

Bayesian design of experiments for generalised linear models and dimensional analysis with industrial and scientific application

David C. Woods[†], Antony M. Overstall[‡], Maria Adamou[†] and Timothy W. Waite^{*}

[†]University of Southampton, Southampton, UK

[‡]University of Glasgow, Glasgow, UK

^{*}University of Manchester, Manchester, UK

Abstract

The design of an experiment can be always be considered at least implicitly Bayesian, with prior knowledge used informally to aid decisions such as the variables to be studied and the choice of a plausible relationship between the explanatory variables and measured responses. Bayesian methods allow uncertainty in these decisions to be incorporated into design selection through prior distributions that encapsulate information available from scientific knowledge or previous experimentation. Further, a design may be explicitly tailored to the aim of the experiment through a decision-theoretic approach using an appropriate loss function. We review the area of decision-theoretic Bayesian design, with particular emphasis on recent advances in computational methods. For many problems arising in industry and science, experiments result in a discrete response that is well described by a member of the class of generalised linear models. Bayesian design for such nonlinear models is often seen as impractical as the expected loss is analytically intractable and numerical approximations are usually computationally expensive. We describe how Gaussian process emulation, commonly used in computer experiments, can play an important role in facilitating Bayesian design for realistic problems. A main focus is the combination of Gaussian process regression to approximate the expected loss with cyclic descent (coordinate exchange) optimisation algorithms to allow optimal designs to be found for previously infeasible problems. We also present the first optimal design results for statistical models formed from dimensional analysis, a methodology widely employed in the engineering and physical sciences to produce parsimonious and interpretable models. Using the famous paper helicopter experiment, we show the potential for the combination of Bayesian design, generalised linear models and dimensional analysis to produce small but informative experiments.

This paper will appear in Quality Engineering

(<http://www.tandfonline.com/doi/full/10.1080/08982112.2016.1246045>).

Keywords: Computer experiments; D -optimality; Gaussian process models; high-dimensional design; nonlinear models; smoothing.

1 Introduction

Design of experiments is an “a priori” activity, taking place before data has been collected, and hence the Bayesian paradigm is a particularly appropriate approach to take. Bayesian methods allow available prior information on the model to be incorporated into both the design of the experiment and the analysis of the resulting data, and produce posterior distributions that are interpretable by scientists. They also reduce reliance on unrealistic assumptions and asymptotic results that may be inappropriate for small- to medium-sized experiments. The Bayesian approach to design enables realistic and coherent accounting for the substantial model and parameter uncertainties that usually exist before an experiment is performed and it is also a natural framework for sequential inference and design.

Contact: David Woods; D.Woods@southampton.ac.uk; Southampton Statistical Sciences Research Institute, University of Southampton, Southampton SO17 1BJ UK

An important problem where Bayesian methods can have substantial impact is optimal design for nonlinear modelling, which relies on some prior information being available about the unknown values of the model parameters (see Atkinson et al., 2007, ch. 17). A Bayesian approach relaxes the requirement of locally optimal design criteria to specify particular values of the parameters. Fully Bayesian design, predicated on using the posterior distributions for inference as outlined below, is also less reliant on the asymptotic assumptions that underpin most classical design for nonlinear models.

A decision-theoretic Bayesian optimal design is found through minimisation of the expectation of a loss function that is chosen to encapsulate the aims of the experiment. Suppose that we require a design for q variables in n points, with the i th point defined as $\mathbf{x}_i = (x_{i1}, \dots, x_{iq})^T \in \mathcal{X} \subset \mathbb{R}^q$. The choice of design space \mathcal{X} determines the quality of design that can be found and should be chosen in conjunction with the statistical model. For a linear model composed of only first-order terms and interactions, a discrete design space containing only the corner points of, for example, the unit cube will suffice for finding an optimal design. For the more complex nonlinear models considered in this paper, a continuous design space is needed to allow an optimal design to be found.

Let $l(\xi, \mathbf{y}, \boldsymbol{\psi})$ be the loss function for a design $\xi = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathcal{X}^n$ producing data $\mathbf{y} = (y_1, \dots, y_n)^T \in \mathcal{Y}$. Assume a statistical model defined via likelihood $p(\mathbf{y} | \boldsymbol{\psi})$, with parameters $\boldsymbol{\psi} \in \Psi$ having prior density $p(\boldsymbol{\psi})$. The vector $\boldsymbol{\psi}$ may include parameters defining both the mean and variance of \mathbf{y} . Then an optimal design ξ^* is defined as

$$\xi^* = \operatorname{argmin}_{\xi \in \mathcal{X}^n} \int_{\mathcal{Y}} \int_{\Psi} l(\xi, \mathbf{y}, \boldsymbol{\psi}) p(\boldsymbol{\psi}, \mathbf{y} | \xi) d\boldsymbol{\psi} d\mathbf{y}. \quad (1)$$

For further details, see the landmark review paper of Chaloner and Verdinelli (1995).

There are a number of challenges in calculating the expected loss in (1):

- (a) the evaluation of l itself is potentially non-trivial, as it may depend on the posterior distribution and only be available numerically;
- (b) the integrals in (1) may be very high dimensional, and are unlikely to be analytically tractable;
- (c) evaluation of the joint density $p(\boldsymbol{\psi}, \mathbf{y} | \xi) = p(\mathbf{y} | \boldsymbol{\psi}, \xi) p(\boldsymbol{\psi})$ may be complicated by the computational expense of calculating the likelihood $p(\mathbf{y} | \boldsymbol{\psi}, \xi)$ for complex models. For example, hierarchical nonlinear models, perhaps resulting from restricted randomisation in the experiment, lead to analytically intractable likelihoods. Alternatively, the likelihood may result from the evaluation of a complex numerical model and hence not be available in closed-form.

Common choices for l include (i) the self-information loss $p(\boldsymbol{\psi}) - \log p(\boldsymbol{\psi} | \mathbf{y}, \xi)$, and (ii) the squared-error loss between (a function of) $\boldsymbol{\psi}$ and its posterior expectation. If the prior density does not depend on the design, minimisation of the expected self-information loss is equivalent to maximisation of the expected Kullback-Leibler(KL) divergence between the prior and posterior distributions (see MacKay, 2003, ch. 2, and Sebastiani and Wynn, 2000). For some experiments, bespoke loss functions may be required. For example, it may be necessary or desirable to incorporate the cost of each run of the experiment. We demonstrate results using both expected self-information loss (SIL), defined as

$$\Phi(\xi)_{\text{SIL}} = \int_{\mathcal{Y}} \int_{\Psi} [\log p(\boldsymbol{\psi}) - \log p(\boldsymbol{\psi} | \mathbf{y}, \xi)] p(\boldsymbol{\psi}, \mathbf{y} | \xi) d\boldsymbol{\psi} d\mathbf{y}, \quad (2)$$

and the expected squared-error loss (SEL) for the prediction of the mean response $\mu(\mathbf{x})$,

$$\Phi(\xi)_{\text{SEL}} = \int_{\mathcal{Y}} \int_{\Psi} \int_{\mathcal{X}} \{\mu(\mathbf{x}) - E[\mu(\mathbf{x}) | \mathbf{y}, \xi]\}^2 p(\boldsymbol{\psi}, \mathbf{y} | \xi) d\mathbf{x} d\boldsymbol{\psi} d\mathbf{y}. \quad (3)$$

We refer to a design minimising (2) or (3) as SIL-optimal or SEL-optimal, respectively.

Until very recently, optimal Bayesian design has not evolved far from the methods reviewed by Chaloner and Verdinelli (1995). Development and application of methods for Bayesian design have lagged behind the progress made in inference and modelling due to the additional complexity introduced by the need to integrate

over the (as yet) unobserved responses, in addition to unknown model parameters. Hence, methodology has been restricted to simple models and fully sequential, one-point-at-a-time, procedures (Ryan et al., 2016).

In this paper, we focus on experiments for multi-variable generalised linear models and present new results for optimal design for dimensional analysis. In Section 2 we introduce the generalised linear models for which optimal designs will be sought and use Box’s helicopter experiment to introduce and demonstrate the concepts of dimensional analysis. In Section 3, we then review the main approaches to overcoming challenges (a)-(c) that have been proposed in the literature. A major focus is the approximate coordinate exchange methodology proposed by Overstall and Woods (2016b). In Section 4 we use approximate coordinate exchange to find Bayesian optimal designs for three examples, including the helicopter experiment. We finish with a short discussion in Section 5, highlighting issues that may hinder the adoption of Bayesian design in practice, and propose some potential remedies.

2 Experiments with generalized linear models

Generalised linear models (GLMs; McCullagh and Nelder, 1989) are an important class of models for scientific and industrial experiments whose response cannot be well described by a normal-theory linear model (see Myers et al., 2010). In addition to standard linear regression, the class of GLMs includes models for binary and count data. A GLM has three components:

1. A distribution for the univariate response $y(\mathbf{x})$ taken from the exponential family.
2. A linear predictor $\eta(\mathbf{x}) = \mathbf{f}^T(\mathbf{x})\boldsymbol{\beta}$, with p -vector $\mathbf{f}(\mathbf{x})$ holding known functions of the explanatory variables and $\boldsymbol{\beta}$ a p -vector of unknown model parameters.
3. A link function $g(\mu(\mathbf{x})) = \eta(\mathbf{x})$ relating \mathbf{x} to the mean $E\{y(\mathbf{x})\} = \mu(\mathbf{x})$.

The variance of $y(\mathbf{x})$ takes the form $\text{Var}\{y(\mathbf{x})\} = \phi V\{\mu(\mathbf{x})\}$, with $\phi > 0$ a dispersion parameter and the form of the function V dependent on the selected exponential family distribution.

Atkinson and Woods (2015) reviewed the upsurge in the development of design methodology for GLMs that has taken place over the last 10 years or so. Key to frequentist, and much Bayesian, optimal design for GLMs is the Fisher information matrix for $\boldsymbol{\beta}$, which for an n -run experiment takes the form

$$M(\boldsymbol{\beta}; \xi) = X^T W X, \tag{4}$$

with X an $n \times p$ model matrix with i th row equal to $\mathbf{f}^T(\mathbf{x}_i)$ and W an $n \times n$ diagonal matrix with i th entry

$$w(\mathbf{x}_i) = [\text{Var}\{y(\mathbf{x}_i)\}]^{-1} \left(\frac{d\mu(\mathbf{x}_i)}{d\eta(\mathbf{x}_i)} \right)^2, \quad i = 1, \dots, n.$$

In general, the information matrix depends on the values of the unknown model parameters $\boldsymbol{\beta}$ through the matrix W . A notable exception is for the linear regression model.

Below we discuss two classes of models: (i) GLMs for discrete responses; and (ii) a GLM for a continuous response with a linear predictor that incorporates physical principles via dimensional analysis.

2.1 Experiments with discrete responses

Perhaps the most common examples of optimal design for GLMs involve discrete responses, for example binary, binomial or count data. Woods et al. (2006) and Woods and van de Ven (2011) described examples from chemistry, food technology and engineering with binary (success/fail) responses. In particular, the potato-packing experiment from Woods et al. (2006) involved measuring the formation, or not, of moisture in a protected atmosphere package. The treatments consisted of the settings of three variables: vitamin concentration in the pre-packing dip and the levels of two gases in the atmosphere. A suitable GLM here might

be a logistic regression. Let $y(\mathbf{x}_i) \sim \text{Bernoulli}\{\rho(\mathbf{x}_i)\}$ be the response from the i th run of the experiment, with variable settings \mathbf{x}_i and

$$\log\left(\frac{\rho(\mathbf{x}_i)}{1-\rho(\mathbf{x}_i)}\right) = \beta_0 + \sum_{j=1}^3 \beta_j x_{ij} + \sum_{j=1}^3 \sum_{k \geq j}^3 \beta_{jk} x_{ij} x_{ik}, \quad (5)$$

where $\beta_0, \dots, \beta_3, \beta_{11}, \beta_{12}, \dots, \beta_{33}$ are unknown parameters to be estimated. Here, $\mu(\mathbf{x}_i) = \rho(\mathbf{x}_i)$ and the variance function is given by $V\{\mu(\mathbf{x}_i)\} = \rho(\mathbf{x}_i)[1-\rho(\mathbf{x}_i)]$ with $\phi = 1$. To illustrate some Bayesian design concepts, Atkinson and Woods (2015) assumed the following independent prior distributions for the parameters:

$$\beta_1, \beta_2 \sim U(2, 6), \quad \beta_0, \beta_3, \beta_{jk} \sim U(-2, 2) \text{ for } j, k = 1, 2, 3. \quad (6)$$

For multi-variable experiments, most theoretical progress on optimal design has been made for Poisson distributed responses, see for example Russell et al. (2009). Poisson regression is often employed in industrial experiments counting numbers of defects (Wu and Hamada, 2009, ch. 14) or in environmental and biological experiments where the response is the count of animal numbers or cell growth. Let $y(\mathbf{x}_i) \sim \text{Poisson}\{\mu(\mathbf{x}_i)\}$, with $V\{\mu(\mathbf{x}_i)\} = \mu(\mathbf{x}_i)$ and $\phi = 1$. McGree and Eccleston (2012) and Atkinson and Woods (2015) presented theoretical constructions of optimal designs robust to the values of the model parameters for log-linear models with linear predictors of the form

$$\log\{\mu(\mathbf{x}_i)\} = \beta_0 + \sum_{j=1}^q \beta_j x_{ij}, \quad i = 1, \dots, n. \quad (7)$$

The latter authors illustrated the construction methods for experiments with $q = 5$ variables with uniform prior distributions assumed for each β_j ,

$$U(1, 1 + \alpha) \text{ for } j = 1, 3, 5 \text{ and } U(-1 - \alpha, -1) \text{ for } j = 2, 4, \quad (8)$$

with $\alpha > 0$. The intercept $\beta_0 = 0$ was assumed known.

In Section 4.1, we find, assess and compare Bayesian optimal designs for both the logistic and log-linear models.

2.2 Dimensional analysis

Dimensional analysis (DA) is a methodology commonly used by engineers and physical scientists to produce parsimonious and dimensionally consistent models (Sonin, 2001). A base set of dimensionless variables are identified via (nonlinear) transformations of the explanatory variables and related to a similarly transformed response variable via a, typically nonlinear, regression model. The model thus formed will satisfy Buckingham's Π theorem (Buckingham, 1914, 1915a,b), which states that physically meaningful relationships must be dimensionally homogeneous. In addition to performing a priori dimension reduction of the input variables, DA provides the possibility of obtaining models that are scale-free and hence, for example, applicable to a range of manufacturing processing scales from lab to production.

Reviews of DA from a statistician's perspective are provided by the recent papers from Albrecht et al. (2013) and Shen et al. (2014). The latter authors applied DA to Box's paper helicopter experiment (Box and Liu, 1999), and we will use this example to demonstrate the potential of the combination of DA, GLMs and Bayesian design.

We use a standard paper helicopter pattern and consider three independent variables: rotor width, rotor length and tail length, see Figure 1. The ranges of these three variables, and other dimensions, are taken from <http://www.paperhelicopterexperiment.com> and are given in Table 1, along with settings of other physical parameters.

Shen et al. (2014) identified the base quantities

$$\Pi_0 = \frac{h}{T\sqrt{gr}}, \quad \Pi_1 = \frac{m}{\rho r^3},$$




Figure 1: Paper helicopter pattern.

Variable/parameter (symbol)	Range/setting & units
Flight time (T)	s
Rotor width (w)	$\in [0.03, 0.09]$ m
Rotor length (r)	$\in [0.07, 0.12]$ m
Tail length (t)	$\in [0.07, 0.12]$ m
Mass (m)	kg
Drop height (h)	Fixed at 2 m
Acceleration due to gravity (g)	Fixed at 9.80665 ms^{-2}
Air density (ρ)	Fixed at 1.20412 kgm^{-3}
Paper density (D)	Fixed at 0.120 kgm^{-2}
Body length (b)	Fixed at 0.025 m
Tail width (d)	Fixed at 0.05 m

Table 1: Variable ranges and physical parameter settings (in SI units) for the paper helicopter experiment. Flight time is the measured response.

and fitted a regression model for flight time on the log scale with

$$\begin{aligned} \log E(T) &= \log \theta_0 + \log \frac{h}{\sqrt{gr}} + \theta_1 \log \left[\frac{\rho r^3}{m} \right] \\ \Rightarrow E(T) &= \frac{\theta_0 h}{\sqrt{gr}} \left[\frac{\rho r^3}{m} \right]^{\theta_1}. \end{aligned} \quad (9)$$

Equation (9) follows a power law and is consistent with Bridgman's principle of absolute significance of relative magnitude (Bridgman, 1931), another fundamental theorem of DA.

In common with Shen et al. (2014), we consider experiments with $n = 4$ runs but adapt these authors' study in two ways:

1. We assume the density of the paper is fixed, and hence helicopter mass (m) is a varying function of the three controllable variables rotor width, rotor length and tail length. Relabeling $T = y(\mathbf{x})$ and $r = x_1$, $w = x_2$ and $t = x_3$ for consistency, equation (9) then becomes

$$\log E\{y(\mathbf{x})\} = \log \mu(\mathbf{x}) = \log \theta_0 + \log \frac{h}{\sqrt{gx_1}} + \theta_1 \log \left\{ \frac{\rho x_1^3}{D[2x_2(x_1 + b) + x_3d]} \right\}. \quad (10)$$

2. Rather than assume additive normal errors on the log scale, we model flight time as a Gamma distributed random variable and assume a GLM with log link and linear predictor (10). Hence $V\{\mu(\mathbf{x})\} = \phi\mu(\mathbf{x})^2$ and ϕ is assumed unknown. The second term on the right-hand side of (10) is treated as an offset, and hence there are three parameters, $\log \theta_0$, θ_1 and ϕ , that require estimation.

We assume prior distributions

$$\log \theta_0 \sim N(0.102, 0.0625), \quad \theta_1 \sim N(0.460, 0.0625), \quad \phi \sim U(0.75, 1.25). \quad (11)$$

The prior means of the physical parameters are equal to estimates from Shen et al. (2014); we set the prior variances equal to 2.5 times the estimated variances from Shen et al. (2014) to obtain more diffuse prior distributions. For the dispersion parameter, our choice of prior leads to the variance of $y(\mathbf{x})$ being between 75% and 125% of the value of $\mu(\mathbf{x})^2$.

In Section 4.2, we find, assess and compare Bayesian optimal designs for this Gamma regression model.

3 A review of approaches to Bayesian design

In this section, we provide an overview of some of the most common approaches to Bayesian design in the literature. We focus on (i) analytical and computational approximations to the expected loss in (1), and (ii) optimization methods for multi-variable experiments.

3.1 Asymptotic approximations

For experiments with large n , the inverse of the expected Fisher information matrix $M(\boldsymbol{\psi}; \xi)$ is an asymptotic approximation to the posterior variance-covariance matrix of the parameters $\boldsymbol{\psi}$. Use of this approximation leads to pseudo-Bayesian “alphabetic” optimality criteria (Atkinson et al., 2007, ch. 10). For example, under pseudo-Bayesian D -optimality, a design is selected to minimise

$$\Phi_D(\xi) = \int_{\Psi} -\log |M(\boldsymbol{\psi}; \xi)| \pi(\boldsymbol{\psi}) d\boldsymbol{\psi}. \quad (12)$$

The integral with respect to $\boldsymbol{\psi}$ is usually of low dimension and amenable to deterministic approximation. Such an approximation to the objective function can then be minimised using a conditional algorithm such as point or coordinate exchange; see, for, example, Gotwalt et al. (2009). For a point prior density on $\boldsymbol{\psi}$, which is equivalent to assuming known parameter values, minimisation of (12) leads to a locally D -optimal design.

3.2 Simulation-based optimisation

In general, the expected loss can be approximated via Monte Carlo integration (e.g. Gentle, 2003, ch. 7) as

$$\tilde{\Phi}(\xi) = \frac{1}{B} \sum_{k=1}^B l(\xi, \mathbf{y}_k, \boldsymbol{\psi}_k),$$

with $(\boldsymbol{\psi}_k, \mathbf{y}_k) \sim p(\boldsymbol{\psi}, \mathbf{y} | \xi)$ a random sample drawn from the joint distribution of parameters $\boldsymbol{\psi}$ and responses \mathbf{y} . Typically, this is obtained by sampling $\boldsymbol{\psi}$ from the prior density $p(\boldsymbol{\psi})$ and then sampling \mathbf{y} from the conditional density $p(\mathbf{y} | \boldsymbol{\psi}, \xi)$. The loss $l(\xi, \mathbf{y}_k, \boldsymbol{\psi}_k)$ often itself requires numerical approximation, necessitating a nested, double-loop, Monte Carlo simulation; see Ryan (2003). Direct optimisation of this approximation requires large B to generate a suitably precise approximation to the objective function and/or expensive stochastic algorithms (e.g. simulated annealing or genetic algorithms), see for example, Hamada et al. (2001) and Huan and Marzouk (2013) who employed polynomial chaos approximations to facilitate sampling from $p(\boldsymbol{\psi}, \mathbf{y} | \xi)$. Alternatively, the optimisation can be embedded within a Markov chain simulation scheme, and an optimal design identified by sampling from an artificial joint distribution for the design, model parameters and data, and then finding the mode of the marginal distribution of the design (Müller, 1999, Müller et al., 2004 and Amzal et al., 2006). Typically, an annealing step is employed to enable easier identification of the optimal design. This approach is most effective for small experiments (few variables and runs). Recent extensions to this algorithm have allowed designs to be found for (i) models with intractable likelihoods using Approximate Bayesian Computation (see Drovandi and Pettitt, 2013, and also Hainy et al., 2013 for other ABC design methods) and (ii) dynamic models with numerous sampling times using dimension reduction (Ryan et al., 2014), importance sampling and Laplace approximations (Ryan et al., 2015).

3.3 Sequential design

Most experiments are part of a sequence, where a Bayesian approach, with sequential updating from prior to posterior distributions, is natural. For point-sequential designs, with one point at a time added to the design, approximation of the expected loss is greatly simplified by the resulting reduction in the dimension of the integral. Recent methods have been suggested for estimation of, and discrimination between, nonlinear models, see Drovandi et al. (2013, 2014). A growing area is Bayesian optimisation of expensive black-box functions (e.g. in computer experiments), using Gaussian process surrogates to reduce the number of

required function evaluations, following the seminal work of Jones et al. (1998). The computational efficiency of sequential design can be greatly aided through the use of sequential Monte Carlo for the necessary inference (see Gramacy and Polson, 2011).

3.4 Smoothing-based optimisation

Smoothing-based design methods (Müller and Parmigiani, 1996) evaluate a computationally expensive, typically Monte Carlo, approximation to the expected loss (1) for a limited number of designs and then smooth these approximated losses to generate a surrogate function which can then be optimised in place of the true expected loss. Recent research includes (i) extension of such methods to employ Gaussian process smoothing and Bayesian optimisation (Weaver et al., 2016) and (ii) use of surrogates to enable Bayesian D -optimal design for generalised linear mixed models (Waite and Woods, 2015).

A key challenge for the application of smoothing-based methods to design problems with large numbers of runs or variables is the high-dimensional smoothing that is required. An extension that addresses this challenge using conditional smoothing and optimisation is described in the next section.

3.5 Approximate coordinate exchange

Overstall and Woods (2016b) proposed the first general purpose methodology for high-dimensional, multi-variable Bayesian design that does not rely on normal approximations to the posterior density. Their approximate coordinate exchange (ACE) algorithm is a conditional optimisation algorithm that makes use of surrogates, or emulators, for the expected loss as a function of a single design coordinate (i.e. the value of a single variable for a particular run).

In the field of computer experiments (Sacks et al., 1989), an emulator is a statistical model built to approximate the output from an expensive computer model or code. The most common emulator is a Gaussian process model (GP; see, for example, Rasmussen and Williams, 2006), which non-parametrically smooths or interpolates the computer model output. In the ACE algorithm, GP models are used to smooth the Monte Carlo approximation to the expected loss.

Algorithm 1 gives the basic steps of ACE. Let $\tilde{\Phi}_{ij}(x|\xi)$ be a Monte Carlo approximation to the expected loss for design ξ with ij th coordinate replaced by $x \in \mathcal{X}_j$ ($i = 1, \dots, n; j = 1, \dots, q$). Here, \mathcal{X}_j is the projection of the design space \mathcal{X} onto the j th dimension. The algorithm steps through each coordinate of the design, and for each coordinate constructs a one-dimensional emulator, $\hat{\Phi}(x)$, for $\tilde{\Phi}_{ij}(x|\xi)$. The value of x that minimises this emulator is found, and an accept/reject step performed in order to decide whether to swap the current design coordinate with this proposed coordinate. This accept/reject step is described in Algorithm 2.

The ACE algorithm would typically be repeated multiple times (perhaps exploiting parallel computing) to avoid local optima. Overstall and Woods (2016b) gave more details of the implementation and application of the algorithm, including its combination with a point-exchange algorithm to consolidate clusters of similar design points.

A GP model is employed in line 10 of Algorithm 1. The emulator is given by the posterior mean function of the GP

$$\hat{\Phi}(x) = \hat{\mu}_{ij} + \hat{\sigma}_{ij} \mathbf{a}^T(x, \zeta_{ij}) A(\zeta_{ij})^{-1} \mathbf{z}_{ij}, \quad (13)$$

with $\hat{\mu}_{ij} = \sum_{k=1}^Q \tilde{\Phi}_{ij}(x_{ij}^k|\xi)/Q$, $\hat{\sigma}_{ij}^2 = \sum_{k=1}^Q (\tilde{\Phi}_{ij}(x_{ij}^k|\xi) - \hat{\mu}_{ij})^2 / (Q - 1)$, \mathbf{z}_{ij} a Q -vector having k th entry equal to the standardised approximate expected loss

$$(\tilde{\Phi}_{ij}(x_{ij}^k|\xi) - \hat{\mu}_{ij}) / \hat{\sigma}_{ij},$$

and $\zeta_{ij} = \{x_{ij}^1, \dots, x_{ij}^Q\}$ being points from a one-dimensional space-filling design in \mathcal{X}_j (see Algorithm 1). Under the common assumption of a squared exponential correlation structure, the correlation between the responses at two points is inversely proportional to the exponential of the squared difference between the points, leading to the Q -vector \mathbf{a} and $Q \times Q$ matrix A having entries

$$\mathbf{a}(x, \zeta_{ij})_u = \exp\{-\rho(x - x_{ij}^u)^2\}, \quad A(\zeta_{ij})_{uv} = \exp\{-\rho(x_{ij}^u - x_{ij}^v)^2\} + \eta \mathbf{I}(u = v), \quad u, v = 1, \dots, Q,$$

Input : Initial (randomly chosen) design ξ
Output: Φ -optimal design

```

1 begin
2   repeat
3     for  $i = 1 : n$  do
4       for  $j = 1 : q$  do
5         Generate a 1d space-filling design  $\zeta_{ij} = \{x_{ij}^1, \dots, x_{ij}^Q\}$  in  $\mathcal{X}_j \subset \mathbb{R}$ ;
6         for  $k = 1 : Q$  do
7           Evaluate  $\tilde{\Phi}_{ij}(x_{ij}^k | \xi)$  using Monte Carlo sample size  $B$ 
8         end
9         ;
10        Construct a 1d emulator  $\hat{\Phi}(x)$  via (13);
11        Set the  $ij$ th coordinate,  $x_{ij}$ , of  $\xi$  equal to  $\operatorname{argmin}_{x \in \mathcal{X}_j} \hat{\Phi}(x)$  with probability  $p^*$  obtained
        from Algorithm 2;
12      end
13    end
14  until convergence;
15 end

```

Algorithm 1: The approximate coordinate exchange (ACE) algorithm.

Input : Current design ξ and proposed new coordinate x
Output: Posterior probability p^* that $\tilde{\Phi}_{ij}(x|\xi) < \tilde{\Phi}(\xi)$

```

1 begin
2   Let  $\xi_p$  be the design formed by replacing the  $ij$ th coordinate of  $\xi$  with  $x$ ;
3   for  $k = 1 : \tilde{B}$  do
4     Sample  $\tilde{\psi}$  from  $\pi(\psi)$ ;
5     Sample  $\mathbf{y}_1 \sim \pi(\mathbf{y} | \psi, \xi_p)$  and  $\mathbf{y}_2 \sim \pi(\mathbf{y} | \psi, \xi)$ ;
6     Set  $L_{1k} = l(\xi_p, \mathbf{y}_1, \tilde{\psi})$  and  $L_{2k} = l(\xi, \mathbf{y}_2, \tilde{\psi})$ ;
7   end
8   Assume  $L_{1k} \sim N(b_1 + b_2, a)$  and  $L_{2k} \sim N(b_1, a)$ ;
9   Calculate the posterior probability,  $p^*$ , that  $b_2 < 0$  using “data”  $L_{1k}$  and  $L_{2k}$ ;
10 end

```

Algorithm 2: Accept/reject step from line 11 of the ACE algorithm.

where I is the indicator function and $\eta > 0$ is the nugget. Adding η to the diagonal elements of the correlation matrix ensures the emulator will smooth, rather than interpolate, the stochastic Monte Carlo approximation $\tilde{\Phi}$, and improves the numerical stability of the emulator (Gramacy and Lee, 2012). We estimate ρ and η via maximum likelihood.

The minimisation in line 11 of Algorithm 1 is subject to both Monte Carlo error and emulator error. To remove the emulator error when making the decision whether to accept the exchange, the steps in Algorithm 2 are performed using \tilde{B} independent Monte Carlo samples to assess the improvement in the design. Algorithm 2 essentially describes a Bayesian t -test based on simulated data from the existing and proposed designs (cf Wang and Zhang, 2006). If the assumption of normality that underpins this test is invalid, a nonparametric procedure may be used instead.

Convergence in Algorithm 1 is assessed graphically, in a similar spirit to convergence diagnostics for Markov chain Monte Carlo; see Overstall and Woods (2016b) for examples.

4 Bayesian designs for generalized linear models

In this section, we find Bayesian optimal designs for the three GLMs outlined in Section 2. The new designs are found using ACE, with $B = 1000$ and $Q = 20$ in Algorithm 1 and $\tilde{B} = 20,000$ in Algorithm 2, unless otherwise stated, and assessed against various competitors from the literature.

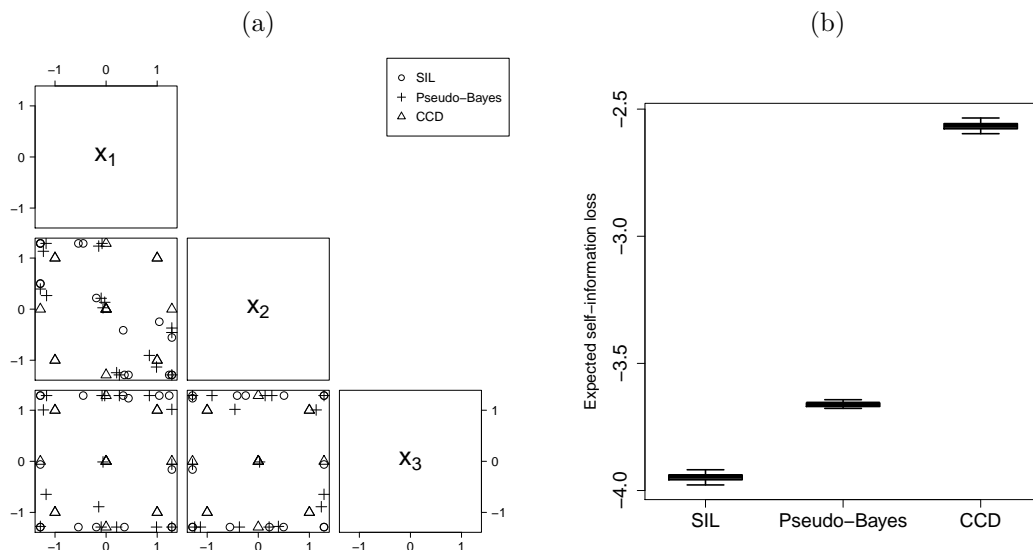


Figure 2: Logistic regression example: (a) Two-dimensional projections of the SIL-optimal, pseudo-Bayesian D -optimal and central composite designs. (b) Boxplots of 20 Monte Carlo approximations of the negative expected KL divergence (2) using $B = 20,000$ simulations for the SIL-optimal, pseudo-Bayesian D -optimal and central composite designs.

4.1 Optimal design for experiments with discrete response

We now demonstrate Bayesian design for logistic and Poisson regression. As most previous design work for these models has focussed on D -optimality, we find SIL-optimal designs and compare with pseudo-Bayesian D -optimal designs. These latter designs are an asymptotic approximation to the SIL-optimal designs (see Overstall and Woods, 2016b).

4.1.1 Logistic regression

We start by finding $n = 16$ run SIL-optimal designs for logistic regression, minimising a Monte Carlo approximation to (2) and using $Q = 10$ in Algorithm 1, with linear predictor (5), $\mathcal{X} = [-1.2872, 1.2872]^3$ and prior distribution (6). The choice of design space \mathcal{X} allows comparison to an orthogonal central composite design (CCD) with $n = 16$ points. The CCD is a common design for a linear model with a second-order linear predictor (see Dean and Voss, 1999, ch. 16). Figure 2(a) gives two-dimensional projections of the SIL-optimal design, along with the CCD and a pseudo-Bayesian D -optimal design with $n = 16$ points that minimises a quadrature approximation to (12). The main qualitative difference between the SIL-optimal and D -optimal designs is the greater concentration of points at the extremes of the design region for the SIL-optimal design, especially for x_3 . These differences in the distribution of the design points can be clearly seen in Figure 3, which shows histograms of the one-dimensional projections of the three designs. The concentration of points at the extremes of x_3 for both optimal designs is consistent with literature results on locally optimal design with parameter values that are small in absolute value (see Cox, 1988); here, $E(\beta_3) = 0$.

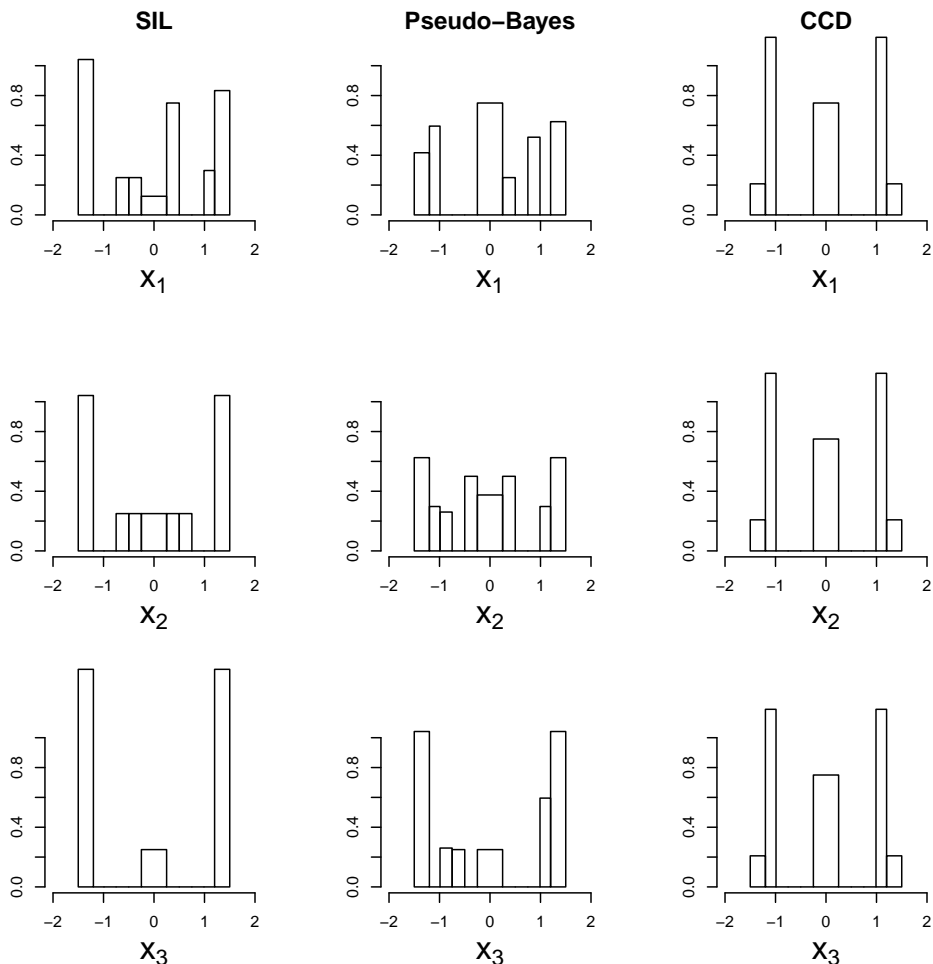


Figure 3: Logistic regression example: one-dimensional projections of the SIL-optimal, pseudo-Bayesian D -optimal and central composite designs.

To quantitatively compare the three designs, we performed 20 repeated Monte Carlo approximations

of (2), each using $B = 20,000$. Boxplots of these results are given in Figure 2(b). The SIL-optimal design naturally performs best, and has around 10% lower negative expected KL divergence compared to the pseudo-Bayesian design. Both are substantially better than the central composite design, which of course makes no use of prior information and should only be employed for the estimation of a second-order response surface model which is linear in the unknown parameters.

4.1.2 Poisson regression

McGree and Eccleston (2012) and Atkinson and Woods (2015) presented analytical construction methods for pseudo-Bayesian optimal designs for log-linear regression with linear predictor (7). The resulting designs minimise (12) amongst the class of minimally supported designs, that is, designs for which the number of distinct design points is equal to the number of parameters in the linear predictor. The construction method uses the algorithm of Russell et al. (2009). For $-1 \leq x_{ij} \leq 1$ ($i = 1, \dots, q + 1$; $j = 1, \dots, q$) and $|E(\beta_j)| > 1$, the unreplicated minimally supported pseudo-Bayesian D -optimal (MSPBD) design has points $\mathbf{x}_i = \mathbf{c} - 2\mathbf{e}_i/E(\beta_i)$ for $i = 1, \dots, q$ and $\mathbf{x}_{q+1} = \mathbf{c}$, where \mathbf{e}_i is the i th column of the $q \times q$ identity matrix and $\mathbf{c} = (c_1, \dots, c_q)^T$ with $c_i = 1$ if $E(\beta_i) > 0$ and $c_i = -1$ if $E(\beta_i) < 0$. The minimally supported optimal design for $q = 5$ variables and prior distribution (8) is given in Table 2.

Using ACE, we find SIL-optimal designs, again minimising a Monte Carlo approximation to (2), with $n = q + 1 = 6$ design points under prior distribution (8). The designs are given in Tables 3(a) and 3(b) for $\alpha = 0.5$ and $\alpha = 0.75$, respectively. These designs have the same structure as the MSPBD-optimal design, with each variable only taking two values, one value with $x_{ij} = \pm 1$ and one value with $-1 < x_{ij} < 1$. Unlike the MSPBD-optimal designs, this latter value is not constant in absolute value across the variables, although it does always have the same sign as the corresponding value in the MSPBD-optimal design.

Run	x_1	x_2	x_3	x_4	x_5
1	$-\gamma$	-1	1	-1	1
2	1	γ	1	-1	1
3	1	-1	$-\gamma$	-1	1
4	1	-1	1	γ	1
5	1	-1	1	-1	$-\gamma$
6	1	-1	1	-1	1

Table 2: Log-linear regression example: minimally-supported Pseudo-Bayesian D -optimal design under uniform prior distribution (8); $\gamma = 0.6$ for $\alpha = 0.5$ and $\gamma = 0.455$ for $\alpha = 0.75$.

(a)						(b)					
Run	x_1	x_2	x_3	x_4	x_5	Run	x_1	x_2	x_3	x_4	x_5
1	-0.500	-1	1	-1	1	1	-0.220	-1	1	-1	1
2	1	0.555	1	-1	1	2	1	0.222	1	-1	1
3	1	-1	-0.309	-1	1	3	1	-1	-0.323	-1	1
4	1	-1	1	0.334	1	4	1	-1	1	0.110	1
5	1	-1	1	-1	-0.381	5	1	-1	1	-1	-0.308
6	1	-1	1	-1	1	6	1	-1	1	-1	1

Table 3: Log-linear regression example: SIL-optimal designs under uniform prior distribution (8) for (a) $\alpha = 0.5$ and (b) $\alpha = 0.75$.

A quantitative comparison of the designs is given in Figures 4(a) for $\alpha = 0.5$ and 4(b) for $\alpha = 0.75$, which display 20 repeated Monte Carlo approximations of (2), each using $B = 20,000$. The expected loss is

lower for the less diffuse prior distribution ($\alpha = 0.5$). For both values of α , the SIL-optimal and MSPBD-optimal designs perform very similarly, showing that the asymptotic approximation (12) is considerably more effective for this problem than for the binary response example, even though the experiment size n is smaller. This finding is unsurprising as the normal distribution, from which (12) is derived, is a more effective approximation to the Poisson distribution than it is to the Bernoulli distribution.




Figure 4: Log-linear regression example: boxplots of 20 Monte Carlo approximations of the negative expected KL divergence (2) using $B = 20,000$ simulations for the SIL-optimal and minimally supported Pseudo-Bayesian D -optimal designs (MSPBD) under uniform prior distributions (8) for (a) $\alpha = 0.5$ and (b) $\alpha = 0.75$.

4.2 Optimal designs for dimensional analysis

For the paper helicopter experiment, we find a SEL-optimal design, minimising a Monte Carlo approximation to (3), for the Gamma regression model with linear predictor (10) and prior distribution (11). The integral with respect to \mathbf{x} in (3) is approximated by summation across a 4^3 grid of values constructed from $x_1, x_3 \in \{0.07, 0.087, 0.103, 0.12\}$ (rotor and tail length) and $x_2 \in \{0.03, 0.05, 0.07, 0.09\}$ (rotor width). The design is given in Table 4(a) in terms of both the original three variables and the base quantity, $\Pi_1 \in [-1.903, 0.351]$. Clearly, any design that results in the same values of Π_1 will have the same expected posterior variance. A goal of equally weighted prediction over \mathcal{X} leads to greater weight being given to smaller values of Π_1 , leading to the design including only points with smaller values of the base quantity.

We compare the SEL-optimal design to four competitors:

- A V -optimal design: the Fisher information matrix for β for Gamma regression with the log link is equivalent to that for a linear model. Hence, classical optimal designs for the linear model (Atkinson et al., 2007, ch. 10) can be employed with this example. A V -optimal design, that minimizes the average prediction variance, has equally replicated design points with $-\log \Pi_1 = -1.904$ and $-\log \Pi_1 = 0.351$. One such design is given in Table 4(b).
- The two regular $n = 2^{3-1}$ fractional factorial designs with defining relation $I = x_1 x_2 x_3$.
- A maximin Latin hypercube (LH) design (Santner et al., 2003, ch. 5) with $n = 4$ points: twenty such designs were generated algorithmically using different starting designs. Each will be an approximation to the (globally) optimal maximin LH design.

(a) SEL-optimal design				
Run	x_1	x_2	x_3	$-\log \Pi_1$
1	0.070	0.079	0.095	-1.753
2	0.070	0.087	0.102	-1.843
3	0.070	0.076	0.116	-1.769
4	0.076	0.089	0.073	-1.593

(b) V -optimal design				
Run	x_1	x_2	x_3	$-\log \Pi_1$
1	0.070	0.090	0.120	-1.903
2	0.070	0.090	0.120	-1.903
3	0.120	0.030	0.070	0.351
4	0.120	0.030	0.070	0.351

Table 4: Helicopter experiment: (a) SEL-optimal and (b) V -optimal designs.




Figure 5: Helicopter experiment: boxplots of 20 Monte Carlo approximations of the average expected posterior variance (3) from $B = 20,000$ simulations.

A quantitative comparison of these designs is given in Figure 5 which displays boxplots from 20 Monte Carlo approximations to (3) for the SEL-optimal, V -optimal, fractional factorial and LH designs. The SEL-optimal design has a performance advantage over all the other designs, having average expected posterior variance around 8% smaller than the V -optimal design, 11-12% smaller than the fractional factorial designs, and 11-17% smaller than the LH designs.

5 Discussion

Optimal Bayesian design is challenging for high-dimensional problems with multi-variable models and/or many design points and there are few literature examples of such designs being used in practice. Reasons for this include the lack of scalable algorithms for design selection, the complexity of available software for Bayesian design, an occasional unwillingness to “bias” designs through the use of prior information (that is, to focus design performance on a certain set of parameter scenarios or models) and, in many application areas, a lack of appreciation that design of experiments can go beyond standard factorial designs. However, Bayesian design is a powerful tool for a variety of experiments. Here, we have focussed on using new methodology to find designs for generalised linear models where some prior information is *necessary* to design informative experiments. We have also demonstrated for the first time how algorithmic Bayesian design can combine empirical and physical modelling principles via generalised linear models and dimensional analysis.

Methodology such as ACE removes some of the barriers to the implementation of Bayesian design, both by widening the scope of models and experiments that can be addressed, and by facilitating the provision of greater evidence for the effectiveness of the methods through rigorous scientific studies. Although the methodology is still computationally challenging for larger examples, an increase in statistical efficiency that allows even slightly smaller experiments to be run can lead to considerable cost savings in expensive industrial experimentation. These savings will often more than offset the additional computational time and resource used to find the designs.

Clearly, in practice it is not always sensible to choose a design based solely on a one-number summary of design performance, particularly if it is obtained from a generic loss function that may not capture the aims of the experiment. However, being able to find optimal, or near-optimal, designs under suitable loss functions enables a short-list of competing designs to be compared on other merits. The methodology demonstrated

in this paper allows the experimenter to understand any trade-offs resulting from the incorporation of other practical considerations.

More details of the methodology demonstrated in this paper can be found in Overstall and Woods (2016b) and also in Overstall et al. (2015) who discussed optimal designs for uncertainty quantification of physical models, an application area of increasing importance. We demonstrated ACE using straightforward Monte Carlo approximations to the expected loss. However, the methodology can be applied with a variety of different approximations to the loss function, including asymptotic approximations to find pseudo-Bayesian designs. The ACE algorithm has been implemented in R package `acebayes` (Overstall and Woods, 2016a).

Acknowledgements

Woods was supported by Fellowship EP/J018317/1 from the UK Engineering and Physical Sciences Research Council. We are grateful to a reviewer for comments that improved the paper.

References

- Albrecht, M. C., C. J. Nachtsheim, T. A. Albrecht, and R. D. Cook (2013). Experimental design for engineering dimensional analysis (with discussion). *Technometrics* 55, 257–295.
- Amzal, B., F. Y. Bois, E. Parent, and C. Robert (2006). Bayesian optimal design via interacting particle systems. *Journal of the American Statistical Association* 101, 773–785.
- Atkinson, A. C., A. N. Donev, and R. D. Tobias (2007). *Optimum Experimental Design, with SAS* (2nd ed.). Oxford: Oxford University Press.
- Atkinson, A. C. and D. C. Woods (2015). Designs for generalized linear models. In A. M. Dean, M. D. Morris, J. Stufken, and D. R. Bingham (Eds.), *Handbook of Design and Analysis of Experiments*. Boca Raton: Chapman & Hall/CRC.
- Box, G. E. P. and P. Y. T. Liu (1999). Statistics as a catalyst to learning by scientific method part I - an example. *Journal of Quality Technology* 31, 1–15.
- Bridgman, P. (1931). *Dimensional Analysis* (2nd ed.). Yale University Press.
- Buckingham, E. (1914). On physically similar systems; illustrations of the use of dimensional equations. *Physical Review* 4, 345–376.
- Buckingham, E. (1915a). Model experiments and the forms of empirical equations. *Transactions of the American Society of Mechanical Engineers* 37, 263–296.
- Buckingham, E. (1915b). The principle of similitude. *Nature* 96, 396–397.
- Chaloner, K. and I. Verdinelli (1995). Bayesian experimental design: a review. *Statistical Science* 10, 273–304.
- Cox, D. R. (1988). A note on design when response has an exponential family distribution. *Biometrika* 75, 161–164.
- Dean, A. M. and D. T. Voss (1999). *Design and Analysis of Experiments*. New York: Springer.
- Drovandi, C. C., J. M. McGree, and A. N. Pettitt (2013). Sequential Monte Carlo for Bayesian sequentially designed experiments for discrete data. *Computational Statistics and Data Analysis* 57, 320–335.
- Drovandi, C. C., J. M. McGree, and A. N. Pettitt (2014). A sequential Monte Carlo algorithm to incorporate model uncertainty in Bayesian sequential design. *Journal of Computational and Graphical Statistics* 23, 3–24.

- Drovandi, C. C. and A. N. Pettitt (2013). Bayesian experimental design for models with intractable likelihoods. *Biometrics* 69, 937–948.
- Gentle, J. E. (2003). *Random Number Generation and Monte Carlo Methods* (2nd ed.). New York: Springer.
- Gotwalt, C. M., B. A. Jones, and D. M. Steinberg (2009). Fast computation of designs robust to parameter uncertainty for nonlinear settings. *Technometrics* 51, 88–95.
- Gramacy, R. B. and H. K. H. Lee (2012). Cases for the nugget in modeling computer experiments. *Statistics and Computing* 22, 713–722.
- Gramacy, R. B. and N. G. Polson (2011). Particle learning of Gaussian process models for sequential design and optimization. *Journal of Computational and Graphical Statistics* 20, 102–118.
- Hainy, M., W. G. Müller, and H. P. Wynn (2013). Approximate Bayesian computational design (abcd), an introduction. In D. Ucinski, A. Atkinson, and M. Patan (Eds.), *MODA 10 - Advances in Model-Oriented Design and Analysis*. Springer.
- Hamada, M., H. F. Martz, C. S. Reese, and A. G. Wilson (2001). Finding near-optimal Bayesian experimental designs via genetic algorithms. *The American Statistician* 55, 175–181.
- Huan, X. and Y. M. Marzouk (2013). Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics* 232, 288–317.
- Jones, D., M. Schonlau, and W. Welch (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* 13, 455–492.
- MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press.
- McCullagh, P. and J. A. Nelder (1989). *Generalized Linear Models* (2nd ed.). London: Chapman and Hall.
- McGree, J. M. and J. A. Eccleston (2012). Robust designs for Poisson regression models. *Technometrics* 54, 64–72.
- Müller, P. (1999). Simulation-based optimal design. In J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, and A. F. M. Smith (Eds.), *Bayesian Statistics 6*. Oxford.
- Müller, P. and G. Parmigiani (1996). Optimal design via curve fitting of Monte Carlo experiments. *Journal of the American Statistical Association* 90, 1322–1330.
- Müller, P., B. Sanso, and M. De Iorio (2004). Optimal Bayesian design by inhomogeneous Markov chain simulation. *Journal of the American Statistical Association* 99, 788–798.
- Myers, R. H., D. C. Montgomery, G. G. Vining, and T. J. Robinson (2010). *Generalized Linear Models with Applications in Engineering and the Sciences* (2nd ed.). Hoboken, New Jersey: Wiley.
- Overstall, A. M. and D. C. Woods (2016a). *acebayes: Optimal Bayesian experimental design using the ACE algorithm*. R package version 1.3.
- Overstall, A. M. and D. C. Woods (2016b). Bayesian design of experiments using approximate coordinate exchange. *Technometrics in press*.
- Overstall, A. M., D. C. Woods, and B. M. Parker (2015). Bayesian optimal design for ordinary differential equation models. *arXiv:1509.04099*.
- Rasmussen, C. E. and C. K. I. Williams (2006). *Gaussian Processes for Machine Learning*. Cambridge, MA.: MIT Press.

- Russell, K. G., D. C. Woods, S. M. Lewis, and J. A. Eccleston (2009). D-optimal designs for Poisson regression models. *Statistica Sinica* 19, 721–730.
- Ryan, E. G., C. C. Drovandi, J. M. McGree, and A. N. Pettitt (2016). A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review* 84, 128–154.
- Ryan, E. G., C. C. Drovandi, and A. N. Pettitt (2015). Fully Bayesian experimental design for pharmacokinetic studies. *Entropy* 17, 1063–1089.
- Ryan, E. G., C. C. Drovandi, M. H. Thompson, and A. N. Pettitt (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, 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., W. J. Welch, T. J. Mitchell, and H. P. Wynn (1989). Design and analysis of computer experiments (with discussion). *Statistical Science* 4, 409–435.
- Santner, T. J., B. J. Williams, and W. I. Notz (2003). *The Design and Analysis of Computer Experiments*. New York: Springer.
- Sebastiani, P. and H. P. Wynn (2000). Maximum entropy sampling. *Journal of the Royal Statistical Society B* 62, 145–157.
- Shen, W., T. Davis, D. K. J. Lin, and C. J. Nachtsheim (2014). Dimensional analysis and its applications in statistics. *Journal of Quality Technology* 46, 185–198.
- Sonin, A. A. (2001). *The Physical Basis of Dimensional Analysis* (2nd ed.). Cambridge, MA.: Department of Mechanical Engineering, Massachusetts Institute of Technology.
- Waite, T. W. and D. C. Woods (2015). Designs for generalized linear models with random block effects via information matrix approximations. *Biometrika* 102, 677–693.
- Wang, L. and L. Zhang (2006). Stochastic optimization using simulated annealing with hypothesis test. *Applied Mathematics and Computation* 174, 1329–1342.
- Weaver, B. P., B. J. Williams, C. M. Anderson-Cook, and D. M. Higdon (2016). Computational enhancements to Bayesian design of experiments using Gaussian processes. *Bayesian Analysis* 11, 191–213.
- Woods, D. C., S. Lewis, J. A. Eccleston, and K. G. Russell (2006). Designs for generalised linear models with several variables and model uncertainty. *Technometrics* 48, 284–292.
- Woods, D. C. and P. van de Ven (2011). Blocked designs for experiments with non-normal response. *Technometrics* 53, 173–182.
- Wu, C. F. J. and M. Hamada (2009). *Experiments: Planning, Analysis and Optimization* (2nd ed.). Hoboken, New Jersey: Wiley.