

# Squeeze-and-Breathe Evolutionary Monte Carlo Optimisation with Local Search Acceleration and its application to parameter fitting

Mariano Beguerisse-Díaz<sup>1,2,\*</sup>, Baojun Wang<sup>1</sup>, Radhika Desikan<sup>2</sup>, Mauricio Barahona<sup>1,\*</sup>

<sup>1</sup>Department of Mathematics and <sup>2</sup>Department of Life Sciences,  
Imperial College London, London SW7 2AZ, United Kingdom

December 16, 2011

## Abstract

Estimating parameters from data is a key stage of the modelling process, particularly in biological systems where many parameters need to be estimated from sparse and noisy data sets. Over the years, a variety of heuristics have been proposed to solve this complex optimisation problem, with good results in some cases yet with limitations in the biological setting.

In this work, we develop an algorithm for model parameter fitting that combines ideas from evolutionary algorithms, sequential Monte Carlo and direct search optimisation. Our method performs well even when the order of magnitude and/or the range of the parameters is unknown. The method refines iteratively a sequence of parameter distributions through local optimisation combined with partial resampling from a historical prior defined over the support of all previous iterations. We exemplify our method with biological models using both simulated and real experimental data and estimate the parameters efficiently even in the absence of *a priori* knowledge about the parameters.

## 1 Introduction

The increasing drive towards quantitative technologies in Biology has brought with it a renewed interest in the modeling of biological systems. Models of biological systems and other complex phenomena are generally nonlinear with uncertain parameters, many of which are often unknown and/or unmeasurable [1, 2]. Crucially, the values of the parameters dictate not only the quantitative but also the qualitative behaviour of such models [3, 4]. A fundamental task in quantitative and systems biology is to use experimental data to infer parameter values that minimise the discrepancy between the behaviour of the model and experimental observations. The parameters thus obtained can then be cross-validated against unused data before employing the fitted model as a predictive tool [2]. Ideally, this process could help close the modelling-experiment loop by: suggesting specific experimental measurements; identifying relevant parameters to be measured; or discriminating between alternative models [5, 6, 7].

The problem of parameter estimation and data fitting is classically posed as the minimisation of a cost function (i.e., the error) [8]. In the case of overdetermined linear systems with quadratic error functions, this problem leads to least-square solutions, convex optimisations that can be solved efficiently and globally based on the singular value decomposition of the covariance matrix of the data [9]. However, data fitting in nonlinear systems with small amounts of data remains difficult, as it usually leads to non-convex optimisations with many local minima [10].

A classic case in biological modeling is the description of the time evolution of a system through ordinary differential equations (ODEs), usually based on mechanistic functional forms. Examples include models of biochemical reactions, infectious spread and neuronal dynamics [1, 11]. Typically, optimal parameters of the nonlinear ODEs must be inferred from experimental time courses but the associated optimisation is far from straightforward. Standard optimisation techniques that require an explicit cost

---

\*m.beguerisse-diaz08, m.barahona@imperial.ac.uk

function are unsuitable for this problem due to the difficulty to obtain full analytical solutions for nonlinear ODEs [4, 12, 13]. Spline-based methods, which approximate the solution through an implicit integration of the differential equation [10], require linearity in the parameters and are therefore not applicable to models with nonlinear parameter dependencies, e.g. Michaelis-Menten and Hill kinetics.

Implicit techniques, such as direct search methods [14], Simulated Annealing [15], Evolutionary Algorithms [16, 17] or Sequential Monte Carlo [18], do not require an explicit cost function. However, if as is usually the case, the cost function is a complicated (hyper)surface in parameter space with many local minima, gradient and direct search methods tend to get trapped in local minima due to their use of local information. Although still a local method, Simulated Annealing alleviates some of the problems related to local minima through the use of stochasticity. However, this comes at the cost of high computational overhead and slow convergence and, yet, with no guarantee of finding the global minimum.

Instead of an optimisation based on local criteria, Evolutionary Algorithms (EA) produce an ensemble of possible answers and evolve them globally through random mutation and cross-over followed by ranking and culling of the worst solutions [16, 17, 19]. This heuristic has been shown to provide an efficient protocol for parameter fitting in the life sciences [20, 21]. However, EA methods can be inefficient when the feasible region in parameter space is too large, a case typical of models with large uncertainty in the parameters.

Probabilistic methods, such as Sequential Monte-Carlo (SMC) [18], propose a different conceptual framework. Rather than finding a *unique* optimal parameter set, SMC maps a prior probability distribution of the parameters onto a posterior constructed from samples with low errors until reaching a converged posterior. Recently, SMC has been combined with Approximate Bayesian Computation (ABC) and applied to data fitting and model selection [22]. However, methods such as ABC-SMC are not only computationally expensive but also require that the starting prior include the *true* value of the parameters. This requirement dents its applicability to many biological models, in which not even the order of magnitude of the parameters is known. In that case, the support of the starting priors must be made overly large (leading to extremely slow convergence) in order to avoid the risk of excluding the true parameter value from the search space.

In this work, we present a novel optimisation algorithm for data fitting that takes inspiration from EA, SMC and direct search optimisation. Our method iterates and refines samples from a probability distribution of the parameters in a ‘squeeze-and-breathe’ sequence. At each iteration the probability distribution is ‘squeezed’ by the consecutive application of local optimisation followed by ranking and culling of the local optima. The parameter distribution is then allowed to ‘breathe’ through a random update from a historical prior that includes the union of all past supports of the solutions (Fig. 1). This iteration proceeds until convergence of the distribution of solutions and their average error. A key, distinctive feature of our algorithm is the accelerated step-to-step convergence through a combination of local optimisation and of culling of local solutions. Importantly, the method can also find parameters that lie outside of the range of the initial prior, and can deal with parameter values that extend across several orders of magnitude. We now provide definitions and a full description of our algorithm and showcase its applicability to different biological models of interest.

## 2 Algorithm

### 2.1 Formulation of the problem

Let  $\mathbf{X}(t) = [x_1(t), \dots, x_d(t)]$  denote the state of a system with  $d$  variables at time  $t$ . The time evolution of the state is described by a system of (possibly nonlinear) ODEs:

$$\dot{\mathbf{X}} = f(\mathbf{X}, t; \boldsymbol{\theta}). \quad (1)$$

Here  $\boldsymbol{\theta} = [\theta_1, \dots, \theta_N]$  is the vector of  $N$  parameters of our model.

The experimental data set is formed by  $M$  observations of some of the variables of the system:

$$\mathcal{D} = \left\{ \tilde{\mathbf{X}}(t_i) \mid i = 1, \dots, M \right\}, \quad (2)$$

where  $\tilde{\mathbf{X}}(t_i)$  corresponds to the real value of the system plus observational error. Ideally,  $M > 2N + 1$  since  $2N + 1$  experiments are enough for unequivocal identification of an ODE model with  $N$  parameters when no measurement error is present [23].

The *cost function* (i.e., the error) to be minimised is:

$$E_{\mathcal{D}}(\boldsymbol{\theta}) = \sum_{i=1}^M \left\| \mathbf{X}(t_i; \boldsymbol{\theta}) - \tilde{\mathbf{X}}(t_i) \right\|, \quad (3)$$

where  $\|\cdot\|$  is a relevant vector norm. A standard choice is the Euclidean norm (or 2-norm) which corresponds to the sum of squared errors:

$$E_{\mathcal{D}}^{(2)}(\boldsymbol{\theta}) = \sum_{i=1}^M \sum_{j=1}^{d'} \left( X_j(t_i; \boldsymbol{\theta}) - \tilde{X}_j(t_i) \right)^2, \quad (4)$$

where we assume that  $d'$  variables are observed. The cost function  $E_{\mathcal{D}} : \mathbb{R}^N \rightarrow \mathbb{R}_+$  maps a  $N$ -dimensional parameter vector onto its corresponding error, thus quantifying how far the data and the model predictions are for that particular parameter set.

The aim of the data fitting procedure is to find the parameter vector  $\boldsymbol{\theta}^{**}$  that minimises the error globally subject to restrictions dictated by the problem of interest:

$$\boldsymbol{\theta}^{**} = \min_{\boldsymbol{\theta}} E_{\mathcal{D}}(\boldsymbol{\theta}), \quad \text{subject to constraints on } \boldsymbol{\theta}. \quad (5)$$

## 2.2 Definitions

- *Data set:*  $\mathcal{D}$ , a set of  $M$  observations, as defined in Eq. (2).
- *Parameter set:*  $\boldsymbol{\theta} = [\theta_1, \dots, \theta_N] \in \mathbb{R}_+^N$ . Due to the nature of the models considered,  $\theta_i \geq 0, \forall i$ .
- *Objective function:*  $E_{\mathcal{D}}(\boldsymbol{\theta})$ , the error function to be minimised, as defined in Eq. (4).
- *Set of local minima of  $E_{\mathcal{D}}(\boldsymbol{\theta})$ :*  $\mathbb{M} = \{\boldsymbol{\theta}^* \mid E_{\mathcal{D}}(\boldsymbol{\theta}^*) \leq E_{\mathcal{D}}(\boldsymbol{\theta}), \forall \boldsymbol{\theta} \in \mathcal{N}(\boldsymbol{\theta}^*)\}$  where  $\mathcal{N}(\boldsymbol{\theta}^*)$  is a neighbourhood of  $\boldsymbol{\theta}^*$ .
- *Global minimum of  $E_{\mathcal{D}}(\boldsymbol{\theta})$ :*  $\boldsymbol{\theta}^{**}$ , a parameter set such that  $E_{\mathcal{D}}(\boldsymbol{\theta}^{**}) \leq E_{\mathcal{D}}(\boldsymbol{\theta}), \forall \boldsymbol{\theta}$ . Clearly,  $\boldsymbol{\theta}^{**} \in \mathbb{M}$ .
- *Local minimisation mapping:*  $L : \mathbb{R}_+^N \rightarrow \mathbb{M}$ . Local minimisation maps  $\boldsymbol{\theta}$  onto a local minimum:  $L(\boldsymbol{\theta}) = \boldsymbol{\theta}^* \in \mathbb{M}$ .
- *Ranking and culling of local minima:*  $\{\boldsymbol{\theta}^\dagger\}_1^B = \mathcal{RC}_B(\{\boldsymbol{\theta}\}_1^J)$ . This operation ranks  $J$  parameter sets and selects the  $B$  parameter sets with the lowest  $E_{\mathcal{D}}$ .
- *Joint probability distributions of the parameters at iteration  $k$ :*  $\pi_k(\boldsymbol{\theta})$  (prior) and  $\varpi_k(\boldsymbol{\theta})$  (posterior).
- *Marginal probability distribution of the  $i^{\text{th}}$  component of  $\boldsymbol{\theta}$ :* For instance,  $\pi(\theta_i) = \int \pi(\boldsymbol{\theta}) \prod_{r \neq i} d\theta_r$ .
- *Historical prior at iteration  $k$ :*  $\zeta_k(\boldsymbol{\theta}) = \prod_{i=1}^N \zeta_k(\theta_i)$  where

$$\zeta_k(\theta_i) \sim U(\min(\mathfrak{Z}_k(\theta_i)), \max(\mathfrak{Z}_k(\theta_i))). \quad (6)$$

Here  $U(a, b)$  is a uniform distribution with support in  $[a, b]$  and  $\mathfrak{Z}_k(\theta_i) = \zeta_{k-1}^{-1} \cup \varpi_k^{-1}$  is the union of the supports of  $\varpi_k(\theta_i)$  and  $\zeta_{k-1}(\theta_i)$ .

- *Update of the prior at iteration  $k$ :*  $\pi_k(\boldsymbol{\theta}) = \prod_{i=1}^N \pi_k(\theta_i)$  with

$$\pi_k(\theta_i) \sim p_m \varpi_k(\theta_i) + (1 - p_m) \zeta_k(\theta_i), \quad (7)$$

that is, a convex mixture of the posterior and the historical prior with weight  $p_m$ , from which a new population is sampled in iteration  $k + 1$ .

---

**Algorithm 1** Squeeze-and-Breathe optimisation.

---

Set running parameters of algorithm:

$B, J \in \mathbb{N}, p_m \in [0, 1], Tol$

Choose initial priors  $\pi_0(\boldsymbol{\theta})$  and  $\zeta_0(\boldsymbol{\theta})$ .

Set  $\mathcal{H}_0 = \emptyset$  and  $k \leftarrow 1$ .

**repeat**

Let  $\mathcal{H}_k = \mathcal{H}_{k-1}$ .

Simulate  $J$  points from  $\pi_{k-1}(\boldsymbol{\theta})$  through re-population.

**for**  $\ell = 1 \rightarrow J$  **do**

Obtain local minimum  $\boldsymbol{\theta}_\ell^* = L(\boldsymbol{\theta}_\ell)$ .

Store the pair  $[\boldsymbol{\theta}_\ell^*, E_{\mathcal{D}}(\boldsymbol{\theta}_\ell^*)]$  in  $\mathcal{H}_k$ .

**end for**

Rank and cull the set of local minima:

$\mathcal{H}_k = \mathcal{RC}_B(\mathcal{H}_k)$

Define the posterior  $\varpi_k(\boldsymbol{\theta})$  from the sample  $\mathcal{H}_k$ .

Update  $\zeta_k(\boldsymbol{\theta})$  from  $\zeta_{k-1}(\boldsymbol{\theta})$  and  $\varpi_k(\boldsymbol{\theta})$ .

Update the prior:

$\pi_k(\boldsymbol{\theta}) \sim p_m \varpi_k(\boldsymbol{\theta}) + (1 - p_m) \zeta_k(\boldsymbol{\theta})$ .

$k \leftarrow k + 1$ .

**until**  $\phi_k < Tol$  and  $\mathcal{MW}(\varpi_k(\boldsymbol{\theta}), \varpi_{k-1}(\boldsymbol{\theta})) = 0$

---

- *Re-population:* Obtain population of  $J$  random points simulated from the prior  $\pi_{k-1}(\boldsymbol{\theta})$ .
- *Convergence criterion for the error:* The difference between the means of the errors of the posteriors in consecutive iterations is smaller than the pre-determined tolerance:

$$\phi_k = \overline{E_{\mathcal{D}}(\varpi_{k-1}(\boldsymbol{\theta}))} - \overline{E_{\mathcal{D}}(\varpi_k(\boldsymbol{\theta}))} < Tol. \quad (8)$$

- *Convergence criterion for the empirical distributions:* The samples of the posteriors in consecutive iterations are indistinguishable at the 5% significance level according to the nonparametric Mann-Whitney rank sum test:

$$\mathcal{MW}(\varpi_k(\boldsymbol{\theta}), \varpi_{k-1}(\boldsymbol{\theta})) = 0. \quad (9)$$

### 2.3 Description of the algorithm

Algorithm 1 presents the pseudo-code for our method using the definitions above. The iterations produce progressively more refined distributions of the parameter vector. At each iteration  $k$ , a population simulated from the prior distribution  $\pi_{k-1}(\boldsymbol{\theta})$  is locally minimised followed by ranking and culling of the local minima to create a posterior distribution  $\varpi_k(\boldsymbol{\theta})$  (squeeze step). This distribution is then combined with an encompassing historical prior to generate the updated prior  $\pi_k(\boldsymbol{\theta})$  (breathe step). The iteration loop terminates when the difference between the mean errors of consecutive posteriors is smaller than the tolerance and the samples of the posteriors are indistinguishable. We now explain these steps in detail (Fig. 1) through the BPM model (see Sec. 3.1).

1. *Formulation of the optimisation:* The data set  $\mathcal{D}$  and the model equations parameterised by  $\boldsymbol{\theta}$  allow us to define an error function  $E_{\mathcal{D}}(\boldsymbol{\theta})$  whose global minimum corresponds to the best model.

In our illustrative example, the BPM model (10) has the parameter vector  $\boldsymbol{\theta} = [\alpha, \beta]$  and the error function is depicted in Fig. 1(a). The global optimisation on the rugged landscape of this function is computationally hard.

2. *Initialisation:*

- Set the running parameters of the algorithm: the size of the simulated population,  $J$ ; the size of the surviving population after culling,  $B$ ; the update probability,  $p_m$ ; and the tolerance,  $Tol$ . In this example,  $J = 500$ ,  $B = 50$ ,  $p_m = 0.95$  and  $Tol = 10^{-5}$ .

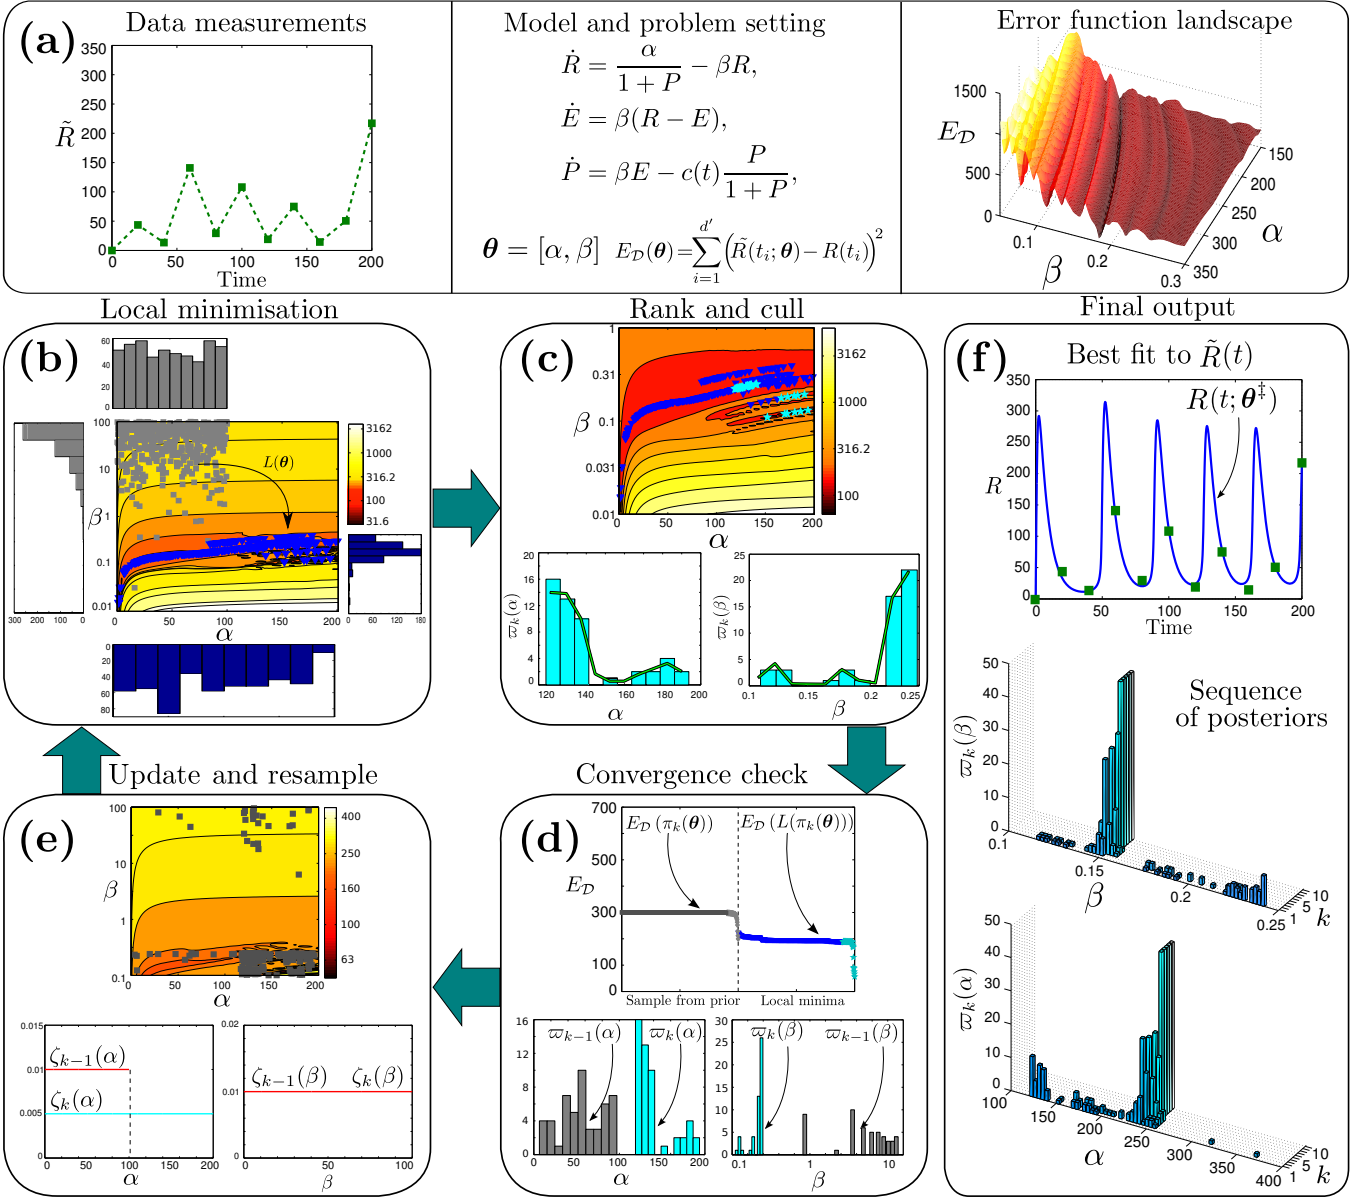


Figure 1: (Colour online) Steps of Algorithm 1 exemplified through the BPM model (10). **(a)**: The problem is defined by the data set, the model and the error function to be minimised. Note the rugged landscape of the error function in the parameter plane  $(\alpha, \beta)$ , with many local minima. **(b)**: In the first iteration, we simulate  $J$  points in parameter space from the uniform initial prior  $\pi_0(\theta)$  (squares in plot, top and left histograms) which are then minimised locally with a Nelder-Mead algorithm  $L(\theta)$  (triangles in plot, bottom and right histograms). The local optimisation aligns the parameter sets onto the level curves of  $E_{\mathcal{D}}$ . **(c)**: The  $B$  best local minima (top, light stars) are selected and considered to be samples from the posterior distribution (bottom, light histograms). **(d)**: Convergence of the error of the samples (top plot on the right,  $B$  lowest minima are the light stars) and of the posterior distributions (bottom, lighter histograms) are checked against the errors of the sample (top plot on the left) and the priors (bottom, darker histograms). **(e)**: If convergence is not achieved, the historical prior is updated (previous historical prior in bold, updated in light) and a new set of  $J$  points are simulated from the posterior with probability  $p_m$  and from the historical prior with probability  $1 - p_m$  (squares in plot). This new sample is fed back to the local minimisation step **(b)**. **(f)**: The algorithm stops when convergence is reached (after nine iterations, in this case) providing an optimal parameter set  $\theta^{\ddagger}$  and a time course (top) and the sequence of optimised posteriors at each iteration (bottom).

- Choose  $\pi_0(\boldsymbol{\theta})$ , the initial prior distribution of the parameter vector.  
In this case, we take  $\alpha$  and  $\beta$  to be independent and uniformly distributed:  $\pi_0(\boldsymbol{\theta}) \sim U(0, 100) \times U(0, 100)$ .
  - Initialise  $\zeta_0(\boldsymbol{\theta}) = \pi_0(\boldsymbol{\theta})$ , the historical prior of the parameters.
  - Simulate  $J$  points from  $\pi_0(\boldsymbol{\theta})$  to generate the initial sample  $\{\widehat{\boldsymbol{\theta}}_0\}_1^J$ .
3. *Iteration (step k)*: Repeated until termination criterion is satisfied. Figure 1 shows the first iteration of our method applied to the BPM example.

- (a) *Local minimisation*: Apply local minimisation to the simulated parameters from the ‘prior’  $\{\widehat{\boldsymbol{\theta}}_{k-1}\}_1^J$  and map them onto local minima of  $E_{\mathcal{D}}(\boldsymbol{\theta})$  to generate  $\{L(\widehat{\boldsymbol{\theta}}_{k-1})\}_1^J \in \mathbb{M}$ .  
Here we use the Nelder-Mead simplex method [24], though others can be used. Figure 1(b) shows the simulated points from  $\pi_0(\boldsymbol{\theta})$  (squares in plot) and its corresponding histograms (top and left). After local minimisation, this sample is mapped onto the dark triangles in Fig. 1(b) (dark histograms bottom and right). Note how the local minima align with the level curves of  $E_{\mathcal{D}}$  with a markedly different distribution to the uniform prior. Note also that many of the optimised values of  $\alpha$  lie outside the range of the prior (0, 100) and are now distributed over the interval (0, 200). On the other hand, the values of  $\beta$  have collapsed inside (0, 1).
- (b) *Ranking and culling*: Rank the  $J + B$  local minima from the  $k - 1$  and  $k$  iterations, select the  $B$  points with the lowest  $E_{\mathcal{D}}$  and cull (discard) the rest:

$$\mathcal{R}C_B \left( \{L(\widehat{\boldsymbol{\theta}}_{k-1})\}_1^J \cup \{\widehat{\boldsymbol{\theta}}_{k-1}^\dagger\}_1^B \right) = \{\widehat{\boldsymbol{\theta}}_k^\dagger\}_1^B.$$

We consider  $\{\widehat{\boldsymbol{\theta}}_k^\dagger\}_1^B$  to be a sample from the optimised (‘posterior’) distribution,  $\varpi_k(\boldsymbol{\theta})$  and we denote the best parameter vector of this set as

$$\boldsymbol{\theta}_k^\ddagger = \min_{E_{\mathcal{D}}} \left( \{\widehat{\boldsymbol{\theta}}_k^\dagger\}_1^B \right).$$

The  $B = 50$  best parameter sets are shown (light stars in plot) in Fig. 1(c) (bottom histograms).

- (c) *Termination criterion*: Check that the difference between the mean errors of the consecutive optimised samples is smaller than the tolerance:  $\phi_k \leq Tol$ . We also gauge the ‘convergence’ of the posteriors through the Mann-Whitney (MW) test to determine if the samples from consecutive posteriors are distinguishable:

$$\begin{aligned} & \mathcal{MW}(\varpi_{k-1}(\boldsymbol{\theta}), \varpi_k(\boldsymbol{\theta})) \\ & \equiv \mathcal{MW} \left( \{\widehat{\boldsymbol{\theta}}_{k-1}^\dagger\}_1^B, \{\widehat{\boldsymbol{\theta}}_k^\dagger\}_1^B \right), \end{aligned}$$

where  $\mathcal{MW}$  is a 0-1 flag. The MW test gives additional information about the change of the optimised posteriors from one iteration to the next.

Figure 1(d) shows the convergence check for the first iteration of the BPM model: (i) top, errors of the sampled prior (left) with errors of the local minima (right) and the  $B$  surviving points (light stars); (ii) bottom, histograms of the prior and the posterior. Clearly, in this iteration neither the error nor the distributions have converged so the algorithm does not stop.

- (d) *Update of historical prior and generation of new sample*: If convergence is not achieved, update the historical prior  $\zeta_k(\boldsymbol{\theta})$  as a uniform distribution over the union of the supports of the existing historical prior and the calculated posterior (6). Equivalently, the support of the historical prior extends over the union of the sequence of all historical priors  $\{\zeta_0(\boldsymbol{\theta}), \dots, \zeta_{k-1}(\boldsymbol{\theta})\}$  and of all posteriors  $\{\varpi_1(\boldsymbol{\theta}), \dots, \varpi_k(\boldsymbol{\theta})\}$ .

As shown in Fig. 1(e) for the BPM example, the marginal of the historical prior for  $\alpha$  is expanded to  $U(0, 200)$ , since the optimised parameter sets have reached values as high as 200. Meanwhile,

$k$	Min.	$\alpha_k^\dagger$	$\beta_k^\dagger$	Conv.	Conv.	$\phi_k$
	Error			$\varpi_k(\alpha)$	$\varpi_k(\beta)$	
1	56.0941	193.7447	0.1304	-	-	-
2	28.2735	246.7510	0.1528	No	No	133.9020
3	27.2083	248.7557	0.1532	No	No	6.8542
4	26.9838	250.3593	0.1536	No	No	0.6532
5	26.6504	251.7189	0.1538	No	No	0.3281
6	26.6504	251.7189	0.1538	No	No	0.1963
7	26.6504	251.7189	0.1538	Yes	Yes	0.0118
8	26.6504	251.7189	0.1538	No	No	0.0131
9	26.6504	251.7189	0.1538	Yes	Yes	$1.414 \times 10^{-6}$

Table 1: Results of the fitting of the BPM model with Algorithm 1: smallest error of iteration  $k$ ; the best values  $\alpha_k^\dagger$  and  $\beta_k^\dagger$ ; whether the distributions have converged; and the difference of the mean errors of the optimised population.

the  $\beta$  marginal of the historical prior remains unchanged as  $U(0, 100)$  because there has been no expansion of the support.

The historical prior is used to mutate the updated prior before the next iteration by constructing a weighted mixture of the posterior and the historical prior with weight  $p_m$ , as shown in (7). We re-populate from this updated prior by simulating from the posterior with probability  $p_m = 0.95$  and from the historical prior with probability  $(1 - p_m)$  to generate the new sample  $\{\theta_k\}_1^J$  and iterate back.

Figure 1(e) shows the sample of  $J$  points simulated from the new prior. The  $\alpha$ -components of most points are between 100 and 200 and the  $\beta$ -components are between 0.1 and 1.0, but there are a few that lie outside the support of the posterior. The process in panels (b), (c), (d), and (e) of Fig. 1 is iterated for this new set of points.

4. *Output of the algorithm:* When the convergence criteria have been met, the iteration stops at iteration  $k^*$  and the minimum of this last iteration,  $\theta_{k^*}^\dagger$ , is presented as the optimal parameter set for the model (i.e., the estimation of the global minimum  $\theta^{**}$  provided by the algorithm). We can also examine the sequence of optimised parameter distributions  $\{\varpi_1(\theta), \dots, \varpi_{k^*}(\theta)\}$  obtained for all iterations (Fig. 1(f)).

### 3 Application to biological examples

We apply our algorithm to three biological examples of interest. The first two correspond to simulated data from models in the literature, while in the third example we apply our algorithm to unpublished experimental data of the dynamical response of an inducible genetic promoter constructed for an application in Synthetic Biology.

#### 3.1 BPM model of gene-product regulation

The Bliss-Painter-Marr (BPM) model [25] describes the behaviour of a gene-enzyme-product control unit with a negative feedback loop:

$$\begin{aligned}
 \dot{R} &= \frac{\alpha}{1+P} - \beta R, \\
 \dot{E} &= \beta(R - E), \\
 \dot{P} &= \beta E - c(t) \frac{P}{1+P}.
 \end{aligned} \tag{10}$$

Here,  $R, E$  and  $P$  are the concentrations (in arbitrary units) of mRNA, enzyme and product, respectively. The degradation rate of the product has an explicit time dependence, which in this case has the form of a ramp saturation:

$$c(t) = \begin{cases} 5 + 0.2t & 0 \leq t < 50, \\ 15 & t \geq 50. \end{cases}$$

The model represents a gene that codes for an enzyme which in turn catalyses a product that inhibits the transcription of the gene. This self-inhibition can lead to oscillations, which have been shown to occur in the tryptophan operon in *E. coli* [25].

We construct a data set from simulations of this model with  $\theta_{\text{real}} = [\alpha, \beta] = [240, 0.15]$  and initial conditions  $R(0) = E(0) = P(0) = 0$ . The data set  $\mathcal{D}$  consists of 10 measurements of  $R(t)$  at particular times with added gaussian noise drawn from  $\mathcal{N}(0, 15^2)$  (Table 1, Supplemental Information). The error function  $E_{\mathcal{D}}(\theta)$  (4) corresponds to a non-convex optimisation landscape<sup>1</sup>: a complex rugged surface with many local minima making global optimisation hard (Fig. 1(a)).

We use Algorithm 1 to estimate the ‘unknown’ parameter values from the ‘measurements’ of  $R$ , as illustrated in Sec. 2.3 and Fig. 1. Feigning ignorance of the true values, we choose a uniform prior distribution with range  $[0, 100]$  for both parameters:  $\pi_0(\theta) \sim [U(0, 100), U(0, 100)]$ . The rest of the parameters are set to:  $J = 500$ ,  $B = 50$ ,  $p_m = 0.95$  and  $Tol = 10^{-5}$ . Note that the *true* value of  $\alpha$  falls outside of the assumed range of our initial prior, while the range of  $\beta$  in our initial prior is two orders of magnitude larger than its true value. This level of uncertainty about parameter values is typical in data fitting for biological models.

Figure 1 highlights a key aspect of our algorithm: the local minimisation can lead to local minima outside of the range of the initial prior. Furthermore, our definition of the historical prior ensures that successive iterations can find solutions within the largest hypercube of optimised solutions in parameter space. In this example, the algorithm moves away from the  $U(0, 100)$  prior for  $\alpha$  and finds a distribution around 240 (the true value) after three iterations, while in the case of  $\beta$ , the distribution collapses to values around 0.15 after one iteration. Although the algorithm finds the minimum  $\theta^\ddagger$  after 5 iterations, the algorithm is terminated after 9 iterations, when the posterior distributions are similar (according to the MW test) and the mean errors have also converged (Table 1). The estimated parameters for this noisy data set are  $\theta_{k^*}^\ddagger = [251.7189, 0.1530]$ . In fact, the error of the estimated parameter set is lower than that of the real parameters:  $E_{\mathcal{D}}(\theta^\ddagger) = 26.65 < E_{\mathcal{D}}(\theta_{\text{real}}) = 28.26$ , due to the noise introduced in the data. When a data set without noise is used, the algorithm finds the true value of the parameters to 9 significant digits (not shown).

### 3.2 SIR epidemics model

Susceptible-Infected-Recovered (SIR) models are widely used in epidemiology to describe the evolution of an infection in a population [11]. In its simplest form, the SIR model has three variables: the susceptible population  $S$ , the infected population  $I$  and the recovered population  $R$ :

$$\begin{aligned} \dot{S} &= \alpha - (\gamma I + d)S, \\ \dot{I} &= (\gamma S - v - d)I, \\ \dot{R} &= vI - dR. \end{aligned} \tag{11}$$

The first equation describes the change in the susceptible population, growing with birth rate  $\alpha$  and decreasing by the rate of infection  $\gamma IS$  and the rate of death  $dS$ . The infected population grows by the rate of infection  $\gamma IS$  and decreases by the rate of recovery  $vI$  and the rate of death  $dI$ . The recovered population grows by the rate of recovery  $vI$  and decreases by the death rate  $dR$ . Here we use the same form of the equations as [22].

The data generated from the model (11) (Table 2, Supplemental Information) was obtained directly from Ref. [22]. Hence the original parameter values were not known to us and further we assumed the

<sup>1</sup>We thank Markus Owen of the University of Nottingham for suggesting this example.



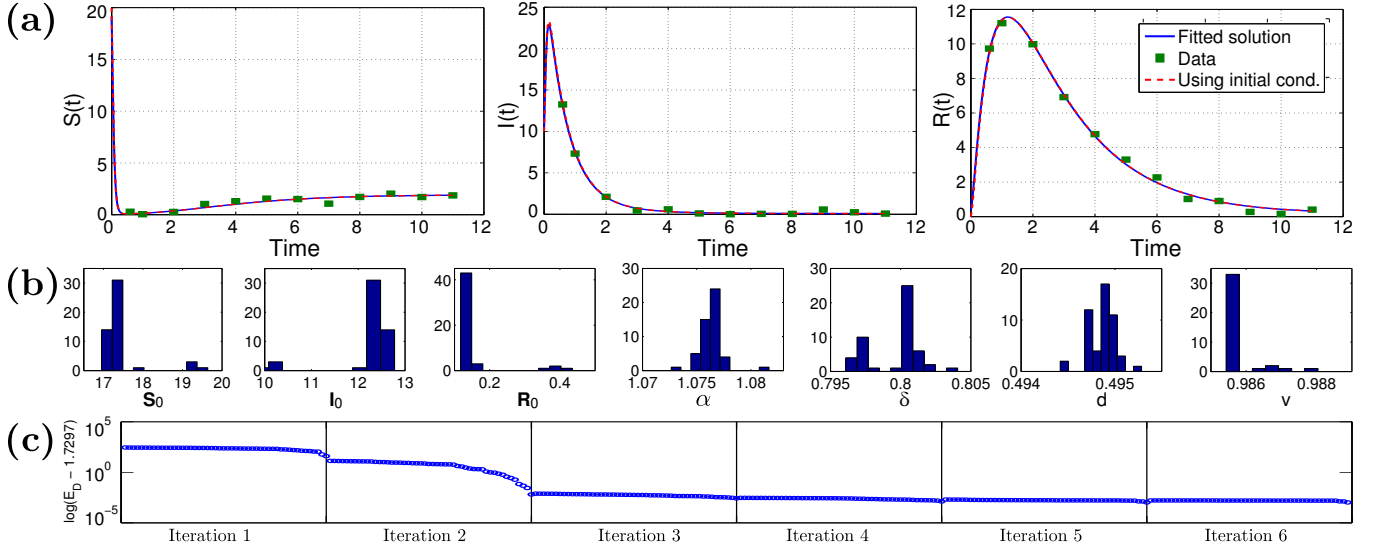


Figure 2: (Colour online) (a): Time courses of the SIR model (11). Squares are simulated ‘data’ points (Table 2, Supplemental Information) and bold lines are the model fit with the best parameters  $\alpha^\ddagger = 1.0726$ ,  $\gamma^\ddagger = 0.7964$ ,  $d^\ddagger = 0.4945$ , and  $v^\ddagger = 0.9863$  and the best fit initial conditions  $S_0^\ddagger = 19.1591$ ,  $I_0^\ddagger = 10.3016$ , and  $R_0^\ddagger = 0.3861$ . Dashed lines use the best fit parameters and the real initial conditions. The minimum error is  $E_{\mathcal{D}}(\theta^\ddagger) = 1.7297$ . (b): Histogram of the values of the 50 best parameters and initial conditions of the model obtained after convergence at six iterations. (c): Convergence of the error of the optimised samples at every iteration relative to the final error.

initial conditions also to be unknown and fitted them as parameters. We used Algorithm 1 to estimate  $\alpha$ ,  $\gamma$ ,  $v$ , and  $d$  and initial conditions  $S_0$ ,  $I_0$ , and  $R_0$ . The prior marginal distributions for all parameters were set as  $U(0, 100)$ . The other parameters were set to:  $J = 1000$ ,  $B = 50$ ,  $p_m = 0.95$  and  $Tol = 10^{-5}$ . The algorithm converged after six iterations. Figure 2(a) shows the prediction of the model (11) with the best parameters estimated by our algorithm. The fit is good with little difference between the curves obtained using the real initial conditions and the ones estimated by our method.

The posterior distributions after six iterations of the algorithm are shown on Fig. 2(b). The errors obtained after each local minimisation in a decreasing order on each iteration are shown on a semilogarithmic scale in Fig. 2(c). We can observe how the errors decrease several orders of magnitude over the first three iterations and converge steadily during the last three iterations until  $\phi_k \leq Tol$ .

### 3.3 An inducible genetic switch from Synthetic Biology

The use of inducible genetic switches is widespread in synthetic biology and bioengineering as building blocks for more complicated gene circuit architectures. An example is shown schematically in the inset of Fig. 3(a). This environment-responsive switch is used to control the expression of a target gene  $G$  (usually tagged with green fluorescent protein or *gfp*) through the addition of an exogenous small molecule  $I_1$  (e.g., isopropyl thiogalactopyranoside or IPTG). The input-output behaviour of this system can be described by the following ordinary differential equation [2, 26]:

$$\dot{G} = \alpha k_1 + \frac{k_1 I_1^{n_1}}{K_1^{n_1} + I_1^{n_1}} - dG. \quad (12)$$

Here,  $\alpha k_1$  is the basal activity of the promoter  $P_1$  and  $dG$  is the linear degradation term. The second term is a Hill function that models the cooperative transcription activation in response to the inducer  $I_1$  with maximum expression rate  $k_1$ , constant  $K_1$  and Hill coefficient  $n_1$ .

The *lacI*- $P_{lac}$  switch has been characterised experimentally in response to different doses of IPTG in Refs. [27, 28]. Equation (12) can be solved explicitly and one can use nonlinear least squares and the analytical solution to fit data at stationarity (i.e., at long times) and estimate  $\alpha$ ,  $n_1$ ,  $K_1$ , and the ratio

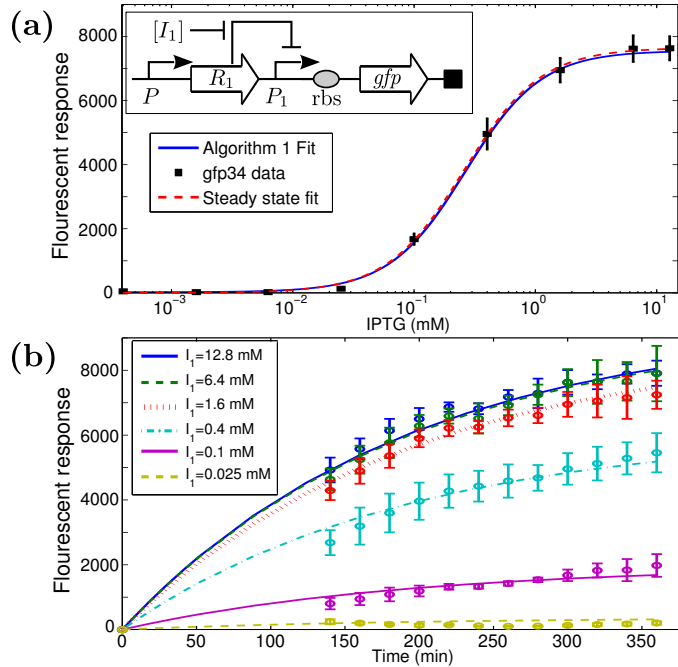


Figure 3: (Colour online) (a): *Inset*: An inducible genetic switch consisting of  $P_1$ , a negatively regulated environment-responsive promoter. The repressor  $R_1$  promoted by  $P$  regulates  $P_1$ . The switch is responsive to an exogenous inducer  $I_1$ , which binds to  $R_1$  to relieve its repression on  $P_1$  and to turn on the transcription of the downstream target gene, such as a  $gfp$ . The ribosome binding site (rbs) is used to tune the translation efficiency of the downstream gene. *Plot*: Fluorescent response of the switch with  $gfp$ -34 to different doses of IPTG (squares). Solution of Eq. (12) using the parameters obtained with Algorithm 1 (solid line), and stationary solution (dashed line). (b): Time course of the fluorescent response of the switch with  $gfp$ -34 to several doses of IPTG (circles) and time-dependent solutions of Eq. (12) using the parameters obtained with Algorithm 1 (solid lines). Similarly good fits were obtained for responses to  $I_1 = 0.0063, 0.0016, 0.0004$ , and  $0.0$  mM (not shown).

$k_1/d$ . These estimates have been obtained assuming equilibrium ( $\dot{G} = 0$ ) and initial condition  $G(0) = 0$  in Ref. [28] (Table 2).

In fact, the experiments measured time series of the expression of  $G$  every 20 minutes from  $t = 140$  to 360 min. for different doses of inducer

$$I_1 = 0.0, 3.9 \times 10^{-4}, 1.6 \times 10^{-3}, 6.3 \times 10^{-3}, 2.5 \times 10^{-2}, 0.1, 0.4, 1.6, 6.4, 12.8 \text{ mM},$$

with two different reporters ( $gfp$ -30 and  $gfp$ -34, see Tables 3 and 4 in the Supplemental Information). Instead of assuming equilibrium and using only the data for  $t > 300$  min as done previously [28], we apply Algorithm 1 to all the data with the full dynamical equation (12) to estimate  $\theta = [\alpha, k_1, n_1, K_1, d]$ . In this case, we used initial priors  $U(0, 1)$  for  $\alpha$  and  $n_1$ ; and  $U(0, 20)$  for  $k_1$ ,  $K_1$  and  $d$ . The other parameters were set to:  $J = 1000$ ,  $B = 50$ ,  $p_m = 0.95$ , and  $Tol = 10^{-5}$ .

Our algorithm converged after five iterations to the parameter values in Table 2. The parameter estimates provide good fits to both the time courses (Fig. 3(b)) and to the dose response data (Fig. 3(a)). The values of  $K_1^\dagger$  and  $n_1^\dagger$  obtained here are similar those obtained in Ref. [27] by using only stationary data. This is reassuring since these parameters are related to the dose threshold to half maximal response and to the steepness of the sigmoidal response, both static properties. On the other hand, the values of  $\alpha$  and the ratio  $k_1/d$  differ to some extent due to the (imperfect) assumption in Ref. [27] that steady state had been reached at  $t = 300$  min. As Fig. 3(b) shows,  $G$  is not at steady state then. Hence the parameter values obtained with our method should give a more faithful representation of the true dynamical response of the switch.

Parameter	Ref. [27]		Algorithm 1	
	<i>gfp-30</i>	<i>gfp-34</i>	<i>gfp-30</i>	<i>gfp-34</i>
$\alpha^\ddagger$	$0.0012 \pm 0.027$	$1.4720 \times 10^{-9}$	0.0043	0.0024
$k_1^\ddagger$	N/A	N/A	76.1354	63.6650
$n_1^\ddagger$	$1.3700 \pm 0.270$	$1.3690 \pm 0.021$	1.4832	1.3879
$K_1^\ddagger$	$0.2280 \pm 0.039$	$0.2590 \pm 0.021$	0.2467	0.2641
$d^\ddagger$	N/A	N/A	0.0069	0.0052
$k_1^\ddagger/d^\ddagger$	$9456 \pm 487$	$7648 \pm 152$	10983.34	12163.04

Table 2: Parameter values obtained from *gfp-30* and *gfp-34* data. In Ref. [27], only the steady state solution was used. Hence only the ratio of  $k_1$  and  $d$  can be estimated.

## 4 Discussion

In this work, we have presented an optimisation algorithm that brings together ingredients from Evolutionary Algorithms, local optimisation and Sequential Monte Carlo. The method is particularly useful for determining parameters of ordinary differential equation models from data. Our approach can also be used in other contexts where an optimisation problem has to be solved on complex landscapes, or when the objective function cannot be written explicitly. The algorithm proceeds by generating a population of solutions through Monte Carlo sampling from a prior distribution and refining those solutions through a combination of local optimisation and culling. A new prior is then created as a mixture of a historical prior (which records the broadest possible range of solutions found) and the distribution of the optimised population. This iterative process induces a strong concentration of the Monte Carlo sampling through local optimisation which accelerates convergence and increases precision, while at the same time the presence of the historical prior allows the possibility that solutions can be found outside of the initial presumed ranges for the parameter values.

We have illustrated the application of the algorithm to ODE models of biological interest and have found it to perform efficiently. The algorithm also works well when applied to larger problems with tens of parameters in a signal transduction model (paper in preparation). The efficiency of the algorithm hinges on selecting appropriate running parameters and priors. For instance, the number of samples from the prior  $J$  should be large enough to allow for significant sampling of the parameter space while small enough to limit the computational cost. We have found that simulating  $J = 350 - 500$  points in models of up to 10 parameters and keeping the best 15% of the local minima leads to termination within fewer than 20 iterations. In our implementation, the Nelder-Mead minimisation is capped at 300 evaluations. These guidelines would result in up to 300,000 evaluations of the objective function per iteration. Therefore our method can become computationally costly if the objective function is expensive to evaluate, e.g. in stiff models that are difficult to solve numerically. In essence, our algorithm proposes a trade-off: fewer but more costly iterations. It is important to remark that, as with any other optimisation heuristic for non-convex problems, there are no strict guarantees of convergence to the global minimum. Therefore, it is always advisable to run the method with different starting points and different settings with enough sampling points in parameter space to check for consistency of the solutions obtained.

The generation of iterative samples of the parameters draws inspiration from Monte Carlo methods [6, 18, 22] but without pursuing the strict guarantees that the nested structure of the distributions in ABC-SMC provides. Our evolutionary approach adopts a highly focused Monte Carlo sampling driven by a sharp local search with culling. Hence our iterative procedure generates samples that only reflect properties of the set of local minima (up to numerical cutoffs) without any focus on the global convergence of the distributions. As noted in Ref. [22], the distributions of the parameters (both their sequence and the final distributions) give information about the sensitivity of the parameters: parameters with narrow support will be more sensitive than those with wider support. Future developments of the method will focus on establishing a suitable theoretical framework that facilitates its use in model selection. Broadening

the choice of historical priors may be a way of establishing such framework. Currently, we make no assumptions about the parameter space, hence we use uniform distributions over the support of all the posteriors. However, other distributions (e.g., exponential or log-normal) may be considered as a way to bias the historical prior towards regions of particular interest. Other work will consider the possibility of incorporating a stochastic ranking strategy in the selection of solutions, similar to the one present in the SRES algorithm [17], in order to solve more general optimisation problems with complex feasible regions.

## Acknowledgements

The authors would like to thank C. Barnes, T. Ellis, E. Garduño, H. Harrington, M. Owen, M. Stumpf, and S. Yaliraki for their comments and suggestions.

**Funding:** MBD is supported by a BBSRC-Microsoft Research Dorothy Hodgkin Postgraduate Award. This work was partly supported by the US Office of Naval Research, and by BBSRC through LoLa grant BB/G020434/1 (MB) and SABR grant BB/F005210/2 (MB), and by EPSRC through grant EP/I017267/1 under the *Mathematics underpinning the Digital Economy* program (MB).

## References

- [1] Edelstein-Keshet L. *Mathematical models in biology*. Classics in applied mathematics. SIAM; 1988.
- [2] Alon U. *An introduction to systems biology: design principles of biological circuits*. Chapman and Hall/CRC mathematical & computational biology series. Chapman & Hall/CRC; 2007.
- [3] Strogatz SH. *Nonlinear Dynamics And Chaos. With Applications to Physics, Biology, Chemistry, and Engineering*. Studies in nonlinearity. Perseus Books Group; 1994.
- [4] Brown KS, Sethna JP. Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E*. 2003 Aug;68(2):021904.
- [5] Gutenkunst RN, Waterfall JJ, Casey FP, Brown KS, Myers CR, Sethna JP. Universally Sloppy Parameter Sensitivities in Systems Biology Models. *PLoS Comput Biol*. 2007 10;3(10):e189.
- [6] Toni T, Stumpf MPH. Simulation-based model selection for dynamical systems in systems and population biology. *Bioinformatics*. 2010 Jan;26(1):104–110.
- [7] Yates A, Chan CCW, Callard RE, George AJT, Stark J. An approach to modelling in immunology. *Brief Bioinform*. 2001;2(3):245–257.
- [8] Gershenfeld NA. *The nature of mathematical modeling*. Cambridge University Press; 1999.
- [9] Lawson CL, Hanson RJ. *Solving least squares problems*. Classics in applied mathematics. SIAM; 1995.
- [10] Brewer D, Barenco M, Callard R, Hubank M, Stark J. Fitting ordinary differential equations to short time course data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2008;366(1865):519–544.
- [11] Anderson RM, May RM. *Infectious Diseases of Humans Dynamics and Control*. Oxford University Press; 1992.
- [12] Chen WW, Niepel M, Sorger PK. Classic and contemporary approaches to modeling biochemical reactions. *Genes Dev*. 2010 Sep;24(17):1861–1875.

- [13] Papachristodoulou A, Recht B. Determining Interconnections in Chemical Reaction Networks. In: American Control Conference, 2007.; 2007. p. 4872–4877.
- [14] Powell MJD. Direct Search Algorithms for Optimization Calculations. *Acta Numerica*. 1998;7(-1):287–336.
- [15] Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by simulated annealing. *Science*. 1983 May;220(4598):671–680.
- [16] Mitchell TM. *Machine Learning*. New York: McGraw-Hill; 1997.
- [17] Runarsson TP, Yao X. Stochastic ranking for constrained evolutionary optimization. *IEEE Transactions on Evolutionary Computation*. 2000 sep;4(3):284 –294.
- [18] Sisson SA, Fan Y, Tanaka MM. Sequential Monte Carlo without likelihoods. *Proc Natl Acad Sci USA*. 2007 Feb;104(6):1760–1765.
- [19] Schwefel HP. *Evolution and optimum seeking*. Sixth-generation computer technology series. Wiley; 1995.
- [20] Moles CG, Mendes P, Banga JR. Parameter Estimation in Biochemical Pathways: A Comparison of Global Optimization Methods. *Genome Research*. 2003;13(11):2467–2474.
- [21] Zi Z, Klipp E. SBML-PET: a Systems Biology Markup Language-based parameter estimation tool. *Bioinformatics*. 2006;22(21):2704–2705.
- [22] Toni T, Welch D, Strelkowa N, Ipsen A, Stumpf MPH. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of The Royal Society Interface*. 2009;6(31):187–202.
- [23] Sontag E. For Differential Equations with  $r$  Parameters,  $2r+1$  Experiments Are Enough for Identification. *Journal of Nonlinear Science*. 2002;12:553–583.
- [24] Nelder JA, Mead R. A Simplex Method for Function Minimization. *The Computer Journal*. 1965;7(4):308–313.
- [25] Bliss RD, Painter PR, Marr AG. Role of feedback inhibition in stabilizing the classical operon. *Journal of Theoretical Biology*. 1982;97(2):177 – 193.
- [26] Szallasi Z, Stelling J, Periwal V. *System modeling in cell biology: from concepts to nuts and bolts*. Bradford Book. MIT Press; 2006.
- [27] Wang B. *Design and Functional Assembly of Synthetic Biological Parts and Devices*. Imperial College London; 2010.
- [28] Wang B, Kitney RI, Joly N, Buck M. Engineering modular and orthogonal genetic logic gates for robust digital-like synthetic biology. *Nat Commun*. 2011;2:508.