

# Evaluation of combinatorial optimisation algorithms for c-optimal experimental designs with correlated observations

Watson, Samuel I.; Pan, Yi

DOI:

[10.1007/s11222-023-10280-w](https://doi.org/10.1007/s11222-023-10280-w)

License:

Creative Commons: Attribution (CC BY)

*Document Version*

Publisher's PDF, also known as Version of record

*Citation for published version (Harvard):*

Watson, SI & Pan, Y 2023, 'Evaluation of combinatorial optimisation algorithms for c-optimal experimental designs with correlated observations', *Statistics and Computing*, vol. 33, no. 5, 112.  
<https://doi.org/10.1007/s11222-023-10280-w>

[Link to publication on Research at Birmingham portal](#)

## General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

## Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact [UBIRA@lists.bham.ac.uk](mailto:UBIRA@lists.bham.ac.uk) providing details and we will remove access to the work immediately and investigate.



# Evaluation of combinatorial optimisation algorithms for c-optimal experimental designs with correlated observations

Samuel I. Watson<sup>1</sup> · Yi Pan<sup>1</sup>

Received: 19 July 2022 / Accepted: 3 July 2023  
© The Author(s) 2023

## Abstract

We show how combinatorial optimisation algorithms can be applied to the problem of identifying c-optimal experimental designs when there may be correlation between and within experimental units and evaluate the performance of relevant algorithms. We assume the data generating process is a generalised linear mixed model and show that the c-optimal design criterion is a monotone supermodular function amenable to a set of simple minimisation algorithms. We evaluate the performance of three relevant algorithms: the local search, the greedy search, and the reverse greedy search. We show that the local and reverse greedy searches provide comparable performance with the worst design outputs having variance < 10% greater than the best design, across a range of covariance structures. We show that these algorithms perform as well or better than multiplicative methods that generate weights to place on experimental units. We extend these algorithms to identifying modle-robust c-optimal designs.

**Keywords** Optimal design · Experimental design · Algorithms · Optimisation · GLMM

## 1 Introduction

We consider the question of how to identify a c-optimal design when the observations are correlated. In particular, we assume the data generating process can be described using a generalised linear mixed model (GLMM). For a  $N$ -length vector of outcomes  $y$ , with an  $N \times P$  matrix  $X$  of covariates and a  $N \times Q$  ‘design matrix’ for random effects  $Z$ , a GLMM can be written as:

$$y \sim F(h^{-1}(X\beta + Z\mathbf{u}), \phi) \quad (1)$$

where  $\beta$  are mean function parameters,  $h(\cdot)$  is a link function,  $F(\cdot)$  is a distribution function with scale parameter(s)  $\phi$ , and  $\mathbf{u} \sim N(0, D)$  is a vector of random effects with covariance matrix  $D$ . Such models provide a flexible parametric approach for the estimation of covariate effects when the observations are correlated, such as longitudinal designs (Zeger et al. 1988), cluster randomised trials (Hussey and

Hughes 2007), and geospatial statistical modelling (Diggle et al. 1998).

To set up the design problem, we assume the rows  $i = 1, \dots, N$  of matrices  $X$  and  $Z$  enumerate and completely define all the possible observations that could be made as part of our experimental study. Discrete design points could be formed by a uniform lattice over a continuous design space (e.g. Yang et al. 2013), for example, by determining a discrete set of potential sampling locations across an area of interest. Alternatively, observations may be possible with discrete groups or clusters at different times, such as in a cluster randomised trial. Indeed, GLMMs are frequently used in blocked designs or designs with groups or clusters of observations; we refer to such a group as an ‘experimental unit’. An experimental unit is defined here as a subset of the indices  $1, \dots, N$ :

$$e_j \subset \{1, \dots, N\}$$

for  $j = 1, \dots, J$ . Where we refer to the ‘observations’ in an experimental unit, we mean the rows of  $X$  and  $Z$  indexed by the indices in  $e_j$ . We assume  $\bigcup_j e_j = \{1, \dots, N\}$  and  $e_j \cap e_{j'} = \emptyset$  for  $j \neq j'$ , that is all observations (indices) are in one and only one experimental unit. Without loss of generality we assume that all the experimental units have the

---

The work was funded by MRC grant MR/V038591/1.

✉ Samuel I. Watson  
s.i.watson@bham.ac.uk

<sup>1</sup> Institute of Applied Health Research, University of Birmingham, Birmingham, UK

same size  $r$ . A *design space* consists of all the experimental units:

$$D := \{e_j : j = 1, \dots, J\}$$

A *design*  $d \subset D$  of size  $m < J$  is defined in terms of its experimental units:

$$d := \{e_j : e_j \in D; |d| = m\}$$

The total number of observations in the design is  $n = mr$ . The design problem we consider is how to find the design  $d^* \subset D$  of size  $m < J$  that minimises some objective function. There are multiple criteria used in the experimental design literature, here we only discuss c-optimality.

### 1.1 Information matrix

We consider GLMMs where  $F(\cdot)$  is in the exponential family, including Gaussian, binomial, and Poisson models. The likelihood is

$$L(\beta, \phi, \theta|y) = \int \prod_{i=1}^n f_{y|u}(y_i|\mathbf{u}, \beta, \phi) f_{\mathbf{u}}(\mathbf{u}|\theta) d\mathbf{u} \quad (2)$$

where  $f_{y|u}(y_i|\mathbf{u}, \beta, \phi) = \exp(y_i \eta_i - c(\eta_i))/a(\phi) + d(y_i, \phi)$ ,  $f_{\mathbf{u}}$  is the multivariate Gaussian density,  $\eta_i = x_i \beta + z_i \mathbf{u}$  is the linear predictor,  $x_i$  and  $z_i$  are the  $i$ th rows of matrices  $X$  and  $Z$ , and  $\theta$  are the parameters that define the covariance matrix  $D$ .

The information matrix for the model parameters  $\beta$  is:

$$M = X^T E_y \left[ \left( \frac{\partial \log L(\beta, \phi, \theta|y)}{\partial \beta} \right) \left( \frac{\partial \log L(\beta, \phi, \theta|y)}{\partial \beta} \right)^T | \beta, \theta, \phi \right] X \quad (3)$$

For the generalised least squares estimator, which is the best linear unbiased estimator of the GLMM, the information matrix is equivalent to

$$M = X^T \Sigma^{-1} X \quad (4)$$

where  $\Sigma = \text{Cov}(y)$ . We use  $M_d$  to denote the information matrix for design  $d$  formed by the rows of  $X$  and  $Z$  indexed by the indices in the experimental conditions in the design.

The c-optimal objective function  $g : 2^d \rightarrow \mathbb{R}_{\geq 0}$  we consider is then:

$$g(d) = \begin{cases} c^T M_d^{-1} c & \text{for } M_d \text{ positive semi-definite} \\ \infty & \text{otherwise} \end{cases} \quad (5)$$

where  $c$  is a  $p$ -length vector such that  $c \in \text{range}(M)$  to ensure estimability of  $c^T \beta$  (Pukelsheim 1980). The vector  $c$  often

has all elements equal to zero except for a one at the position corresponding to the treatment effect parameter of interest. It is possible for a design not to produce a positive semi-definite information matrix, such as when the matrix  $X$  is not of full rank, which will occur if there are too few observations, for example. For these cases we assume that the objective function takes the value of infinity, i.e. the design provides no information about the parameters. The consequences of this choice are discussed later. Our optimisation problem is the to find the design  $d \subset D$  of size  $m$  that minimises  $g(d)$ .

### 1.2 Previous literature

It is generally not possible to identify exact c-optimal designs in this context. For designs with correlated observations perhaps the most common approximate method is to identify the optimal ‘weights’ to place on each experimental unit or design point. These weights can be interpreted as an ‘amount of effort’ to place on the observations in the design. Elfving’s Theorem is a classic result in theory of optimal designs that underlies this method (Elfving 1952; Ford et al. 1992; Studden 2005). For independent identically distributed observations (i.e. when  $\Sigma = \sigma^2 I$ ) we can write the information matrix as the sum  $M = \sum_i x_i^T x_i$  for all observations  $i = 1, \dots, N$ . We can place a probability measure over the observations  $\rho = \{(\rho_i, x_i) : i = 1, \dots, N; \rho_k \in [0, 1]\}$  where  $\rho_i$  are the weights on each observation. The information matrix of the approximate design is  $M(\rho) = \sum_i x_i^T x_i \rho_i$ . Elfving’s theorem provides a geometric characterisation of c-optimality and shows that an optimal design  $\rho$  lies at the intersection of the convex hull of the  $x_i$  and the vector defined by  $c$ . The values of  $\rho$  for the optimal design can then be obtained using linear programming methods (Harman and Jurik 2008).

Holland-Letz et al. (2011) and Sagnol (2011) extended Elfving’s theorem to design spaces with experimental units in which the observations may be correlated, but where there is no correlation *between* experimental units. In this case, the information matrix (4) can be written as the sum of the information matrices for each experimental unit:

$$M = \sum_{j=1}^J X_{e_j}^T \Sigma_{e_j}^{-1} X_{e_j} \quad (6)$$

where we use  $X_{e_j}$  to represent the rows of  $X$  indexed by the indices in  $e_j$ , and similarly  $\Sigma_{e_j}$  is the submatrix of  $\Sigma$  given by the indices in  $e_j$ . The weights  $\rho$  are placed on each experimental unit so that a design can be represented by the pairs  $\{(e_1, \rho_1), \dots, (e_J, \rho_j)\}$  and the problem reduces to identifying the optimal weights using the generalised Elfving theorem. Sagnol (2011) shows that the optimal weights for each experimental condition can be solved using conic optimisation methods with a second order cone program. We refer to

these approaches generally as ‘multiplicative’. Holland-Letz et al. (2012) provide an estimate on the lower bound of the efficiency of multiplicative methods in a article examining optimal assignment of individuals to different dose schedules in a pharmacokinetic study. Prior approaches to optimal experimental designs with correlated observations relied on asymptotic arguments (Sacks and Ylvisaker 1968; Muller and Pázman 2003; Näther 1985).

One potential limitation of multiplicative methods is that there are several different approaches to rounding weights to integer totals of experimental units (Balinski and Young 2002). Pukelsheim and Rieder (1992) determine the optimal rounding scheme when at least one of each type of experimental unit is required. However, for many design problems this restriction is not necessary. As such, different rounding methods may produce different designs, which may not necessarily be optimal. Multiplicative methods currently also have the limitation that they cannot be extended to designs where there may be correlation *between* experimental units. We may also wish to accommodate restrictions in the design space, such as a maximum or minimum number (or weight) on particular experimental conditions given practical restrictions.

In this article, we show that a set of simple algorithms are applicable to the problem of identifying c-optimal experimental designs with correlated observations both within and between experimental units, and we compare their performance with a set of example study designs. These algorithms are combinatorial optimisation methods, which aim to select the optimal set of items from a larger finite discrete set. Where relevant, we refer to these methods as ‘combinatorial algorithms’ to differentiate them from the multiplicative methods. These algorithms can identify local minima, however, as we discuss, combinatorial approach cannot guarantee a global minimum is found. Results from the optimisation literature show that the difference between the solutions from the algorithms and global minima can be bounded though. We discuss the relevant combinatorial algorithms in Sect. 2 and show that they can be applied to the c-optimal design problem. We consider approximations to the information matrix to accommodate non-Gaussian models. In Sect. 3 we compare the performance of these algorithms across a set of example problems. Sect. 4 compares the performance of multiplicative and combinatorial approaches, and Sect. 5 extends the discussion to robust optimisation.

## 2 Monotone supermodular function minimisation

A function  $g$  is called *supermodular* if:

$$g(d \cup \{e_j\}) - g(d) \geq g(d' \cup \{e_j\}) - g(d') \tag{7}$$

for all  $d' \subseteq d$ . That is, there are diminishing marginal reductions in the function with increasing size of the design. The function is *monotone decreasing* if  $d' \subseteq d \rightarrow g(d') \geq g(d)$ . A function is called *submodular* if in Eq. (7) the inequality is reversed.

Equation (6) shows that when the observations in different experimental units are independent, then the information matrix can be written as a sum of information matrices for each unit. We can derive a more general expression for the marginal change to the information matrix when observations are added. Let  $d'$  and  $d$  be two designs such that  $d' \subset d \subset D$  and  $d = d' \cup d''$ . We let  $X_1$  and  $X_2$  be the covariate matrices for designs  $d'$  and  $d''$ , respectively, and  $\Sigma_1$  and  $\Sigma_2$  be their covariance matrices.  $\Sigma_{12}$  is the covariance between the observations in designs  $d'$  and  $d''$ . Then,

$$\begin{aligned} M_d &= \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}^T \begin{bmatrix} \Sigma_1 & \Sigma_{12} \\ \Sigma_{12}^T & \Sigma_2 \end{bmatrix}^{-1} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} \\ &= \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}^T \begin{bmatrix} \Sigma_1^{-1} - \Sigma_1^{-1} \Sigma_{12}^T S^{-1} \Sigma_{12} \Sigma_1^{-1} & -\Sigma_1^{-1} \Sigma_{12}^T S^{-1} \\ -S^{-1} \Sigma_{12} \Sigma_1^{-1} & S^{-1} \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} \\ &= X_1^T \Sigma^{-1} X_1 + X_2^T S^{-1} X_2 - X_1^T \Sigma_1^{-1} \Sigma_{12}^T S^{-1} X_2 \\ &\quad - X_2^T S^{-1} \Sigma_{12} \Sigma_1^{-1} X_1 + X_1^T \Sigma_1^{-1} \Sigma_{12}^T S^{-1} \Sigma_{12} \Sigma_1^{-1} X_1 \\ &= M_{d'} + [X_2 - \Sigma_{12}^T \Sigma_1^{-1} X_1]^T S^{-1} [X_2 - \Sigma_{12}^T \Sigma_1^{-1} X_1] \\ &= M_{d'} + \delta M_{d',d''} \end{aligned} \tag{8}$$

where  $S = (\Sigma_2 - \Sigma_{12}^T \Sigma_1^{-1} \Sigma_{12})$  is the Schur complement, which we assume is invertible, and we use  $\delta M_{d',d''}$  to represent the marginal change in the information matrix of  $M_{d'}$  when the additional observations in  $d''$  are added. It is evident that if  $\Sigma_{12} = 0$  then  $\delta M_{d',d''}$  reduces to  $X_2^T \Sigma_2^{-1} X_2$  as in Eq. (6). We also note that  $\delta M_{d',d''}$  is positive definite such that  $M_d \geq M_{d'}$  (where  $A \geq B$  means  $A - B$  is positive semi-definite), which implies that  $g(d) \leq g(d')$  for  $d' \subseteq d$ . Thus, the function  $g$  is monotone decreasing.

For our specific c-optimality problem, we can re-express (7) using Equation (8) as:

$$\begin{aligned} &c^T ([M_{d'} + \delta M_{d',d''} + \delta M_{d,e_j}]^{-1} - [M_{d'} + \delta M_{d',d''}]^{-1})c \\ &\geq c^T ([M_{d'} + \delta M_{d',e_j}]^{-1} - M_{d'}^{-1})c \\ &c^T [M_{d'} + \delta M_{d',d''} + (M_{d'} \\ &\quad + \delta M_{d',d''})\delta M_{d,e_j}^{-1}(M_{d'} + \delta M_{d',d''})]^{-1}c \\ &\leq c^T [M_{d'} + M_{d'}\delta M_{d',e_j}^{-1}M_{d'}]^{-1}c \end{aligned} \tag{9}$$

where the second line follows from Hua’s identity. The function is then supermodular if:

$$M_d[M_d^{-1} + \delta M_{d,e_j}^{-1}]M_d \geq M_{d'}[M_{d'}^{-1} + \delta M_{d',e_j}^{-1}]M_{d'} \tag{10}$$

It can be shown that this condition is satisfied if  $\delta M_{d,e_j}$  is symmetric, positive semidefinite for all  $d$ , which Eq. (8)

shows to be the case if the covariance matrix and Schur complement are symmetric and invertible. Thus, the c-optimal design problem for the GLMM under the GLS estimator is monotone supermodular.

## 2.1 Algorithms

### 2.1.1 Local search algorithm

Algorithm 1 shows the *local search* algorithm. We start with a random design of size  $m$  and at each step of the algorithm the best swap of an experimental unit in the design with one not in the design is made until there are no more swaps that improve the design.

For monotone supermodular function minimisation in general, there is no guarantee the local search will converge to the globally optimal design. However, the algorithm does have a provable ‘constant factor approximation’. If  $d$  is the output of the local search algorithm and  $d^*$  is the global minimiser of the function, then the constant factor approximation is the upper bound of  $g(d)/g(d^*)$ . Fisher et al. (1978) showed that under a cardinality constraint (such as  $|d| \leq n$ ) the optimality bound is  $3/2$ . Filmus and Ward (2014) improved this bound to  $(1 + 1/e)$  by using an auxiliary function in place of  $g$  that excludes poor local optima. Feige (1998) showed that further improving these bounds is an NP-hard problem and Nemhauser and Wolsey (1978) showed that algorithms that improve on these bounds require an exponential number of function evaluations, rather than the polynomial number of the local search. In practice the local algorithm, and those discussed below, may perform significantly better than their lower bounds would suggest, however there exists little empirical evidence for the types of design we consider. We also note that Fedorov (1972) developed the first local search algorithm for D-optimal designs (i.e. a design that maximises  $\det(M_d)$ ) with independent observations; several variants were later proposed (Nguyen and Miller 1992). Although there is presently no proof that such an approach converges to a D-optimal design.

**Data:**  $X, Z, \Omega, \beta, \phi, D$   
**Result:** An optimal design  $d^*$   
 Let  $d_0$  be size  $m$  design ;  
 Set  $\delta = 1$  and  $d \leftarrow d_0$  ;  
**while**  $\delta < 0$  **do**  
   **foreach**  $e_j \in d$  and  $e_{j'} \in D/d$  **do** Calculate  
    $g(d/\{e_j\} \cup \{e_{j'}\})$ ; Set  $d' \leftarrow \operatorname{argmin}_{j, j'} g(d/\{e_j\} \cup \{e_{j'}\})$  ;  
    $\delta = g(d') - g(d)$  ;  
   **if**  $\delta > 0$  **then**  
   |  $d \leftarrow d'$   
   **end**  
**end**

### 2.1.2 Greedy and reverse greedy search algorithms

Algorithm 2 shows the ‘greedy algorithm’. We start from the empty set ( $d = \emptyset$ ) and at each step of the algorithm add the experimental unit with the smallest marginal increase in the objective function. Rather than sequentially adding experimental units, one can start from the complete design space  $D$  and sequentially remove units. This is the ‘reverse greedy algorithm’, which is shown in Algorithm 3.

**Data:**  $X, Z, \Omega, \beta, \phi, D, n$   
**Result:** An optimal design  $d^*$   
 Let  $d$  be a non-degenerate design of size  $p \leq s < m$  ;  
 Set  $k = 0$  ;  
**while**  $k < m$  **do**  
   **foreach**  $e_j \in D/d$  **do** Calculate  $g(d \cup \{e_j\})$ ; Set  
    $d \leftarrow d \cup \operatorname{argmin}_{e_j} g(d \cup \{e_j\})$  ;  
    $k \leftarrow k + 1$   
**end**

**Data:**  $X, Z, \Omega, \beta, \phi, D, n$   
**Result:** An optimal design  $d^*$   
 Let  $d = D$  be the design containing all experimental conditions ;  
 Set  $k = J$  ;  
**while**  $k > m$  **do**  
   **foreach**  $e_j \in d$  **do** Calculate  $g(d/\{e_j\})$ ; Set  
    $d \leftarrow d/\operatorname{argmin}_{e_j} g(d/\{e_j\})$  ;  
    $k \leftarrow k - 1$   
**end**

The constant factor approximations for the greedy and reverse greedy algorithms are more complex than the local search case. In the case of submodular function maximisation, a famous result is that the constant factor approximation is  $1 + 1/e$  (Nemhauser and Wolsey 1978). However, this result does not carry over to minimising a supermodular function (Il’ev 2001). Indeed, it is not possible to implement the greedy algorithm for the design problems we discuss, as all designs with fewer than  $p$  observations, and many with more than  $p$  observations, will result in a non-positive semidefinite information matrix. We can start the algorithm from a random small design, as Algorithm 2 describes, but this of course would sacrifice any theoretical guarantees. Il’ev (2001) discusses the approximation factor for the reverse greedy algorithm in the case of minimising a supermodular function. The result depends on the ‘steepness’ or curvature of the function  $g$ , which is defined as:

$$\max_{e \in D} \frac{(g(e) - g(\emptyset)) - (g(D) - g(D/e))}{(g(e) - g(\emptyset))}$$

In Eq. (5) we specified that the function had infinite variance for the empty set, in which case the steepness would be one, which is equivalent to an unbounded curvature. In these cases the reverse greedy search does not have an approximation factor (Il'ev 2001; Sviridenko et al. 2017). Specifying the value of the function to be undefined would also fail to provide a bound.

Greedy algorithms have been used in the experimental design literature previously. Yang et al. (2013) developed a continuous sequential/greedy algorithm to identify optimal designs for generalised linear models under a range optimality criteria, although not for c-optimality. They discretised the design space using a regular lattice and showed convergence to optimal designs as the number of lattice cells grows. Variants and combinations of these methods have also been proposed, for example, the ‘Cocktail algorithm’ combines a sequential algorithm with two other algorithms in each step to find D-optimal designs (Yu 2011). Fedorov (1972) proposed a variant of this algorithm for D-optimal designs, often called a sequential algorithm, in which observations are sequentially added to an existing design until a convergence criterion identifying D-optimality is reached. Fedorov (1972) showed this algorithm produced a D-optimal design for linear models without cardinality constraint. Accelerated (or adaptive) greedy algorithms provide significant computational improvements on the standard greedy algorithm by avoiding recomputation of the objective function (Robertazzi and Schwartz 1989). Accelerated greedy algorithms have been used in experimental design, designing sensor networks, and other problems (Yang et al. 2019; Zou et al. 2016; Guo et al. 2019). However, again, there are few applications for c-optimality.

Given the lack of theoretical guarantees, one may consider these algorithms irrelevant to the c-optimal design problem. However, for some of the areas we use as examples below, they have been used anyway. For example, there has been growing interest in methods to identify c-optimal designs for cluster randomised trials (e.g. Girling and Hemming (2016); Hooper et al. (2020)). Several recent articles have used an algorithmic approach that involves sequential removal of observations from a design space to identify c-optimal cluster randomised trial designs (Hooper et al. 2020), or using the change in variance of treatment effect estimators when experimental units are removed to identify efficient designs (Kasza and Forbes 2019). While these algorithms lack theoretical guarantees, they may empirically still perform adequately for these design problems. They also run faster than the local search. So we include them in the empirical comparisons below.

## 2.2 Computation and approximation

### 2.2.1 Information matrix approximations

The greatest limitation on executing these algorithms is the evaluation of the information matrix (4) as it requires calculation and inversion of  $\Sigma$  or calculation of the gradient of the log likelihood. As we discuss in Sect. 2.2.2, once the covariance matrix is obtained, updating its inverse after adding or removing an observation can be done relatively efficiently. However, an efficient means of generating  $\Sigma$  is still required for non-linear models. Breslow and Clayton (1993) used the marginal quasilielihood of the GLMM to propose the following first-order approximation:

$$\Sigma \approx W^{-1} + ZDZ^T \tag{11}$$

where  $W$  is a diagonal matrix with entries  $W_{ii} = \left( \frac{\partial h^{-1}(\eta)}{\partial \eta} \right)^2 \text{Var}(y|\mathbf{u})$ , which are recognisable as the GLM iterated weights (McCullagh and Nelder 1989). The approximation is exact for the Gaussian model with identity link. Higher order approximations exist in the literature, however, their use has not been found to improve the quality of optimal designs, at least in the case of D-optimality. Waite and Woods (2015) consider the case of D-optimal designs and compare (11) with the GEE working covariance matrix. They find it does not perform as well as the approximation based on the marginal quasilielihood. We do not consider the GEE covariance in this article, as we aim to use explicit covariance functions with different parameterisations.

Zeger et al. (1988) suggests that when using the marginal quasilielihood a better approximation to the marginal mean can be found by “attenuating” the linear predictor in non-linear models. For example, with the log link the “attenuated” mean is  $E(y_i) \approx h^{-1}(x_i\beta + z_i D z_i^T / 2)$  and for the logit link  $E(y_i) \approx h^{-1}(x_i\beta \det(a D z_i^T z_i + I)^{-1/2})$  with  $a = 16\sqrt{3}/(15\pi)$ . Waite and Woods (2015) find that using attenuated parameters with the approximation (11) can achieve more efficient designs for D-optimality. For the non-linear models below we compare with and without attenuation. Waite and Woods (2015) also propose approximations based on (3). They consider blocked designs where there is no correlation between experimental units, and so the information matrix can be computed as the sum of information matrices of the units in the design as Eq. (8). For a binomial-logistic model, the information matrix can then be calculated using (3) by completely enumerating the outcome space if the size of the experimental unit is relatively small. However, we

consider designs where there may be correlation between experimental units, limiting the computational tractability of such an approach as a complete enumeration of the outcome space would be infeasible.

### 2.2.2 Algorithm efficiency

Equation (8) also shows how we can achieve some computational efficiency with correlated observations. A naive optimisation approach that recalculated the information matrix for each design would require at least  $O(n^3)$  operations by needing to invert the covariance matrix each time. We can iteratively add or remove single observations at a time, i.e. moving from  $d$  to  $d/\{i\} \cup \{i'\}$  and  $i \neq i'$ , so the calculation only requires  $O(rn^2)$  operations to add or remove an experimental unit through rank-1 up/down dates of the inverse covariance matrix (see ‘‘Appendix’’). The total running time of the local search algorithm scales with  $O(m^4r^3(J - m))$ , as we have to evaluate swapping  $m$  experimental units with  $J - m$  remaining units of size  $r$  up to  $m$  times. The experimental units are not always unique, in these cases we can detect any duplicated experimental conditions and only evaluate any swap involving it once, while keeping track of the number of copies in the design space and design, to reduce running time. As such the values of  $m$  and  $J$  in the expression for complexity can be interpreted as the numbers of *unique* experimental units in the design and design space, respectively.

The computational complexity of the greedy search algorithm scales as  $O(m^3r^3(J - m))$ , however it generally runs much faster than the local search since most function evaluations are of designs smaller than  $m$ . The complexity of the reverse greedy algorithm scales as  $O(J^3r^3(J - m))$ .

## 3 Comparative performance

### 3.1 Comparison of algorithms

We consider several examples to compare the three algorithms described above. We compare them in two areas: quality of solution and computational time. As the local and greedy search algorithms have a random starting set, we run them each 100 times; the reverse greedy search is deterministic and so is run only once per example.

We use the approximation (11) for all the analyses. For the examples using a Gaussian-identity model, the approximation is exact. For non-linear models (binomial-logistic, binomial-log, and poisson-log) we make the additional comparison between an approximation with and without attenuation as described in Sect. 2.2.1.

For each example we calculate the ‘relative efficiency’ as the ratio (expressed as a percentage) of the variance (i.e.

the value of  $c^T M_d^{-1} c$ ) of the design(s) from each algorithm compared to the variance of the best design from all algorithms. For the non-linear models, we only evaluate the single best design from each algorithm, with and without attenuated parameters, using Eq. (6). Enumerating the complete outcome space to evaluate the expectations in (6) would not be possible, so we use Monte Carlo integration with 100,000 iterations to estimate the relative variance.

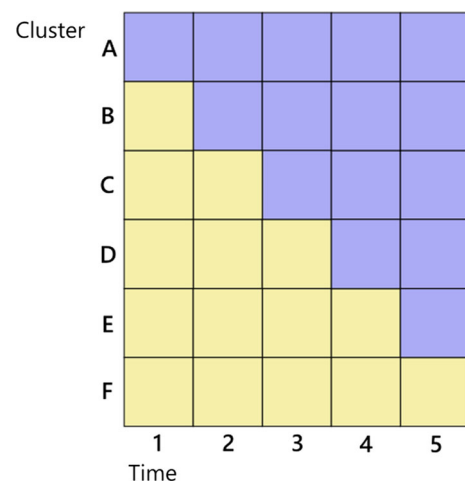
We also report the approximate running time of each algorithm. Timings were made on a computer with Intel Core i7-9700K, 32GB RAM, Windows 10. We make all these algorithms available as part of the `glmmlOptim` package for R. Code to reproduce the analyses in this article is included as a supplementary file.

### 3.2 Applied examples

Our examples are derived from two study types that motivated this article. Identifying optimal cluster randomised trial designs, and determining optimal sampling locations to estimate treatment effects in a geospatial setting. We describe each of these in turn along with their associated examples.

#### 3.2.1 Cluster randomised trial

A cluster randomised trial is a type of randomised trial design in which groups, or ‘clusters’, of individuals are randomly allocated to receive either a treatment or control. Cluster trials are typically used to evaluate interventions that are applied to groups of people rather than individuals, for example, quality improvement initiatives for healthcare clinics, or educational interventions in classrooms. Figure 1 describes a design space for a cluster randomised trial with repeated measures. Each



**Fig. 1** An illustration of a cluster randomised trial design space. Each cell represents a cluster-period and is comprised of observations from individuals within the cell. Yellow = control status with no intervention; blue = with intervention. (Color figure online)

**Table 1** Model specifications for the comparison examples for the cluster randomised trial

	Model			Parameters	
	<i>F</i>	<i>h</i>	Cov	$\beta$	$\theta$
<i>Cluster trial—cross-sectional</i>					
A	Gaussian	Identity	Exc	N/A	$\sigma_1 = 0.25 \sigma_2 = 0.1 \sigma_e = 1$
B	Gaussian	Identity	Exc	N/A	$\sigma_1 = 0.1 \sigma_2 = 0.1 \sigma_e = 1$
C	Gaussian	Identity	AR1	N/A	$\sigma_1 = 0.25 \lambda = 0.6 \sigma_e = 1$
D	Gaussian	Identity	AR1	N/A	$\sigma_1 = 0.1 \lambda = 0.9 \sigma_e = 1$
E	Binomial	Logit	Exc	$\beta_0 = 0.1$ $\beta_1 = [-0.5, -0.3, -0.1, 0.1, 0.3]$	$\sigma_1 = 0.25 \sigma_2 = 0.1$
F	Binomial	Logit	Exc	$\beta_0 = 0.1$ $\beta_1 = [-0.5, -0.3, -0.1, 0.1, 0.3]$	$\sigma_1 = 0.1 \sigma_2 = 0.1$
G	Binomial	Log	AR1	$\beta_0 = 0.1$ $\beta_1 = [-1.5, -1.3, -1.1, -0.9, -0.7]$	$\sigma_1 = 0.25 \lambda = 0.6$
H	Binomial	Log	AR1	$\beta_0 = 0.1$ $\beta_1 = [-1.5, -1.3, -1.1, -0.9, -0.7]$	$\sigma_1 = 0.1 \lambda = 0.9$
<i>Cluster trial—cohort</i>					
I	Gaussian	Identity	Exc	N/A	$\sigma_1 = 0.25 \sigma_2 = 0.1 \sigma_c^2 = 0.8 \sigma_e^2 = 0.2$
J	Gaussian	Identity	Exc	N/A	$\sigma_1 = 0.1 \sigma_2 = 0.1 \sigma_c^2 = 0.8 \sigma_e^2 = 0.2$
K	Gaussian	Identity	AR1	N/A	$\sigma_1 = 0.25 \lambda = 0.6 \sigma_c^2 = 0.8 \sigma_e^2 = 0.2$
L	Gaussian	Identity	AR1	N/A	$\sigma_1 = 0.1 \lambda = 0.9 \sigma_c^2 = 0.8 \sigma_e^2 = 0.2$
<i>Geospatial sampling</i>					
M	Gaussian	Identity	Exp	N/A	$\sigma_1 = 0.25 \lambda = 0.25$
N	Binomial	Logit	Exp	$\beta_0 = 0 \beta_1 = \log(2), \beta_2 = 4$	$\sigma_1 = 0.25 \lambda = 0.25$

Exc. exchangeable covariance function, AR1 autoregressive covariance function, Exp. exponential covariance function

cell is a cluster-period within which we can observe multiple individuals. There are  $K = 6$  clusters and  $T = 5$  total time periods. The linear predictor for an observation  $i$  in cluster  $k$  at time  $t$  is:

$$\eta_{ikt} = \beta_0 \Delta_{kt} + \beta_1 \tau_t + \epsilon_{ikt} \tag{12}$$

where  $\Delta_{kt}$  is a treatment indicator equal to one if the cell has the intervention and zero otherwise as shown in Fig. 1,  $\tau_t$  are  $T$  time-period indicators (so we do not include an intercept). We consider two covariance function specifications for  $\epsilon_{ikt}$  at the cluster-level and examine both cross-sectional and cohort designs. The covariance functions represent the most widely used specifications for these studies (Li et al. 2021). First, an exchangeable covariance function with cluster and cluster-period random effects with cross-sectional sampling in each time period (Hussey and Hughes 2007; Hemming et al. 2015):

$$Cov(\epsilon_{ikt}, \epsilon_{i'k't'}) = \begin{cases} \sigma_1^2 + \sigma_2^2 & \text{if } k = k', t = t' \\ \sigma_1^2 & \text{if } k = k', t \neq t' \\ 0 & \text{otherwise} \end{cases}$$

Second, an autoregressive covariance function:

$$Cov(\epsilon_{ikt}, \epsilon_{i'k't'}) = \begin{cases} \sigma_1^2 \lambda^{|t-t'|} & \text{if } k = k' \\ 0 & \text{otherwise} \end{cases}$$

For models with cohort effects where the same individuals appear in each cluster in every period, we modify these covariance functions by adding an additional term  $\sigma_c^2$  to the covariance if the individual is the same  $i = i'$ . For Gaussian models we notate the observation-level variance of the error term as  $\sigma_e^2$ .

Table 1 lists the different model specifications we examined for the cluster randomised trial and other examples. We include both linear and non-linear models with differing covariance structures and parameters. The choice of covariates generally represents a range from ‘high’ to ‘low’ levels of between-cluster correlation in these settings (Hemming et al. 2015, 2020). We specify a maximum number of individuals per cluster-period of 10, and aim to identify an optimal design of  $m = 100$  individuals (out of a possible 300); an experimental unit is a single observation from an individual. We set  $c = (1, 0, \dots, 0)^T$ .



**Table 2** Relative efficiency and approximate running time of the three algorithms for Examples using Gaussian-Identity model

	Local		Greedy		Reverse greedy	
	Rel. eff.	Time (s)	Rel. eff.	Time (s)	Rel. eff.	Time (s)
A	100.0–100.2	1–2	102.9–109.1	0–1	100.0	2
B	100.0–100.4	1–2	102.2–109.5	0–1	100.0	2
C	100.0–100.2	1–2	101.4–107.0	0–1	100.1	2
D	100.0–100.8	1–2	103.9–109.5	0–1	100.0	2
I	101.2–108.2	40–50	116.0–345.6	1–2	100.0	10
J	100.3–106.0	40–50	114.0–201.0	1–2	100.0	10
K	100.9–106.2	40–50	111.8–184.1	1–2	100.0	10
L	101.9–112.7	40–50	130.3–161.3	1–2	100.0	10
M	100.0–100.0	90–100	100.3–101.9	2–3	100.0	3

### 3.2.2 Geospatial sampling

There is a broad literature on selecting the optimal sampling locations (and times) to draw samples across an area of interest (e.g. Chipeta et al. 2017). However, these sampling patterns are generally designed to estimate a statistic like the prevalence of a disease across an area and its spatio-temporal distribution. A related, but possibly more complex, question asks where across an area one should sample to provide the most efficient estimates of point-source interventions with spatially-heterogeneous effects. A geospatial statistical model can be represented as a GLMM (Diggle et al. 1998), thus if the possible sampling locations are discretised, the design problem is amenable to the methods in this article.

Our design space is a unit-square  $A = [0, 1]^2$ . The space is divided into a regular  $15 \times 15$  lattice, where observations are made at the cell centroids  $a \in A$ . An intervention is located at the point  $z = (0.5, 0.5)$ . The mean function is specified as:

$$\eta_i(a) = \beta_0 + \beta_1 \exp(-\beta_2|a - z|) + \epsilon_a \tag{13}$$

To accommodate the non-linear mean function in the framework described above we use an additional first-order approximation to the information matrix [following Holland-Letz et al. (2011, 2012) and others]:

$$M_d = F^T \Sigma^{-1} F \tag{14}$$

where the first column of  $F$  is a vector of ones, the second column is  $\partial\mu/\partial\beta_1 = \exp(-\beta_2|a - z|)$ , and the third column is  $\partial\mu/\partial\beta_2 = -\beta_1|a - z|\exp(-\beta_2|a - z|)$ . We set  $\beta_0 = 1$ ,  $\beta_1 = \ln(2)$ , and  $\beta_2 = 4$ . We specify a Poisson distribution with log link function. Finally, we specify an exponential covariance function:

$$Cov(\epsilon_a, \epsilon_{a'}) = \sigma_1^2 \exp(-\lambda|a - a'|) \tag{15}$$

which is commonly used in geospatial applications. Our aim is to find an optimal design of size  $m = 80$  (of a total possible

325). We set  $c = (0, 1, 0.1)^T$ . Table 1 lists the parameter and model specifications for these examples.

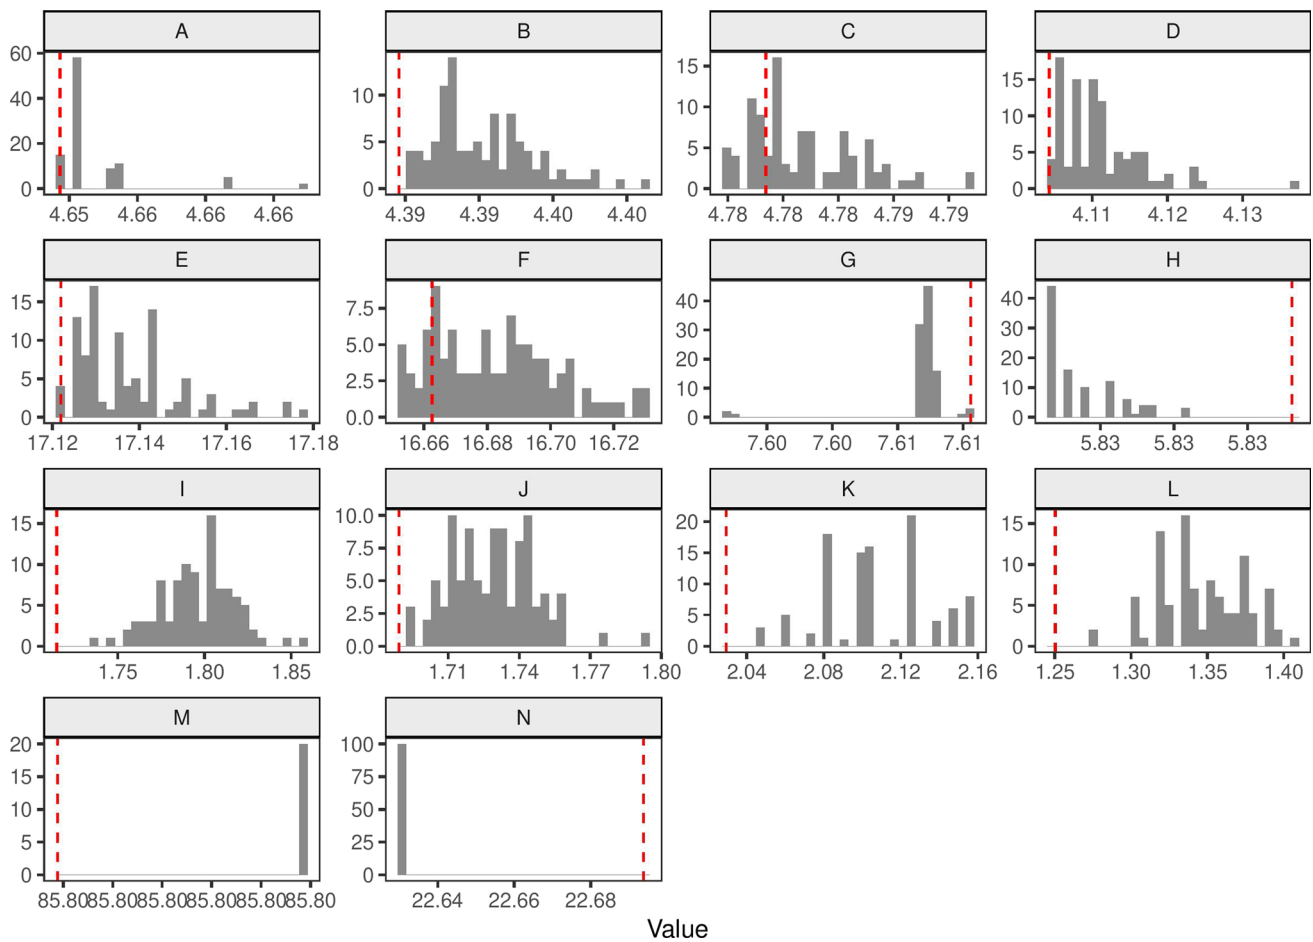
### 3.3 Results

Table 2 reports the relative efficiency and approximate computational times for the Gaussian examples. The greedy search performed the worst of the three algorithms in all examples. The local and reverse greedy searches both often found the optimal design, although the reverse greedy algorithm was more consistent with the worst design having a variance only 0.1% greater than the local search. Figure 2 shows the distribution of variances of the designs from the local search compared to the reverse greedy search. For the cross-sectional cluster trial designs the worst design produced by the local search had a variance less than 1% greater than the best design, but with the cohort cluster trial designs this rose to 12%. The same design was produced on every iteration for the geospatial example. Running times for these examples ranged from 1 to 100s. The local search scaled poorly. The cross-sectional cluster trial examples had 30 unique experimental units compared to 60 for the cohort design and 325 for the geospatial example.

Table 3 reports the results for the non-Gaussian examples. Similarly to the Gaussian examples, the local and reverse greedy searches performed the best with either one or both finding the best design. Attenuation appears to make little difference to the quality of the solutions in these examples. Figures 3 and 4 show the best design for examples E, H, and M. Notably, the optimal design for Example H is not symmetric owing to the heteroskedastic variance of the binomial-log model.

### 4 Comparison with multiplicative methods

An approximate ‘multiplicative’ design is characterised by a probability measure over unique experimental units  $\phi =$



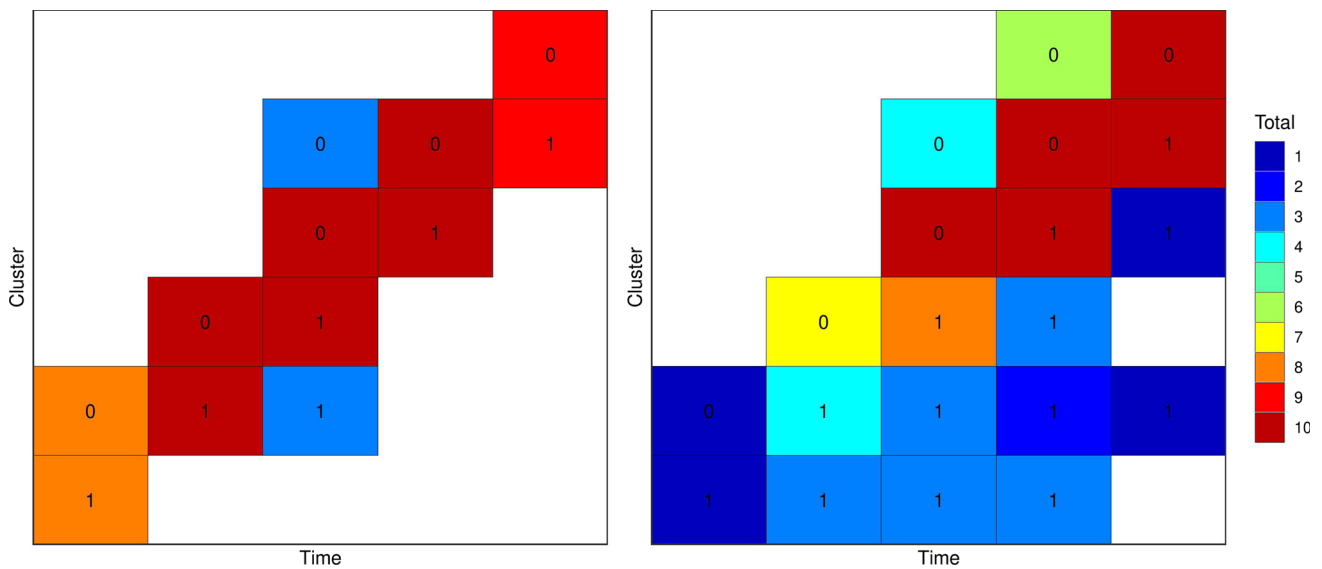
**Fig. 2** Histogram showing the values of the c-optimal criterion from the local search for each of the examples. The dashed red line indicates the value from the reverse greedy search. X-axis values are multiplied by 100

**Table 3** Relative efficiency and approximate running of the best designs from the three algorithms with and without attenuation (Att.) for the examples with non-Gaussian models

	Att	Local		Greedy		Reverse greedy	
		Rel.eff.	Time (s)	Rel.eff.	Time (s)	Rel.eff.	Time (s)
E	N	100.0	1–2	104.4	0–1	100.0	2
	Y	100.7	1–2	101.1	0–1	100.0	2
F	N	100.0	1–2	101.9	0–1	100.0	2
	Y	100.0	1–2	100.0	0–1	100.0	2
G	N	100.0	1–2	100.4	0–1	100.2	2
	Y	100.0	1–2	100.3	0–1	100.2	2
H	N	100.0	1–2	100.5	0–1	100.0	2
	Y	100.0	1–2	100.5	0–1	100.0	2
N	N	100.0	90–100	100.2	2–3	100.0	3
	Y	100.0	90–100	100.1	2–3	100.0	3

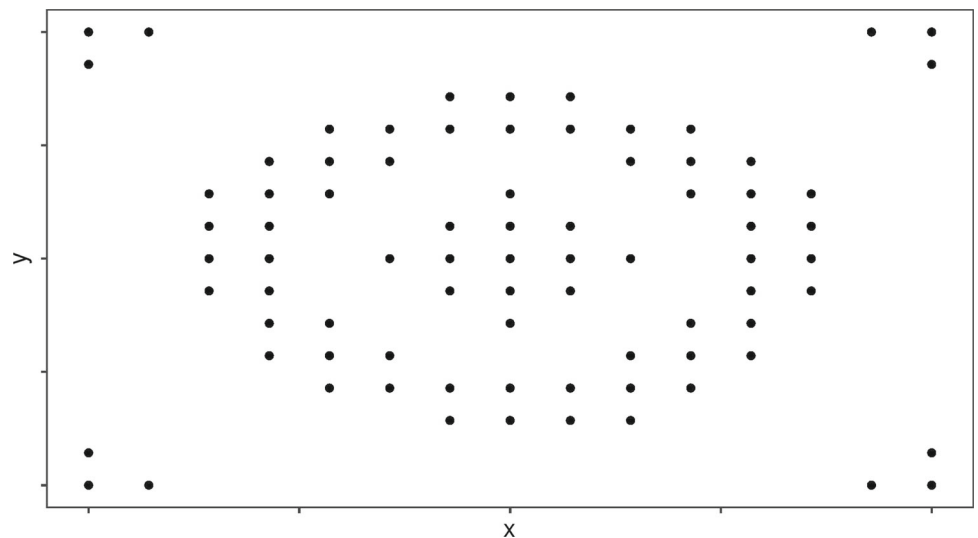
$\{(\rho_j, e_j) : j = 1, \dots, J; \rho_j \in [0, 1]; \sum_j \rho_j = 1\}$  where  $\rho_j$  are the weights associated with each experimental units. These weights can be identified for design spaces with uncorrelated experimental units (Holland-Letz et al. 2012; Sagnol 2011). However, there are several methods for rounding proportions to integer counts that sum to a given total (Balinski

and Young 2002). Briefly, for a given set of weights  $\rho_j$  and a target total of number of experimental units  $m$  in the design, the methods for determining the number of each experimental units  $m_j$  such that  $\sum_j m_j = m$ , are:



**Fig. 3** Optimal designs for Examples E and G. The 0 and 1 labels in the top row indicate the treatment status of the cluster period

**Fig. 4** Optimal geospatial sampling pattern for Example I



1. *Hamilton's method* Let  $\pi_j = n * \rho_j$ , we assign  $\lfloor \pi_j \rfloor$  of each experimental condition. The remaining total is filled by the experimental conditions with the largest remainders  $|\pi_j - \lfloor \pi_j \rfloor|$  until we have  $n$  experimental conditions.
2. *Divisor methods* Start with all  $n_j = 0$  and let  $\pi_j = n * \phi_j$ . Proceeding iteratively, we choose the next experimental condition in the design to be that with  $\max_j \pi_j / \alpha(n_j)$ , for which we update the total until the condition  $\sum_j n_j = n$  is met.
  - (a) *Jefferson's method*  $\alpha(n_j) = n_j + 1$
  - (b) *Webster's method*  $\alpha(n_j) = n_j + 0.5$
  - (c) *Adam's method*  $\alpha(n_j) = n_j$ . Initially we include one of each experimental condition with  $\pi_j > 0$ .

Pukelsheim and Rieder (1992) showed that the optimal rounding method was a variant of Adam's method, although under the assumption that there is at least one of each experimental unit. However, for many experimental design problems this assumption is not required, including the example we examine below. To compare the performance of the multiplicative methods discussed in Sect. 1.2, we take the best design from the three combinatorial algorithms, and compare it to the design with the lowest variance from across all rounding methods. To identify the optimal approximate design  $\rho$  we use the second-order cone program proposed by Sagnol (2011), which is also implemented in the R package `glmnrOptim` along with the rounding methods.

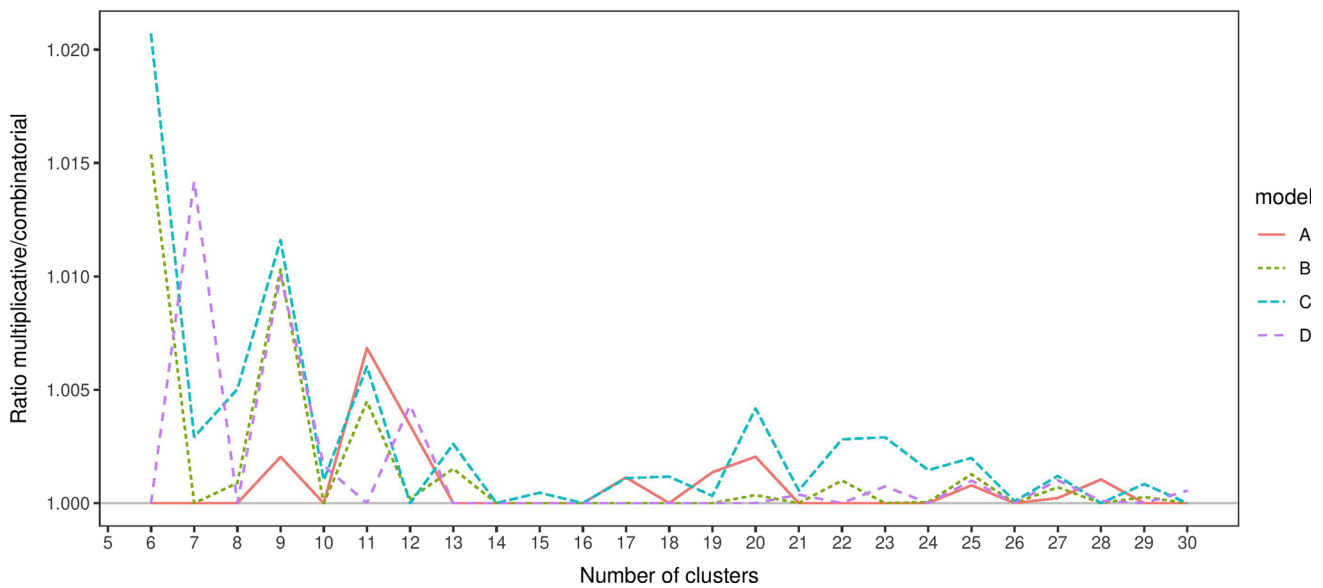


Fig. 5 Ratio of the variance of the design from each method to the highest variance for 6 to 30 clusters in example D

### 4.1 Example

We return to the cluster randomised trial examples, setting the experimental unit to a whole cluster sequence, i.e. a row consisting of five time periods each with ten individuals in Fig. 1. We wish to include  $m$  clusters in the design, each of which may be assigned to any of the experimental conditions, and consider  $m = 6$  to  $m = 30$ . We use the models specified as Examples A–D in Table 1.

Figure 5 presents the ratio of the variance of the best design from the multiplicative algorithm to the variance of the best design from the combinatorial algorithms. For smaller sample sizes, particularly with an odd sample size, the multiplicative algorithm performed worst than the combinatorial approach, although the worst design had a variance only 2% larger than the best design. For larger sample sizes the two methods generally produced the same optimal design.

## 5 Robust optimal designs

The analysis and discussion so far has been based on the assumption that the correct model specification is known. However, this is typically an unrealistic assumption, and it is well known that an optimal design for one model may perform very poorly for an alternative model. We consider a model robust optimal design criterion that is amenable to combinatorial optimisation. Following Dette (1993), we assume that the “true” model belongs to a known class of GLMMs. We can define a model with the collection  $G = (F, h, \beta, \theta)$ . The class of models is then

$$\mathcal{G}_U = \{G_1, \dots, G_U\} \tag{16}$$

The vector  $\rho = (\rho_1, \dots, \rho_U)$  with  $\rho > 0$  is then called a prior for the class  $\mathcal{G}$  with the values reflecting the belief about the relative probability or adequacy of each model. The objective function (5) can be written for a specific model as  $g_u(d) = g(d; G_u)$  and we can define the robust objective function as:

$$h(d) = \sum_{u=1}^U \rho_u g_u(d) \tag{17}$$

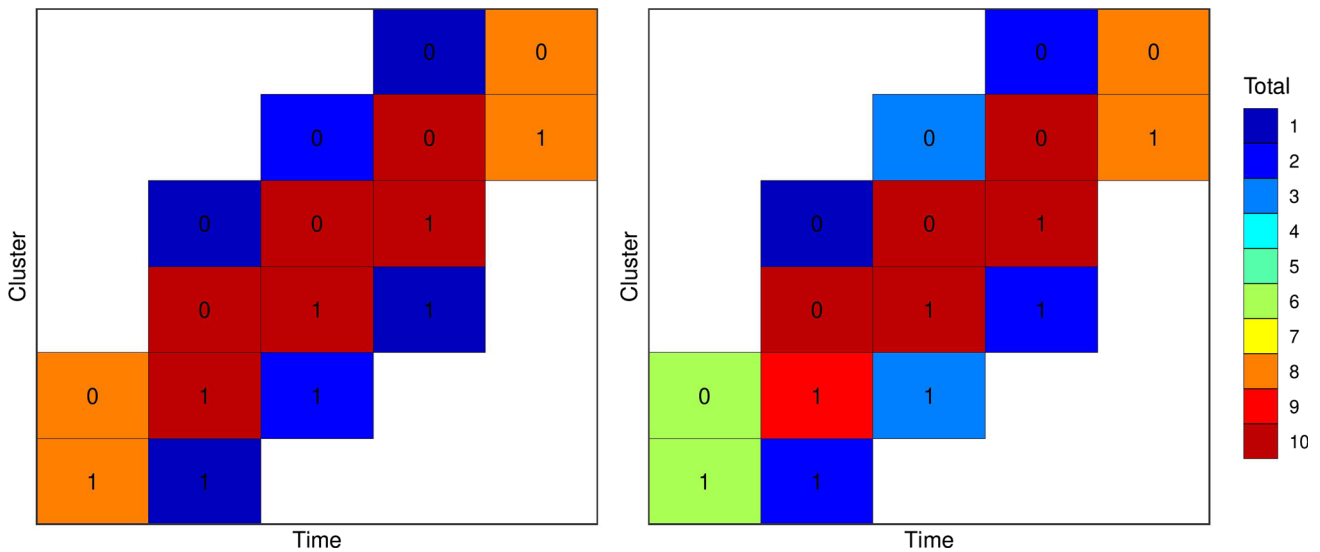
The objective function (17) is also monotone supermodular since  $\sum_{u=1}^U \rho_u g_h(d \cup \{e\}) - \sum_{u=1}^U \rho_u g_h(d) = \sum_{u=1}^U \rho_u [g_h(d \cup \{e\}) - g_h(d)] \leq \sum_{u=1}^U \rho_u [g_h(d' \cup \{e\}) - g_h(d')]$  and  $\sum_{u=1}^U \rho_u g_h(d) \geq \sum_{u=1}^U \rho_u g_h(d')$  for  $d' \subseteq d$  if all the  $g_u$  are themselves monotone submodular. We can therefore use the algorithms described above. A design minimising (17) is then said to be optimal for  $\mathcal{G}_U$  over the prior  $\rho$ . We note that other robust specifications such as minimax ( $h(d) = \max_u g_u(d)$ ) are not supermodular, so we do not consider them here.

Dette (1993) provides a geometric characterisation of the model robust criterion for c-optimal designs with uncorrelated observations, using the objective function  $h(d) = \sum_u \rho_u \log(g_u(d))$ , building on similar work for D-optimal designs. However, this has not yet been extended to correlated experimental units to permit use of multiplicative methods in this context.

We examine two examples for the robust optimal design. The examples A–D and E–H are taken as two classes of models. We assign equal weight to each design in each class

**Table 4** Relative efficiency of the best designs from the three algorithms with and without attenuation for Examples E–H and J with non-Gaussian models

	Att	Local		Greedy		Reverse greedy	
		Rel.eff.	Time (s)	Rel.eff.	Time (s)	Rel.eff.	Time (s)
A–D	N/A	100.0–100.3	2–4	102.9–109.0	0–1	100.0	7
E–H	N	100.0	2–4	104.3	0–1	100.0	6
	Y	100.0	2–4	104.1	0–1	100.0	6



**Fig. 6** Approximate robust and Bayesian optimal designs for examples D–G. The 0 and 1 labels in the top row indicate the treatment status of the cluster period

as the prior. For the class E–H, we use both attenuated and non-attenuated linear predictors for the approximation.

### 5.1 Results

Table 4 shows the results for the two robust optimal design examples. The results reflect those from all the previous examples: the greedy search performs relatively poorly with variances up to 10% larger than the best design. Both the local and reverse greedy searches identify the best design in each class. Figure 6 shows the model-robust optimal designs.

## 6 Conclusion

In this article, we have showed that the  $c$ -optimal design criterion is a monotone supermodular function for GLMMs using the GLS information matrix. We evaluated the performance of three supermodular function minimisation algorithms to identify  $c$ -optimal experimental designs with correlated experimental units. The theoretical upper bound on the relative variance of a design from the local search algorithm is 1.5 and no bound exists for the greedy and reverse greedy algorithms, however, for the examples we considered the performance is significantly better than 1.5 times the best

design. The greedy algorithm performed the worst, which was to be expected given that it cannot be executed fully as it cannot start from the empty set. The local search and reverse greedy searches performed comparably in terms of their best designs, although the local search could produce designs with variance more than 10% larger than the output of the reverse greedy search. Thus, the local search needs running multiple times to provide a reliable output. The local search also had poorer scaling in terms of computation time than the reverse greedy. Thus, while the reverse greedy search lacks a theoretical guarantee, it would be favoured empirically for the types of study design considered here.

We showed that the algorithms could also be applied for model robust optimal design identification using a weighted average design criterion. The method uses a prior, specifying the weights to place on each possible model. This specification suggests a way of applying these algorithms for use in Bayesian optimal design. Chaloner and Verdinelli (1995) comprehensively review Bayesian experimental design criteria and show the Bayesian  $c$ -optimality function to be:

$$\int \int (c^T [M_d + V_0]^{-1} c) p(\theta) p(\beta) d\theta d\beta$$

where  $V_0$  is the prior covariance of the  $\beta$  parameters and  $p(\theta)$  and  $p(\beta)$  the prior density function for the covariance and

linear predictor parameters, respectively. One can approximate the integral above using a Riemann sum, which would discretise the parameter space and provide a set of weights to place on each model. Recent advances have generated general algorithms for Bayesian optimal design problems with non-linear models, in particular Overstall and Woods (2017). Further research is needed to determine whether an approximation using the simple algorithms in this paper provide a viable or useful alternative to more advanced approaches.

We cannot guarantee that the optimal design was included in the output of any of the algorithms. However, our comparison with other multiplicative methods provides some reassurance. For designs with correlation within but not between experimental units, deriving weights using multiplicative methods for each experimental unit provides one method of approximating an optimal design (Holland-Letz et al. 2011; Sagnol 2011). Combinatorial approaches produced the same or better designs in the examples we considered.

Optimal designs may sometimes be impractical or difficult to implement. The design in the right panel of Fig. 3 is highly unlikely to ever be implemented. However, being able to identify approximately optimal designs provides a benchmark against which to justify proposed experiments. Many types of study that can be described by GLMMs, such as cluster randomised trials or spatio-temporal sampling across an area, can be significant and expensive undertakings. The combinatorial algorithms provide a means of identifying near-optimal or optimal designs to support their planning. Many design problems are not inherently discrete; however, we can discretise the design space by specifying a set of design points (Yang et al. 2013). Thus, the methods evaluated in this article provide a useful set of tools to support study design.

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1007/s11222-023-10280-w>.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## Appendix

### Rank-1 down/up dating to remove/add observations

#### Removing an observation

For a design  $d$  with  $m$  observations with inverse covariance matrix  $\Sigma_d^{-1}$  we can obtain the inverse of the covariance matrix of the design with one observation removed  $d' = d/\{i\}$ ,  $\Sigma_{d'}^{-1}$  as follows. Without loss of generality we assume that the observation to be removed is the last row/column of  $\Sigma_d^{-1}$ . We can write  $\Sigma_d^{-1}$  as

$$\Sigma_d^{-1} = \begin{pmatrix} C & d \\ d^T & e \end{pmatrix} \tag{18}$$

where  $C$  is the  $(m - 1) \times (m - 1)$  principal submatrix of  $B$ ,  $d$  is a column vector of length  $(m - 1)$  and  $e$  is a scalar. Then,

$$G = \Sigma_{d/(i)}^{-1} = C - dd^T/e \tag{19}$$

#### Adding an observation

For a design  $d$  with  $m$  observations with inverse covariance matrix  $\Sigma_d^{-1}$ , we aim now to obtain the inverse covariance matrix of the design  $d' = d \cup \{i'\}$ . Recall that  $Z$  is a  $R \times Q$  design effect matrix with each row corresponding to a possible observation. We want to generate  $H^{-1} = \Sigma_{d'}^{-1}$ . Note that:

$$H = \Sigma_{d'} = \begin{pmatrix} G^{-1} & f \\ f^T & h \end{pmatrix} \tag{20}$$

where  $f = Z_{i \in d} D Z_{i'}$  is the column vector corresponding to the elements of  $\Sigma = W^{-1} + Z D Z^T$  with rows in the current design and column corresponding to  $i'$ , and  $h$  is the scalar  $W_{i',i'}^{-1} + Z_{i'} D Z_{i'}^T$ . Also now define:

$$H^* = \begin{pmatrix} \Sigma_d & 0 \\ 0 & h \end{pmatrix} \tag{21}$$

so that

$$H^{*-1} = \begin{pmatrix} \Sigma_d^{-1} & 0 \\ 0 & 1/h \end{pmatrix} \tag{22}$$

and

$$H^{**} = \begin{pmatrix} \Sigma_d & f \\ 0 & h \end{pmatrix} \tag{23}$$

and  $u = (f^T, 0)^T$  and  $v = (0, \dots, 0, 1)^T$ , both of which are length  $m$  column vectors. So we can get  $H^{**}$  from  $H^*$

using a rank-1 update as  $H^{**} = H^* + uv^T$  and similarly  $H = H^{**} + vu^T$ . Using the Sherman-Morison formula:

$$H^{**^{-1}} = H^{*^{-1}} - \frac{H^{*^{-1}}uv^T H^{*^{-1}}}{1 + v^T H^{*^{-1}}u} \quad (24)$$

and

$$H^{-1} = H^{**^{-1}} - \frac{H^{**^{-1}}vu^T H^{**^{-1}}}{1 + u^T H^{**^{-1}}v} \quad (25)$$

So we have calculated the updated inverse with only matrix-vector multiplication, which is  $O(n^2)$ .

## References

- Balinski, M., Young, P.: Fair Representation: Meeting the Ideal of One Man, One Vote, 2nd edn. Brookings Institution Press, Washington (2002)
- Breslow, N.E., Clayton, D.G.: Approximate inference in generalized linear mixed models. *J. Am. Stat. Assoc.* **88**(421), 9–25 (1993). <https://doi.org/10.1080/01621459.1993.10594284>
- Chaloner, K., Verdinelli, I.: Bayesian experimental design: a review. *Stat. Sci.* **10**, 273–304 (1995)
- Chipeta, M., Terlouw, D., Phiri, K., Diggle, P.: Inhibitory geostatistical designs for spatial prediction taking account of uncertain covariance structure. *Environmetrics* **28**(1), e2425 (2017). <https://doi.org/10.1002/env.2425>
- Dette, H.: Elfving's theorem for  $\$D\$$ -optimality. *Ann. Stat.* (1993). <https://doi.org/10.1214/aos/1176349149.full>
- Diggle, P.J., Tawn, J.A., Moyeed, R.A.: Model-based geostatistics (with discussion). *J. R. Stat. Soc. Ser. C* **47**(Part 3), 299–350 (1998)
- Elfving, G.: Optimum allocation in linear regression theory. *Ann. Math. Stat.* **23**(2), 255–262 (1952)
- Fedorov, V.: Theory of Optimal Experiments. Academic Press, New York (1972)
- Feige, U.: A threshold of  $\ln n$  for approximating set cover. *J. ACM* **45**(4), 634–652 (1998). <https://doi.org/10.1145/285055.285059>
- Filmus, Y., Ward, J.: Monotone submodular maximization over a matroid via non-oblivious local search. *SIAM J. Comput.* **43**(2), 514–542 (2014). <https://doi.org/10.1137/130920277>
- Fisher, M.L., Nemhauser, G.L., Wolsey, L.A.: An analysis of approximations for maximizing submodular set functions-II. In: *Polyhedral Combinatorics. Mathematical Programming Studies*, vol. 8. Springer (1978)
- Ford, I., Torsney, B., Wu, C.F.J.: The use of a canonical form in the construction of locally optimal designs for non-linear problems. *J. R. Stat. Soc. Ser. B (Methodol.)* **54**, 569–583 (1992)
- Girling, A.J., Hemming, K.: Statistical efficiency and optimal design for stepped cluster studies under linear mixed effects models. *Stat. Med.* **35**(13), 2149–2166 (2016). <https://doi.org/10.1002/sim.6850>
- Guo, Y., Dy, J., Erdogmus, D., Kalpathy-Cramer, J., Ostmo, S., Campbell, J.P., Chiang, M.F., Ioannidis, S.: Accelerated experimental design for pairwise comparisons. In: *Proceedings of the 2019 SIAM International Conference on Data Mining*, Society for Industrial and Applied Mathematics, Philadelphia, PA, pp. 432–440 (2019). <https://doi.org/10.1137/1.9781611975673.49>
- Harman, R., Jurik, T.: Computing  $\epsilon$ -optimal experimental designs using the simplex method of linear programming. *Comput. Stat. Data Anal.* **53**(2), 247–254 (2008)
- Hemming, K., Kasza, J., Hooper, R., Forbes, A., Taljaard, M.: A tutorial on sample size calculation for multiple-period cluster randomized parallel, cross-over and stepped-wedge trials using the Shiny CRT calculator. *Int. J. Epidemiol.* (2020). <https://doi.org/10.1093/ije/dyz237/5748155>
- Hemming, K., Lilford, R., Girling, A.J.: Stepped-wedge cluster randomised controlled trials: a generic framework including parallel and multiple-level designs. *Stat. Med.* **34**, 181–196 (2015). <https://doi.org/10.1002/sim.6325>
- Holland-Letz, T., Dette, H., Pepelyshev, A.: A geometric characterization of optimal designs for regression models with correlated observations. *J. Stat. Soc. Ser. B (Stat. Methodol.)* **73**(2), 239–252 (2011). <https://doi.org/10.1111/j.1467-9868.2010.00757.x>
- Holland-Letz, T., Dette, H., Renard, D.: Efficient algorithms for optimal designs with correlated observations in pharmacokinetics and dose-finding studies. *Biometrics* **68**, 138–145 (2012)
- Hooper, R., Kasza, J., Forbes, A.: The hunt for efficient, incomplete designs for stepped wedge trials with continuous recruitment and continuous outcome measures. *BMC Med. Res. Methodol.* **20**(1), 279 (2020). <https://doi.org/10.1186/s12874-020-01155-z>
- Hussey, M.A., Hughes, J.P.: Design and analysis of stepped wedge cluster randomized trials. *Contemp. Clin. Trials* **28**(2), 182–191 (2007)
- Il'ev, V.P.: An approximation guarantee of the greedy descent algorithm for minimizing a supermodular set function. *Discrete Appl. Math.* **114**(1–3), 131–146 (2001)
- Kasza, J., Forbes, A.B.: Information content of cluster-period cells in stepped wedge trials. *Biometrics* **75**(1), 144–152 (2019). <https://doi.org/10.1111/biom.12959>
- Li, F., Hughes, J.P., Hemming, K., Taljaard, M., Melnick, E.R., Hegerter, P.J.: Mixed-effects models for the design and analysis of stepped wedge cluster randomized trials: an overview. *Stat. Methods Med. Res.* **30**(2), 612–639 (2021). <https://doi.org/10.1177/0962280220932962>
- McCullagh, P., Nelder, J.A.: *Generalized Linear Models*, 2nd edn. Routledge, London (1989)
- Muller, W.G., Pázman, A.: Measures for designs in experiments with correlated errors. *Biometrika* **90**(2), 423–434 (2003). <https://doi.org/10.1093/biomet/90.2.423>
- Näther, W.: Exact designs for regression models with correlated errors. *Statistics* **16**(4), 479–484 (1985). <https://doi.org/10.1080/02331888508801879>
- Nemhauser, G.L., Wolsey, L.A.: Best algorithms for approximating the maximum of a submodular set function. *Math. Oper. Res.* **3**(3), 177–188 (1978). <https://doi.org/10.1287/moor.3.3.177>
- Nguyen, N.-K., Miller, A.J.: A review of some exchange algorithms for constructing discrete D-optimal designs. *Comput. Stat. Data Anal.* **14**(4), 489–498 (1992)
- Overstall, A.M., Woods, D.C.: Bayesian design of experiments using approximate coordinate exchange. *Technometrics* **59**(4), 458–470 (2017). <https://doi.org/10.1080/00401706.2016.1251495>
- Pukelsheim, F.: On linear regression designs which maximize information. *J. Stat. Plan. Inference* **4**, 339–364 (1980)
- Pukelsheim, F., Rieder, S.: Efficient rounding of approximate designs. *Biometrika* **79**(4), 763 (1992)
- Robertazzi, T.G., Schwartz, S.C.: An accelerated sequential algorithm for producing D-optimal designs. *SIAM J. Sci. Stat. Comput.* **10**(2), 341–358 (1989). <https://doi.org/10.1137/0910022>
- Sacks, J., Ylvisaker, D.: Designs for regression problems with correlated errors: many parameters. *Ann. Math. Stat.* **39**(1), 49–69 (1968)
- Sagnol, G.: Computing optimal designs of multiresponse experiments reduces to second-order cone programming. *J. Stat. Plan. Inference* **141**(5), 1684–1708 (2011)
- Studden, W.: Elfving's theorem revisited. *J. Stat. Plan. Inference* **130**(1–2), 85–94 (2005)
- Sviridenko, M., Vondrák, J., Ward, J.: Optimal approximation for submodular and supermodular optimization with bounded curvature.

- Math. Oper. Res. **42**(4), 1197–1218 (2017). <https://doi.org/10.1287/moor.2016.0842>
- Waite, T.W., Woods, D.C.: Designs for generalized linear models with random block effects via information matrix approximations. *Biometrika* **102**(3), 677–693 (2015)
- Yang, J., Ban, X., Xing, C.: Using greedy random adaptive procedure to solve the user selection problem in mobile crowdsourcing. *Sensors* **19**(14), 3158 (2019)
- Yang, M., Biedermann, S., Tang, E.: On optimal designs for nonlinear models: a general and efficient algorithm. *J. Am. Stat. Assoc.* **108**(504), 1411–1420 (2013). <https://doi.org/10.1080/01621459.2013.806268>
- Yu, Y.: D-optimal designs via a cocktail algorithm. *Stat. Comput.* **21**(4), 475–481 (2011). <https://doi.org/10.1007/s11222-010-9183-2>
- Zeger, S.L., Liang, K.-Y., Albert, P.S.: Models for longitudinal data: a generalized estimating equation approach. *Biometrics* **44**(4), 1049–1060 (1988)
- Zou, Z.-Q., Li, Z.-T., Shen, S., Wang, R.-C.: Energy-efficient data recovery via greedy algorithm for wireless sensor networks. *Int. J. Distrib. Sens. Netw.* **12**(2), 7256396 (2016). <https://doi.org/10.1155/2016/7256396>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.