



Contents lists available at ScienceDirect

Computational Statistics and Data Analysis

journal homepage: www.elsevier.com/locate/csda

Model selection via Bayesian information capacity designs for generalised linear models

David C. Woods^{a,*}, James M. McGree^b, Susan M. Lewis^a^a University of Southampton, UK^b Queensland University of Technology, Australia

ARTICLE INFO

Article history:

Received 29 January 2016

Received in revised form 22 October 2016

Accepted 26 October 2016

Available online 2 November 2016

Keywords:

Bayesian *D*-optimality

Factorial experiments

Generalised information criterion

Screening

ABSTRACT

The first investigation is made of designs for screening experiments where the response variable is approximated by a generalised linear model. A Bayesian information capacity criterion is defined for the selection of designs that are robust to the form of the linear predictor. For binomial data and logistic regression, the effectiveness of these designs for screening is assessed through simulation studies using all-subsets regression and model selection via maximum penalised likelihood and a generalised information criterion. For Poisson data and log-linear regression, similar assessments are made using maximum likelihood and the Akaike information criterion for minimally-supported designs that are constructed analytically. The results show that effective screening, that is, high power with moderate type I error rate and false discovery rate, can be achieved through suitable choices for the number of design support points and experiment size. Logistic regression is shown to present a more challenging problem than log-linear regression. Some areas for future work are also indicated.

© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

An important problem in scientific discovery is to find those variables (or factors) that have a substantive influence on an observed response through experiments on a possibly large set of potentially important variables. There has been much research into such variable screening, or model selection, focussed on the design and analysis of experiments in which the response variable is adequately approximated by a linear model (see Draguljić et al., 2014 and Woods and Lewis, 2016, and references therein). Such experiments are used increasingly in scientific research and product development, for example, in the pharmaceutical and chemical industries.

In many practical applications, for example when binary or count data are observed, a generalised linear model (GLM; McCullagh and Nelder, 1989) may be needed to describe a response. Previous research on designs for model selection for GLMs has focussed on experiments involving only a few variables through pairwise comparisons of a small number of models (see, for example, López-Fidalgo et al., 2007 and Waterhouse et al., 2008). Hence, such methods are not applicable to, or easily generalisable for, the screening problem. In the literature, the majority of multi-variable experimentation with GLMs has employed (fractional) factorial designs, including examples on solder-joint defects (Hamada and Nelder, 1997), windshield moulding, non-conforming tiles and semi-conductor defects (see Lewis et al., 2001). Although such designs are effective for both model selection and estimation for normal-theory linear models, they have been shown to be inefficient for experiments that provide non-normal data (Woods et al., 2006).

* Correspondence to: Mathematical Sciences, University of Southampton, Southampton, SO17 1BJ, UK.
E-mail address: D.Woods@southampton.ac.uk (D.C. Woods).

In common with other non-linear models, for the GLMs considered in this paper the performance of a design depends on the unknown values of the parameters in the model. One approach to overcoming this problem is to assume a particular value for each parameter and hence obtain a “locally optimal” design; that is, a design that is optimal under a given criterion provided the assigned parameter values are correct. We adopt the alternative approach of making the less stringent assumption of a prior distribution for each model parameter from which we obtain a “pseudo-Bayesian” design (Atkinson and Woods, 2015).

In this paper, we investigate variable screening for GLMs with q independent variables, labelled x_1, \dots, x_q . In the j th run ($j = 1, \dots, N$) of the experiment, a treatment or combination of variable values $\mathbf{x}_j = (x_{1j}, \dots, x_{qj})^T$ is applied to an experimental unit and a univariate response, y_j , is observed. We assume that $|x_{ij}| \leq 1$ for $i = 1, \dots, q$; $j = 1, \dots, N$.

The aim of the experiment is to identify those *active* variables having a substantial effect on the response variable and to estimate efficiently a GLM involving those variables alone. For $j = 1, \dots, N$, the y_j have independent exponential family distributions with expectation μ_j related to a linear predictor $\eta_j = f(\mathbf{x}_j)^T \boldsymbol{\beta}$ via a link function, $g(\mu_j) = \eta_j$. The vectors $f(\mathbf{x})$ and $\boldsymbol{\beta}$ are $p \times 1$ vectors of known functions of \mathbf{x} and unknown model parameters, respectively. We also assume that the experimental units are exchangeable, in the sense that the distribution of the response to a treatment does not depend on the unit to which the treatment is applied.

For canonical link functions, the log-likelihood may be written as

$$l(\boldsymbol{\beta}; \mathbf{y}) = \sum_{j=1}^N [y_j \eta_j - b(\eta_j) + c(y_j)], \quad (1)$$

where $b(\cdot)$ and $c(\cdot)$ are known functions of the linear predictor and response, respectively. For the binomial distribution and the logistic link, $b(\eta_j) = -n_j \log(1 + e^{\eta_j})$ and $c(y_j) = \log(n_j! / [y_j!(n_j - y_j)!])$, with n_j the number of Bernoulli trials made at the j th run. For the Poisson distribution and the log link, $b(\eta_j) = e^{\eta_j}$ and $c(y_j) = -\log(y_j!)$.

Maximum likelihood estimators (MLEs) $\hat{\boldsymbol{\beta}}$ can be found via (numerical) maximisation of (1). For small data sets, however, the MLEs may have considerable bias. For sparse data, such as binomial data with small numbers, n_j , of trials for each run, one or more maximum likelihood estimates may be infinite, for example, as the result of separation of the responses into zeros and ones via a hyperplane in the linear predictor (Silvapulle, 1981). To remove this bias and guarantee the existence of estimates for GLMs with a canonical link function, Firth (1993) defined penalised maximum likelihood estimators $\tilde{\boldsymbol{\beta}}$ as maximisers of

$$l^*(\boldsymbol{\beta}; \mathbf{y}) = l(\boldsymbol{\beta}; \mathbf{y}) + \frac{1}{2} \log \det \{X^T W X\}, \quad (2)$$

where X is the $N \times p$ model matrix with j th row $f(\mathbf{x}_j)^T$ and $W = \text{diag}\{\text{var}(y_j)\}$ (see also Kosmidis and Firth, 2009). This estimation procedure is equivalent to finding the posterior mode of $\boldsymbol{\beta}$ assuming the Jeffreys prior distribution.

The information matrix $X^T W X$, which is the asymptotic inverse variance–covariance matrix for both $\hat{\boldsymbol{\beta}}$ and $\tilde{\boldsymbol{\beta}}$, is used to define the D -optimality criterion. This criterion specifies selection of a design that maximises the objective function

$$\phi_D(\xi) = \frac{1}{p} \log \det \{X^T W X\}, \quad (3)$$

where

$$\xi = \begin{Bmatrix} \mathbf{x}_1 & \cdots & \mathbf{x}_n \\ \omega_1 & \cdots & \omega_n \end{Bmatrix}, \quad (4)$$

$\mathbf{x}_1, \dots, \mathbf{x}_n$ are the distinct treatments in the design (assumed, without loss of generality, to be applied to the first n runs of the experiment), $\omega_k > 0 \in \mathbb{N}$, and $\sum_{k=1}^n \omega_k = N$, the total number of runs. For the GLMs considered in this paper, (3) depends on $\boldsymbol{\beta}$ through the matrix W and hence selection of a D -optimal design requires knowledge of the values of these parameters. Thus a locally optimal design is obtained.

The relative performance of two designs, ξ_1 and ξ_2 , under D -optimality may be assessed using relative D -efficiency, defined as

$$\text{DEff}(\xi_1, \xi_2) = \exp \{\phi_D(\xi_1) - \phi_D(\xi_2)\}, \quad (5)$$

where $0 \leq \text{DEff}(\xi_1, \xi_2)$. If ξ_2 is a D -optimal design that maximises (3), then (5) provides an absolute measure of the performance of design ξ_1 .

In this paper, we address the screening problem of model selection and estimation of parameters in the selected model. We define, in Section 2, a Bayesian information capacity criterion that generalises D -optimality to provide model-robust designs for GLMs. We also present and discuss a model selection strategy that uses all-subsets regression and suitable penalties for model complexity. Sections 3 and 4 describe simulation studies of logistic and log-linear regression modelling, respectively, which demonstrate and assess the effectiveness of the methods. In Section 5, we present some avenues for future work to further develop methodology for screening experiments with non-normal data.

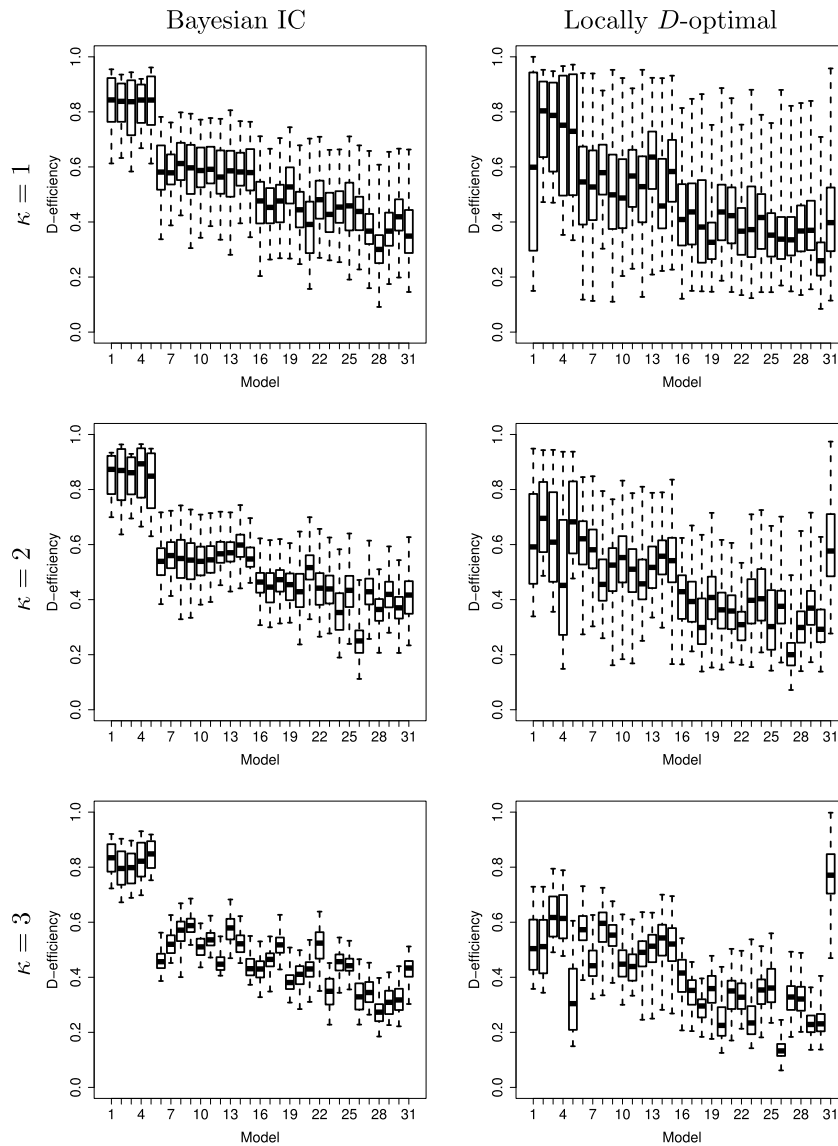


Fig. 1. Boxplots of D -efficiencies for Bayesian information capacity designs (left column) and locally D -optimal designs for the maximal model (right column) for logistic regression with $n = 6$ support points and three prior distributions ($\kappa = 1, 2, 3$) for the model parameters.

2. Information capacity designs and model selection

Consider a set \mathcal{M} of $M = |\mathcal{M}|$ distinct candidate models, each of which have the same link function. The linear predictor for the m th model and j th run is given by

$$\eta_j^m = \beta_{0m} + \sum_{i=1}^q \beta_{im} x_{ij} I(i, m), \quad (6)$$

where the β_{im} are the values of the parameters in model m , $I(i, m) = 1$ if variable i is in model m , and $I(i, m) = 0$ otherwise. Hence, the number of parameters in model m is $p_m = 1 + \sum_{i=1}^q I(i, m)$.

2.1. Bayesian information capacity

Information capacity (IC) was introduced as a linear-model design selection criterion by Sun (1993). It has been further developed and applied by, for example, Wu (1993) (supersaturated designs), and Li and Nachtsheim (2000) (model-robust factorial designs). In essence, this criterion seeks a design whose projections onto subsets of the variables produce

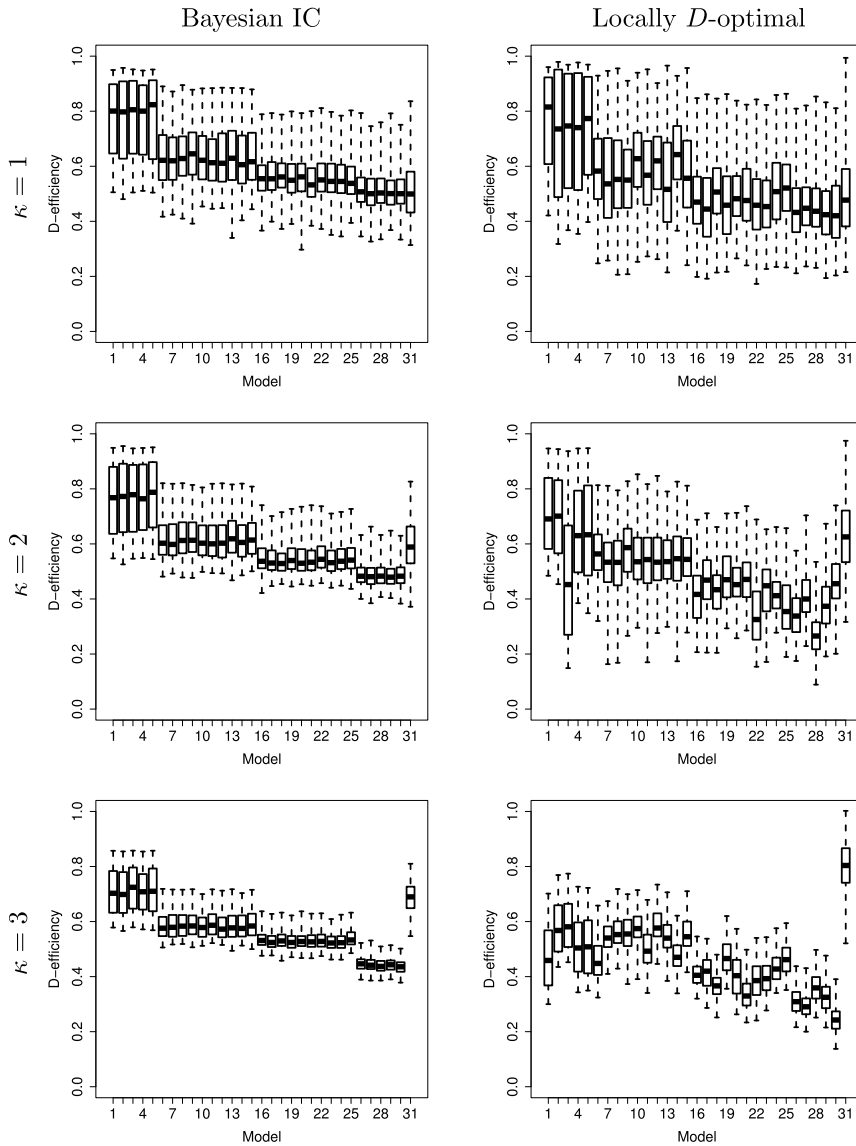


Fig. 2. Boxplots of D -efficiencies for Bayesian information capacity designs (left column) and locally D -optimal designs for the maximal model (right column) for logistic regression with $n = 30$ support points and three prior distributions ($\kappa = 1, 2, 3$) for the model parameters.

sub-designs having good estimation properties for the corresponding submodels. This is achieved by selecting a design that maximises a weighted average of the D -criterion objective function for each submodel.

For GLMs, Woods (2010) employed the criterion of Woods et al. (2006) to find locally optimal information capacity designs for an example having five variables. Designs were found that maximised

$$\Psi(\xi) = \sum_{m=1}^M \frac{1}{p_m} \log \det \{X_m^T W_m X_m\}, \tag{7}$$

where X_m and W_m are the respective model and weight matrices for the m th model in \mathcal{M} .

We define the Bayesian IC criterion which incorporates into (7) uncertainty in the parameter values assumed for each model. This criterion selects a design that maximises the objective function

$$\Phi(\xi) = \sum_{m=1}^M \frac{1}{p_m} \int_{\mathcal{B}_m} \log \det \{X_m^T W_m X_m\} \pi_m(\beta_m) d\beta_m, \tag{8}$$

where $\mathcal{B}_m \subset \mathbb{R}$ is the parameter space for model m , $\beta_m = (\beta_{0m}, \dots, \beta_{qm})^T$, and $\pi_m(\beta_m)$ is the prior distribution for β_m . The choice of $\pi(\beta_m)$ and \mathcal{B}_m , and the evaluation of (8), are discussed in Section 3.1 for logistic regression. For log-linear

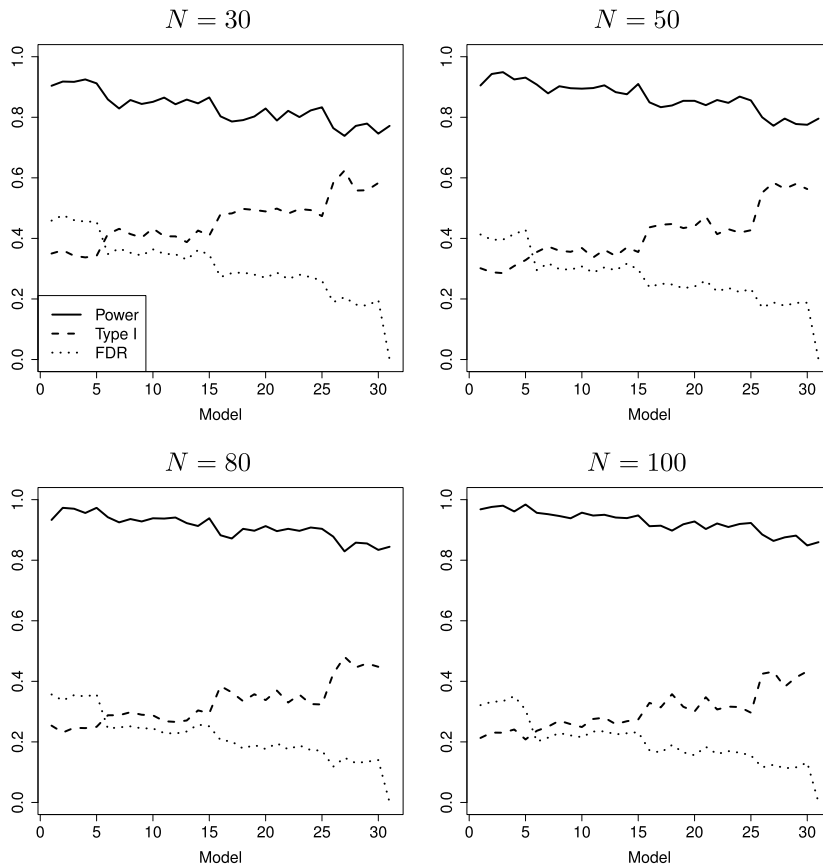


Fig. 3. Average power, type I error rate and FDR for logistic regression with $\kappa = 1$ and $N = 30, 50, 80, 100$ runs.

regression, we make use of results in the literature that enable analytical construction of minimally-supported D - and Bayesian D -optimal designs, see Section 4.1.

2.2. Model selection

A variety of model selection procedures exist for determining the most appropriate GLM from a set of models, including Bayesian (Chen et al., 2008) and shrinkage methods (Park and Hastie, 2007). To focus investigations on the impact of design selection, we restrict attention to all-subsets regression and use an information criterion to adjust for the bias inherent from in-sample estimation of the prediction error (see Burnham and Anderson, 2002, ch. 2). When maximum likelihood estimation is employed, we use the Akaike information criterion (AIC; Akaike, 1974) as the model selection criterion, and choose a model that minimises

$$AIC(m; \hat{\beta}) = -2l_m(\hat{\beta}; \mathbf{y}) + 2p_m, \tag{9}$$

where $l_m(\cdot; \cdot)$ is the log-likelihood function (1) for model m .

When β is estimated via penalised maximum likelihood (see (2)), AIC is no longer an appropriate criterion. This is because the effective number of parameters is reduced, equivalent to the inclusion of prior information (Gelman et al., 2014). The reduction depends on the number (N) of runs, with a smaller number of effective parameters for smaller N . Hence when N is small, use of AIC will over-penalise larger models. To avoid this problem, we use a generalised information criterion (GIC; Konishi and Kitagawa, 1996) that relaxes the assumptions of (i) estimation via maximum likelihood, and (ii) inclusion of the true model in \mathcal{M} . Hence, we select the model that minimises

$$GIC(m; \tilde{\beta}) = -2l_m(\tilde{\beta}; \mathbf{y}) + 2\text{tr} \left\{ J^{-1}(\tilde{\beta}) I(\tilde{\beta}) \right\}, \tag{10}$$

where

$$J(\tilde{\beta}) = -\frac{1}{N} \frac{\partial^2 l_m^*(\tilde{\beta}; \mathbf{y})}{\partial \beta \partial \beta^T} \Big|_{\tilde{\beta}}, \quad I(\tilde{\beta}) = \frac{1}{N} \frac{\partial l_m^*(\tilde{\beta}; \mathbf{y})}{\partial \beta} \frac{\partial l_m(\tilde{\beta}; \mathbf{y})}{\partial \beta^T} \Big|_{\tilde{\beta}},$$

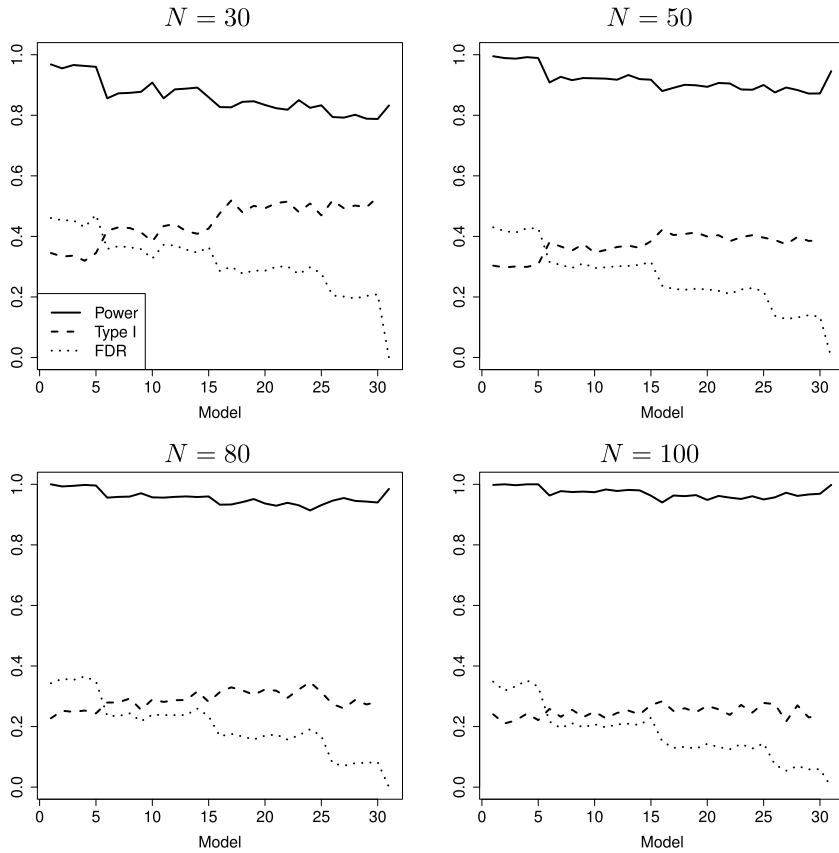


Fig. 4. Average power, type I error rate and FDR for logistic regression with $\kappa = 3$ and $N = 30, 50, 80, 100$ runs.

with $l_m^*(\cdot; \cdot)$ the penalised log-likelihood function (2) for model m ; see also Murata et al. (1994) and Zhang et al. (2010). The evaluation of $J(\hat{\beta})$ and $I(\hat{\beta})$ is straightforward for the GLMs and penalised likelihood estimation method used in this paper. The performance of the GIC is investigated in Section 3.2.

Following analysis of the data from an experiment, those variables found to be involved in the selected model are deemed to be active. In simulation studies to assess the performance of the model selection strategies, we use three summary measures: (i) *power*: the proportion of truly active variables that are correctly identified as active by the model selection strategy; (ii) *type I error rate*: the proportion of inactive variables (i.e. those not included in the true model) that are incorrectly identified as active by the model selection strategy; and (iii) *false discovery rate* (FDR): the proportion of variables identified as active by the model selection strategy that are truly inactive (i.e. not in the true model).

3. Designs and model selection for binomial response and logistic regression

To investigate the performance of the methodology for logistic regression we study a five variable example, with linear predictors of the form (6). We assume that any subset of these variables may be the set of active variables. Therefore there are 31 possible models. The models are ordered lexicographically within each model size and assigned labels $1, \dots, 31$. Models $1, \dots, 5$ have linear predictors that contain a single variable, $1, \dots, 5$, respectively; models $6, \dots, 15$ have two variables, $1, 2; 1, 3; \dots, 4, 5$. Similarly, models $16, \dots, 25$ are three variable models, $26-30$ are four-variable models and model 31 contains all five variables.

To find optimal designs and perform subsequent simulation studies, the model parameters β_{im} are assumed to have independent prior distributions of the form

$$\beta_{im} \sim \begin{cases} \text{Uniform}(\kappa, 5) & \text{for } i = 1, 3, 4 \text{ and } I(i, m) = 1, \\ \text{Uniform}(-5, -\kappa) & \text{for } i = 2, 5 \text{ and } I(i, m) = 1, \end{cases} \quad (11)$$

where $\kappa = 1, 2, 3$ and we assume $\beta_{0m} = 0$ and $\beta_{im} = 0$ if $I(i, m) = 0$ for $m = 1, \dots, M$. The adoption of bounded uniform prior distributions prevents the occurrence of parameter vectors in the support of the prior for which no design has a non-singular information matrix (cf. Waite, 2015).

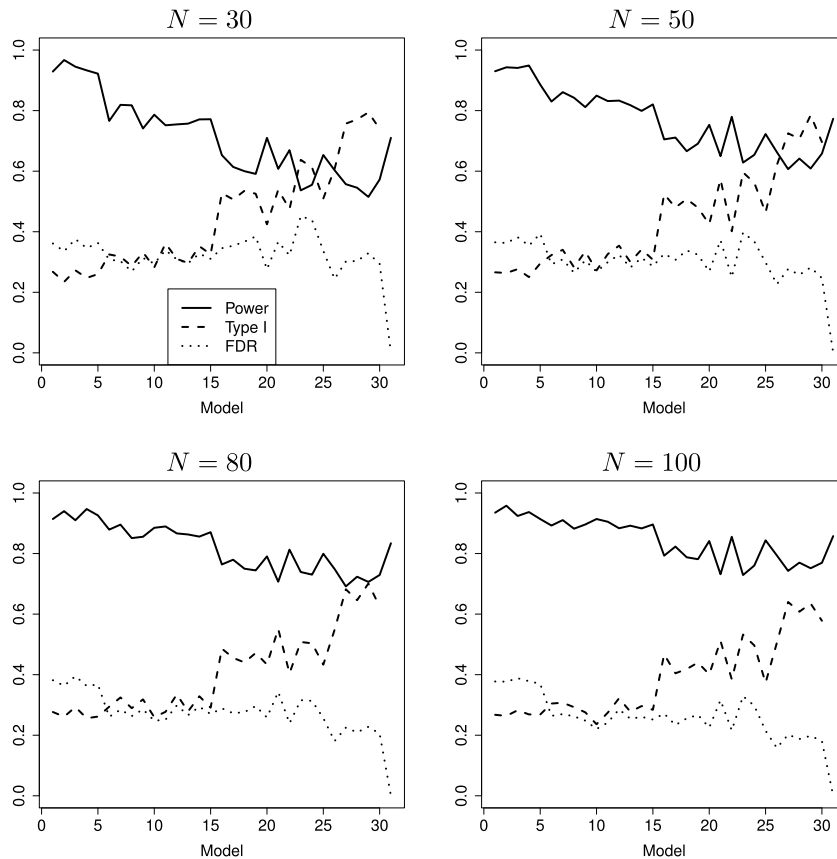


Fig. 5. Average power, type I error rate and FDR for logistic regression with $\kappa = 1$, $N = 30, 50, 80, 100$ runs and a locally D -optimal design for the maximal model.

3.1. Information capacity designs

We relax the assumption in (4) that ω_k is integer ($k = 1, \dots, n$) and consider approximate designs (e.g. Atkinson et al., 2007, ch. 9). An approximate Bayesian IC design for logistic regression maximises

$$\Phi^\dagger(\xi) = \sum_{m=1}^M \frac{1}{p_m} \int_{\mathcal{B}_m} \log \det \left\{ \sum_{k=1}^n \tilde{\omega}_k \text{var}(y_k) f_m(\mathbf{x}_k) f_m(\mathbf{x}_k)^\top \right\} \pi_m(\boldsymbol{\beta}_m) d\boldsymbol{\beta}_m, \quad (12)$$

where $0 < \tilde{\omega}_k = \omega_k/N \leq 1$ and $f_m(\mathbf{x}_k)^\top \boldsymbol{\beta}_m$ is the linear predictor for the m th model. Clearly, an optimal choice of approximate Bayesian IC design can be made independently of the total experiment size N . Finding approximate designs also substantially reduces the computational burden of the design optimisation. We found designs using simulated annealing (Haines, 1987) where the integral in (12) was evaluated numerically as a summation across a quasi-Monte Carlo sample (Lemieux, 2009, ch. 5). The simulated annealing algorithm employed was a cyclic descent algorithm that proposed, evaluated and accepted moves for one coordinate of the design at a time; see Woods (2010). Such ‘‘coordinate exchange’’ algorithms are standard in the design of experiments, and solve difficult, high-dimensional, optimisation problems via a series of one-dimensional optimisations. Use of a stochastic optimisation algorithm such as annealing has the advantage of helping to escape local optima which we have found is a particular issue when each coordinate can take values in a continuous range.

Figs. 1 and 2 (left columns) summarise the D -efficiencies of the Bayesian IC designs for $n = 6$ and $n = 30$ support points, respectively. For each choice of prior distribution for $\kappa = 1, 2, 3$ from (11), the plots are obtained using (i) 500 random draws of $\boldsymbol{\beta}_m$ values from distribution (11); (ii) the locally D -optimal design for each value of $\boldsymbol{\beta}_m$, again found using simulated annealing; and (iii) calculation of efficiency (5) to compare the Bayesian IC design with the locally D -optimal design. In general, the efficiencies decrease with model size. For $n = 6$, the efficiencies are highly variable between models, even for models of the same size. It is not uncommon for a Bayesian design for logistic regression to require a large number of support points (see Chaloner and Larntz, 1989, and Woods and Lewis, 2011). This variability is also evident in the simulation results for model selection. Hence, in the next section, we present only assessments of the designs with $n = 30$ support points.

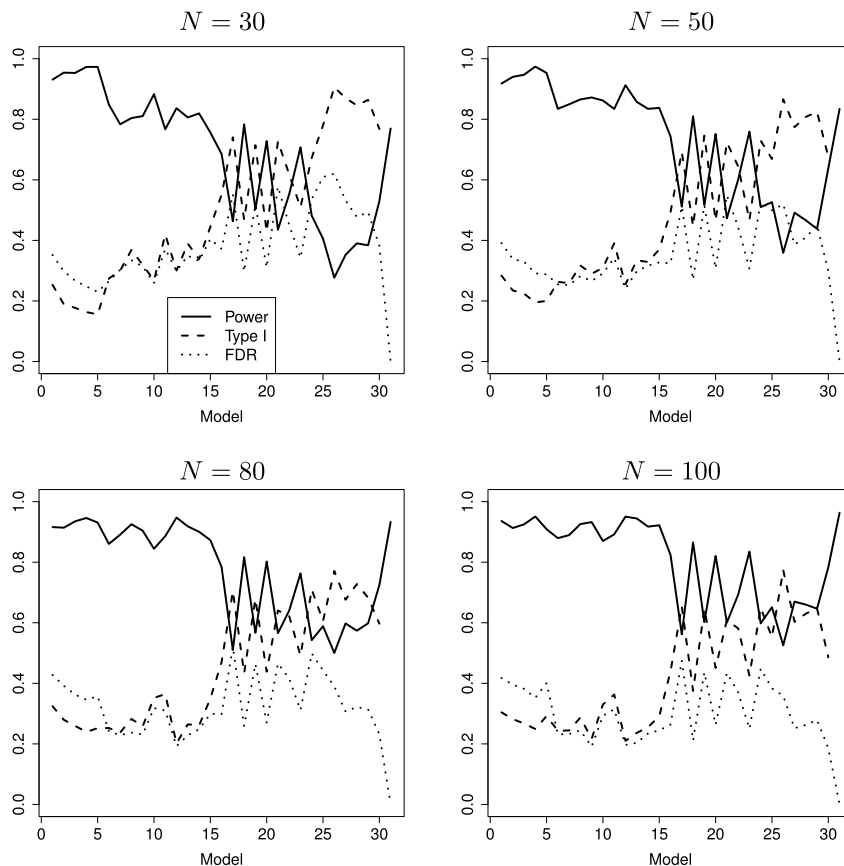


Fig. 6. Average power, type I error rate and FDR for logistic regression with $\kappa = 3$, $N = 30, 50, 80, 100$ runs and a locally D -optimal design for the maximal linear model.

For comparison, Figs. 1 and 2 (right columns) also present the D -efficiencies of locally D -optimal designs for the maximal model containing all five variables (model 31) and with each parameter set to its prior expectation. With the obvious, and expected, exception of model 31, the D -efficiencies are more variable and generally lower than those obtained from the Bayesian IC design.

3.2. Model selection results

To assess the performance of the designs for model selection, a simulation study was performed for each Bayesian IC design in which (i) each of the models, in turn, was used as the true model for the data generating process; (ii) 1000 data sets were generated independently by simulating values of β_m from the prior distribution, followed by simulation of responses \mathbf{y} from a Binomial distribution; (iii) for each data set, each of the models in \mathcal{M} was fitted using maximum penalised likelihood, and the model selected that minimises GIC (10); and (iv) power, type I error rate and FDR were calculated for each simulated data set. The optimal approximate designs were converted into exact designs via rounding ω_k to the nearest integer. Results are provided for $N = 30, 50, 80, 100$ in Figs. 3 and 4, for $\kappa = 1, 3$ respectively. The results for $\kappa = 2$ (not shown) are similar to those for $\kappa = 3$.

For all three prior distributions, high power is achieved for all experiment sizes: greater than 80% for $N = 30, 50$, and greater than 90% for $N = 80, 100$. Generally, there is a slight downward trend in power as the size of the true model increases. With the exception of the model including all five variables, a similar trend was observed for D -efficiency. The type I error rate is, unsurprisingly, an increasing function of the number of variables in the data-generating (true) model, as there are fewer inactive variables (smaller denominator) for larger true models. The maximum type I error rate of about 0.5 occurs for those true models involving four variables and corresponds to identifying, on average, less than one additional active variable. In contrast, FDR is a decreasing function of true model size, as again there are fewer inactive variables for larger true models. The maximum FDR of approximately 0.5 occurs when $N = 30$ for true models that contain a single variable, and corresponds to identifying one fewer additional active variable on average.

The results are fairly similar for the different prior distributions. The major difference is lower type I error rates when $\kappa = 3$, where there is a greater distinction between the sizes of the model coefficients for active and inactive variables. For

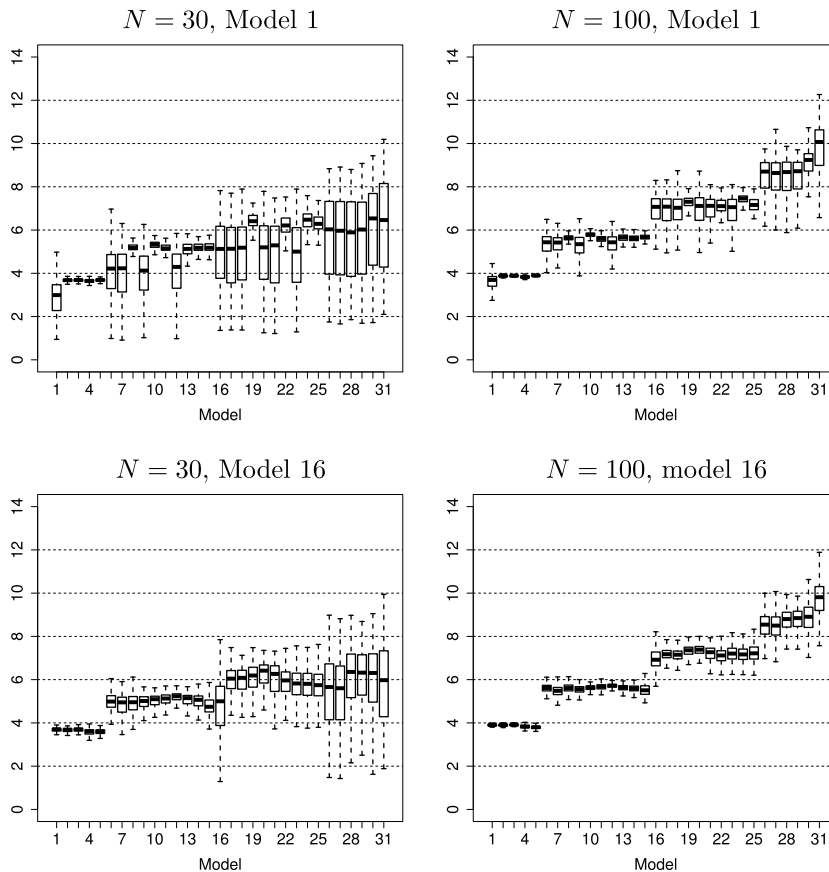


Fig. 7. Boxplots of the GIC penalty from (10) with $\kappa = 1$ and data generated from models 1 (variable 1 only) and 16 (variables 1 and 2 only).

comparison, we present results obtained from using the locally D -optimal design for model 31 (Fig. 5 for $\kappa = 1$ and Fig. 6 for $\kappa = 3$). In general, the average power is lower and average Type I error and false discovery rates higher for this design than for the Bayesian IC designs. For $\kappa = 3$, the locally optimal design has particularly poor performance for some models having three or four variables.

Another obvious comparator for the Bayesian IC designs are locally optimal IC designs, i.e. designs that maximise (7) with parameters β_m set equal to the mean of prior distribution (11) ($m = 1, \dots, M$). For this example, the Bayesian IC and locally optimal IC designs have almost identical D -efficiency distributions and model selection results, with the Bayesian IC designs displaying very slightly less variable D -efficiencies when there are $n = 6$ support points. These results (which are not shown) further illustrate the generally good performance of IC designs for model selection. We anticipate greater differences between Bayesian and locally optimal IC designs for examples where the prior distribution does not specify a known sign for each model parameter.

A key determinant of the model selection findings is the size of the GIC penalty term in (10). Our numerical studies have shown that this depends not only on the size, p_m , of the model but also on the estimated model parameters and the goodness of fit, with better-fitting models having a smaller penalty. Fig. 7 shows the distributions of the penalties obtained when model 1 (variable 1) and model 16 (variables 1 and 2) are true for $N = 30, 100$ and $\kappa = 1$. In general, the penalty is somewhat less than $2p_m$, although it increases with N and hence does not penalise larger models to the same degree as AIC. Models that include the correct variables have smaller penalty than other models. Further research on the use of this penalty is needed.

4. Design and model selection for Poisson response and log-linear regression

To investigate the performance of the methodology for log-linear regression, simulation studies were performed for two examples. Both assume the log link, $g(\mu) = \eta$, and linear predictors of the form (6). In the first example, there are again $q = 5$ variables that may affect the response (31 possible models) and prior distribution (11) is assumed. In the second example, there are $q = 10$ variables but, in line with factor sparsity (Box and Meyer, 1986), we consider only linear predictors including at most three active variables (175 possible models). Prior distributions for β_{im} are given by (11) for

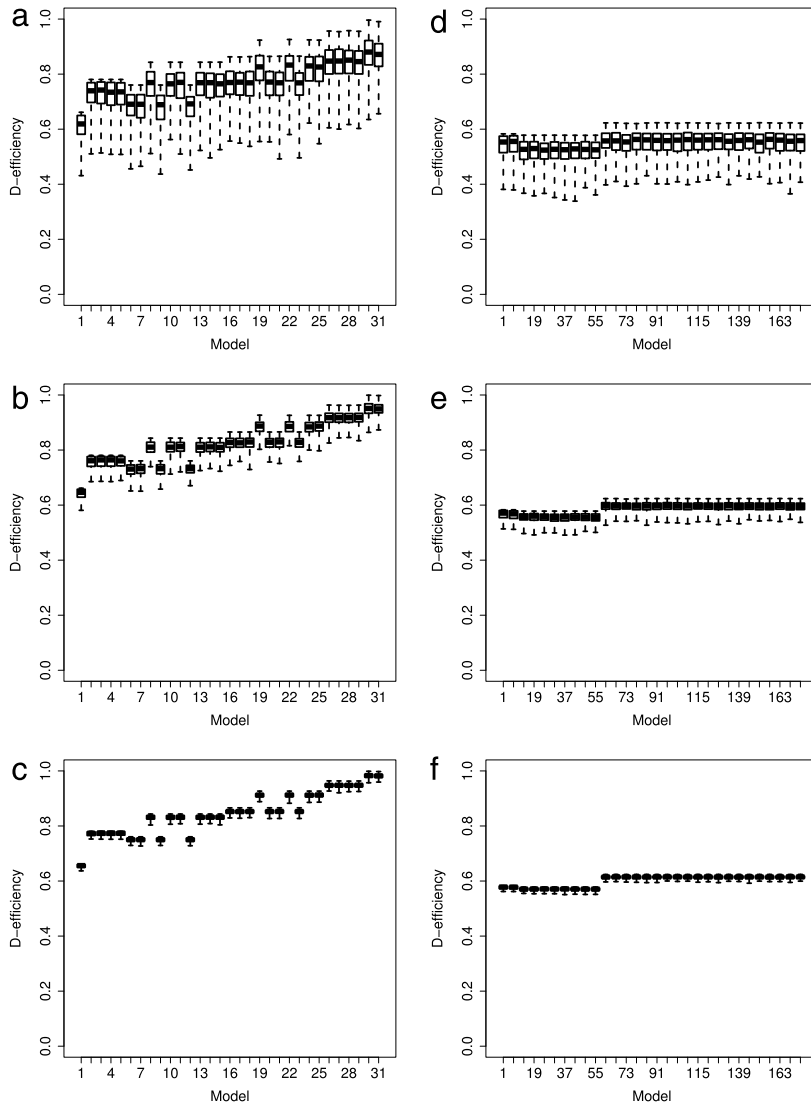


Fig. 8. Boxplots of D -efficiencies for robust designs for log-linear regression: five variables (a) $\kappa = 1$, (b) $\kappa = 2$ and (c) $\kappa = 3$; 10 variables (d) $\kappa = 1$, (e) $\kappa = 2$ and (f) $\kappa = 3$. For the 10 variable case, only results for every sixth model are displayed.

$i = 1, \dots, 5$ and for the remaining parameters by

$$\beta_{im} \sim \begin{cases} \text{Uniform}(\kappa, 5) & \text{for } i = 6, 8, 9 \text{ and } I(i, m) = 1, \\ \text{Uniform}(-5, -\kappa) & \text{for } i = 7, 10 \text{ and } I(i, m) = 1. \end{cases}$$

Again, $\beta_{0m} = 0, \beta_{im} = 0$ if $I(i, m) = 0$ and $\kappa = 1, 2, 3$.

4.1. Minimally-supported designs

We restrict attention to designs that are minimally supported with respect to the maximal model, that is, where the number, n , of distinct support points is $q + 1$. For this class of designs, [Russell et al. \(2009\)](#) and [McGree and Eccleston \(2012\)](#) presented analytical design construction methods. [Atkinson and Woods \(2015\)](#) showed that for these designs with $-1 \leq x_{ij} \leq 1$ and $E(\beta_{im}) \geq 1$, for any $m = 1, \dots, M$,

$$\int_{\mathcal{B}_M} \log \det \{X_m^T W_m X_m\} \pi_m(\beta_m) d\beta_m = \log \det \{X_m^T W_m^* X_m\},$$

where $W_m^* = \text{diag}\{e^{\eta_{jm}^*}\}$ with $\eta_{jm}^* = f(x_j)^T \beta_m^*$ and $\beta_m^* = E(\beta_m)$, the prior expectation of β_m . Hence, numerical integration is no longer required for design evaluation. To exploit the available theory, we find designs that maximise

$$\Phi_D(\xi) = \log \det \{X_M^T W_M^* X_M\}. \tag{13}$$

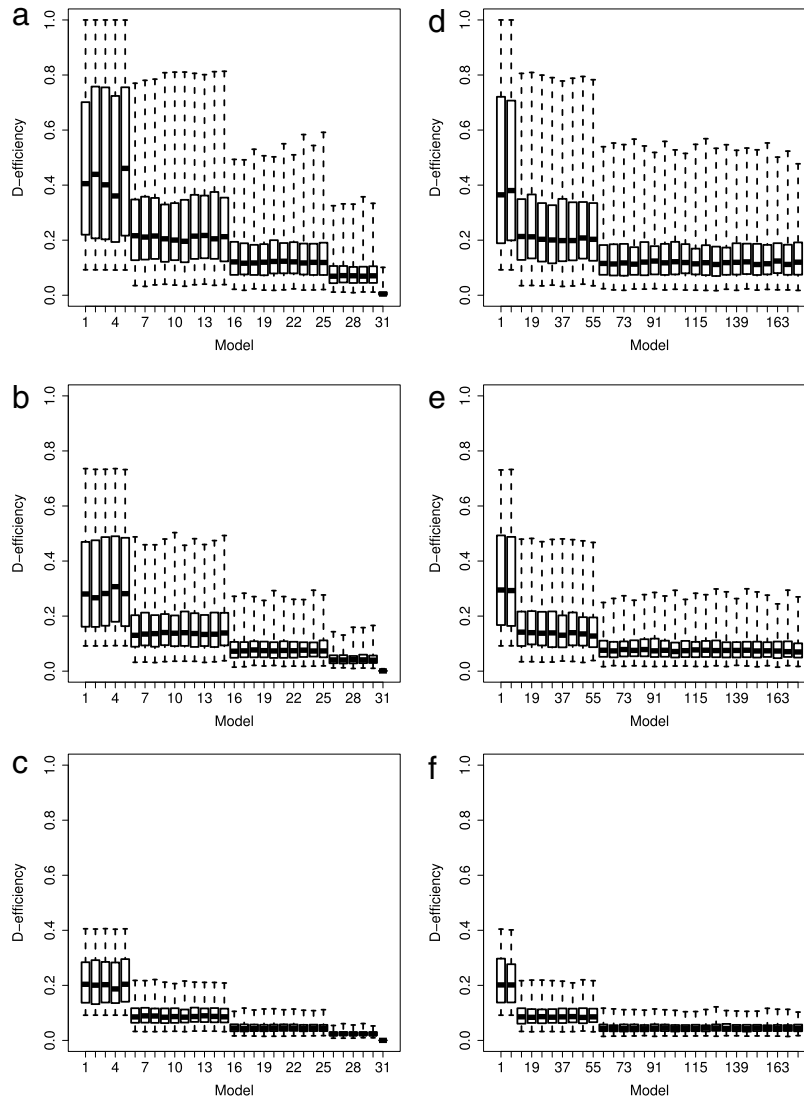


Fig. 9. Boxplots of D -efficiencies for fractional factorial designs for log-linear regression: five variables (a) $\kappa = 1$, (b) $\kappa = 2$ and (c) $\kappa = 3$; 10 variables (d) $\kappa = 1$, (e) $\kappa = 2$ and (f) $\kappa = 3$. For the 10 variable case, only results for every sixth model are displayed.

Maximisation of (13) defines a (pseudo-) Bayesian D -optimality criterion for the maximal model. A heuristic justification for using this criterion to find model-robust designs was given by McGree and Eccleston (2012) who pointed out that, assuming common prior distributions, the levels included for each variable in the minimally-supported Bayesian D -optimal design for each individual model m are the same. Only the numbers of replications of each variable value differ between the designs. Hence the sub-designs defined as projections of the minimally-supported design for the maximal model into a subset of the variables will contain the same values of the variables as a minimally-supported optimal design for that subset of variables but with different replication. Typically, designs defined in this way display less balance in the variable levels than the D -optimal designs for the sub-models. The advantage of maximising (13) is that no numerical optimisation is required for design selection, and hence large examples (e.g. 10 variables) can be investigated.

We replicate minimally-supported Bayesian D -optimal designs that maximise (13) to obtain designs with $N = 16$ runs for five variables and $N = 32$ runs for 10 variables. Fig. 8 summarises, for each model, the D -efficiencies (5) for the five variable and 10 variable designs, calculated as described in Section 3.1 except that, in step (ii), the locally D -optimal designs obtained from the theorem of Russell et al. (2009) are used. In general, the efficiencies are somewhat higher than those achieved by the equivalent five variable designs for logistic regression.

There are three main points of interest: (1) the D -efficiencies are higher for the five variable design due to the smaller number of variables in the maximal model leading to less imbalance in the variable values in the sub-designs; (2) for both the five-variable and 10-variable designs, the D -efficiency increases with the size of the model, reflecting the construction method of maximising (13) for the maximal model; and (3) the spread of the D -efficiencies decreases as κ increases, making




Fig. 10. Power, type I error rate and FDR for log-linear regression for $\kappa = 1$: (a) five variables and $N = 16$; (b) 10 variables and $N = 32$.

the prior distribution more concentrated. In both examples and for all κ values, the minimum efficiency is greater than ~ 0.4 , and the mean efficiency is greater than ~ 0.55 . For smaller κ and models with larger numbers of variables, the designs often have much higher D -efficiencies.

For this example, we also assess the performance of minimum aberration fractional factorial designs of resolution V (with $N = 16$ runs for five variables) and resolution IV (with $N = 32$ runs for 10 variables), see Fig. 9. Although these designs are D -optimal for the linear model, they perform uniformly poorly under log-linear regression, and much worse than the robust minimally-supported designs. Their efficiencies are particularly low for the larger values of κ , where the variance of the response is least constant.

4.2. Model selection results

Simulations to assess the performance of the designs for model selection were conducted as described in Section 3.2 except that, in step (iii), the model parameters were estimated using maximum likelihood and a model was chosen using AIC. Fig. 10 shows the results for five variable and 10 variable studies with $\kappa = 1$. Results in both cases are very encouraging, with almost uniformly high power and low type I error rates (< 0.2). For data-generating models with only one active variable (models 1–5 for the five variable experiment and models 1–10 for the 10 variable experiment), the truly active variable is occasionally missed, and another variable is identified as active. These errors lead to slightly lower power for these models, and non-zero FDR. For models with larger numbers of variables, all active variables are successfully identified (power equal to 1). For the five variable study, the FDR is consistently just below 0.3, corresponding to a maximum of about one non-active variable being incorrectly identified as active. For the 10 variable study, no screening errors are made for models containing three active variables (model 56 onwards). For both studies, the somewhat counter-intuitive result that performance improves for true models containing more active variables is explained by the construction method of the design (see Section 4.1), which focusses on the model containing the maximum number of variables.

For this example, the model selection performance of the two fractional factorial designs was very similar to that of the robust designs. Hence, the factorial designs would be effective for discrimination between the competing models but would provide poor estimation of the selected model.

5. Discussion and further research

This paper provides the first investigation of designs for screening variables under a generalised linear model. The results demonstrate that effective screening (high power with only moderate type I error rate and FDR) is achievable. For a binomial response and logistic regression, both design and model selection are more challenging than for a Poisson response, and larger designs are required to achieve good model selection results. For a binomial response, the results presented here can easily be extended to linear predictors that include products of variables representing interactions.

Future work is needed to investigate in more detail the use of the GIC penalty with maximum penalised likelihood. In some experiments, it may also be necessary to choose the link function in addition to the linear predictor. Compromise designs for this situation were found by Woods et al. (2006). In this paper, we have restricted the size of the model space under consideration by applying the principle of factor sparsity or by restricting the number of variables. For larger model spaces, the curse of dimensionality may prevent an all-subsets approach to model selection and alternative methods, such as sampling the model space (Smucker and Drew, 2015) or shrinkage regression (Friedman et al., 2010), would then need to be employed. Clearly, the choice of design, and any resultant “confounding” of model effects, will have an impact on any model selection procedure, including Bayesian and shrinkage methods. Investigations into the performance of these methods is another area for future work.

In the binomial and Poisson examples, we chose designs that ensured all models were estimable. This strategy is in contrast to the use of a design criterion tailored to model discrimination alone such as T -optimality (Atkinson and Fedorov, 1975), where the requirement of model estimability is often not met for nested models. An alternative approach is to generalise to multiple models those design selection methods that focus on both estimation and discrimination, such as the use of compound criteria (Atkinson, 2008) or hybrid designs (Waterhouse et al., 2008).

Acknowledgement

D. C. Woods was supported by Fellowship EP/J018317/1 from the UK Engineering and Physical Sciences Research Council. The authors thank an Associate Editor and two referees for comments that improved the paper.

References

- Akaike, H., 1974. A new look at the statistical model identification. *IEEE Trans. Automat. Control* 19, 716–723.
- Atkinson, A.C., 2008. DT -optimum designs for model discrimination and parameter estimation. *J. Statist. Plann. Inference* 138, 56–64.
- Atkinson, A.C., Donev, A.N., Tobias, R.D., 2007. *Optimum Experimental Designs*, with SAS, second ed. Oxford University Press, Oxford.
- Atkinson, A.C., Fedorov, V.V., 1975. The design of experiments for discriminating between two rival models. *Biometrika* 62, 57–70.
- Atkinson, A.C., Woods, D.C., 2015. Designs for generalized linear models. In: Dean, A.M., Morris, M.D., Stufken, J., Bingham, D.R. (Eds.), *Handbook of Design and Analysis of Experiments*. Chapman & Hall/CRC, Boca Raton.
- Box, G.E.P., Meyer, R.D., 1986. An analysis for unreplicated fractional factorials. *Technometrics* 28, 11–18.
- Burnham, K.P., Anderson, D.R., 2002. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*, second ed. Springer, New York.
- Chaloner, K., Larntz, K., 1989. Optimal Bayesian design applied to logistic regression experiments. *J. Statist. Plann. Inference* 21, 191–208.
- Chen, M.H., Huang, L., Ibrahim, J.G., Kim, S., 2008. Bayesian variable selection and computation for generalized linear models with conjugate priors. *Bayesian Anal.* 3, 585–613.
- Draguljić, D., Woods, D.C., Dean, A.M., Lewis, S.M., Vine, A.E., 2014. Screening strategies in the presence of interactions (with discussion). *Technometrics* 56, 1–28.
- Firth, D., 1993. Bias reduction of maximum likelihood estimates. *Biometrika* 80, 27–38.
- Friedman, J., Hastie, T., Tibshirani, R., 2010. Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* 33, URL <http://www.jstatsoft.org/index.php/jss/article/view/v033i01>.
- Gelman, A., Hwang, J., Vehtari, A., 2014. Understanding predictive information criteria for Bayesian models. *Stat. Comput.* 24, 997–1016.
- Haines, L.M., 1987. The application of the annealing algorithm to the construction of exact optimal designs for linear regression models. *Technometrics* 29, 439–447.
- Hamada, M., Nelder, J.A., 1997. Generalized linear models for quality-improvement experiments. *J. Qual. Technol.* 29, 292–304.
- Konishi, S., Kitagawa, G., 1996. Generalised information criteria in model selection. *Biometrika* 83, 875–890.
- Kosmidis, I., Firth, D., 2009. Bias reduction in exponential family nonlinear models. *Biometrika* 96, 793–804.
- Lemieux, C., 2009. *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer, New York.
- Lewis, S.L., Montgomery, D.C., Myers, R.H., 2001. Examples of designed experiments with nonnormal responses. *J. Qual. Technol.* 33, 265–278.
- Li, W., Nachtsheim, C., 2000. Model-robust factorial designs. *Technometrics* 42, 345–352.
- López-Fidalgo, J., Tommasi, C., Trandafir, P.C., 2007. An optimal experimental design criterion for discriminating between non-normal models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 69, 231–242.
- McCullagh, P., Nelder, J.A., 1989. *Generalized Linear Models*, second ed. Chapman & Hall/CRC, Boca Raton.
- McGree, J.M., Eccleston, J.A., 2012. Robust designs for Poisson regression models. *Technometrics* 54, 64–72.
- Murata, N., Yoshizawa, S., Amari, S., 1994. Network information criterion - determining the number of hidden units for an artificial neural network model. *IEEE Trans. Neural Netw.* 5, 865–872.
- Park, M.Y., Hastie, T., 2007. l_1 -regularization path algorithm for generalized linear models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 19, 659–677.
- Russell, K.G., Woods, D.C., Lewis, S.M., Eccleston, J.A., 2009. D -optimal designs for Poisson regression models. *Statist. Sinica* 19, 721–730.
- Silvapulle, M.J., 1981. On the existence of maximum likelihood estimates for the binomial response models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 43, 310–313.
- Smucker, B.J., Drew, N.M., 2015. Approximate model spaces for model-robust experiment design. *Technometrics* 57, 54–63.
- Sun, D.X., 1993. *Estimation capacity and related topics in experimental designs* (Ph.D. thesis), University of Waterloo, Canada.
- Waite, T.W., 2015. Singular prior distributions in Bayesian D -optimal design for nonlinear models. arXiv:1506.02916.
- Waterhouse, T.H., Woods, D.C., Eccleston, J.A., Lewis, S.M., 2008. Design selection criteria for discrimination/estimation for nested models and a binomial response. *J. Statist. Plann. Inference* 138, 132–144.
- Woods, D.C., 2010. Robust designs for binary data: applications of simulated annealing. *J. Stat. Comput. Simul.* 80, 29–40.
- Woods, D.C., Lewis, S.M., 2011. Continuous optimal designs for generalized linear models under model uncertainty. *J. Stat. Theory Pract.* 5, 137–145.
- Woods, D.C., Lewis, S.M., 2016. Design of experiments for screening. In: Ghanem, R., Hidgon, D., Owhadi, H. (Eds.), *Handbook of Uncertainty Quantification*. Springer, New York.
- Woods, D.C., Lewis, S.M., Eccleston, J.A., Russell, K.G., 2006. Designs for generalized linear models with several variables and model uncertainty. *Technometrics* 48, 284–292.
- Wu, C.F.J., 1993. Construction of supersaturated designs through partially aliased interactions. *Biometrika* 80, 661–669.
- Zhang, Y., Li, R., Tsai, C.L., 2010. Regularization parameter selections via generalized information criteria. *J. Amer. Statist. Assoc.* 105, 312–323.