Properties of Fisher information gain for Bayesian design of experiments

Antony M. Overstall

Southampton Statistical Sciences Research Institute,
University of Southampton,
Southampton SO17 1BJ UK
(A.M.Overstall@southampton.ac.uk)

Abstract

The Bayesian decision-theoretic approach to design of experiments involves specifying a design (values of all controllable variables) to maximise the expected utility function (expectation with respect to the distribution of responses and parameters). For most common utility functions, the expected utility is rarely available in closed form and requires a computationally expensive approximation which then needs to be maximised over the space of all possible designs. This hinders practical use of the Bayesian approach to find experimental designs. However, recently, a new utility called Fisher information gain has been proposed. The resulting expected Fisher information gain reduces to the prior expectation of the trace of the Fisher information matrix. Since the Fisher information is often available in closed form, this significantly simplifies approximation and subsequent identification of optimal designs. In this paper, it is shown that for exponential family models, maximising the expected Fisher information gain is equivalent to maximising an alternative objective function over a reduced-dimension space, simplifying even further the identification of optimal designs. However, if this function does not have enough global maxima, then designs that maximise the expected Fisher information gain lead to non-identifiablility.

Keywords: Bayesian design of experiments; non-identifiability; parameter redundancy; under-supported designs

1 Introduction

Suppose an experiment is to be conducted to learn the relationship between a series of controllable variables and a measurable response for some physical system. The experiment consists of a number of runs where each run involves specification of a series of controllable variables and subsequent observation of a response. Typically, on completion of the experiment, it is assumed that the observed responses are realisations of a random variable with a fully specified probability distribution apart from a vector of unknown parameters. That is, a statistical model is assumed for the responses. The observed responses are used to estimate the parameters thus allowing estimation of the relationship between controllable variables and response.

A question of interest is how should the controllable variables for all runs (i.e. the design) be specified to facilitate as precise as possible estimation of the unknown parameters. The Bayesian decision-theoretic approach (see, for example, Chaloner and Verdinelli, 1995) is to select the design to maximise the expectation of a utility function. Expectation is taken with respect to the joint distribution of all unknown quantities, i.e. the responses and the parameters, with this joint distribution following from the specification of the statistical model and a prior distribution for the parameters. The advantages of this approach are as follows. (1) The choice of utility allows bespoke experimental aims to be incorporated (for example, experimental aims of prediction and/or model selection can be considered instead of estimation although these are not the focus of this paper). (2) By taking expectation with respect to unknown quantities, all known sources of uncertainty are incorporated. (3) The framework fits into the iterative nature in which knowledge is accumulated in science, i.e. the distribution of unknown quantities for the current iteration is based on the

results of the previous iteration. With regards to point (3) above, the approach can even be extended to account for the fact that the expected utility will be maximised in subsequent iterations (see, for example, Huan and Marzouk, 2016), i.e. so-called Bayesian sequential or adaptive design. However this extension is not considered in this paper.

The choice of utility allows different ways of specifying estimation precision. Ideally, the choice of utility should be uniquely tailored to the experiment. However, a more pragmatic choice is often made from certain default utility functions. Two common default utilities for estimation precision are Shannon information gain (Lindley, 1956) and negative squared error loss (see, for example, Robert, 2007, pages 77-79). In either case, typically, the utility function depends on the responses through the posterior distribution of the parameters. Even for relatively simple statistical models, the posterior distribution is not of a known form meaning that the utility and expected utility are not available in closed form and require approximation.

Recently, Walker (2016) instead proposed the Fisher information gain utility. This is the difference between the squared Euclidean length of the gradient of the log posterior and log prior densities of the parameters. The appealing feature of Fisher information gain is that the resulting expected utility is the prior expectation of the trace of the Fisher information matrix. For many models, the Fisher information matrix is available in closed form, significantly simplifying approximation of the expected utility and subsequent identification of an optimal design.

However, in this paper it is shown that, under a broad class of models, a design that maximises the expected Fisher information gain can lead to a non-identifiable model. Under the classical approach to statistical inference, this means the parameters are not uniquely estimable (see, for example, Catchpole and Morgan, 1997). Under the Bayesian approach, it means that the posterior and prior distributions for a subset of parameters (conditional on the complement of the subset) are identical (see, for example, Gelfand and Sahu, 1999). These are undesirable consequences for a default utility function.

The paper is organised as follows. Section 2 provides a background to Bayesian decision-theoretic design of experiments, introduces the Fisher information gain utility, describes the exponential family of models and discusses the notion of non-identifiability. Section 3 explores the consequences of designing experiments under Fisher information gain with examples provided in Section 4.

2 Background

2.1 Bayesian decision-theoretic design of experiments

Suppose there are k controllable variables denoted $\mathbf{d} = (d_1, \ldots, d_k)^{\mathrm{T}} \in \mathcal{D}$, where $\mathcal{D}$ denotes a set of possible values for the controllable variables, and let y denote the response. Suppose the experiment consists of a fixed number n of runs. For the ith run, the controllable variables are denoted $\mathbf{d}_i = (d_{i1}, \ldots, d_{ik})^{\mathrm{T}}$ and the response by y_i, for $i = 1, \ldots, n$. Let $\mathbf{y} = (y_1, \ldots, y_n)^T$ denote the $n \times 1$ vector of responses and $\mathbf{D} = (\mathbf{d}_1^{\mathrm{T}}, \ldots, \mathbf{d}_n^{\mathrm{T}})^{\mathrm{T}} \in \mathcal{D}^n$ denote the $nk \times 1$ vector of controllable variables for all runs (termed the design). The vectors in the set $\{\mathbf{d}_1, \ldots, \mathbf{d}_n\}$ are referred to as designs points and the distinct vectors as support points.

Throughout it is assumed that the elements of $\mathbf{y}$ are independent and $y_i \sim \mathrm{F}(\boldsymbol{\theta}, \mathbf{d}_i)$ where $\mathrm{F}(\boldsymbol{\theta}, \mathbf{d})$ denotes a probability distribution depending on a $p \times 1$ vector of unknown parameters $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_p)^T \in \Theta$ and controllable variables $\mathbf{d}$, with Θ denoting a p-dimensional parameter space. It is assumed that $p \leq n$.

Bayesian decision-theoretic design of experiments (see, for example, Chaloner and Verdinelli, 1995) proceeds by specifying a utility function denoted $u(\boldsymbol{\theta}, \mathbf{y}, \mathbf{D})$. The utility gives the precision in estimating $\boldsymbol{\theta}$ given responses $\mathbf{y}$ obtained via design $\mathbf{D}$, where different choices of utility allow different ways of measuring precision. Typically the utility depends on the responses through the posterior distribution of $\boldsymbol{\theta}$ given by Bayes' theorem $\pi(\boldsymbol{\theta}|\mathbf{y}, \mathbf{D}) \propto \pi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{D})\pi(\boldsymbol{\theta})$, where $\pi(\boldsymbol{\theta})$ is the probability density function (pdf) of the prior distribution of $\boldsymbol{\theta}$ (assumed independent of the design $\mathbf{D}$) and $\pi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{D}) = \prod_{i=1}^{n} \pi(y_i|\boldsymbol{\theta}, \mathbf{d}_i)$ is the likelihood function with $\pi(y_i|\boldsymbol{\theta}, \mathbf{d}_i)$ the pdf or mass function of $F(\boldsymbol{\theta}, \mathbf{d}_i)$.

Let
$$U(\mathbf{D}) = \mathrm{E}_{\mathbf{y}, \boldsymbol{\theta}|\mathbf{D}} \left[u(\boldsymbol{\theta}, \mathbf{y}, \mathbf{D}) \right],$$

be the expected utility where expectation is with respect to the joint distribution of responses $\mathbf{y}$ and parameters $\boldsymbol{\theta}$. Assuming that $U(\mathbf{D})$ exists, an optimal Bayesian design maximises $U(\mathbf{D})$ over the space of all designs $\mathcal{D}^n$, where this design is not necessarily unique.

It should be noted that the use of Bayesian methods to determine a design does not restrict the post-experiment analysis to be Bayesian (see, for example, Atkinson et al., 2007, Chapter 18).

Ideally, the utility should be chosen to represent the experimental aim. Instead, often a more pragmatic approach is taken where the utility is chosen from certain default utility functions. These utilities aim to represent standard aims such as precision of point estimation or gain in information from prior to posterior. Common default choices for the utility are Shannon information gain (SIG) $u_{SIG}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{D}) = \log \pi(\boldsymbol{\theta}|\mathbf{y}, \mathbf{D}) - \log \pi(\boldsymbol{\theta})$ and negative squared error loss (NSEL) $u_{NSEL}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{D}) = -\|\boldsymbol{\theta} - \mathrm{E}_{\boldsymbol{\theta}|\mathbf{y}, \mathbf{D}}(\boldsymbol{\theta})\|$, where $\|\mathbf{v}\| = \mathbf{v}^{\mathrm{T}}\mathbf{v}$ is the squared Euclidean length of vector $\mathbf{v}$. A SIG-optimal design that maximises the expected SIG utility equivalently maximises the expected Kullback-Liebler divergence between the posterior and prior distributions of $\boldsymbol{\theta}$, where expectation is with respect to the marginal distribution of $\mathbf{y}$. A NSEL-optimal design equivalently minimises the trace of the expected posterior variance matrix of $\boldsymbol{\theta}$.

The dependence of the utility function on the posterior distribution of $\boldsymbol{\theta}$, and the fact that this distribution is typically not of known form for non-trivial statistical models, is a significant hurdle to the practical implementation of Bayesian design of experiments. For example, under SIG and NSEL, the utility function is not available in closed form, not to mention the expected SIG and NSEL utilities. Even with the ongoing development of new methodology (see, for example, Ryan et al., 2016; Woods et al., 2017) to approximate and maximise the expected utility, finding designs in practice is a computationally expensive task.

2.2 Fisher information gain

Walker (2016) recently proposed the Fisher information gain (FIG) utility

$$u_{FIG}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{D}) = \left\| \frac{\partial \log \pi(\boldsymbol{\theta}|\mathbf{y}, D)}{\partial \boldsymbol{\theta}} \right\| - \left\| \frac{\partial \log \pi(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right\|,$$

the difference between the squared Euclidean length of the $p \times 1$ gradient vectors of the posterior and prior log pdfs (assuming that these gradients exist). The length of the gradient vector of the log pdf of a continuous distribution provides a measure of information provided by the distribution. The FIG utility gives the gain in this information provided by the posterior compared to the prior. A design that maximises the expected Fisher information gain $U_{FIG}(\mathbf{D}) = \mathrm{E}_{\mathbf{y}, \boldsymbol{\theta}|\mathbf{D}}[u_{FIG}(\boldsymbol{\theta}, \mathbf{y}, \mathbf{D})]$ is termed FIG-optimal and denoted $\tilde{\mathbf{D}}$.

Under standard regularity conditions (see, for example, Schervish, 1995, page 111), it can be shown (Walker, 2016) that the expected FIG utility can be written

$$U_{FIG}(\mathbf{D}) = \mathrm{E}_{\boldsymbol{\theta}}\left[\mathrm{tr}\left\{\mathcal{I}(\boldsymbol{\theta}, \mathbf{D})\right\}\right],$$

where

$$\mathcal{I}(\boldsymbol{\theta}, \mathbf{D}) = \mathrm{E}_{\mathbf{y}|\boldsymbol{\theta}, \mathbf{D}}\left[\frac{\partial \log \pi(\mathbf{y}|\boldsymbol{\theta}, D)}{\partial \boldsymbol{\theta}} \frac{\partial \log \pi(\mathbf{y}|\boldsymbol{\theta}, D)}{\partial \boldsymbol{\theta}^{\mathrm{T}}}\right]$$

is the $p \times p$ Fisher information matrix. That is, the expected FIG utility is the prior expectation of the trace of the Fisher information matrix where it is assumed that both the Fisher information, and the prior expectation of its trace, exists.

The Fisher information matrix is available in closed form for many classes of model (for example, generalised linear models) for which the corresponding posterior distribution (under any prior distribution) is not of known form. Therefore, approximating the expected FIG utility is a significantly simpler task than, for example, approximating the expected SIG or NSEL utilities. This represents a major advantage in the use of FIG over other utility functions.

Furthermore, the expected FIG appears to parallel common objective functions for pseudo-Bayesian classical design criteria, e.g. D- and A-optimality (see, for example, Atkinson et al., 2007, Chapter 10), where in each case, the objective function is given as the prior expectation of a scalar function of the Fisher

information matrix. For D-optimality the scalar function is the log-determinant and, for A-optimality, it is the negative trace of the inverse. Under these classical criteria, the Fisher information appears since its inverse is an asymptotic approximation to the variance of the maximum likelihood estimator of $\boldsymbol{\theta}$. To overcome dependence on $\boldsymbol{\theta}$, the scalar function of $\mathcal{I}(\boldsymbol{\theta}, D)$ is averaged with respect to a prior distribution on $\boldsymbol{\theta}$. The term "pseudo-Bayesian" refers to the use of Bayesian machinery for a classical procedure. However, the objective functions for D- and A-optimality can also be derived as asymptotic approximations to the expected utility under the SIG and NSEL utilities, respectively (see, for example, Chaloner and Verdinelli, 1995). As pointed out by Walker (2016), the FIG utility is fully Bayesian and does not result from an asymptotic approximation. For the objective function for D- and A-optimality to be greater than $-\infty$, the inverse of the Fisher information must exist. There is no such requirement for the expected FIG.

2.3 Exponential family models

We now present a broad class of models for which the results in this paper apply. It is assumed that the elements of $\mathbf{y}$ are distributed according to a member of the exponential family of distributions (see, for example Kosmidis et al., 2020). This family of distributions include the normal, Poisson, binomial and gamma. The pdf or mass function of $F(\boldsymbol{\theta}, \mathbf{d}_i)$ is

$$\pi(y_i | \zeta_i, \gamma) = \exp\left(\frac{y_i \zeta_i - b(\zeta_i) - c_1(y_i)}{\gamma} - \frac{1}{2} a\left(-\frac{1}{\gamma}\right) + c_2(y_i) \right),$$

for $i = 1, \ldots, n$, where $a(\cdot)$, $b(\cdot)$, $c_1(\cdot)$ and $c_2(\cdot)$ are sufficiently smooth specified functions, $\gamma > 0$ is a common scale parameter and $\zeta_1, \ldots, \zeta_n$ are natural parameters, depending on $\boldsymbol{\theta}$ and $\mathbf{d}_1, \ldots, \mathbf{d}_n$, respectively. It follows that $\mathrm{E}_{y_i | \zeta_i, \gamma}(y_i) = \mathrm{E}_{y_i | \boldsymbol{\theta}, \mathbf{d}_i, \gamma}(y_i) = b'(\zeta_i)$ (see, for example Kosmidis et al., 2020), where it is assumed that $\mathrm{E}_{y_i | \zeta_i, \gamma}(y_i) = \mathrm{E}_{y_i | \boldsymbol{\theta}, \mathbf{d}_i, \gamma}(y_i) = \mu(\boldsymbol{\theta}, \mathbf{d}_i)$: a differentiable mean function of parameters $\boldsymbol{\theta}$ and controllable variables $\mathbf{d}_i$. Additionally, $\mathrm{var}_{y_i | \zeta_i, \gamma}(y_i) = \mathrm{var}_{y_i | \boldsymbol{\theta}, \mathbf{d}_i, \gamma}(y_i) = \gamma b''(\zeta_i)$.

The distribution of y_i is uniquely determined by $\mu(\boldsymbol{\theta}, \mathbf{d}_i)$ through the natural parameter ζ_i. This means that the difference in distribution of elements y_i and y_j is entirely determined by $\mathbf{d}_i$ and $\mathbf{d}_j$, for $i \neq j$.

The decision to focus on exponential family models is due to 1) a large number of statistical models used to analyse experimental responses belong to this family; 2) the Fisher information matrix exists and is available in closed form for models in this family making the use of Fisher information gain particularly attractive; and 3) the existence of results to determine whether or not an exponential family model is identifiable (see Section 2.4).

If the scale parameter γ is known, the Fisher information is the $p \times p$ matrix

$$\mathcal{I}(\boldsymbol{\theta}, \mathbf{D}) = M(\boldsymbol{\theta}, \mathbf{D})^{\mathrm{T}} W(\boldsymbol{\theta}, \mathbf{D}) M(\boldsymbol{\theta}, \mathbf{D})$$

(see, for example Kosmidis et al., 2020) where $M(\boldsymbol{\theta}, \mathbf{D})$ is an $n \times p$ matrix with ijth element given by

$$M_{ij}(\boldsymbol{\theta}, \mathbf{D}) = \frac{\partial \mu(\boldsymbol{\theta}, \mathbf{d}_i)}{\partial \theta_j}, \tag{1}$$

and $W(\boldsymbol{\theta}, \mathbf{D})$ is an $n \times n$ diagonal matrix with ith diagonal element

$$W_{ii}(\boldsymbol{\theta}, \mathbf{D}) = \frac{1}{\gamma b''(\zeta_i)} = \frac{1}{\mathrm{var}_{y_i | \zeta_i, \gamma}(y_i)} = \frac{1}{\mathrm{var}_{y_i | \boldsymbol{\theta}, \mathbf{d}_i, \gamma}(y_i)},$$

for $i = 1, \ldots, n$.

If γ is unknown and of interest, an extra row and column can be included in the Fisher information matrix. However, it is arguably more realistic to allow γ to be unknown but not of direct interest, i.e. it is a nuisance parameter. The correct approach here is to obtain the marginal likelihood by marginalising over γ with respect to its prior distribution, i.e.

$$\pi(\mathbf{y} | \boldsymbol{\theta}, \mathbf{D}) = \int_0^\infty \pi(\mathbf{y} | \boldsymbol{\theta}, \gamma, \mathbf{D}) \pi(\gamma) \mathrm{d}\gamma. \tag{2}$$

4

where $\pi(\mathbf{y}|\boldsymbol{\theta}, \gamma, \mathbf{D}) = \prod_{i=1}^{n} \pi(y_i|\zeta_i, \gamma)$ and $\pi(\gamma)$ is the pdf of the prior distribution of γ (assumed independent of $\boldsymbol{\theta}$). The Fisher information matrix can then be obtained from the marginal likelihood (2). However, the exact form of the marginal likelihood depends on the form of $a(\cdot)$, $b(\cdot)$, $c_1(\cdot)$ and $c_2(\cdot)$. In Section 3, we consider the special case of a normal distribution.

2.4   Non-identifiability and parameter redundancy

Catchpole and Morgan (1997) define identifiability to mean that there do not exist two parameter values (separated by at least δ under some distance measure) such that the probability distribution for the responses is identical under both parameter values. An obvious cause of non-identifiability is parameter redundancy, i.e. the statistical model can be equivalently re-expressed using a subset of the parameters. For the exponential family models of Section 2.3, Catchpole and Morgan (1997) show that the model is parameter redundant if and only if the $n \times p$ matrix $M(\boldsymbol{\theta}, \mathbf{D})$ is rank deficient, i.e. the number of linearly independent rows (or columns) of $M(\boldsymbol{\theta}, \mathbf{D})$ is less than p.

For classical inference, a consequence of parameter redundancy is that the parameters are not uniquely estimable since the likelihood can exhibit a ridge, i.e. a curve of parameter values that all maximise the likelihood function.

For Bayesian inference, the issue is more nuanced. Under a proper prior distribution, the posterior distribution for $\boldsymbol{\theta}$ will exist for a parameter redundant model. However, parameter redundancy does result in Bayesian non-identifiability (see, for example, Gelfand and Sahu, 1999). This is where $\boldsymbol{\theta}$ can be decomposed as $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^{\mathrm{T}}, \boldsymbol{\theta}_2^{\mathrm{T}})^{\mathrm{T}}$ and the posterior distribution of $\boldsymbol{\theta}_1$, conditional on $\boldsymbol{\theta}_2$, is equal to the prior distribution of $\boldsymbol{\theta}_1$, conditional on $\boldsymbol{\theta}_2$. This does not mean that no learning has taken place about $\boldsymbol{\theta}_1$, rather that learning comes indirectly through $\boldsymbol{\theta}_2$.

3   Maximising expected Fisher information gain

In this section we show that the design points of a FIG-optimal design, for exponential family models, are given by the solutions to maximising a function $\phi(\mathbf{d})$ over $\mathcal{D}$. That is, a FIG-optimal design is given by $\tilde{\mathbf{D}} = \left(\tilde{\mathbf{d}}_1, \ldots, \tilde{\mathbf{d}}_n\right)^{\mathrm{T}}$ where $\tilde{\mathbf{d}}_i$ maximises $\phi(\mathbf{d})$ over $\mathcal{D}$, for $i = 1, \ldots, n$. Ostensibly, this is appealing since an nk-dimensional maximisation problem has been converted into a k-dimensional maximisation problem. However, the number of global maxima of $\phi(\mathbf{d})$ over $\mathcal{D}$ determines the maximum number of support points for a FIG-optimal design. If this number is less than p (the number of parameters), then the design is under-supported, and we discuss how this is an undesirable feature for a default utility function. Note that for the remainder of this paper, the term maxima is used to mean global maxima.

3.1   Form of $\phi(\mathbf{d})$

3.1.1   Known scale parameter

We begin by considering the case where the scale parameter γ is known. The following lemma provides the form for $\phi(\mathbf{d})$.

Lemma 3.1. *For the exponential family models described in Section 2.3 with scale parameter γ known, then*

$$\phi(\mathbf{d}) = \sum_{j=1}^{p} \mathrm{E}_{\boldsymbol{\theta}} \left[\frac{1}{\mathrm{var}_{y|\boldsymbol{\theta}, \mathbf{d}}(y)} \left(\frac{\partial \mu(\boldsymbol{\theta}, \mathbf{d})}{\partial \theta_j} \right)^2 \right]. \tag{3}$$

The proof is given in Section SM1 of the Supplementary Material.

3.1.2 Unknown scale parameter

When the scale parameter γ is unknown and of direct interest, we consider an extended $(p+1) \times 1$ parameter vector $(\boldsymbol{\theta}^{\mathrm{T}}, \gamma)^{\mathrm{T}}$. The following results gives the form of $\phi(\mathbf{d})$ in this case.

Lemma 3.2. *For the exponential family models described in Section 2.3 with scale parameter γ unknown and of direct interest, then $\phi(\mathbf{d})$ is given by*

$$\phi(\mathbf{d}) = \sum_{j=1}^{p} \mathrm{E}_{\boldsymbol{\theta}, \gamma} \left[\frac{1}{\mathrm{var}_{y|\boldsymbol{\theta}, \gamma, \mathbf{d}}(y)} \left(\frac{\partial \mu(\boldsymbol{\theta}, \mathbf{d})}{\partial \theta_j} \right)^2 \right]. \tag{4}$$

The proof is given in Section SM2 of the Supplementary Material.

When the scale parameter γ is an unknown nuisance parameter, the situation is not as straightforward since the form of the marginal likelihood (2) will depend on the form of $a(\cdot)$, $b(\cdot)$, $c_1(\cdot)$ and $c_2(\cdot)$. We focus on the case when the member of the exponential family is the normal distribution, with the variance of y_i given by γ, for $i = 1, \ldots, n$. An inverse-gamma (IG) prior distribution for γ is assumed giving a closed form expression for the marginal likelihood. If $\gamma \sim \mathrm{IG}(s_1/2, s_2/2)$ then the pdf is

$$\pi(\gamma) = \frac{s_2^{\frac{s_1}{2}}}{2^{\frac{s_1}{2}} \Gamma\left(\frac{s_1}{2}\right)} \gamma^{-\frac{s_1}{2} - 1} \exp\left(-\frac{s_2}{2\gamma} \right)$$

for known constants s_1 (shape) and s_2 (scale). For such a model, the following lemma provides the form for $\phi(\mathbf{d})$.

Lemma 3.3. *Under a normal distribution for the response with unknown scale parameter assumed to have prior distribution $\gamma \sim \mathrm{IG}(s_1/2, s_2/2)$, then*

$$\phi(\mathbf{d}) = \sum_{j=1}^{p} \mathrm{E}_{\boldsymbol{\theta}} \left[\left(\frac{\partial \mu(\boldsymbol{\theta}, \mathbf{d})}{\partial \theta_j} \right)^2 \right]. \tag{5}$$

The proof is given in Section SM3 of the Supplementary Material.

3.2 Parameter redundancy in FIG-optimal designs

A consequence of Lemmas 3.1 to 3.3 is that to find the design points of any FIG-optimal design, we merely need to find vectors of controllable variables that maximise $\phi(\mathbf{d})$ over $\mathcal{D}$. Indeed, a FIG-optimal design exists with just one support point, i.e. $\tilde{\mathbf{D}} = \left(\tilde{\mathbf{d}}, \ldots, \tilde{\mathbf{d}} \right)^{\mathrm{T}}$ with $\tilde{\mathbf{d}}$ any vector such that $\tilde{\mathbf{d}} = \arg\max_{\mathbf{d} \in \mathcal{D}} \phi(\mathbf{d})$. The number of maxima of $\phi(\mathbf{d})$ over $\mathcal{D}$ is the maximum number of support points q of any FIG-optimal design. If $q < p$, then all FIG-optimal designs are under-supported, i.e. they have less support points than the number of unknown parameters. For such under-supported designs, the number of linearly independent rows of $M(\boldsymbol{\theta}, \mathbf{D})$ is at most $q < p$, meaning that $M(\boldsymbol{\theta}, \mathbf{D})$ is rank deficient, the model is parameter redundant and is therefore non-identifiable.

In general, for an arbitrary statistical model (i.e. specification of response distribution and mean function $\mu(\boldsymbol{\theta}, \mathbf{d})$) and prior distribution for $\boldsymbol{\theta}$, analytically determining the number of maxima of $\phi(\mathbf{d})$, and therefore, whether or not all FIG-optimal designs are under-supported, will not be possible. In these cases, a numerical search of the maxima of $\phi(\mathbf{d})$ will be necessary.

Note that even if a FIG-optimal design exists that is not under-supported, it can still result in parameter redundancy. This is because even though $\tilde{\mathbf{D}}$ has at least p support points, there still may be less than p linearly independent rows of $M(\boldsymbol{\theta}, \mathbf{D})$. This issue is demonstrated in the context of (generalised) linear models in Sections 4.1 and 4.2, where sufficient conditions are provided for parameter redundancy.

4 Examples

In this section we consider special cases of exponential family models such that Lemmas 3.1 and 3.3 can be applied, namely normal linear, Poisson, and logistic regression models and a non-linear compartmental model. These models have been frequently used in the literature to demonstrate developments of design of experiments methodology. In some cases, it is possible to determine conditions that result in parameter redundancy, and in others it is not (necessitating a numerical search).

4.1 Normal linear regression model

Design of experiments for the normal linear regression model is ubiquitous in the literature due to the applicability of such models for a wide range of applications and the availability of closed form results (see, for example, Morris, 2011). For such models, the exponential family distribution is normal and $\mu(\boldsymbol{\theta}, \mathbf{d}) = f(\mathbf{d})^{\mathrm{T}} \boldsymbol{\theta}$ for a regression function $f : \mathcal{D} \to \mathbb{R}^p$. In this section, the elements of the regression function $f(\mathbf{d}) = (f_1(\mathbf{d}), \ldots, f_p(\mathbf{d}))^{\mathrm{T}}$ are characterised as a monomial of the elements of $\mathbf{d}$, i.e.

$$ f_j(\mathbf{d}) = \prod_{r=1}^{k} d_r^{u_{rj}}, $$

for $j = 1, \ldots, p$, where u_{rj} is a non-negative integer. This general formulation allows incorporation of an intercept, main effects, interactions and powers of the controllable variables. The columns of the resulting matrix $[u]_{rj}$ should be unique. It is assumed that the row sums of $[u]_{rj}$ are positive which is equivalent to at least one element of $f(\mathbf{d})$ being dependent on each element of $\mathbf{d}$. This is a weak assumption since it would be unwise to be able to control a variable but automatically exclude it from any terms in the model.

Suppose that the design space is constrained so that $d_r \in [a_r, b_r]$, with $a_r < b_r$, for $r = 1, \ldots, k$, and $\mathcal{D} = [a_1, b_1] \otimes [a_2, b_2] \otimes \cdots \otimes [a_k, b_k]$ where $\otimes$ denotes the Cartesian product.

As an example, consider the second order response surface model in $k = 2$ controllable variables in Atkinson et al. (2007, pages 170-171). Here $f(\mathbf{d}) = \left(1, d_1, d_2, d_1^2, d_2^2, d_1 d_2\right)^{\mathrm{T}}$, corresponding to a model with an intercept, main effects and squared terms for both controllable variables and an interaction, with $p = 6$. In this case

$$ [u]_{rj} = \left(\begin{array}{cccccc} 0 & 1 & 0 & 2 & 0 & 1 \\ 0 & 0 & 1 & 0 & 2 & 1 \end{array} \right). $$

The design space is constrained such that $a_1 = a_2 = -1$ and $b_1 = b_2 = 1$, with $\mathcal{D} = [-1, 1]^2$.

Under the above setup, FIG-optimal designs can result in parameter redundancy and, therefore, non-identifiability. To see this, first consider the following result.

Proposition 4.1. *Under the normal linear model described above with an assumed inverse-gamma prior distribution for γ, any FIG-optimal design is $\tilde{\mathbf{D}} = \left(\tilde{\mathbf{d}}_1, \ldots, \tilde{\mathbf{d}}_n\right)^{\mathrm{T}}$ where $\tilde{\mathbf{d}}_i = \left(\tilde{d}_{i1}, \ldots, \tilde{d}_{ik}\right)^{\mathrm{T}}$ with*

$$ \tilde{d}_{ir} = \left\{ \begin{array}{ll} a_r & \text{if } |a_r| > |b_r|; \\ b_r & \text{if } |a_r| < |b_r|; \\ \pm|b_r| & \text{if } |a_r| = |b_r|; \end{array} \right. \tag{6} $$

for $i = 1, \ldots, n$ and $r = 1, \ldots, k$. The proof is given in Section SM4 of the Supplementary Material.

A consequence of Proposition 4.1 is that the maximum number of support points for any FIG-optimal design is $q = 2^C$ where $C = \sum_{r=1}^{k} I(|a_r| = |b_r|)$ with $I(A)$ denoting the indicator function of event A. Therefore, all FIG-optimal designs are under-supported if $p > 2^C$. Commonly, $a_r = -b_r$, for $r = 1, \ldots, k$ and all FIG-optimal designs are under-supported if $p > 2^k$.

Consider the second-order response model example above with $k = 2$ controllable variables. Since $6 = p > 2^k = q = 4$, all FIG-optimal designs for this model are under-supported resulting in parameter redundancy. Specifically, the design points are $\tilde{\mathbf{d}}_i = (\pm 1, \pm 1)^{\mathrm{T}}$.

Even if $q \geq p$, and therefore a FIG-optimal design exists which is not under-supported, all FIG-optimal designs can still result in parameter redundancy. The following result provides a sufficient condition on the form of the model which results in this outcome.

Proposition 4.2. *Under the normal linear model described above with an assumed inverse-gamma prior distribution for γ, suppose there exists $j_1, j_2 \in \{1, \ldots, k\}$, with $j_1 \neq j_2$, such that for all $r = 1, \ldots, k$, $u_{rj_1} = 0$, and there exists an $\bar{r} \in \{1, \ldots, k\}$ such that $u_{\bar{r}j_2} = 2$ and $u_{rj_2} = 0$ for $r \in \{1, \ldots, k | r \neq \bar{r}\}$. Then the model will be parameter redundant. The proof is given in Section SM5 of the Supplementary Material.*

Proposition 4.2 means that if a normal linear regression model includes an intercept and a term given by the square of any of the controllable variables, then any FIG-optimal design results in a parameter redundant model, irrespective of whether or not it is under-supported. An intuitive explanation for this result is as follows. From Proposition 4.1, it can be seen that the design points of a FIG optimal design are at the extremes of the design space. When a controllable variable is constrained to be in an interval symmetric about zero, i.e. $|a_r| = |b_r|$, to be able to uniquely estimate the parameters of a model which includes a squared term for this controllable variable (and an unknown intercept), design points in the interior of the interval are required to be able to estimate curvature.

As a particular example, suppose there is $k = 1$ controllable variable $d \in [-1, 1]$ and the model has

$$f(d) = (1, d)^{\mathrm{T}} \qquad [u]_{rj} = \begin{pmatrix} 0 & 2 \end{pmatrix},$$

with $p = 2$. All FIG-optimal designs have the form $\tilde{\mathbf{D}} = \left\{ \tilde{d}_1, \ldots, \tilde{d}_n \right\}$ with $\tilde{d}_i = \pm 1$. Now suppose that at least one design point is -1 and at least one design point is $+1$. Then any design of this form is FIG-optimal but not under-supported (since it has two support points). However, by Proposition 4.2, the model is parameter redundant. To see this, note that the two columns of the $n \times 2$ matrix $M(\boldsymbol{\theta}, \mathbf{D})$ are both $\mathbf{1}_n$ (an $n \times 1$ vector of ones), so the matrix is rank deficient.

4.2 Poisson regression model

Under the Poisson generalised linear model, the exponential family distribution is Poisson and the mean function is $\mu(\boldsymbol{\theta}, \mathbf{d}) = \exp\left(f(\mathbf{d})^{\mathrm{T}}\boldsymbol{\theta}\right)$, where $f(\mathbf{d})$ is a regression function as defined in Section 4.1. The scale parameter is $\gamma = 1$ (i.e. known) and $\text{var}(y|\boldsymbol{\theta}, \mathbf{d}) = \mu(\boldsymbol{\theta}, \mathbf{d})$. From Lemma 3.1,

$$\phi(\mathbf{d}) = \mathrm{E}_{\boldsymbol{\theta}} \left[\exp\left(f(\mathbf{d})^{\mathrm{T}}\boldsymbol{\theta}\right)\right] \sum_{j=1}^{p} f_j(\mathbf{d})^2.$$

For an arbitrary prior distribution it will not be possible to determine the number of maxima of $\phi(\mathbf{d})$.

Instead, we assume that the elements of $\boldsymbol{\theta}$ are, a-priori, independent with $\theta_j \sim \mathrm{N}\left(0, \sigma_j^2\right)$. In this case,

$$\phi(\mathbf{d}) = \exp\left(\frac{1}{2} \sum_{j=1}^{p} \sigma_j^2 f_j(\mathbf{d})^2\right) \sum_{j=1}^{p} f_j(\mathbf{d})^2.$$

If the design space is constrained so that $d_{ir} \in [a_r, b_r]$, for $i = 1, \ldots, n$, then using a similar argument to the proof of Proposition 4.1, a FIG-optimal design is $\tilde{\mathbf{D}} = \left(\tilde{\mathbf{d}}_1, \ldots, \tilde{\mathbf{d}}_n\right)^{\mathrm{T}}$ where $\tilde{\mathbf{d}}_i = \left(\tilde{d}_{i1}, \ldots, \tilde{d}_{ik}\right)^{\mathrm{T}}$ and $\tilde{d}_{ir}$ is given by (6), for $i = 1, \ldots, n$ and $r = 1, \ldots, k$. The consequences of using such a design for a normal linear regression model apply here, i.e. parameter redundancy if i) the design is under-supported when $2^C < p$ (see Proposition 4.1), and/or ii) inclusion of an intercept and squared term (see Proposition 4.2).

4.3 Logistic regression model

Consider a logistic regression model where $y_i \sim \text{Bernoulli}\left(\mu(\boldsymbol{\theta}, \mathbf{d}_i)\right)$, for $i = 1, \ldots, n$ where $\mu(\boldsymbol{\theta}, \mathbf{d}) = 1/\left(1 + \exp\left(-f(\mathbf{d})^{\mathrm{T}}\boldsymbol{\theta}\right)\right)$. The scale parameter is $\gamma = 1$ and $\text{var}(y|\boldsymbol{\theta}, \mathbf{d}) = \mu(\boldsymbol{\theta}, \mathbf{d})\left(1 - \mu(\boldsymbol{\theta}, \mathbf{d})\right)$. From

Lemma 3.1,

$$\phi(\mathbf{d}) = \mathrm{E}_{\boldsymbol{\theta}}\left[\mu(\boldsymbol{\theta}, \mathbf{d})\left(1 - \mu(\boldsymbol{\theta}, \mathbf{d})\right)\right] \sum_{j=1}^{p} f_j(\mathbf{d})^2. \tag{7}$$

Similar to the Poisson regression model in Section 4.2, under an arbitrary prior distribution for $\boldsymbol{\theta}$, it will not be possible to evaluate $\phi(\mathbf{d})$ in closed form and thus determine the number of maxima.

Instead we consider a particular example from Woods et al. (2006) which is frequently used to demonstrate design of experiments methodology (see, for example, Gotwalt et al., 2009). In this example, there are $k = 4$ controllable variables $\mathbf{d} = (d_1, \dots, d_4)^{\mathrm{T}}$. The regression function is $f(\mathbf{d}) = (1, d_1, \dots, d_4)^{\mathrm{T}}$ leading to an intercept, a main effect for each controllable variable and $p = 5$ parameters. Independent uniform prior distributions are assumed for the elements of $\boldsymbol{\theta}$ as follows

$$\begin{aligned} \theta_1 &\sim \mathrm{U}[-3, 3], & \theta_2 &\sim \mathrm{U}[4, 10], & \theta_3 &\sim \mathrm{U}[5, 11], \\ \theta_4 &\sim \mathrm{U}[-6, 0], & \theta_5 &\sim \mathrm{U}[-2.5, 3.5]. \end{aligned}$$

The elements of $\mathbf{d}$ are constrained such that $a_r = -1$ and $b_r = 1$, for $r = 1, \dots, 4$.

The prior expectation involved in evaluating the function $\phi(\mathbf{d})$ in (7) is approximated using Gauss-Legendre quadrature (see, for example, Davis and Rabinowitz, 1984, Chapter 3), i.e.

$$\hat{\phi}(\mathbf{d}) = \left[\sum_{c=1}^{C} \omega_c \mu(\mathbf{t}_c, \mathbf{d})\left(1 - \mu(\mathbf{t}_c, \mathbf{d})\right)\right] \sum_{j=1}^{p} f_j(\mathbf{d})^2,$$

where $\{\mathbf{t}_1, \dots, \mathbf{t}_C\}$ are quadrature points with weights $\{\omega_1, \dots, \omega_C\}$. The number of quadrature points is chosen to be $C = 10,000$.

To determine the number of maxima of $\phi(\mathbf{d})$, the function $\hat{\phi}(\mathbf{d})$ is numerically maximised using a quasi-Newton scheme, independently from $B = 1000$ different starting points, generated uniformly over $\mathcal{D}$. From this procedure, it is found that $\hat{\phi}(\mathbf{d})$ has only two maxima, at $\tilde{d} = \pm(-0.8701, 1, 1, 1)^{\mathrm{T}}$ (to 4 decimal places). That means that any FIG-optimal design for the above logistic regression model has a maximum of $q = 2$ support points, thus resulting in parameter redundancy.

4.4 Compartmental non-linear model

Compartmental models are frequently used in pharmacokinetics to model the uptake and excretion of a substance from an organism. Here the first-order compartmental model from Gotwalt et al. (2009) is considered. There is $k = 1$ controllable variable denoted $d \in \mathcal{D} = [0, 24]$, representing the time (in hours) from a substance first entering the organism until the response (y; the concentration of substance in the organism) is measured. The response is assumed to be normally distributed with mean function $\mu(\boldsymbol{\theta}, d) = \theta_3\left(\exp(-\theta_1 d) - \exp(-\theta_2 d)\right)$, where $\theta_2 > \theta_1$. Following Gotwalt et al. (2009), a-priori the elements of $\boldsymbol{\theta}$ are assumed independent with $\theta_1 \sim U(0.01884, 0.09884)$, $\theta_2 \sim U(0.298, 8.298)$, and θ_3 having a point mass at 21.8. Finally, the scale parameter γ (the common variance of the response) is assumed to have an inverse gamma prior distribution.

From Lemma 3.3,

$$\phi(d) = \mathrm{E}_{\boldsymbol{\theta}}\left[\sum_{j=1}^{p}\left(\frac{\partial \mu(\boldsymbol{\theta}, d)}{\partial \theta_j}\right)^2\right].$$

Similar to Section 4.3, $\phi(d)$ is approximated using Gauss-Legendre quadrature. Figure SM1 in the Supplementary Material shows the resulting approximation $\hat{\phi}(d)$ plotted against $d \in [0, 24]$. From inspection, $\hat{\phi}(d)$ has one maximum at $\tilde{d} = 24$. Therefore for a design of n sampling times, the unique FIG-optimal design is $\tilde{D} = \tilde{d}\mathbf{1}_n$. This design is under-supported leading to parameter redundancy.

5   Concluding remarks

In this paper it is shown that, for exponential family models the design points of FIG-optimal designs can be found by maximising a function $\phi(\mathbf{d})$ over $\mathcal{D}$. The advantage is that this significantly simplifies the identification of a FIG-optimal design. However, the disadvantage is that if the number of maxima of $\phi(\mathbf{d})$ is less than the number of parameters then all FIG-optimal designs will be under-supported resulting in parameter redundancy and, therefore, a non-identifiable model. Furthermore, even when a FIG-optimal design exists that is not under-supported, it can still result in parameter redundancy for (generalised) linear models on the inclusion of higher-order terms.

The conclusions from the paper are as follows.

- If the experimental aim is accurately represented by Fisher information gain, and a Bayesian analysis pursued post-experiment, then the findings are inconsequential. Under a proper prior distribution, the posterior distribution for $\boldsymbol{\theta}$ will exist for a parameter-redundant model and a FIG-optimal design maximises expected gain in information as measured by the FIG utility.

- However, if a default utility is to be considered, with the possibility of a non-Bayesian analysis post-experiment, then Fisher information gain should be used with caution.

In the latter case, it should be verified that the FIG-optimal design does not result in parameter redundancy. Alternatively, a different approach to design of experiments whilst still incorporating Fisher information gain could be employed. Prangle et al. (2020) empirically observed some of the disadvantages of FIG-optimal designs explored in this paper. To counter this, Prangle et al. (2020) developed a game theoretic approach to design of experiments which is still based on Fisher information gain, but circumvents the disadvantages.

Supplementary Material

SM1   Proof of Lemma 3.1

The expected FIG utility is

$$
\begin{aligned}
U_{FIG}(\mathbf{D}) &= \mathrm{E}_{\boldsymbol{\theta}}\left[\mathrm{tr}\left\{\mathcal{I}(\boldsymbol{\theta},\mathbf{D})\right\}\right] \\
&= \mathrm{E}_{\boldsymbol{\theta}}\left[\sum_{j=1}^{p}\sum_{i=1}^{n}\frac{1}{\mathrm{var}_{y_i|\boldsymbol{\theta},\gamma,\mathbf{d}_i}(y_i)}\left(\frac{\partial\mu(\boldsymbol{\theta},\mathbf{d}_i)}{\partial\theta_j}\right)^2\right]
\end{aligned}
$$

where the second line follows from the assumption that the elements of $\mathbf{y}$ follow an exponential family distribution. Interchanging the order of the summations and the expectation,

$$
\begin{aligned}
U_{FIG}(\mathbf{D}) &= \sum_{i=1}^{n}\sum_{j=1}^{p}\mathrm{E}_{\boldsymbol{\theta}}\left[\frac{1}{\mathrm{var}_{y_i|\boldsymbol{\theta},\gamma,\mathbf{d}_i}(y_i)}\left(\frac{\partial\mu(\boldsymbol{\theta},\mathbf{d}_i)}{\partial\theta_j}\right)^2\right] \\
&= \sum_{i=1}^{n}\phi(\mathbf{d}_i),
\end{aligned}
$$

with $\phi(\mathbf{d})$ given by (3).

SM2   Proof of Lemma 3.2

The Fisher information matrix for the extended parameter vector $(\boldsymbol{\theta}^{\mathrm{T}},\gamma)^{\mathrm{T}}$ is a $(p+1)\times(p+1)$ matrix. The $p\times p$ sub-matrix corresponding to $\boldsymbol{\theta}$ is identical to the case when γ is known. The element on the diagonal

of the Fisher information matrix corresponding to γ is $na''(-1/\gamma)/2\gamma^4$ (see, for example, Kosmidis et al., 2020). Therefore,

$$U_{FIG}(\mathbf{D}) = \mathrm{E}_{\boldsymbol{\theta},\gamma}\left[\sum_{j=1}^{p}\sum_{i=1}^{n}\frac{1}{\mathrm{var}_{y_i|\boldsymbol{\theta},\gamma,\mathbf{d}_i}(y_i)}\left(\frac{\partial\mu(\boldsymbol{\theta},\mathbf{d}_i)}{\partial\theta_j}\right)^2\right] + \mathrm{E}_{\boldsymbol{\theta},\gamma}\left[n\frac{a''\left(-\frac{1}{\gamma}\right)}{2\gamma^4}\right].$$

In the first term, the order of the summations and the expectation are interchanged. The second term does not depend on the design $\mathbf{D}$. Therefore, up to an additive constant not depending on the design,

$$\begin{aligned}
U_{FIG}(\mathbf{D}) &= \sum_{i=1}^{n}\sum_{j=1}^{p}\mathrm{E}_{\boldsymbol{\theta},\gamma}\left[\frac{1}{\mathrm{var}_{y_i|\boldsymbol{\theta},\gamma,\mathbf{d}_i}(y_i)}\left(\frac{\partial\mu(\boldsymbol{\theta},\mathbf{d}_i)}{\partial\theta_j}\right)^2\right] \\
&= \sum_{i=1}^{n}\phi(\mathbf{d}_i),
\end{aligned}$$

with $\phi(\mathbf{d})$ given by (4).

SM3   Proof of Lemma 3.3

The marginal likelihood (2) is

$$\begin{aligned}
\pi(\mathbf{y}|\boldsymbol{\theta},\mathbf{D}) &= \int\pi(\mathbf{y}|\boldsymbol{\theta},\gamma,\mathbf{D})\pi(\gamma)\mathrm{d}\gamma \\
&\propto \int\gamma^{-n/2}\exp\left[-\frac{\sum_{i=1}^{n}(y_i-\mu(\boldsymbol{\theta},\mathbf{d}_i))^2}{2\gamma}\right]\gamma^{-s_1/2-1}\exp\left(-\frac{s_2}{2\gamma}\right)\mathrm{d}\gamma
\end{aligned}$$

$$\hspace{12cm}\text{(SM3.1)}$$

$$= \frac{2^{\frac{s_1+n}{2}}\Gamma\left(\frac{s_1+n}{2}\right)}{\left[s_2+\sum_{i=1}^{n}(y_i-\mu(\boldsymbol{\theta},\mathbf{d}_i))^2\right]^{\frac{s_1+n}{2}}} \hspace{4cm}\text{(SM3.2)}$$

$$\propto \left[1+\frac{1}{s_2}(\mathbf{y}-\boldsymbol{\mu}(\boldsymbol{\theta},\mathbf{D}))^{\mathrm{T}}(\mathbf{y}-\boldsymbol{\mu}(\boldsymbol{\theta},\mathbf{D}))\right]^{-\frac{s_1+n}{2}}, \hspace{2cm}\text{(SM3.3)}$$

where $\boldsymbol{\mu}(\boldsymbol{\theta},\mathbf{D}) = (\mu(\boldsymbol{\theta},\mathbf{d}_1),\ldots,\mu(\boldsymbol{\theta},\mathbf{d}_n))^{\mathrm{T}}$. The expression in (SM3.2) follows since the integrand in (SM3.1) is proportional to the pdf of an inverse-gamma distribution with shape parameter s_1+n and scale parameter $s_2+\sum_{i=1}^{n}(y_i-\mu(\boldsymbol{\theta},\mathbf{d}_i))^2$. The final expression in (SM3.3) is the pdf of the multivariate t-distribution (see, for example, Kotz and Nadarajah, 2004, page 1), with location vector $\boldsymbol{\mu}(\boldsymbol{\theta},\mathbf{D})$, scale matrix $(s_2/s_1)I_n$ and s_1 degrees of freedom. From Lange et al. (1989) the Fisher information under this marginal model is

$$\mathcal{I}(\boldsymbol{\theta},\mathbf{D}) = \frac{s_1(s_1+n)}{s_2(s_1+n+2)}M(\boldsymbol{\theta},\mathbf{D})^{\mathrm{T}}M(\boldsymbol{\theta},\mathbf{D}). \hspace{2cm}\text{(SM3.4)}$$

Using the same argument as in the proof of Lemmas 3.1 and 3.2, it follows that the expected FIG utility is

$$U_{FIG}(\mathbf{D}) = \frac{s_1(s_1+n)}{s_2(s_1+n+2)}\sum_{i=1}^{n}\phi(\mathbf{d}_i),$$

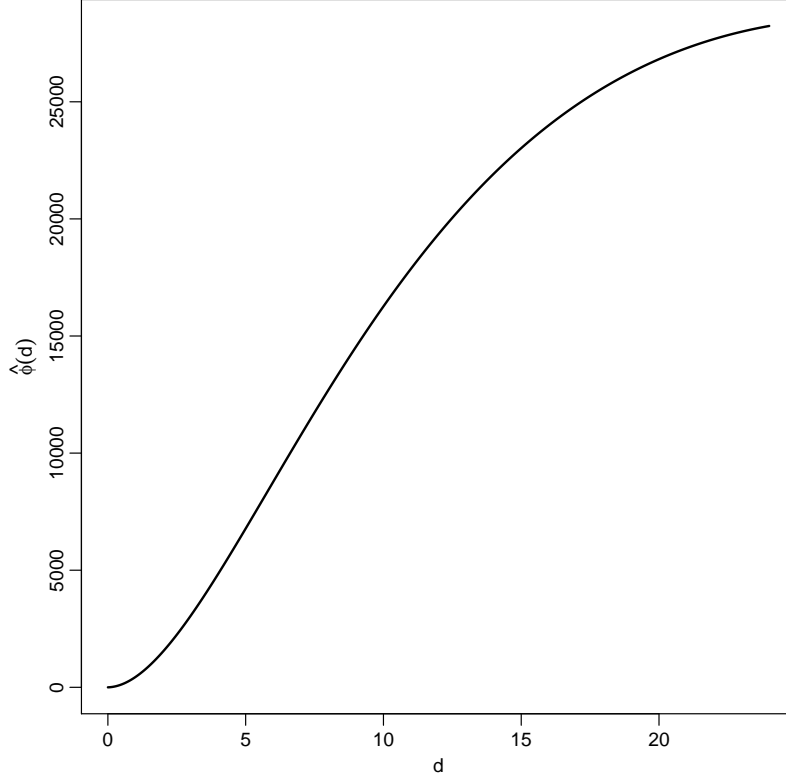with $\phi(\mathbf{d})$ given by (5) in the main manuscript.

Figure SM1: Plot of $\hat{\phi}(d)$ against d for the compartmental non-linear model.

SM4   Proof of Proposition 4.1

Since $\partial\mu(\boldsymbol{\theta},\mathbf{d})/\partial\theta_j = f_j(\mathbf{d})$, independent of $\boldsymbol{\theta}$, it follows from Lemma 3.3 that

$$\phi(\mathbf{d}) = \sum_{j=1}^{p} \prod_{r=1}^{k} \left(d_r^{u_{rj}}\right)^2.$$

The positive row sum assumption, along with with $u_{rj} \geq 0$, means that $\phi(\mathbf{d})$ is maximised by

$$\tilde{d}_r = \begin{cases} a_r & \text{if } |a_r| > |b_r|; \\ b_r & \text{if } |a_r| < |b_r|; \\ \pm|b_r| & \text{if } |a_r| = |b_r|; \end{cases}$$

for $r = 1, \ldots, k$.

SM5   Proof of Proposition 4.2

Under a FIG optimal design, the $n \times p$ matrix $M(\boldsymbol{\theta}, \tilde{\mathbf{D}})$ has ijth element given by $f_j(\tilde{\mathbf{d}}_i)$. From Proposition 4.1, the j_1th column of $M(\boldsymbol{\theta}, \tilde{\mathbf{D}})$ is $\mathbf{1}_n$ (where $\mathbf{1}_n$ denotes the $n \times 1$ vector of ones) and the j_2th column is

$$\left(f_{j_2}(\tilde{\mathbf{d}}_1), \ldots, f_{j_2}(\tilde{\mathbf{d}}_n)\right)^{\mathrm{T}} = \begin{cases} a_{\bar{r}}^2 \mathbf{1}_n & \text{if } |a_{\bar{r}}| > |b_{\bar{r}}|; \\ b_{\bar{r}}^2 \mathbf{1}_n & \text{if } |a_{\bar{r}}| < |b_{\bar{r}}|; \\ b_{\bar{r}}^2 \mathbf{1}_n & \text{if } |a_{\bar{r}}| = |b_{\bar{r}}|. \end{cases}$$

In all cases, columns j_1 and j_2 are linearly dependent. Therefore $M(\boldsymbol{\theta}, \tilde{\mathbf{D}})$ has less than p linearly independent columns and is rank deficient.

SM6   Plot of $\hat{\phi}(d)$ against d for the compartmental non-linear model

Figure SM1 shows the approximate objective function $\hat{\phi}(d)$, for the compartmental non-linear model, plotted against $d \in [0, 24]$.

References

Atkinson, A., Donev, A. and Tobias, R. (2007) *Optimum Experimental Designs, with SAS.* Oxford University Press.

Catchpole, E. A. and Morgan, B. J. T. (1997) Detecting parameter redundancy. *Biometrika* **84**, 187–196.

Chaloner, K. and Verdinelli, I. (1995) Bayesian experimental design: A review. *Statistical Science* **10**, 273–304.

Davis, P. and Rabinowitz, P. (1984) *Methods of Numerical Integration.* Academic Press 2nd edn.

Gelfand, A. and Sahu, S. (1999) Identifiability, improper priors, and Gibbs sampling for generalized linear models. *Journal of the American Statistical Association* **94**, 247–253.

Gotwalt, C. M., Jones, B. A. and Steinberg, D. M. (2009) Fast computation of designs robust to parameter uncertainty for nonlinear settings. *Technometrics* **51**, 88–95.

Huan, X. and Marzouk, Y. (2016) Sequential bayesian optimal experimental design via approximate dynamic programming. *arXiv:1604.08320.*

Kosmidis, I., Kenne Pagui, E. C. and Sartori, N. (2020) Mean and median bias reduction in generalized linear models. *Statistics and Computing* **30**, 43–59.

Kotz, S. and Nadarajah, S. (2004) *Multivariate t Distributions and Their Applications.* Cambridge.

Lange, K., Little, R. and Taylor, J. (1989) Robust statistical modeling using the t distribution. *Journal of the American Statistical Association* **84**, 881–896.

Lindley, D. (1956) On a measure of the information provided by an experiment. *Annals of Mathematical Statistics* **27**, 986–1005.

Morris, M. (2011) *Design of Experiments: an introduction based on linear models.* Chapman & Hall.

Prangle, D., Harbisher, S. and Gillespie, C. (2020) Bayesian experimental design without posterior calculations: an adversarial approach. *arXiv:1904:05703.*

Robert, C. (2007) *The Bayesian Choice.* Springer 2nd edn.

Ryan, E., Drovandi, C., McGree, J. and Pettitt, A. (2016) A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review* **84**, 128–154.

Schervish, M. (1995) *Theory of Statistics.* Springer-Verlag, New York.

Walker, S. (2016) Bayesian information in an experiment and the Fisher information distance. *Statistics and Probability Letters* **112**, 5–9.

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

Woods, D. C., Lewis, S. M., Eccleston, J. A. and Russell, K. G. (2006) Designs for generalized linear models with several variables and model uncertainty. *Technometrics* **48**, 284–292.