

RESEARCH ARTICLE

Nonmyopic and pseudo-nonmyopic approaches to optimal sequential design in the presence of covariates

Mia Sato Tackney ^a and Dave Woods ^b and Ilya Shpitser ^c

^a London School of Hygiene and Tropical Medicine, UK; ^b University of Southampton, UK; ^c Johns Hopkins University, USA

ARTICLE HISTORY

Compiled August 13, 2022

Abstract

In sequential experiments, subjects become available for the study over a period of time, and covariates are often measured at the time of arrival. We consider the setting where the sample size is fixed but covariate values are unknown until subjects enrol. Given a model for the outcome, a sequential optimal design approach can be used to allocate treatments to minimize the variance of the estimator of the treatment effect. We extend existing optimal design methodology so it can be used within a nonmyopic framework, where treatment allocation for the current subject depends not only on the treatments and covariates of the subjects already enrolled in the study, but also the impact of possible future treatment assignments within a specified horizon. The nonmyopic approach requires recursive formulae and suffers from the curse of dimensionality. We propose a pseudo-nonmyopic approach which has a similar aim to the nonmyopic approach, but does not involve recursion and instead relies on simulating trajectories of future possible decisions. Our simulation studies show that, for the simple case of a logistic regression with a single binary covariate and a binary treatment, and a more realistic case with four binary covariates, binary treatment and treatment-covariate interactions, the nonmyopic and pseudo-nonmyopic approaches provide no competitive advantage over the myopic approach, both in terms of the size of the estimated treatment effect and also the efficiency of the designs. Results are robust to the size of the horizon used in the nonmyopic approach, and the number of simulated trajectories used in the pseudo-nonmyopic approach.

KEYWORDS

design of experiments; optimal design; dynamic programming; sequential design; coordinate exchange

1. Introduction

How treatments should be allocated in sequential experiments in the presence of covariates is a highly debated topic, particularly within the clinical trials community [1,2]. We find and compare designs for experiments where subjects become available sequentially, covariates are measured at the time of arrival, and treatment is assigned soon after. We assume that a response is measured before the next subject arrives, and we assume a fixed sample size. At any point in the experiment, the covariate values for the subjects yet to enrol in the experiment are unknown. Such a set-up is often

characteristic of large Phase III trials, but is also common in experiments in the social sciences, such as political psychology lab experiments [3]. A specific example is an obstetrics trial to investigate three different techniques for pain relief during labour and their impact on normal vaginal delivery rates [4]. They recruited nulliparous women who requested epidural for pain relief during labour in two maternity units between August 1997 and April 2000. Age (categorized into five groups) and ethnicity (categorized into three groups) of the mothers were important covariates, and the size of trial ($n = 1054$) was determined at the start to achieve an anticipated power of 80% for estimating the change in rate of normal vaginal delivery. In such settings, covariates should be included in the analysis as their omission can result in bias [1].

From an optimal design point of view, the allocation of treatments should be done to maximize precision, or equivalently minimize variance, of the parameter estimators [5], which often results in balance, or equal replication, across groups defined by the distinct combinations of the covariate values. The optimal design approach aims to minimize the variance of the estimator of the treatment effect under the statistical model which is assumed to describe the relationship between the treatments, covariates and response. Atkinson [5] developed sequential optimal design methods for a linear model using D_A -optimality to make decisions for treatment allocation. In Section 2, we extend this approach to the generalized linear model case, which can be applied using any standard variance-based optimality criterion.

Minimization [6,7] is an alternative approach that directly targets equal replication across covariate groups, and is used extensively in clinical trials. It has received some criticism for being based on measures of imbalance of covariates which are not theoretically grounded [8] and it does not directly consider the precision of the parameter estimators in a statistical model.

The aforementioned design approaches are myopic in the sense that treatment allocation decisions are made using information about past subjects' covariates, treatments and, sometimes, response and the current subject's covariates. Treatment allocation for the current subject is made assuming that the experiment will terminate after its response is obtained, ignoring the fact that there are further subjects which will enter the trial, and that the estimators of interest will be based on data from the entire experiment. In contrast, nonmyopic approaches incorporate the potential impact of the current treatment decision on future possible decisions [9] in terms of efficiency of the final estimators. In this paper, we assess whether there is an efficiency benefit to taking into account the impact of future possible decisions. Dynamic programming is used to evaluate some expected loss, where the expectation is taken over unknown quantities from future subjects [10, p. 323]. Most applications of nonmyopic approaches in clinical trials aim to maximize a measure which combines efficiency and benefit of the treatment to the subject. An example is the Gittens index, which is a deterministic rule for allocating treatments to patients that aims to balance learning and efficiency with patient benefit [11,12,13,14]. Nonmyopic approaches in general and the specific case of a nonmyopic approach to logistic regression is described in Section 3.

The nonmyopic approach is computationally expensive which limits its use in practical settings. We propose a pseudo-nonmyopic approach in Section 4, which has a similar aim but does not require recursive formulae for the evaluation of the expected utility. We compare how it fares against the myopic and nonmyopic approaches in two simulation studies in Section 5. We discuss our findings and potential extensions of our work in Section 6.

2. Myopic Sequential Design for Generalized Linear Models

Suppose there are n subjects in total in an experiment, fixed in advance. For $i = 1, \dots, n$, we observe values,

$$\mathbf{z}_i = (z_{i,1}, \dots, z_{i,s})^T, \quad (1)$$

for the s covariates associated with unit i and we select a treatment t_i from a set of possible treatments \mathcal{T} . Following application of the treatment, we observe response y_i . For subjects $1, \dots, i$, define

$$\mathbf{Z}_i = \begin{pmatrix} \mathbf{z}_1^T \\ \mathbf{z}_2^T \\ \vdots \\ \mathbf{z}_i^T \end{pmatrix},$$

$$\mathbf{t}_i = (t_1, t_2, \dots, t_i)^T,$$

$$\mathbf{y}_i = (y_1, y_2, \dots, y_i)^T,$$

to be the $i \times s$ matrix of covariate values, the i -vector of treatments and i -vector of responses, respectively.

For $i = 1, \dots, n$, we assume a generalized linear model (GLM) with the y_i following independent exponential family distributions, with expectation μ_i related to a linear predictor $\eta_i = \mathbf{x}_i^T \boldsymbol{\beta}$ via a link function $g(\mu_i) = \eta_i$. The q -vector \mathbf{x}_i includes known functions of the treatment t_i and covariates \mathbf{z}_i , while the q -vector $\boldsymbol{\beta}$ includes the unknown model parameters. In later examples, y_i will be binary, with zero denoting the favourable response and \mathbf{x}_i will hold a constant, linear effects of the treatment and covariates and sometimes treatment-covariate products.

The information matrix for a GLM up to the allocation of the i th treatment can be written as $\mathbf{M}_i = \mathbf{X}_i^T \mathbf{W}_i \mathbf{X}_i$, with \mathbf{X}_i the $i \times q$ model matrix with j th row \mathbf{x}_j and $\mathbf{W}_i = \text{diag}\{\text{var}(y_j)\}$, $j = 1, \dots, i$; $i = 1, \dots, n$. The information matrix, which is the asymptotic inverse variance-covariance matrix for the maximum likelihood estimator $\hat{\boldsymbol{\beta}}$, can be used to define the class of D -optimality design selection criteria.

In a sequential setting, for the i th enrolled subject, a standard D -optimal design chooses treatment t_i to minimize the determinant of the inverse information matrix given covariate values \mathbf{z}_i and the current state $S_{i-1} = (\mathbf{Z}_{i-1}, \mathbf{t}_{i-1}, \mathbf{y}_{i-1})$, the treatments \mathbf{t}_{i-1} , covariates \mathbf{Z}_{i-1} and responses \mathbf{y}_{i-1} from subjects $1, \dots, i-1$. That is, a D -optimal design minimizes

$$\Psi_D(t_i | \mathbf{z}_i, S_{i-1}) = \left| (\mathbf{X}_i^T \mathbf{W}_i \mathbf{X}_i)^{-1} \right|, \quad (2)$$

where $|\cdot|$ denotes matrix determinant. A D -optimal design minimizes the volume of the confidence ellipsoid for $\boldsymbol{\beta}$ [15, p. 53]. The dependence of the design on responses \mathbf{y}_{i-1} is indirect, typically through estimates of $\boldsymbol{\beta}$ which are required to obtain the entries $\text{var}(y_j)$ of matrix \mathbf{W}_i .

Suppose interest is confined to a subset of the model parameters, e.g. corresponding to treatment effects, or in some linear combination of the parameters. In either case, we can express the target of inference as $\mathbf{A}^T \boldsymbol{\beta}$, where \mathbf{A} is a $q \times m$ matrix with $m < q$, and define a D_A -optimal design as

$$\Psi_{D_A}(t_i | \mathbf{z}_i, S_{i-1}) = \left| \mathbf{A}^T (\mathbf{X}_i^T \mathbf{W}_i \mathbf{X}_i)^{-1} \mathbf{A} \right|, \quad (3)$$

see Atkinson [15, p.137]. For precise estimation of a single treatment effect, \mathbf{A} is a q -vector with a single non-zero entry one corresponding to the treatment effect.

We adopt the common “biased coin” approach to optimal design in a clinical setting, extending the work of Atkinson [16] for binary treatments, and assign treatment $t \in \mathcal{T}$ to subject i with probability

$$P(t_i = t | \mathbf{z}_i, S_{i-1}) = \frac{\Psi(t | \mathbf{z}_i, S_{i-1})^{-1}}{\sum_{u \in \mathcal{T}} \Psi(u | \mathbf{z}_i, S_{i-1})^{-1}}, \quad (4)$$

where $\Psi(\cdot)$ denotes any variance-based optimality criterion. The probability of selecting treatment t is a decreasing function of $\Psi(\cdot)$, meaning designs with lower generalised variance will be preferred. The random element to selection avoids any suspicion of selection bias [5].

There are two practical issues that arise in the application of sequential design for GLMs. Firstly, for count or, especially, binary data, separation can occur where a linear combination of covariates perfectly predicts the response. Separation can result in the likelihood function becoming monotonic and maximum likelihood estimates of the regression coefficients tending to plus or minus infinity [17]. It can be particularly prevalent for small experiments, as in the early stages of a sequential study, and when treatment factors and covariates are binary. A common approach to overcome this issue is the introduction of a prior distribution for $\boldsymbol{\beta}$ to shrink parameter estimates towards zero and hence reduce the bias introduced by separation. Common choices include Jeffreys’ prior [18] and independent Cauchy priors [19]. Under either of these choices, the inverse information matrix still provides a measure of dispersion of the estimator. We adopt the latter approach.

Secondly, the objective functions (2) and (3) depend on values of the model parameters through matrix \mathbf{W}_i [20]. To overcome this issue, for the initial $n_0 < n$ subjects we assign treatments under the assumption that $\boldsymbol{\beta}$ is a vector of zeros. Responses from these first n_0 subjects are then used to obtain estimates of the model parameters which are used for the selection of t_{n_0+1} . For the selection of subsequent treatments, all responses available up to that point are used to obtain estimates $\hat{\boldsymbol{\beta}}_i = (\hat{\beta}_{0,i}, \dots, \hat{\beta}_{q-1,i})^T$ ($i = n_0 + 1, \dots, n$) which are used to evaluate probability (4). Algorithm 2 in Appendix D outlines the steps in constructing a sequential optimal design.

3. Nonmyopic Approach

A nonmyopic approach to treatment allocation assesses the impact of a proposed allocation for subject i using the impact on inference at the current stage and also across future stages accounting for yet unmade allocation decisions. The number of future subjects considered is called the horizon, with the number of subjects in the hori-

zon denoted by N . A nonmyopic approach balances two conflicting aims in treatment allocation:

- (1) Exploitation: choose the treatment that leads to the most precise estimates of β at the current state.
- (2) Exploration: choose treatments which may not be optimal at the current state, but leads to model-based exploration to increase information about the parameters.

Dynamic programming is an approach for solving multistage optimization problems (see, for example, Powell [21]). The overall problem is broken into different stages, which often correspond to points in time, and each stage of the problem can be optimized conditionally on past states. The key idea is that the overall sequence of decisions for treatment selection will be optimal for the entire experiment [10, p. 320]. The optimal design can be obtained by forward or backward induction. We focus on backward induction since it is the approach that is usually most appropriate in problems involving uncertainties [10, p. 328]. In backward induction, we start by finding the optimal decision at the end of the sequence of decisions, taking into account all possible treatments and covariates that may have been observed up until that point. Then, one can work backwards and obtain the optimal design taking expectations of unknown quantities [10, p.330]. See Huan and Marzouk [9] for a recent overview of approximate dynamic programming in the context of Bayesian experimental design. Dynamic programming has been used in some clinical trials applications where one wishes to balance the aim of estimating the parameters (exploration) with the aim of giving subjects the best possible treatment or obtaining maximum total benefit to the patient (exploitation). See, for example, [22], [23], [24] or [25].

We now describe the nonmyopic approach for the binary response. To keep notation simple, we assume that we have linear predictor of the form $1 + z_i + t_i$ where t_i and z_i are binary treatments and binary covariates, respectively. We begin by constructing an initial design \mathbf{X}_{n_0} with n_0 subjects using the exchange algorithm. We assume $\beta = \mathbf{0}$ as an initial guess for evaluating the objective function in the construction of \mathbf{X}_{n_0} . We then obtain responses for the first n_0 subjects, \mathbf{y}_{n_0} , and fit the model to obtain the initial maximum likelihood estimates of the model parameters, $\hat{\beta}_0$.

Now, supposing that we have a design for $i - 1$ subjects, and we have obtained parameter estimates $\hat{\beta}_{i-1}$ as a result of that design. We observe covariate value z_i for the i th subject and wish to evaluate the impact of assigning treatment t_i on decisions on future possible subjects. For example, for horizon $N = 1$, we consider the expected value of the objective function after $i + 1$ subjects. Suppose treatment t_i is assigned to subject i . Since the objective function depends on y_i , we need to consider the two possible responses that y_i may take, and then consider the possible values that z_{i+1} can take. For a given covariate value z_{i+1} for subject $i + 1$, we denote by $t_{i+1}^*(z_{i+1}, t_i, y_i | z_i, S_{i-1})$ the optimal choice of treatment for subject $i + 1$ given z_{i+1} and t_i :

$$t_{i+1}^*(z_{i+1}, t_i, y_i | z_i, S_{i-1}) = \arg \min_{t_{i+1}} \Psi(t_{i+1} | z_{i+1}, S_i), \quad (5)$$

where t_{i+1}^* denotes a hypothetical future treatment decision, as opposed to t_i , which

denotes an actual treatment decision in the design. From here on, we suppress the conditioning and write $t_{i+1}^*(z_{i+1}, t_i, y_i)$ for simplicity. Now, we take the expectation of the objective function over two possible responses which may be obtained to find an expected value of the objective function $\Psi(\cdot)$ over the unknown response:

$$\begin{aligned}\mathbb{E}_{y_i} \Psi(t_{i+1} | z_{i+1}, S_i) &= \mathbb{P}(y_i = 0 | z_i, t_i, S_{i-1}) \Psi(t_{i+1} | z_{i+1}, S_i) \\ &\quad + \mathbb{P}(y_i = 1 | z_i, t_i, S_{i-1}) \Psi(t_{i+1} | z_{i+1}, S_i),\end{aligned}$$

where $y_i \sim \text{Bernoulli}(\pi_i)$ with π_i given by:

$$\pi_i = \frac{\exp(\mathbf{x}_i \hat{\boldsymbol{\beta}}_{i-1})}{1 + \exp(\mathbf{x}_i \hat{\boldsymbol{\beta}}_{i-1})}, \quad (6)$$

where $\mathbf{x}_i = (1, z_i, t_i)$ is the i th row of the design matrix. Now, we consider the possible covariate values that we may observe for the next subject. We denote by $\mathbb{P}(z_i = z)$ the probability that the i th subject has covariate value z . In some cases, the distribution of the covariates may be known; if not, the distribution can be estimated by the empirical distribution of the covariates of the first i subjects. We denote by $\Psi_i(t_i | z_i, S_{i-1})$ the expected value of the objective function when treatment t_i is assigned to subject i , taking into account the impact of the decision on one further decision in the future. We obtain an expectation over the possible covariate combinations of the optimality criterion:

$$\begin{aligned}\Psi_i(t_i | z_i, S_{i-1}) &= \mathbb{E}_{\mathbf{z}_{i+1}} \mathbb{E}_{y_i} \Psi(t_{i+1}^*(z_{i+1}, z_i, y_i, t_i, S_{i-1}) | z_{i+1}, z_i, y_i, t_i, S_{i-1}) \\ &= \sum_z \mathbb{P}(z_{i+1} = z) \sum_y \mathbb{P}(y_i = y) \Psi(t_{i+1}^*(z_{i+1}, S_i) | z_{i+1}, S_i),\end{aligned} \quad (7)$$

where $t_{i+1}^*(z_{i+1}, S_i) = \arg \min_t \Psi(t | z_{i+1}, S_i)$ and $S_i = (\mathbf{Z}_i, \mathbf{t}_i, \mathbf{y}_i)$.

For a horizon N greater than 1, we can use the following recursive relationship to find the optimal treatment for subject i . The expected value of the objective function after $i + N$ subjects, when treatment t_i has been assigned, is given as follows, for $k \in \{i, i + 1, \dots, i + N - 1\}$:

$$\begin{aligned}\Psi_k(t_k | z_k, S_{k-1}) &= \mathbb{E}_{\mathbf{z}_{k+1}} \mathbb{E}_{y_k} \Psi_{k+1}(t_{k+1}^*(z_{k+1}, S_{k-1}) | z_{k+1}, S_{k-1}) \\ &= \sum_z \mathbb{P}(z_{k+1} = z) \sum_y \mathbb{P}(y_k = y) \Psi_{k+1}(t_{k+1}^*(z_{k+1}, S_k) | z_{k+1}, S_k)\end{aligned} \quad (8)$$

$$\Psi_{i+N}(t_{i+N} | z_{i+N}, S_{i+N-1}) = \Psi(t_{i+N} | z_{i+N}, S_{i+N-1}), \quad (9)$$

where $t_{i+1}^*(z_{i+1}, S_i) = \arg \min_t \Psi_{i+1}(t | z_{i+1}, S_i)$ and $S_i = (\mathbf{Z}_i, \mathbf{t}_i, \mathbf{y}_i)$.

The recursion the non-myopic approach makes it computationally expensive, in particular with increasing N . For illustrative purposes, we compare the CPU time for the myopic and non-myopic approach (with horizon ranging from 1 to 5) for constructing a D_A -optimal design with 25 patients in Table 1. The design has one binary covariate and a binary response, and the initial design is 10 patients. We observe that the non-myopic approach is more time consuming than the myopic approach, and the CPU time increases dramatically with increasing N .

4. Pseudo-nonmyopic approach

We now explore a *pseudo-nonmyopic* approach which involves evaluating an objective function with a similar aim to that of the nonmyopic approach, without the use of recursion which leads to the curse of dimensionality. In the pseudo-nonmyopic approach, in order to make a decision about the treatment of the i th patient, we generate M possible *trajectories* of covariate values for patient $i + 1$ until patient n . For each of the M trajectories, we construct a *pseudo-design* in which we have the i patients and $(n - i - 1)$ patients in the trajectory, and treatments allocated using an approach that we describe below. We look at the average value of the objective function of the M pseudo-designs where we assign $t_i = 1$, and compare it to the average value of objective function the M pseudo-designs when $t_i = -1$; we select t_i according to a probability that is weighted by these averages. The computational burden is reduced as nested expectations and minimizations are not necessary but we are still able to incorporate information about future possible decisions. We describe this novel approach for the logistic model case and provide a simulation to show how it compares to the myopic approach.

This approach takes averages over simulated values of the covariates for subjects $i + 1$ up to n . Optimization based on Monte Carlo simulations of unknown quantities is typically conducted in a Bayesian setting for design of experiments [26], where values of the unknown parameters may be simulated from a prior distribution. See Gentle [27] for an overview of Monte Carlo methods and Ryan [28] for an application to Bayesian design of experiments.

In order to create a design using the pseudo-nonmyopic approach for the logistic model, just like in the sequential myopic and nonmyopic algorithms, we begin by constructing an initial design \mathbf{X}_{n_0} . This involves an exchange algorithm where we assume $\boldsymbol{\beta} = \mathbf{0}$ as an initial guess. We then generate responses \mathbf{y}_{n_0} , and fit the model to obtain the initial estimates $\hat{\boldsymbol{\beta}}_{n_0}$.

Then, to select a treatment for subject i , for $i \in \{n_0 + 1, \dots, n\}$, we observe \mathbf{z}_i . Based on the assumed covariate distribution $f_{\mathbf{z}}$, we generate M possible trajectories for the covariates, $\mathbf{Z}_{(i+1):n}^1, \mathbf{Z}_{(i+1):n}^2, \dots, \mathbf{Z}_{(i+1):n}^m$, where

$$\mathbf{Z}_{(i+1):n}^m = \begin{pmatrix} \mathbf{z}_{i+1}^{mT} \\ \mathbf{z}_{i+2}^{mT} \\ \vdots \\ \mathbf{z}_n^{mT} \end{pmatrix}, \quad (10)$$

for $m \in \{1, 2, \dots, M\}$. We assume, as for the nonmyopic approach, that we have a distribution $f_{\mathbf{z}}$ for the covariate \mathbf{z} . This may be the true distribution in the population (if it is known), or an empirical approximation based on the subjects in the trial up until the i th subject. The covariate distribution may depend on time, in which case we refer to it as a dynamic covariate. We then allocate treatments sequentially along each trajectory.

Given the first subject in the trajectory, \mathbf{z}_{i+1}^m , we choose the treatment t_{i+1}^{*m} which minimizes the objective function Ψ given t_i , and the treatments and covariates of previous subjects and estimates $\hat{\boldsymbol{\beta}}_{i-1}$ based on the responses of the previous subjects, \mathbf{y}_{i-1} :

$$t_{i+1}^{*m}(\mathbf{z}_{i+1}^m, t_i \mid z_i, S_{i-1}) = \arg \min_{t_{i+1}} \Psi(t_{i+1} \mid z_i, \mathbf{z}_{i+1}^m, t_i, S_{i-1}). \quad (11)$$

To allocate a treatment for the next subject in the trajectory with covariate values \mathbf{z}_{i+2}^m , we then assume that t_{i+1}^{*m} has been allocated to subject \mathbf{z}_{i+1}^m and choose the treatment t_{i+2}^{*m} which minimizes the objective function. We make the assumption that the future decisions are independent of the future responses. This means that we assume the same estimate for $\boldsymbol{\beta}$ as in Equation (11) and do not update it, as it would seem circular to use responses generated from a particular estimate of the parameters to re-estimate the same parameters. We continue in this way until all subjects in the trajectory have been allocated a treatment:

For each j in $\{i+2, i+4, \dots, n\}$, we define:

$$\begin{aligned} & t_j^{*m}(\mathbf{z}_j^m, t_{j-1}^* \mid z_i, \mathbf{Z}_{(i+1):(j-1)}^m, t_i, \mathbf{t}_{(i+1):(j-2)}^*, S_{i-1}) \\ &= \arg \min_{t_j} \Psi(t_j \mid z_i, \mathbf{Z}_{(i+1):(j)}^m, t_i, \mathbf{t}_{(i+1):(j-1)}^*, S_{i-1}). \end{aligned} \quad (12)$$

For the m th trajectory, we obtain a pseudo-design with n subjects where the i th treatment is 1, as well as a pseudo-design where the i th subject receives treatment -1 . We denote the objective function of the two designs as follows:

$$\Psi(t_n \mid z_i, \mathbf{Z}_{(i+1):n}^m, t_i = 1, \mathbf{t}_{(i+1):(n-1)}^{*m}, S_{i-1}), \quad (13)$$

$$\Psi(t_n \mid z_i, \mathbf{Z}_{(i+1):n}^m, t_i = -1, \mathbf{t}_{(i+1):(n-1)}^{*m}, S_{i-1}). \quad (14)$$

We define the average objective function for $i = n_0 + 1, \dots, n - 1$ across the M designs, for $t_i = t$:

$$\bar{\Psi}(t_i = t) = \frac{1}{M} \sum_{m=1}^M \Psi(t_n \mid z_i, \mathbf{Z}_{(i+1):n}^m, t_i = t, \mathbf{t}_{(i+1):(n-1)}^{*m}, S_{i-1}), \quad (15)$$

For $i = n$, we do not generate any future covariates so we have:

Algorithm		CPU time (seconds)
Myopic		0.124
Nonmyopic	$N = 1$	0.552
	$N = 2$	3.862
	$N = 3$	31.454
	$N = 4$	272.821
	$N = 5$	2165.024
Pseudo-nonmyopic	$M = 10$	2.307
	$M = 50$	10.787

Table 1.: CPU time in seconds for myopic, nonmyopic (with horizon N ranging from 1 to 5) and pseudo-nonmyopic (with the number of trajectories M equal to 10 or 50) approaches to constructing a D_A -optimal design with 25 patients, where there is one binary covariate and binary response. The initial design is 10 patients. Simulations were performed on a Linux machine with a 64-bit processor and 16 GB of memory.

$$\bar{\Psi}(t_i = t) = \Psi(t_n = t \mid z_n, S_{n-1}). \quad (16)$$

for $t \in \{-1, 1\}$.

We sample t_i from the set $\{-1, 1\}$ where the probability of selecting 1 is given by

$$\frac{\bar{\Psi}(t_i = 1)^{-1}}{\bar{\Psi}(t_i = 1)^{-1} + \bar{\Psi}(t_i = -1)^{-1}}. \quad (17)$$

We then observe the response y_i and refit the model to obtain $\hat{\beta}_i$.

We compare the CPU time for the myopic, nonmyopic and pseudo-nonmyopic approaches to constructing a D_A -optimal design in Table 1. The design has 25 patients, one binary covariate and a binary response, and the initial design is 10 patients. The pseudo-nonmyopic approaches with 10 or 50 trajectories take into account decisions up until the end of the experiment, but since they require no recursion, run faster than the nonmyopic approaches with horizon 2 or greater.

5. Simulations

We conduct two simulation studies to compare the myopic, nonmyopic and pseudo-nonmyopic approaches to constructing sequential D_A -optimal designs. Specifically, we aim to compare estimates of the model parameters across the three methods, as well as the efficiencies of the nonmyopic and pseudo-nonmyopic approaches relative to the myopic approach. We define the D_A -efficiencies of a design \mathbf{X}_i relative to another design \mathbf{X}_i^* with parameter values $\hat{\beta}_i$ in the logistic model case as

$$\text{Eff}_{D_A} = \left\{ \frac{\Psi_{D_A}(\mathbf{X}_i^*, \hat{\boldsymbol{\beta}}_i)}{\Psi_{D_A}(\mathbf{X}_i, \hat{\boldsymbol{\beta}}_i)} \right\}^{1/m}, \quad (18)$$

where m is the number of parameters of interest.

The first simulation explores a simple setting where there is a single covariate and no treatment-covariate interactions. We compare a number of settings for the sequential design approaches, including the size of the initial design, the horizon for the nonmyopic approach and the number of trajectories for the pseudo-nonmyopic approach. The second simulation explores a more realistic setting with four covariates, where treatment-covariate interactions are of interest. A few select settings for the nonmyopic and pseudo-nonmyopic approaches are chosen and the applicability of these approaches to this more complex setting is demonstrated.

5.1. Single covariate setting with no interaction

In the first simulation, we have a simple set-up with 250 units, where a single covariate is observed. There are two parameters of interest: the effect of treatment and the effect of the covariate. The covariate can take values in $\{-1, 1\}$ and is generated such that $\mathbb{P}(z_i = 1) = 0.5$ and $\mathbb{P}(z_i = -1) = 0.5$ for all i . We assume the true model for the response is $y_i \sim \text{Bernoulli}(\pi_i)$ with $\text{logit}(\pi_i) = z_i + t_i$, and generate responses according to this model. We fit the models using the R function `bayesglm` in the `arm` package [29], with Cauchy prior distribution with center zero and scale given by 2.5 for both the treatment and covariate parameters. We generate responses ensuring that the data generating mechanism is the same across simulations comparing the myopic, nonmyopic and pseudo-nonmyopic designs as described in Appendix A.

Since the responses are generated with a logistic model, the information matrix and the objective function depend on values of the model parameters, so estimates of parameters are needed in order to design the experiment aimed to estimate these parameters in the first place. An initial design is constructed with the exchange algorithm with D_A optimality as the objective function, under the assumption that $\boldsymbol{\beta}$ is a vector of zeros. We consider four possible sizes for the initial design: 10 units, 20 units, 50 units and 100 units.

We consider the following approaches for constructing a sequential D_A -optimal design:

- A myopic D_A -optimal design.
- A nonmyopic D_A -optimal design with horizon $N = 1$, with the correct covariate distribution assumed.
- A nonmyopic D_A -optimal design with horizon $N = 2$, with the correct covariate distribution assumed.
- A nonmyopic D_A -optimal design with horizon $N = 3$, with the correct covariate distribution assumed.
- A nonmyopic D_A -optimal design with horizon $N = 1$, with the empirical covari-

- ate distribution assumed.
- A nonmyopic D_A -optimal design with horizon $N = 2$, with the empirical covariate distribution assumed.
 - A nonmyopic D_A -optimal design with horizon $N = 3$, with the empirical covariate distribution assumed.
 - A pseudo-nonmyopic D_A -optimal design with $M = 10$.
 - A pseudo-nonmyopic D_A -optimal design with $M = 50$.

Designs are evaluated using the efficiency relative to the myopic design, given by Equation (18), at each sample size between the initial sample size and 250, inclusive. The true values of the parameters are used to calculate Ψ_{D_A} . The simulation is repeated 100 times.

5.1.1. Results

Figure 1 displays the estimates of β_i for all methods when initial sample size is 10. Across all methods, as sample size increases, the estimates converge towards the true value and the variability of the estimates reduces. We observe that the MC error bars include the true value of the estimates at a smaller sample size for the myopic approach compared to the other approaches. This is particularly evident for the treatment effect.

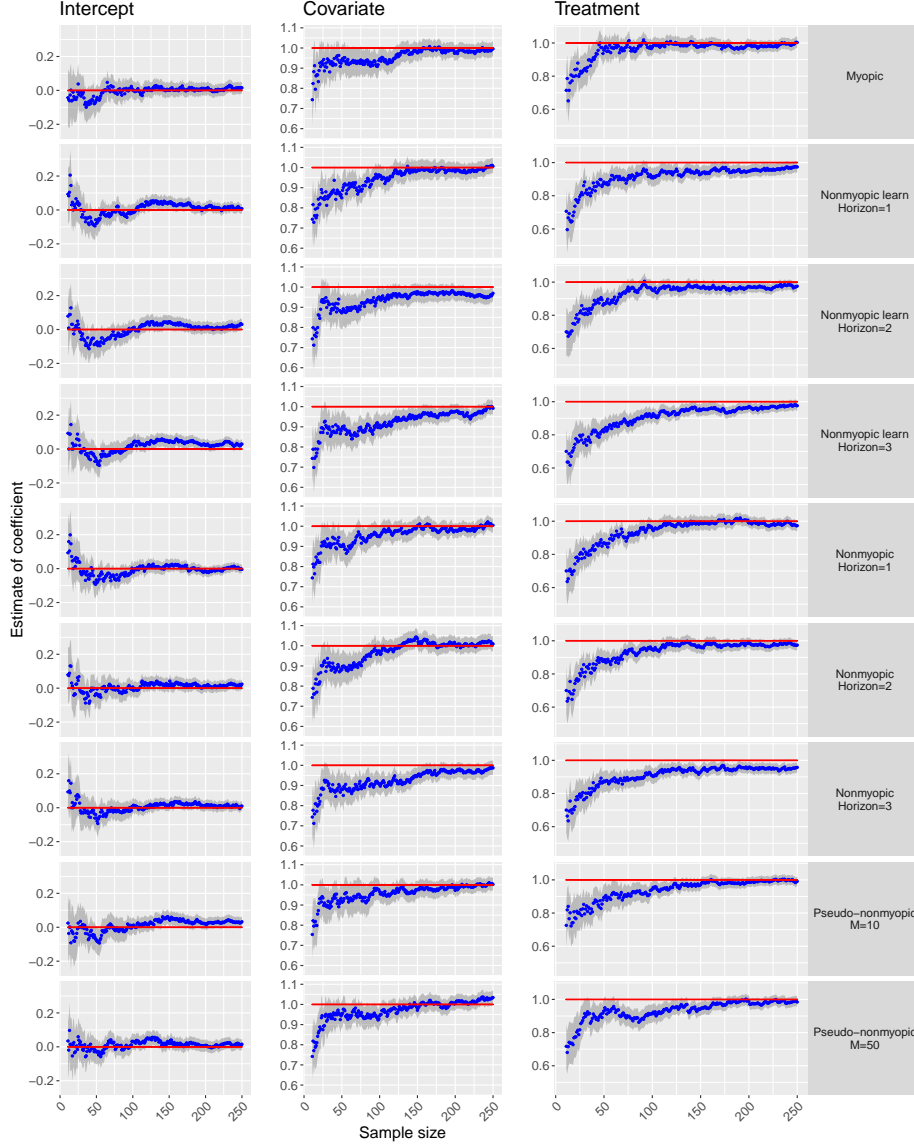


Figure 1.: Estimates of $\hat{\beta}_i$ for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 10. Results for the myopic approach, the nonmyopic approach and pseudo-nonmyopic approach to constructing D_A -optimal designs are shown. For the nonmyopic approach, we consider the horizon equal to $N = 1, 2$ and 3 , and we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used (nonmyopic learn). For the pseudo-nonmyopic approach, we consider $M = 10$ and 50 . Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The true values of β are indicated in red.

In Figure 2, we plot the efficiencies of the nonmyopic and pseudo-nonmyopic approaches relative to the myopic approach with initial sample size of 10. We observe that initially, the nonmyopic approach and pseudo-nonmyopic approaches are less efficient than the myopic approach. When sample size is roughly 150, there is no noticeable difference between the nonmyopic approaches and the myopic approach. Further, there

efficiencies for the nonmyopic approaches appear to be similar regardless of the choice of N and whether the true or empirical distribution of the covariates is used. For the pseudo-nonmyopic approach, however, the performance is significantly lower than the myopic approach until sample size is over 230. Increasing the number of trajectories from 10 to 50 has very little effect on the efficiency.

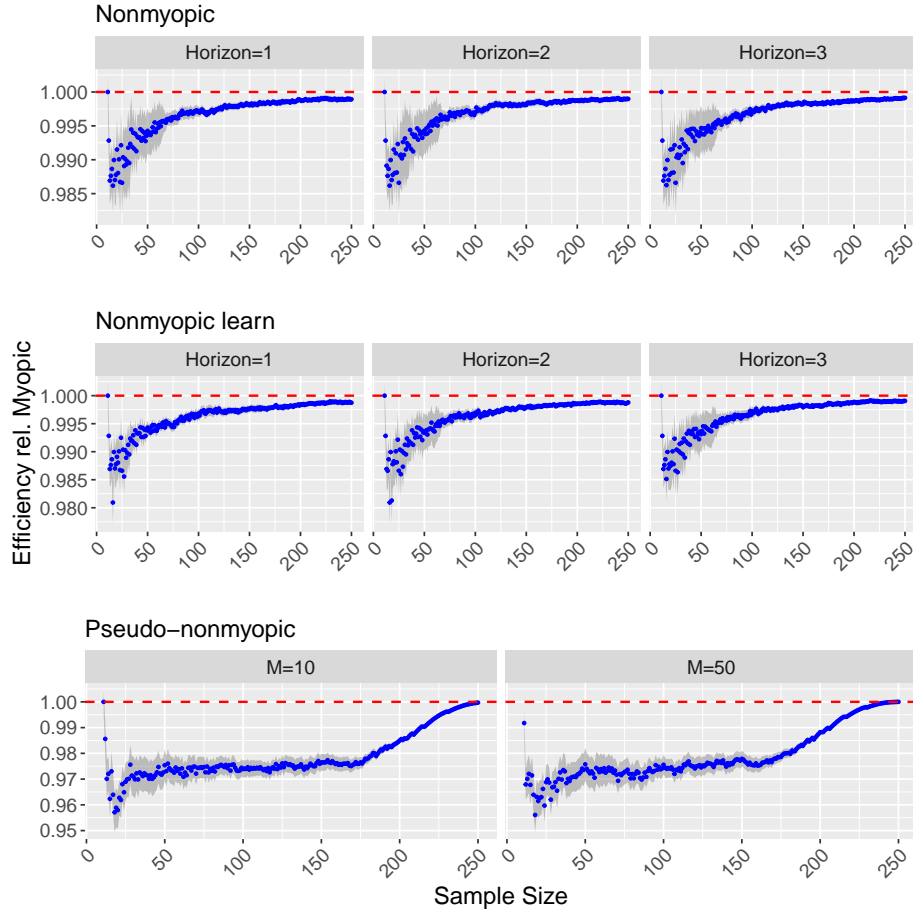


Figure 2.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 10. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.

We observe in this simulation that there appears to be no benefit to the nonmyopic and pseudo-nonmyopic approaches in this setting where we have one binary treatment and one binary covariate, and the covariate is generated such that $\mathbb{P}(z_i) = 0.5$ for all i . Results for initial sample sizes of 20, 50 and 100 are shown in Appendix B. With a larger initial sample size, the initial estimates of β are closer to their true values and the initial efficiencies are closer to 1.

5.2. Multiple covariates and treatment-covariate interactions

In the second simulation, we explore a more complex setting which could realistically be

observed in a clinical trial scenario, to demonstrate the practicality of this method. In this simulation, there are 250 patients and four binary biomarkers are measured on each patient. Such a setting could realistically be observed in a clinical trial; for example, the FOCUS 4 trial recruited patients registered with newly-diagnosed, advanced or metastatic disease from colorectal cancer, from four different molecular cohorts [30]. There are nine parameters in the model: the effect of treatment, the effect of the four biomarkers and the four treatment-biomarker interactions. We assume that the effect of treatment and the four treatment-biomarker interactions are of interest. The covariates take values in $\{1, -1\}$ and generated such that:

$$\begin{aligned}\mathbb{P}(z_{i,1} = 1) &= 0.5 \\ \mathbb{P}(z_{i,2} = 1) &= 0.4 \\ \mathbb{P}(z_{i,3} = 1) &= 0.7 \\ \mathbb{P}(z_{i,4} = 1) &= 0.2.\end{aligned}$$

We assume the true model for the response is $y_i \sim \text{Bernoulli}(\pi_i)$ with $\text{logit}(\pi_i) = 3z_{i,2} + 4z_{i,4} + 2z_{i,4}t_i$ and generate responses according to this model. We fit the model with `bayesglm` and try to control sources of variability as described in Appendix A. The initial design is constructed with an exchange algorithm to allocate treatments to 50 units, under the assumption that β is a vector of zeros.

We consider the following different approaches for constructing a sequential D_A optimal design:

- A myopic D_A -optimal design.
- A nonmyopic D_A -optimal design with horizon $N = 3$, with the correct covariate distribution assumed.
- A pseudo-nonmyopic D_A -optimal design with $M = 10$.

5.2.1. Results

In Figure 3, we see the estimates of β for the myopic approach, the nonmyopic approach with $N = 3$ and the pseudo-nonmyopic approach with $M = 10$, where the initial sample size is 50. We observe that, across all three methods, the second and fourth biomarkers, which have a strong effect on the response, are poorly estimated; the effect is attenuated towards zero. The myopic approach is able to estimate the interactions better than the nonmyopic and pseudo-nonmyopic approaches; the interaction between the treatment and fourth biomarker is estimated well by the myopic approach, but it is attenuated towards zero for the nonmyopic and pseudo-nonmyopic approaches. Further, the interaction between the treatment and second biomarker is not estimated well for all approaches, although the myopic approach achieves a value closer to the true value than the nonmyopic or pseudo-nonmyopic approaches.

In Figure 4, we plot the D_A -efficiencies of the nonmyopic and pseudo-nonmyopic approaches relative to the myopic approach. The nonmyopic approach and pseudo-nonmyopic approaches are generally less efficient than the myopic approach, and we observe that the pseudo-nonmyopic approach achieves a higher efficiency than the nonmyopic approach at the end of the experiment.

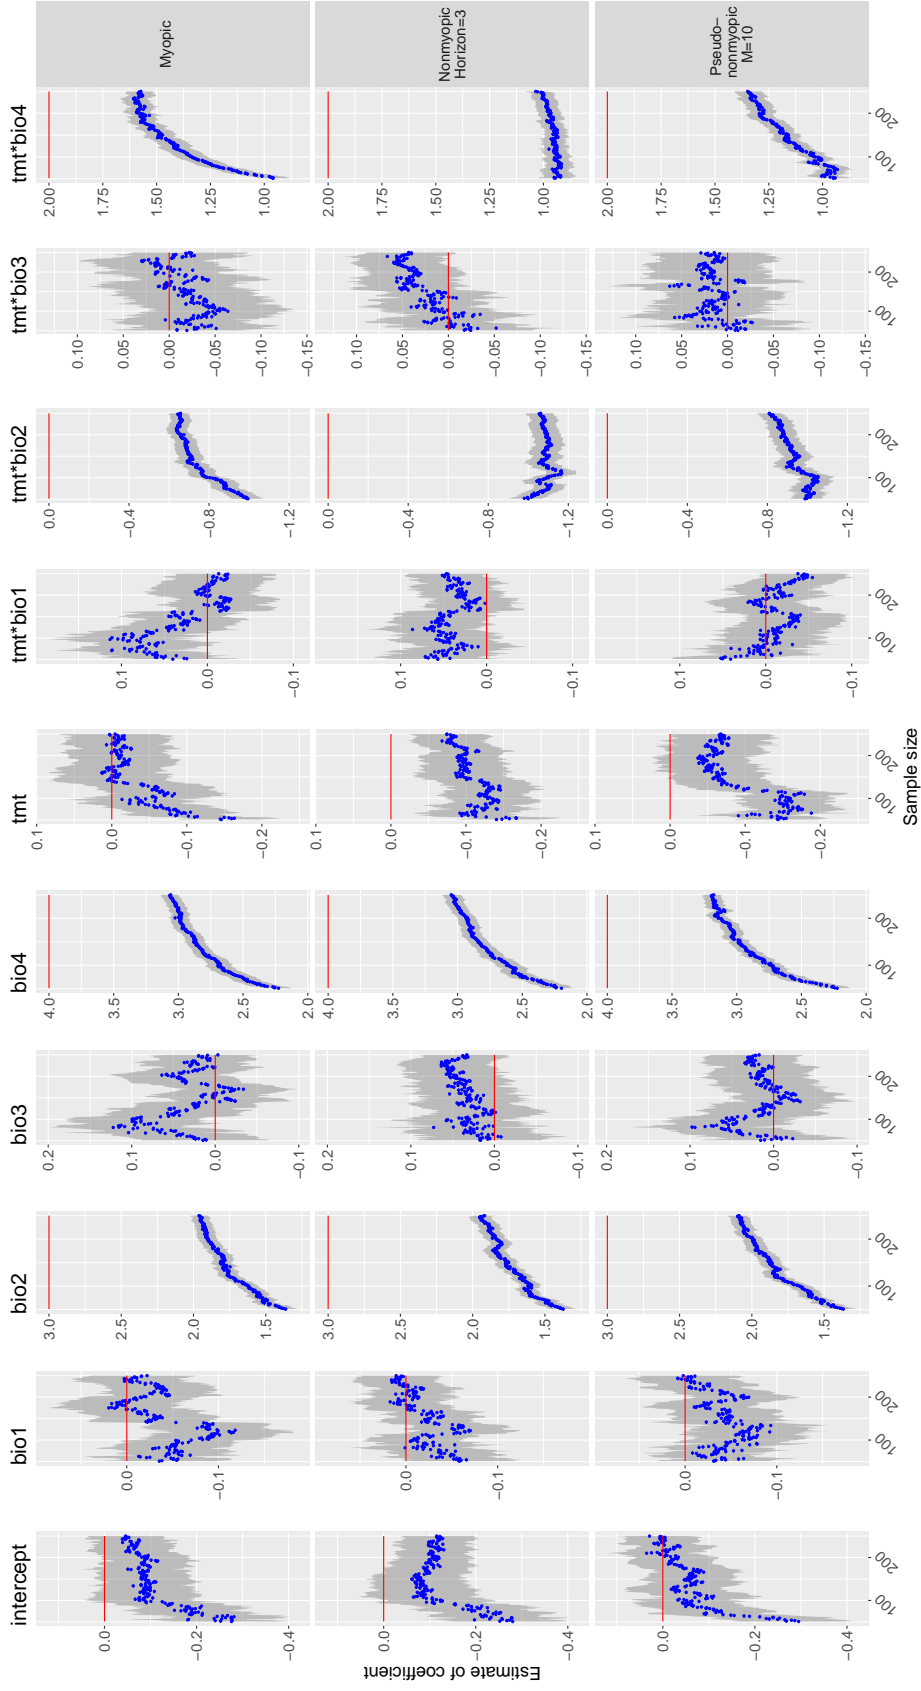


Figure 3.: Estimates of $\hat{\beta}_i$ for the simulation with a four covariates and treatment-biomarker interactions and initial sample size of 50. Results for the myopic approach, the nonmyopic approach with $N = 3$ and pseudo-nonmyopic approach with $M = 10$ to constructing D_A -optimal designs are shown. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times$ MC error. The true values of β are indicated in red.

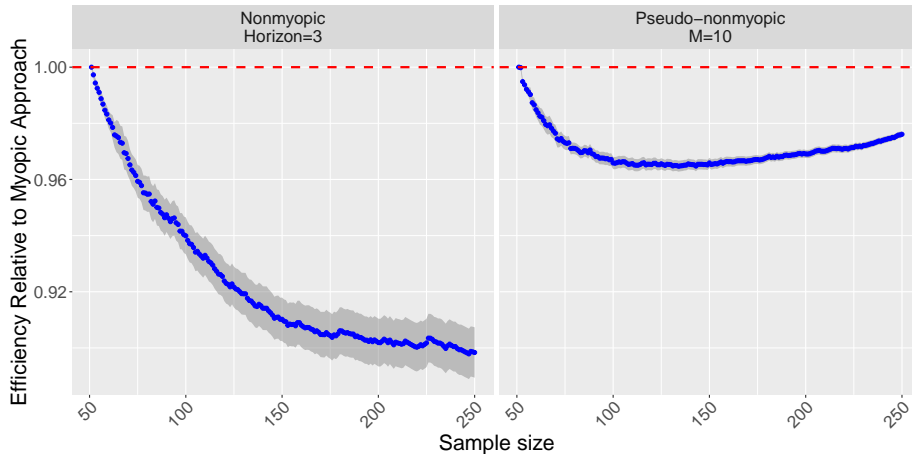


Figure 4.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a four covariates, treatment-covariate interactions and initial sample size of 50. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.

In this more complex example, we observe that the myopic approach is better able to estimate the interaction effects, and is more efficient than the nonmyopic and pseudo-nonmyopic approaches. We found that, under this particular setting, the pseudo-nonmyopic approach is more efficient than the nonmyopic approach. Simulations with other covariate settings and choices in optimality criteria were explored by Tackney [31], which showed consistently that the myopic approach is more efficient than the nonmyopic and pseudo-nonmyopic approaches.

6. Discussion

This paper extended the sequential optimal design approach first proposed by Atkinson [5] for the logistic model case and for any optimality criterion. We then placed this approach in a nonmyopic framework. We then developed a novel methodology called the pseudo-nonmyopic approach which is still able to take into account future possible subjects, but is less computationally expensive than the nonmyopic approach. Simulations showed that the nonmyopic approach and pseudo-nonmyopic approaches do not offer competitive advantage to the myopic approach for the logistic model case with a binary treatment where there is a single covariate, and also when there are four binary covariates with treatment-covariate interactions. The presented method could easily be applied to settings with a different number of covariates; for example, the obstetrics trial by the COMET Study Group UK used Age (categorized into five categories) and Ethnic Group (categorized into three groups) as covariates [4].

There are a number of possible extensions to this work which would improve its ability to be directly applicable to clinical trials and other experiments involving human subjects. Firstly, we assume responses are measured soon after treatments are given to subjects. In the particular example of the obstetrics trial, where the treatment is pain relief during labour and the outcome is mode of delivery, this assumption may be realistic. However, in many medical settings, there is a long

period of time between the onset of treatment and measurement of response. Some method to allow for a delay between treatment allocation and response could be useful. One modification would be to allow for the method to be batch sequential; instead of allocating treatments to one subject at a time, a group of subjects may be given optimal treatments by using the exchange algorithm. It is also possible to incorporate delay in adaptive designs. Hardwick et al. achieve this by assuming that subjects arrive according to a Poisson process [32].

Secondly, we do not consider toxicity in our work. We assume that the treatment which leads to a better response is the more desirable treatment, but it is possible that such a treatment has unsafe toxicity levels [33]. In our algorithms for treatment assignment, if the optimality criterion is equal for treatment $t_i = 1$ and $t_i = -1$, we would assign the treatment at random. In clinical trials, this is less likely to happen as relative efficiency of the treatments need to be considered in conjunction with relative toxicity [34]. In general, Rosenberger [33] recommended that adaptive designs should be considered after previous experiments have been able to establish low toxicity of the treatments.

Thirdly, while we have assumed a total of 250 subjects in our simulations, clinical trials typically have stopping rules which determine when the trial should terminate [35]. See [36] for a frequentist perspective and [37] and [38] for a Bayesian perspective on stopping rules in interim analysis. Including this element into our designs would mean that our methodology is more generally applicable to clinical trials. Further, we may be able to make statements about relative numbers of subjects and costs required in order to detect a significant difference in treatment effect for each method.

We have presented the nonmyopic and pseudo-nonmyopic algorithms in the setting where the response and treatments are binary. Natural extensions include allowing for more complex treatment structures, such as factorial designs, or allowing for a continuous response. Computing the expected objective function for a continuous response would require Monte-Carlo simulations. Extending our framework for the nonmyopic approach to allow for a more general response will require greater computational efficiency in our algorithms. This is also true of the pseudo-nonmyopic approach.

Appendix A. Generation of simulated data

In the simulations in Section 5.1 and 5.2, the generation of responses can be a source of unnecessary variability. When comparing the myopic and nonmyopic designs, or the myopic and pseudo-nonmyopic designs, we generate the responses in the following way:

- (1) Generate a deviate u_i from the $\text{Unif}(0, 1)$ distribution.
- (2) Set

$$y_i = \begin{cases} 1 & \text{if } u_i \geq \pi_i \\ 0 & \text{if } u_i < \pi_i \end{cases}. \quad (\text{A1})$$

The deviates u_i kept the same for the approaches being compared to try to

minimize sources of random variability in the simulation.

Appendix B. Further Results




Figure B1.: Estimates of $\hat{\beta}_i$ for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 20. Results for the myopic approach, the nonmyopic approach and pseudo-nonmyopic approach to constructing D_A -optimal designs are shown. For the nonmyopic approach, we consider the horizon equal to $N = 1, 2$ and 3 , and we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used (nonmyopic learn). For the pseudo-nonmyopic approach, we consider $M = 10$ and 50 . Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The true values of β are indicated in red.

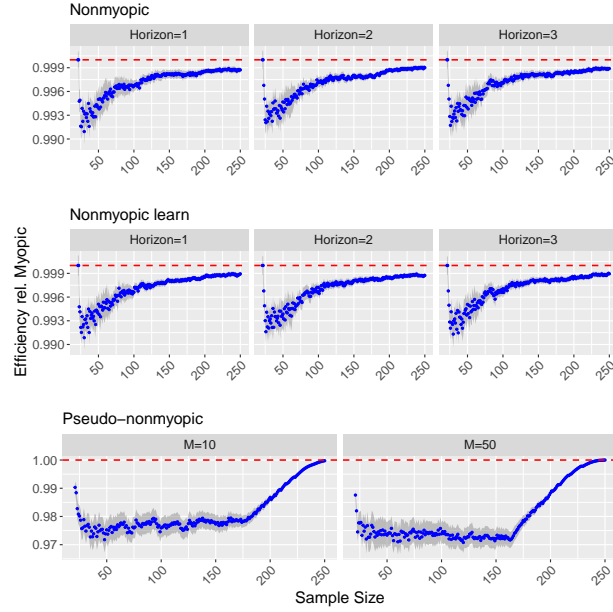


Figure B2.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 20. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.

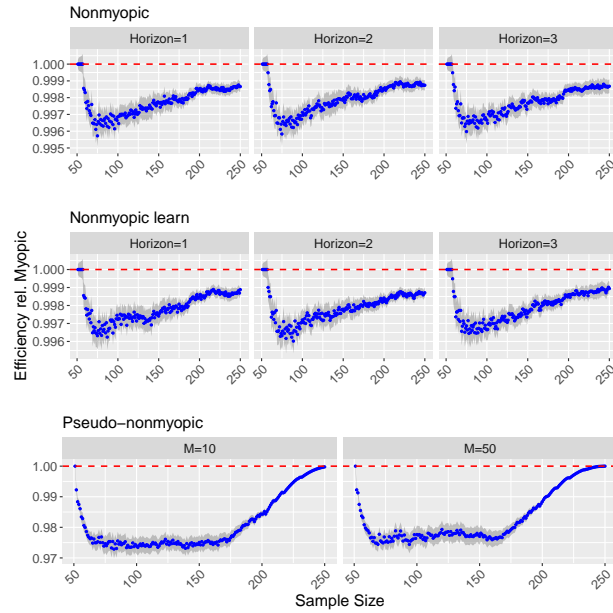


Figure B3.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 50. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.




Figure B4.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 100. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.

Appendix C. R code

We refer readers to Chapter 8 of [31] for a vignette of the R code and examples of how to implement the the myopic, nonmyopic and pseudo-nonmyopic approaches to optimal design. The function for the nonmyopic approach to optimal design for logistic regression is described in Section 8.2.2, and the function for the pseudo-nonmyopic approach is described in Section 8.3.

Appendix D. Myopic design algorithm

Algorithm 1 Myopic sequential optimal design: returns a design matrix for a sequentially conducted experiment. Inputs are the covariate values, \mathbf{Z}_n , and the number of subjects in initial design, n_0 .

```
1: function SEQOPTL( $\mathbf{z}_n, n_0$ )  
  
2:   Initialization  
3:   Construct initial design  $\mathbf{X}_{n_0}$  using an exchange algorithm for the first  $n_0$  subjects assuming  $\boldsymbol{\beta} = \mathbf{0}$ .  
4:   Observe responses  $\mathbf{y}_{n_0} = (y_1, y_2, \dots, y_{n_0})^T$ .  
5:   Fit the required generalized linear model to obtain a penalised MLE  $\hat{\boldsymbol{\beta}}_{n_0}$ .  
  
6:   for  $i$  in  $n_0 + 1$  to  $n$  do  
7:     Observe  $\mathbf{z}_i = (z_{i,1}, \dots, z_{i,s})^T$   
8:     Calculate  $\Psi(t \mid \mathbf{t}_{i-1}, \mathbf{Z}_i, \mathbf{y}_{i-1})$  for each treatment  $t \in \mathcal{T}$ .  
9:     Sample treatment allocated to subject  $i$  with probability (4).  
10:    Observe response  $y_i$ .  
11:    Refit model with responses  $\mathbf{y}_i$  and updated design matrix  $\mathbf{X}_i$ .  
12:    Obtain updated parameter estimates  $\hat{\boldsymbol{\beta}}_i$ .  
13:  end for  
  
14:  return  $\mathbf{X}_n$   
  
15: end function
```

Disclosure statement

No potential competing interest was reported by the authors

Funding

M. S. Tackney was supported by the Economic and Social Research Council Doctoral Training Centre (grant number ES/J500161/1). D. C. Woods was supported by the UK Engineering and Physical Sciences Research Council (grant number EP/J018317/1). I. Shpitser was supported by Defense Advanced Research Projects Agency (grant number HR0011-18-C-0049), National Institutes of Health (R01 AI127271-01A1) and Office of Naval Research (grant number N00014-18-1-2760).

Appendix E. References

References

- [1] Senn S. Seven myths of randomisation in clinical trials. *Statistics in Medicine*. 2013; 32:1439–1450.
- [2] Rosenberger WF, Sverdlov O. Handling Covariates in the Design of Clinical Trials. *Statistical Science*. 2008;23:404–419.
- [3] Moore RT, Moore SA. Blocking for sequential political experiments. *Political Analysis*. 2013;21:507–523.
- [4] Comparative Obstetric Mobile Epidural Trial (COMET) Study Group UK. Effect of low-dose mobile versus traditional epidural techniques on mode of delivery: a randomised controlled trial. *Lancet*. 2001;358:19–23.
- [5] Atkinson AC. Optimum biased coin designs for sequential clinical trials with prognostic factors. *Biometrika*. 1982;69:61–67.
- [6] Pocock S, Simon R. Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics*. 1975;31:691–694.
- [7] Taves DR. Minimization: a new method of assigning patients to treatment and control groups. *Clinical Pharmacology and Therapeutics*. 1974;15:443–453.
- [8] Senn S, Anisimov VV, Fedorov VV. Comparisons of minimization and Atkinson’s algorithm. *Statistics in Medicine*. 2010;29:721–730.
- [9] Huan X, Marzouk YM. Sequential Bayesian optimal experimental design via approximate dynamic programming. arXiv:160408320. 2016; Available from: <http://arxiv.org/abs/1604.08320>.
- [10] Bradley SP, Hax AC, Magnati TL. *Applied mathematical programming*. Boston: Addison Wesley; 1977.
- [11] Gittins JC, Jones DM. A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*. 1979;66:561–565.
- [12] Smith AL, Villar SS. Bayesian adaptive bandit-based designs using the gittins index for multi-armed trials with normally distributed endpoints. *Journal of Applied Statistics*. 2018;45(6):1052–1076. Available from: <https://doi.org/10.1080/02664763.2017.1342780>.
- [13] Williamson SF, Jacko P, Villar SS, et al. A Bayesian adaptive design for clinical trials in rare diseases. *Computational Statistics and Data Analysis*. 2017;113:136–153.
- [14] Villar SS, Rosenberger WF. Covariate-adjusted response-adaptive randomization for multi-arm clinical trials using a modified forward looking Gittins index rule. *Biometrics*. 2018;74:49–57.

- [15] Atkinson AC, Donev AN, Tobias RD. Optimum experimental designs, with sas. 2nd ed. Oxford: Oxford University Press; 2007.
- [16] Atkinson AC. Optimum biased-coin designs for sequential treatment allocation with covariate information. *Statistics in Medicine*. 1999;18:1741–1752.
- [17] Rainey C. Dealing with separation in logistic regression models. *Political Analysis*. 2016; 24:339–355.
- [18] Firth D. Bias reduction of maximum likelihood estimates. *Biometrika*. 1993;80:27–38.
- [19] Gelman A, Jakulin A, Pittau MG, et al. A weakly informative default prior distribution for logistic and other regression models. *Annals of Applied Statistics*. 2008;2:1360–1383.
- [20] Atkinson AC, Woods DC. Designs for generalized linear models. In: Dean A, Morris M, Stufken J, et al., editors. *Handbook of design and analysis of experiments*. Chapman & Hall/CRC; 2015. p. 471–514.
- [21] Powell WB. What you should know about approximate dynamic programming. *Naval Research Logistics*. 2009;56:239–249.
- [22] Cheng Y, Berry DA. Optimal adaptive randomized designs for clinical trials. *Biometrika*. 2007;94:673–687.
- [23] Ondra T, Jobjornsson S, Beckman RA, et al. Optimized adaptive enrichment designs. *statistical methods in medical research*. 2019;28:2096–2111.
- [24] Mueller P, Berry DA, Grieve AP, et al. Simulation-based sequential Bayesian design. *Journal of Statistical Planning and Inference*. 2007;137:3140–3150.
- [25] Bartroff J, Lai TL. Approximate dynamic programming and its applications to the design of phase I cancer trials. *Statistical Science*. 2010;25:245–257.
- [26] Woods DC, Overstall AM, Adamou M, et al. Bayesian design of experiments for generalized linear models and dimensional analysis with industrial and scientific application (with discussion). *Quality Engineering*. 2017;29:91–118.
- [27] Gentle JE. *Random number generation and monte carlo methods*. 2nd ed. New York: Springer-Verlag; 2003.
- [28] Ryan KJ. Estimating expected information gains for experimental designs with application to the random fatigue-limit model. *Journal of Computational and Graphical Statistics*. 2003;12:585–603.
- [29] Gelman A, Su YS. *arm: Data analysis using regression and multilevel/hierarchical models*; 2016. R package version 1.9-3; Available from: <https://CRAN.R-project.org/package=arm>.
- [30] Kahan BC, Forbes AB, Doré CJ, et al. A re-randomisation design for clinical trials. *BMC Medical Research Methodology*. 2015;15(1):96. Available from: <http://bmcmedresmethodol.biomedcentral.com/articles/10.1186/s12874-015-0082-2>.
- [31] Tackney M. *Design of sequential experiments with covariate information [dissertation]*. University of Southampton; 2019.
- [32] Hardwick J, Oehmke R, Stout QF. New adaptive designs for delayed response models. *Journal of Statistical Planning and Inference*. 2006;136:1940–1955.
- [33] Rosenberger WF. Randomized play-the-winner clinical trials: Review and recommendations. *Controlled Clinical Trials*. 1999;20:328–342.
- [34] Simon R. Adaptive treatment assignment methods and clinical trials. *International Biometric Society*. 1977;21:721–732.
- [35] Stallard N, Whitehead J, Todd S, et al. Stopping rules for phase II studies. *British Journal of Clinical Pharmacology*. 2001;51:523–529.
- [36] Whitehead J. Interim analyses and stopping rules in cancer clinical trials. *British Journal of Cancer*. 1993;68:1179–1185.
- [37] Berry DA. Monitoring accumulating data in a clinical trial. *Biometrics*. 1989;45:1197–1211.
- [38] Freedman LS, Spiegelhalter DJ. Comparison of Bayesian with group sequential methods for monitoring clinical trials. *Controlled Clinical Trials*. 1989;10:357–367.

Appendix A. Generation of simulated data

In the simulations in Section 5.1 and 5.2, the generation of responses can be a source of unnecessary variability. When comparing the myopic and nonmyopic designs, or the myopic and pseudo-nonmyopic designs, we generate the responses in the following way:

- (1) Generate a deviate u_i from the Unif(0, 1) distribution.
- (2) Set

$$y_i = \begin{cases} 1 & \text{if } u_i \geq \pi_i \\ 0 & \text{if } u_i < \pi_i \end{cases}. \quad (\text{A1})$$

The deviates u_i kept the same for the approaches being compared to try to minimize sources of random variability in the simulation.

Appendix B. Further Results

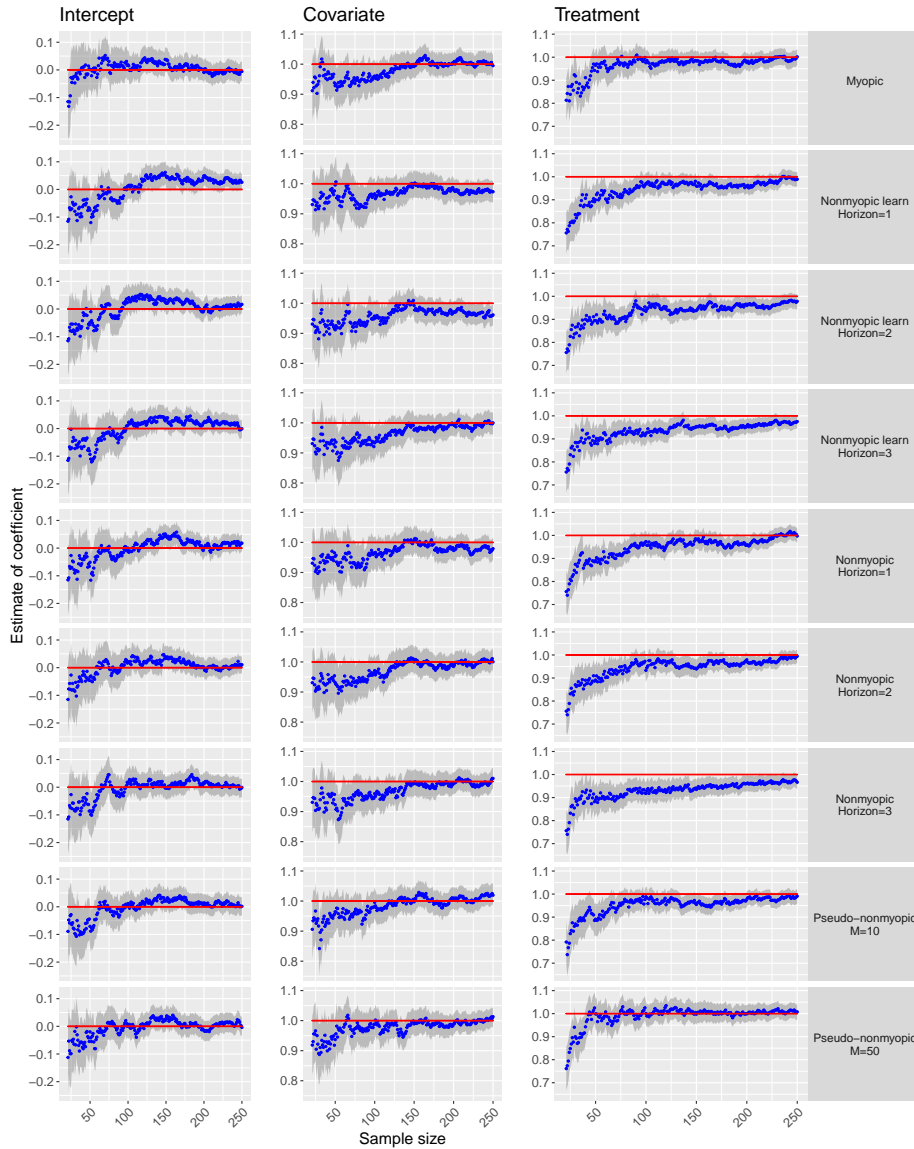


Figure B1.: Estimates of $\hat{\beta}_i$ for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 20. Results for the myopic approach, the nonmyopic approach and pseudo-nonmyopic approach to constructing D_A -optimal designs are shown. For the nonmyopic approach, we consider the horizon equal to $N = 1, 2$ and 3, and we consider both the case where the correct covariate distribution is known, and when it is unknown so the empirical covariate distribution is used (nonmyopic learn). For the pseudo-nonmyopic approach, we consider $M = 10$ and 50. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The true values of β are indicated in red.

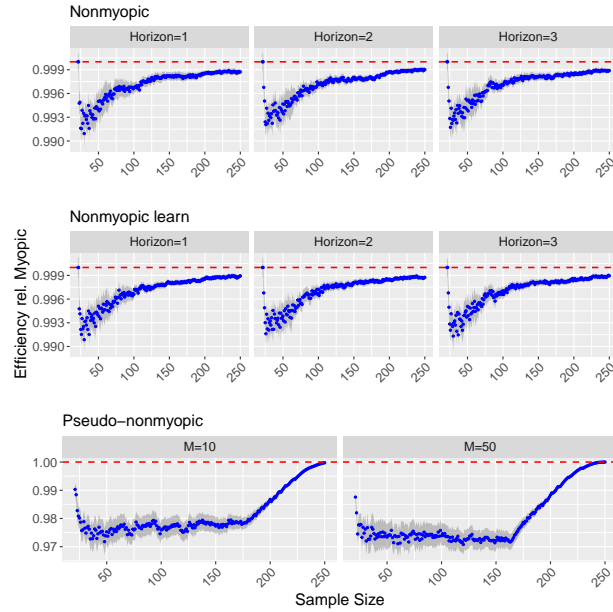


Figure B2.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 20. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.

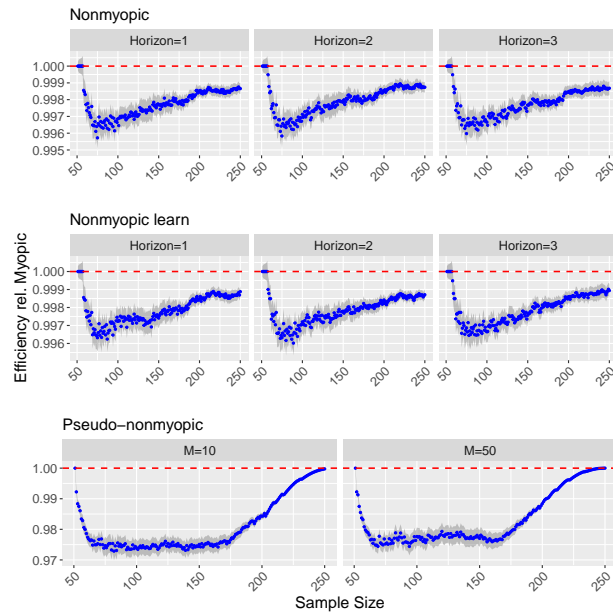


Figure B3.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 50. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.

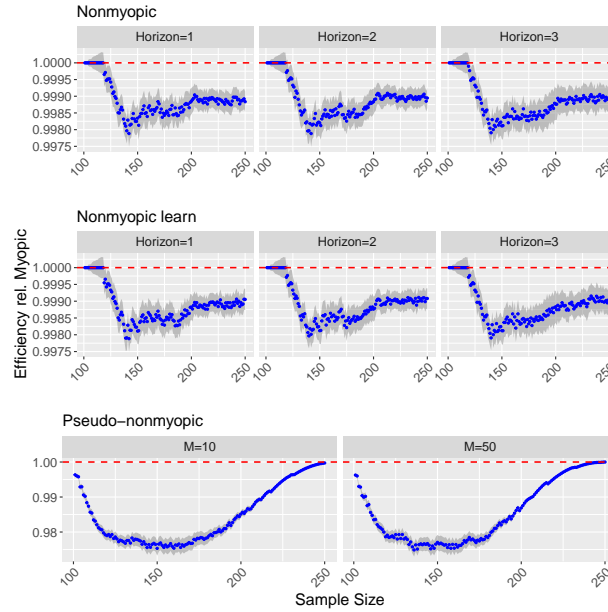


Figure B4.: Relative efficiencies of the nonmyopic and pseudo-nonmyopic D_A -optimal designs relative to the myopic D_A -optimal designs for the simulation with a single covariate, no treatment-covariate interactions and initial sample size of 100. Blue dots indicate the median of the estimate across simulations, and the grey area indicates $\pm 1.96 \times \text{MC error}$. The red line indicates equal efficiency to the myopic approach.

Appendix C. R code

We refer readers to Chapter 8 of [31] for a vignette of the R code and examples of how to implement the the myopic, nonmyopic and pseudo-nonmyopic approaches to optimal design. The function for the nonmyopic approach to optimal design for logistic regression is described in Section 8.2.2, and the function for the pseudo-nonmyopic approach is described in Section 8.3.

Appendix D. Myopic design algorithm

Algorithm 2 Myopic sequential optimal design: returns a design matrix for a sequentially conducted experiment. Inputs are the covariate values, \mathbf{Z}_n , and the number of subjects in initial design, n_0 .

```
1: function SEQOPTL( $\mathbf{z}_n, n_0$ )  
  
2:   Initialization  
3:   Construct initial design  $\mathbf{X}_{n_0}$  using an exchange algorithm for the first  $n_0$  subjects assuming  $\boldsymbol{\beta} = \mathbf{0}$ .  
4:   Observe responses  $\mathbf{y}_{n_0} = (y_1, y_2, \dots, y_{n_0})^T$ .  
5:   Fit the required generalized linear model to obtain a penalised MLE  $\hat{\boldsymbol{\beta}}_{n_0}$ .  
  
6:   for  $i$  in  $n_0 + 1$  to  $n$  do  
7:     Observe  $\mathbf{z}_i = (z_{i,1}, \dots, z_{i,s})^T$   
8:     Calculate  $\Psi(t \mid \mathbf{t}_{i-1}, \mathbf{Z}_i, \mathbf{y}_{i-1})$  for each treatment  $t \in \mathcal{T}$ .  
9:     Sample treatment allocated to subject  $i$  with probability (4).  
10:    Observe response  $y_i$ .  
11:    Refit model with responses  $\mathbf{y}_i$  and updated design matrix  $\mathbf{X}_i$ .  
12:    Obtain updated parameter estimates  $\hat{\boldsymbol{\beta}}_i$ .  
13:  end for  
  
14:  return  $\mathbf{X}_n$   
  
15: end function
```
