



Bayesian prediction for physical models with application to the optimization of the synthesis of pharmaceutical products using chemical kinetics

Antony M. Overstall^{a,*}, David C. Woods^a, Kieran J. Martin^b

^a University of Southampton, Southampton SO17 1BJ, UK

^b Roche, Welwyn Garden City AL7 1TW, UK

ARTICLE INFO

Article history:

Received 15 November 2017

Received in revised form 22 October 2018

Accepted 22 October 2018

Available online 2 November 2018

Keywords:

Approximate coordinate exchange

Multivariate Gaussian process

Gibbs sampling

Parallel tempering

Riemann manifold Langevin

Metropolis–Hastings

ABSTRACT

Quality control in industrial processes is increasingly making use of prior scientific knowledge, often encoded in physical models that require numerical approximation. Statistical prediction, and subsequent optimization, is key to ensuring the process output meets a specification target. However, the numerical expense of approximating the models poses computational challenges to the identification of combinations of the process factors where there is confidence in the quality of the response. Recent work in Bayesian computation and statistical approximation (emulation) of expensive computational models is exploited to develop a novel strategy for optimizing the posterior probability of a process meeting specification. The ensuing methodology is motivated by, and demonstrated on, a chemical synthesis process to manufacture a pharmaceutical product, within which an initial set of substances evolve according to chemical reactions, under certain process conditions, into a series of new substances. One of these substances is a target pharmaceutical product and two are unwanted by-products. The aim is to determine the combinations of process conditions and amounts of initial substances that maximize the probability of obtaining sufficient target pharmaceutical product whilst ensuring unwanted by-products do not exceed a given level. The relationship between the factors and amounts of substances of interest is theoretically described by the solution to a system of ordinary differential equations incorporating temperature dependence. Using data from a small experiment, it is shown how the methodology can approximate the multivariate posterior predictive distribution of the pharmaceutical target and by-products, and therefore identify suitable operating values.¹

© 2018 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Pharmaceutical products are often manufactured by chemical synthesis, with an initial set of chemical substances evolving over time via a series of chemical reactions into a substance or substances of interest. One (or more) of these substances will be the target pharmaceutical product while the others will be unwanted or harmful *by-product* substances. In addition to a time dependence, the chemical synthesis process may be hypothesized to depend on various controllable factors. The aim of the chemical engineers will be to manipulate the controllable factors such that specification limits on the

* Correspondence to: Southampton Statistical Sciences Research Institute, University of Southampton, Southampton, SO17 1BJ, UK.

E-mail address: A.M.Overstall@southampton.ac.uk (A.M. Overstall).

¹ Materials to replicate the analysis can be found at [www.github.com/amo105/chemicalkinetics](https://github.com/amo105/chemicalkinetics).

final substances are satisfied. For example, the quantity of pharmaceutical product may be required to exceed some specified level whilst the amounts of unwanted by-products are correspondingly kept below a certain level.

It is common for a physical model, for example derived from scientific theory as the solution to a set of ordinary differential equations (ODEs), to be postulated to approximately describe the dependence of the synthesis process on the controllable factors (and time); see chapters in [am Ende \(2011\)](#). In addition, often an experiment can be performed within which the amounts of the substances of interest are measured for several different specifications of the controllable factors. These observed responses can then be used to inform the relationship between controllable factors and the synthesis process, e.g. through the estimation of unknown tuning parameters.

In this paper, we develop and apply methodology for the statistical modelling of such experiments, motivated by a pharmaceutical exemplar which measured three substances of interest: the pharmaceutical product and two unwanted by-products. The dynamics behind the synthesis process are approximately described by a system of non-linear ODEs, the solution to which forms a physical model for the amounts of the three substances at a given time. This model depends on a set of controllable factors (including time) and unknown parameters. A small experiment has been conducted where the amounts of the substances of interest have been observed for different controllable factors at certain time points. The aim is to use this observed data to estimate the unknown parameters and then to use these estimates to maximize the probability of the predicted amounts satisfying known specification limits as a function of the controllable factors and time. The approach used here will be Bayesian, giving a coherent method of propagating uncertainty on the unknown parameters (given the observed data) through to the predicted amounts. This approach is closely aligned with the concept of “pharmaceutical design space”,¹ a set of combinations of values of the controllable factors that have been demonstrated to provide assurance of quality. See, for example, [Peterson \(2008\)](#), [Peterson and Yahyah \(2009\)](#) and [Lebrun et al. \(2013\)](#) for related Bayesian approaches to the definition of a design space using linear statistical modelling methods. Beyond pharmaceutical design space, in the wider field of process optimization, see, for example, [Chiao and Hamada \(2001\)](#), [Peterson \(2004\)](#) and [Del Castillo \(2007\)](#).

Despite the apparent simplicity and appealing nature of the Bayesian modelling approach, several computational challenges remain. Firstly, the solution to the system of ODEs is analytically intractable and so a computationally expensive numerical solution is required (e.g. [Iserles, 2009](#)). The computational expense of this approximation will impact the computing time required both for the estimation of the unknown parameters through the evaluation of their posterior distribution and the predictions for new combinations of the controllable factors. To reduce computational expense, we make use of (multivariate) Gaussian process emulators to accelerate both the Markov chain Monte Carlo (MCMC) algorithm employed and the prediction of future responses. Secondly, some of the observations fall below a minimum threshold, χ , under which the actual amount cannot be measured reliably, resulting in left-censored observations. Instead of assuming that the true observations are equal to zero or χ with certainty, we assume that the true observations are unknown, subject to being in the interval $(0, \chi)$, within a natural Bayesian formulation. Thirdly, we know that the physical model derived from the solution to the ODEs only provides an approximation to the true process. We account for any potential mismatch by introducing an explicit model discrepancy term (c.f. [Kennedy and O’Hagan, 2001](#)). Fourthly, the function giving the probability of satisfying the specification limits for given controllable factors will be noisy and computationally expensive to evaluate making optimization non-trivial.

The contribution of this paper is the novel application of a combination of modern computational and statistical methodology to address these challenges, and the demonstration of this approach on an important practical problem.

The paper is organized as follows. In Section 2 we describe the motivating experiment, and the physical and statistical models for the observations in more detail. In Section 3 we describe the methodology including construction of Gaussian process emulators. In Section 4 we present the results of applying the methodology to the motivating experiment, before concluding with a discussion in Section 5.

2. Background

2.1. Chemical reactions

The motivating chemical synthesis process involves nine substances, labelled A to I. The substances of interest are E, F and H, where F constitutes the pharmaceutical product with E and H being unwanted by-product substances. Let $A(t)$ denote the amount (in mols) of substance A at time t (in seconds), with similar notation for the remaining substances. At time $t = 0$, initial amounts of A, D and E, denoted by $A_0 = A(0)$, $D_0 = D(0)$ and $E_0 = E(0)$, respectively, are introduced into the laboratory apparatus at a specified temperature (λ in Kelvins) and volume (V , in litres). Let $\mathbf{x} = (x_1, \dots, x_5) = (A_0, D_0, E_0, \lambda, V)$ be an $F \times 1$ vector ($F = 5$) and assume that $\mathbf{x} \in \mathcal{X} = \prod_{f=1}^F [x_{f,\min}, x_{f,\max}]$. The upper and lower limits, $x_{f,\min}$ and $x_{f,\max}$, for each factor are given in [Table 1](#).

The chemical synthesis process is governed by the following series of chemical reactions:



¹ <https://www.fda.gov/downloads/drugs/guidances/ucm073507.pdf>.

Table 1
Lower and upper limits of the controllable factors in the motivating experiment.

Definition	Name	Lower limit, $x_{f,min}$	Upper limit, $x_{f,max}$	Units
Initial amount of A	$x_1 = A_0$	22.5	45	mols
Initial amount of D	$x_2 = D_0$	91.4	91.59	mols
Initial amount of E	$x_3 = E_0$	26.42	26.47	mols
Temperature	$x_4 = \lambda$	298.15	313.15	Kelvins (K)
Volume	$x_5 = V$	31.28	32.56	litres (l)

where k_1 , k_2 , and k_3 denote the chemical reaction rates. The second line denotes a reversible reaction and thus k_2 is split into k_2^- and k_2^+ , for the forward and backward reactions, respectively.

Let $[A](t) = A(t)/V$ be the concentration (in mol/litre) at time t of A, let $\dot{A}(t) = d[A](t)/dt$ be the corresponding time derivative, and assume similar definitions for the remaining substances; B to I. Chemical reactions (1) lead to the following system of non-linear ODEs (suppressing the dependence on t):

$$\left. \begin{aligned} \dot{A} &= -k_1[A] \\ \dot{B} &= 2k_1[A] - k_3[F][B] \\ \dot{C} &= k_1[A] \\ \dot{D} &= -k_2^- [B][D][E] + k_2^+ [B][G][F] \\ \dot{E} &= -k_2^- [B][D][E] + k_2^+ [B][G][F] \\ \dot{F} &= k_2^- [B][D][E] - k_2^+ [B][G][F] - k_3[F][B] \\ \dot{G} &= k_2^- [B][D][E] - k_2^+ [B][G][F] \\ \dot{H} &= k_3[F][B] \\ \dot{I} &= k_3[F][B] \end{aligned} \right\} \text{for } t \in \mathcal{T} = [0, 3000] \text{ seconds.} \quad (2)$$

The initial concentrations of all nine substances are zero, except for A, D and E which are $[A](0) = A_0/V$, $[D](0) = D_0/V$ and $[E](0) = E_0/V$, respectively.

For given chemical reaction rates, the solution $([A](t), \dots, [I](t))$ from (2) provides the concentrations of the nine substances for given combinations of the controllable factors, \mathbf{x} , and reaction rates. The amounts of the nine substances are then given by $(A(t), \dots, I(t)) = V \times ([A](t), \dots, [I](t))$. However, the solution is analytically intractable and hence we employ a computationally expensive numerical method to find an approximate solution. In particular, we use the variable coefficient ordinary differential equation solver (Brown et al., 1989) as implemented in the R (R Core Team, 2018) package `deSolve` (Soetaert et al., 2010). However the methodology described in Section 3 can be used in conjunction with any numerical method.

The first reaction is assumed to be instantaneous, which the chemical engineers model by assuming $k_1 = 10000$. The remaining reaction rates, k_2^- , k_2^+ and k_3 , are unknown, and the temperature dependence is incorporated into the system of ODEs via the Arrhenius equation (see, for example, Laidler (1984)). Hence, the general form for a reaction rate is

$$k_i = k_i^{(r)} \exp\left(\frac{E_i}{G} \left(\frac{1}{\lambda} - \frac{1}{\lambda^{(r)}}\right)\right), \quad i = 2, 3,$$

where $\lambda^{(r)}$ is the reference temperature (here 298.15 K), $k_i^{(r)} > 0$ is the reaction rate at the reference temperature, $E_i > 0$ is the activation energy and $G = 8.31445 \text{ J mol}^{-1} \text{ K}^{-1}$ is the gas constant. By definition, the activation energies for the reversible reactions are equal, i.e. $E_2^+ = E_2^- = E_2$. This means there are $p = 5$ unknown physical parameters, $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_5) = (k_2^{(r)-}, k_2^{(r)+}, k_3^{(r)}, E_2, E_3)$. To increase computational efficiency of the methodology described in Section 3, we operate on the logarithm scale, i.e. consider $\theta_k = \log \gamma_k$, for $k = 1, \dots, p$, and let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p) \in \Theta = \mathbb{R}^p$.

The quality control question of interest is to make choices of initial amounts of A, D and E, temperature, volume and time, such that the amount of the three substances of interest satisfy the following specification limits (in mols):

$$\begin{aligned} E(t) &< 3, \\ F(t) &> 20, \\ H(t) &< 3. \end{aligned} \quad (3)$$

Recall that F is the pharmaceutical product and its amount needs to be sufficiently large to make the production process economically viable. However E and H represent unwanted by-products with upper limits for safety reasons.

2.2. Experiment and physical model

We estimate the values of the unknown physical parameters $\boldsymbol{\theta}$ using observations from a small experiment. The chemical reactions (1) are observed for $l' = 6$ combinations of the controllable factors \mathbf{x} , given in Table 2. For run $i = 1, \dots, l'$, we observe the amounts of E, F and H at a series of n_i times t_{i1}, \dots, t_{in_i} . The times range from 0.5 to 2902 s, with n_i ranging from

Table 2
The six combinations of the controllable factors in the experiment.

Run	Symbol	Initial amount (in mols)			Temperature (λ , Kelvin)	Volume (V , litres)
		A_0	D_0	E_0		
1	\mathbf{x}_1	22.50	91.59	26.47	298.15	31.31
2	\mathbf{x}_2	45.00	91.59	26.47	298.15	32.56
3	\mathbf{x}_3	22.50	91.50	26.45	313.15	31.28
4	\mathbf{x}_4	45.00	91.50	26.45	313.15	32.53
5	\mathbf{x}_5	33.75	91.40	26.42	305.65	31.88
6	\mathbf{x}_6	33.75	91.40	26.42	305.65	31.88

17 to 20 time points per run. In total, there are $n = \sum_{i=1}^{I'} n_i = 109$ observations of the amounts over all I' runs. Note how run 6 is a repetition of run 5, thus there are $l = I' - 1 = 5$ unique treatments. Runs 5 and 6 have observations taken at 18 different time points. However the last two time points for run 5 are at 1620 and 1825 s, whereas the last two time points for run 6 are taken four seconds later. Therefore there are $m = 93$ unique combinations of controllable factors and time points. Let (\mathbf{x}'_j, t_j) be the values of the controllable factors for each of these unique combinations ($j = 1, \dots, m$). The experiment was not designed to be statistically optimal for either the estimation of parameters θ or prediction for new treatments.

Let $K = 3$ be the number of substances of interest. We denote by $\mu(\theta; \mathbf{x}, t) = \log(E(t), F(t), H(t))$ the $K \times 1$ vector giving the log of the theoretical amount of the three substances of interest (E, F and H) obtained from the numerical solution to the system of ODEs at \mathbf{x} , time t and parameters θ .

Let \mathbf{Y}_i be the $n_i \times K$ matrix containing the log observed amounts of E, F and H for the i th run, for $i = 1, \dots, I'$. The j th row of \mathbf{Y}_i , denoted by \mathbf{y}_{ij} is a $K \times 1$ vector containing the observed amounts for the j th time point of the i th run, for $i = 1, \dots, I'$ and $j = 1, \dots, n_i$. Let \mathbf{Y} be the $n \times K$ matrix formed by stacking the matrices $\mathbf{Y}_1, \dots, \mathbf{Y}_{I'}$, i.e.

$$\mathbf{Y} = \begin{pmatrix} \mathbf{Y}_1 \\ \vdots \\ \mathbf{Y}_{I'} \end{pmatrix}.$$

The experimental apparatus cannot detect substance amounts lower than $\chi = 0.01$ mols. Therefore, we have censored values when the observed amount is less than χ . In total, there are zero censored observations in the first column of \mathbf{Y} , and 5 and 45 in the second and third columns, respectively. We take account of the censoring in our statistical modelling strategy described in Section 2.3.

2.3. Statistical model

We assume that

$$\mathbf{Y} = \mathbf{G}(\mathbf{M}(\theta) + \mathbf{D}) + \mathbf{E}. \tag{4}$$

In (4), $\mathbf{M}(\theta)$ is the $m \times K$ matrix of unique model solutions, given by

$$\mathbf{M}(\theta) = \begin{pmatrix} \mathbf{M}_1(\theta) \\ \vdots \\ \mathbf{M}_l(\theta) \end{pmatrix},$$

where $\mathbf{M}_i(\theta)$, for $i = 1, \dots, l - 1$, is the $n_i \times K$ matrix with j th row given by $\mu(\theta; \mathbf{x}_i, t_{ij})$, and $\mathbf{M}_l(\theta)$ is the $(n_l + 2) \times k$ matrix where the j th row is given by $\mu(\theta; \mathbf{x}_l, t_{lj})$, for $j = 1, \dots, n_l$, and the $(n_l + 1)$ th and $(n_l + 2)$ th rows are given by $\mu(\theta; \mathbf{x}_l, 1624)$ and $\mu(\theta; \mathbf{x}_l, 1829)$, respectively (corresponding to the two time points that differ between runs 5 and 6).

Furthermore, \mathbf{D} is the $m \times K$ matrix of model discrepancy errors given by

$$\mathbf{D} = \begin{pmatrix} \mathbf{D}_1 \\ \vdots \\ \mathbf{D}_l \end{pmatrix},$$

which represents the discrepancy between the physical model and the true mean value of the process (Kennedy and O'Hagan, 2001). The $n \times m$ binary incidence matrix \mathbf{G} identifies the rows of $\mathbf{M} + \mathbf{D}$ corresponding to the rows of \mathbf{Y} . Finally, \mathbf{E} is the $n \times K$ matrix of observational errors.

The distribution of \mathbf{E} is assumed to be

$$\mathbf{E} \sim \text{MN}(\mathbf{0}, \Omega, \mathbf{T}),$$

where MN denotes the matrix normal distribution (e.g. Dawid, 1981), Ω is an unknown $K \times K$ column covariance matrix, and \mathbf{T} is an $n \times n$ row correlation matrix. The matrix $\mathbf{T} = \text{diag}\{\mathbf{T}_1, \dots, \mathbf{T}_{I'}\}$ is block-diagonal, with j th element of \mathbf{T}_i

given by

$$T_{i,jl} = \exp\left(-\rho (t_{ij} - t_{il})^2\right) \quad j, l = 1, \dots, n_i,$$

with time correlation parameter $\rho > 0$ assumed unknown. This covariance structure for \mathbf{E} assumes that the correlation between observations from the same run is dependent on the difference in time, but observations from different runs are independent.

We complete the model specification by specifying prior distributions for the discrepancy \mathbf{D} and the model parameters. The discrepancy is assumed to follow a matrix normal distribution,

$$\mathbf{D} \sim \text{MN}(\mathbf{0}, \Sigma, \mathbf{S}), \quad (5)$$

where Σ is an unknown $K \times K$ covariance matrix and \mathbf{S} is an $m \times m$ correlation matrix. The jl th element of \mathbf{S} is given by

$$S_{jl} = \exp\left(-\psi_1 \sum_{f=1}^F (d_{jf} - d_{lf})^2 - \psi_2 (t_j - t_l)^2\right), \quad (6)$$

for $j, l = 1, \dots, m$ and where the model discrepancy correlation parameters $\psi_1 > 0$ and $\psi_2 > 0$ are unknown. In (6), $d_{jf} = (x'_{jf} - x_{f,\min}) / (x_{f,\max} - x_{f,\min})$ is the scaled value of the f th controllable factor, for $f = 1, \dots, F$ and $j = 1, \dots, m$, where x'_{jf} is the f th element of \mathbf{x}'_j . This correlation structure allows the differences in model discrepancy between different treatments to depend on the “distance” between the controllable factors and time points for each treatment.

The chemical engineers have some limited prior knowledge on the location of θ . It is believed that the reaction rates are likely to lie between 10^{-8} and 10^{-4} and the activation energies between 10^2 and 10^6 . We encode this information by assuming that the reaction rates and activation energies have independent log-normal prior distributions with hyperparameters chosen so that there is probability 0.95 that the value of interest lies between the two limits above. This is achieved by setting the 0.025 and 0.975 quantiles to be the lower and upper limits respectively. In summary, we assume

$$\log \theta \sim N(\boldsymbol{\mu}, \Delta),$$

where $\boldsymbol{\mu} = (-13.8, -13.8, -13.8, 9.21, 9.21)^T$ and $\Delta = 5.52\mathbf{I}_p$. There is no prior information available for the remaining parameters, so we specify vague prior distributions that contribute negligible information. The variance matrices Ω and Σ are both assumed to have inverse-Wishart prior distributions with $\nu = 4$ degrees of freedom and an identity scale matrix (Gelman et al., 2014, Chapter 3). The correlation parameters ρ , ψ_1 and ψ_2 are given independent exponential prior distributions with mean equal to one following the arguments of, for example, Overstall and Woods (2016). In Section 4 we discuss the results of a sensitivity analysis to assess the robustness to the choice of these vague prior distributions.

Let $\mathbf{y} = \text{vec}(\mathbf{Y})$, $\mathbf{m}(\theta) = \text{vec}(\mathbf{M}(\theta))$, $\mathbf{d} = \text{vec}(\mathbf{D})$ and $\mathbf{e} = \text{vec}(\mathbf{E})$, with the vec operator stacking columns of a matrix. Then a model specification equivalent to (4) is given by

$$\mathbf{y} = \mathbf{H}(\mathbf{m}(\theta) + \mathbf{d}) + \mathbf{e}, \quad (7)$$

where $\mathbf{H} = \mathbf{I}_k \otimes \mathbf{G}$,

$$\mathbf{e} \sim N(\mathbf{0}, \Omega \otimes \mathbf{T}), \quad (8)$$

$$\mathbf{d} \sim N(\mathbf{0}, \Sigma \otimes \mathbf{S}), \quad (9)$$

and \otimes denotes the Kronecker product.

Let \mathbf{y}_s and \mathbf{y}_c denote the elements of \mathbf{y} which are fully observed (i.e. greater than $\log \chi$) and censored (i.e. less than $\log \chi$), respectively. Therefore, we need to evaluate the posterior distribution (conditional on \mathbf{y}_s) of the model parameters and censored observations. This distribution is given by

$$\pi(\theta, \Omega, \rho, \mathbf{d}, \Sigma, \boldsymbol{\psi}, \mathbf{y}_c | \mathbf{y}_s) \propto \pi(\mathbf{y} | \theta, \mathbf{d}, \Omega, \rho) \pi(\mathbf{d} | \Sigma, \boldsymbol{\psi}) \pi(\theta) \pi(\Omega) \pi(\rho) \pi(\Sigma) \pi(\boldsymbol{\psi}), \quad (10)$$

where $\pi(\mathbf{y} | \theta, \mathbf{d}, \Omega, \rho)$ is the complete data likelihood given by (7) and (8), $\pi(\mathbf{d} | \Sigma, \boldsymbol{\psi})$ is the density of the vectorized model discrepancy (9), and $\pi(\theta)$, $\pi(\Omega)$, $\pi(\rho)$, $\pi(\Sigma)$ and $\pi(\boldsymbol{\psi})$ are prior densities for the model parameters.

2.4. Prediction

Let \mathbf{y}_0 be a $K \times 1$ vector denoting the predicted log amounts of E, F and H for arbitrary controllable factors $\mathbf{x}_0 \in \mathcal{X}$ and $t_0 \in \mathcal{T}$. The posterior predictive distribution of \mathbf{y}_0 is given by

$$\pi(\mathbf{y}_0 | \mathbf{y}_s) = \int \pi(\mathbf{y}_0 | \theta, \mathbf{D}, \boldsymbol{\psi}, \Omega, \Sigma) \pi(\theta, \mathbf{D}, \boldsymbol{\psi}, \Omega, \Sigma | \mathbf{y}_s) d\theta d\mathbf{D} d\boldsymbol{\psi} d\Omega d\Sigma, \quad (11)$$

where

$$\mathbf{y}_0 | \theta, \mathbf{D}, \boldsymbol{\psi}, \Omega, \Sigma \sim N(\boldsymbol{\mu}(\theta; \mathbf{x}, t) + \mathbf{D}^T \mathbf{S}^{-1} \mathbf{s}, (1 - \mathbf{s}^T \mathbf{S}^{-1} \mathbf{s}) \Sigma + \Omega),$$

and \mathbf{s} is an $m \times 1$ vector with j th element given by

$$s_j = \exp \left(-\psi_1 \sum_{f=1}^F (d_{jf} - d_{0f})^2 - \psi_2 (t_j - t_0)^2 \right),$$

with $d_{0f} = (x_{0f} - x_{f,min}) / (x_{f,max} - x_{f,min})$ being the scaled value of each element of \mathbf{x}_0 , for $f = 1, \dots, F$.

The probability of E, F and H satisfying the specification limits (3) at point (\mathbf{x}_0, t_0) is given by

$$P(\mathbf{y}_0 \in \mathcal{Y} | \mathbf{y}_S) = \int_{\mathcal{Y}} \pi(\mathbf{y}_0 | \mathbf{y}_S) d\mathbf{y}_0,$$

where $\mathcal{Y} = \{\boldsymbol{\eta} \in \mathbb{R}^K | \eta_1 < \log 3, \eta_2 > \log 20, \eta_3 < \log 3\}$.

3. Methodology

The posterior and predictive distributions given by (10) and (11), respectively, will be analytically intractable. Our general approach to numerically approximating these distributions uses two phases. The **Sampling Phase** involves generating a sample from the posterior distribution using MCMC methods. The **Prediction Phase** uses this MCMC sample to generate a sample from the posterior predictive distribution and thus estimate the probability $P(\mathbf{y}_0 \in \mathcal{Y} | \mathbf{y}_S)$ by the proportion of sampled values satisfying the specification limits. These phases make use of Gibbs sampling and parallel tempering algorithms, and multivariate Gaussian process emulators for the numerical solution to the ODEs.

3.1. Gibbs sampling and parallel tempering

We use Gibbs sampling in conjunction with parallel tempering to generate a sample from the joint posterior distribution (10) of model parameters and censored observations. To improve the efficiency of the Gibbs sampler, we employ hierarchical centering (Papaspiliopoulos et al., 2003) and reparameterize model (7) as

$$\mathbf{y} = \mathbf{H}\mathbf{d}^* + \mathbf{e}, \tag{12}$$

where

$$\mathbf{d}^* \sim N(\mathbf{m}(\boldsymbol{\theta}), \boldsymbol{\Sigma} \otimes \mathbf{S}). \tag{13}$$

Sampled values of \mathbf{d}^* can easily be transformed to values of \mathbf{d} using $\mathbf{d} = \mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})$. The posterior distribution of model parameters and censored observations is now given by

$$\pi(\boldsymbol{\theta}, \boldsymbol{\Omega}, \rho, \mathbf{d}^*, \boldsymbol{\Sigma}, \boldsymbol{\psi}, \mathbf{y}_C | \mathbf{y}_S) \propto \pi(\mathbf{y} | \mathbf{d}^*, \boldsymbol{\Omega}, \rho) \pi(\mathbf{d}^* | \boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\psi}) \pi(\boldsymbol{\theta}) \pi(\boldsymbol{\Omega}) \pi(\rho) \pi(\boldsymbol{\Sigma}) \pi(\boldsymbol{\psi}),$$

where $\pi(\mathbf{y} | \mathbf{d}^*, \boldsymbol{\Omega}, \rho)$ is the complete-data likelihood defined by (12) and $\pi(\mathbf{d}^* | \boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\psi})$ is defined by (13). The full conditional posterior distributions of \mathbf{d}^* , $\boldsymbol{\Omega}$ and $\boldsymbol{\Sigma}$ are of known form and are given in Appendix A. The full conditional posterior distribution of \mathbf{y}_C is a truncated multivariate normal distribution (see Appendix A) and there exist efficient methods for generating values from such distributions (Geweke, 1991).

For the remaining parameters $\boldsymbol{\theta}$, ρ and $\boldsymbol{\psi}$, the full conditional distributions are not of known form, so a Metropolis-within-Gibbs step is employed. Specifically, we exploit the Riemann manifold Langevin Metropolis–Hastings (RMLMH) algorithm (Girolami and Calderhead, 2011). This method uses derivative information on the posterior surface to construct an efficient proposal distribution. The algorithmic details are briefly described in Appendix B and the components necessary for its application to the full conditional posterior distributions of $\boldsymbol{\theta}$, ρ and $\boldsymbol{\psi}$ are given in Appendix C.

Parallel tempering (see, for example, Gelman et al., 2014, pp. 299–300) is used to improve the efficiency of sampling from a potentially multi-modal posterior distribution. It involves defining R distributions by a series of increasing “temperatures”, where the lowest temperature distribution corresponds to the posterior distribution. An MCMC chain is initialized under each distribution, and the parallel tempering scheme allows for swaps between the states of the chains for distributions at different temperatures. The lowest temperature MCMC chain is a sample from the posterior distribution. Parallel tempering allows larger moves (e.g. between different areas of high posterior density) made at the higher temperature chains to be passed down to the lower temperature chains, hence improving mixing.

Let $\boldsymbol{\delta} = (\boldsymbol{\theta}, \boldsymbol{\Omega}, \rho, \mathbf{d}^*, \boldsymbol{\Sigma}, \boldsymbol{\psi}, \mathbf{y}_C)$ be the collection of unknown parameters and censored observations. For $r = 1, \dots, R$, define the distribution for the r th chain to have density proportional to

$$U_r(\boldsymbol{\delta}) = \exp \left(\frac{\log \pi(\mathbf{y}_S | \boldsymbol{\delta}) + \log \pi(\boldsymbol{\delta})}{\tau_r} \right),$$

where $1 = \tau_1 < \dots < \tau_R$ are a series of temperatures. A parallel tempering iteration involves applying either (i) a sampling step (with probability ω); or (ii) a swap step (with probability $1 - \omega$). In the sampling step, the current values $\boldsymbol{\delta}_{(r)}^c$ in the r th

chain are updated to $\delta_{(r)}^{c+1}$, for $r = 1, \dots, R$, using the Gibbs sampling algorithm. In the swap step, two neighbouring chains (r and $r + 1$) are chosen at random. With probability

$$\min \left[1, \frac{U_{r+1}(\delta_{(r)}^c)U_r(\delta_{(r+1)}^c)}{U_{r+1}(\delta_{(r+1)}^c)U_r(\delta_{(r)}^c)} \right],$$

set $\delta_{(r)}^{c+1} = \delta_{(r+1)}^c$ and $\delta_{(r+1)}^{c+1} = \delta_{(r)}^c$. To complete sampling step (i) for each distribution at temperature τ_r , the full conditional and RMLMH components need to be adjusted (see [Appendices A and B](#) for details).

3.2. Amended Gibbs sampling and parallel tempering

One iteration of the RMLMH step in the Gibbs sampling algorithm outlined in Section 3.1 requires one evaluation of $\mathbf{m}(\theta)$ to calculate the acceptance probability and one evaluation each of sensitivities $\partial \mathbf{m}(\theta)/\partial \theta$ and $\partial^2 \mathbf{m}(\theta)/\partial \theta \partial \theta_k$ to form the proposal distribution. To evaluate $\mathbf{m}(\theta)$ we are required to evaluate $\mu(\theta; \mathbf{x}_i, t_{ij})$, for $i = 1, \dots, I$ and $j = 1, \dots, n_i$. The derivative terms require evaluation of $\partial \mu(\theta; \mathbf{x}_i, t_{ij})/\partial \theta$ and $\partial^2 \mu(\theta; \mathbf{x}_i, t_{ij})/\partial \theta \partial \theta_k$, for $i = 1, \dots, I, j = 1, \dots, n_i$ and $k = 1, \dots, p$.

The system of ODEs can be augmented with additional equations to numerically solve for the first- and second-order sensitivities ([Valko and Vajda, 1984](#); [Girolami and Calderhead, 2011](#)). However this will substantially increase the computational burden of the numerical solution. Note that to generate a sample of size B using parallel tempering with R chains will require BR evaluations of the three functions for each \mathbf{x}_i and t_{ij} , for $i = 1, \dots, I, j = 1, \dots, n_i$ and $k = 1, \dots, p$. Once we have generated a sample of size B from the posterior distribution, to estimate the probability $P(\mathbf{y}_0 \in \mathcal{D} | \mathbf{y}_S)$, we require B further evaluations of $\mu(\theta; \mathbf{x}, t)$, for each $\mathbf{x} \in \mathcal{X}$ and $t \in \mathcal{T}$ at which we wish to evaluate this probability.

Making such a large number of evaluations of the numerical solution is clearly computationally infeasible. Instead we form an approximation to the numerical solution by the construction of a statistical emulator through a computer experiment ([Sacks et al., 1989](#)). The numerical solution to the ODEs is evaluated at a selection of N combinations of parameters θ ; controllable variables \mathbf{x} ; and times t . These evaluations are treated as data to which a statistical model, termed an emulator, is fitted. The emulator is essentially a predictive distribution (denoted by $Q(\theta, \mathbf{x}, t)$) from which fast predictions of the numerical solution can be obtained. For example, the predictive mean, $\hat{\mu}(\theta; \mathbf{x}, t)$, will be used as a point prediction. For more details on computer experiments and statistical emulators, see, for example, [Santner et al. \(2003\)](#), [Fang et al. \(2006\)](#) and [Dean et al. \(2015, Section V\)](#).

In the Sampling Phase, we apply the Gibbs sampling and parallel tempering algorithm with the following amendments:

1. All evaluations of sensitivities $\partial \mu(\theta; \mathbf{x}, t)/\partial \theta$ and $\partial^2 \mu(\theta; \mathbf{x}, t)/\partial \theta \partial \theta_k$ are replaced by $\partial \hat{\mu}(\theta; \mathbf{x}, t)/\partial \theta$ and $\partial^2 \hat{\mu}(\theta; \mathbf{x}, t)/\partial \theta \partial \theta_k$, respectively, for $k = 1, \dots, p$.
2. $R + 1$ parallel chains are constructed for temperatures $1 = \tau_0 = \tau_1 < \tau_2 < \dots < \tau_R$. In chains $r \geq 1$, all evaluations of $\mu(\theta; \mathbf{x}, t)$ in the evaluation of $U_r(\delta)$ in the acceptance probability are replaced by evaluation of $\hat{\mu}(\theta; \mathbf{x}, t)$. For evaluation of $U_0(\delta)$ in chain $r = 0$ with temperature $\tau_0 = 1$, evaluation of the actual $\mu(\theta; \mathbf{x}, t)$ is retained.

This modification results in the sample generated in chain $r = 0$ being from the true posterior distribution of δ but the overall sampling scheme requiring only B evaluations of $\mu(\theta; \mathbf{x}_i, t_{ij})$ for $i = 1, \dots, I$ and $j = 1, \dots, n_i$.

In the Prediction Phase, to estimate the probability of satisfying the specification limits, all evaluations of $\mu(\theta; \mathbf{x}, t)$ are replaced by a value generated from $Q(\theta, \mathbf{x}, t)$. Sampling from this predictive distribution means the estimate of the probability takes account of additional uncertainty introduced by using an emulator approximation.

3.3. Multivariate Gaussian process emulator

In this paper, we employ the multivariate Gaussian process (MGP; [Conti and O'Hagan 2010](#)) emulator. The numerical solution to the ODEs, $\mu(\theta; \mathbf{x}, t)$, is evaluated at each of the elements of the set (or meta-design) $\zeta = \{(\theta_1^*, \mathbf{x}_1^*, t_1^*), \dots, (\theta_N^*, \mathbf{x}_N^*, t_N^*)\}$. The superscript * notation has been introduced to distinguish the meta-design from the design of the physical experiment. Let $\mathbf{z}_i = \mu(\theta_i^*; \mathbf{x}_i^*, t_i^*)$, \mathbf{Z} be the $N \times K$ matrix with i th row given by \mathbf{z}_i^T , and $\mathbf{w}_i = (\theta_i^*; \mathbf{x}_i^*, t_i^*)$ be the $U \times 1$ vector of all inputs, where $U = p + F + 1 = 11$. The central assumption of the MGP is that \mathbf{Z} follows a matrix normal distribution,

$$\mathbf{Z} \sim \text{MN}(\mathbf{1}_N \boldsymbol{\beta}^T, \boldsymbol{\Phi}, \mathbf{P}),$$

where $\boldsymbol{\beta}$ is a $K \times 1$ vector of common column means of \mathbf{Z} , $\boldsymbol{\Phi}$ is a $K \times K$ unstructured column covariance matrix, \mathbf{P} is an $N \times N$ row correlation matrix and $\mathbf{1}_N$ is the $N \times 1$ vector of ones. We model the ij th element of \mathbf{P} as

$$P_{ij} = \exp \left(- \sum_{u=1}^U \alpha_u (w_{iu} - w_{ju})^2 \right),$$

where w_{iu} is the u th element of \mathbf{w}_i and $\alpha_u > 0$ are correlation parameters, for $u = 1, \dots, U$.

Suppose we wish to predict $\mathbf{z}_0 = \boldsymbol{\mu}(\boldsymbol{\theta}_0; \mathbf{x}_0, t_0)$. Let

$$\hat{\boldsymbol{\beta}} = \frac{\mathbf{1}_N^T \mathbf{P}^{-1} \mathbf{Z}}{\mathbf{1}_N^T \mathbf{P}^{-1} \mathbf{1}_N}, \tag{14}$$

$$\hat{\boldsymbol{\Psi}} = \mathbf{Z}^T \mathbf{P}^{-1} \left(\mathbf{I}_N - \frac{\mathbf{1}_N \mathbf{1}_N^T \mathbf{P}^{-1}}{\mathbf{1}_N^T \mathbf{P}^{-1} \mathbf{1}_N} \right) \mathbf{Z}, \tag{15}$$

and $\mathbf{w}_0 = (\boldsymbol{\theta}_0, \mathbf{x}_0, t_0)$, with \mathbf{I}_N the $N \times N$ identity matrix.

It can be shown that the predictive distribution of \mathbf{z}_0 , i.e. $Q(\boldsymbol{\theta}, \mathbf{x}, t)$, conditional on outputs $\mathbf{z}_1, \dots, \mathbf{z}_N$ is given by

$$N(\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t), \hat{\boldsymbol{\Pi}}(\boldsymbol{\theta}; \mathbf{x}, t)), \tag{16}$$

where

$$\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t) = \hat{\boldsymbol{\beta}} + \mathbf{p}_0^T \mathbf{P}^{-1} (\mathbf{Z} - \mathbf{1}_N \hat{\boldsymbol{\beta}}^T), \tag{17}$$

$$\hat{\boldsymbol{\Pi}}(\boldsymbol{\theta}; \mathbf{x}, t) = (1 - \mathbf{p}_0^T \mathbf{P}^{-1} \mathbf{p}_0) \hat{\boldsymbol{\Psi}}, \tag{18}$$

with $N \times 1$ vector \mathbf{p}_0 having i th element

$$p_{0i} = \exp \left(- \sum_{u=1}^U \alpha_u (w_{iu} - w_{0u})^2 \right),$$

and w_{0u} being the u th element of \mathbf{w}_0 . The predictive distribution given by (16) is conditional on the correlation parameters, $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_U)$. We replace these values by their marginal posterior mode (e.g. Overstall and Woods, 2016) although a fully Bayesian approach could be taken with these parameters integrated out with respect to their posterior distribution, conditional on \mathbf{Z} .

3.3.1. Meta-design for the sampling and prediction phases

The two emulators formed prior to the Sampling and Prediction Phases depend on the choice of the meta-design ζ . The numerical solution, $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$, has three different inputs: $\boldsymbol{\theta}$, \mathbf{x} and t . We construct the meta-design as a cartesian product of designs for each of these three input types,

$$\zeta = \zeta_1 \times \zeta_2 \times \zeta_3,$$

where $\zeta_1 = \{\boldsymbol{\theta}_1^*, \dots, \boldsymbol{\theta}_{N_1}^*\}$, $\zeta_2 = \{\mathbf{x}_1^*, \dots, \mathbf{x}_{N_2}^*\}$, $\zeta_3 = \{t_1^*, \dots, t_{N_3}^*\}$, and $N = \prod_{i=1}^M N_i$ where $M = 3$ for our experiment.

Separate designs in t and $\boldsymbol{\theta}$ and \mathbf{x} allow us to exploit computational efficiencies in the computation of the numerical solution to the ODEs; computing $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t_1)$ and $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t_2)$, for $t_1 \neq t_2$, requires a negligible increase in computational effort over just computing $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t_1)$.

This structure of meta-design also allows the emulator row correlation matrix, \mathbf{P} , to be decomposed as

$$\mathbf{P} = \bigotimes_{i=1}^M \mathbf{P}_i = \mathbf{P}_1 \otimes \mathbf{P}_2 \otimes \mathbf{P}_3,$$

where the ij th elements of \mathbf{P}_1 , \mathbf{P}_2 and \mathbf{P}_3 are

$$P_{1,ij} = \exp \left(- \sum_{u=1}^p \alpha_u (\theta_{iu}^* - \theta_{ju}^*)^2 \right) \quad \text{for } i, j = 1, \dots, N_1,$$

$$P_{2,ij} = \exp \left(- \sum_{u=1}^F \alpha_{p+u} (x_{iu}^* - x_{ju}^*)^2 \right) \quad \text{for } i, j = 1, \dots, N_2,$$

$$P_{3,ij} = \exp \left(-\alpha_U (t_i^* - t_j^*)^2 \right) \quad \text{for } i, j = 1, \dots, N_3,$$

respectively. Hence,

$$\mathbf{P}^{-1} = \bigotimes_{i=1}^M \mathbf{P}_i^{-1},$$

significantly reducing the computational burden of constructing the emulator and improving numerical stability. Similarly, for a meta-design with this structure, (14), (15), (17) and (18) also simplify:

$$\hat{\boldsymbol{\beta}} = \frac{\left(\bigotimes_{i=1}^M \mathbf{1}_{N_i}^T \mathbf{P}_i^{-1} \right) \mathbf{Z}}{\prod_{i=1}^M \mathbf{1}_{N_i}^T \mathbf{P}_i^{-1} \mathbf{1}_{N_i}},$$

$$\hat{\Psi} = \mathbf{Z}^T \left[\bigotimes_{i=1}^M \mathbf{P}_i^{-1} - \bigotimes_{i=1}^M \mathbf{P}_i^{-1} \mathbf{1}_{N_i} \mathbf{1}_{N_i}^T \mathbf{P}_i^{-1} \right] \mathbf{Z},$$

$$\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t) = \hat{\boldsymbol{\beta}} + \left(\bigotimes_{i=1}^M \mathbf{p}_i^T \mathbf{P}_i^{-1} \right) (\mathbf{Z} - \mathbf{1}_N \hat{\boldsymbol{\beta}}^T),$$

$$\hat{\boldsymbol{\Pi}}(\boldsymbol{\theta}; \mathbf{x}, t) = \left(1 - \prod_{i=1}^M \mathbf{p}_i^T \mathbf{P}_i^{-1} \mathbf{p}_i \right) \hat{\Psi},$$

where vectors \mathbf{p}_1 , \mathbf{p}_2 and \mathbf{p}_3 have j th elements given by

$$p_{1j} = \exp \left(- \sum_{u=1}^p \alpha_u (\theta_{ju}^* - \theta_u)^2 \right), \quad \text{for } j = 1, \dots, N_1,$$

$$p_{2j} = \exp \left(- \sum_{u=1}^p \alpha_{p+u} (x_{ju}^* - x_u)^2 \right), \quad \text{for } j = 1, \dots, N_2,$$

$$p_{3j} = \exp \left(-\alpha_U (t_j^* - t)^2 \right), \quad \text{for } j = 1, \dots, N_3,$$

respectively.

When we choose ζ to construct the emulator for the **Sampling Phase**, we choose the values of \mathbf{x} and t to coincide with the design of the physical experiment, i.e. to be equal to $\mathbf{x}_1, \dots, \mathbf{x}_l$ and t_1, \dots, t_m . Hence uncertainty in the prediction will only result from the different choice of parameter $\boldsymbol{\theta}$. For our experiment, this means fixing $\zeta_2 = \{\mathbf{x}_1, \dots, \mathbf{x}_5\}$ and $\zeta_3 = \{t_1, \dots, t_m\}$, i.e. $N_2 = l = 5$ and $N_3 = m$.

We use the following exploratory algorithm to adaptively choose ζ_1 . This uses a hybrid of the methodologies proposed by Rasmussen (2003), Fielding et al. (2011) and Overstall and Woods (2013). Rasmussen (2003) first proposed using MCMC methods to adaptively improve an emulator to an unnormalized posterior density based on a computationally expensive function in light of observed data. Fielding et al. (2011) extended this methodology by using parallel tempering, and Overstall and Woods (2013) proposed to emulate the computationally expensive function, as opposed to the posterior density, to make model criticism more efficient.

The following *exploratory algorithm* iteratively improves the accuracy of the emulator in the region of the parameter space Θ corresponding to high posterior density.

1. Generate the set $\zeta_1^0 = \{\boldsymbol{\theta}_1^*, \dots, \boldsymbol{\theta}_{N_1^0}^*\}$ as a sample of size N_1^0 from the prior distribution of $\boldsymbol{\theta}$.
2. Let $\zeta^0 = \zeta_1^0 \times \zeta_2 \times \zeta_3$ and fit the MGP to the resulting evaluations from $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$.
3. Let $\zeta = \zeta^0$ and repeat the following steps until $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$ as been evaluated a total of N times.
 - (a) Perform L iterations of the Gibbs sampling and parallel tempering scheme (for chains $r = 1, \dots, R$) where evaluation of $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$ is replaced by evaluation of $\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t)$ in all instances.
 - (b) Evaluate $\tilde{\mathbf{z}}_{rij} = \boldsymbol{\mu}(\boldsymbol{\theta}_{(r)}; \mathbf{x}_i, t_j)$, for $r = 1, \dots, R, i = 1, \dots, l$ and $j = 1, \dots, n_i$, where $\boldsymbol{\theta}_{(r)}$ is the current value of $\boldsymbol{\theta}$ in the r th chain. Augment the matrix \mathbf{Z} by $\tilde{\mathbf{z}}_{rij}$ for $r = 1, \dots, R, i = 1, \dots, l$ and $j = 1, \dots, n_i$. Refit the MGP.

The meta-design for the **Prediction Phase** is constructed as the cartesian product of three space-filling designs. Firstly, ζ_1 is constructed as a space-filling (e.g. Johnson et al., 1990) sub-sample of size N_1 , selected using the `cover.design` function in the R package `Fields` (Nychka et al., 2015) using the sample generated from the posterior distribution of $\boldsymbol{\theta}$ in the Sampling Phase as a candidate list. Secondly, ζ_2 and ζ_3 are chosen as Latin hypercube designs in \mathcal{X} and \mathcal{T} of sizes N_2 and N_3 , respectively (Santner et al., 2003, Ch. 5).

4. Results

4.1. Sampling phase

To create the meta-design for the Sampling Phase, we set $N_1^0 = 50, N_1 = 100$ and set up $R = 5$ parallel chains. This results in ten iterations of the exploratory algorithm in Section 3.3.1 with $L = 50$. We adapt the temperatures of the parallel chains using the methodology of Miasojedow et al. (2013). Fig. 1 shows the resulting values of $\boldsymbol{\theta}^*$ in meta-design ζ_1 . It is clear that the exploratory algorithm selects points from a concentrated region of Θ .

We then obtain a sample of size $B = 50\,000$ from the posterior distribution of $\boldsymbol{\delta}$ using the amended Gibbs sampling and parallel tempering algorithm presented in Section 3.2. Convergence was assessed informally via trace plots (not shown) which showed that the chains had mixed adequately and each posterior was unimodal. Fig. 2 shows plots of the prior and estimated posterior densities for each element of $\boldsymbol{\theta}$ as well as the elements of $\boldsymbol{\psi}$ and ρ . Clearly the data has led to an increase in information from the prior to posterior distributions.

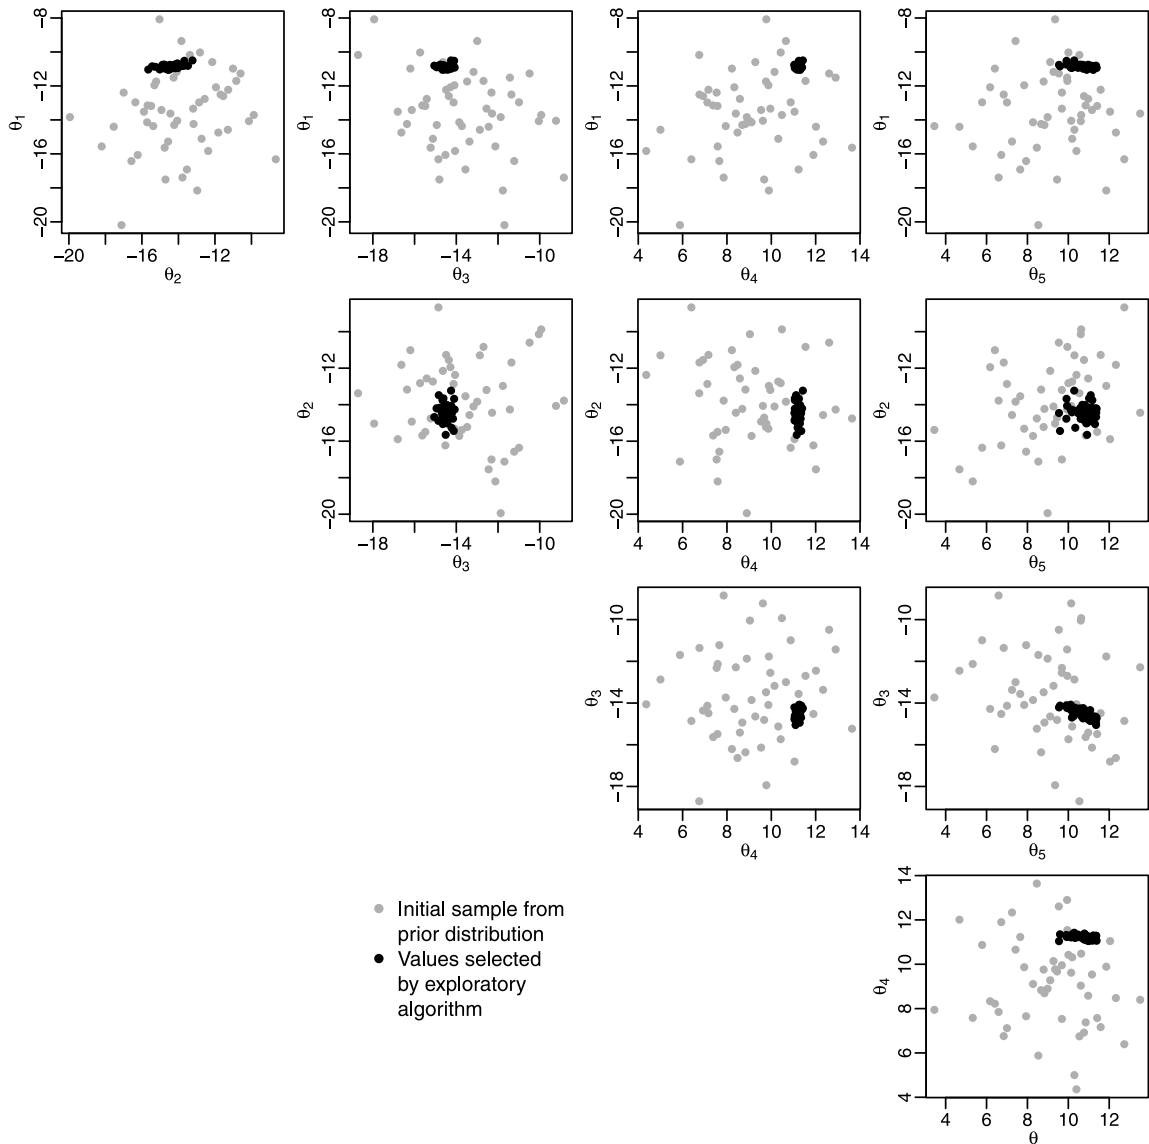


Fig. 1. Values of θ in the meta-design ζ_1 for the Sampling Phase as found by the exploratory algorithm in Section 3.3.1: grey points are the N_1^0 initial values generated from the prior distribution and black points are the values selected in step 3 of the exploratory algorithm.

Diagnostics for assessing the adequacy of the fitted model (see, for example, Gelman et al., 2014, Chapter 6) indicated that there were no reasons to believe that the fit was inadequate.

4.2. Prediction phase

We now use the sample generated from the posterior distribution of δ , and, in particular, θ , to construct an emulator $Q(\theta, \mathbf{x}, t)$, as defined in (16), for the numerical solution to the ODEs. We let $N_1 = 20$, $N_2 = 50$ and $N_3 = 50$. It could be noted that this value of N_1 is quite small compared to the rule of thumb of Loepky et al. (2009) that the sample size should be approximately ten times the number of input dimensions (in this case $p = 5$). However, we found that, due to the concentration of the posterior, this value was sufficiently large to produce an adequate emulator and avoid numerical instability in the inversion of \mathbf{P}_1 . To assess the adequacy of the emulator, we created a test meta-design of the same size (as above, a space filling sub-sample with a candidate list excluding the points from the original meta-design) and implemented the diagnostic methods of Overstall and Woods (2016). In particular, the root mean squared errors between the numerical solution and the predictive mean were 3.4×10^{-2} , 0.19 and 0.40 for each of the $K = 3$ dimensions of $\mu(\theta; \mathbf{x}, t)$, respectively. The overall coverage of the 95% predictive intervals, as assessed on the test design, was 89.0%.

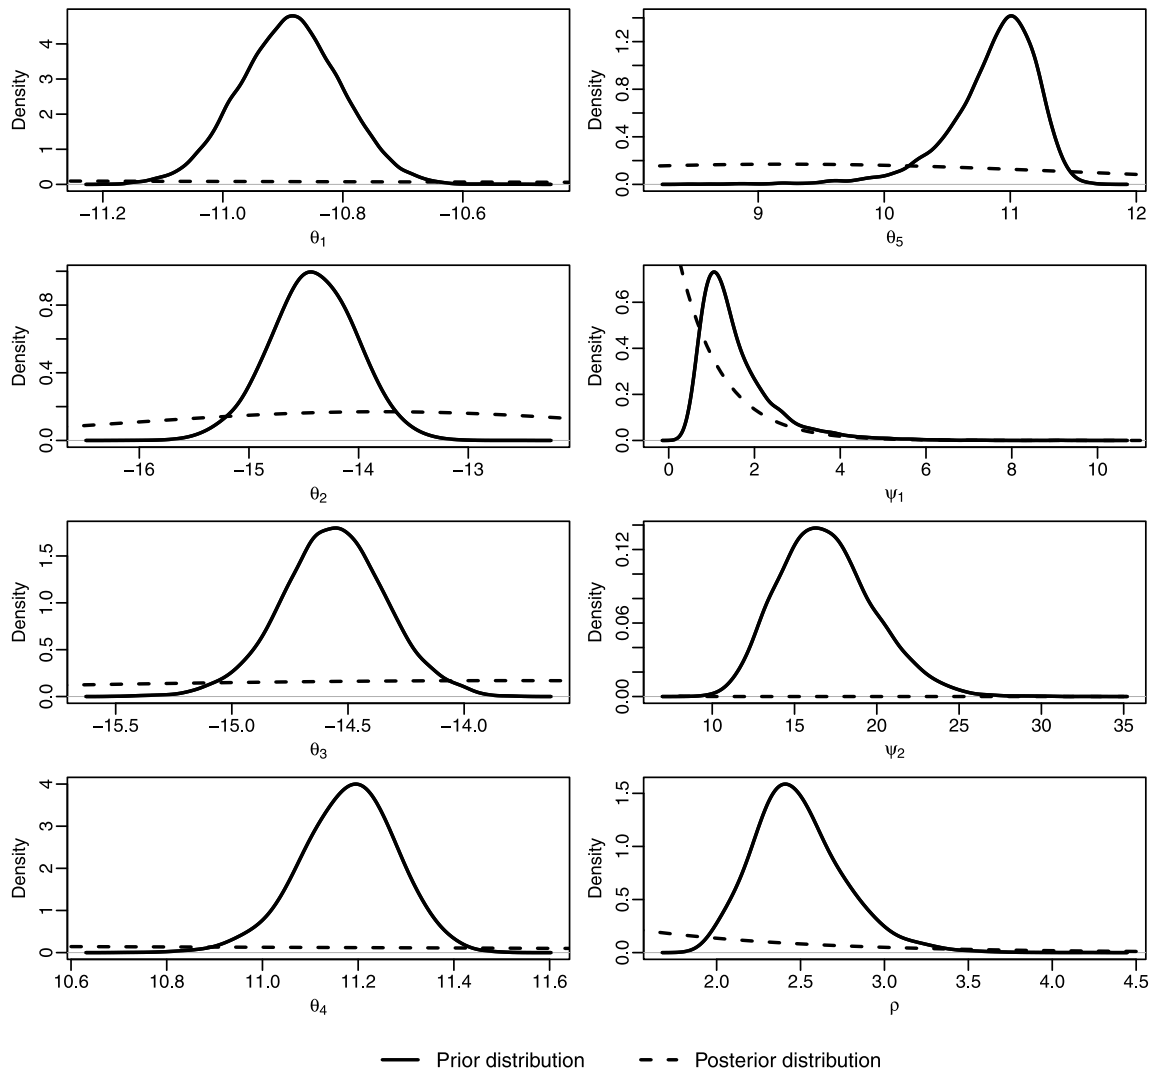


Fig. 2. Estimated posterior and prior densities for each element of θ , ψ_1 , ψ_2 and ρ .

Table 3

Optimum values of the controllable factors, \mathbf{x} and t , as found by ACE.

Name	Symbol		Value	Unit
Initial amount of A	A_0	x_1	30.52	mols
Initial amount of D	D_0	x_2	91.51	mols
Initial amount of E	E_0	x_3	26.47	mols
Temperature	λ	x_4	313.15	K
Volume	V	x_5	31.28	litres
Time	t		199.10	s

Now we can use emulator $Q(\theta, \mathbf{x}, t)$ to approximate the probability, $P(\mathbf{y}_0 \in \mathcal{Y} | \mathbf{y}_S)$, of satisfying the specification limits for given \mathbf{x} and t . We need to maximize this probability over the space $S = \mathcal{X} \times \mathcal{T}$. To do this we use the approximate coordinate exchange (ACE; Overstall and Woods, 2017) algorithm. ACE was originally developed for finding Bayesian optimal experimental designs where the objective function is maximized over a design space. In these situations, the objective function is usually analytically intractable. ACE uses a cyclic ascent algorithm (e.g. Lange, 2013, Chapter 7) to maximize the objective function via a univariate Gaussian process emulator based on evaluations of a Monte Carlo approximation to the objective function. We use the implementation of ACE given by the R package *acebayes* (Overstall et al., 2017). We use 100 random starts of ACE, where the initial value of (\mathbf{x}, t) is uniformly generated in $S \times \mathcal{T}$. The optimum values of \mathbf{x} and t , as found by ACE, are shown in Table 3. The approximate maximum value of $P(\mathbf{y}_0 \in \mathcal{Y} | \mathbf{y}_S)$ attained was 0.77 (subject to Monte




Fig. 3. The probability of satisfying each of the three marginal constraints and the joint constraints for time (in s; x-axis) against initial amount of A (in mols; y-axis) for five different initial amounts of D. The colours indicate the magnitude of the probability, with black being 0 and white being 1.

Carlo error). We investigated the sensitivity of the above approach to the choice of prior distributions for ρ , ψ , Ω and Σ . We did this by repeating the above analysis to find the optimal settings for \mathbf{x} and t under different prior specifications. We found the results to be broadly robust to choice of prior distribution and details of this are provided in [Appendix D](#).

To further explore how $P(\mathbf{y}_0 \in \mathcal{Y} | \mathbf{y}_s)$ depends on \mathbf{x} and t , we fix the initial amount of E, temperature and volume at 26.47 mols, 31.28 litres, and 313.5 K, respectively, i.e. the optimum values as found by ACE. We then vary the initial amounts of A and D, and time over the ranges identified in [Table 1](#). The last row of [Fig. 3](#) shows $P(\mathbf{y}_0 \in \mathcal{Y} | \mathbf{y}_s)$ for time against the initial amount of A, with different ranges for different initial amounts of D. The first three rows show corresponding plots for the marginal probability of satisfying each of the specification limits (3) on E, F and H, respectively. The trade-off between the objectives given by the specification limits can be clearly seen. To satisfy the limit $E(t) < 3$ mols, t has to be large. This is intuitively obvious since the initial amount of E is 26.47 mols so the process will need to progress for some time before the amount of E has decreased enough to satisfy the limit. The opposite is true for H (initial amount of 0 mols), where the optimum time to satisfy $H(t) < 3$ mols is for t to be small. This means that there is a very narrow window of values of t

that will satisfy all three constraints, as shown in the last row of Fig. 3. Note that the optimum values of \mathbf{x} and t (as shown in Table 3) lie within this narrow window.

5. Discussion

This paper has developed and applied methodology for choosing combinations of values of the controllable factors that produce a high probability of meeting process specification when the process is at least approximately determined by a physical model. Our application was a chemical synthesis process used to produce a pharmaceutical product when specification was determined by (a) the amount of the substance constituting the pharmaceutical product being above some level; and (b) the amounts of two unwanted by-product substances being below some levels. The relationship between the controllable factors and the amounts of substances of interest is hypothesized to be governed by the analytically intractable solution to a system of ODEs. Responses from a physical experiment are used to refine the knowledge on this relationship and the probability of satisfying the constraints is maximized. Statistical emulators for the numerical solution are constructed to accelerate the model-fitting process and the estimation of the probability of satisfying the constraints. The methodology produces computationally feasible results that take account of different uncertainties involved in the experimental and modelling processes (measurement error, model inadequacy, emulator error). The methods have the potential to be applied to a variety of quality control and pharmaceutical design space problems that involve the numerically expensive approximation to physical models.

The modelling and computational approach taken in this paper can be modified in a variety of different ways to suit the application of interest. For example, we have explicitly taken account of the model discrepancy. If the experimenters believed the model given by the solution to the ODEs was a very close approximation to the true physical process then the model discrepancy could be discarded by setting $\mathbf{D} = \mathbf{0}$ and omitting the steps for sampling from the full conditional distributions of Σ and ψ . Another example of modification is the introduction of noise into the specification of the controllable factors. In this paper, we assumed that the chemical engineers have complete control over the specification of the controllable factors when we perform the maximization of the probability of satisfying the specification limits given by (3). In some cases, the chemical engineer would not be able to control these exactly. Following the approach in this paper, it would be straightforward to introduce variability as an intermediate step between specifying the controllable variables and evaluating the Gaussian process emulator. Lastly, we employed the RMLMH algorithm to sample from the full conditional distributions of θ , ρ and ψ . It would also be possible to use the methodology in this paper to employ Riemann Manifold Hamiltonian Monte Carlo methodology (Girolami and Calderhead, 2011).

Acknowledgements

The authors would like to thank Drs John Peterson, Mohammad Yahyah and Neil Hodnett (GlaxoSmithKline) for providing the original kinetics modelling problem and data set. Additionally, the authors are grateful to the editor, associate editor and two referees who provided valuable feedback on earlier versions of the paper. DCW was supported by UK Engineering and Physical Sciences Research Council (EPSRC) Fellowship EP/J018317/1 and KJM was supported by a PhD award from EPSRC, UK and GlaxoSmithKline, UK.

Appendix A. Full-conditional distributions for Gibbs sampling

A.1. Full conditional distribution for model discrepancy

For the r th chain with temperature τ_r , the full-conditional distribution of \mathbf{d}^* is

$$\mathbf{d}^* | \mathbf{y}, \theta, \Sigma, \Omega, \rho, \psi \sim N(\boldsymbol{\mu}_{\mathbf{d}^*}, \tau_r \mathbf{V}_{\mathbf{d}^*}),$$

where

$$\begin{aligned} \mathbf{V}_{\mathbf{d}^*} &= (\mathbf{H}^T (\Omega^{-1} \otimes \mathbf{T}^{-1}) \mathbf{H} + (\Sigma^{-1} \otimes \mathbf{S}^{-1}))^{-1}, \\ \boldsymbol{\mu}_{\mathbf{d}^*} &= \mathbf{V}_{\mathbf{d}^*} (\mathbf{H}^T (\Omega^{-1} \otimes \mathbf{T}^{-1}) \mathbf{y} + (\Sigma^{-1} \otimes \mathbf{S}^{-1}) \mathbf{m}(\theta)). \end{aligned}$$

A.2. Full conditional distributions for covariance matrices

For the r th chain with temperature τ_r , the full-conditional distributions of Ω and Σ are given by

$$\Omega | \mathbf{Y}, \mathbf{D}^*, \rho \sim \text{IW} \left(\frac{\nu + n + K + 1}{\tau_r} - K - 1, \frac{1}{\tau_r} \left(\mathbf{I}_K + (\mathbf{Y} - \mathbf{G}\mathbf{D}^*)^T \mathbf{T}^{-1} (\mathbf{Y} - \mathbf{G}\mathbf{D}^*) \right) \right),$$

and

$$\Sigma | \mathbf{D}^*, \theta, \psi \sim \text{IW} \left(\frac{\nu + m + K + 1}{\tau_r} - K - 1, \frac{1}{\tau_r} \left(\mathbf{I}_K + (\mathbf{D}^* - \mathbf{M}(\theta))^T \mathbf{S}^{-1} (\mathbf{D}^* - \mathbf{M}(\theta)) \right) \right),$$

respectively, where IW denotes the inverse-Wishart distribution.

A.3. Full conditional distribution for censored observations

Let \mathbf{A} be the $n \times n$ permutation matrix which re-orders the elements of \mathbf{y} such that we can write

$$\mathbf{A}\mathbf{y} = \begin{pmatrix} \mathbf{y}_S \\ \mathbf{y}_C \end{pmatrix} \sim N(\mathbf{A}\mathbf{H}\mathbf{d}^*, \mathbf{A}(\boldsymbol{\Omega} \otimes \mathbf{T})\mathbf{A}^T).$$

Define $\mathbf{b} = \mathbf{A}\mathbf{H}\mathbf{d}^*$ and $\mathbf{L} = \mathbf{A}(\boldsymbol{\Omega} \otimes \mathbf{T})\mathbf{A}^T$. Let $\mathbf{b} = (\mathbf{b}_S, \mathbf{b}_C)^T$ where \mathbf{b}_S and \mathbf{b}_C are the $n_S \times 1$ and $n_C \times 1$ sub-vectors corresponding to \mathbf{y}_S and \mathbf{y}_C , respectively. Similarly, partition \mathbf{L} as

$$\mathbf{L} = \begin{pmatrix} \mathbf{L}_{SS} & \mathbf{L}_{SC} \\ \mathbf{L}_{CS} & \mathbf{L}_{CC} \end{pmatrix}.$$

The full conditional distribution of \mathbf{y}_C is

$$\mathbf{y}_C \sim N(\mathbf{b}_C + \mathbf{L}_{CS}\mathbf{L}_{SS}^{-1}(\mathbf{y}_S - \mathbf{b}_S), \tau_r(\mathbf{L}_{CC} - \mathbf{L}_{CS}\mathbf{L}_{SS}^{-1}\mathbf{L}_{SC}))$$

truncated to the hypercube $[-\infty, \log \chi]^{n_C}$.

Appendix B. Riemann manifold Langevin Metropolis–Hastings (RMLMH) algorithm

Consider generating a sample from the distribution of $\delta \in \mathbb{R}^p$, conditional on \mathbf{y} , having density $\pi(\delta|\mathbf{y}) \propto (\pi(\mathbf{y}|\delta)\pi(\delta))^{1/\tau_r}$, for temperature τ_r , using the Metropolis–Hastings algorithm. Let $h(\delta) = \log \pi(\mathbf{y}|\delta) + \log \pi(\delta)$ and define

$$\begin{aligned} \nabla h(\delta) &= \frac{\partial \log \pi(\mathbf{y}|\delta)}{\partial \delta} + \frac{\partial \log \pi(\delta)}{\partial \delta}, \\ \mathbf{G}(\delta) &= -E_{\mathbf{y}|\delta} \left[\frac{\partial \log \pi(\mathbf{y}|\delta)}{\partial \delta} \frac{\partial \log \pi(\mathbf{y}|\delta)}{\partial \delta^T} \right] - \frac{\partial^2 \log \pi(\delta)}{\partial \delta \partial \delta^T}. \end{aligned}$$

If δ^c is the current value of δ , then under the RMLMH algorithm, the proposal distribution is $N(\boldsymbol{\mu}_\delta, \tau_r \mathbf{V}_\delta)$, where

$$\begin{aligned} \boldsymbol{\mu}_\delta &= \delta^c + \frac{\epsilon^2}{2} \mathbf{G}(\delta^c)^{-1} \nabla h(\delta^c) - \epsilon^2 \mathbf{n}^{(1)} + \frac{\epsilon^2}{2} \mathbf{n}^{(2)}, \\ \mathbf{V}_\delta &= \epsilon^2 \mathbf{G}(\delta^c)^{-1}, \end{aligned}$$

with $\mathbf{n}^{(1)}$ and $\mathbf{n}^{(2)}$ having i th elements

$$\begin{aligned} n_i^{(1)} &= \sum_{j=1}^p \left\{ \mathbf{G}(\delta^c)^{-1} \frac{\partial \mathbf{G}(\delta)}{\partial \delta_j} \Big|_{\delta=\delta^c} \mathbf{G}(\delta^c)^{-1} \right\}_{ij}, \\ n_i^{(2)} &= \sum_{j=1}^p \left\{ \mathbf{G}(\delta^c)^{-1} \right\}_{ij} \text{tr} \left\{ \mathbf{G}(\delta^c)^{-1} \frac{\partial \mathbf{G}(\delta)}{\partial \delta_j} \Big|_{\delta=\delta^c} \right\}, \end{aligned}$$

respectively.

Appendix C. Components required for Riemann manifold MH algorithm

C.1. Physical parameters

The log density of the full conditional distribution of $\boldsymbol{\theta}$ is given by

$$h(\boldsymbol{\theta}) \propto -\frac{1}{2} (\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta}))^T (\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}) (\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})) - \frac{1}{2} (\boldsymbol{\theta} - \boldsymbol{\mu})^T \boldsymbol{\Delta}^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}).$$

The derivative with respect to $\boldsymbol{\theta}$ is

$$\nabla h(\boldsymbol{\theta}) = (\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta}))^T (\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}) \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} - \boldsymbol{\Delta}^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}).$$

The matrix tensor for $\boldsymbol{\theta}$ is

$$\mathbf{G}(\boldsymbol{\theta}) = \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T} (\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}) \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} + \boldsymbol{\Delta}^{-1},$$

with derivatives

$$\frac{\partial \mathbf{G}(\boldsymbol{\theta})}{\partial \theta_k} = 2 \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T} (\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}) \frac{\partial^2 \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta} \partial \theta_k},$$

for $k = 1, \dots, p$.

C.2. Correlation parameter for time dependency

Consider a log transformation of ρ , i.e. $a = \log \rho$. The log density of the full conditional distribution of a is given by

$$h(a) \propto -\frac{K}{2} \sum_{i=1}^{I'} \log |\mathbf{T}_i| - \frac{1}{2} (\mathbf{y} - \mathbf{Hd}^*)^\top (\boldsymbol{\Omega}^{-1} \otimes \text{diag}_{i=1, \dots, I'} \{\mathbf{T}_i^{-1}\}) (\mathbf{y} - \mathbf{Hd}^*) + a - \exp(a).$$

The derivative with respect to a is

$$\frac{\partial h(a)}{\partial a} = -\frac{K}{2} \sum_{i=1}^{I'} \text{tr} \{\mathbf{T}_i^{-1} \mathbf{T}_{ia}\} + \frac{1}{2} (\mathbf{y} - \mathbf{Hd}^*)^\top (\boldsymbol{\Omega}^{-1} \otimes \text{diag}_{i=1, \dots, I'} \{\mathbf{T}_i^{-1} \mathbf{T}_{ia} \mathbf{T}_i^{-1}\}) (\mathbf{y} - \mathbf{Hd}^*) + 1 - \exp(a),$$

where $\mathbf{T}_{ia} = \partial \mathbf{T}_i / \partial a$ with jl th element $-\exp(a)(t_{ij} - t_{il})^2 T_{i,jl}$. The tensor matrix for a is

$$\mathbf{G}(a) = \frac{K}{2} \text{tr} \{\mathbf{T}_i^{-1} \mathbf{T}_{ia} \mathbf{T}_i^{-1} \mathbf{T}_{ia}\} - \exp(a),$$

with derivative

$$\frac{\partial \mathbf{G}(a)}{\partial a} = K \text{tr} \{\mathbf{T}_i^{-1} \mathbf{T}_{ia} \mathbf{T}_i^{-1} (\mathbf{T}_{iaa} - \mathbf{T}_{ia} \mathbf{T}_i^{-1} \mathbf{T}_{ia})\} - \exp(a),$$

where $\mathbf{T}_{iaa} = \partial^2 \mathbf{T}_i / \partial a^2$ with jl th element $\exp(a)(t_{ij} - t_{il})^2 (\exp(a)(t_{ij} - t_{il})^2 - 1) T_{i,jl}$.

C.3. Correlation parameters for model discrepancy

Consider a log transformation of each element of $\boldsymbol{\psi}$, i.e. $b_i = \log \psi_i$ for $i = 1, 2$. The log density of the full conditional distribution of $\mathbf{b} = (b_1, b_2)$ is given by

$$h(\mathbf{b}) \propto -\frac{K}{2} \log |\mathbf{S}| - \frac{1}{2} (\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta}))^\top (\boldsymbol{\Sigma} \otimes \mathbf{S}^{-1}) (\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})) + \sum_{i=1}^2 b_i - \exp(b_i).$$

The derivative with respect to b_i is

$$\frac{\partial h(\mathbf{b})}{\partial b_i} = -\frac{K}{2} \text{tr} \{\mathbf{S}^{-1} \mathbf{S}_i\} + \frac{1}{2} (\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta}))^\top (\boldsymbol{\Sigma} \otimes \mathbf{S}^{-1} \mathbf{S}_i \mathbf{S}^{-1}) (\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})) + 1 - \exp(b_i),$$

where $\mathbf{S}_i = \partial \mathbf{S} / \partial b_i$ with jl th element given by

$$\begin{cases} -\exp(b_1) \left[\sum_{f=1}^F (d_{jf} - d_{lf})^2 \right] S_{jl}, & \text{if } i = 1, \\ -\exp(b_2)(t_j - t_l)^2 S_{jl}, & \text{if } i = 2. \end{cases}$$

The tensor matrix, $\mathbf{G}(\mathbf{b})$ has ij th element

$$\mathbf{G}(\mathbf{b})_{ij} = \frac{K}{2} \text{tr} \{\mathbf{S}^{-1} \mathbf{S}_i \mathbf{S}^{-1} \mathbf{S}_j\} - I(i = j) \exp(b_i).$$

The derivatives of $\mathbf{G}(\mathbf{b})$ are

$$\frac{\partial \mathbf{G}(\mathbf{b})_{ij}}{\partial b_k} = \frac{K}{2} \text{tr} \{\mathbf{S}^{-1} \mathbf{S}_i \mathbf{S}^{-1} \mathbf{S}_{jk} - \mathbf{S}^{-1} \mathbf{S}_i \mathbf{S}^{-1} \mathbf{S}_k \mathbf{S}^{-1} \mathbf{S}_j + \mathbf{S}^{-1} \mathbf{S}_j \mathbf{S}^{-1} \mathbf{S}_{ik} - \mathbf{S}^{-1} \mathbf{S}_i \mathbf{S}^{-1} \mathbf{S}_j \mathbf{S}^{-1} \mathbf{S}_k\} - I(i = j = k) \exp(b_i),$$

where $\mathbf{S}_{ik} = \partial^2 \mathbf{S} / \partial b_i \partial b_k$ with jl th element given by

$$\begin{cases} \exp(b_1) \left[\sum_{f=1}^F (d_{jf} - d_{lf})^2 \right] \left(\exp(b_1) \left[\sum_{f=1}^F (d_{jf} - d_{lf})^2 \right] - 1 \right) S_{jl}, & \text{if } i = 1 \text{ and } k = 1, \\ \exp(b_2)(t_j - t_l)^2 (\exp(b_2)(t_j - t_l)^2 - 1) S_{jl}, & \text{if } i = 2 \text{ and } k = 2, \\ \exp(b_1 + b_2) \left[\sum_{f=1}^F (d_{jf} - d_{lf})^2 \right] (t_j - t_l)^2 S_{jl}, & \text{if } i = 1 \text{ and } k = 2. \end{cases}$$

Table 4

For Priors 0, 1, 2 and 3: (a) the values of α , β and ν specifying the prior distributions for ψ , Σ and Ω ; (b) the optimum values of \mathbf{x} and t for maximizing the probability of satisfying the specification limits; and (c) the probability of satisfying the specification limits given by (3) under each prior distribution for the optimum values also found under each prior distribution.

Prior	0	1	2	3
α	1	0.500	0.250	0.125
β	1	0.500	0.250	0.125
ν	4	2	1	0.5
Optimum values for controllable factors and time				
Initial amount of A	30.52	22.50	39.58	37.75
Initial amount of D	91.51	91.49	91.59	91.59
Initial amount of E	26.47	26.47	26.47	26.47
Temperature	313.15	313.15	313.15	313.15
Volume	31.28	31.28	31.28	31.28
Time	199.10	453.36	169.59	169.59
Probability of satisfying specification limits				
Prior 0	0.77	0.75	0.78	0.73
Prior 1	0.74	0.75	0.75	0.73
Prior 2	0.76	0.74	0.78	0.73
Prior 3	0.76	0.74	0.77	0.76

Appendix D. Prior sensitivity analysis

Suppose

$$\begin{aligned} \rho &\sim \text{Gamma}(\alpha, \beta) & \Sigma &\sim \text{IW}(\nu, \mathbf{I}_3) \\ \psi_i &\sim \text{Gamma}(\alpha, \beta) & \Omega &\sim \text{IW}(\nu, \mathbf{I}_3) \end{aligned}$$

for $i = 1, 2$, where the Gamma distribution is parameterized such that the expectation and variance are α/β and α/β^2 , respectively, and the inverse-Wishart such that the mode is \mathbf{I}_3/ν . Let the prior distribution for ρ , ψ , Σ and Ω outlined in Section 2.3 be denoted as Prior 0 where $\alpha = \beta = 1$ and $\nu = 4$. We consider three different prior specifications (denoted Prior 1, 2 and 3) given by different combinations of α , β and ν , corresponding to a sequence of increasingly diffuse prior distributions. For each prior, the analysis described in Section 4 was undertaken and the optimum values of \mathbf{x} and t were found by ACE. Table 4 shows the values of α , β and ν for Priors 1, 2 and 3 as well as the optimum values of \mathbf{x} and t . Also included is the probability of satisfying the specification limits given by (3) under each prior distribution for the optimum values also found under each prior distribution. There is some variability with respect to initial amounts of A and D and, particularly, with respect to time. However the probability of satisfying the specification limits remain broadly insensitive to these changes. We conclude that the results are robust to choice of prior distribution.

References

- Brown, P., Byrne, G., Hindmarsh, A., 1989. Vode: A variable coefficient ode solver. *SIAM J. Sci. Stat. Comput.* 10, 1038–1051.
- Chiao, C., Hamada, M., 2001. Analyzing experiments with correlated multiple responses. *J. Qual. Technol.* 33, 451–465.
- Conti, S., O'Hagan, A., 2010. Bayesian emulation of complex multi-output and dynamic computer models. *J. Statist. Plann. Inference* 140, 640–651.
- Dawid, A., 1981. Some matrix-variate distribution theory: Notational considerations and a bayesian application. *Biometrika* 68, 265–274.
- Dean, A., Morris, M., J., S., Bingham, D.E. (Eds.), 2015. *Handbook of Design and Analysis of Experiments*. CRC Press, Boca Raton.
- Del Castillo, E., 2007. *Process Optimization – A Statistical Approach*. Springer, New York.
- am Ende, D.J. (Ed.), 2011. *Chemical Engineering in the Pharmaceutical Industry*. Wiley, Hoboken, NJ.
- Fang, K., Li, R., Sudianto, A., 2006. *Design and Modelling for Computer Experiments*. Chapman and Hall, Boca Raton.
- Fielding, M., Nott, D.J., Liong, S.Y., 2011. Efficient MCMC schemes for computationally expensive posterior distributions. *Technometrics* 53, 16–28.
- Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B., 2014. *Bayesian Data Analysis*, volume third ed. Chapman and Hall.
- Geweke, J.F., 1991. Efficient simulation from the multivariate normal and student-t distributions subject to linear constraints. In: *Computer Science and Statistics. Proceedings of the 23rd Symposium on the Interface*.
- Girolami, M., Calderhead, B., 2011. Riemann manifold langevin and hamiltonian monte carlo methods (with discussion). *J. Roy. Statist. Soc. B* 73, 123–214.
- Iserles, A., 2009. *A First Course in the Numerical Analysis of Differential Equations*, second ed. Cambridge University Press.
- Johnson, M.E., Moore, L.M., Ylvisaker, D., 1990. Minimax and maximin distance designs. *J. Statist. Plann. Inference* 26, 131–148.
- Kennedy, M.C., O'Hagan, A., 2001. Bayesian calibration of computer models (with discussion). *J. R. Stat. Soc. Ser. B Stat. Methodol.* 63, 425–464.
- Laidler, K.J., 1984. The development of the arrhenius equation. *J. Chem. Educ.* 61, 494.
- Lange, K., 2013. *Optimization*, second ed. Springer, New York.
- Lebrun, P., Boulanger, B., Debrus, B., Lambert, P., Hubert, P., 2013. A Bayesian design space for analytical methods based on multivariate models and predictions. *J. Biopharm. Statist.* 23, 1330–1351.
- Loepky, J.L., Sacks, J., Welch, W.J., 2009. Choosing the sample size of a computer experiment: a practical guide. *Technometrics* 51, 366–376.
- Miasojedow, B., Moulines, E., Vihola, M., 2013. An adaptive parallel tempering algorithm. *J. Comput. Graph. Statist.* 22, 649–664.
- Nychka, D., Furrer, R., Paige, J., Sain, S., 2015. *fields: Tools for spatial data*. URL: www.image.ucar.edu/fields. R package version 8.15.
- Overstall, A.M., Woods, D.C., 2013. A strategy for bayesian inference for computationally expensive models with application to the estimation of stem cell properties. *Biometrics* 69, 458–468.

- Overstall, A.M., Woods, D.C., 2016. Multivariate emulation of computer simulators: model selection and diagnostics with application to a humanitarian relief model. *J. Roy. Statist. Soc. C* 65, 483–505.
- Overstall, A.M., Woods, D.C., 2017. Bayesian design of experiments using approximate coordinate exchange. *Technometrics* 59, 458–470.
- Overstall, A.M., Woods, D.C., Adamou, M., 2017. acebayes: An R package for Bayesian optimal design of experiments via approximate coordinate exchange. arXiv:1705.08096.
- Papaspiliopoulos, O., Roberts, G.O., Skold, M., 2003. Non-centered parameterizations for hierarchical models and data augmentation. In: Bernardo, J.M., Bayarri, M.J., Berger, J.O., Dawid, A.P., Heckerman, D., Smith, A.F.M., West, M. (Eds.), *Bayesian Statistics*, Vol. 7. Oxford.
- Peterson, J., 2004. A posterior predictive approach to multiple response surface optimization. *J. Qual. Technol.* 36, 139–153.
- Peterson, P.J., 2008. A Bayesian approach to the ICH Q8 definition of design space. *J. Biopharm. Statist.* 18, 959–975.
- Peterson, P.J., Yahyah, M., 2009. A Bayesian design space approach to robustness and system suitability for pharmaceutical assays and other processes. *Stat. Biopharm. Res.* 1, 441–449.
- R Core Team, 2018. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>.
- Rasmussen, C.E., 2003. Gaussian processes to speed up hybrid monte carlo for expensive bayesian integrals. In: Bernardo, J.M., Bayarri, M.J., Berger, J.O., Dawid, A.P., Heckerman, D., Smith, A.F.M., West, M. (Eds.), *Bayesian Statistics*, Vol. 7. Oxford.
- Sacks, J., Welch, W., Mitchell, T., Wynn, H., 1989. Design and analysis of computer experiments. *Statist. Sci.* 4, 409–435.
- Santner, T.J., Williams, B.J., Notz, W.I., 2003. *The Design and Analysis of Computer Experiments*. Springer, New York.
- Soetaert, K., Petzoldt, T., Woodrow Setzer, R., 2010. Solving differential equations in R: Package deSolve. *J. Stat. Softw.* 33, 1–25, URL <http://www.jstatsoft.org/v33/i09>.
- Valko, P., Vajda, S., 1984. An extended ODE solver for sensitivity calculations. *Comput. Chem.* 8, 255–271.