomena mathematically [3, 4, 5]. For CME-based models, Approximate Bayesian Computation (ABC) is becoming an increasingly popular method of inferring model parameters [6, 7, 8].

Parameter inference has traditionally been seen as the preserve of Bayesian statistics. Any existing knowledge of a parameter, θ , is encoded as a *prior* distribution, $\pi(\theta)$; the probability of a parameter, θ , given data D , $\mathcal{P}[\theta \mid D]$ is then known as the *posterior* distribution [9].

*Mathematical Institute, Woodstock Road, Oxford, OX2 6GG, UK. Email: `lesterc AT maths.ox.ac.uk`.

Following Bayes' theorem, the *likelihood*, $\mathbb{P}[D \mid \theta]$, relates the *posterior* distribution to the *prior*:

$$\mathbb{P}[\theta \mid D] \propto \mathbb{P}[D \mid \theta] \times \pi(\theta).$$

For many mechanistic models, the likelihood is intractable, rendering many Bayesian inference methods infeasible [8]. As a likelihood-free method, ABC is unaffected and is often able to infer parameters with a high degree of accuracy. ABC uses stochastic simulation to infer parameters: after sampling parameters from a prior distribution, sample paths or realisations of the model of interest are generated. By studying the sample paths that have been generated, the posterior distribution is estimated.

In this work, we describe and implement an efficient multi-level ABC method for investigating model parameters. In short, we generate sample paths of CME-based models with varying time resolutions. We will first generate low resolution sample paths, by which we mean sample paths with few time-steps (and therefore low-quality approximations of the model dynamics), and we will later generate high resolution sample paths, by which we mean sample paths with more time-steps (but which require more computational resources to generate). The multi-level ABC ('ML-ABC') method that we set out in this manuscript proceeds to:

1. Start by generating low-resolution sample paths to infer model parameters.
2. Choose a subset of sample paths, with a view to improving their time resolution. The subset is chosen to include more of the sample paths likely to influence the posterior distribution. The sample paths not included in the subset, are discarded.
3. Step 2 is then recursively repeated with the remaining sample paths, until all sample paths have a sufficiently high resolution. The posterior distribution is then estimated.

The result is that we are able to discard many (if not most) sample paths after a quickly-generated low-resolution version is generated. The bulk of the computational effort is thus expended on improving the time resolution of those sample paths that most likely to contribute to the posterior distribution.

This work is predicated on, and seeks to unify, three principal previous works:

1. Whilst rejection sampling had previously been used to sample a posterior distribution [10], Toni et al. [8] provide perhaps the most widely-read account of the ABC inference framework, and we build on their framework.

2. Prangle [11] describes a ‘lazy’ ABC approach. In this work, we apply the theory developed by Prangle [11] in a novel way.
3. Giles [12] describes a multi-level method for calculating Monte Carlo estimates. A recent work [13] discusses the means to re-use random inputs to generate sample paths with different time resolutions. We follow this approach.

This manuscript is arranged as follows: in Section 2 we provide background material covering biochemical reaction networks, and the ABC algorithm. In Section 3 our new framework is set out. The framework is implemented and tested in Section 4, and a discussion of our results then follows.

2 Inference for stochastic biochemical networks

This section is in two parts: first, we explain the CME modelling framework in Section 2.1, and in Section 2.2, the Approximate Bayesian Computation methodology is outlined.

2.1 Biochemical reaction networks

We study a biochemical network comprising N species, $S_1, \dots, S_N$, that may interact through M reaction channels, $R_1, \dots, R_M$. In our case, the dynamics of a biochemical reaction network are governed by the Chemical Master Equation (CME) framework [14]. At time t , the population, or copy number, of species S_i is denoted by $X_i(t)$, and the state vector, $\mathbf{X}(t)$, is given by

$$\mathbf{X}(t) := [X_1(t), \dots, X_N(t)]^T. \quad (1)$$

We associate two quantities with each reaction channel R_j . The first is the stoichiometric or state-change vector,

$$\boldsymbol{\nu}_j := [\nu_{1j}, \dots, \nu_{Nj}]^T, \quad (2)$$

where ν_{ij} is the change in the copy number of S_i caused by reaction R_j taking place. The second quantity is the propensity function, $p_j(\mathbf{X}(t))$. For infinitesimally small dt , the rate $p_j(\mathbf{X}(t))$ is defined as follows:

$$p_j(\mathbf{X}(t))dt := \mathbb{P}[R_j \text{ occurs in } [t, t + dt)].$$

We assume that the system is well-stirred, so the reaction activity can be modelled with mass action kinetics. In this case, the propensity function of reaction R_j , p_j , is proportional to the number of possible combinations of reactant particles in the system [14]. The constants of proportionality are known as rate constants and must be inferred from experimental data [5].

For our purposes, it is convenient to use the Random Time Change Representation (RTCR) of the CME framework, which was first described by Kurtz [15]. The RTCR describes the dynamics of a biochemical network by using a set of unit-rate Poisson processes. The number of times reaction R_j (for $j = 1, \dots, M$) takes place (‘fires’) over the time interval $(0, T]$ is given by a Poisson counting process

$$\mathcal{Y}_j \left(0, \int_0^T p_j(\mathbf{X}(t)) dt \right),$$

where $\mathcal{Y}_j$ is a unit-rate Poisson process, and $\mathcal{Y}_j(\alpha, \beta)$ is defined as

$$\mathcal{Y}_j(\alpha, \beta) := \# \text{ of arrivals in } (\alpha, \beta]. \quad (3)$$

Every time reaction R_j occurs, the state vector (see Equation (1)), is updated by adding the appropriate stoichiometric vector (see Equation (2)) to it. Therefore, by considering all possible reactions over the time interval $(0, T]$, we can determine the state vector at time T as

$$\mathbf{X}(T) = \mathbf{X}(0) + \sum_{j=1}^M \mathcal{Y}_j \left(0, \int_0^T p_j(\mathbf{X}(t)) dt \right) \cdot \nu_j. \quad (4)$$

We can think of different sets of Poisson processes, $\{\mathcal{Y}_j \text{ for } j = 1, \dots, M\}$, as providing the randomness for different sample paths of our biochemical reaction network. Each random sample path is thus uniquely associated with a set of Poisson processes (and each reaction channel is associated with a specific Poisson process).

It is possible to directly simulate the sample path described by Equation (4) by using the Modified Next Reaction Method (the MNRM) [16]. The MNRM is set out as Algorithm 1; the sample paths produced by this method are statistically indistinguishable from the sample paths generated with the famed Gillespie Algorithm. As with the Gillespie Algorithm, the MNRM is a serial algorithm: it uses a relatively high level of computational resources to generate each sample path as it ‘fires’ only a single reaction at each step of the algorithm. As alluded to above, a feature of the MNRM is that it associates a specific Poisson process with each reaction channel; we will rely on this feature throughout the remainder of this manuscript.

Sample paths of $\mathbf{X}(t)$ can be generated more efficiently by using the tau-leap assumption.

Algorithm 1 The MNRM. A single sample path is generated by using Poisson processes, $\mathcal{Y}_j$ (for $j = 1, \dots, M$), to ‘fire’ reactions through different reaction channels.

Require: initial conditions, $\mathbf{X}(0)$, and terminal time, T .

```

1: set  $\mathbf{X} \leftarrow \mathbf{X}(0)$ , and set  $t \leftarrow 0$ 
2: for each  $R_j$ , set  $P_j \leftarrow 0$ , generate  $T_j \leftarrow \text{Exp}(1)$   $\triangleright T_j$  is the first inter-arrival time of  $\mathcal{Y}_j$ .
3: loop
4:   for each  $R_j$ , calculate propensity values  $p_j(\mathbf{X})$  and calculate  $\Delta_j$  as
      
$$\Delta_j = \frac{T_j - P_j}{p_j}$$

5:   set  $\Delta \leftarrow \min_j \Delta_j$ , and  $k \leftarrow \text{argmin}_j \Delta_j$ 
6:   if  $t + \Delta > T$  then
7:     break
8:   end if
9:   set  $\mathbf{X}(t + \Delta) \leftarrow \mathbf{X}(t) + \nu_j$ , set  $t \leftarrow t + \Delta$ , and for each  $R_j$ , set  $P_j \leftarrow P_j + p_j \cdot \Delta$ 
10:  generate  $u \sim \text{Exp}(1)$ , then set  $T_k \leftarrow T_k + u$   $\triangleright u$  is the next inter-arrival time of  $\mathcal{Y}_k$ .
11: end loop
```

In this case, the state vector, $\mathbf{X}(t)$, is only updated at fixed times: for example at times $t = \tau, 2 \cdot \tau, \dots$, where τ is an appropriate time-step. The tau-leap assumption means the following substitution can be made in Equation (4):

$$\int_0^T p_j(\mathbf{X}(t)) dt \rightarrow \sum_{k=0}^K p_j(\mathbf{X}(k \cdot \tau)) \cdot \tau. \quad (5)$$

Equation (4) can then be rearranged into an update formula:

$$\mathbf{X}(K \cdot \tau) = \mathbf{X}((K-1) \cdot \tau) + \sum_{j=1}^M \mathcal{Y}_j \left(\sum_{k=0}^{K-1} p_j(\mathbf{X}(k \cdot \tau)) \cdot \tau, \sum_{k=0}^K p_j(\mathbf{X}(k \cdot \tau)) \cdot \tau \right) \cdot \nu_j. \quad (6)$$

Effectively, Equation (6) states that, to advance from time $t = (K-1) \cdot \tau$ to time $t = K \cdot \tau$, we ‘read-off’ the number of times each Poisson process, $\mathcal{Y}_j$ (for $j = 1, \dots, M$), has ‘fired’ during the time-interval $\left(\sum_{k=0}^{K-1} p_j(\mathbf{X}(k \cdot \tau)) \cdot \tau, \sum_{k=0}^K p_j(\mathbf{X}(k \cdot \tau)) \cdot \tau \right]$. As expected, the length of each interval is $p_j(\mathbf{X}((K-1) \cdot \tau)) \cdot \tau$.

As explained above, each sample path is associated with a unique set of Poisson processes, $\{\mathcal{Y}_j$ for $j = 1, \dots, M\}$. The same set of Poisson processes can be used, but with a different value of τ , to change the resolution at which a sample path is constructed.

We illustrate with an example:

Case study. Consider a reaction network comprising a single species, X , together with a

reaction channel R_1 , which we specify as:

We take $X(0) = 10$ and $\theta = 0.3$. In Figure 1, we use a single Poisson process, $\mathcal{Y}$, to generate a single sample path, but at different resolutions. As $\tau \downarrow 0$, the sample path corresponds with a sample path generated by the MNRM.

Figure 1: A single sample path of System (7) is shown, but which has been generated by using different time-steps in Equation (6). We show sample paths with $\tau = 1.0$ (10 steps) and $\tau = 0.2$ (50 steps), together with the MNRM sample path ($\tau \downarrow 0$). The same Poisson process is used throughout.

2.2 Approximate Bayesian Computation

We now detail how ABC infers the parameters of the CME-based stochastic model described in Section 2.1. In Section 1 we explained that the ABC method is an example of a likelihood-free inference method that allows us to directly estimate the posterior, $\mathbb{P}[\theta \mid D]$:

- The prior, $\pi(\cdot)$, encodes any existing knowledge regarding the parameters. An uninformative prior – for example, a uniform distribution – can be used if little is known about the model.
- A parameter¹, θ , is sampled from the prior. Then, based on this parameter choice, a sample path, $\mathbf{X}$, is generated. We write $\mathbf{X} \sim f(\cdot \mid \theta)$ to indicate this. This step can be repeated N times to generate many sample paths.

¹Whilst we refer to ‘parameter’ as singular, indeed, the parameter could be a vector.

- The sample paths, $\{\mathbf{X}_1, \dots, \mathbf{X}_N\}$, are compared with the observed data, which we label $\widehat{\mathbf{X}}$. As the CME model is stochastic, the observed data are very unlikely to co-incide with a generated sample path.

Therefore, summary statistics are specified, and the summary statistics of sample paths are compared with the summary statistics of the raw data. For example, we might take $s(\mathbf{X}) = [\mathbf{X}(t_1), \dots, \mathbf{X}(t_K)]$, and then compare summary statistics with the Frobenius norm for matrices, $\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\|_2$.

- A tolerance, $\varepsilon > 0$, is specified. Those sample paths that match the data – for example, where $\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon$ – are *accepted*, whilst the remainder are *rejected*. The parameters associated with the accepted sample paths are used to empirically generate an approximate posterior distribution, $\mathbb{P}[\theta \mid \|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon]$. Where the summary statistics are *sufficient*, as $\varepsilon \downarrow 0$, the true posterior, $\mathbb{P}[\theta \mid \mathbf{X}]$, can be recovered [6].

This is the rejection ABC algorithm, and it is stated as Algorithm 2. This basic algorithm can be improved through the use of a Markov-Chain Monte Carlo (MCMC) approach [17], or through a Sequential Monte Carlo sampler method [8].

Algorithm 2 Rejection ABC. The loop is repeated to generate sample paths, $\mathbf{X} \sim f(\cdot \mid \cdot)$, until sufficiently many values of θ have been accepted.

Require: data, $\widehat{\mathbf{X}}$, tolerance, ε , summary statistic, $s(\cdot)$, distance metric $\|\cdot\|$, and prior, $\pi(\cdot)$

```

1: loop
2:   sample  $\theta \sim \pi(\cdot)$ 
3:   generate sample path  $\mathbf{X} \sim f(\cdot \mid \theta)$ 
4:   if  $\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon$  then
5:     accept  $\theta$ 
6:   end if
7: end loop

```

Before proceeding to Section 3, we point out that Algorithm 2 can be viewed as a weighted sampling algorithm. In particular, after each parameter, θ , is sampled, it can be given a weight, ω , according to

$$\omega = \begin{cases} 1 & \text{if } \mathbf{X} \sim f(\cdot \mid \theta) \text{ is accepted.} \\ 0 & \text{if } \mathbf{X} \sim f(\cdot \mid \theta) \text{ is rejected.} \end{cases}$$

This opens up at least two possibilities: firstly, one can sample θ as $\theta \sim g(\cdot)$, where $g(\cdot)$ has been carefully chosen to explore the parameter space most likely to result in a parameter being accepted. The weight of θ is then scaled by a factor of $\pi(\theta)/g(\theta)$. The second possibility involves terminating the simulation of $\mathbf{X}$ prematurely: we explore this idea in Section 3. The

use of non-integer weights means that an *effective sample size* needs to be calculated. Kish’s *effective sample size* is given by

$$\text{ESS} = \frac{\left[\sum_{r=1}^N \omega_r \right]^2}{\sum_{r=1}^N \omega_r^2}. \quad (8)$$

Having discussed the CME modelling framework, and described the ABC method, we are now in a position to set out our new multi-level ABC (ML-ABC) method.

3 Multi-level Approximate Bayesian Computation

This section is in three parts: first, in Section 3.1, we discuss and adapt Prangle [11]’s ‘lazy’ ABC method; then, in Section 3.2 we refer to an earlier work [13] to explain how to generate the requisite sample paths. Finally, Section 3.3 spells out our method.

3.1 Early-rejection ABC

Prangle [11] introduces what we will call an ‘early-rejection’ ABC method, which we now summarise. Let $\mathbf{X}$ refer to the sample paths of the model of interest, and recall that $\mathbf{X}$ is sampled as $\mathbf{X} \sim f(\cdot \mid \theta)$. Let $\mathbf{Y}$ refer to a *partial* sample path of the model that meets the following requirements:

- $\mathbf{Y}$ can be sampled as $\mathbf{Y} \sim g(\cdot \mid \theta)$ cheaply.
- Given the partial sample, $\mathbf{Y}$, a full sample can be deduced as $\mathbf{X} \sim f(\cdot \mid \theta, \mathbf{Y})$.

The early-rejection procedure is as follows: a sample path $\mathbf{Y}$ is generated. Then, a decision function $\alpha(\theta, \mathbf{Y})$ provides a *continuation probability*: with probability $\alpha(\theta, \mathbf{Y})$, the partial sample $\mathbf{Y}$ is upgraded to a full sample, $\mathbf{X}$, by simulating $\mathbf{X} \sim f(\cdot \mid \theta, \mathbf{Y})$. The weight of the full sample is then re-scaled by a factor of $1/\alpha(\theta, \mathbf{Y})$, and the ABC algorithm proceeds. Alternatively, if the partial sample, $\mathbf{Y}$, is not upgraded, it is discarded.

Prangle [11] envisage a range of scenarios where early-rejection ABC can be implemented. For example, suppose that the sample path $\mathbf{X}$ is generated over a time-interval $[0, t)$. Then the partial sample path, $\mathbf{Y}$, could be generated over a shorter time-interval $[0, s)$ (with $s \ll t$, or with s a random stopping time). Based on the shorter sample path, the probability of the

complete sample path meeting the ABC acceptance criterion, $\mathbb{P} \left[\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon \mid \theta, \mathbf{Y} \right]$, can be estimated. Unlikely sample paths can thus be discarded after using only a small amount of computational resources, but at the cost an increased variance. Asymptotically, the optimal efficiency is achieved where the ratio of the expected effective sample size (see Equation (8)) to the expected CPU time per sample, $\mathbb{E}[T]$, is maximised:

$$\frac{\mathbb{E}[\omega]^2}{\mathbb{E}[\omega^2] \cdot \mathbb{E}[T]}. \quad (9)$$

Let $\phi(\theta, \mathbf{Y})$ be a summary statistic² that describes θ and the partial sample path $\mathbf{Y}$. Then, let T_1 be the CPU time required to simulate $\mathbf{Y} \sim g(\cdot \mid \theta)$, and T_2 the CPU time required for simulating $\mathbf{X} \sim g(\cdot \mid \theta, \mathbf{Y})$. It can be shown that taking

$$\alpha(\phi) \propto \left[\frac{\mathbb{P} \left[\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon \mid \phi \right]}{T_2(\phi)} \right]^{1/2} \quad (10)$$

is optimal [11]. The early-rejection ABC method is stated as Algorithm 3.

Algorithm 3 Early-rejection ABC. The loop is repeated to return pairs of (θ, ω) .

Require: data, $\widehat{\mathbf{X}}$, tolerance, ε , summary statistic, $s(\cdot)$, distance metric $\|\cdot\|$, and prior, $\pi(\cdot)$
Require: summary $\phi(\theta, \mathbf{Y})$ and continuation probability, $\alpha(\phi)$

```

1: loop
2:   sample  $\theta \sim \pi(\cdot)$ 
3:   generate sample path  $\mathbf{Y} \sim g(\cdot \mid \theta)$ 
4:   with probability  $\max(1, \alpha(\phi(\theta, \mathbf{Y})))$  continue
5:   generate sample path  $\mathbf{X} \sim f(\cdot \mid \theta, \mathbf{Y})$ 
6:   calculate  $\omega$ 

```

$$\omega = \begin{cases} \frac{1}{\alpha} \mathbb{I} \left\{ \|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon \right\} & \text{with probability } \max(1, \alpha(\phi(\theta, \mathbf{Y}))), \\ 0 & \text{otherwise.} \end{cases}$$

```

7:   return  $(\theta, \omega)$ 
8: end loop

```

3.2 Generating sampling paths with different resolutions

In Section 1 we explained that the stochastic process, defined by Equation (4),

$$\mathbf{X}(T) = \mathbf{X}(0) + \sum_{j=1}^M \mathcal{Y}_j \left(0, \int_0^T p_j(\mathbf{X}(t)) dt \right) \cdot \nu_j,$$

²In Section 4, we will simply take $\phi(\theta, \mathbf{Y}) = \|s(\mathbf{Y}) - s(\widehat{\mathbf{X}})\|$.

can be simulated more efficiently if the tau-leap method is used. The tau-leap can be seen as making the substitution $\int_0^T p_j(\mathbf{X}(t))dt \rightarrow \sum_{k=0}^K p_j(\mathbf{X}(k \cdot \tau)) \cdot \tau$. Thus, let the tau-leap process, $\mathbf{Y}$, be defined by Equation (6), which is,

$$\mathbf{Y}(K \cdot \tau) = \mathbf{Y}((K-1) \cdot \tau) + \sum_{j=1}^M \mathcal{Y}_j \left(\sum_{k=0}^{K-1} p_j(\mathbf{Y}(k \cdot \tau)) \cdot \tau, \sum_{k=0}^K p_j(\mathbf{Y}(k \cdot \tau)) \cdot \tau \right) \cdot \nu_j.$$

We let $\mathbf{Y}$ take the role of the partial sample path in Algorithm 3. More generally, we let $\{\mathbf{Y}_\tau : \tau \in \{\tau_1 > \tau_2 > \dots\}\}$ be a family of partial sample paths. From Equation (6) it is clear that simulating a sample path $\mathbf{Y}_\tau$ is a matter of generating the Poisson processes, $\mathcal{Y}_j$ (for $j = 1, \dots, M$). A Poisson process can be simulated in two ways:

- The Poisson *random number*, $\mathcal{P}(\Delta)$, indicates the number of arrivals in the Poisson process over an interval length Δ (but does not indicate the specific arrival times).
- The exponential *random number*, $\text{Exp}(1)$, indicates the waiting time between successive arrivals in the Poisson process.

The tau-leap method generates sample paths quickly by using Poisson random numbers to ‘fire’ multiple reactions at once; there is no time-saving (over the MNRM) if exponential random variables are simulated.

As such, we start by generating $\mathbf{Y}_\tau$ with time-step $\tau \leftarrow \tau_1$. This means that, to advance from time $t = 0$ to time $t = K \cdot \tau$ in Equation (6), we need to compute the number of arrivals in the time intervals $(0, p_j(\mathbf{Y}(0)) \cdot \tau]$, $(p_j(\mathbf{Y}(0)) \cdot \tau, p_j(\mathbf{Y}(\tau)) \cdot \tau]$, $(p_j(\mathbf{Y}(\tau)) \cdot \tau, p_j(\mathbf{Y}(2\tau)) \cdot \tau]$, $\dots$, for the appropriate $j = 1, \dots, M$. Poisson random numbers can be used to generate these quantities.

The resolution of the sample path, $\mathbf{Y}_\tau$, can be increased if we change the time-step to $\tau' \leftarrow \tau_2$. In this case, we need to compute the number of arrivals in the time intervals $(0, p_j(\mathbf{Y}(0)) \cdot \tau']$, $(p_j(\mathbf{Y}(0)) \cdot \tau', p_j(\mathbf{Y}(\tau')) \cdot \tau']$, $(p_j(\mathbf{Y}(\tau')) \cdot \tau', p_j(\mathbf{Y}(2\tau')) \cdot \tau']$, $\dots$, for the appropriate $j = 1, \dots, M$. These do not correspond with the intervals previously generated, and so the existing Poisson process needs to be *interpolated*. The following rule is used to interpolate such processes:

Rule 1. If there are K arrivals in a Poisson process over an interval (α, γ) then, for β between α and γ , the number of arrivals over the interval (α, β) is binomially distributed³ as

$$\mathcal{B}\left(K, \frac{\beta - \alpha}{\gamma - \alpha}\right). \quad (11)$$

³This follows as the K arrivals are uniformly distributed over (α, γ) .

A return to our earlier case study illustrates the interpolation of a Poisson process:

Case study. A sample path of System (7) was generated with multiple resolutions, and the results detailed in Figure 1. In Figure 2 we show how a single Poisson process is repeatedly interpolated to generate the sample path with different resolutions.

Figure 2: A unit-rate Poisson process is shown. The cumulative number of arrivals at different time-points is indicated. Initially, the Poisson process is not fully resolved as this is unnecessary. Where the number of arrivals needs to be calculated at different time-points, the process is interpolated per Rule 1.

3.3 The ML-ABC algorithm

We are in a position to set out the ML-ABC method. First, we settle on a choice of approximate sample resolutions, and therefore the set $\{\mathbf{Y}_1, \dots, \mathbf{Y}_L : \text{where } \mathbf{Y}_\ell \text{ has time-step } \tau_\ell\}$. Then,

- Sample $\theta \sim \pi(\cdot)$ and fix the Poisson processes, $\mathcal{Y}_j$ for $j = 1, \dots, M$.
- Generate $\mathbf{Y}_1$ with time-step τ_1 .

With probability $\alpha_1(\phi_1(\mathbf{Y}_1))$ proceed to (using the aforementioned Poisson processes):

- Generate $\mathbf{Y}_2$ with time-step τ_2 .

With conditional probability $\alpha_2(\phi_2(\mathbf{Y}_1, \mathbf{Y}_2))$ proceed to (using the Poisson processes):

- Generate $\mathbf{Y}_3$ with time-step τ_3 .

Proceed recursively, until, with conditional probability $\alpha_L(\phi_L(\mathbf{Y}_1, \dots, \mathbf{Y}_L))$ we proceed to (using the Poisson processes):

- Generate $\mathbf{X}$ with the MNRM. $\mathbf{X}$ is the *exact* or complete sample path.

A pair (θ, ω) is then returned; the weight, ω , is given by

$$\omega = \begin{cases} \left[\prod \alpha_i \right]^{-1} \mathbb{I} \left\{ \|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon \right\} & \text{if } \mathbf{X} \text{ is generated,} \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

In practice, we suggest the following procedure for choosing α_ℓ . A small number of survey simulations are generated, with $\alpha_1 = \dots = \alpha_L = 1$. Then, we take ϕ_ℓ to be the current ‘error’, given by:

$$\phi_\ell(\mathbf{Y}_1, \dots, \mathbf{Y}_\ell) = \|s(\mathbf{Y}_\ell) - s(\widehat{\mathbf{X}})\|. \quad (13)$$

For clarity, we label the quantity (13) as $\text{err}(\mathbf{Y}_\ell)$. Then, following Equation (10), we take

$$\alpha_1(\mathbf{Y}_1) \propto \mathbb{P} \left[\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon \mid \|s(\mathbf{Y}_1) - s(\widehat{\mathbf{X}})\| \right]^{1/2}. \quad (14)$$

Note that we have suppressed the usage of ϕ , as it no longer is necessary. Carrying on this way, we take

$$\alpha_\ell(\mathbf{Y}_\ell) \propto \mathbb{P} \left[\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon \mid \|s(\mathbf{Y}_\ell) - s(\widehat{\mathbf{X}})\| \right]^{1/2}. \quad (15)$$

Note that the probabilities are to be weighted by the probabilities of level ℓ being reached.

The probability $\mathbb{P} \left[\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| < \varepsilon \mid \|s(\mathbf{Y}_\ell) - s(\widehat{\mathbf{X}})\| \right]$ indicates the chance that, based on the current error $\text{err}(\ell)$, that the associated full (exact) sample path, $\mathbf{X}$ would be within an ε -distance of the test data, $\widehat{\mathbf{X}}$, and would consequently be *accepted*. The probabilities, $\mathbb{P}[\cdot]$ can be estimated in many ways – for example, it could be sufficient to fit a decision function $p_\ell(\xi) = \exp \left[\sum_i \lambda_i (\xi - \xi_0)^i \right]$, together with a suitable regularisation, to the survey simulation data.

4 Case studies

In this section, we present two case studies; a discussion of our results then follows in Section 4.3. The first case study is concerned with a birth process (Section 4.1), and the second case

study models the spread of a disease through an S-I-S compartment model (Section 4.2). Figure 3 presents the *in-silico* data that we have generated to test the ML-ABC algorithm on. For each of Case studies 1 and 2, a single time-series is shown; all our sample paths were generated on an Intel Core i5 CPU rated at 2.5 GHz.

Figure 3: On the left, an *in-silico* data path of a birth process is shown. The data path on the right is of an S-I-S model. The data were randomly generated, with a random seed set to the author’s birth-date. The proposed summary statistics are shown in red.

4.1 Case study 1: a birth process

The first case study is concerned with a birth process. As in Section 2, this process comprises a single species, X , together with a reaction channel R_1 :

$$R_1 : X \xrightarrow{\theta} X + X,$$

with $X(0) = 10$. We seek to infer the parameter θ . The data, $\widehat{\mathbf{X}}$, are shown in Figure 3 (note that $X(0)$ is chosen in accordance with Figure 3). In particular, for our purposes it is sufficient to let the summary statistic, s , be given by the final population, $s(\mathbf{X}) = \mathbf{X}(10)$.

We use a uniform prior distribution, taking $\pi(\cdot) \sim \mathcal{U}[0.01, 1.00]$. We then generate $N = 10^5$ sample paths using the MNRM (see Algorithm 1). The error associated with each sample path is given by

$$\text{err}(\mathbf{X}) := \|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\| = |\mathbf{X}(10) - \widehat{\mathbf{X}}(10)|. \quad (16)$$

Taking a tolerance of $\varepsilon = 35$, an empirical posterior distribution is plotted, and the results detailed in Figure 4. The maximum of this empirical posterior distribution is $\theta = 0.304$, which compares very favourably with the choice of $\widehat{\theta} = 0.3$ used to generate the *in-silico* data. In

addition, Figure 4 indicates the empirical posterior distributions associated with the use of tau-leap sample paths, with τ taken as $\tau = 1.0$ and $\tau = 0.2$. This figure is consistent with the tau-leap sample paths introducing an additional bias.

Figure 4: Empirical posterior distributions for the parameter θ of System (7) are shown. Different simulation methods are used; a parameter of $\varepsilon < 35$ is used to decide whether to accept a sample. Approximately 3.1% of MNRM sample paths are accepted; where tau-leaping is used, 4.2% of sample paths are accepted where $\tau = 1.0$, and where $\tau = 0.2$, the rate is 3.3%.

The ML-ABC method is now implemented. To illustrate the method, we take $\tau_1 = 1.0$ and $\tau_2 = 0.2$. As outlined in Section 3, we sample $\theta \sim \pi(\cdot)$, and we then generate a sample path with time-step $\tau_1 = 1.0$, $\mathbf{Y}_1$. If $\mathbf{Y}_1$ is such that it is likely to contribute to the empirical posterior distribution, then we re-generate $\mathbf{Y}_1$, but with a time-step of $\tau_2 = 0.2$. As the same Poisson processes are used, $\mathbf{Y}_2$ is a higher-resolution version of $\mathbf{Y}_1$. Based on $\mathbf{Y}_2$, we decide whether the exact (or complete) sample path, $\mathbf{X}$ should be generated. At each stage, we re-weight our sample to avoid introducing an additional bias.

In order to decide which sample paths to refine, we first generate a number of survey simulations to calibrate the model. We take $N' = 1000$, and therefore generate 1000 triples $\{\mathbf{Y}_1, \mathbf{Y}_2, \mathbf{X}\}$. Any suitable machine-learning classifier can then be used to probabilistically associate the error, $\text{err}(\mathbf{Y}_1)$ (see Equation (16)) with the final outcome, $\mathbb{I}\{\text{err}(\mathbf{X}) < \varepsilon\}$, that is the probabilistic relationship

$$\text{err}(\mathbf{Y}_1) = |s(\mathbf{Y}_1) - s(\widehat{\mathbf{X}})| \quad \text{to the class} \quad \mathbb{I}\{|s(\mathbf{X}) - s(\widehat{\mathbf{X}})| < \varepsilon\}. \quad (17)$$

To this end, we use a simple, Gaussian-form probability function, which is fitted to the test data. The function, ψ_1 , is indicated on the left side of Figure 5. Following from Equation (10),

it is sufficient⁴ for our continuation probability, α_1 , to be taken as

$$\alpha_1 := \lambda_1 [\psi_1(\text{err}(\mathbf{Y}_1))]^{1/2}. \quad (18)$$

Having specified α_1 , our attention turns to α_2 . Whilst Section 3 indicated more general formulations are possible, we seek to relate

$$\text{err}(\mathbf{Y}_2) = |s(\mathbf{Y}_2) - s(\widehat{\mathbf{X}})| \quad \text{to the class} \quad \mathbb{I}\{|s(\mathbf{X}) - s(\widehat{\mathbf{X}})| < \varepsilon\}. \quad (19)$$

We use another Gaussian-form probability function, which is fitted to the test data. This decision function is labelled ψ_2 . The data are weighted in proportion to the previous continuation probability, α_1 . This means that samples that are unlikely to continue to the second stage have a relatively small effect on ψ_2 . The right side of Figure 5 indicates the second decision function. As before, we now take

$$\alpha_2 := \lambda_2 [\psi_2(\text{err}(\mathbf{Y}_2))]^{1/2}. \quad (20)$$

The values of λ_1 and λ_2 are now optimised on the set of $N' = 1000$ sample points.

Figure 5: On the left, we plot an indicator function $\mathbb{I}\{|s(\mathbf{X}) - s(\widehat{\mathbf{X}})| < \varepsilon\}$ against the error $\text{err}(\mathbf{Y}_1)$ in blue. A moving average is shown in green, and a Gaussian-form curve is fitted in orange. The curve indicates the probability that, given $\text{err}(\mathbf{Y}_1)$, that if the complete sample $\mathbf{X}$ is generated from $\mathbf{Y}_1$, that the complete sample will be within an ε -distance of the test data. On the right, an equivalent graph for $\text{err}(\mathbf{Y}_2)$ is shown.

It is now possible to run the full algorithm, with $N = 10^5$ initial samples of $\mathbf{Y}_1$. Of the $N = 10^5$ samples, 10,435 are selected to be refined into $\mathbf{Y}_2$. Of those 10,435 samples, only 2,712 samples are further refined to the exact sample path, $\mathbf{X}$. Of the 2,712 samples, 1,273 are accepted (as $\text{err}(\mathbf{X}) < \varepsilon$), and contribute to the empirical posterior distribution. The ESS (see Equation

⁴As highlighted in Equation (10), a continuation probability can also take into account the expected CPU-time of refining a sample path. This is not necessary for this particular case study.

(8)) takes into account sample weights, meaning that the ESS is 1,079.8. The effective samples per CPU-second is 487.2 units.

If the regular, ABC-rejection method (see Algorithm 2) is used, with $N = 10^5$ initial samples and $\varepsilon = 35$, then 3,107 samples are accepted for inclusion in the posterior distribution. This is a rate of 16.9 effective samples per CPU-second. The result is that the ML-ABC method is approximately 28.8 times more efficient than rejection ABC.

4.2 Case study 2: an S-I-S model

In this second case study, we consider an S-I-S infection model [18]. The model comprises two species, S (the ‘susceptible’ individuals) and I (the ‘infected’ individuals). Initially, $S = 950$ and $I = 50$ (these quantities can be deduced from Figure 3). There are two reaction channels:

We seek the values of parameters θ_1 and θ_2 . The data, $\widehat{\mathbf{X}}$, are shown in Figure 3. In particular, for our purposes it is sufficient to let the summary statistic, s , be given by the vector, $s(\mathbf{X}) = [S(1), \dots, S(4), I(1), \dots, I(4)]$.

We use a uniform prior distribution, taking $\pi(\cdot) \sim \mathcal{U}[0.0001, 0.02] \times \mathcal{U}[0.0001, 5.0]$. We generate $N = 10^5$ sample paths using the MNRM (see Algorithm 1). The error associated with each sample path is given by the ℓ^2 -distance

$$\text{err}(\mathbf{X}) := \|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\|_2. \quad (22)$$

Taking a tolerance of $\varepsilon = 250$, an empirical posterior distribution is plotted, and the results detailed in Figure 6. The maximums of the point-wise empirical posterior distributions are $(\theta_1, \theta_2) = (0.0035, 1.207)$, which compares with the choice of $(\widehat{\theta}_1, \widehat{\theta}_2) = (0.003, 1.000)$ used to generate the *in-silico* data.

The ML-ABC method is now implemented. Given the larger domain spanned by the prior, $\pi(\cdot)$, we will use an adaptive tau-leap method to generate $\mathbf{Y}_1$. This means that, instead of specifying τ directly, we specify a parameter, ξ , that parameterises a tau-selection algorithm. In this case, the precise value of τ is chosen according to ξ and the state vector (see Equation (1)). We use the popular tau-selection method of Cao et al. [19]. We use only one approximation level, so to generate $\mathbf{Y}_1$ we fix $\xi = 0.2$. If $\mathbf{Y}_1$ is such that it is likely to contribute to the empirical

Figure 6: Empirical posterior distributions for the parameters θ_1 and θ_2 of System (21) are shown. A parameter of $\varepsilon < 250$ is used to decide whether to accept a sample. Approximately 1.7% of the sample paths have been accepted.

posterior distribution, then we generate the complete sample path, $\mathbf{X}$.

As before, we must first generate a small number of survey simulations to calibrate the model. We take $N' = 1000$, and therefore generate 1000 pairs $\{\mathbf{Y}_1, \mathbf{X}\}$. As before, any suitable machine-learning classifier can then be used to associate the error, $\text{err}(\mathbf{Y}_1)$ (see Equation (16)) with the final outcome, $\mathbb{I}\{\text{err}(\mathbf{X}) < \varepsilon\}$. In this case, we use logistic regression to estimate the probability that, given $\text{err}(\mathbf{Y}_1)$, that the full path, $\mathbf{X}$ will be within a ε -distance of the test data, $\widehat{\mathbf{X}}$. Our logistic regression is shown in Figure 7.

Figure 7: We plot an indicator function $\mathbb{I}\{|s(\mathbf{X}) - s(\widehat{\mathbf{X}})| < \varepsilon\}$ against the error $\text{err}(\mathbf{Y}_1)$ in blue. A moving average is shown in green, and a logistic regression curve is fitted in orange.

Following from Equation (10), it is sufficient for our continuation probability, α , to be taken as

$$\alpha = \lambda[\psi(\text{err}(\mathbf{Y}_1))]^{1/2}. \quad (23)$$

The value of λ is optimised on the set of $N' = 1000$ sample points.

It is now possible to run the full algorithm, with $N = 10^5$ initial samples of $\mathbf{Y}_1$. Of the $N = 10^5$ samples, 3,695 are selected to be refined into exact samples, $\mathbf{X}$. Of the 3,695 samples, 1,426 are accepted, and contribute to the empirical posterior distribution. The ESS (see Equation (8)) takes into account sample weights; the ESS is recorded as 1,176.3. The effective samples per CPU-second is 159.0 units.

If the regular, ABC-rejection method (see Algorithm 2) is used, with $N = 10^5$ initial samples and $\varepsilon = 250$, then 1,677 samples are accepted for inclusion in the posterior distribution. This is a rate of 17.0 effective samples per CPU-second. The result is that the ML-ABC method is approximately 9.3 times more efficient than rejection ABC.

4.3 Discussion

Stochastic models of complicated real-world phenomena often rely on carefully chosen parameters. As such, the design and development of efficient inference algorithms plays a pivotal role in underpinning the work of the modelling community. In this work, we presented the ML-ABC method, and demonstrated its efficacy at inferring model parameters. In particular, we demonstrated a multi-resolution inference method, where a machine-learning-led approach is able to select sample paths for further refinement, without introducing an additional bias.

Decision functions. We make only one further point. In Section 4.1, we presented a case study of a birth process. In Figure 8, we consider, $\text{err}(\mathbf{Y}_1)$ and $[\text{err}(\mathbf{Y}_2) - \text{err}(\mathbf{Y}_1)]$. The first quantity indicates the distance between the summary statistics of $\mathbf{Y}_1$ and the data, $\widehat{\mathbf{X}}$, whilst the second quantity indicates the *change* in the error, as τ is reduced from $\tau = 1.0$ to $\tau = 0.2$. We show how the aforementioned quantities are related to (a) the parameter that has been drawn from the prior, θ ; and (b) the final error, that is $\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\|$. Under certain circumstances, it might be favourable to generate *both* $\mathbf{Y}_1$ and $\mathbf{Y}_2$, and then, make a single decision as to whether one should continue to generate the complete sample, $\mathbf{X}$.

In our view, the future of Approximate Bayesian Computation is inextricably tied to machine learning approaches that probabilistically select sample paths and parameters for further investigation.

Figure 8: On the x -axis, we indicate $\text{err}(\mathbf{Y}_1)$, and on the y -axis, we show $[\text{err}(\mathbf{Y}_2) - \text{err}(\mathbf{Y}_1)]$. The left-side graph shows a data point for each triple, $\{\mathbf{Y}_1, \mathbf{Y}_2, \mathbf{X}\}$, with the colour of the data-point corresponding to the parameter, θ , sampled from the prior. The right-side shows the same data points, but with the colour corresponding to the final error, $\|s(\mathbf{X}) - s(\widehat{\mathbf{X}})\|$. The colour-bars have been truncated; approximately 43% of data points are within the axes limits.

Acknowledgements. The author would like to thank Ruth E. Baker for very helpful discussions.

References

- [1] Baker, R. E., Peña, J.-M., Jayamohan, J., and Jérusalem, A. Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biology Letters*, **14**(5):20170660, 2018.
- [2] Weinan, E. *Principles of Multiscale Modeling*. Cambridge University Press, 2011.
- [3] Allen, L. J. *An Introduction to Stochastic Processes with Applications to Biology*. CRC Press, 2010.
- [4] Van Kampen, N. G. *Stochastic Processes in Physics and Chemistry*, volume 1. Elsevier, 1992.
- [5] Wilkinson, D. J. *Stochastic Modelling for Systems Biology*. CRC Press, 2006.
- [6] Csilléry, K., Blum, M. G., Gaggiotti, O. E., and François, O. Approximate Bayesian computation (ABC) in practice. *Trends in Ecology & Evolution*, **25**(7):410–418, 2010.
- [7] Beaumont, M. A. Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution, and Systematics*, **41**:379–406, 2010.

- [8] Toni, T., Welch, D., Strelkowa, N., Ipsen, A., and Stumpf, M. P. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, **6**(31):187–202, 2009.
- [9] Garthwaite, P. H., Jolliffe, I. T., Jolliffe, I., and Jones, B. *Statistical inference*. Oxford University Press, 2002.
- [10] Tavaré, S., Balding, D. J., Griffiths, R. C., and Donnelly, P. Inferring coalescence times from DNA sequence data. *Genetics*, **145**(2):505–518, 1997.
- [11] Prangle, D. Lazy ABC. *Statistics and Computing*, **26**(1-2):171–185, 2016.
- [12] Giles, M. B. Multilevel Monte Carlo path simulation. *Operations Research*, **56**(3):607–617, 2008.
- [13] Lester, C., Yates, C. A., and Baker, R. E. Robustly simulating biochemical reaction kinetics using multi-level Monte Carlo approaches. *Journal of Computational Physics*, **375**:1401–1423, 2018.
- [14] Gillespie, D. T., Hellander, A., and Petzold, L. R. Perspective: Stochastic algorithms for chemical kinetics. *Journal of Chemical Physics*, **138**(17), 2013.
- [15] Kurtz, T. G. Representations of Markov processes as multiparameter time changes. *Annals of Probability*, pages 682–715, 1980.
- [16] Anderson, D. F. A modified next reaction method for simulating chemical systems with time-dependent propensities and delays. *Journal of Chemical Physics*, **127**(21):214107, 2007.
- [17] Marjoram, P., Molitor, J., Plagnol, V., and Tavaré, S. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, **100**(26):15324–15328, 2003.
- [18] Diekmann, O. and Heesterbeek, J. A. P. *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*, volume 5. John Wiley & Sons, 2000.
- [19] Cao, Y., Gillespie, D. T., and Petzold, L. R. Efficient step size selection for the tau-leaping simulation method. *Journal of Chemical Physics*, **124**(4):044109, 2006.