Sequential Experimental Designs for Generalized Linear Models

Hovav A. DROR and David M. STEINBERG

We consider the problem of experimental design when the response is modeled by a generalized linear model (GLM) and the experimental plan can be determined sequentially. Most previous research on this problem has been limited either to one-factor, binary response experiments or to augmenting the design when there are already sufficient data to compute parameter estimates. We suggest a new procedure for the sequential choice of observations that offers five important advantages: (1) It can be applied to multifactor experiments and is not limited to the one-factor setting; (2) it can be used with any GLM, not just binary responses; (3) both fully sequential and group sequential settings are treated; (4) it enables efficient design from the outset of the experiment; and (5) the experimenter is not constrained to specify a single model and can use the prior to reflect uncertainty as to the link function and the form of the linear predictor. Our procedure is based on a D-optimality criterion and on a Bayesian analysis that exploits a discretization of the parameter space to efficiently represent the posterior distribution. In the one-factor setting, a simulation study shows that our method is superior in efficiency to commonly used procedures, such as the "Bruceton" test, Neyer's procedure, or Wu's improved Robbins–Monro method. We also present a comparison of results obtained with the new algorithm versus the "Bruceton" method on an actual sensitivity test conducted recently at an industrial plant. Source code for the algorithms and examples throughout the article is available at *http://www.math.tau.ac.il/~dms/GLM_Design*.

KEY WORDS: Bayesian; Binary response; D-optimal; Dose-response; Poisson; Sensitivity test.

1. INTRODUCTION

Efficient experimental designs for generalized linear models (GLMs) depend on the unknown coefficients, so two experiments having the same model but different coefficient values typically will require different designs. Khuri, Mukherjee, Sinha, and Ghosh (2006) surveyed design issues for generalized linear models and note that the research is still very much at a developmental stage, especially for multifactor designs. In recent work (Dror and Steinberg 2006), we suggested a method for constructing robust D-efficient experimental designs for generalized linear models, which is based on clustering locally D-optimal designs.

In this article develop ideas for carrying out efficient sequential designs for GLMs. The idea of sequential design for such problems is natural. An efficient design requires knowledge of the parameters, so there should be a benefit to using all current data to choose the next design points. Our work is similar in spirit to that of Chaudhuri and Mykland (1993), who showed that sequential design in general nonlinear settings, including GLMs, could lead to fully efficient designs and asymptotically efficient maximum likelihood estimators. Our ideas differ from those of Chaudhuri and Mykland in our sample size focus; we are concerned with small samples, and thus rapid progress toward an efficient design, whereas Chaudhuri and Mykland emphasized only asymptotic properties. We provide an algorithm for efficient design beginning with the first observation; Chaudhuri and Mykland gave only general conditions for initial designs, which could be quite large. We use a Bayesian approach to get good small-sample designs, taking advantage of a computationally efficient representation of the posterior distribution of the coefficients.

Sequential design for binary response data has a rich history, going back to the work of Dixon and Mood (1948) and the sequential approximation scheme of Robbins and Monro (1951). There also has been considerable recent work, including that of Haines, Perevozskaya, and Rosenberger (2003), Ivanova and Wang (2004), Biedermann, Dette, and Zhu (2006), and Karvanen, Vartiainen, Timofeev, and Pekola (2007). We propose a more general treatment than those authors, who focused on single-factor experiments. As in previous work (Dror and Steinberg 2006), here we consider multifactor experiments. We allow the prior distribution to describe uncertainty over possible coefficient values and also ambiguity of the proper linear predictor, allowing the design to assist in determining the necessity of certain interactions, or between higher- and lower-order models. Designs can be either fully sequential or group sequential, in which the experimental plan is revised after each batch of k observations. We limit our discussion to parametric models and to the estimation of the model's coefficients. We assess designs through the D-optimality criterion, that is, the determinant of the Fisher information matrix for the parameters.

2. PREVIOUS WORK AND ITS LIMITATIONS

Here we review existing sequential design methods for GLMs. We begin with work on binary response data with a single design factor, which has been studied in some depth. Then we describe methods that can handle multifactor experiments. A common feature of all of the proposals is the need to address the fact that efficient designs depend on the parameters that they aim to estimate.

2.1 Design for Quantal Response

There is considerable interest in the quantal response setting, often under the name of "sensitivity tests" or "dose-response" studies. Typical applications are experiments aimed at learning about the sensitivity of a new explosive, as a function of the strength of a shock, or the potency or toxicity of a drug administered at different doses. We review both "static" and sequential designs.

The "Bruceton" up-and-down method (Dixon and Mood 1948) is still commonly used in sensitivity tests for explosives.

Hovav A. Dror is recent Ph.D. graduate (E-mail: *hovavdror@gmail.com*) and David M. Steinberg is Professor (E-mail: *dms@post.tau.ac.il*), Department of Statistics and Operations Research, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Ramat Aviv 69978, Israel. The authors thank the referees and an associate editor for their helpful comments. They also thank Benny Yoskovich and Yehoshua Borkowsky for letting them try the method at IMI.

Journal of the American Statistical Association
March 2008, Vol. 103, No. 481, Theory and Methods
DOI 10.1198/016214507000001346

This method was designed with the probit model in mind, with a normal distribution describing the sensitivity to the input stress of the population of explosive devices. The algorithm takes the first observation at a guessed value of the mean stress μ. Subsequent stresses are shifted up or down by a step size equal to a guessed value of the standard deviation σ.

The logistic regression model is a standard tool in quantal response. Abdelbasit and Plackett (1983) showed that the locally D-optimal design for the one-factor logistic model should divide observations equally between the two x values with expected responses of approximately .18 and .82. Because these input values will not be known before the experiment, the authors also suggested a sequential approach, beginning with a locally D-optimal design for a previous guess of the parameter values. At each stage, the parameters are estimated by maximum likelihood and the design procedure is repeated, using the revised estimates.

A drawback to the foregoing method is its reliance on maximum likelihood estimates (MLEs) of the parameter. Quite a few initial observations may be required before the conditions for existence of the MLEs are satisfied (Silvapulle 1981), especially if the initial parameter guesses are not good. Even when MLEs can be computed, they may suffer from severe bias in small samples, resulting in an inefficient choice of experimental sites.

Neyer (1994) addressed these issues, suggesting a D-optimality–based sensitivity test. He suggested a three-part procedure, in which the first part is designed to "close in" on the region of interest, the second part is devoted to determining unique estimates of the parameters, and only in its third and final part does the method use a local D-optimal design based on the MLEs for the parameters. Cases in which the MLEs are "wild" estimates are handled by restricting the values of the parameters at each step. Thus the strategy of Neyer's (1994) procedure is to minimize the time until we can get good estimates from the data, then use a local D-optimal design.

For nonsequential experiments, Chaloner and Larntz (1989) adopted a fully Bayesian approach, assigning proper prior distributions to the parameters. Their design criterion was a Bayesian extension of D-optimality, averaging the logarithm of the determinant of the Fisher information matrix over the prior. Using a prior distribution for the parameters leads to designs that are robust against a poor initial guess and that typically include more than just two input sites. (See Chaloner and Verdinelli 1995 for justification of the criterion.)

Beginning with Robbins and Monro (1951), much work on sequential designs was stimulated by experiments for estimating the quantile x_p associated with a particular probability, p, of response. Although the Robbins–Monro procedure is nonparametric, it also can be derived from a parametric model in which the response function, after suitable transformation, is linear near x_p. Wu (1985) exploited this correspondence to develop an improved version of the Robbins–Monro procedure. At each step, Wu's (1985) logit-MLE method uses all current data to fit a logistic regression model. The next input is the value of x for which the fitted logistic gives the desired probability p. Wu added rules to ensure that step sizes between successive inputs were not too large. The logit-MLE method requires data that are sufficiently rich to compute MLEs for the logistic regression. To generate initial data, Wu suggested either using the original Robbins–Monro method, which amounts to adopting a fixed slope, or using an ad hoc design of about 10 points, symmetric about the guessed value for $x_{.5}$. Wu (1986) extended the idea to regression settings with other GLM responses. Sitter and Wu (1999) looked at designs with two stages in the context of dosing trials for pharmaceuticals. Joseph (2004) showed how to further improve the efficiency of the Robbins–Monro scheme, especially for estimating extreme quantiles. Joseph, Tian, and Wu (2007) developed a Bayesian extension of Wu's approach. The primary extension covers uncertainty in the functional form of the response function. The authors also provided a Bayesian approach that can be applied when the response function is assumed to be correct and is specified up to a set of parameters. The Bayesian approach uses MAP estimates of the logistic parameters rather than MLEs, as proposed by Wu (1985).

There is an important difference between our work and the Robbins–Monro procedure and its descendants. Our goal is to precisely estimate the coefficients in a parametric GLM; the Robbins–Monro procedure is nonparametric. Choosing a parametric setting has important implications for the design. We are naturally led to D-optimality as a design criterion. The goal for Robbins and Monro was to generate a sequence of inputs that would converge to x_p. Thus a Robbins–Monro sequence can be expected to place many points close to one another, clearly a poor design if the goal is to estimate, say, a logistic or probit regression model. Moreover, completely different designs are needed to estimate different quantiles of the response curve, whereas we use a single design for all quantiles. Of course, we need to make assumptions (as to the form of the quantal response curve); the Robbins–Monro methods enjoy a robustness to those assumptions and converge, as Wu (1985) demonstrated, even when based on an incorrect approximation to the response.

2.2 Design for Multifactor Experiments

Woods, Lewis, Eccleston, and Russell (2006) were the first to seriously address the practical design of multifactor experiments for GLMs. They extended the Bayesian approach of Chaloner and Larntz (1989). In earlier work (Dror and Steinberg 2006), we derived similar designs but using a quicker computational scheme that involved clustering the points in locally optimal designs. Although our method does not attempt to directly optimize the Bayesian D-optimality criterion, it produces designs that perform very well.

Chaudhuri and Mykland (1993) proposed a completely sequential scheme for optimal design in parametric nonlinear problems, including GLMs. Their method involves three basic steps: (a) Run an initial design from which the parameters can be estimated by maximum likelihood; (b) find the next design site as the D-optimal augmentation of the design thus far, using the local D-optimality criterion corresponding to the current parameter estimates; and (c) obtain the next observation and update the MLE. They showed that under fairly general conditions, this procedure converges to the locally D-optimal design for the true parameter values and that the MLE would be asymptotically normal and efficient. Sinha and Wiens (2002) extended the ideas of Chaudhuri and Mykland. They showed how one also might incorporate some uncertainty as to the nature of the parametric model as well as heteroscedastic errors.

For practical experimenters, the weak link in Chaudhuri and Mykland's scheme is the initial design. Little guidance is given on how to set up the initial design or on how many runs might be needed. The reliance of the procedure on MLEs compounds this problem. We have already pointed out that even in the one-factor setting, GLMs may require 10 or more observations before the MLE can be computed. In multifactor experiments, many more observations may be needed.

3. METHODOLOGY

Our methodology uses Bayesian methods to jump start the sequential design process and to achieve a good initial design. Following Chaloner and Larntz (1989), we begin with a proper prior for the parameters in the model. However, full computation of the posterior distribution and of the Bayesian D-optimality criterion for augmenting the design is not feasible. Instead, we exploit a discrete representation of the posterior and a corresponding approximation to the design criterion. We begin by presenting our overall design strategy, followed by our method for representing the prior and the posterior. Finally, we discuss the selection of new design points.

To set the stage, let $\mathbf{x} = (x_1, \ldots, x_p)$ denote a potential design point for the experiment. We assume that $\mathbf{x}$ can take on values in a closed subset of R^p and, to be concrete, also assume that the actual factors have been scaled so that the experimental region is given by $[-1, 1]^p$. The response $Y(\mathbf{x})$ follows an exponential family with $\mathrm{E}\{Y(\mathbf{x})\} = \mu(\mathbf{x})$ and $h[\mu(\mathbf{x})] = \mathbf{f}^T(\mathbf{x})\boldsymbol{\beta}$, where h is the link function for the model and $\mathbf{f}^T(\mathbf{x})\boldsymbol{\beta}$ is the linear predictor.

The Fisher information matrix for a design d with observations at $\mathbf{x}_1, \ldots, \mathbf{x}_n$ is

$$\mathbf{I}(\boldsymbol{\beta}; d) = \sum_{i=1}^{n} \mathbf{f}(\mathbf{x}_i)\mathbf{f}^T(\mathbf{x}_i)w_i = \mathbf{F}^T\mathbf{W}\mathbf{F}, \tag{1}$$

where $w_i = 1/\{\mathrm{var}(\mu_i)[h'(\mu_i)]^2\}$ is a weight associated with the ith observation. The first term in the weight is the variance of the ith observation, as a function of its expectation, and the second term is related to the link function. In the final expression in (1), $\mathbf{F}$ is the regression matrix whose ith row is $\mathbf{f}^T(\mathbf{x}_i)$ and $\mathbf{W}$ is a diagonal matrix with the weights. Note that the information matrix depends on the parameter values through the weights.

The local D-optimality criterion for a particular parameter vector $\boldsymbol{\beta}$ is $|\mathbf{I}(\boldsymbol{\beta}; d)|$, where $|\mathbf{A}|$ denotes the determinant of the matrix $\mathbf{A}$. The Bayesian D-optimality criterion of Chaloner and Larntz (1989) is

$$\phi(d) = \int \log\bigl(|\mathbf{I}(\boldsymbol{\beta}; d)|\bigr)\, d\pi(\boldsymbol{\beta}), \tag{2}$$

where $\pi(\boldsymbol{\beta})$ is the prior distribution on $\boldsymbol{\beta}$.

We need a proper prior distribution on the parameters, so choice of the prior is obviously important. In the interest of robustness, we advise making the spread of the prior large. It is much better to exaggerate the spread of the prior than to risk a situation in which the true parameter values are at the extremes, or even outside the support of the prior. In any event, we anticipate that the posterior should be dominated by the data obtained in the experiment.

3.1 The Design Augmentation Strategy

Our strategy for adding a new site (or sites) to a given design includes three basic steps:

1. Determine the "augmentation horizon" as the number of observations, m, needed for a highly efficient locally D-optimal design at the prior median.
2. Generate an m-point augmentation to the current design.
3. Select an augmentation site (or sites) from the points found in the previous step, or their median.

The algorithm can be run in a fully sequential mode, adding one new site at each step, or in a group-sequential mode, adding a fixed number of sites. The number of sites added in step 3 usually will be determined by practical issues in running the experiment and so is set by the user. We use the same approach to find the first run(s) of an experiment, augmenting a "null" design.

The first two steps are important in the early stages of an experiment. With a small number of design points, information matrices will be singular, and design criteria such as the determinant give no useful information. Steps 1 and 2 provide a device that lets us circumvent the singularity problem.

We provide details on each step in our algorithm later in this section. First, we provide some ideas on reducing computation by approximating the posterior distribution of the parameters and the design criterion $\phi(d)$.

3.2 A Discrete Representation of the Posterior

Our setup could be used to implement a fully Bayesian approach. Based on the data at hand, compute the posterior distribution of $\boldsymbol{\beta}$ and use it as the basis to find the next design point. But computing the exact posterior distribution at each iteration of the design is not a trivial task; it requires substantial computation. Moreover, it is not necessary; a much cheaper (computationally) alternative is sufficient.

We represent the posterior using a large (say, $N = 10{,}000$) discrete set of vectors sampled from the prior, $\boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_N$. These vectors could be a random sample or a quasi-random sample (such as in Niederreiter 1988). At any stage of the experiment, the likelihood $L(\boldsymbol{\beta}_u)$ can be easily and rapidly computed for each of these vectors. We normalize the likelihoods across our sample, generating weights $r_u = L(\boldsymbol{\beta}_u)/\sum_{v=1}^{N} L(\boldsymbol{\beta}_v)$. We now can estimate functionals of the posterior as weighted summaries of the vectors sampled from our prior; for example, the posterior mean vector can be estimated as $\sum r_u\boldsymbol{\beta}_u$. This is essentially an importance sampling scheme, with the prior serving as the base sampling distribution and the importance weights coming from the fact that the posterior divided by the prior is proportional to the likelihood.

The same scheme can easily handle uncertainty in the form of the linear predictor. Suppose that the experimenters believe that a first-order logistic regression model might be appropriate for their data but are concerned that several two-factor interactions might be present. The prior distribution then can be set out in two stages. At the first stage, assign prior probabilities to each of the models that are entertained. At the second stage, conditional on each model assign prior distributions to its coefficients. The sampling scheme proceeds as before, with the

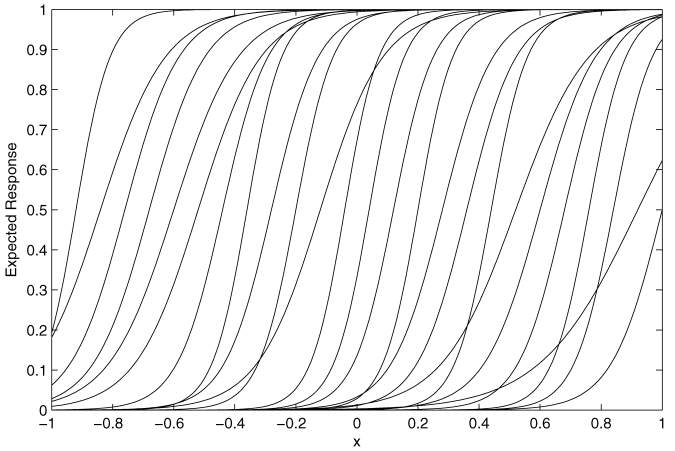

Figure 1. A random sample of 25 logistic probability curves from a prior in which the LD50 has a uniform distribution on the interval $[-1, 1]$ and the slope has a uniform prior on $[6, 18]$.

probabilities from the first stage used to dictate what fraction of vectors are sampled from each of the prior distributions on the coefficients.

Figures 1 and 2 illustrate our ideas for representing the posterior, in the context of a simple logistic regression model with the input factor limited to the interval $[-1, 1]$. Following Chaloner and Larntz (1989), we parameterize the model in terms of the slope θ and the LD50, μ, so that $p(x; \mu, \theta) = 1/[1 + \exp\{-\theta(x - \mu)\}]$. We adopt independent uniform priors on both coefficients, with $\mu \sim \mathrm{U}[-1, 1]$ and $\theta \sim \mathrm{U}[6, 18]$. Figure 1 presents the expected response $p(x; \mu, \theta)$ as a function of x for a random sample of 25 vectors from the prior. There is great variety in the prior response curves, in accordance with our recommendations to choose a broad prior. One choice to begin the experiment is to use the locally D-optimal design for the prior mean $\mu = 0$ and $\theta = 12$. That design takes half of the observations at each of $x = \pm.13$. Now suppose that the first run is made with $x = .13$ and results in $Y = 0$. Many of the logistic curves in Figure 1 considered positive response at $x = .13$ as "a sure thing," so this single observation is enough to cast strong doubt on that part of the prior. Figure 2 shows how the result is reflected in our representation of the posterior distribution. Panel (a) shows a quasi-random sample of 10,000 points from the prior; in panel (b), we have simply deleted all of those points in the prior with likelihood $<.05$. (Of course, we never actually delete such points; we just downweight them.) There are two important lessons here. First, although we clearly cannot compute MLEs at this point, Figure 2 shows how much we can learn from this one observation. Second, having observed a "nonresponse" at $x = .13$, the wisdom of proceeding with the next observation at $x = -.13$ is dubious indeed.

3.3 Approximate Design Criteria

The fully Bayesian approach for adding points to an existing design is to use $\phi(d)$, from (2), averaging at each step with respect to the posterior distribution. As noted in the preceding section, we can approximate this average by $\sum_{u=1}^{N} r_u \log|\mathbf{I}(\boldsymbol{\beta}_u; d)|$. Thus we have the criterion

$$\phi_1(d) = \sum_{u=1}^{N} r_u \log|\mathbf{I}(\boldsymbol{\beta}_u, d)|. \tag{3}$$

Optimizing this criterion is not trivial, and we would like something that can be computed even faster. Our suggestion is to replace the average by $\log|\mathbf{I}(\boldsymbol{\beta}; d)|$ at a single point, as done by Chaudhuri and Mykland (1993). We evaluate at the posterior median for each of the parameters, although other measures of the center of the distribution also could be used. Again, we use the weighted representation of the posterior to estimate the median. For example, we estimate the median of β_0 by $\tilde{\beta}_0$, where half of the weight belongs to values $\leq \tilde{\beta}_0$ and half to larger values; if $\beta_{u,0},\ u = 1, \ldots, N$, are the N discrete β_0 values in our sample, and if r_u is the weight attributed to β_u, then $\tilde{\beta}_0$ is found by sorting $(\beta_{u,0}, r_u)$ in ascending $\beta_{u,0}$ order and choosing $\tilde{\beta}_0 = \beta_{(g),0}$ such that $\sum_{u=1}^{g} r_{(u)} \geq \frac{1}{2}$ and $\sum_{u=g}^{N} r_{(u)} \geq \frac{1}{2}$. Thus we obtain the criterion

$$\phi_2(d) = \log|\mathbf{I}(\tilde{\boldsymbol{\beta}}, d)|. \tag{4}$$

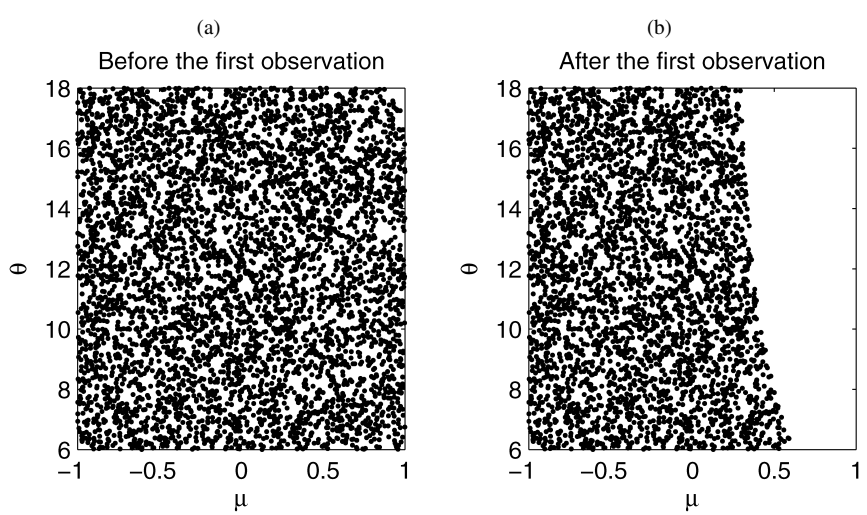

Figure 2. Discrete approximation to our distribution for the parameters of a one-factor logistic model, in which μ is the LD50 and θ is the slope. (a) A large quasi-random sample of points from the uniform prior. (b) The sample with the prior points that have very low weight following a "no response" at $x = .13$ deleted.

3.4 The Augmentation Horizon

When the number of existing design points plus the number of augmentation sites is less than the number of parameters in the model, the corresponding information matrices will be singular. Thus none of our criteria will provide any information that can be used to decide which design sites will be the most informative ones. To avoid such problems, we first determine an *augmentation horizon*, m. Then, at each augmentation step, we begin by finding an m-run augmentation. The horizon itself is determined at the start of the experiment, as the number of observations required for a design that has high local D-efficiency at the prior median of the parameters, $\tilde{\boldsymbol{\beta}}^{(0)}$. If the prior includes more than one linear predictor, then the horizon is determined using the predictor with the largest number of regression functions.

We now find locally D-optimal designs, maximizing $\phi_2(d)$ at the prior median for sample sizes $n = p, \ldots, P$, where p is the number of terms in the linear predictor and P is about $4p$. Let $d_n^{(0)}$ denote the design with n runs. These designs are found using our earlier algorithm (Dror and Steinberg 2006), which applies an exchange algorithm (Federov 1972) to a transformed regression matrix, $\tilde{\mathbf{F}} = \mathbf{F}\mathbf{W}^{1/2}$. This algorithm is suitable for constructing and augmenting local D-optimal designs for models of high dimension with any GLM response.

To compare the designs, we make a standard modification to the ϕ_2 criterion that removes its dependence on sample size. We regard $d_n^{(0)}$ as a probability measure, with mass $1/n$ on each design point. We also take the pth root of the determinant of the information matrix, which scales its size proportional to the number of observations. The resulting criterion is

$$\phi_3(d) = (1/p)\phi_2(d) - \log(n), \tag{5}$$

where n is the actual number of observations in d. Let d^* denote the design (from among $d_p^{(0)}, \ldots, d_P^{(0)}$) that maximizes $\phi_3(d)$. Define the efficiency of $d_n^{(0)}$ as

$$\mathrm{Eff}\big(d_n^{(0)}\big) = \exp\big\{\phi_3\big(d_n^{(0)}\big) - \phi_3(d^*)\big\}. \tag{6}$$

We choose the horizon m as the smallest value of n for which $\mathrm{Eff}(d_n^{(0)})$ is at least 99%.

Figure 3 illustrates the idea of the augmentation horizon in the context of a crystallography experiment presented by Woods et al. (2006). The goal was to design a four-factor experiment, with binary outcomes, to estimate a first-order logistic regression model. The experimenters' best guess for the intercept was 0, and prior guesses for the four slopes were 7, 8, −3, and .5. Using this parameter vector, locally D-optimal designs with $n = 5, \ldots, 24$ runs were generated. Figure 3 plots the corresponding D-efficiencies and shows that $m = 8$ is a good choice for the augmentation horizon.

3.5 One-Point Augmentations

Here we describe a fully sequential design approach, in which new observation sites are chosen one at a time. The algorithm depends on the augmentation horizon m found at the start of the experiment. The exact determination also depends on whether or not sufficient runs have been made to achieve a nonsingular information matrix.

The algorithms is as follows:

1. Find a locally D-optimal m-run augmentation to the existing design, maximizing ϕ_2 at the current parameter median. If the prior includes more than one linear predictor, then use the linear predictor with the highest posterior probability.
2. Generate a candidate set for the augmentation consisting of the m points found in the previous step and their coordinatewise median.

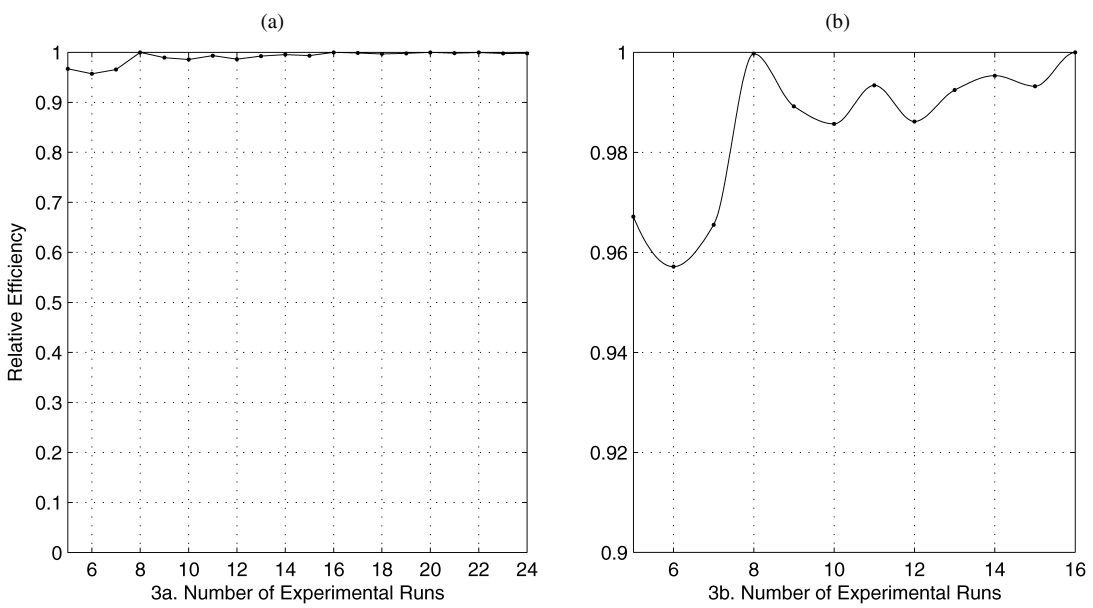

Figure 3. Choosing the augmentation horizon for a four-factor logistic regression model. (a) The relative efficiency for design measures corresponding to locally optimal designs with $n = 5, \ldots, 24$ runs. (b) Results for $n = 5, \ldots, 16$ runs highlighted by magnifying the relative efficiency. Using eight or more runs gives very high efficiency, so we set the augmentation horizon to 8.

3. If the design points run thus far provide a nonsingular information matrix, then choose the next design point as the candidate that gives the best result for ϕ_1 when added to the current design.
4. If the design thus far does not provide a nonsingular information matrix, then choose the next design point from among the candidates by comparing values of ϕ_1 for designs that consist of the points run thus far, the m-run augmentation, and the candidate. Note that these designs may include two copies of the augmentation candidate.

Step 1 uses the computationally fast approximate design criterion ϕ_2 to produce a limited set of candidate points. The better, but more computationally intensive ϕ_1 is used at step 3 or 4 only to evaluate this small set. In step 2 we add the median of the points found in step 1 to our candidate set. Locally D-optimal designs often tend to push design points away from the center of the design region, but sometimes the center is a better compromise choice for the design.

Step 4 provides a simple fix for early stages in the experiment when the information matrix is singular. Note that singularity of the information matrix is a function only of the regression matrix F, so it is sufficient to check for singularity at the posterior median.

3.6 Group-Sequential Designs

Between the two extremes of fixing the entire experiment in advance and the fully sequential approach in which we update the plan after each observation, there is the possibility of a group-sequential design. The term "group-sequential" refers to cases for which we obtain data in groups of $k > 1$ observations. We proceed in much the same way as for one-point augmentations. However, we now want to add further assurance that we find design points that are robust with respect to our uncertainty about the parameters. Choosing the points by directly optimizing ϕ_1 is slow computationally, so we again adopt the tactic of optimizing ϕ_2. The difference is that we now do so with respect to k different points representing the parameter space rather than just the posterior median. We choose these points using a weighted K-means clustering algorithm, as in the robust procedure for nonsequential designs in our previous work (Dror and Steinberg 2006).

The algorithm thus proceeds as follows:

1. Run K-means clustering (with $K = k$) on the points representing the posterior distribution of the parameters. Weight the points in the analysis by the normalized likelihoods, for determining both the clusters and the cluster centroids. If the prior includes more than one linear predictor, then represent each one proportional to its posterior weight, with clustering done separately.
2. For each cluster centroid, find a locally D-optimal m-run augmentation to the existing design, maximizing ϕ_2.
3. Generate a candidate set for the augmentation consisting of the mk points found in the previous step and of the k coordinatewise medians.
4. If the design points run thus far provide a nonsingular information matrix, then choose the next k design points as the candidates that give the best result for ϕ_1 when added to the current design. This can be done efficiently, beginning with a random choice of k candidates and proceeding with an exchange algorithm.
5. If the design thus far does not provide a nonsingular information matrix, then choose the next k design points from among the candidates by comparing values of ϕ_1 for designs that consist of the points run thus far, the m-run augmentation at the median, and k of the candidates. The best k candidates can again be chosen by an exchange algorithm.

We have found that the computational burden can be further reduced without harming design efficiency by an additional simplification. When making exchanges, require that the new point be one that came from the same centroid as the point that it replaces.

4. APPLICATION TO A SENSITIVITY EXPERIMENT

In this section we depict an application of our sequential design algorithm to a sensitivity experiment conducted in June 2006 at an industrial plant. The experiment's objectives were to estimate the sensitivity curve in general, and in particular to verify a manufacturer's statement that the explosives will not detonate at 12 V (being a safe voltage) and will detonate ("all fire") at 25 V. Quantitatively, the requirement was to show, using probit regression, that the probability of detonation at 12 V is $<5\%$ and the probability of detonation at 25 V is $>95\%$; that is, the 95% probit confidence interval for the expected response should be $<5\%$ at 12 V and $>95\%$ at 25 V.

The plant engineers agreed to use our methods, but only after first running the experiment following their standard format, which was based on the Bruceton method of Dixon and Mood (1948). Although the standard method was geared toward estimating quantiles with about 50% response probability rather than more extreme quantiles, it was the method of choice for all sensitivity tests in the plant.

The experimenters began with limited prior knowledge. They could not say what voltage would provoke a response from half of the observations or even whether this value was within the specified range of 12–25 V; not much was known about the dispersion as well, with the possibility of a very slow increase from the "no-fire" zone to the "all-fire" zone to a very steep curve.

Together with the plant engineers, we formulated the following prior distribution, which reflects this (lack of) information. The prior takes $(x - \mu)/\sigma$ as the parametrization and a lognormal distribution for both the mean and dispersion, with $\mu \sim \text{lognormal}(\log(17), .5^2)$ and $\sigma \sim \text{lognormal}(\log(.7), 1^2)$. Figure 4 presents a sample of 150 curves given possible values for (μ, σ) sampled from the prior distribution.

It can be seen that there are both very steep curves (with a fraction of a volt separating the "no-fire" and "all-fire" zones) and very flat curves offering the possibility that there is a nonnegligible likelihood of detonation at 12 V and of no detonation at 25 V. The industrial plant's procedure began with a "screening" phase following the Bruceton procedure with a rather large step size of 1 V. After 9 observations, when the experimenters felt they knew the region, they began a new Bruceton procedure with a smaller step size of .25 V; this second Bruceton continued for 31 trials, so a total of 40 runs was performed. For the new algorithm, 20 trials were allocated. The locations at which the observations were taken by each method, and their outcomes are presented in Figure 5.

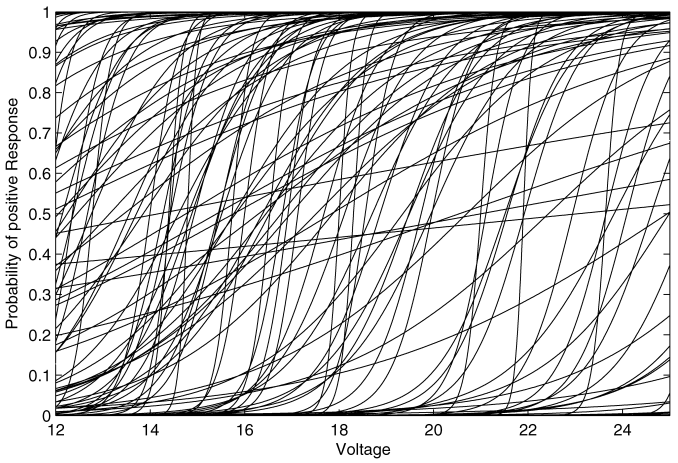

Figure 4. Sensitivity experiment. A sample of possible curves from the prior distribution elicited from the plant engineers.

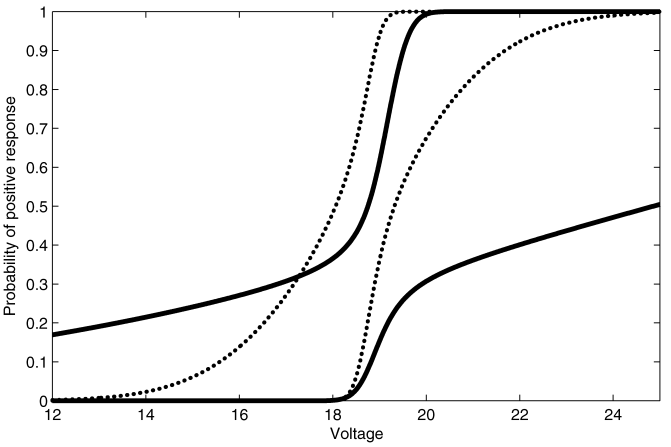

Figure 6. Comparison of the plant format (——) and our algorithm (·····) after 20 observations from each. The lines are pointwise 95% confidence intervals for the probability of response.

Figure 6 compares the analysis of outcomes of both methods after 20 trials. The curves in the figure represent the 95% confidence intervals for the expected response as a function of voltage value, as produced by a probit regression. Figure 7 is similar to Figure 6 but compares the analysis based on all 40 observations from the Bruceton method with that using the 20 observations from the new algorithm. Even in the extreme comparison where the Bruceton method benefits from twice as many observations as the new method, the latter provides better results. In the strip where the probit 95% confidence interval is $>.05$ and $<.95$ for the probability of detonation, the Bruceton method yields a strip of 15–23.5 V, compared with the narrower strip 15–22.5 V yielded by the new algorithm.

5. COMPARISON WITH EXISTING METHODS

Here we compare our design algorithm with other methods that have been proposed. In one-factor experiments, two leading approaches are the method of Neyer (1994), which is also directed toward precise estimation of the parameters in a GLM, and the collection of Robbins–Monro type procedures.

We first present simulation results comparing our method and Neyer's method. We also include the Bruceton method (Dixon and Mood 1948) in these simulations. Our experience has been that the Bruceton method remains in widespread use, despite Neyer's results showing that his method is more efficient.

We also present simulations comparing our method with the RM scheme of Wu (1985). As we noted earlier, direct comparison of our method with RM schemes is problematic, because the latter are nonparametric, and their goal is to converge to a specified quantile. We limit our comparisons here to the RM goal of quantile determination.

For multifactor experiments, we are not aware of any competing methods for the sequential design of experiments for GLMs. Both Chaudhuri and Mykland (1993) and Sinha and Wiens (2002) addressed this problem, but only for augmenting designs that are already sufficiently rich to allow for computation of MLEs. Chaudhuri and Mykland did not discuss the choice of initial designs, and Sinha and Wiens did so only in the context of examples, where they provided ad hoc solutions. We present some simulation results comparing our algorithm for sequential designs with our earlier algorithm (Dror and Steinberg 2006) for robust nonsequential experiments.

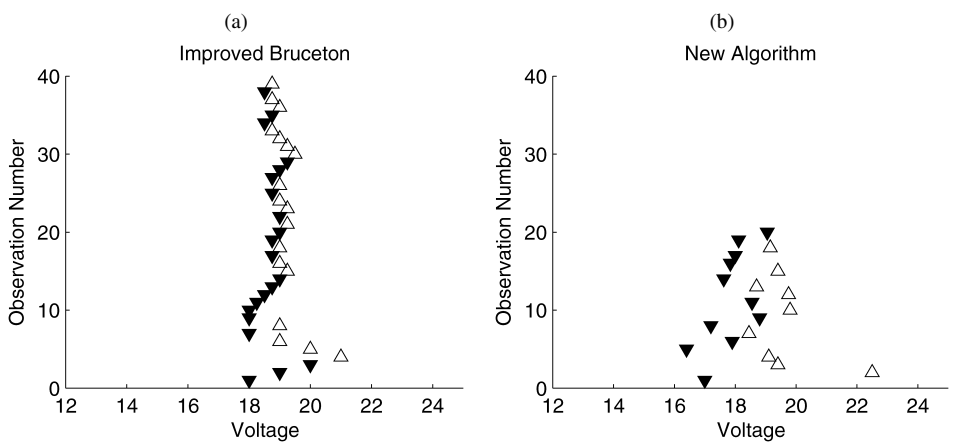

Figure 5. Sensitivity experiment; record of observations' location and outcome. Positive responses are designated by empty triangles pointing up; failures to respond, by solid triangles pointing down. (a) The 40 runs using the standard experimental format in the plant. (b) The 20 runs using our design algorithm.

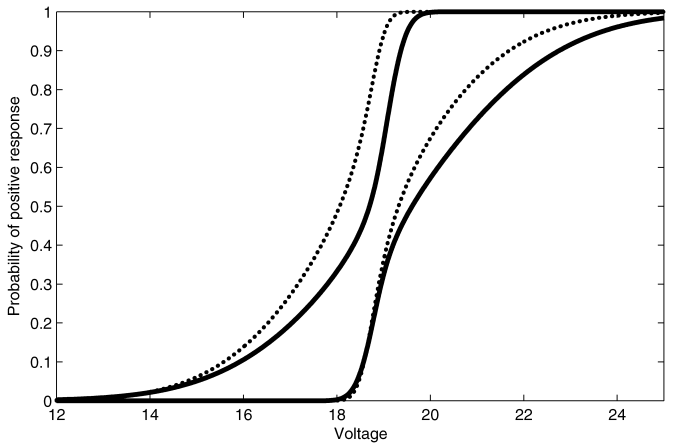

Figure 7. Comparison of the Bruceton method after 40 observations (——) our method with 20 observations (·····). The lines are pointwise 95% confidence intervals for the probability of response.

5.1 Comparison With Neyer and Bruceton

We conducted simulations to compare our method with that of Neyer and to the Bruceton method for a one-factor setting like the sensitivity experiment in the previous section. Figure 8 presents the different "true" models that were used and shows that we consider cases with both true mean <12 V or >25 V, very steep curves, and very flat curves. The bold curve emphasizes the case where the true model is at the center of our prior distribution.

Each of the methods requires a slightly different input from the user. For our method, a prior distribution is needed. We used the same prior as before: $\mu \sim \text{lognormal}(\log(17), .5^2)$ and $\sigma \sim \text{lognormal}(\log(.7), 1^2)$. This broad prior is appropriate when prior knowledge is weak so that robustness to true parameter values is required.

The Bruceton method (Dixon and Mood 1948) requires an initial guess for μ and a step size, based on a probit model guess for σ. We use the center of our prior distribution, 17, as the guess for μ, and the chosen step size is 1.28 V, after adjusting the difference of σ for the logit and probit models. Neyer's (1994) algorithm requires guesses for a lower bound and an upper bound for the mean, which we take to be $\mu_{\min} = 12$, $\mu_{\max} = 25$, and a guess of the standard deviation, for which we use the center of the prior distribution, .7.

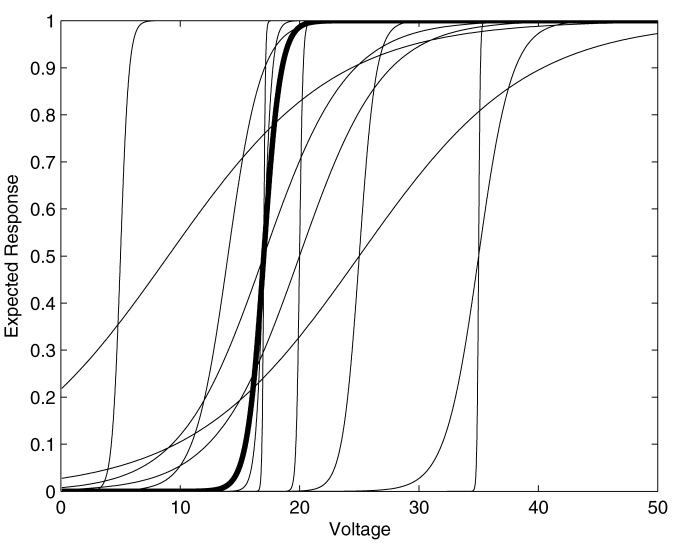

Figure 8. "True" probit models for the Monte Carlo comparison of techniques. The bold curve represents the case where the best guess is in fact the true model.

Table 1. Median D-efficiency for the new algorithm, Neyer's method, and the Bruceton method, based on 100 simulated experiments of a one-factor sensitivity study, for 13 true probit models

	True	True	16 observations			48 observations		
Case	μ	σ	Bruceton	Neyer	New	Bruceton	Neyer	New
1	17	.07	.002	.40	.31	.002	.74	.63
2	35	.07	.0005	.11	.22	.002	.32	.45
3	20	.14	.18	.59	.48	.19	.78	.76
4	5	.35	.37	.58	.59	.70	.82	.77
5	17	.35	.67	.77	.69	.67	.86	.84
6	17	.7	.85	.80	.77	.86	.88	.88
7	25	.7	.64	.69	.74	.79	.84	.85
8	14	1.4	.77	.74	.83	.79	.86	.89
9	35	1.4	.25	.61	.67	.64	.83	.85
10	17	3.5	.41	.61	.72	.50	.86	.88
11	20	3.5	.43	.49	.68	.54	.86	.86
12	9	7	.31	.51	.59	.42	.86	.85
13	25	7	.33	.44	.57	.43	.85	.83

We conducted experiments of length 16 and 48 observations. For each model, sample size, and method, we carried out 100 repetitions and computed the D-efficiency of the resulting design. Tables 1 and 2 summarize the D-efficiencies by their medians and 5% quantiles.

Figure 9 offers a graphical representation of the median D-efficiencies for the three methods for experiments with 16 observations. The lower part of Figure 9 shows the values of the true parameters, in the same order as in Table 1; for example, case 1 is the parameter vector in the first row; $(\mu, \sigma) = (17, .07)$, and case 6 is the best guess, or center of the prior distribution, with $(\mu, \sigma) = (17, .7)$. The upper part shows the median D-efficiencies.

Table 2. The 5% quantile of the D-efficiency for the new algorithm, Neyer's method, and the Bruceton method, based on 100 simulated experiments of a one-factor sensitivity study, for 13 true probit models

	True	True	16 observations			48 observations		
Case	μ	σ	Bruceton	Neyer	New	Bruceton	Neyer	New
1	17	.07	.002	.37	.25	.002	.65	.50
2	35	.07	.0001	.11	.18	.002	.28	.41
3	20	.14	.15	.47	.36	.18	.63	.56
4	5	.35	.28	.45	.45	.64	.71	.64
5	17	.35	.63	.61	.48	.65	.69	.66
6	17	.7	.76	.67	.61	.83	.76	.76
7	25	.7	.51	.55	.57	.74	.67	.72
8	14	1.4	.57	.62	.60	.69	.64	.73
9	35	1.4	.14	.51	.51	.57	.62	.72
10	17	3.5	.30	.27	.36	.40	.54	.63
11	20	3.5	.27	.25	.41	.40	.49	.71
12	9	7	.15	.16	.23	.30	.51	.61
13	25	7	.16	.15	.25	.30	.39	.60

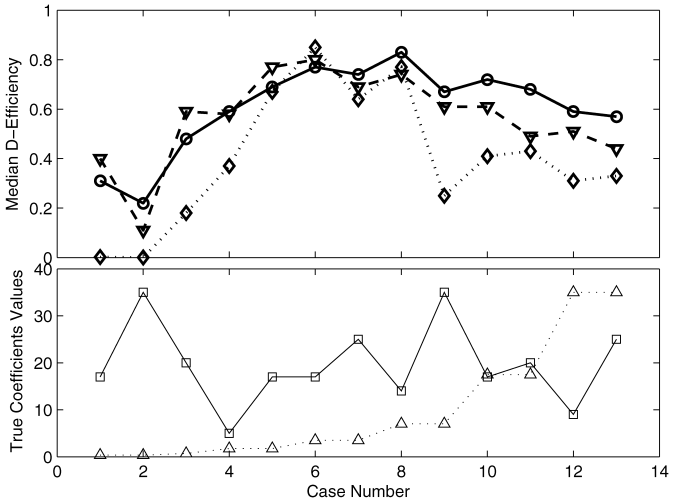

Figure 9. Comparison of median D-efficiencies for 16-run sensitivity experiments, for 13 true probit models (—o—, new algorithm; - ▽ -, Neyer's algorithm; ··◇··, Bruceton). The model parameters are plotted in the bottom panel (—□—, true μ; ··△··, 5* true σ).

Compared with Neyer's (1994) procedure, our algorithm handles better the uncertainty in the parameter values. D-efficiencies for Neyer's (1994) algorithm are roughly equal to or slightly better than the new algorithm when the true σ is smaller than its guess (or at the center of the prior distribution), but the proposed algorithm is significantly superior when the true σ is larger than its best guess. Using the new algorithm with a stochastically smaller prior for σ, such as $\sigma \sim \text{lognormal}(\log(\frac{7}{2}), 1^2)$), gives results that are generally superior to those of Neyer's (1994) algorithm for both large and small σ values. But this choice does not exploit the efficiency for large true σ values to its maximum potential, and we do not recommend it. Our results show that the Bruceton method is inferior to the other techniques.

The comparison so far has assumed vague knowledge of the parameters as expressed by the prior distribution. Often higher-fidelity prior information is available, say reducing the variance value in the prior so that $\mu \sim \text{lognormal}(\log(17), .2^2)$ and $\sigma \sim \text{lognormal}(\log(.7), .75^2)$. Note that this prior still has non-negligible mass for μ <12 and >25 and also allows various σ values. Results for this prior, given in Table 3, show that the proposed algorithm is then even more efficient compared with Neyer's (1994) algorithm.

Table 3. Median D-efficiency for 16- and 48-run sensitivity experiments using a narrower prior for the new algorithm

Observations	True μ	True σ	Bruceton	Neyer	New algorithm	
					Wide prior	Narrow prior
16	18	.37	.72	.76	.77	**.83**
	17	.7	.85	.80	.75	**.83**
	3	1.75	.71	.71	.76	**.71**
48	18	.37	.74	.86	.87	**.89**
	17	.7	.86	.89	.88	**.91**
	3	1.75	.74	.86	.89	**.85**

5.2 Comparison With the Logit-MLE Method

We now consider the problem of estimating quantiles of the response curve for binary data and compare the estimates obtained with our design algorithm to those found by the improved Wu logit-MLE method (Wu 1985) and the Bayesian version of that method proposed by Joseph et al. (2007). In one simulation, the true response curve is a linear logistic model, so that $p(x; \mu, \theta) = 1/[1 + \exp\{-\theta(x - \mu)\}]$. We set $\mu = 0$ and $\theta = 5$, so that values of x outside the interval $[-.5, .5]$ have extreme probabilities. In the second simulation, the true response curve is a skewed logistic model, for which $p(x) = 1/[1 + \exp\{-5x\}]^2$. The skewed logistic model was also studied by Wu (1985) and is useful for examining robustness of the methods when the true model does not match assumptions. For both methods, experiments were limited to $[-1, 1]$.

We ran our design for both simulations with a linear logistic regression model. Our prior for μ was uniform on $[-1, 1]$, and our prior for θ was lognormal with a mean of log(5) and a standard deviation of 1.5. The priors express much uncertainty as to the response curve. We used the same prior for the Bayesian version of the logit-MLE (Joseph et al. 2007).

We ran Wu's method with an upper bound of 500, in accordance with the recommendations in that article, and adjusting to our slope of 5. As a starting design, we took 10 observations equally spaced between -1 and 1. If the MLE for a linear logistic model could be computed from the initial data, then we began Wu's iterative scheme with the 11th observation. Otherwise, we took eight more equally spaced observations on $[-1, 1]$ before beginning the iterations. The estimate of the quantile from an n-run design is the recommended setting for the next run.

We considered estimation of $x_{.5}$ and $x_{.8}$ for the logit model. For the skewed logistic, we also considered $x_{.2}$. Sample sizes of 20 and of 50 were examined. Tables 4 and 5 summarize the results of 1,000 simulated experiments for the logistic and skewed logistic response models. For both models, the best quantile estimates were achieved by the Bayesian logit-MLE scheme, and the worst estimates were those from the original logit-MLE scheme. Our method gave intermediate results. Furthermore, as noted earlier, our method estimates *all* of the quantiles from a single experiment, whereas the results from the logit-MLE method are based on separate experiments for each quantile.

Table 4. Simulation results of the new algorithm and Wu's logit-MLE (l-MLE) method for estimating quantiles from a logistic response curve

Observations	Quantile method	$x_{.5}$		$x_{.8}$	
		Bias	SD	Bias	SD
20	l-MLE	−.001	.127	−.073	.201
	Bayes-l-MLE	.005	.115	.026	.141
	New	−.004	.128	.006	.193
50	l-MLE	.002	.089	−.061	.156
	Bayes-l-MLE	.003	.073	.007	.086
	New	.009	.078	.012	.110

NOTE: SD, standard deviation.

Table 5. Simulation results of the new algorithm and Wu's logit-MLE (l-MLE) method for estimating quantiles from a skewed logistic response curve

Observations	Quantile method	$x_{.2}$ Bias	$x_{.2}$ SD	$x_{.5}$ Bias	$x_{.5}$ SD	$x_{.8}$ Bias	$x_{.8}$ SD
20	l-MLE	.074	.159	.013	.126	−.059	.188
	Bayes-l-MLE	−.015	.115	.007	.100	.020	.129
	New	−.018	.148	.015	.105	.014	.175
50	l-MLE	.075	.135	.003	.085	−.070	.143
	Bayes-l-MLE	−.002	.062	.001	.055	.015	.080
	New	.002	.079	.022	.067	.010	.103

NOTE: SD, standard deviation.

5.3 Multifactor Experiment

Here we return to the crystallography experiment with four factors and binary outcomes introduced by Woods et al. (2006), who considered 16-run and 48-run designs for estimating a first-order logistic regression model. Earlier (Dror and Steinberg 2006), suggested a robust one-stage experiment for the crystallography experiment. We adopted independent uniform priors on each of the 5 model parameters and evaluated the design using a sample of 10,000 vectors from the prior, calculating the local D-efficiency of the design versus each of them and reporting the median and minimum local D-efficiencies. We applied the sequential algorithm with the same prior, using a Monte Carlo evaluation as before, with a small modification; each of the 10,000 iterations was performed with a different true coefficient value, sampled randomly from the prior. This allows a legitimate comparison of the median local D-efficiency achieved through the robust one-stage design and the proposed sequential design. Table 6 displays the comparison.

As expected, performing the experiment sequentially considerably improves the efficiency. Note that the 5% quantile D-efficiency for the sequential design is better than the median D-efficiency for the robust one-stage designs.

5.4 Group-Sequential Designs

We continue with the crystallography example, presenting in Table 7 the addition of a group-sequential design, with each group including 16 observations.

We see that even for a large group size of 16 observations, the group-sequential procedure's D-efficiency is closer to the fully sequential results than to the one-stage robust design's efficiency.

For the 16-observation experiment, the group-sequential design is in fact a one-stage design. It is desired that the proposed sequential procedure provide an efficiency as close as possible to that of a one-stage design if all of the experiments are grouped together into a single batch. Indeed, Table 7 shows that our group-sequential algorithm succeeded in reproducing the one-stage results we obtained earlier (Dror and Steinberg 2006).

Table 6. D-efficiency comparison of robust one-stage design versus the proposed multivariate sequential procedure, for 16- and 48-run designs, with a 4-factor logistic regression model

	16 observations		48 observations	
	One-stage	Sequential	One-stage	Sequential
Median D-efficiency	.42	.64	.42	.83
5% quantile D-efficiency	.26	.46	.31	.73

6. CONCLUSIONS

We have suggested an algorithm for efficient sequential design of experiments with GLM responses. This algorithm has a number of valuable features. First, it is immediately applicable to multifactor experiments, whereas most of the design algorithms for GLMs have been limited to the one-factor case or to augmenting designs. Second, it provides for efficient initial designs, with a full specification of how to generate design points from the outset of the experiment. Third, the method is robust to uncertainty about the parameter values and even to uncertainty about which terms are needed in the linear predictor. No other method combines all of these advantages.

Our ideas are related to those of Chaudhuri and Mykland (1993), who provided a sequential augmentation scheme but no systematic approach for finding initial designs. For large samples, our posterior distribution should be highly concentrated around the MLEs of the model parameters, and our algorithm then will behave just like the algorithm of Chaudhuri and Mykland. We thus conjecture that our algorithm will have the same asymptotic optimality properties as that of Chaudhuri and Mykland. But we emphasize that our focus is not on the asymptotics, but rather on getting efficient designs with small samples. Our simulation results show that we succeed in this goal for one-factor experiments, relative to existing methods, and that we achieve clear improvement over one-stage designs for multifactor experiments.

Our method takes advantage of prior information through an explicitly Bayesian approach. We also exploit a computationally cheap approximation to the posterior distribution through a discrete point set. For the examples that we have examined, we found that 10,000 points gives an effective summary of the posterior. However, the number of points should increase with both the number of parameters in the model and the projected sample size. Further research may be useful in guiding the choice of the size of this point set.

[Received July 2006. Revised August 2007.]

Table 7. D-efficiencies for the group-sequential procedure versus one-stage and fully sequential designs, for 16- and 48-run designs, with a 4-factor logistic regression model

	16 observations			48 observations		
	One-stage	Group-sequential	Sequential	One-stage	Group-sequential	Sequential
Median D-efficiency	.42	.42	.64	.42	.73	.83
5% quantile D-efficiency	.26	.26	.46	.31	.62	.73

REFERENCES

Abdelbasit, K. M., and Plackett, R. L. (1983), "Experimental Designs for Binary Data," *Journal of the American Statistical Association*, 78, 90–98.

Biedermann, S., Dette, H., and Zhu, W. (2006), "Optimal Designs for Dose-Response Models With Restricted Design Spaces," *Journal of the American Statistical Association*, 101, 747–759.

Chaudhuri, P., and Mykland, P. A. (1993), "Nonlinear Experiments: Optimal Design and Inference Based on Likelihood," *Journal of the American Statistical Association*, 88, 538–546.

Chaloner, K., and Larntz, K. (1989), "Optimal Bayesian Design Applied to Logistic Regression Experiments," *Journal of Statistical Planning and Inference*, 21, 191–208.

Chaloner, K., and Verdinelli, I. (1995), "Bayesian Experimental Design: A Review," *Statistical Science*, 10, 273–304.

Dixon, J. W., and Mood, A. M. (1948), "A Method for Obtaining and Analyzing Sensitivity Data," *Journal of the American Statistical Association*, 43, 109–126.

Dror, H. A., and Steinberg, D. M. (2006), "Robust Experimental Design for Multivariate Generalized Linear Models," *Technometrics*, 48, 520–529.

Federov, V. V. (1972), *Theory of Optimal Experiments*, New York: Academic Press.

Haines, L. M., Perevozskaya, I., and Rosenberger, W. F. (2003), "Bayesian-Optimal Designs for Phase I Clinical Trials," *Biometrics*, 59, 591–600.

Ivanova, A., and Wang, K. (2004), "A Non-Parametric Approach to the Design and Analysis of Two-Dimensional Dose-Finding Trials," *Statistics in Medicine*, 23, 1861–1870.

Joseph, V. R. (2004), "Efficient Robbins–Monro Procedure for Binary Data," *Biometrika*, 91, 461–570.

Joseph, V. R., Tian, Y., and Wu, C. F. J. (2007), "Adaptive Designs for Stochastic Root-Finding," *Statistica Sinica*, 17, 1549–1565.

Karvanen, J., Vartiainen, J. J., Timofeev, A., and Pekola, J. (2007), "Experimental Designs for Binary Data in Switching Measurements on Superconducting Josephson Junctions," *Applied Statistics*, 56, 167–181.

Khuri, A., Mukherjee, B., Sinha, B., and Ghosh, M. (2006), "Design Issues for Generalized Linear Models: A Review," *Statistical Science*, 21, 376–399.

Langlie, H. J. (1965), "A Reliability Test Method for 'One-Shot' Items," Technical Report U-1792, Ford Motor Company, Aeronutronic Division, Newport Beach, CA.

Neyer, B. (1994), "A D-Optimality–Based Sensitivity Test," *Technometrics*, 36, 61–70.

Niederreiter, H. (1988), "Low-Discrepancy and Low-Dispersion Sequences," *Journal of Number Theory*, 30, 51–70.

Robbins, H., and Monro, S. (1951), "A Stochastic Approximation Method," *Annals of Mathematical Statistics*, 29, 400–407.

Silvapulle, M. J. (1981), "On the Existence of Maximum Likelihood Estimators for the Binomial Response Model," *Journal of the Royal Statistical Society*, Ser. B, 43, 310–313.

Sinha, S., and Wiens, D. P. (2002), "Robust Sequential Designs for Nonlinear Regression," *Canadian Journal of Statistics*, 30, 601–618.

Sitter, R. R., and Wu, C. F. J. (1999), "Two-Stage Design of Quanta1 Response Studies," *Biometrics*, 55, 396–402.

Woods, D. C., Lewis, S. M., Eccleston, J. A., and Russell, K. G. (2006), "Designs for Generalized Linear Models With Several Variables and Model Uncertainty," *Technometrics*, 48, 284–292.

Wu, C. F. J. (1985), "Efficient Sequential Designs With Binary Data," *Journal of the American Statistical Association*, 80, 974–984.

——— (1986), "Maximum Likelihood Recursion and Stochastic Approximation in Sequential Designs," in *Adaptive Statistical Procedures and Related Topics*, Vol. 8, ed. J. Van Ryzin, IMS Monograph Series, pp. 298–314.