

Bayesian optimal design for ordinary differential equation models with application in biological science

Antony M. Overstall^{*†}, David C. Woods[†] and Ben M. Parker[‡]

[†]Southampton Statistical Sciences Research Institute,
University of Southampton, United Kingdom

[‡]School of Computing and Engineering,
University of West London, United Kingdom

Abstract

Bayesian optimal design is considered for experiments where the response distribution depends on the solution to a system of non-linear ordinary differential equations. The motivation is an experiment to estimate parameters in the equations governing the transport of amino acids through cell membranes in human placentas. Decision-theoretic Bayesian design of experiments for such nonlinear models is conceptually very attractive, allowing the formal incorporation of prior knowledge to overcome the parameter dependence of frequentist design and being less reliant on asymptotic approximations. However, the necessary approximation and maximization of the, typically analytically intractable, expected utility results in a computationally challenging problem. These issues are further exacerbated if the solution to the differential equations is not available in closed-form. This paper proposes a new combination of a probabilistic solution to the equations embedded within a Monte Carlo approximation to the expected utility with cyclic descent of a smooth approximation to find the optimal design. A novel precomputation algorithm reduces the computational burden, making the search for an optimal design feasible for bigger problems. The methods are demonstrated by finding new designs for a number of common models derived from differential equations, and by providing optimal designs for the placenta experiment.

Keywords: Approximate coordinate exchange algorithm; decision-theoretic design; Gaussian process emulation; nonlinear design.

^{*}CONTACT Antony M. Overstall; A.M.Overstall@southampton.ac.uk; Southampton Statistical Sciences Research Institute, University of Southampton, Southampton, SO17 1BJ, UK.

1 Introduction

The dynamics behind a complex physical process can often be described by a set of non-linear ordinary differential equations, where the solution to these equations represents the evolution of system states with respect to time. It is common for the system of equations to depend on some unknown physical properties (parameters) of the process in question and, potentially, on some additional controllable variables. In this paper, new methods are presented for designing experiments for the estimation of statistical models built on the solution to such a system of equations; that is, choosing the most informative combinations of time points and values of the controllable (design) variables at which observations of the physical process should be made. A decision-theoretic approach is adopted, and hence the quality of a design is measured via the expectation of a utility function chosen to encapsulate the aims of the experiment.

We assume equations with s system states $\mathbf{u}(t; \mathbf{x}, \boldsymbol{\theta}) = [u_1(t; \mathbf{x}, \boldsymbol{\theta}), \dots, u_s(t; \mathbf{x}, \boldsymbol{\theta})]^T$ modeled as a function of time t and v design variables with values held in the treatment vector $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^v$. The p -vector $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^p$ holds the physical parameters requiring estimation. For notational simplicity, the dependence of the system states on \mathbf{x} and $\boldsymbol{\theta}$ is usually suppressed, with $\mathbf{u}(t) = \mathbf{u}(t; \mathbf{x}, \boldsymbol{\theta})$, unless multiple treatments or parameter vectors are being considered. We mostly find designs for initial value problems, with $\mathbf{u}(t)$ defined via equations

$$\begin{cases} \dot{\mathbf{u}}(t) &= \mathbf{f}(\mathbf{u}(t), t, \mathbf{x}; \boldsymbol{\theta}) & \text{for } t \in \mathcal{T} = [T_0, T_1]; 0 \leq T_0 < T_1, \\ \mathbf{u}(T_0) &= \mathbf{u}_0, \end{cases} \quad (1)$$

where $\dot{\mathbf{u}}(t)$ is the gradient vector of $\mathbf{u}(t)$ with respect to time t , $\mathbf{u}_0 = (u_{01}, \dots, u_{0s})^T \in \mathbb{R}^s$ denotes initial conditions and, for given $\boldsymbol{\theta}$, $\mathbf{f} : \mathbb{R}^s \times \mathcal{T} \times \mathcal{X} \rightarrow \mathbb{R}^s$ is a continuous function satisfying the Lipschitz condition (see Iserles, 2009, p. 3). This latter assumption ensures equation (1) has a unique solution.

Our research is motivated by experiments to study the transport of serine, an amino acid, within a human placenta. Specifically, interest is in the movement of serine across a placental cell membrane (called a vesicle). In the experiments, initial amounts (μl) of both radioactive and non-radioactive serine are placed exterior and interior to the vesicle, and then the amount of radioactive serine interior to the vesicle is measured at a series of time points. The experimenters have control over initial amounts of both the interior and exterior

non-radioactive serine for each experiment, and the times (in seconds) at which observations are taken. The theoretical interior amounts of radioactive and non-radioactive serine at time t form the $s = 2$ system states, $\mathbf{u}(t) = [u_1(t), u_2(t)]^T$, with the $v = 2$ design variables, $\mathbf{x} = (x_1, x_2)^T \in [0, 1000]^2$, being, respectively, the exterior amounts of radioactive and non-radioactive serine at time $t = 0$. The equations governing the evolution of the system states are

$$\left. \begin{aligned} \dot{u}_1(t) &= \frac{x_1(u_2(t) + \theta_2\theta_4) - u_1(t)(x_2 + \theta_2\theta_3)}{u^*(\mathbf{u}(t), t, \boldsymbol{\theta}, \mathbf{x})}, \\ \dot{u}_2(t) &= \frac{x_2(u_1(t) + \theta_2\theta_4) - u_2(t)(x_1 + \theta_2\theta_3)}{u^*(\mathbf{u}(t), t, \boldsymbol{\theta}, \mathbf{x})}, \\ u_1(0) &= u_{01}, \\ u_2(0) &= u_{02}, \end{aligned} \right\} t \in [0, 600], \quad (2)$$

where

$$u^*(\mathbf{u}(t), t, \boldsymbol{\theta}, \mathbf{x}) = \frac{1}{\theta_1} \{2x_{12}^* u_{12}^*(t) + (1 + \theta_2) [\theta_4 x_{12}^* + \theta_3 u_{12}^*(t)] + 2\theta_3\theta_4\},$$

$u_{12}^*(t) = u_1(t) + u_2(t)$, $x_{12}^* = x_1 + x_2$, and initial conditions $\mathbf{u}_0 = (u_{01}, u_{02})^T \in [0, 1000]^2$ are the amounts of radioactive and non-radioactive serine interior to the vesicle at time $t = 0$. Here, the four physical parameters correspond to the maximum uptake (θ_1), the proportion of the reaction occurring through active transport (θ_2) and two reaction rates (θ_3 and θ_4). The values of these parameters are of scientific interest. See Panitchob et al. (2015) and Widdows et al. (2017) for further details of the model and experiment.

To model experimental data from a physical process governed by (1), we build a statistical model linking the physical parameters to noisy observations of the system states, or functions thereof, via an assumed data-generating process dependent on the solution to the equations (see, for example, Ramsay et al. 2007). We also assume that an experiment can be conducted where these observations are collected at various different times and, possibly, from multiple runs of the experiment with different combinations of values of the design variables. Let n denote the number of runs in the experiment, with the j th run being made for treatment $\mathbf{x}_j = (x_{1j}, \dots, x_{vj})^T$ and initial conditions $\mathbf{u}_{0j} = (u_{j01}, \dots, u_{j0s})^T$, with observations being made at time points $\mathbf{t}_j = (t_{j1}, \dots, t_{jn_j})^T$ ($j = 1, \dots, n$). At each time point, observations $\mathbf{y}_{jl} \in \mathcal{Y}_{jl} \subset \mathbb{R}^c$ are taken on $c \leq s$ different responses. Let $\mathbf{y}_j^T = (\mathbf{y}_{j1}^T, \dots, \mathbf{y}_{jn_j}^T)$ be the cn_j -vector of observations from the j th run, and $\mathbf{y} = (\mathbf{y}_1^T, \dots, \mathbf{y}_n^T)^T$ be the vector of observations from the whole experiment.

We describe the experimental data \mathbf{y} using statistical model

$$\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\gamma}, \mathbf{d} \sim F(\boldsymbol{\theta}, \boldsymbol{\gamma}; \mathbf{d}), \quad (3)$$

with F a specified probability distribution, $\boldsymbol{\gamma} \in \Gamma \subset \mathbb{R}^q$ a q -vector of nuisance parameters, and $\mathbf{d} \in \mathcal{D}$ a vector specifying the design, chosen from the space of possible designs \mathcal{D} . The dependence of (3) on physical parameters $\boldsymbol{\theta}$ and design \mathbf{d} is through the solution to equations (1). The most common form of this dependence, assumed in this paper, is via the expected response,

$$E(\mathbf{y}_{jl}|\boldsymbol{\theta}, \mathbf{x}_j, t_{jl}) = g(\mathbf{u}(t_{jl}), \boldsymbol{\theta}),$$

with $g: \mathbb{R}^s \times \Theta \rightarrow \mathcal{Y}_{jl}$ an assumed function. However, the methodology developed here is also immediately applicable to other types of dependency.

Here, we find designs for experiments where one or more of the treatments $\mathbf{x}_1, \dots, \mathbf{x}_n$, observation times t_{j1}, \dots, t_{jn_j} , for $j = 1, \dots, n$, and initial conditions $\mathbf{u}_{01}, \dots, \mathbf{u}_{0n}$ are under the experimenters' control. In practice, some of these may be fixed by the protocol of the experiment. We also find designs where the initial conditions are unknown, and included in the vector of parameters.

In the human placenta experiment, the initial quantities of non-radioactive serine interior (u_{02}) and exterior (x_2) to the vesicle can be varied, with the initial quantities of radioactive serine (u_{01} and x_1) fixed by the experimental protocol. The $c = 1$ observed response, y_{jl} , is the amount of interior radioactive serine at time t_{jl} ($j = 1, \dots, n; l = 1, \dots, n_j$). A statistical model is assumed where $E(y_{jl}|\boldsymbol{\theta}, \mathbf{x}_j, t_{jl}) = u_1(t_{jl}; \mathbf{x}_j, \boldsymbol{\theta})$. Hence, for this experiment $g(\mathbf{u}, \boldsymbol{\theta}) = u_1$. The design consists of n combinations of initial quantities of exterior and interior non-radioactive serine, x_{2j} and u_{02j} , along with corresponding observation times t_{j1}, \dots, t_{jn_j} ; that is, $\mathbf{d} = [(x_{21}, u_{021})^T, \dots, (x_{2n}, u_{02n})^T, \mathbf{t}_1^T, \dots, \mathbf{t}_n^T]^T$.

Previous research on optimal design for models formed as the solution of ordinary differential equations has focussed on frequentist methods for models with additive normally distributed errors, with a design selected that maximizes a function of the Fisher information matrix for $\boldsymbol{\theta}$ (e.g. Atkinson and Bogacka, 2002 and Rodríguez-Díaz and Sánchez-León, 2014). The inverse of the Fisher information matrix provides an asymptotic approximation to the variance-covariance matrix for maximum likelihood estimators of $\boldsymbol{\theta}$. As is usual for nonlinear models, the information matrix depends on the value of $\boldsymbol{\theta}$, which is uncertain

prior to the experiment. The most common methodology to overcome this dependence is the adoption of pseudo-Bayesian techniques, where a design is found that maximizes the expectation of the function of the information matrix with respect to a prior distribution for θ . Numerical methods are used to obtain the derivatives of the expected response with respect to θ that are necessary to obtain the information matrix. Most commonly, the “direct method” (Valko and Vajda, 1984) is employed, with an additional set of differential equations being defined that then also require numerical solution. Many developments in this area have occurred in the chemical engineering literature, labeled “model-based design of experiments”; see Franceschini and Macchietto (2008) for a review.

In contrast to the above approaches, in this paper, we present and apply the first methods for decision-theoretic Bayesian optimal design for models formed from ordinary differential equations. Although straightforward in principle, Bayesian optimal design faces a number of practical difficulties. Firstly, assessment of a given design requires evaluation of an expected utility depending on high-dimensional and typically intractable integrals. Secondly, maximization of the expected utility presents a high-dimensional and stochastic optimization problem. See Ryan et al. (2016) and Woods et al. (2017) for recent reviews.

To address the high-dimensional optimization problem, we extend and apply the approximate coordinate exchange (ACE) algorithm recently proposed by Overstall and Woods (2017). A brief description of the algorithm is provided in Section 4.1 and Appendix A.

The computational burden of optimal Bayesian design is exacerbated when the model evaluations (systems states) are only available as the numerical solution to the differential equations. In addition to increasing the computational expense of evaluating the expected utility, numerical solutions introduce an additional source of uncertainty through the numerical errors that result from finite discretization of the time interval \mathcal{T} . We evaluate the expected utility by embedding within a Monte Carlo approximation scheme an adaption of the probabilistic solution to systems of differential equations proposed by Chkrebtii et al. (2016); see Section 2. In essence, this approach accounts for uncertainty due to discretization error by placing a joint Gaussian process prior on both the system states and time derivatives, and predicts future system states by conditioning on the derivatives. In Section 3, after introducing the foundations of Bayesian design, we propose innovative precomputation of variance and covariance quantities that substantially reduces the computational burden of

incorporating the probabilistic solution into a Bayesian design strategy. Our approach makes it possible to search for multi-variable designs which would otherwise be computationally infeasible.

We demonstrate the effectiveness for optimal design of the combination of Monte Carlo approximation, probabilistic numerics and cyclic descent for a variety of exemplar models in Section 4. The differing complexities of the problems addressed showcase the flexibility of the methodology. In Section 5 we apply the methodology to a realistic statistical model for the human placenta example, based on the solution to (2), and compare to designs proposed by the experimenters. We find designs for the goals of parameter estimation and model selection, where the aim is to determine if a simpler model with $\theta_3 = \theta_4$ (i.e. the two reaction rates equal) is an adequate description for the data.

2 Probabilistic solutions to ordinary differential equations

When working with numerical models implemented via computer code, it has become standard to build statistical approximations, or emulators, by performing a computer experiment to obtain model outputs at carefully selected input combinations. Most commonly, a Gaussian process (GP) prior is assumed to describe the output from the model, with the emulator formed from the updated posterior GP (conditioned on the model output from the computer experiment); see Sacks et al. (1989) and Santner et al. (2003). In contrast, central to the Chkrebtii et al. (2016) methodology is the adoption of a GP prior for the h th derivative function $\dot{u}_h(\cdot)$, $h = 1, \dots, s$, defined via mean and covariance functions $\dot{m}_{0h}(\cdot)$ and $\dot{C}_0(\cdot, \cdot)$, where we assume a common covariance function for each of the s derivatives. Such a prior implies that for any finite collection of times $\mathbf{t} = (t_1, \dots, t_w)^\top$, the joint distribution of $\dot{u}_h(\mathbf{t}) = [\dot{u}_h(t_1), \dots, \dot{u}_h(t_w)]^\top$ will be multivariate normal $N(\dot{\mathbf{m}}_{0h}(\mathbf{t}), \dot{\mathbf{C}}_0(\mathbf{t}, \mathbf{t}))$, with n -vector $\dot{\mathbf{m}}_{0h}(\mathbf{t})$ having l th entry $\dot{m}_{0h}(t_l)$ and, for vector $\mathbf{t}' = (t'_1, \dots, t'_{w'})^\top$, $w \times w'$ matrix $\dot{\mathbf{C}}_0(\mathbf{t}, \mathbf{t}')$ having lk th entry $\dot{C}_0(t_l, t'_k)$. A joint Gaussian process prior for both $\dot{u}_h(\cdot)$ and the solution function $u_h(\cdot)$ then follows directly, implying the joint distribution

$$\begin{pmatrix} \dot{u}_h(\mathbf{t}) \\ u_h(\mathbf{t}') \end{pmatrix} \sim N \left(\begin{pmatrix} \dot{\mathbf{m}}_{0h}(\mathbf{t}) \\ \mathbf{m}_{0h}(\mathbf{t}') \end{pmatrix}, \begin{pmatrix} \dot{\mathbf{C}}_0(\mathbf{t}, \mathbf{t}) & \bar{\mathbf{C}}_0(\mathbf{t}, \mathbf{t}') \\ \bar{\mathbf{C}}_0(\mathbf{t}', \mathbf{t}) & \mathbf{C}_0(\mathbf{t}', \mathbf{t}') \end{pmatrix} \right),$$




Figure 1: Plots showing 1000 draws from the probabilistic solution of $u_1(t)$ and $u_2(t)$ against t for system of equations (2) that describe the transport of serine in a human placenta.

with w -vector $\mathbf{m}_{0h}(\mathbf{t})$ having l th entry $m_{0h}(t_l) = \int_0^{t_l} \dot{m}_{0h}(z) dz + u_{0h}$, $w \times w'$ matrix $\mathbf{C}_0(\mathbf{t}, \mathbf{t}')$ having lk th entry $C_0(t_l, t'_k) = \int_0^{t_l} \int_0^{t'_k} \dot{C}_0(z, z') dz dz'$, and $w \times w'$ cross-covariance matrix $\bar{\mathbf{C}}_0(\mathbf{t}, \mathbf{t}')$ having lk th entry $\bar{C}_0(t_l, t'_k) = \int_0^{t'_k} \dot{C}_0(t_l, z) dz$; see also Solak et al. (2003) and Holsclaw et al. (2013). Hence, solution vector $u_h(\mathbf{t}) = [u_h(t_1), \dots, u_h(t_n)]^T$ follows the multivariate normal distribution $N(\mathbf{m}_{0h}(\mathbf{t}), \mathbf{C}_0(\mathbf{t}, \mathbf{t}'))$. Note that definition of the covariance function of $u_h(\mathbf{t})$ via integration ensures $C_0(0, 0) = 0$ and hence enforces the boundary condition $u_h(0) = u_{0h}$.

For a given \mathbf{x} and $\boldsymbol{\theta}$, this prior distribution can be updated using derivative evaluations on a grid $\boldsymbol{\tau} = (\tau_1, \dots, \tau_N)^T$ of time points via Algorithm 1 by sequentially conditioning on $f(\mathbf{u}, \tau_{r+1}, \mathbf{x}; \boldsymbol{\theta})$ calculated for solution state u_h sampled from the posterior distribution at point τ_r . The final marginal Gaussian process for $u_h(t)$ has mean and covariance functions given by

$$m_{Nh}(t) = u_{0h} + \bar{\mathbf{C}}_0(t, \boldsymbol{\tau}) \mathbf{B}_N \mathbf{F}_N \mathbf{e}_h, \quad C_N(t, t') = C_0(t, t') - \bar{\mathbf{C}}_0(t, \boldsymbol{\tau}) \mathbf{B}_N \bar{\mathbf{C}}_0(\boldsymbol{\tau}, t'),$$

for $h = 1, \dots, s$, where \mathbf{e}_h is the h th unit vector, and the $N \times s$ matrix of derivative evaluations \mathbf{F}_N and the updated $N \times N$ derivative covariance matrix \mathbf{B}_N are defined as in Algorithm 1.

Chkrebtii et al. (2016) allowed covariance function $\dot{C}_0(t, t')$ to depend on hyperparameters controlling the scale and length of the covariances. Given experimental data, a joint posterior distribution for the model parameters and hyperparameters can be sampled by embedding the probabilistic solution to the differential equations within a Markov chain Monte Carlo scheme. Chkrebtii et al. (2016) also suggested possible fixed values for the

Algorithm 1: Sequential updating and sampling for time points $\mathbf{t} = (t_1, \dots, t_w)^\top$ of the joint Gaussian process for the derivative and solution for the s system states for initial values \mathbf{u}_0 , treatment vector \mathbf{x} , physical parameters $\boldsymbol{\theta}$ and evaluation grid $\boldsymbol{\tau} = (\tau_1, \dots, \tau_N)^\top$, with $\tau_1 = T_0$. (Adapted from Chkrebtii et al., 2016).

1 Set $\boldsymbol{\Lambda}_1 = 0$ and $\mathbf{f}_1 = \mathbf{f}(\mathbf{u}_0, T_0, \mathbf{x}; \boldsymbol{\theta})$

2 for $r = 1, \dots, N - 1$ do

(a) Set $\boldsymbol{\tau}_r = (\tau_1, \dots, \tau_r)^\top$

(b) Compute

$$\mathbf{B}_r = (\dot{\mathbf{C}}_0(\boldsymbol{\tau}_r, \boldsymbol{\tau}_r) + \boldsymbol{\Lambda}_r)^{-1}$$

$$\mathbf{a}_r = \mathbf{B}_r \bar{\mathbf{C}}_0(\boldsymbol{\tau}_r, \tau_{r+1})$$

$$\mathbf{C}_r = \mathbf{C}_0(\tau_r, \tau_r) - \bar{\mathbf{C}}_0(\tau_{r+1}, \boldsymbol{\tau}_r) \mathbf{B}_r \bar{\mathbf{C}}_0(\boldsymbol{\tau}_r, \tau_{r+1})$$

$$\dot{\mathbf{C}}_{r+1} = \dot{\mathbf{C}}_0(\tau_{r+1}, \tau_{r+1}) - \dot{\mathbf{C}}_0(\tau_{r+1}, \boldsymbol{\tau}_r) \mathbf{B}_r \dot{\mathbf{C}}_0(\boldsymbol{\tau}_r, \tau_{r+1})$$

$$\boldsymbol{\Lambda}_{r+1} = \text{diag}\{\boldsymbol{\Lambda}_r, \dot{\mathbf{C}}_{r+1}\}$$

(c) Compute

$$\mathbf{m}_r = \mathbf{u}_0 + \mathbf{F}_r^\top \mathbf{a}_r, \text{ where } \mathbf{F}_r \text{ is the } r \times s \text{ matrix with } k\text{th row } \mathbf{f}_k \text{ (} k = 1, \dots, N - 1 \text{)}$$

(d) Sample

$$\mathbf{u}(\tau_{r+1}) \sim N(\mathbf{m}_r, \mathbf{C}_r \mathbf{I}_s)$$

and compute

$$\mathbf{f}_{r+1} = \mathbf{f}(\mathbf{u}(\tau_{r+1}), \tau_{r+1}, \mathbf{x}; \boldsymbol{\theta})$$

3 Compute

$$\mathbf{B}_N = (\dot{\mathbf{C}}_0(\boldsymbol{\tau}_N, \boldsymbol{\tau}_N) + \boldsymbol{\Lambda}_N)^{-1}$$

$$\mathbf{A}_N(\mathbf{t}) = \mathbf{B}_N \bar{\mathbf{C}}_0(\boldsymbol{\tau}, \mathbf{t})$$

$$\mathbf{M}_N(\mathbf{t}) = \mathbf{1}_n \otimes \mathbf{u}_0^\top + \mathbf{A}_N^\top(\mathbf{t}) \mathbf{F}_N, \text{ with } \mathbf{1}_n \text{ the } n\text{-vector with all entries equal to one and } \mathbf{F}_N \text{ the } N \times s \text{ matrix with } k\text{th row } \mathbf{f}_k \text{ (} k = 1, \dots, N \text{)}$$

$$\mathbf{C}_N(\mathbf{t}, \mathbf{t}) = \mathbf{C}_0(\mathbf{t}) - \bar{\mathbf{C}}_0(\mathbf{t}, \boldsymbol{\tau}) \mathbf{B}_N \bar{\mathbf{C}}_0(\boldsymbol{\tau}, \mathbf{t})$$

4 For $h = 1, \dots, s$, sample

$$u_h(t_1), \dots, u_h(t_n) \sim N(\mathbf{M}_N(t) \mathbf{e}_h, \mathbf{C}_N(\mathbf{t}, \mathbf{t})), \text{ where } \mathbf{e}_h \text{ is the } h\text{th unit vector}$$

covariance hyperparameters. In Section 3.2 we demonstrate the computational savings that can be achieved for optimal design via precomputing of various posterior quantities when these parameters are fixed.

Figure 1 presents 1000 draws from probabilistic solutions for the placenta example following equations (2). Updated Gaussian processes for $u_1(t)$ and $u_2(t)$ were generated using Algorithm 1 assuming a squared exponential covariance function for $\dot{C}(t, t')$ (see Rasmussen and Williams, 2006, p. 83). An evaluation grid $\boldsymbol{\tau}$ with $N = 501$ evenly spaced time points was used, and the solution sampled for time $t \in [T_0, T_1] = [0, 600]$ seconds with physical parameters $\boldsymbol{\theta} = (200, 0.05, 100, 100)^\top$, initial values $\mathbf{u}_0 = (0, 1000)^\top$ and treatment $\mathbf{x} = (7.5, 1000)^\top$. Note how the uncertainty in the solution increases as t increases away from $t = T_0 = 0$ where we know, in this example, the true value of $\mathbf{u}(t)$.

3 Bayesian design for ordinary differential equation models

3.1 Decision-theoretic Bayesian optimal design

Design of experiments fits naturally within a Bayesian framework, with the decision on what design \mathbf{d} to employ made before the data is collected. Hence it is natural to use available prior information to inform this choice. This information includes the form of statistical model (3) including any underpinning physical theory, for example, as encapsulated in equations such as (1). It also includes any prior information on the values of the parameters $\boldsymbol{\theta}, \boldsymbol{\gamma}$, captured via a prior density $\pi(\boldsymbol{\theta}, \boldsymbol{\gamma})$, which we assume is independent of the design.

A decision-theoretic Bayesian optimal design, \mathbf{d}^* , maximizes the expectation of a specified utility function $\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d})$ with respect to the unknowns prior to experimentation,

$$\begin{aligned} \Phi(\mathbf{d}^*) &= \max_{\mathbf{d} \in \mathcal{D}} \mathbb{E} [\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) | \mathbf{d}] \\ &= \max_{\mathbf{d} \in \mathcal{D}} \int_{\boldsymbol{\theta}, \mathbf{y}} \phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) \pi(\boldsymbol{\theta}, \mathbf{y} | \mathbf{d}) \, \mathrm{d}\boldsymbol{\theta} \, \mathrm{d}\mathbf{y}, \end{aligned}$$

where the joint distribution of the unknown physical parameters and responses, conditional on the design used for data collection, can be decomposed as

$$\pi(\boldsymbol{\theta}, \mathbf{y} | \mathbf{d}) = \int_{\Gamma} \pi(\mathbf{y} | \boldsymbol{\theta}, \boldsymbol{\gamma}, \mathbf{d}) \pi(\boldsymbol{\theta}, \boldsymbol{\gamma}) \, \mathrm{d}\boldsymbol{\gamma},$$

and hence, when regarded as a function of $\boldsymbol{\theta}$ alone, is proportional to the posterior density. See the seminal review paper by Chaloner and Verdinelli (1995).

The function $\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d})$ quantifies the utility, relative to the aims of the experiment, for choosing design \mathbf{d} when we obtain data \mathbf{y} under physical parameters $\boldsymbol{\theta}$. Its choice should reflect the goals of the experiment. Here, we apply the following exemplar utility functions:

1. Negative squared error loss (NSEL) for estimation of $\boldsymbol{\theta}$:

$$\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = -\|\boldsymbol{\theta} - \mathbb{E}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})\|_2^2,$$

with $\|\cdot\|_p$ denoting the l_p -norm and $\mathbb{E}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})$ the posterior mean, where expectation is taken with respect to the marginal density $\pi(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d}) = \int_{\Gamma} \pi(\boldsymbol{\theta}, \boldsymbol{\gamma}|\mathbf{y}, \mathbf{d}) d\boldsymbol{\gamma}$. It can be shown that the expected utility simplifies to

$$\Phi(\mathbf{d}) = - \int_{\mathcal{Y}} \text{tr} \{ \text{var}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d}) \} \pi(\mathbf{y}|\mathbf{d}) d\mathbf{y},$$

the negative expected value of the posterior variance-covariance matrix for $\boldsymbol{\theta}$ with respect to the marginal distribution of the response.

2. Negative absolute error loss (NAEL) for estimation of $\boldsymbol{\theta}$:

$$\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = -\|\boldsymbol{\theta} - \text{Med}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})\|_1,$$

with $\text{Med}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})$ the vector of marginal posterior medians of the physical parameters.

3. Shannon information gain (SIG) for $\boldsymbol{\theta}$:

$$\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = \log \pi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{d}) - \log \pi(\mathbf{y}|\mathbf{d}), \tag{4}$$

where

$$\pi(\mathbf{y}|\mathbf{d}) = \int_{\Theta} \pi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{d})\pi(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad \pi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{d}) = \int_{\Gamma} \pi(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\gamma}, \mathbf{d})\pi(\boldsymbol{\gamma}) d\boldsymbol{\gamma}.$$

Maximizing the expectation of (4) is equivalent to maximizing the expected Kullback-Liebler divergence between the prior and posterior distributions (Chaloner and Verdinelli, 1995).

For the human placenta example, we also employ two bespoke utility functions tailored to the problems of point estimation and model selection.

4. 0-1 utility for estimation of $\boldsymbol{\theta}$:

$$\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = \mathbb{1}_{\check{\Theta}} [\mathbb{E}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})] ,$$

with $\mathbb{1}_{\check{\Theta}}$ the indicator function for the product set

$$\check{\Theta} = \prod_{i=1}^p \check{\Theta}_i = \{(\check{\theta}_1, \dots, \check{\theta}_p) \mid \check{\theta}_i \in \check{\Theta}_i \forall i \in \{1, \dots, p\}\} ,$$

where $\check{\Theta}_i = \{\check{\theta} \mid \theta_i - \delta_i \leq \check{\theta} \leq \theta_i + \delta_i\}$, and $\boldsymbol{\delta} = (\delta_1, \dots, \delta_p)^\text{T}$ is a specified tolerance vector. That is, the utility is equal to 1 if, for all $i = 1, \dots, p$, the i th element of the posterior mean vector $\mathbb{E}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})$ lies within δ_i of the corresponding element of $\boldsymbol{\theta}$.

For the final utility function considered we redefine the utility as a function of the chosen model $m \in \mathcal{M}$.

5. 0-1 utility for model selection:

$$\phi(m, \mathbf{y}, \mathbf{d}) = \mathbb{1}_m(m^*) ,$$

where $\mathbb{1}_m$ is the indicator function for the singleton set with element m and $m^* \in \arg \max_{m \in \mathcal{M}} \pi(m|\mathbf{y})$ is the model with maximum posterior probability. For this utility, the expected utility is given by

$$\Phi(\mathbf{d}) = \sum_{m \in \mathcal{M}} \pi(m) \int_{\mathbf{y}} \phi(m, \mathbf{y}, \mathbf{d}) \pi(\mathbf{y}|m, \mathbf{d}) \, \mathrm{d}\mathbf{y} ,$$

with $\pi(\mathbf{y}|m, \mathbf{d}) = \int_{\Theta^{(m)}} \int_{\Gamma^{(m)}} \pi(\mathbf{y}|\boldsymbol{\theta}^{(m)}, \boldsymbol{\gamma}^{(m)}, m, \mathbf{d}) \pi(\boldsymbol{\theta}^{(m)}, \boldsymbol{\gamma}^{(m)}|m) \, \mathrm{d}\boldsymbol{\gamma}^{(m)} \, \mathrm{d}\boldsymbol{\theta}^{(m)}$, and $\boldsymbol{\theta}^{(m)} \in \Theta^{(m)}$ and $\boldsymbol{\gamma}^{(m)} \in \Gamma^{(m)}$ physical and nuisance parameters, respectively, for model m .

A barrier to the application of Bayesian design for most nonlinear models, including those considered in this paper, is the analytic intractability of both the utility function (which typically depends on posterior quantities) and expected utility. Numerical methods are therefore required, with a double-loop Monte Carlo approximation being commonly employed (Ryan, 2003). Such an approach uses an “inner” Monte Carlo sample of size \tilde{B} to approximate any necessary posterior quantities, and then an “outer” Monte Carlo sample of size B to approximate the expected utility with respect to the joint distribution of \mathbf{y} and $\boldsymbol{\theta}$; see also Overstall and Woods (2017).

We use the approximation

$$\hat{\Phi}(\mathbf{d}) = \frac{1}{B} \sum_{b=1}^B \hat{\phi}(\boldsymbol{\theta}_b, \mathbf{y}_b, \mathbf{d}), \quad (5)$$

with $\{\boldsymbol{\theta}_b, \mathbf{y}_b\}_{b=1}^B$ a first (outer) sample from the joint distribution of the physical parameters and response, and $\hat{\phi}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d})$ a further Monte Carlo approximation to the utility function.

Each of the utility functions above can be approximated using a second (inner) Monte Carlo sample $\{\tilde{\boldsymbol{\theta}}_{\tilde{b}}, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}\}_{\tilde{b}=1}^{\tilde{B}}$ from distribution with density $\pi(\boldsymbol{\theta}, \boldsymbol{\gamma})$:

1. NSEL:

$$\hat{\phi}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = -\|\boldsymbol{\theta} - \hat{\mathbf{E}}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})\|_2^2,$$

for an importance sampling estimate of $\mathbf{E}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})$,

$$\hat{\mathbf{E}}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d}) = \sum_{\tilde{b}=1}^{\tilde{B}} w_{\tilde{b}} \tilde{\boldsymbol{\theta}}_{\tilde{b}}, \quad (6)$$

with

$$w_{\tilde{b}} = \frac{\pi(\mathbf{y}|\tilde{\boldsymbol{\theta}}_{\tilde{b}}, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}, \mathbf{d})}{\sum_{\tilde{b}=1}^{\tilde{B}} \pi(\mathbf{y}|\tilde{\boldsymbol{\theta}}_{\tilde{b}}, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}, \mathbf{d})}. \quad (7)$$

See Ryan et al. (2016) and references therein.

2. NAEL:

$$\hat{\phi}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = -\|\boldsymbol{\theta} - \widehat{\text{Med}}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})\|_1,$$

with vector $\widehat{\text{Med}}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})$ having i th entry $\widehat{\text{Med}}_i(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d}) = (\tilde{\theta}_{i(z)} + \tilde{\theta}_{i(z+1)})/2$ ($i = 1, \dots, p$), where $\tilde{\theta}_{i(1)} \leq \dots \leq \tilde{\theta}_{i(\tilde{B})}$ are the ordered values taken by the i th element of the sample $\{\tilde{\boldsymbol{\theta}}_{\tilde{b}}\}_{\tilde{b}=1}^{\tilde{B}}$, $z = \max\{l = 1, \dots, \tilde{B} | \sum_{\tilde{b}=1}^l w_{i(\tilde{b})} \leq 0.5\}$ and the $w_{i(\tilde{b})}$ are the weights (7) ordered according to $\theta_{i(\tilde{b})}$.

3. SIG:

$$\hat{\phi}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = \log \hat{\pi}(\mathbf{y}|\boldsymbol{\theta}, \mathbf{d}) - \log \hat{\pi}(\mathbf{y}|\mathbf{d}),$$

with

$$\hat{\pi}(\mathbf{y}|\mathbf{d}) = \frac{1}{\tilde{B}} \sum_{\tilde{b}=1}^{\tilde{B}} \pi(\mathbf{y}|\tilde{\boldsymbol{\theta}}_{\tilde{b}}, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}, \mathbf{d}), \quad \hat{\pi}(\mathbf{y}|\boldsymbol{\theta}, \mathbf{d}) = \frac{1}{\tilde{B}} \sum_{\tilde{b}=1}^{\tilde{B}} \pi(\mathbf{y}|\boldsymbol{\theta}, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}, \mathbf{d}).$$

4. 0-1 estimation:

$$\hat{\phi}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d}) = \mathbb{1}_{\hat{\Theta}} \left[\hat{\mathbb{E}}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d}) \right],$$

for $\hat{\mathbb{E}}(\boldsymbol{\theta}|\mathbf{y}, \mathbf{d})$ once again the importance sampling estimate (6) of the posterior mean.

5. 0-1 model selection:

$$\hat{\phi}(m, \mathbf{y}, \mathbf{d}) = \mathbb{1}_m(\hat{m}^*),$$

where $\hat{m}^* \in \arg \max_{m \in \mathcal{M}} \pi(m) \sum_{\tilde{b}=1}^{\tilde{B}} \pi(\mathbf{y}|\tilde{\boldsymbol{\theta}}_{\tilde{b}}^{(m)}, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}^{(m)}, m, \mathbf{d}) / \tilde{B}$ for $\left\{ \tilde{\boldsymbol{\theta}}_{\tilde{b}}^{(m)}, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}^{(m)} \right\}_{\tilde{b}=1}^{\tilde{B}}$ a sample from the prior distribution under model m with density $\pi(\boldsymbol{\theta}^{(m)}, \boldsymbol{\gamma}^{(m)}|m)$.

The above Monte Carlo approximations $\hat{\phi}$ to the utility functions will introduce some bias into the approximation of the expected utility, as the utilities are nonlinear functions of posterior quantities. In general, this bias will be inversely proportional to the value of \tilde{B} , and hence can be made negligible for large inner samples.

3.2 Extensions to ordinary differential equation models

To apply the methodology outlined in the previous section to models built from systems of ordinary differential equations requires incorporation of further steps to account for discretization errors in the numerical solution to the equations, and to mitigate the additional computational cost of multiple evaluations of the numerical solution. The approximations to the expected utilities require repeated sampling of \mathbf{y} from distribution (3), and evaluation of the corresponding density function $\pi(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\gamma}, \mathbf{d})$. When the distribution of \mathbf{y} depends on the solution vector, the approximations require at least $B + \tilde{B}$ evaluations of a numerical solution to $\mathbf{u}(t_{jl}; \mathbf{x}_j, \boldsymbol{\theta})$ for each $j = 1, \dots, n$ and $l = 1, \dots, n_j$. In addition to the computational cost of these repeated evaluations, the necessary discretization of the time domain by a numerical solver introduces an additional source of uncertainty that should be accounted for in both the design of the experiment and the subsequent inference.

The probabilistic solution of Chkrebtii et al. (2016), outlined in Section 2, fits naturally within a Monte Carlo approximation of the expected utility; for each generated value of the physical parameters $\boldsymbol{\theta}$, a solution path for $\mathbf{u}(t)$ is generated from an updated Gaussian process. The uncertainty introduced by the discretization of time is quantified, and updated, via the joint Gaussian process prior for the time derivatives and solution. Algorithm 2

outlines the steps in generating an approximation to a general utility function ϕ using double loop Monte Carlo. As given, Algorithms 1–2 depend on the initial values \mathbf{u}_{0j} for the j th treatment, i.e. the initial values are assumed known. In some situations, learning unknown initial values may be part of the inference problem, i.e. prior distributions are assumed and updated to a posterior distribution in light of the experimental responses. This case can be incorporated into these algorithms by replacing all occurrences of \mathbf{u}_{0j} by a value \mathbf{u}_{0jb} generated from the prior distribution in Algorithm 2, in an analogy to how the physical parameters $\boldsymbol{\theta}$ are handled.

Algorithm 2: Evaluation of the approximate expected utility $\hat{\Phi}(d)$ when the distribution of the response depends on the solution to an ordinary differential equation.

```

1 for  $\tilde{b} = 1, \dots, \tilde{B}$  do
    Sample  $(\tilde{\boldsymbol{\theta}}_{\tilde{b}}^T, \tilde{\boldsymbol{\gamma}}_{\tilde{b}}^T)^T \sim \pi(\boldsymbol{\theta}, \boldsymbol{\gamma})$  (the prior distribution)
    for  $j = 1, \dots, n$  do
        for  $l = 1, \dots, n_j$  do
            Sample  $u_s(t_{jl}; \mathbf{x}_j, \tilde{\boldsymbol{\theta}}_{\tilde{b}})$  using Algorithm 1
2 for  $b = 1, \dots, B$  do
    Sample  $(\boldsymbol{\theta}_b^T, \boldsymbol{\gamma}_b^T)^T \sim \pi(\boldsymbol{\theta}, \boldsymbol{\gamma})$  (the prior distribution)
    for  $j = 1, \dots, n$  do
        for  $l = 1, \dots, n_j$  do
            Sample  $u_s(t_{jl}; \mathbf{x}_j, \boldsymbol{\theta}_b)$  using Algorithm 1
        Sample  $\mathbf{y}_j | \boldsymbol{\theta}_b, \boldsymbol{\gamma}_b, d \sim F(\boldsymbol{\theta}_b, \boldsymbol{\gamma}_b; d)$ 
    Calculate  $\hat{\phi}(\boldsymbol{\theta}_b, \mathbf{y}_b, d)$  using the inner sample generated in step 1
3 Calculate  $\hat{\Phi}(d) = \frac{1}{B} \sum_{b=1}^B \hat{\phi}(\boldsymbol{\theta}_b, \mathbf{y}_b, d)$ 

```

Naive implementation of Algorithm 2 for approximating the expected utility presents a considerable computational challenge, with the matrix computations in steps 2(b) and 3 of Algorithm 1 being undertaken $\tilde{n}(B + \tilde{B})$ times, with $\tilde{n} = \sum_{j=1}^n n_j$. In particular, calculation of matrix \mathbf{B}_N requires inversion of an $N \times N$ matrix. This leads to an algorithm with computational complexity $\mathcal{O}(\tilde{n}N^3(B + \tilde{B}))$.

To reduce the computational cost of the algorithm, we can compromise on the choice of covariance function $\dot{C}_0(t, t')$. Rather than tune the covariance through the selection of different parameter values for each choice of \mathbf{x} and $\boldsymbol{\theta}$, we can fix these parameters (e.g. fol-

lowing recommendations in Chkrebtii et al. (2016); see Section 4 for our choices). This allows precomputation of various covariance matrices and vectors, see Algorithm 3. Such precomputation alleviates the need to invert \mathbf{B}_N when sampling $u(t)$, reducing the computational complexity of the approximation to $\mathcal{O}(N^3 + \tilde{n}N^2(B + \tilde{B}))$.

In fact, this precomputation can be performed just once, *prior* to any optimization routine being called. Hence for large experiments and Monte Carlo sample sizes, the computational complexity of the precomputation is essentially fixed, and the complexity of the approximation within the optimization becomes $\mathcal{O}(\tilde{n}N^2(B + \tilde{B}))$. This computational savings makes the optimization feasible for experiment sizes, evaluation grids and Monte Carlo sample sizes for which designs could not otherwise be found.

Algorithm 3: Precomputation of variances C_r , \dot{C}_{r+1} , \mathbf{B}_r and covariances \mathbf{a}_r for evaluation grid $\boldsymbol{\tau} = (\tau_1, \dots, \tau_N)^\top$

- 1 Set $\boldsymbol{\Lambda}_1 = 0$
 - 2 **for** $r = 1, \dots, N - 1$ **do**
 - (a) Set $\boldsymbol{\tau}_r = (\tau_1, \dots, \tau_r)^\top$
 - (b) Compute

$$\mathbf{B}_r = (\dot{C}_0(\boldsymbol{\tau}_r, \boldsymbol{\tau}_r) + \boldsymbol{\Lambda}_r)^{-1}$$

$$\mathbf{a}_r = \mathbf{B}_r \bar{C}_0(\boldsymbol{\tau}_r, \tau_{r+1})$$

$$C_r = C_0(\tau_r, \tau_r) - \bar{C}_0(\tau_{r+1}, \boldsymbol{\tau}) \mathbf{B}_r \bar{C}_0(\boldsymbol{\tau}_r, \tau_{r+1})$$

$$\dot{C}_{r+1} = \dot{C}_0(\tau_{r+1}, \tau_{r+1}) - \dot{C}_0(\tau_{r+1}, \boldsymbol{\tau}_r) \mathbf{B}_r \dot{C}_0(\boldsymbol{\tau}_r, \tau_{r+1})$$

$$\boldsymbol{\Lambda}_{r+1} = \text{diag}\{\boldsymbol{\Lambda}_r, \dot{C}_{r+1}\}$$
 - 3 Compute $\mathbf{B}_N = \left(\dot{C}_0(\boldsymbol{\tau}_N, \boldsymbol{\tau}_N) + \boldsymbol{\Lambda}_N \right)^{-1}$
-

4 Examples

4.1 Preliminaries

In this section we demonstrate the Bayesian design methodology for three common examples of models formed from the solution of ordinary differential equations:

1. a compartmental model (Section 4.2);
2. a model formed from the FitzHugh-Nagumo equations (Section 4.3);

3. a model of the JAK-STAT mechanism (Section 4.4).

For each, we use the methodology in Section 3.2 to approximate expected utilities for parameter estimation. Bayesian optimal (or near optimal) designs are found by embedding these Monte Carlo approximations within the ACE algorithm (Overstall and Woods, 2017). The ACE algorithm is a cyclic descent, or coordinate exchange, algorithm (see Meyer and Nachtsheim, 1995 and Lange, 2013, p. 171) that performs a sequence of conditional maximizations for each element (coordinate) of \mathbf{d} in turn, keeping all other elements fixed. Each of these one-dimensional maximizations is performed by constructing a Gaussian process smoother, or emulator, for the Monte Carlo approximation as a function of the coordinate. Use of an emulator alleviates both the computational burden and lack of smoothness associated with the Monte Carlo approximations. This algorithm extends the optimal design via curve fitting methods originally presented by Müller and Parmigiani (1996) to high-dimensional design problems. The ACE algorithm is outlined in Appendix A and implemented in the `acebayes` R package (Overstall et al., 2018b, Overstall et al., 2018c), available on CRAN.

To employ the probabilistic solution to the ordinary differential equations, a choice of covariance function is required for the Gaussian process prior on the derivative functions. The choice of covariance function should be determined by the assumed smoothness of the solutions $u_h(t)$. Chkrebtii et al. (2016) suggested two covariance functions, the squared exponential covariance

$$\hat{C}_0(t, t') = \sqrt{\pi} \alpha^{-1} \lambda \exp \left\{ -(t - t')^2 / 4\lambda^2 \right\}, \quad (8)$$

which is infinitely differentiable and hence suitable for smooth solutions, and the piecewise linear uniform covariance

$$\hat{C}_0(t, t') = \begin{cases} \alpha^{-1} \{ \min(t, t') - \max(t, t') + 2\lambda \} & \text{for } \{ \max(t, t') - \min(t, t') \} / 2 > \lambda, \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

where $\alpha, \lambda > 0$. This latter function is non-differentiable and hence suited to non-smooth solutions. We employ these two functions, with fixed values of α and λ to facilitate the precomputation outlined in Section 3.2. Throughout, we assume the probabilistic solution is calculated on a grid $\boldsymbol{\tau} = (\tau_1, \dots, \tau_N)^T$ of equally-spaced points and, unless otherwise stated, set $\alpha = N$ and $\lambda = 4(\tau_N - \tau_1)/N$.

The Supplementary Material contains an R package called `aceodes` and a vignette. The vignette describes how `aceodes` can be used to reproduce the designs found in the remainder of this section and in Section 5.

4.2 Compartmental model

In pharmacokinetics studies, compartmental models are used to describe the distribution of a drug inside a living body. Such models have been routinely used to demonstrate optimal experimental design methodology (see, for example, Atkinson et al., 1993, Ryan et al., 2014, and Overstall and Woods 2017). To compare designs found using the probabilistic solution to designs found using an exact solution, we use a simple example where an analytical solution to the differential equations is available. An open one-compartment model is considered with first-order absorption, described by the following system of $s = 2$ ordinary differential equations for $t \in [0, 24]$ hours:

$$\begin{aligned} \dot{u}_1(t) &= -\theta_1 u_1(t), \\ \dot{u}_2(t) &= (\theta_2/\theta_3)u_1(t) - \theta_2 u_2(t), \\ \mathbf{u}(0) &= (D, 0)^T, \end{aligned}$$

where $u_1(t)$ and $u_2(t)$ are respectively the amounts of drug outside and inside the body, D is the known initial dose, and $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)^T$ are unknown parameters.

These equations define a homogeneous linear system with constant coefficients, resulting in the analytical solution

$$\begin{aligned} u_1(t) &= D \exp(-\theta_1 t), \\ u_2(t) &= \frac{D\theta_2}{\theta_3(\theta_2 - \theta_1)} (\exp(-\theta_1 t) - \exp(-\theta_2 t)). \end{aligned} \quad (10)$$

Following Ryan et al. (2014), we assume $D = 400$ and $\log \theta_i \sim N(\mu_i, 0.05)$, independently, for $l = 1, 2, 3$, with $(\mu_1, \mu_2, \mu_3)^T = (\log 0.1, \log 1, \log 20)^T$. The amount of drug inside the body, y_l , is observed at observation time t_l , and is modeled through assuming $y_l \sim N(u_2(t_l), \sigma^2 + \tau^2 u_2(t_l)^2)$, independently, where $\sigma^2 = 0.1$ and $\tau^2 = 0.01$. The choice of design here only involves selecting $n = 15$ observation times: t_1, \dots, t_n . We impose the practically realistic constraint that the observation times have to be at least 15 minutes apart. Such a constraint is straightforward to incorporate into the ACE algorithm (see Overstall and Woods, 2017).

When applying the probabilistic solution, we assume squared exponential covariance (8) as the functions $\mathbf{u}(t)$ are known to be smooth and a discrete evaluation grid, $\boldsymbol{\tau}$, with $N = 501$.

For each of the NSEL, NAEL and SIG utility functions from Section 3.1, we compare designs found under the exact and probabilistic solutions using ACE to a uniform design with $n = 15$ equally-spaced time points in $[0, 24]$ hours. Figure 2 presents boxplots of twenty evaluations of the Monte Carlo approximation to the expected utility for the uniform design and the optimal design found for each utility. There is negligible difference between the designs found under the exact and probabilistic solutions, and these designs are clearly superior to the uniform design. Figure 2 also gives the observation time points from each design being compared. The optimal designs appear to favor observation times near the peak of $u_2(t)$, at $t \approx 2.5$ hours, and then a series of observation times towards the end of the time interval. The optimal design under SIG has two distinct sets of points just before and after the maximum of $u_2(t)$, whereas the designs under NSEL and NAEL have just one set of points, generally occurring just after the peak response.

4.3 FitzHugh-Nagumo equations

The FitzHugh-Nagumo equations (FitzHugh, 1961 and Nagumo et al., 1962) describe the behavior of spike potential in the giant axon of squid neurons:

$$\begin{aligned}\dot{u}_1(t) &= \theta_3 [u_1(t) - u_1(t)^3/3 + u_2(t)] , \\ \dot{u}_2(t) &= -[u_1(t) - \theta_1 + \theta_2 u_2(t)] / \theta_3 , \\ \mathbf{u}(0) &= (-1, 1)^T ,\end{aligned}$$

where $u_1(t)$ is the voltage across the axon membrane, $u_2(t)$ is the recovery variable giving a summary of outward current, $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)^T$, and $t \in [0, 20]$ ms. These equations cannot be solved analytically.

We assume an experiment that measures the voltage, y_l , at time t_l , for $l = 1, \dots, n$. Following Ramsay et al. (2007), $y_i \sim \text{N}(u_1(t_i), \sigma^2)$, independently, where $\sigma \sim \text{Uniform}[1/2, 1]$. A priori, we assume $\theta_1, \theta_2 \sim \text{Uniform}[0, 1]$ and $\theta_3 \sim \text{Uniform}[1, 5]$.

As noted by Ramsay et al. (2007), the solution to the FitzHugh-Nagumo equations can alternate between smooth evolution and sharp changes of direction. Hence, we employ uniform covariance (9) for the probabilistic solution. The evaluation grid has size $N = 200$.

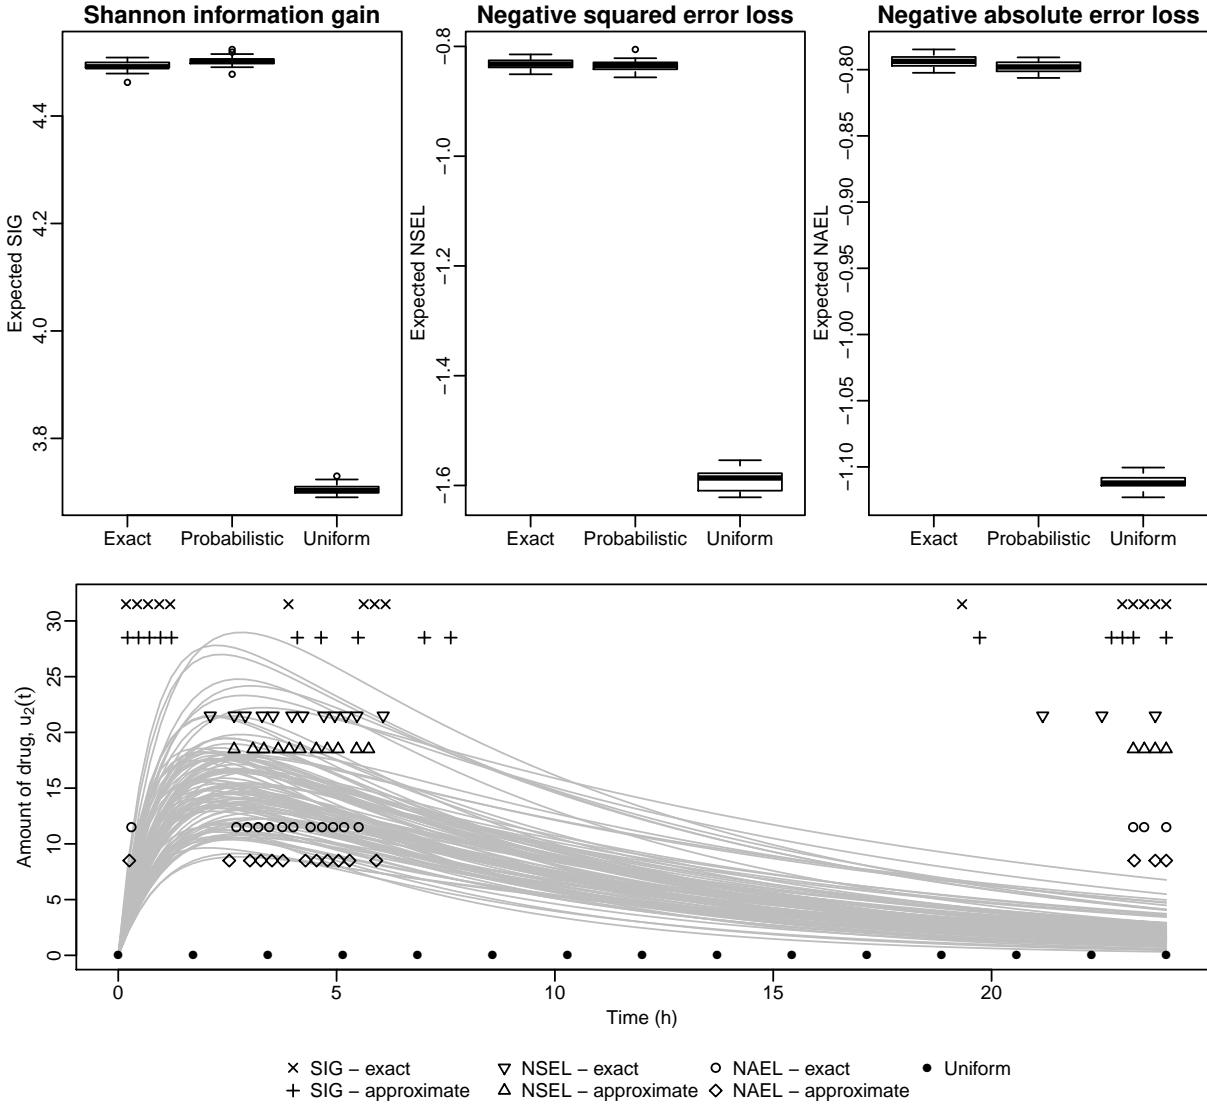


Figure 2: Results from the compartmental model in Section 4.2. Top row: boxplots of 20 evaluations of the Monte Carlo approximation to the expected utility for the uniform design and the optimal designs (for the exact and probabilistic solution) found under three different utility functions. Bottom plot: design points from each of the optimal designs and the uniform design, along with 100 draws from the exact solution, $u_2(t)$, giving the amount of drug at time t , for values drawn from the prior distribution of θ .

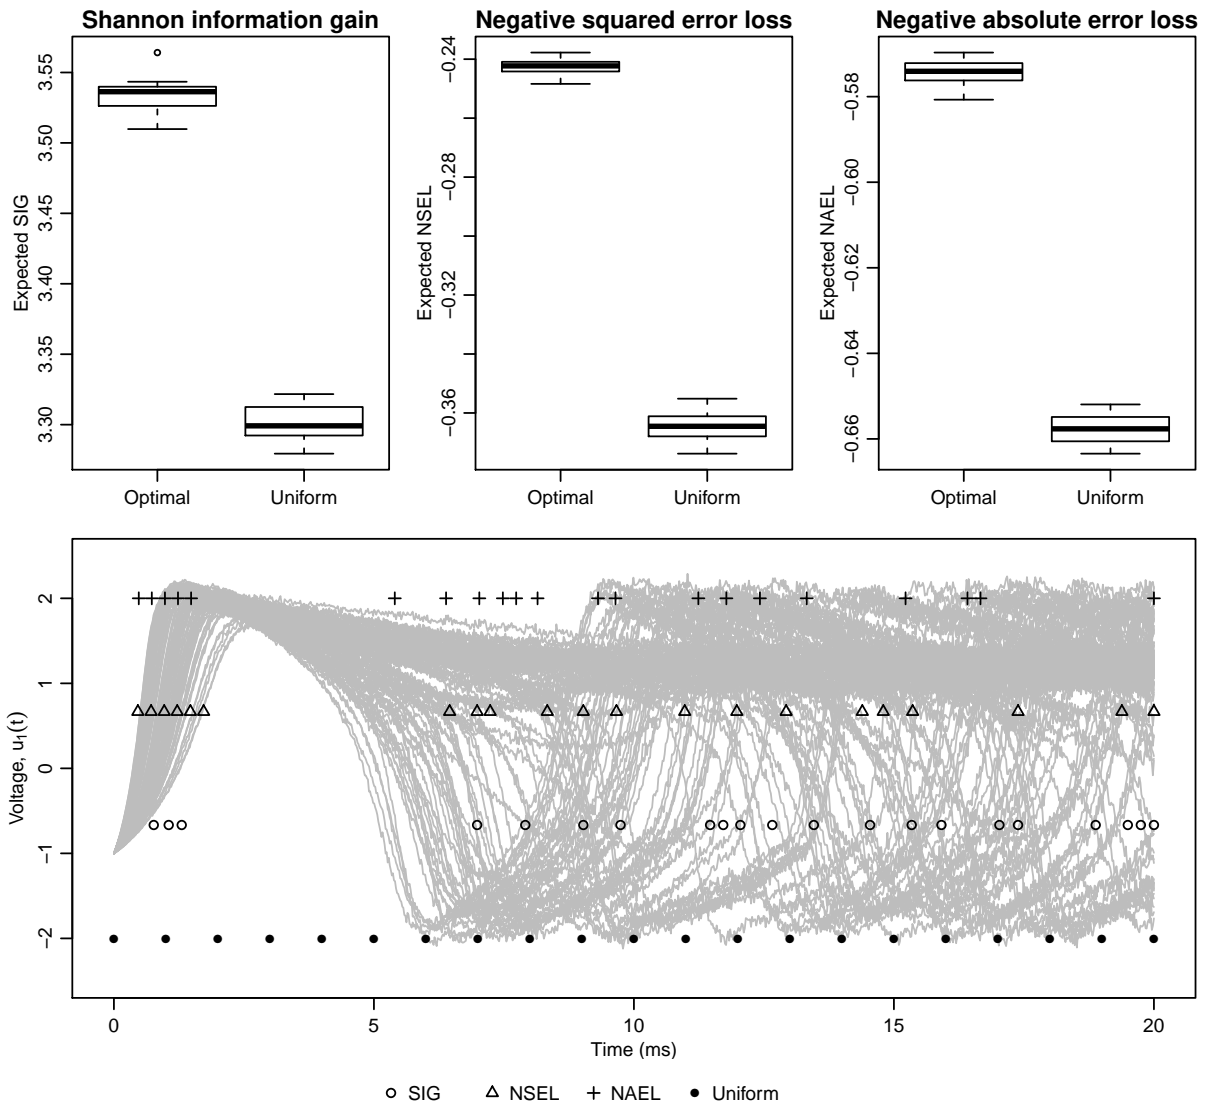


Figure 3: Results from the FitzHugh-Nagumo equations in Section 4.3. Top row: boxplots of 20 evaluations of the Monte Carlo approximation to the expected utility for the uniform design and the optimal designs found under three different utility functions. Bottom plot: design points from each of the optimal designs and the uniform design, along with 100 draws from the probabilistic solution, $u_1(t)$, giving the voltage at time t , for values drawn from the prior distribution of θ .

The design consists of the $n = 21$ observation times, t_1, \dots, t_n . Similarly to Section 4.2, we stipulate that the observation times must be at least 0.25ms apart, and find designs under the NSEL, NAEL and SIG utility functions. We compare these optimal designs to a uniform design with n equally spaced points in $[0, 20]$ ms. Figure 3 presents boxplots of twenty evaluations of the Monte Carlo approximation to the expected utility for the uniform design and the optimal designs found via ACE under each utility function. In each case, there is a clear improvement to be made over using the uniform design. Also shown in Figure 3 are the four designs under comparison, along with realizations drawn from the solution $u_1(t)$. Both the NSEL and NAEL optimal designs have a substantial number of observations near the beginning of the experiment. Both these designs have around one-third of their observation times before 2.5ms; the SIG and uniform designs only make three observations before this time. A feature of all of the optimal designs is that they make no observations between about 2.5 and 6ms, where the voltage is expected to rapidly decrease. The remaining observation times are close to being evenly spaced. The initial phase of high frequency observations provides information about the steep increase in voltage for small t . The remaining observation times aid efficient parameter estimation, occurring within an interval within which different parameter values can produce very different model solutions.

4.4 JAK-STAT mechanism

Chkrebtii et al. (2016), and authors referenced therein, considered Bayesian inference for the JAK-STAT mechanism. A system of $s = 4$ equations describes changes in the biochemical reaction states of STAT-5 transcription factors that occur in response to binding of the Erythropoietin hormone to cell surface receptors (Pellegrini and Dusanter-Fourt, 1997):

$$\left. \begin{aligned} \dot{u}_1(t) &= -\theta_1 u_1(t) \kappa(t) + 2\theta_4 u_4(t - \omega), \\ \dot{u}_2(t) &= \theta_1 u_1(t) \kappa(t) - \theta_2 u_2(t)^2, \\ \dot{u}_3(t) &= -\theta_3 u_3(t) + \frac{1}{2} \theta_2 u_2(t)^2, \\ \dot{u}_4(t) &= \theta_3 u_3(t) - \theta_4 u_4(t - \omega), \\ \mathbf{u}(t) &= (u_{01}, 0, 0, 0)^T, \end{aligned} \right\} \begin{aligned} &t \in [0, 60] \text{ seconds,} \\ &t \in [-\omega, 0], \end{aligned}$$

with $u_{01} \geq 0$ unknown and $\kappa(t)$ an unknown forcing function. The transcription states return to the initial state after gene activation in the cell nucleus, modeled via the unknown time delay $\omega \geq 0$. This system is an example of a delay initial function problem.

Swameye et al. (2003) conducted an experiment that made measurements on the nonlinear transformation of the states given by

$$g(\mathbf{u}, \boldsymbol{\theta}) = \begin{pmatrix} \theta_5(u_2 + 2u_3) \\ \theta_6(u_1 + u_2 + 2u_3) \\ u_1 \\ u_3/(u_2 + u_3) \end{pmatrix} = \begin{pmatrix} g_1(\mathbf{u}, \boldsymbol{\theta}) \\ g_2(\mathbf{u}, \boldsymbol{\theta}) \\ g_3(\mathbf{u}, \boldsymbol{\theta}) \\ g_4(\mathbf{u}, \boldsymbol{\theta}) \end{pmatrix}.$$

The experiment made $n = 16$ (noisy) observations on g_1 and g_2 at times t_1, \dots, t_{16} , one observation on each of g_3 and g_4 at $t = 0$ and $t = t^*$, respectively. The design (choices of time points) used in the experiment reported by Swameye et al. (2003) are given in Figure 4. The following statistical model is assumed

$$\begin{aligned} (y_{1l}, y_{2l})^T &\sim N([g_1(\mathbf{u}(t_l), \boldsymbol{\theta}), g_2(\mathbf{u}(t_l), \boldsymbol{\theta})]^T, \mathbf{A}_l), \\ y_3 &\sim N(g_3(\mathbf{u}(0), \boldsymbol{\theta}), \sigma_3^2), \quad y_4 \sim N(g_4(\mathbf{u}(t^*), \boldsymbol{\theta}), \sigma_4^2), \end{aligned}$$

independently, for $l = 1, \dots, n$, where $\mathbf{A}_l = \text{diag}\{\sigma_{1l}^2, \sigma_{2l}^2\}$.

We design a follow-up experiment using information from this previous study, and choose values of t_1, \dots, t_n and t^* to maximize different expected utilities assuming, for simplicity, a single observation of y_3 will also be made at $t = 0$ (as in the original experiment). We use the posterior distributions from Chkrebtii et al. (2016) as priors for $\boldsymbol{\theta}$, ω and u_{01} . These authors assumed the variance parameters were fixed. Instead, we assume $\sigma_{1l}^2 = \sigma_1^2$, $\sigma_{2l}^2 = \sigma_2^2$, for all $l = 1, \dots, n$, and $\sigma_1, \sigma_2 \sim \text{Uniform}[0, 0.1]$, $\sigma_3 \sim \text{Uniform}[0, 20]$ and $\sigma_4 \sim \text{Uniform}[0, 0.1]$. These prior distributions are consistent with the experimentally determined values used for previous analyses (see Raue et al., 2009). The forcing function $\kappa(t)$ is assumed unknown but has been measured at 16 time points. We follow Chkrebtii et al. (2016) and assume these measurements are made without error and interpolate with a Gaussian process to allow a probabilistic prediction of $\kappa(t)$ for any $t \in [0, 60]$.

The nature of the delay initial function problem introduces an added complexity to our implementation of the probabilistic solution. At the end of step 2 of Algorithm 1, we compute $\mathbf{f}_{r+1} = \mathbf{f}(\mathbf{u}(\tau_{r+1}), \tau_{r+1}, \boldsymbol{\theta}_b)$. For this example, to compute \mathbf{f}_{r+1} , we require $u_4(\tau_{r+1} - \omega_b)$, where ω_b is a value generated from the prior distribution of ω . If $\tau_{r+1} - \omega_b \leq 0$, then $u_4(\tau_{r+1} - \omega_b) = 0$ as specified by the initial conditions of the system of equations. For $\tau_{r+1} - \omega_b > 0$, the conditional distribution of $u_4(\tau_{r+1} - \omega_b)$ can be derived in the probabilistic

solution of Chkrebtii et al. (2016) and a value for u_4 generated. However, this will be computationally expensive to incorporate in the implementation of the probabilistic solution described in Section 3.2 and would prevent the precomputation in Algorithm 3. Hence, if $\tau_{r+1} - \omega_b > 0$, we replace $u_4(\tau_{r+1} - \omega_b)$ by $u_4(\tau_{\bar{r}})$, where $\bar{r} = \arg \min_{r'=1, \dots, r+1} |\tau_{r+1} - \omega_b - \tau_{r'}|$, i.e. from the series of $u_4(\tau_1), \dots, u_4(\tau_{r+1})$ values generated in step 2 thus far, we choose the value for the time $\tau_{\bar{r}}$ that is closest in absolute value to $\tau_{r+1} - \omega_b$.

We employ uniform covariance (9) as the time delay can cause discontinuities in the derivative, as noted by Chkrebtii et al. (2016). The evaluation grid, $\boldsymbol{\tau}$, has size $N = 500$, and the auxiliary parameters are set to $\lambda = 0.085$ and $\alpha = 8000$, consistent with the posterior distribution from the original analysis.

We use the methodology from Section 3.2 and the ACE algorithm to find designs that maximize each of the NSEL, NAEL and SIG utilities. We compare these designs to the original design used by Swameye et al. (2003). As in the previous examples, we introduce the constraint that the observation times need to be at least 1 second apart, a requirement also satisfied by the original experiment. Figure 4 presents boxplots of twenty evaluations of the Monte Carlo approximation to the expected utility for the original design and the optimal designs found under each utility function. Once again, in each case, the optimal designs are considerably more efficient. Also shown in Figure 4 are the four designs under comparison. The optimal designs favor having a dense set of points early in the observation window, and then a smaller set of times near the end of the experiment. This is especially true for the designs under NSEL and NAEL where 75% of the observation times occur before $t = 15$ seconds, compared to about 60% for SIG design and 50% for the original design. Early observation times provide information about the peak in g_1 and the sharp decrease in g_2 at about 10 seconds. For the single observation time, t^* , on g_4 , the optimal designs clearly favor making a very early observation. Note that t^* for each of the optimal designs is between 1 and 2 seconds.

5 Application: transport of serine across human placenta

We now use the methodology in Section 3 to redesign the experiment for the human placenta study introduced in Section 1. The experimental protocol specifies fixed initial amounts of

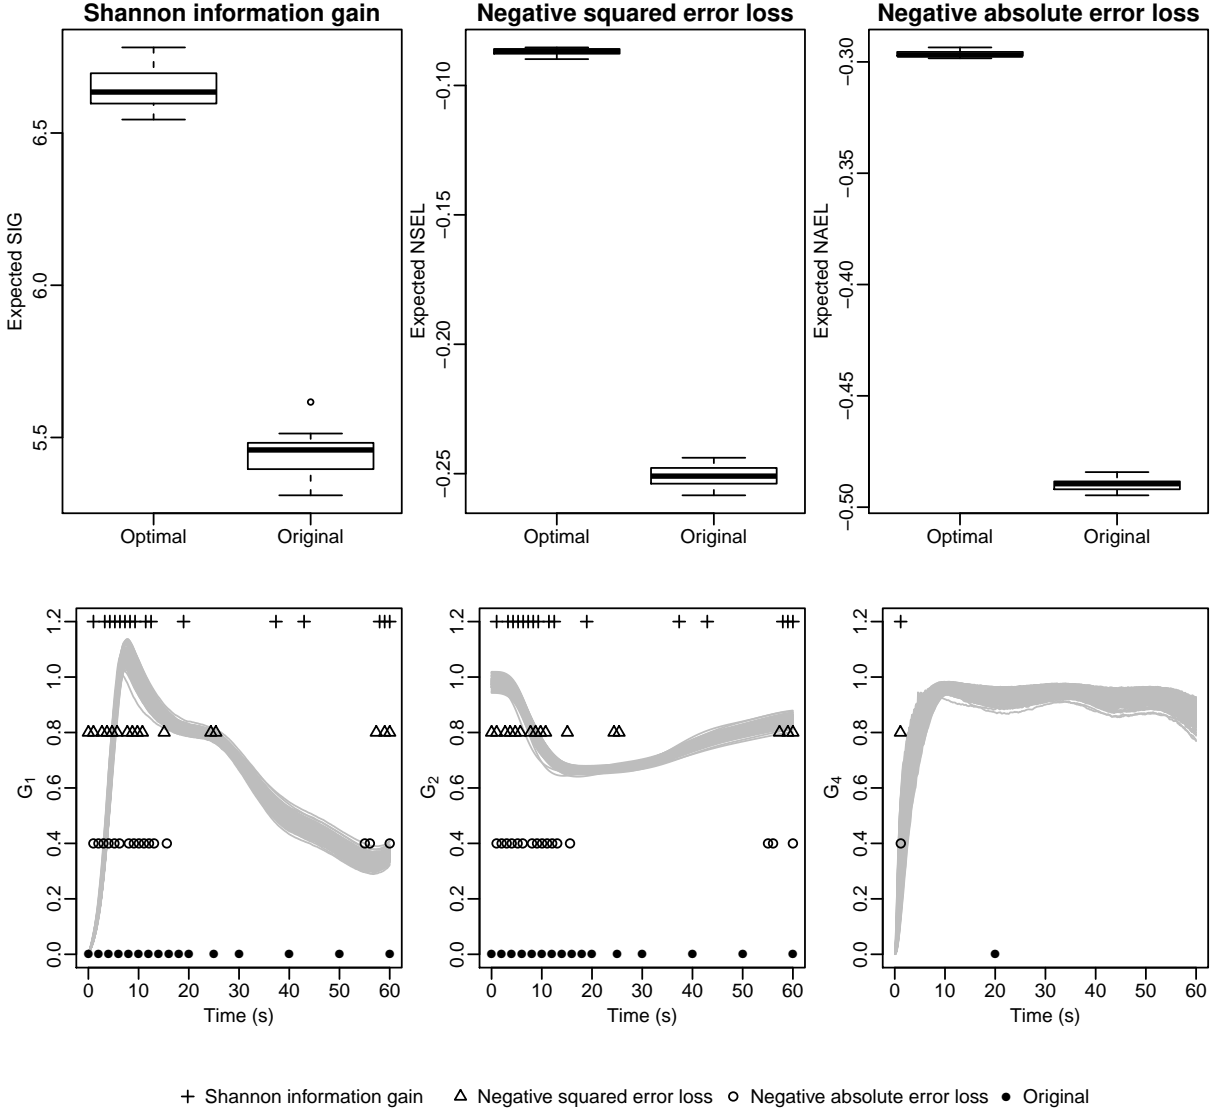


Figure 4: Results from the JAK-STAT example in Section 4.4. Top row: boxplots of 20 evaluations of the Monte Carlo approximation to the expected utility for the original design and the optimal designs found under three different utility functions. Bottom row: design points from each of the optimal designs and the original design at which noisy observations of g_1 (left), g_2 (center), g_4 (right) are made, along with 100 draws from g_1 , g_2 and g_4 , at time t , for values drawn from the prior distribution of θ , u_{01} and ω .

radioactive serine interior (u_{01}) and exterior (x_1) to the placenta (0 and $7.5\mu\text{l}$, respectively). The original design proposed by the experimenters used $n = 7$ placentas (runs) with differing amounts of non-radioactive serine interior (u_{01}) and exterior (x_2) to the placenta, see Table 1. Noisy observations on the amount of interior radioactive serine (u_1) were made at eight times, common to each of the seven placentas. The experimenters expected greater variability in the concentration of interior radioactive serine near the start of the experiment, before convergence to an equilibrium. Therefore, they choose a design containing a large number of early time points. We broadly follow this protocol, but find optimal designs using $n = 2, \dots, 7$ placentas with each having $n_t = 8$ observations taken at common times, t_1, \dots, t_8 , chosen from across the interval $[0, 600]$.

A hierarchical statistical model is assumed for the observed responses:

$$y_{jl} = u_1(t_l; \mathbf{x}_j, \boldsymbol{\theta}_j) + \varepsilon_{jl}, \quad \text{for } j = 1, \dots, n; l = 1, \dots, n_t,$$

where $\mathbf{x}_j = (x_1, x_{2j})^T$, ε_{jl} are independent and identically normally distributed with constant variance σ^2 , and $\boldsymbol{\theta}_j$ holds the $p = 4$ subject-specific parameters for the j th placenta with elements assumed to follow independent uniform distributions

$$\theta_{ji} \sim \text{U}[\theta_i(1 - c_i), \theta_i(1 + c_i)], \quad c_i > 0, \quad i = 1, \dots, p.$$

The goal of the experiment is estimation of the population physical parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)^T$.

A priori, we assume $c_i \sim \text{Uniform}[0, 0.05]$ and $\theta_i \sim \text{Tri}[a_i, b_i]$, where $\text{Tri}[a, b]$ denotes the symmetric triangle distribution on the interval $[a, b]$. Reflecting prior knowledge from previous experiments, we set $a_1 = a_3 = a_4 = 80$, $b_1 = b_3 = b_4 = 120$, $a_2 = 0.02$, $b_2 = 0.08$ and we assume $\sigma^2 \sim \text{U}[0, 1]$ for the response variance.

We expect the solution to system of equations (2) to be smooth, and so use squared exponential covariance (8) for the probabilistic solution. The evaluation grid, $\boldsymbol{\tau}$, has size $N = 601$ and we set auxiliary correlation parameter $\alpha = 10N$.

Specifying a design corresponds to specifying the n experimental conditions x_{21}, \dots, x_{2n} , initial values u_{021}, \dots, u_{02n} , and the common $n_t = 8$ observation times t_1, \dots, t_{n_t} . Hence for $n = 2, \dots, 7$, the design space has between 12 and 22 dimensions. As for the examples in Section 4, we impose a constraint on the observation times and specify that they must be at least 5 seconds apart.

Table 1: Treatments from the optimal and original designs with $n = 7$ runs for the placenta example in Section 5: initial concentrations (to nearest integer) of interior ($u_{02} = u_2(0)$) and exterior (x_2) non-radioactive serine for each run (placenta).

| Placenta | NSEL | | NAEL | | Est01 [†] | | MS01 [*] | | Original | |
|----------|-------|----------|-------|----------|--------------------|----------|-------------------|----------|----------|----------|
| | x_2 | u_{02} | x_2 | u_{02} | x_2 | u_{02} | x_2 | u_{02} | x_2 | u_{02} |
| 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 2 | 0 | 38 | 0 | 0 | 0 | 0 | 0 | 0 | 250 | 0 |
| 3 | 0 | 50 | 0 | 50 | 0 | 56 | 0 | 0 | 250 | 250 |
| 4 | 0 | 68 | 0 | 67 | 0 | 58 | 0 | 0 | 250 | 1000 |
| 5 | 182 | 1000 | 160 | 1000 | 177 | 1000 | 0 | 38 | 1000 | 0 |
| 6 | 185 | 1000 | 175 | 1000 | 196 | 1000 | 0 | 41 | 1000 | 250 |
| 7 | 206 | 1000 | 211 | 1000 | 210 | 1000 | 115 | 62 | 1000 | 1000 |

[†] 0-1 estimation utility; ^{*} 0-1 model selection utility

We find designs for the NSEL, NAEL, 0-1 estimation and 0-1 model selection utility functions defined in Section 3.1. For the 0-1 estimation utility, we set $\boldsymbol{\delta} = (5, 5, 0.01, 5)^T$; for utility $\phi(\boldsymbol{\theta}, \mathbf{y}, \mathbf{d})$ to equal 1, the posterior mean for $\boldsymbol{\theta}$ must lie in the box set $\prod_{i=1}^4 [\theta_i - \delta_i, \theta_i + \delta_i]$, which contains 0.5% of the volume of the prior support. For the model selection utility, we suppose interest is in determining if the reaction rates are equal, i.e. does $\theta_3 = \theta_4$? To answer this question, we define two models: m_1 (where $\theta_3 = \theta_4$) and m_2 (where $\theta_3 \neq \theta_4$).

Figure 5 presents boxplots of twenty evaluations of the Monte Carlo approximation to the expected utility for the optimal design found under each utility function for $n = 2, \dots, 7$. We also present boxplots of the performance of the original design with $n = 7$. Unsurprisingly, the expected utility increases with n , and the optimal designs are clearly superior to the original design. For each utility function, the optimal design with $n = 2$ outperforms the original design with $n = 7$ placentas, with substantial differences in expected utility.

Table 1 gives the treatments for each design found for $n = 7$. Figure 6 shows the observation times for the optimal designs under NSEL, NAEL and 0-1 estimation utilities, along with realizations from the solution to $u_1(t)$, for each run of each design. The designs under NSEL and NAEL utilities have similar treatments and observation times. The initial concentrations in Table 1 lead to three distinct profiles of $u_1(t)$ (labeled placentas 1 and 2; 3

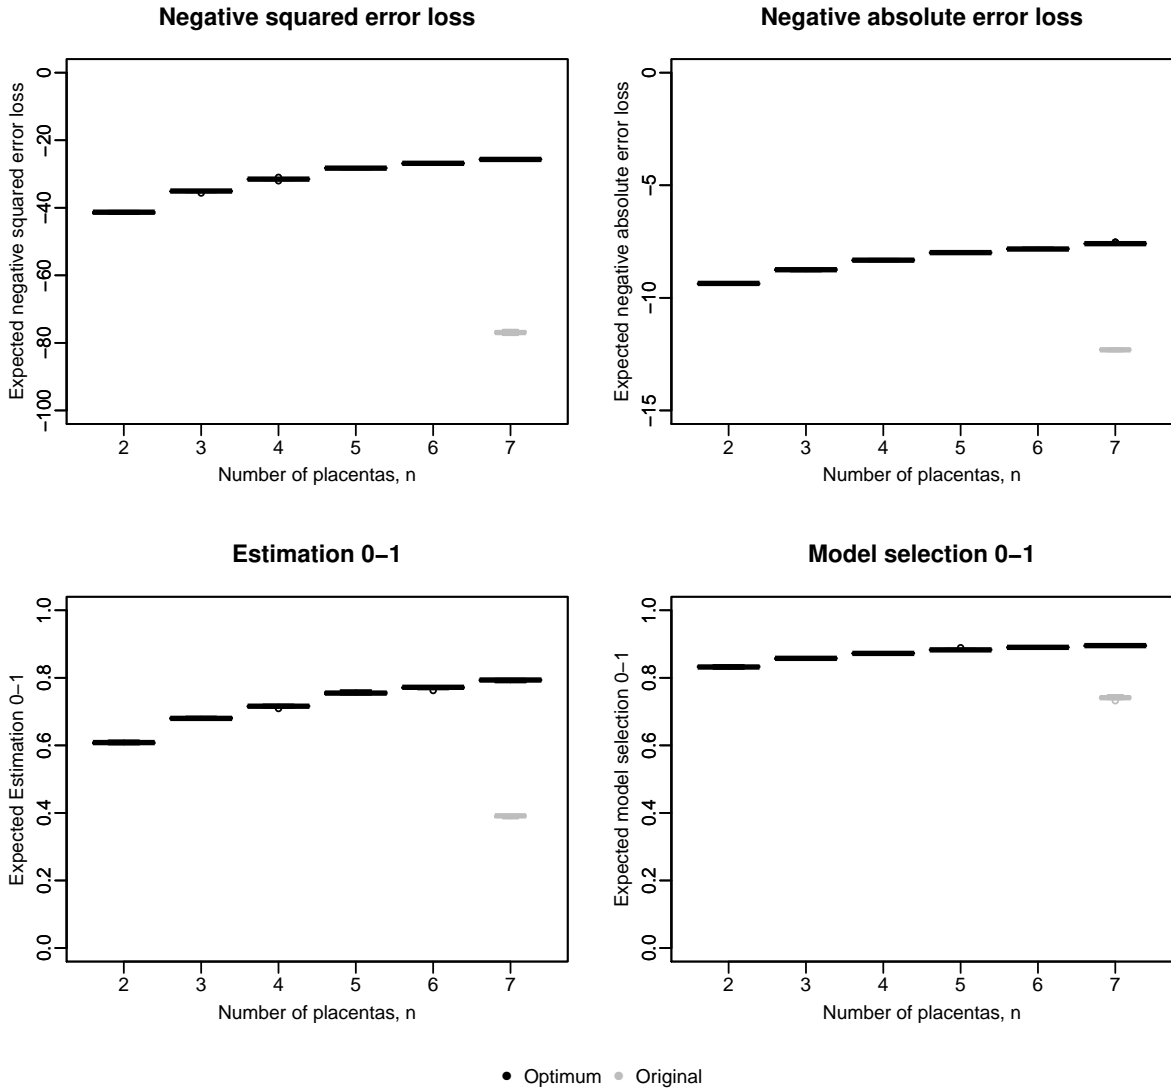


Figure 5: Results from the placenta example in Section 5. Boxplots of 20 evaluations of the Monte Carlo approximation to the expected utility for the original design and the optimal designs found under four different utility functions for $n = 2, \dots, 7$.

and 4; 5, 6 and 7; note though that the placentas are exchangeable). The profile for placentas 1 and 2 has a slow steady increase in $u_1(t)$ with respect to t . Placentas 3 and 4 have a steep initial increase and subsequent decrease in $u_1(t)$ with respect to t . Finally, placentas 5 to 7 have a steep initial increase in $u_1(t)$ with respect to t followed by a slow decrease. The optimal observation times are predominantly at the beginning of the observation window, where $u_1(t)$ is changing most quickly. The designs under the 0-1 estimation utility are also similar, except a non-zero amount (35 μl) of the initial interior non-radioactive serine is applied to placenta 2.

Figure 7 shows the designs from the 0-1 model selection utility, along with realizations of the solutions $u_1(t)$ under models m_1 and m_2 . The treatments for the optimal design under the 0-1 utility result in two distinct profiles of $u_1(t)$. For placentas 1–5, $u_1(t)$ has a slow steady increase in $u_1(t)$ with respect to t . Placentas 6 and 7 have a steep initial increase and subsequent decrease in $u_1(t)$ with respect to t . Unlike the other optimal designs, the observation times are predominantly towards the end of the observation window. The $u_1(t)$ profiles are similar under both models, with the most substantial differences occurring in the inter-profile variability towards the middle of the time interval. This region is where the majority of observation times are located.

The original design proposed by the experimenters had an unequal spacing of observation times across the entire interval $[0, 600]$. There are more observations taken near the start of the interval, and the time points are not dissimilar to those in the optimal designs under NSEL, NAEL and 0-1 estimation. However, the original design has treatments that are very different from any of the optimal designs, with an almost factorial structure and some treatments with high values of x_2 (exterior initial concentration of non-radioactive serine). None of the optimal designs include treatments with high x_2 , demonstrating how it is often difficult to predict by intuition the treatments in a Bayesian optimal design for a complicated nonlinear model. In addition, the designs for point estimation (under NSEL, NAEL and 0-1 estimation utilities) are quite different to the design for model selection.

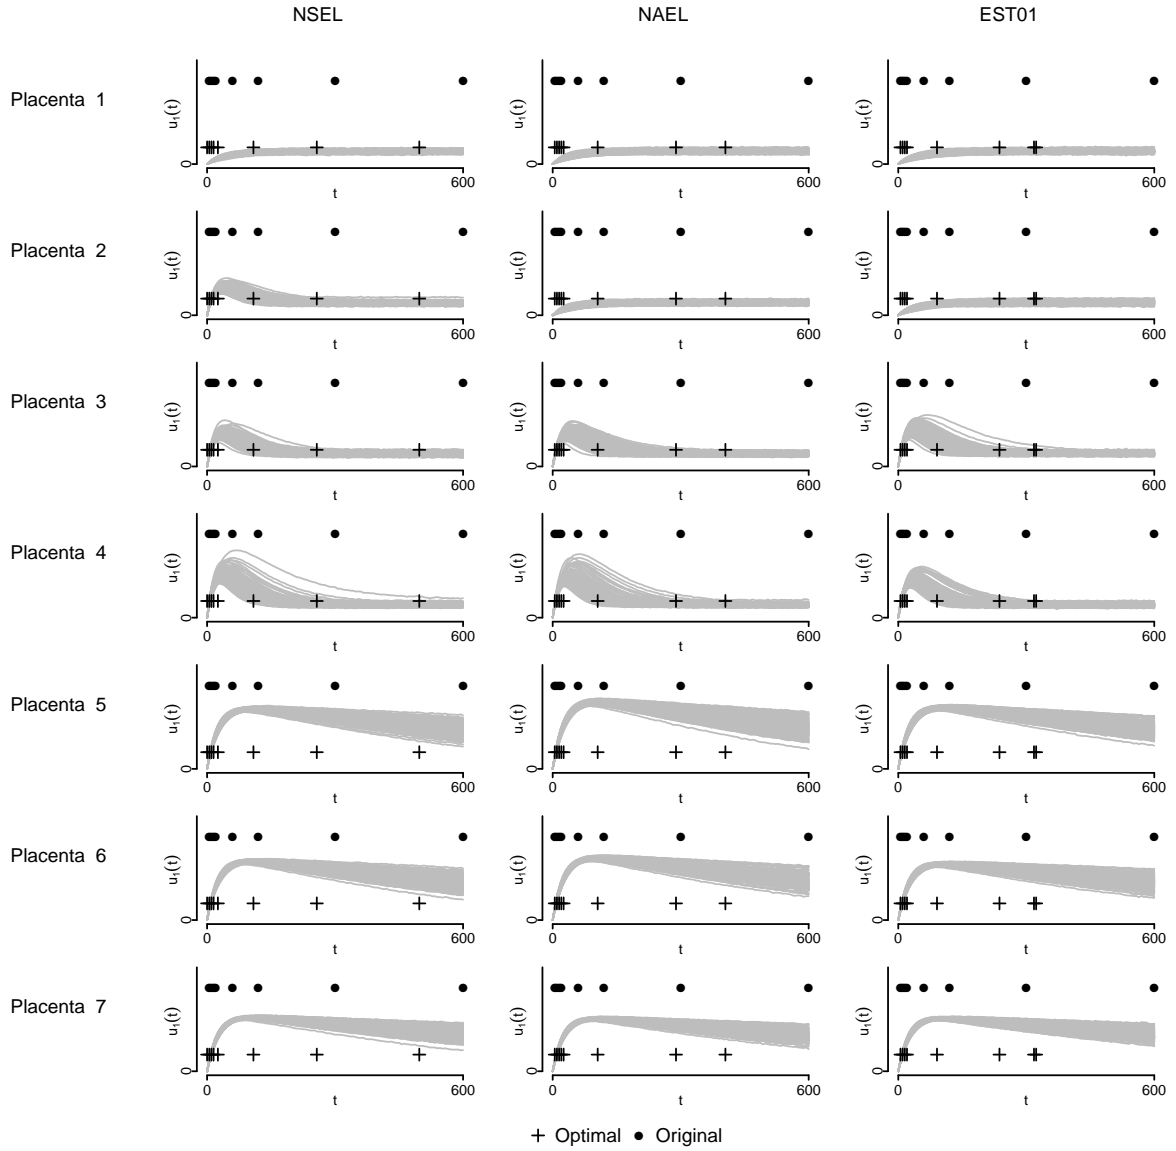


Figure 6: Results from the designs found under SIG, NSEL and NAEL utilities with $n = 7$ placentas in Section 5. Displayed are 100 draws from solution $u_1(t)$ plotted against t for values drawn from the the prior distribution of θ , for each of the $n = 7$ placentas and treatments given in Table 1.

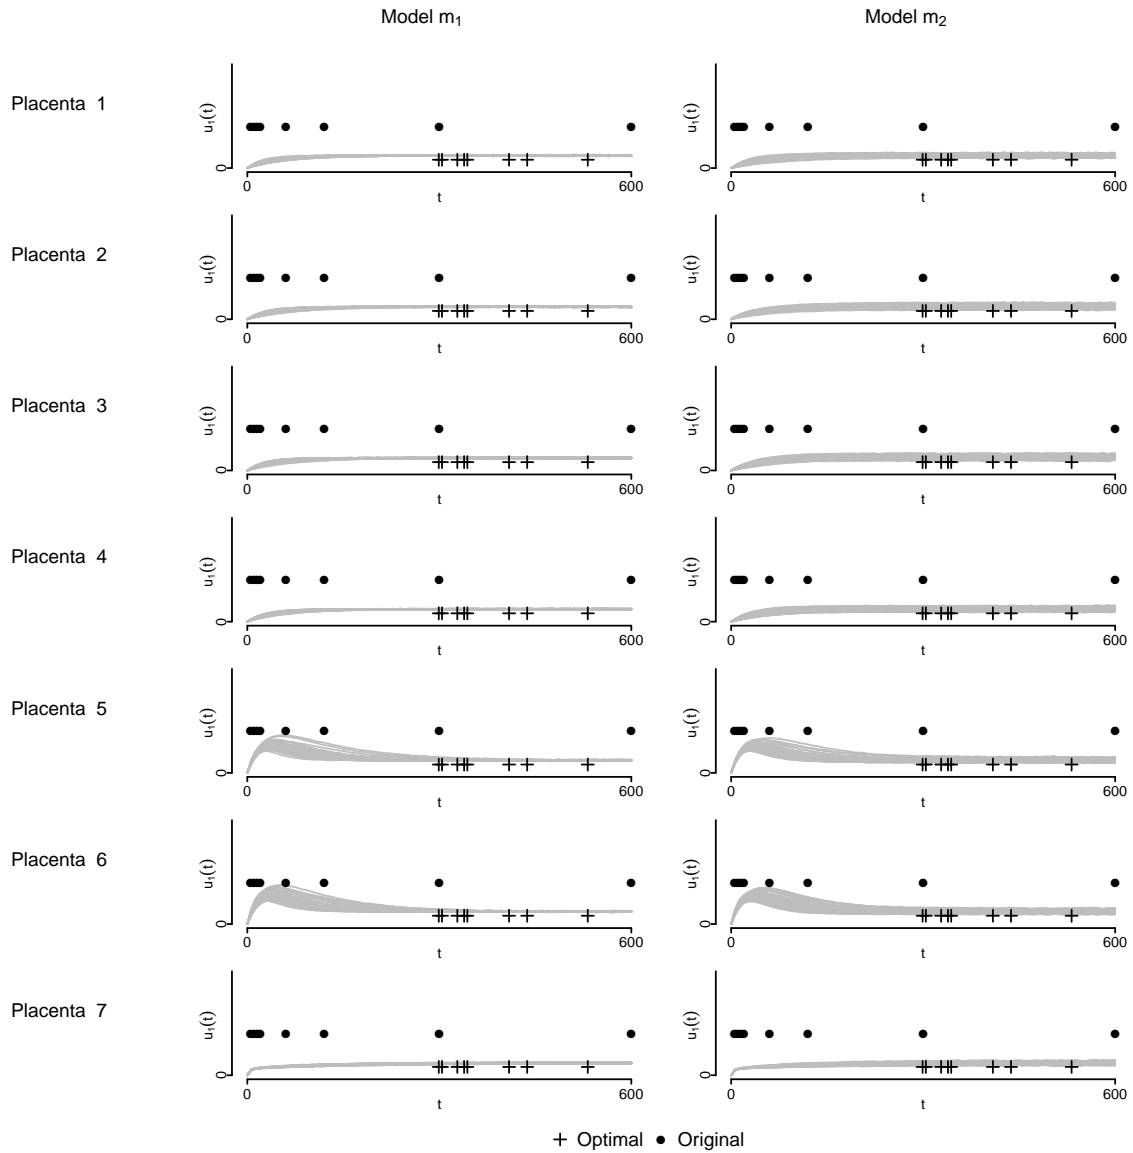


Figure 7: Results from the designs under the 0-1 model selection loss with $n = 7$ placentas in Section 5. Displayed are 100 draws from solution $u_1(t)$ under model m_1 ($\theta_3 = \theta_4$) and model m_2 ($\theta_3 \neq \theta_4$) plotted against t for values drawn from the the prior distribution of θ , for each of the $n = 7$ placentas and treatments given in Table 1.

6 Concluding Remarks

This paper introduces and demonstrates the first practical methodology for Bayesian optimal design of experiments for statistically nonlinear models formed from the solution to intractable ordinary differential equations. The work is motivated by a challenging design problem from the biological sciences, which we address through a combination of probabilistic solutions to the equations, simulation-based approximation to expected utilities and optimization via smoothing and cyclic descent. Our novel adjustments to the Chkrebtii et al. (2016) probabilistic algorithm are key to providing a computationally efficient solution to the optimal design problem. Through demonstration on a number of examples, including the motivating experiment on serine transport across placental membranes, we show the efficiency gains that can be made by use of optimal designs over obvious, and proposed, alternatives. We also show how it is often not possible to “second guess” via intuition the solutions to optimal problems for nonlinear models.

We have adopted the nested integration and optimization methods from Overstall and Woods (2017) to find optimal designs for the differential equation models in this paper (namely the ACE algorithm). The Markov chain simulation schemes of Müller (1999) and Müller et al. (2004), among other authors, would be an interesting alternative approach. Extension and application of such methods to the problems in the current paper is an area for future research.

One key issue not addressed is model discrepancy (see, e.g., Kennedy and O’Hagan, 2001 and Plumlee, 2017); the systematic mis-match between the true physical process and the solution to the ordinary differential equations. Not taking account of this error can lead to significant bias in posterior estimates of the physical parameters (Brynjarsdottir and O’Hagan, 2014). Future work will focus on Bayesian optimal design for physical models subject to model discrepancy.

Some limited insight into the impact of model mis-match can be gained from a simple extension to the compartmental model in Section 4.2. For the purpose of finding designs, we assume the model

$$u_2(t) = \frac{D}{0.9\theta_3} (\exp(-\theta_1 t) - \exp(-\theta_2 t)) . \quad (11)$$

That is, we simplify (10) by setting θ_1 and θ_2 equal to their prior means in the fraction that




Figure 8: Results from the misspecified model example in Section 6. Boxplots of 20 evaluations of the Monte Carlo approximation to the expected utility for the optimal designs found under the correct model (10) and the misspecified model (11) under the SIG, NSEL and NAEL utility functions. In each case, the correct model is assumed for evaluating the expected utility.

multiplies the exponential term. We still assume the exponential depends on unknown θ_1, θ_2 , and assume the same prior distributions for all parameters as in Section 4.2.

To assess the impact of model mis-match, we find optimal designs under the SIG, NSEL and NAEL utilities assuming the misspecified model (11). We then assess these designs under the correct, more complex, model (10) and compare them to designs found under the correct model by evaluating the approximate expected utility under the correct model, see Figure 8. Differences in approximate expected utility between the designs found under the correct and misspecified models are comfortably within Monte Carlo error for SIG. However, assuming a misspecified model under the NSEL utility results in a loss of expected utility of around 6%; the differences are somewhat less for NAEL but still larger than Monte Carlo error. Clearly, the reduction in expected utility from assuming a misspecified model will depend on the models under consideration, the difference between the models and the choice of prior distributions, in addition to the choice of utility function. This is an important area for future research.

Acknowledgements

We thank Rohan Lewis, Bram Sengers and Nont Panitchob from Medicine, Life Sciences and Bioengineering at the University of Southampton for providing details of the placenta experiment. We are grateful to two anonymous reviewers whose comments improved the paper.

Funding

The second author was supported by Fellowship EP/J018317/1 from the United Kingdom Engineering and Physical Sciences Research Council.

A The ACE algorithm

In Algorithm 4, we outline the basic approximate coordinate exchange algorithm. For full details, see Overstall and Woods (2017).

Let M be the total number of coordinates (values taken by each variable in each run) of the design. That is, M is the dimension of the design. In step 9, the probability of the suggested design having higher expected utility is calculated, a posteriori to two independent Monte Carlo samples from the joint distributions of the data and parameters conditional on the current and suggested designs. Calculation of this probability assumes the utility evaluations are well described by a normal distribution. In the case of 0-1 utilities, a similar test based on a Bernoulli likelihood and Beta prior is applied (see Overstall et al., 2018a for details). Convergence in step 11 is assessed informally using trace plots of the evaluations of either $\bar{\phi}^*$, if the proposed design was accepted, or $\bar{\phi}^C$, otherwise, from step 9.

The ACE algorithm should be started from multiple different starting designs \mathbf{d}^0 . From the resulting designs, the one with the lowest value of $\hat{\Phi}(\mathbf{d})$ should be returned.

Algorithm 4: The approximate coordinate exchange (ACE) algorithm.

- 1 Choose an initial design $\mathbf{d}^0 = (d_1^0, \dots, d_M^0)^\top$ and set the current design to be
 $\mathbf{d}^C = (d_1^C, \dots, d_M^C)^\top = \mathbf{d}^0$
 - 2 **for** $i = 1 : M$ **do**
 - 3 Generate a one-dimensional space-filling design $\zeta_i = \{d_i^1, \dots, d_i^R\}$ in $\mathcal{D}_i \subset \mathbb{R}$, the set of possible values for the i th coordinate
 - 4 Let $\mathbf{d}^C(d_i^r)$ equal \mathbf{d}^C with i th coordinate replaced by d_i^r
 for $r = 1 : R$ **do**
 | Evaluate $\hat{\Phi}(\mathbf{d}^C(d_i^r))$, the approximation to the expected utility, i.e. equation (5)
 - 5 Fit a Gaussian process emulator $\tilde{\Phi}(d)$ using “data” $\{d_i^r, \hat{\Phi}(\mathbf{d}^C(d_i^r))\}_{r=1}^R$
 - 6 Set $\mathbf{d}^* = (d_1^C, \dots, d_{i-1}^C, \tilde{d}, d_{i+1}^C, \dots, d_M^C)^\top$, where $\tilde{d} \in \arg \max_{d \in \mathcal{D}_i} \tilde{\Phi}(d)$
 for $j = 1 : B$ **do**
 | Generate $[(\mathbf{y}_j^C)^\top, (\boldsymbol{\theta}_j^C)^\top]^\top \sim \pi(\boldsymbol{\theta}, \mathbf{y} | \mathbf{d}^C)$ and $[(\mathbf{y}_j^*)^\top, (\boldsymbol{\theta}_j^*)^\top]^\top \sim \pi(\boldsymbol{\theta}, \mathbf{y} | \mathbf{d}^*)$
 | Set $\phi_j^C = \phi(\boldsymbol{\theta}_j, \mathbf{y}_j, \mathbf{d}^C)$ and $\phi_j^* = \phi(\boldsymbol{\theta}_j, \mathbf{y}_j, \mathbf{d}^*)$
 - 9 Calculate

$$p^* = 1 - F_{t, 2B-2} \left(-\frac{\sum_{i=j}^B \phi_i^C - \sum_{i=1}^B \phi_j^*}{\sqrt{2B\hat{v}}} \right),$$
 where $F_{t,a}(\cdot)$ is the distribution function of the t -distribution with a degrees of freedom,

$$\hat{v} = \frac{\sum_{j=1}^B (\phi_j^C - \bar{\phi}^C)^2 + \sum_{j=1}^B (\phi_j^* - \bar{\phi}^*)^2}{2B - 2},$$
 and $\bar{\phi}^C = \sum_{j=1}^B \phi_j^C / B$ and $\bar{\phi}^* = \sum_{j=1}^B \phi_j^* / B$
 - 10 Set $\mathbf{d}^C = \mathbf{d}^*$ with probability p^*
 - 11 Return to step 2 until convergence.
-

References

- Atkinson, A. C. and Bogacka, B. (2002), “Compound and other optimum designs for systems of nonlinear differential equations arising in chemical kinetics,” *Chemometrics and Intelligent Laboratory Systems*, 61, 17–33.
- Atkinson, A. C., Chaloner, K., Herzberg, A. M., and Juritz, J. (1993), “Experimental designs for properties of a compartmental model,” *Biometrics*, 49, 325–337.
- Brynjarsdottir, J. and O’Hagan, A. (2014), “Learning about physical parameters: the importance of model discrepancy,” *Inverse Problems*, 30, 114007.
- Chaloner, K. and Verdinelli, I. (1995), “Bayesian experimental design: a review,” *Statistical Science*, 10, 273–304.
- Chkrebtii, O. A., Campbell, D. A., Calderhead, B., and Girolami, M. A. (2016), “Bayesian uncertainty quantification for differential equations,” *Bayesian Analysis*, 11, 1239–1267.
- FitzHugh, R. (1961), “Impulses and physiological states in models of nerve membrane,” *Biophysical Journal*, 1, 445–466.
- Franceschini, G. and Macchietto, S. (2008), “Model-based design of experiments for parameter precision: state of the art,” *Chemical Engineering Science*, 63, 4846–4872.
- Holsclaw, T., Sansó, B., Lee, H. K. H., Heitmann, K., Habib, S., Higdon, D., and Alam, U. (2013), “Gaussian process modeling of derivative curves,” *Technometrics*, 55, 57–67.
- Iserles, A. (2009), *A First Course in the Numerical Analysis of Differential Equations*, Cambridge University Press.
- Kennedy, M. C. and O’Hagan, A. (2001), “Bayesian calibration of computer models (with discussion),” *Journal of the Royal Statistical Society B*, 63, 425–464.
- Lange, K. (2013), *Optimization*, New York: Springer, 2nd ed.
- Meyer, R. K. and Nachtsheim, C. J. (1995), “The coordinate-exchange algorithm for constructing exact optimal experimental designs,” *Technometrics*, 37(1), 60–69.

- Müller, P. (1999), “Simulation-based optimal design,” in *Bayesian Statistics 6*, eds. Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M., Oxford: Oxford University Press, pp. 459–474.
- Müller, P. and Parmigiani, G. (1996), “Optimal design via curve fitting of Monte Carlo experiments,” *Journal of the American Statistical Association*, 90, 1322–1330.
- Müller, P., Sanso, B., and De Iorio, M. (2004), “Optimal Bayesian design by inhomogeneous Markov chain simulation,” *Journal of the American Statistical Association*, 99, 788–798.
- Nagumo, J., Arimoto, S., and Joshizawa, S. (1962), “An active pulse transmission line simulating a nerve axon,” *Proceedings of the Institute of Radio Engineers*, 50, 2061–2070.
- Overstall, A. M., McGree, J. M., and Drovandi, C. C. (2018a), “Fully Bayesian design using the approximate coordinate exchange algorithm and normal-based approximations to posterior quantities,” *Statistics and Computing*, 28, 343–358.
- Overstall, A. M. and Woods, D. C. (2017), “Bayesian design of experiments using approximate coordinate exchange,” *Technometrics*, 59, 458–470.
- Overstall, A. M., Woods, D. C., and Adamou, M. (2018b), “acebayes: An R package for Bayesian optimal design of experiments via approximate coordinate exchange,” *arXiv:1705.08096*.
- (2018c), *acebayes: Optimal Bayesian experimental design using the ACE algorithm*, R package version 1.6.0.
- Panitchob, N., Widdows, K. L., Crocker, I. P., Hanson, M. A., Johnstone, E. D., Please, C. P., Sibley, C. P., Glazier, J. D., Lewis, R. M., and Sengers, B. G. (2015), “Computational modelling of amino acid exchange and facilitated transport in placental membrane vesicles,” *Journal of Theoretical Biology*, 365, 352–364.
- Pellegrini, S. and Dusanter-Fourt, I. (1997), “The structure, regulation and function of the Janus kinases (JAKs) and the signal transducers and activators of transcription (STATs),” *European Journal of Biochemistry*, 248, 615–633.

- Plumlee, M. (2017), “Bayesian calibration of inexact computer models,” *Journal of the American Statistical Association*, 112, 1274–1285.
- Ramsay, J., Hooker, G., Campbell, D., and Cao, J. (2007), “Parameter estimation for differential equations: a generalised smoothing approach (with discussion),” *Journal of the Royal Statistical Society B*, 69, 741–796.
- Rasmussen, C. E. and Williams, C. K. I. (2006), *Gaussian Processes for Machine Learning*, Cambridge, MA.: MIT Press.
- Raue, A., Kreutz, C., Maiwald, T., Bachmann, J., Schilling, M., Klingmüller, U., and Timmer, J. (2009), “Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood,” *Bioinformatics*, 25, 1923–1929.
- Rodríguez-Díaz, J. M. and Sánchez-León, G. (2014), “Design optimality for models defined by a system of ordinary differential equations,” *Biometrical Journal*, 56, 886–900.
- Ryan, E., Drovandi, C., Thompson, M., and Pettitt, A. (2014), “Towards Bayesian experimental design for nonlinear models that require a large number of sampling times,” *Computational Statistics and Data Analysis*, 70, 45–60.
- Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. (2016), “A review of modern computational algorithms for Bayesian optimal design,” *International Statistical Review*, 84, 128–154.
- Ryan, K. J. (2003), “Estimating expected information gains for experimental designs with application to the random fatigue-limit model,” *Journal of Computational and Graphical Statistics*, 12, 585–603.
- Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989), “Design and analysis of computer experiments (with discussion),” *Statistical Science*, 4, 409–435.
- Santner, T. J., Williams, B. J., and Notz, W. I. (2003), *The Design and Analysis of Computer Experiments*, New York: Springer.
- Solak, E., Murray-Smith, R., Leithead, W. E., Leith, D. J., and Rasmussen, C. E. (2003), “Derivative observations in Gaussian process models of dynamic systems,” in *Advances in*

Neural Information Processing Systems, eds. Becker, S., Thrun, S., and Obermayer, K., MIT Press, vol. 15, pp. 1057–1064.

Swameye, I., Muller, T., Timmer, J., Sandra, O., and Klingmuller, U. (2003), “Identification of nucleocytoplasmic cycling as a remote sensor in cellular signaling by databased modeling,” *Proceedings of the National Academy of Sciences*, 100, 1028–1033.

Valko, P. and Vajda, S. (1984), “An extended ODE solver for sensitivity calculations,” *Computers and Chemistry*, 8, 255–271.

Widdows, K. L., Panitchob, N., Crocker, I. P., Please, C. P., Hanson, M. A., Sibley, C. P., Johnstone, E. D., Sengers, B. G., Lewis, R. M., and Glazier, J. D. (2017), “Integration of computational modeling with membrane transport studies reveals new insights into amino acid exchange transport mechanisms,” *The FASEB Journal*, in press.

Woods, D. C., Overstall, A. M., Adamou, M., and Waite, T. W. (2017), “Bayesian design of experiments for generalised linear models and dimensional analysis with industrial and scientific application (with discussion),” *Quality Engineering*, 29, 91–118.