Approximate Local D-optimal Experimental Design for Binary Response

Technical Report RP-SOR-0501

Hovav A. Dror and David M. Steinberg

Department of Statistics and Operations Research

Raymond and Beverly Sackler Faculty of Exact Sciences

Tel Aviv University

Ramat Aviv 69978

Israel

Email: dms@post.tau.ac.il

November 2005

Abstract: A fast and simple method is proposed that produces approximate multivariate local D-optimal designs of high efficiency for models with binary response. The method assumes availability of a D-optimal design for a parallel normal response linear problem that has the same linear predictor, with an assumption of homogenous variance; the change required to transform the standard design into an efficient one for a multivariate logit or probit model is to shift any design point whose probability is very low or very high (and is therefore non-informative) into the nearest feasible point of moderate probability.

KEY WORDS: Generalized Linear Models ; Logistic Response; Logit; Probit; Design of Experiments

1. INTRODUCTION

Construction of an experimental design for a generalized linear model presents a level of complexity greater than that required for a model with a normally distributed error term of constant variance. Finding an optimal design for a linear model is already a numerically intensive optimization problem. However, some linear models can be characterized and have a known trivial optimal design, and designs for more complicated models can be sought through available software such as "gosset" (Hardin and Sloane 1993), the statistical toolbox in MATLAB (The Mathworks, inc), JMP or the SAS Optex procedure. Extension from a linear model to a logistic one does not retain these qualities. Trivial solutions to linear model design problems are often factorial designs, utilizing the corners of the design region; for a logisitic response, some of these corners have probabilities that are close to zero or one, so that the responses are almost deterministic, and are therefore non-informative. Furthermore, optimal design for binary response depends on the unknown coefficients, and hence two experiments having the same model but different coefficient values will typically have different optimal designs. Common software packages assume all observations have homogenous variance and so do not provide a remedy for generalized linear models.

This paper proposes a simple method for constructing an approximate local D-optimal design for a multivariate binary response problem with the logit or probit link.

2. EXPERIMENTAL DESIGN FOR GENERALIZED LINEAR MODELS

Most work on Experimental Design focuses on linear models with a continuous response. A common assumption is that the error term of the model is normally distributed and that its variance is constant over the design region. These assumptions are not met (even asymptotically) with the popular logistic model or with the probit model for binary response data. Standard procedure for the analysis of these models uses iteratively reweighted least squares, which asymptotically leads to weights that enable solving the problem as a simple linear model (McCullagh and Nelder 1989). These weights depend on the parameters of the problem and are generally different for two binary response problems of the same structure but different true (unknown) coefficients.

Frequently, Experimental Design uses optimality criteria based on Fisher's Information Matrix. Of these, D-optimality is the most intensively studied. For the normal response setting with $\mathbf{y} = \boldsymbol{\eta} + \boldsymbol{\varepsilon}$ and $\boldsymbol{\eta} = F\boldsymbol{\beta}$, D-optimality maximizes the determinant of the information matrix F^TF and so minimizes the volume of the confidence ellipsoid for the unknown coefficients $\boldsymbol{\beta}$. Atkinson & Donev (1992) note that the assumptions of normality and constancy of variance enter the optimal design through the information matrix. They indicate that instead of F^TF, other matrices are appropriate for non-normal distributions, and that once the appropriate matrix has been defined the principles and practice of optimal experimental design are similar to continuous response problems. This means that the optimal experimental design is driven from the information matrix of the parameters F^TWF; the matrix of weights, W, depends on the unknown parameters and on the design. This dependence creates a difficulty in constructing an optimal experimental design.

2.1 Local D-Optimal Designs

Because a D-optimal design depends on the coefficients, its derivation require us to make some assumptions on the values of the unknowns we want to estimate. When we use an initial estimate (rather than a Bayesian a-priori distribution, for instance) we produce a design that is optimal only locally, for the coefficient values that were adopted. If the initial estimate is poor, the design's performance may be far from optimal. D-optimal designs that are based on an initial estimate for the unknown coefficient values are designated as local D-optimal designs.

2.2 Local D-Optimal Designs for a logistic model

Abdelbasit & Plackett (1983) discuss the construction of local D-optimal designs for binary response with one explanatory variable. They show that for an unconstrained univariate problem with the model $logit(p) = \beta_0 + \beta_1 x$ the local D-optimal design places half the points where the estimated probability is 0.824 and the other half at the location with probability 0.176. Univariate binary response has been further considered by many authors, see for example Atkinson & Donev (1992), Sitter (1992), Hedayat, Yan & Pezzuto (1997), or Mathew & Sinha (2001).

Although the final coordinates creating the optimal design are found numerically, the function to maximize is derived through analytical procedures. Increasing the number of unknowns complicates both the analytical route and the numerical process. This is probably the reason not much has been published on multivariate design problems. Chipman and Welch (1996) compare D-optimal designs for Generalized Linear Models over a constrained region to linear regression D-optimal designs, including multifactor problems. Their comparison was based on computer generated D-optimal designs, using weighted linear models. This way they show the general phenomenon that points in the linear model optimal design may "move in" from the edges of the design region for the logistic model. Sitter and Torsney (1995) analyze D- and c- optimal designs for binary response experiments with two design variables and various link functions; Smith and Ridout (2003) obtain optimal Bayesian designs for bioassays involving two parallel dose response relationships where the main interest is in doses of one substance; Atkinson, Demetrio & Zocchi (1995) describe a dose response experiment with two variables - one continuous and one indicator variable, considering the special case in which the value of the indicator variable is unknown during the process of the experiment, but is available for the posterior analysis. Atkinson (2005) gives examples of first order D-optimal designs with two variables and discusses usage of standard factorial design as an approximation, similar to the method proposed here; but - assuming one needs to search for the approximate design over irregular design regions Atkinson concludes that "Searching over the original design space to find optimal design for the generalized linear model would both be easier and lead to a more efficient design than would trying to find such a regression approximation".

It is noted that the treatment described so far was limited to a first order model with two variables. Woods, Lewis, Eccleston and Russel (2005) offer a method of creating multivariate compromise designs that are robust to uncertainty in aspects of the model, including the uncertainty in the coefficient values, different choices of a link function and various models, including interactions. As can be expected for such complex designs, their method requires intensive computation.

We will now suggest a simple method of constructing approximate local D-optimal designs. This method is not limited to a small number of covariates, or to first-order designs without interactions. It can produce approximate designs almost instantaneously, as it does not require intensive computation, and it is straightforward to implement. There are two cornerstones for this algorithm; first, one needs an optimal design for an analogous linear model (that is, the optimal design for a problem with the same formulation and a constant variance). Second, we use computation of the probability at design points according to the specific initial coefficient values.

3. REGION OF LINEARITY

We are considering a binary response with the logistic link. That is $p_i = \frac{e^{F_i\boldsymbol{\beta}}}{1+e^{F_i\boldsymbol{\beta}}}$, F_i being the i-th row of the regression matrix F. We attempt to find an approximation for a local optimal design. As explained, the term "local" implies that the design relates to particular values of the coefficients.

Note that while both $\boldsymbol{\beta}$ and feasible points for the design space may be of any dimension, the probability depends on their scalar product $F_i\boldsymbol{\beta}$. Figure 1 displays this probability $p(\boldsymbol{\beta}, \mathbf{x})$ versus $\boldsymbol{\beta}'\mathbf{x}$, where $\mathbf{x}$ is any possible row of the regression matrix F:

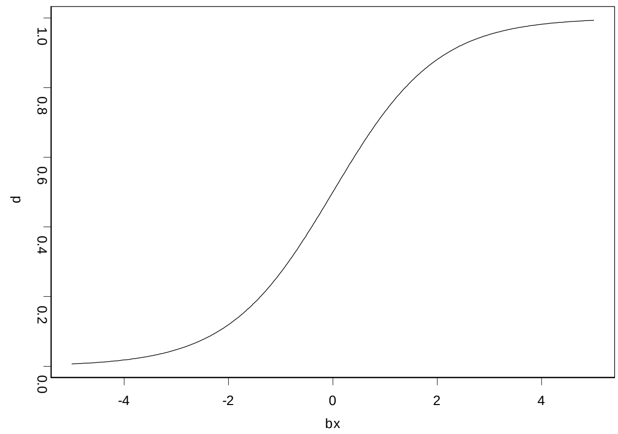

Figure 1: The probability $p(\boldsymbol{\beta}, \mathbf{x})$ versus the product $\boldsymbol{\beta}'\mathbf{x}$, for the Logit link

We see that when $0.2 \leq p \leq 0.8$ the probability p is approximately a linear function of $\boldsymbol{\beta}'\mathbf{x}$. Even in the larger range of $0.1 \leq p \leq 0.9$ it is not far from being linear.

Therefore it is tempting to assume that the local optimal design would be similar to a regular linear response experimental design, for models in which the design region does not contain points with extreme probabilities (very high or very low). Still, having a linear link function is not sufficient to accept that assumption since the error term is not distributed with a constant variance σ^2. It is known that $var\left(\widehat{\boldsymbol{\beta}}\right) \propto (F'WF)^{-1}$, and for any row F_i representing a single observation, $var(\widehat{\eta}_i) \propto F_i'(F'WF)^{-1} F_i$. Luckily, for the logit link the variance (which is clearly proportional to the reciprocal of the weights) is nearly constant within the range suggested above, while its magnitude increases rapidly outside that region (note that we discuss the variance of the estimated linear predictor which is different from $p_i(1-p_i)$, the variance of the observations). This is illustrated in Figure 2.

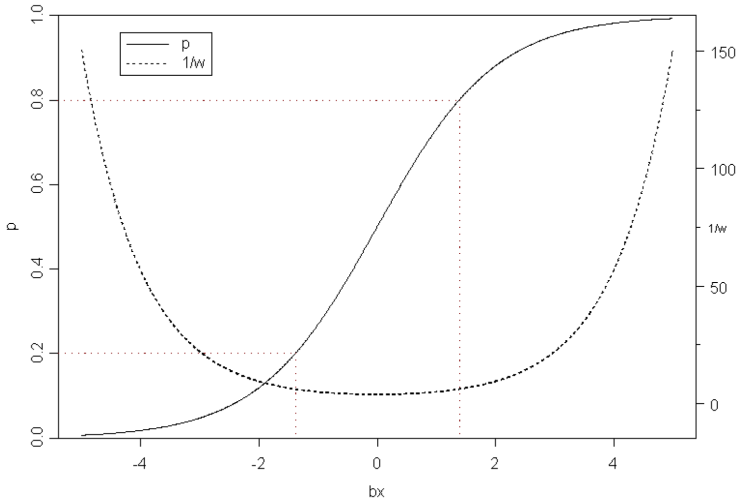

Figure 2: For $0.2 < p < 0.8$ the variance is approximately constant, increasing rapidly outside those limits

Therefore, if our design region is limited to points with an expected probability above 0.1-0.2 and less than

0.8-0.9 (based on the initial parameter values), we can expect the binary response local optimal design to be close to a normally distributed response optimal design.

4. APPROXIMATE LOCAL D-OPTIMAL DESIGN

Avoiding design points with extreme probabilities should be an effective strategy for any optimality criterion. We can further develop this understanding for some specific choices, such as the D-criterion. Basically, it enables us to approximate the D-optimal design even when the design region includes points with probabilities outside the region where $0.2 \leq p \leq 0.8$.

The D-optimality criterion selects experimental points that minimize $\det\left(var\left(\widehat{\boldsymbol{\beta}}\right)\right)$. As noted $var\left(\widehat{\boldsymbol{\beta}}\right) \propto (F'WF)^{-1}$ with W being the weights matrix, a diagonal matrix which for a logit link has values of $w_{ii} = p_i\left(1-p_i\right) = \frac{e^{\mathbf{F}_i\boldsymbol{\beta}}}{\left(1+e^{\mathbf{F}_i\boldsymbol{\beta}}\right)^2}$. Clearly the use of experimental points for which w_{ii} is very small would be unwise. Outside the region where $0.2 \leq p \leq 0.8$ the weights decrease very fast; observe Figure 2 again (noting that the presented curve is $1/w_{ii}$, not w_{ii}). For that reason we should expect that points with very small or very high probabilities will not be included in a D-optimal design. As a result, **an easy way to produce a good approximation to a local D-optimal solution would be to exclude from the design region all points with very low or high probabilities and find a normal response D-optimal design over the constrained region.**

As explained, "very low probabilities" are probabilities under ~0.1-0.2 and "very high" are over ~0.8-0.9. As a rule of thumb we can cut at $p < 0.15$ and $p > 0.85$. Denoting $c \equiv \ln\left(\frac{0.85}{1-0.85}\right) = -\ln\left(\frac{0.15}{1-0.15}\right) = 1.7346$, we add the constraints $-c\mathbf{1} \leq \mathbf{F}\boldsymbol{\beta} \leq c\mathbf{1}$, which are applied using the initial guess for $\boldsymbol{\beta}$.

Finding a D-optimal design for the constrained normal response model can be done with many available algorithms. But for many common models **constructing an approximate D-optimal design can be made even simpler, releasing the user from the need for complicated algorithms**. It is well known that if we are considering a normal response problem constrained to the region $[-1,1]^m$ for a model consisting of first-order terms and all possible interactions, then the D-optimal design is straightforward: it is a full factorial experiment placing an equal number of points at each of the 2^m corners of the region. If all these points are of moderate probability, then this should be the local D-optimal design for a binary response as well. Otherwise we should shift any corner point of extreme probability to the nearest point of the closer surface $\mathbf{F}\boldsymbol{\beta} = \pm\, c\mathbf{1}$.

If the model contains up to 3 variables shifting can be done graphically by drawing the design region and plotting the surfaces of equi-probability 0.15 and 0.85 inside it. For higher dimensions graphic tools are usually not suitable, but the problem of finding the nearest point to a corner on a surface is a familiar optimization task. It is easy to use sequential quadratic programming in order to find it. Many commercial products can perform the optimization process, for instance the function 'fmincon' in MATLAB.

Remark 1 *A familiar result from normal response theory is that the support points of a linear model optimal design are located on the boundary of the design region. Unlike this result, optimal designs for binary response typically contain points in the interior of the design region.*

Remark 2 *For first-order models with no interactions, the familiar result prevails. As can be seen in Sitter & Torsney (1995) and Atkinson (2005) it is possible to present a first-order binary model as an equivalent linear model by replacing the design region with an induced region. In such cases more efficient designs will result if we add a constraint, requiring all points being placed on the boundary of the design region.*

Remark 3 *Possibly there will be more than one point on the boundary that will minimize the distance from an examined corner while achieving the required probability; in such cases one needs to split up the original point into several design points, each of which gets an equal fraction of the weight assigned to the original point* ($1/2^m$).

Remark 4 *Note that although we assume throughout this paper that the model coefficients are known, in practice, most people will find it much easier to specify areas of low and high probabilities than coefficient values. Hence, the algorithm proposed gives the user an easy, fast and intuitive way to exploit process knowledge in deriving approximate local D-optimal design.*

Remark 5 *Figures 1 and 2 presented the values appropriate for the logistic link. The probit link behaves similarly, and the results apply equally for both link functions. The use of the approximation suggested is less suitable for the asymmetric Complementary Log-Log link.*

Some examples follow.

5. EXAMPLES

All the examples compare the approximate design to a local D-optimal design. The efficiency is calculated as the ratio $\left(\left|F_A' W_A F_A\right| / \left|F_O' W_O F_O\right|\right)^{1/p}$ where p is the number of unknown coefficients, and W_A, F_A, W_O, F_O are the weights and design matrices for the approximate and optimal designs, respectfully. Optimal designs were found by numerous methods including a variation on Wynn's algorithm (1972), a variation on the implementation of Federov's (1972) Exchange Algorithm in MATLAB, and a simulated annealing algorithm implemented by Woods, Lewis, Eccleston and Russell (2005).

5.1 Example 1

The first example is a simple case of a symmetrical first order model with two variables of identical influence. The model used is: $\text{logit}(p) = \beta_0 + \beta_x x + \beta_y y$ with $\boldsymbol{\beta} = (0, 2, 2)'$. Throughout the paper we will assume all variables are scaled to the region $[-1, 1]$.

This model could represent a machine that receives data signals at several different frequencies. The amplitude (x and y) of the signal transmitted at each frequency may be determined independently. The integrated effect of the signal may cause the machine to be unstable and as a result to stop production. The goal of the experiment is to model the probability of a production stoppage.

Figure 3 shows the local D-optimal design for this model. The six marked points designate the design; dotted lines are equi-probability contours, with their probability marked as a number on them, the bold rational number near each design point notes its frequency of appearance (its weight) in the design, and stars mark the approximate design constructed as described above.

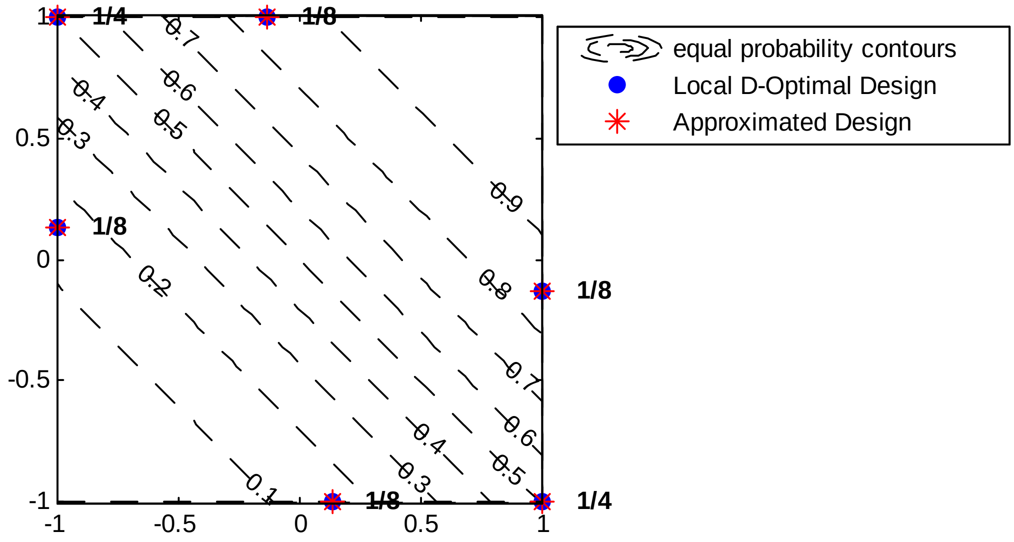

Figure 3: Local D-optimal Design and Approximate design for a first-order model

It is seen that the approximation is almost identical to the true local D-optimal design. Two of the original design points ($[-1, 1]$ and $[1, -1]$) have $p = 0.5$ and are not moved. The other two factorial design points ($[-1, -1]$ and $[1, 1]$) are shifted until the required probability (0.15 / 0.85) is reached; since each of these points has two points on the boundary with the required probability, and both are of equal distance from the corner - these points are split in the design, and each subpoint has a design weight of $\frac{1}{8}$ instead of the original $\frac{1}{4}$. Avoiding the splitting by choosing only one equi-distant point does not harm the efficiency of the design. If $\beta_x \neq \beta_y$ then no splitting is required. This is demonstrated in Figure 4 for $\boldsymbol{\beta} = (0, 1, 2)'$

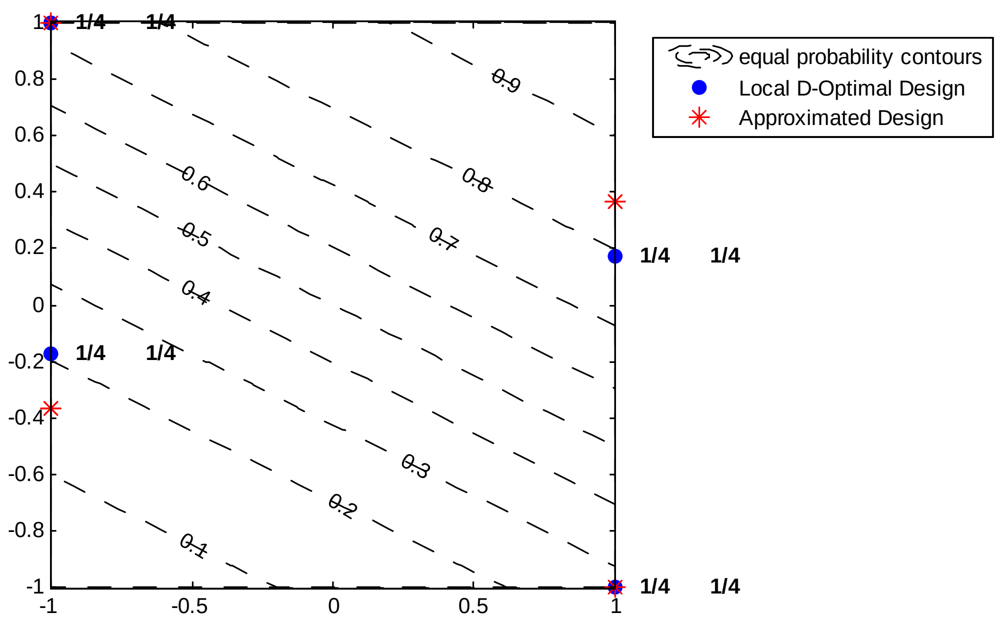

Figure 4: Local D-optimal Design and Approximate design for a first order model with $\beta_x \neq \beta_y$

5.2 Example 1a

A variation of the first model would be to add an interaction with an assumed coefficient value of zero. That is: $\text{logit}(p) = \beta_0 + \beta_x x + \beta_y y + \beta_{xy} xy$ with $\boldsymbol{\beta} = (0, 2, 2, 0)'$. The expected response values, expressed by the equi-probability contours, are the same for both models. But, as noted in remarks 1 and 2, once we consider the possibility for an interaction effect in the model, the optimal design is no longer constrained to the boundary of the design region. Figure 5 shows the local D-optimal design for this model, together with the approximate design.

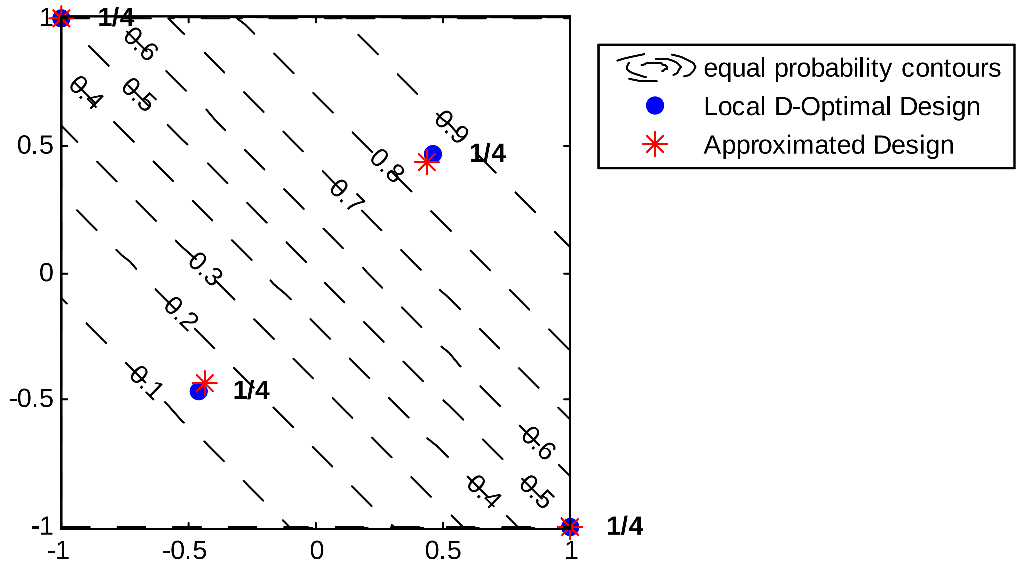

Figure 5: Approximate vs Local D-optimal Design for the same model

It is seen that, unlike the D-optimal design for a normal response or a first order logistic model, in our binary model which contains an interaction, two corner points $([-1, -1], [1, 1])$ "drifted" inside. Once the design is not constrained to be on the boundary, the nearest point to the corner $[-1, -1]$ $([1, 1])$ with probability 15% (85%) is found in the interior of the design space. There is no need to split points and the optimal design has only 4 points. Notice the difference between this model and the first model discussed. Suppose we believe that the two variables (e.g. amplitude at each frequency) are of equal importance and that an interaction is not likely, but might be present. The last model better reflects this by including an interaction term in the model, but setting its coefficient value to zero.

For all three models, the approximate design's efficiency is over 98%, which is to be expected as they are almost identical to the true optimal designs. For comparison, the regular full factorial design achieves efficiency of 78% for the first example of a model with no interaction, 81% for the second example, and 73% in the last example, including an interaction with an anticipated coefficient value of 0.

5.3 Example 2

The same conclusions apply for more complex models as well. The next example uses an asymmetric effect of two variables, together with a strong interaction.

We use the same model: $\text{logit}(p) = \beta_0 + \beta_x x + \beta_y y + \beta_{xy} xy$; this time with $\boldsymbol{\beta} = (0, 1, 2, 3)'$. This could be the case, for example, if one frequency is more dominant than the other, and if there is a synergetic effect when both frequencies are applied.

The results are displayed in Figure 6, maintaining the same notations as before.

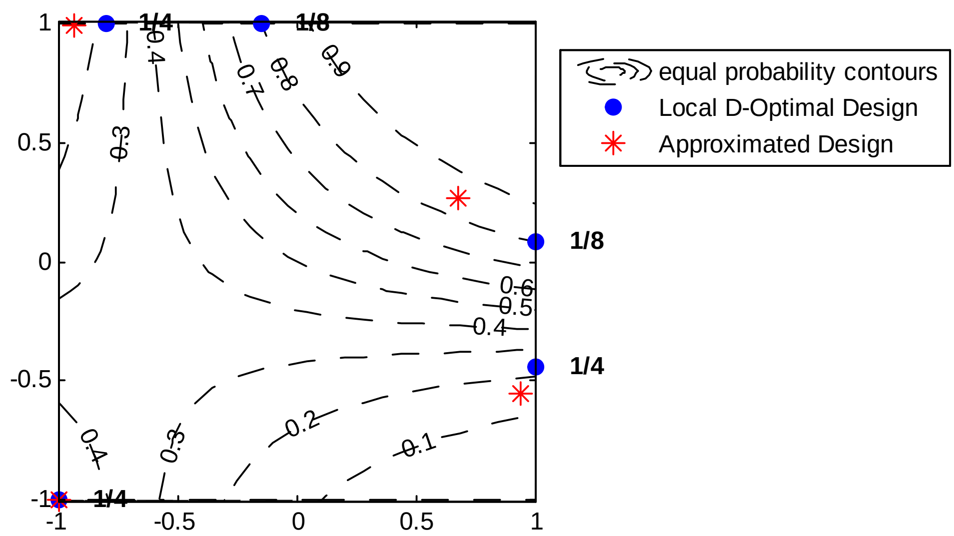

Figure 6: Approximate and Local D-optimal Designs for a 2 variable model with a strong interaction effect

The efficiency of the approximate design was 89% this time. This is still a high value, but smaller than before. A regular factorial design has an efficiency of only 35%. Note that the upper right point of the factorial design not only shifted inward in the local D-optimal design, but was also split into two points with $\frac{1}{8}$ weight - half of the original one-quarter weight of the corner point. Had we sought an approximation, not by shifting the points directly, but instead by finding the D-optimal design for a normal response model constrained to the region $-c\mathbf{1} \leq \mathbf{F}\boldsymbol{\beta} \leq c\mathbf{1}$, then the upper right design point would split as well, and the approximation efficiency would rise over 95%. The drawback to the latter method is the loss of simplicity. An interim method might be to add an algorithm step, after shifting the points inward, of trying to split points when the extreme points on the surface $\mathbf{F}\boldsymbol{\beta} = c\mathbf{1}$ are approximately the same distance from the original corner as the point of minimum distance. The decision whether to split points does not require running a D-optimality computer algorithm, but may still be a complex problem. Hence, it may be preferred to use the slightly less efficient design in order to retain simplicity.

5.4 Example 3

We now further expand the complexity and give an example with three factors.

The model is extended to: $\text{logit}(p) = \beta_0 + \beta_x x + \beta_y y + \beta_z z + \beta_{xy} xy + \beta_{xz} xz + \beta_{yz} yz + \beta_{xyz} xyz$ with the coefficients chosen to be $\boldsymbol{\beta} = (0, 2, 2, 2, 0, 0, 0, 0)'$. Adding interactions with a coefficient value of 0 ensures a design that can examine the hypothesis of no interaction effect.

Figures 7a & 7b and Table 1 compare the local D-optimal design to the approximate design. The two figures differ only in the view angle.

It can be seen that the points comprising the optimal design are all on or near the shaded surfaces of probabilities 20% or 80%, with two points (the ones that would be at $(-1,-1,-1)$ and $(1,1,1)$ in a normal response unconstrained design) approaching the centers of the shaded areas, and the rest of the points placed on the corners of the polygons describing each shaded area.

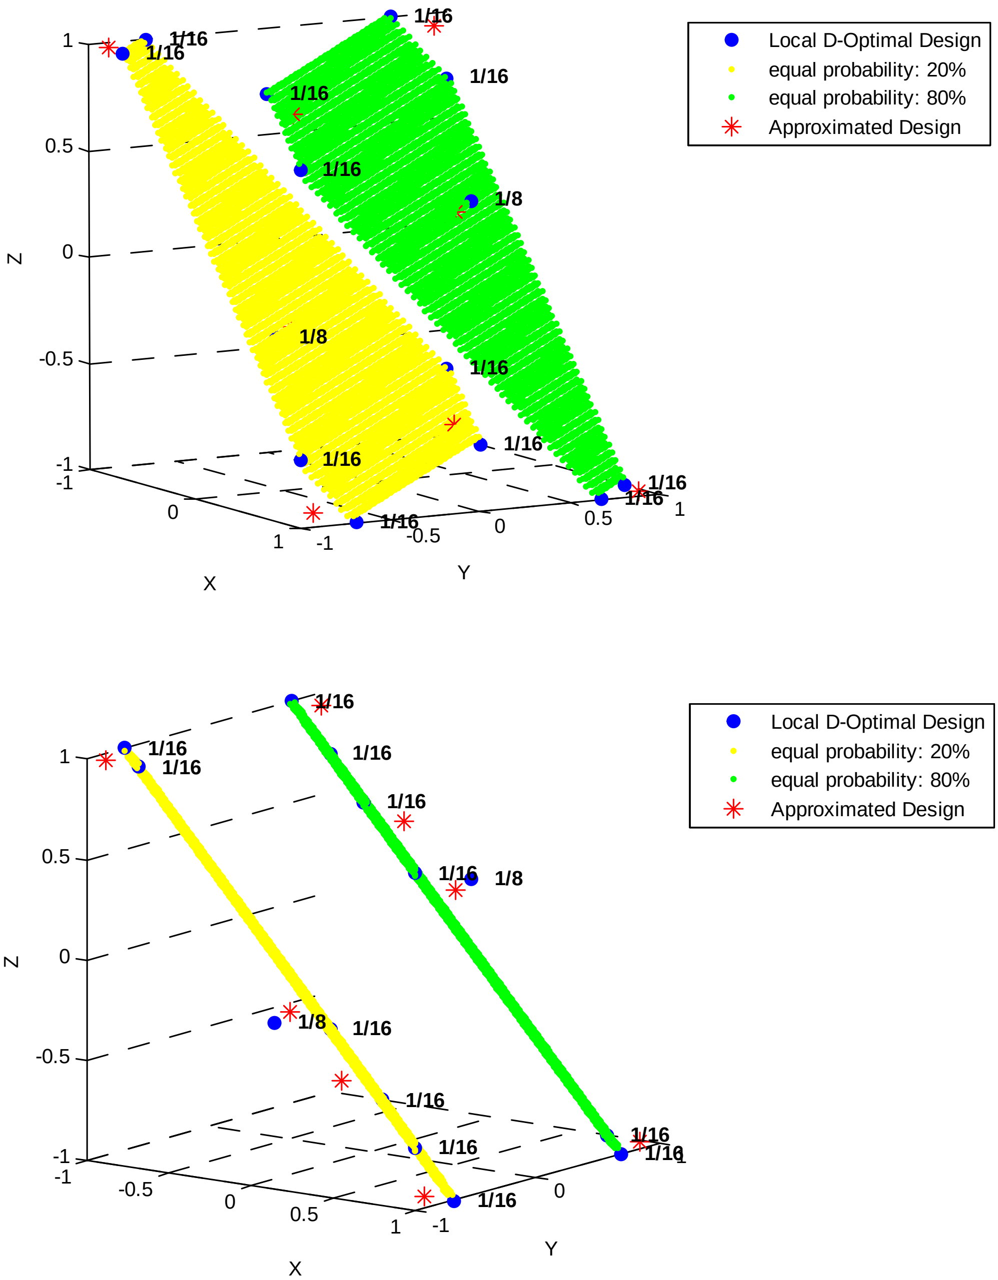

Figure 7a & 7b: Approximate and Local D-optimal Designs for a 3 variable model

Table 1: Approximate and Local D-optimal Designs for a 3 variable model

Optimal Design				Approximate Design			
x	y	z	weight	x	y	z	weight
-1	-0.684	1	1/16	-1	0.93	0.94	1/8
-1	0.684	1	1/16	-0.93	-0.93	0.99	1/8
-1	1	-0.684	1/16	-0.93	1	-0.94	1/8
-1	1	0.684	1/16	-0.29	-0.29	-0.29	1/8
-0.684	-1	1	1/16	0.28	0.29	0.29	1/8
-0.684	1	-1	1/16	0.93	-1	0.94	1/8
-0.344	-0.344	-0.344	1/8	0.93	0.93	0.99	1/8
0.344	0.344	0.344	1/8	1	-0.93	-0.94	1/8
0.684	-1	1	1/16				
0.684	1	-1	1/16				
1	-1	-0.684	1/16				
1	-1	0.684	1/16				
1	-0.684	-1	1/16				
1	0.684	-1	1/16				

Once again we see that the approximate design places its points near the optimal locations, but most original points from the full factorial are split into two in the local D-optimal design. Even with its lack of splitting, the efficiency of the approximate design is above 95%. For comparison, the efficiency of a full factorial design is slightly over 65%.

5.5 Example 4

We now extend to an asymmetric 3 factor problem.

The model is still: $\text{logit}(p) = \beta_0 + \beta_x x + \beta_y y + \beta_z z + \beta_{xy} xy + \beta_{xz} xz + \beta_{yz} yz + \beta_{xyz} xyz$, but this time with $\boldsymbol{\beta} = (1, 2, 3, 4, 5, 6, 0, 0)'$

Even though the model is quite complex, all the qualities shown before are fully preserved, and the approximate design achieved an efficiency of 80% in comparison to the local optimal design. For contrast, the efficiency of a regular factorial design is only 1.5%.

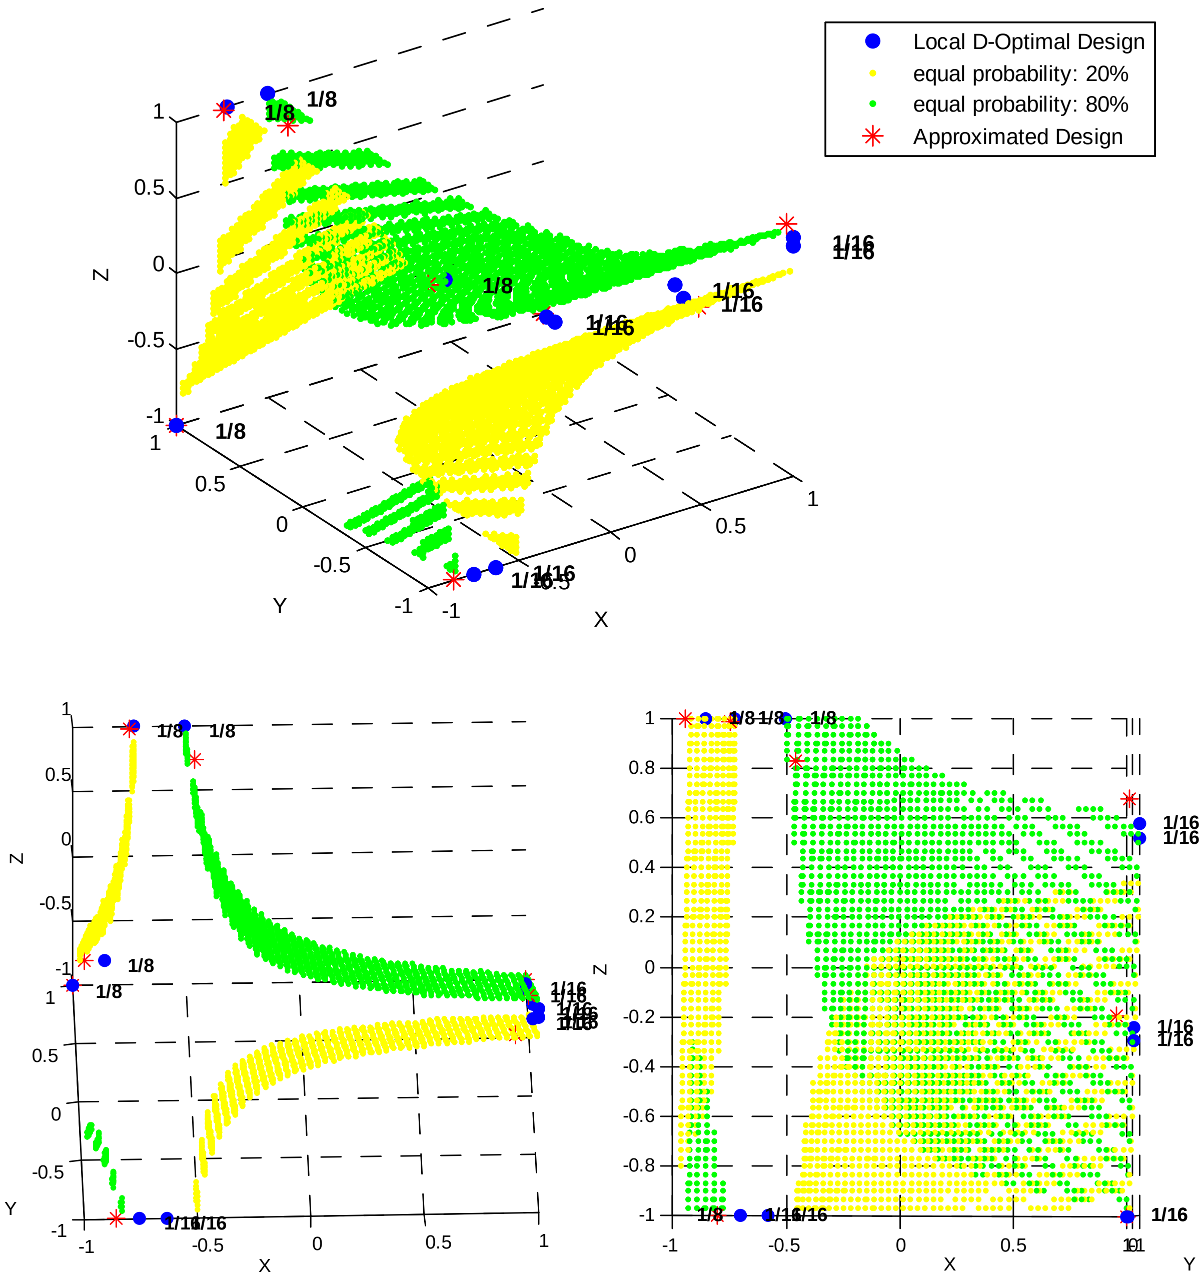

Figure 8a,b,c : Asymmetric 3 dimensions model with $\boldsymbol{\beta} = (1, 2, 3, 4, 5, 6, 0, 0)'$

Table 2: Asymmetric 3 dimensions model with $\boldsymbol{\beta} = (1, 2, 3, 4, 5, 6, 0, 0)$

Optimal Design				Approximate Design			
x	y	z	weight	x	y	z	weight
-1	1	-1	1/8	-1	-1	1	1/8
-0.91	-1	1	1/8	-1	1	-1	1/8
-0.75	-1	-1	1/16	-0.86	-1	-1	1/8
-0.73	1	1	1/8	-0.75	0.99	0.98	1/8
-0.63	-1	-1	1/16	-0.46	0.90	0.83	1/8
-0.50	1	1	1/8	0.92	-0.36	-0.19	1/8
1	-1	0.52	1/16	0.96	-1	0.68	1/8
1	-1	0.58	1/16	1	1	-1	1/8
1	-0.11	-0.30	1/16				
1	-0.05	-0.24	1/16				
1	0.91	-1	1/16				
1	0.97	-1	1/16				

5.6 Example 5

We continue with the same model, and a different choice of coefficients, making it a bit messier: $\boldsymbol{\beta} = (1, 2, 3, 4, 3, 1, 1, 1)'$.

Again, while a regular factorial design has poor efficiency (less than 15%) the approximate design reaches very high efficiency (over 90%) and the complexity of the model did not harm the qualities shown before:

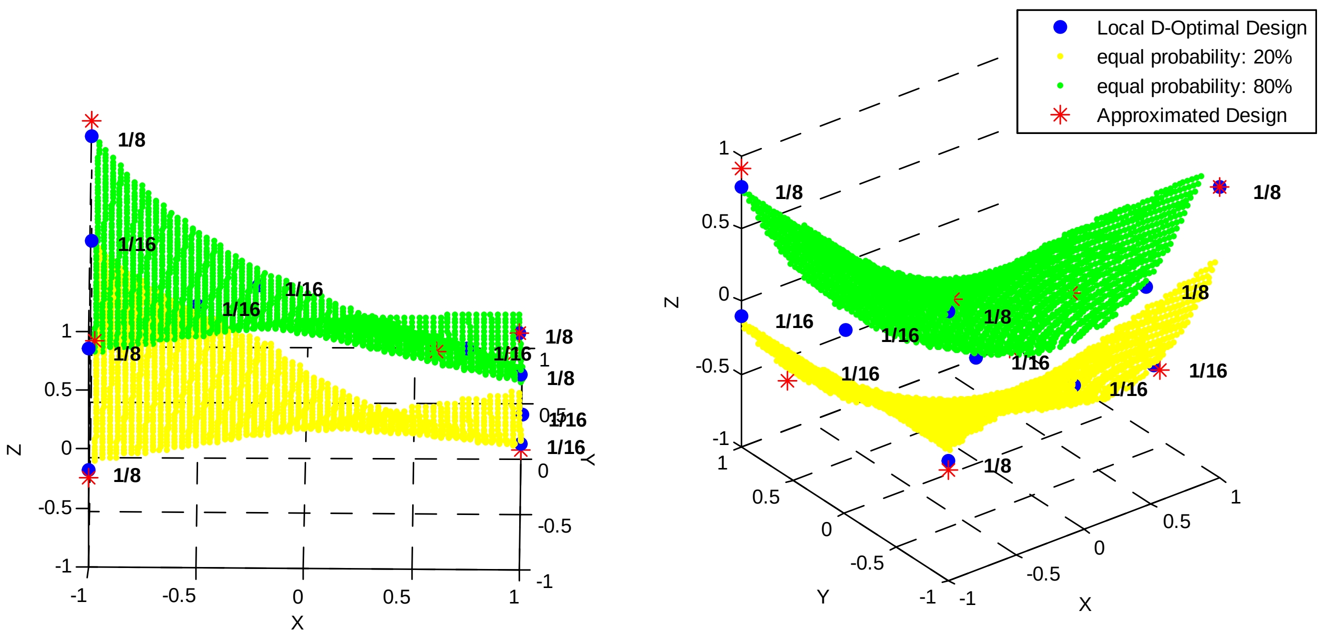

Figure 9a,b : Asymmetric 3 factor model with $\boldsymbol{\beta} = (1, 2, 3, 4, 3, 1, 1, 1)'$

5.7 Example 6

Last, we look at an extension of example 3. This time we include 6 variables, with all possible interactions. Assume a value of 5 for all main effect coefficients, and a value of 0 for all interactions and the intercept. This may represent an experiment to assess the effectiveness as a function of different amplitude choices for 6 different frequencies, under the assumption that all frequencies are equally effective, with the chance of the system becoming unstable determined entirely by the sum of the amplitude values. Note that the experiment is aimed at assessing 64 (!) coefficient values. Figure 10 compares the approximate design to the local optimal design, showing all possible two-dimensional projections.

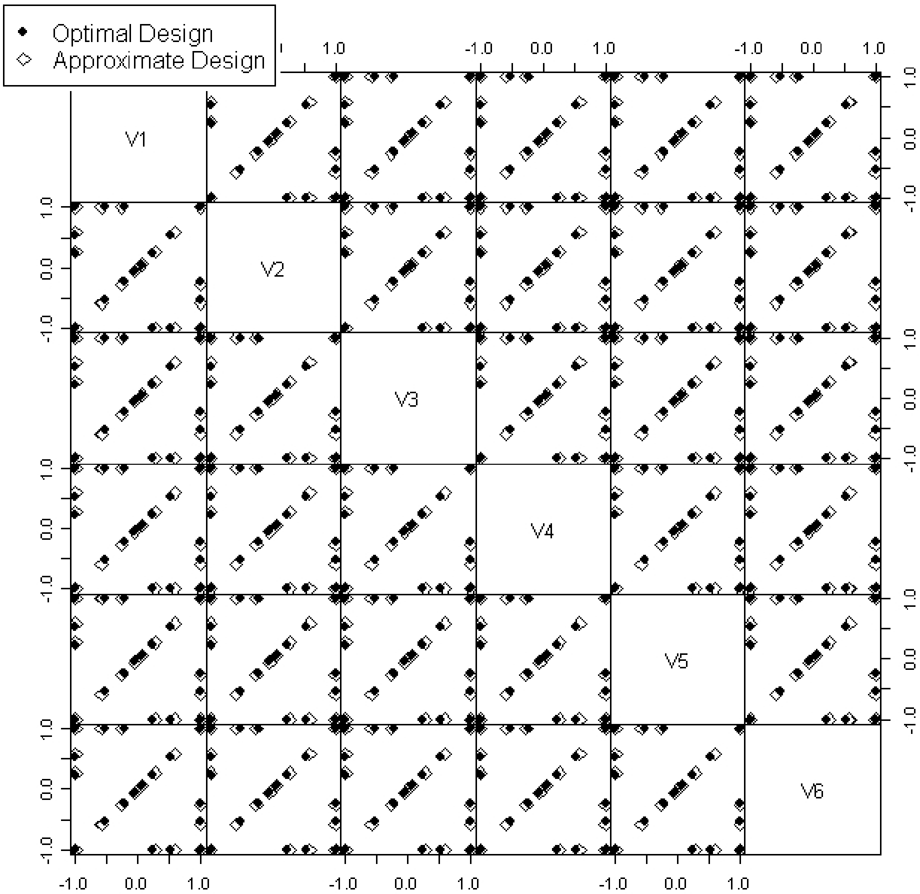

Figure 10: Extension of example 3 into an experiment aimed at assessing 6 factors and 64 coefficients

It is seen that once again the approximate design is visually similar to the optimal design. The approximate design has the obvious advantage of being constructed in negligible time, compared to the highly intensive task of finding an optimal design when the regression matrix has 64 columns. The efficiency of the approximate design is only 55% this time. This value is lower than achieved in all previous examples, but is still relatively high, especially considering the large number of coefficients being estimated. The efficiency of a regular factorial design is 0.2%, meaning it requires more than 250 times as many observations to achieve the same D-criterion value as the approximate design. Note that location-wise the differences between the approximate and optimal designs are almost negligible.

6. DESIGNS BASED ON FRACTIONAL FACTORIALS

If the model includes only main effects, then for normal response linear models it is sufficient to use a fractional factorial design. Each resolution III fraction is as efficient as the full factorial. For non-linear models this result is not valid. Consider a 5 factor main effects logistic model with $\beta = (1, 2, 3, 4, 5, 6)'$. The full factorial design consists of 32 runs, which can be decomposed into four fractional factorial designs of resolution III with 8 runs each. Figure 11 compares the efficiency achieved by an approximate design based on the full factorial (70%, designated as a line) to the efficiency of approximate designs based on the four fractional factorials (77%, 58%, 56%, 72%), designated as bars.

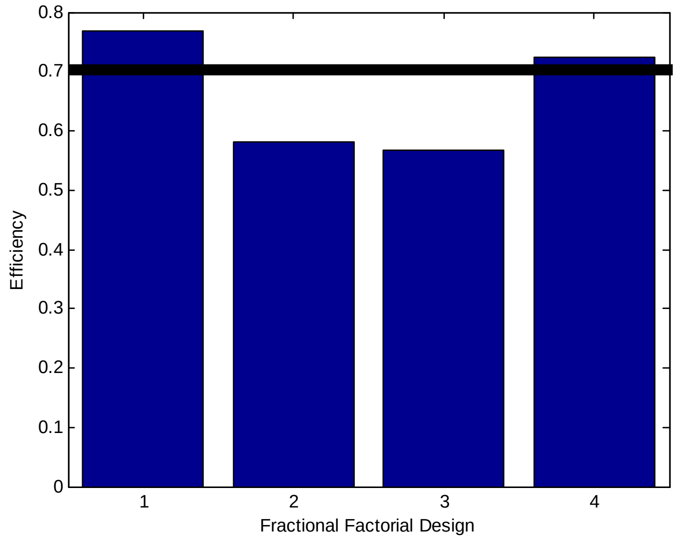

Figure 11: Efficiency of Approximate Designs based on a full factorial (line) and fractional factorials (bars)

The efficiencies of the approximate designs based on the fractional factorials are different from each other and from the full factorial based approximate design. Both best and worst efficiencies are achieved for fractional factorial based approximations. Often the best designs are based on the fraction with the smallest number of points having probability of extreme value. In the last example the first fraction contains only 5 points that have extreme probability, the last fraction has 6 such points, and in the middle two fractions, with the worst efficiency, all 8 points are with probabilities close to 0 or 1.

It is therefore advised, if resources permit, to compare the relative efficiency of all possible approximate designs based on fractional factorials, or at least those fractions with few points of extreme probability.

7. LIMITATIONS

Being a heuristic, the approximation algorithm may sometimes lead to inefficient results. A severe example is a 2 variable model with the coefficients being $\boldsymbol{\beta} = (2, 2, 2, -4)'$, as plotted in Figure 12.

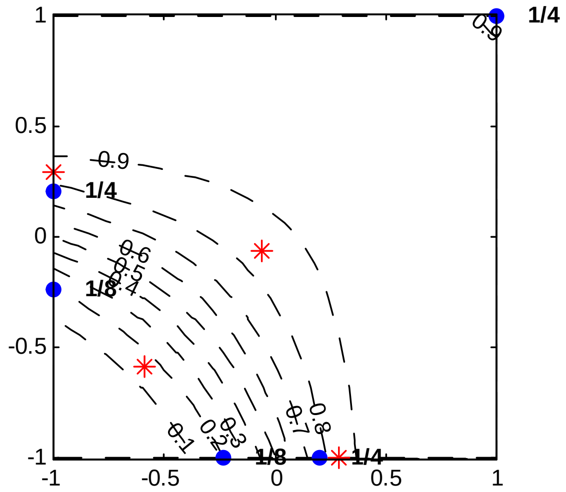

Figure 12: An inefficient approximation

In this example the strong interaction creates a very flat response in the north-east part of the plot, beginning to effectively decrease near the boundary of the design region. The result is a significant distance between the optimal design point $[1, 1]$ and its approximation parallel, which decreases the efficiency to less than 50%.

The method proposed is of limited suitability to models of second or higher order. When several design points for a normal response model are located in an extreme probability region and are close enough to each other to be united in the approximate design, the approximation may not only be inefficient, but even insufficient for the estimation of all required coefficients.

8. CONCLUSION

Multivariate local D-optimal designs for logistic models with main effects and interactions can be approximated without the need for intensive computation. Extensions would be to efficient approximate designs that take account of the uncertainty in the coefficient values, to other design criteria such as I-optimality, and to higher order models.

REFERENCES

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

2. Atkinson, A.C. (2005), "Generalized linear models and response transformation," *Response Surface Methodology and Related Topics*, ed. A. I. Khuri. Singapore: World Scientific. (In press).

3. Atkinson, A. C., Demetrio, C. G. B., Zocchi, S. S. (1995), "Optimum dose levels when males and females differ in response," *Applied Statistics*, 44, 213-226.

4. Atkinson, A.C. and Donev, A. N. (1992), *Optimum Experimental Designs*, Oxford University Press, Oxford.

5. Chipman, H. and Welch, W. (1996), "D-optimal Design for Generalized Linear Models," unpublished manuscript.

6. Federov, V.V. (1972), *Theory of optimal experiments.* New York: Academic Press.

7. Ford, I., Torsney B. and Wu C.F.J. (1992), "The Use of a Canonical Form in the Construction of Locally Optimal Designs for Non-linear Problems," *Journal of the Royal Statistical Society*, Ser. B, 54, 569-583.

8. Hardin, R. H. and Sloane, N. J. A.(1993), "A New Approach to the Construction of Optimal Designs," *Journal of Statistical Planning and Inference*, 37 339-369.

9. Hedayat, A. S., Yan, B. and Pezzuto, J. M. (1997), "Modeling and identifying optimum designs for fitting dose-response curves based on raw optical density data," *Journal of the American Statistical Association* 92, 1132-1140.

10. Mathew, T. and Sinha, B.K. (2001), "Optimal designs for binary data under logistic regression," *Journal of Statistical Planning and Inference* 93, 295-307.

11. McCullagh, P. and J.A. Nelder. (1989), *Generalized Linear Models*, Chapman and Hall: London.

12. Sitter, R.R. (1992), "Robust Designs for Binary Data," *Biometrics*, 48, 1145-1156.

13. Sitter, R.R. and Torsney, B. (1995), "Optimal Designs for Binary Response Experiments with Two Design Variables". *Statistica Sinica*, 5, 405-419.

14. Smith, D.M. and Ridout, M.S. (2003), "Optimal designs for criteria involving log(potency) in comparative bioassays," *Journal of Statistical Planning and Inference*, 113(2), 617-632.

15. Woods, D.C., Lewis, S.M., Eccleston, J.A. and Russell, K.G. (2005), "Designs for generalized linear models with several variables and model uncertainty," *Technometrics*, in press.

16. Wynn, H.P. (1972), "Results in Theory and Construction of D-Optimum Experimental Designs," *Journal of the Royal Statistical Society*, Ser. B, 34, 133-147.