

Penalized and response-adaptive optimal designs. Application to dose finding

Luc Pronzato

Laboratoire I3S, CNRS/Université de Nice-Sophia Antipolis
Bât Euclide, Les Algorithmes
2000 route des lucioles, BP 121
06903 Sophia Antipolis cedex, France
pronzato@i3s.unice.fr

Abstract. Optimal design under a cost constraint is considered, with a scalar coefficient setting the compromise between information (i.e., precision of the estimation of the model parameters) and cost. For suitable cost functions, by increasing the value of the coefficient one can force the support points of an optimal design measure to concentrate around points of minimum cost. When the experiment is constructed sequentially, the choice of each new design point being based on the current estimated value of the model parameters (response-adaptive design), the strong consistency and asymptotic normality of the estimator of the model parameters is obtained under the assumption that the design variables belong to a finite set. An example of adaptive design in a dose-finding problem with a bivariate binary model is presented, showing the effectiveness of the approach.

Keywords. Adaptive design, bivariate binary model, compound optimal design, constrained optimal design, dose finding, optimal experimental design, penalized experimental design, sequential design.

1 Penalized D -optimal design: introduction and motivation

This work is motivated by the recent papers (Dragalin and Fedorov, 2006; Dragalin et al., 2008) where the authors use constrained optimal design to make a compromise between individual and collective ethics in dose-finding studies. Their idea is to use a cost function that accounts for poor efficacy and for toxicity, and to maximize information-per-cost-unit, which can be put in the form of a standard (unconstrained) optimal design problem. Using a parametric model for the dose/efficacy-toxicity responses (Gumbel or Cox model as in (Dragalin and Fedorov, 2006) or a bivariate probit model as in (Dragalin et al., 2008)), the Fisher information matrix can be calculated and optimal designs can be constructed.

In the present paper we introduce some flexibility in setting the balance between the information gained (in terms of precision of parameter estimation) and the cost of the experiment (in terms of poor success for the patients enrolled in the trial) by maximizing information-per-observation under a constraint on the cost or, equivalently, by optimizing a penalized design criterion where the penalty is related to the cost of the experiment.

Let \mathcal{X} , a compact subset of \mathbb{R}^d , denote the admissible domain for the experimental variables x (design points) and $\theta \in \mathbb{R}^p$ denote the (p -dimensional) vector of parameter of interest in a parametric model generating the log-likelihood $l(Y, x; \theta)$ for the observation Y at the design point x . We suppose that $\theta \in \Theta$, a compact subset of \mathbb{R}^p . For N independent observations $\mathbf{Y} = (Y_1, \dots, Y_N)$ at non random design points $\mathbf{X} = (x_1, \dots, x_N)$ the log-likelihood at θ is $l(\mathbf{Y}, \mathbf{X}; \theta) = \sum_{i=1}^N l(Y_i, x_i; \theta)$. Let $\mathbf{M}(X, \theta)$ denote the corresponding Fisher information matrix, $\mathbf{M}(X, \theta) = -\mathbb{E}_\theta \{ \partial^2 l(\mathbf{Y}, \mathbf{X}; \theta) / (\partial \theta \partial \theta^\top) \} = \sum_{i=1}^N \mu(x_i, \theta)$. When $N(x_i)$ denotes the number of observations made at $x = x_i$, we get the following normalized information matrix per observation $\mathbf{M}(\xi, \theta) = (1/N) \mathbf{M}(X, \theta) = \sum_{i=1}^K [N(x_i)/N] \mu(x_i, \theta)$, where K is the number of distinct design points and ξ is the design measure (a probability measure on \mathcal{X}) that puts mass $N(x_i)/N$ at x_i . Following the usual approximate design approach, we shall relax the constraints on design measures and consider ξ as any element of Ξ , the set of probability measures on \mathcal{X} , so that $\mathbf{M}(\xi, \theta) = \int_{\mathcal{X}} \mu(x, \theta) \xi(dx)$.

In a regression model with independent and homoscedastic observations satisfying $\mathbb{E}_\theta(Y|x, \theta) = \eta(x, \theta)$, with $\eta(x, \theta)$ differentiable with respect to θ for any x , we have

$$\mu(x, \theta) = \mathcal{I} \frac{\partial \eta(x, \theta)}{\partial \theta} \frac{\partial \eta(x, \theta)}{\partial \theta^\top} \quad (1)$$

with $\mathcal{I} = \int [\varphi'(t)/\varphi(t)]^2 \varphi(t) dt$ the Fisher information for location, where $\varphi(\cdot)$ is the probability density function of the observation errors and $\varphi'(\cdot)$ its derivative.

In a dose-response problem with single response $Y \in \{0, 1\}$ (efficacy or toxicity response at the dose x for instance) and $\text{Prob}\{Y = 1|x, \theta\} = \pi(x, \theta)$ we have $l(Y, x; \theta) = Y \log[\pi(x, \theta)] + (1 - Y) \log[1 - \pi(x, \theta)]$ so that, assuming $\pi(x, \theta)$ differentiable with respect to θ for any x ,

$$\mu(x, \theta) = \frac{\partial \pi(x, \theta)}{\partial \theta} \frac{\partial \pi(x, \theta)}{\partial \theta^\top} \frac{1}{\pi(x, \theta)[1 - \pi(x, \theta)]}.$$

Bivariate extensions, where both efficacy and toxicity responses are observed at a dose x , are considered in Dragalin and Fedorov (2006) (Gumbel and Cox models) and Dragalin et al. (2008) (bivariate probit model). See also the example in Sect. 2. Besides a few additional technical difficulties, the main difference with the single response case is the fact that $\mu(x, \theta)$ may have rank larger than one, so that less than p support points in ξ may suffice to estimate θ consistently. The same situation occurs for regression models when $\dim(\eta) > 1$ so that (1) may have rank larger than one. We assume that $\mu(x, \theta)$ is bounded on \mathcal{X} .

Local D -optimal design consists in determining a measure ξ_D^* that maximizes $\log \det[\mathbf{M}(\xi, \theta)]$, with $\mathbf{M}(\xi, \theta)$ the Fisher information matrix at a given value of θ . In many circumstances, besides the optimality criterion $\log \det[\mathbf{M}(\xi, \theta)]$, it is desirable to introduce a constraint of the form $\Phi(\xi, \theta) \leq C$ for the design measure. In dose-finding problems, the introduction of such a constraint allows one to take individual ethical concerns into account. For instance, when both the efficacy and toxicity responses are observed, one can relate $\Phi(\xi, \theta)$ to the probability of success (efficacy and no toxicity) for a given dose, as done in Dragalin and Fedorov (2006); Dragalin et al. (2008). See also Sect. 2. We suppose that the cost (or penalty) function $\Phi(\xi, \theta)$ is linear in ξ , that is

$$\Phi(\xi, \theta) = \int_{\mathcal{X}} \phi(x, \theta) \xi(dx),$$

and that $\phi(x, \theta)$ is bounded on \mathcal{X} (see, e.g., Cook and Fedorov (1995) and Fedorov and Hackl (1997, Chap. 4) for extensions to nonlinear constraints). Also, we restrict our attention to the case where a single (scalar) constraint is present.

The fact that a single cost function is present permits to consider the problem of maximizing information per cost-unit, which can be formulated as a design problem without constraint, see (Dragalin and Fedorov, 2006; Dragalin et al., 2008). However, in dose-finding problems this formulation has the important consequence that the prohibition of excessively low (with poor efficacy) or high (with high toxicity) doses can only be obtained by an ad-hoc modification of the cost function $\phi(x, \theta)$. Indeed, this is the only way to modify the optimal design and hopefully to change its support. This can be contrasted with the solution of the constrained design problem that we consider in the present paper.

A direct formulation of the optimal design problem under constraint is as follows:

$$\text{Maximize } \log \det[\mathbf{M}(\xi, \theta)] \text{ with respect to } \xi \in \Xi \text{ under the constraint } \Phi(\xi, \theta) \leq C. \quad (2)$$

We say that a design measure $\xi \in \Xi$ is θ -admissible if $\Phi(\xi, \theta) \leq C$ and we suppose that a strictly θ -admissible measure exists in Ξ ($\Phi(\xi, \theta) < C$ for some $\xi \in \Xi$). A necessary and sufficient condition for the optimality of a θ -admissible $\xi^* \in \Xi$ for (2) then becomes:

$$\exists \lambda^* \geq 0 \text{ such that } \begin{cases} \lambda^* [C - \Phi(\xi^*, \theta)] = 0 \\ \text{and} \\ \forall x \in \mathcal{X}, \text{ trace}[\mu(x, \theta) \mathbf{M}^{-1}(\xi^*, \theta)] \leq p + \lambda^* [\phi(x, \theta) - \Phi(\xi^*, \theta)]. \end{cases} \quad (3)$$

In practice, ξ^* can be determined by maximizing

$$H_\theta(\xi, \lambda) = \log \det[\mathbf{M}(\xi, \theta)] - \lambda \Phi(\xi, \theta) \quad (4)$$

for an increasing sequence (λ_i) of Lagrange coefficients λ , starting at $\lambda_0 = 0$ and stopping at the first λ_i such that the associated optimal design ξ^* satisfies $\Phi(\xi^*, \theta) \leq C$, see, e.g., Mikulecká (1983) (for

C large enough, the unconstrained D -optimal design ξ_D^* is optimal for the constrained problem). The parameter λ can thus be used to set the tradeoff between the maximization of $\log \det[\mathbf{M}(\xi, \theta)]$ (gaining information) and minimization of $\Phi(\xi, \theta)$ (reducing cost). Notice that maximizing $H_\theta(\xi, \lambda)$ for $\lambda \geq 0$ is equivalent to maximizing $(1 - \gamma) \log \det[\mathbf{M}(\xi, \theta)] + \gamma [-\Phi(\xi, \theta)]$ with $\gamma = \lambda / (1 + \lambda) \in [0, 1)$ (one may refer to Cook and Wong (1994) for the equivalence between constrained and compound optimal designs). Similarly to the case of D -optimal design, the optimal matrix $\mathbf{M}(\xi^*, \theta)$ is unique (but the optimal design measure ξ^* is not necessarily unique).

Let $\xi^*(\lambda)$ denote an optimal design for $H_\theta(\xi, \lambda)$ given by (4). One can easily check that both $\log \det\{\mathbf{M}[\xi^*(\lambda), \theta]\}$ and $\Phi[\xi^*(\lambda), \theta]$ are non-increasing functions of λ , see Cook and Wong (1994) for examples. We suppose that $\mu(x, \theta)$ and $\phi(x, \theta)$ are continuous in $x \in \mathcal{X}$, with \mathcal{X} a compact subset of \mathbb{R}^d , and define

$$\phi_\theta^* = \min_{x \in \mathcal{X}} \phi(x, \theta), \quad x^* = x^*(\theta) = \arg \min_{x \in \mathcal{X}} \phi(x, \theta). \quad (5)$$

One can then show that, for suitable penalty functions, the support of an optimal design for (2) depends on C or, equivalently, the support of an optimal design for (4) depends on λ . When x^* is unique, one may then obtain that the supporting points of ξ^* converge to x^* as $\lambda \rightarrow \infty$. For dose-response problems, this property has the important consequence that excessively high or low doses can be prohibited by choosing C small enough or, equivalently, λ large enough. Its effectiveness very much depends on the choice of the penalty function, and in particular on its local behavior around x^* : (contrary to what intuition might suggest, it requires the cost function $\phi(\cdot, \theta)$ to be sufficiently flat around x^* : indeed, in that case a design ξ supported in the neighborhood of x^* can at the same time have a small cost $\Phi(\xi, \theta)$ and be dispersed enough to carry significant information through $\log \det \mathbf{M}(\xi, \theta)$).

2 Example: Cox model for efficacy-toxicity response

The example is taken from (Dragalin and Fedorov, 2006) and concerns a problem with bivariate binary responses. For Y (respectively Z) the binary indicator of efficiency (resp. of toxicity) at dose x for a model with parameters θ , we write $\text{Prob}\{Y = y, Z = z | x, \theta\} = \pi_{yz}(x, \theta)$, $Y, y, Z, z \in \{0, 1\}$, with

$$\begin{aligned} \pi_{11}(x, \theta) &= \frac{e^{a_{11} + b_{11} x}}{1 + e^{a_{01} + b_{01} x} + e^{a_{10} + b_{10} x} + e^{a_{11} + b_{11} x}} \\ \pi_{10}(x, \theta) &= \frac{e^{a_{10} + b_{10} x}}{1 + e^{a_{01} + b_{01} x} + e^{a_{10} + b_{10} x} + e^{a_{11} + b_{11} x}} \\ \pi_{01}(x, \theta) &= \frac{e^{a_{01} + b_{01} x}}{1 + e^{a_{01} + b_{01} x} + e^{a_{10} + b_{10} x} + e^{a_{11} + b_{11} x}} \\ \pi_{00}(x, \theta) &= \left(1 + e^{a_{01} + b_{01} x} + e^{a_{10} + b_{10} x} + e^{a_{11} + b_{11} x}\right)^{-1} \end{aligned}$$

and $\theta = (a_{11}, b_{11}, a_{10}, b_{10}, a_{01}, b_{01})^\top$. The log-likelihood function of a single observation (Y, Z) at dose x is then $l(Y, Z, x; \theta) = YZ \log \pi_{11}(x, \theta) + Y(1 - Z) \log \pi_{10}(x, \theta) + (1 - Y)Z \log \pi_{01}(x, \theta) + (1 - Y)(1 - Z) \log \pi_{00}(x, \theta)$ and elementary calculations show that the contribution to the Fisher information matrix is

$$\mu(x, \theta) = \frac{\partial \mathbf{p}^\top(x, \theta)}{\partial \theta} \left(\mathbf{P}^{-1}(x, \theta) + [1 - \pi_{11}(x, \theta) - \pi_{10}(x, \theta) - \pi_{01}(x, \theta)]^{-1} \mathbf{1} \mathbf{1}^\top \right) \frac{\partial \mathbf{p}(x, \theta)}{\partial \theta^\top}$$

where $\mathbf{p}(x, \theta) = [\pi_{11}(x, \theta), \pi_{10}(x, \theta), \pi_{01}(x, \theta)]^\top$, $\mathbf{P}(x, \theta) = \text{diag}\{\mathbf{p}(x, \theta)\}$ and $\mathbf{1} = (1, 1, 1)^\top$. Note that $\mu(x, \theta)$ is generally of rank 3. As in Dragalin and Fedorov (2006), we take $\theta = (3, 3, 4, 2, 0, 1)^\top$ and \mathcal{X} the finite set $\{x^{(1)}, \dots, x^{(11)}\}$ where the doses $x^{(i)}$ are equally spaced in the interval $[-3, 3]$. The D -optimal design is supported on $x^{(1)}, x^{(4)}, x^{(5)}$ and $x^{(10)}$, with associated weights 0.3318, 0.3721, 0.1259 and 0.1701.

We first choose a cost function related to the probability $\pi_{10}(x, \theta)$ of efficacy and no toxicity and take $\phi(x, \theta) = \pi_{10}^{-1}(x, \theta)$. The Optimal Safe Dose (OSD), minimizing $\phi(x, \theta)$, is $x^{(5)} = -0.6$. Figure 1




Fig. 1. Optimal designs $\xi^*(\lambda)$ as function of $\lambda \in [0, 100]$ for the cost function $\pi_{10}^{-1}(x, \theta)$: each horizontal dotted line corresponds to a point in \mathcal{X} , the thickness of the plot indicates the associated weight.

presents the optimal designs $\xi^*(\lambda)$ for λ varying between 0 and 100 along the horizontal axis. The weight associated with each $x^{(i)}$ on the vertical axis is proportional to the thickness of the plot.

Consider now the cost function

$$\phi(x, \theta) = \{\pi_{10}^{-1}(x, \theta) - [\max_x \pi_{10}(x, \theta)]^{-1}\}^2 \tag{6}$$

which is more flat than $\pi_{10}^{-1}(x, \theta)$ around its minimum (at the OSD $x^{(5)}$). One can show that the optimal designs then concentrate on three doses around the OSD when λ is large enough. Figure 2 indicates that for $\lambda \gtrsim 75$ the optimum designs are supported on $x^{(4)}$ and $x^{(6)}$ only, with weights approximately 1/2 each, that is, all patients in a trial defined by $\xi^*(\lambda)$ receive a dose close to the optimal one, $x^{(5)}$. Note, however, that none receives the OSD (compare with Figure 2). The situation changes for larger values of λ , and numerical calculations show that the optimal design is supported on $\{x^{(4)}, x^{(5)}, x^{(6)}\}$ for $\lambda \gtrsim 160$, with the weight of the optimal dose $x^{(5)}$ increasing with λ .




Fig. 2. Same as Figure 1, but for the cost-function (6).

3 Adaptive designs

In a nonlinear situation, like in the example in the section above, $M(\xi, \theta)$ and $\Phi(\xi, \theta)$ usually depend on θ . A common approach to overcome the issue of dependence of the optimum design in θ consists in

designing the experiment sequentially. In adaptive D -optimal design for instance, the design point after N observations is taken as

$$x_{N+1} = \arg \max_{x \in \mathcal{X}} \text{trace}[\mu(x, \hat{\theta}^N) \mathbf{M}^{-1}(\xi_N, \hat{\theta}^N)], \quad (7)$$

where $\hat{\theta}^N$ is the current estimated value of θ . By alternating between estimation based on previous observations and determination of the next design point where to observe, one forces the empirical design measure to progressively adapt to the correct (true) value of the model parameters. Adaptive design is considered in (Dragalin and Fedorov, 2006; Dragalin et al., 2008), but the convergence of the procedure (strong consistency of the parameter estimator and convergence of the empirical design measure to the optimal non-sequential design for the true value of the model parameters) is left as an open issue. The difficulty of proving the consistency of the estimator when design variables are sequentially determined is usually overcome by considering an initial experiment (non adaptive) that grows in size when the total number of observations increases, see, e.g., Chaudhuri and Mykland (1993). Although this number is often severely limited in practise, especially for clinical trials, we think that it is reassuring to know that, *for a given initial experiment*, adaptive design guarantees suitable asymptotic properties under reasonable conditions. Using simple arguments, one can show that this is indeed the case when *the design space is finite*, which forms a rather natural assumption in the context of clinical trials. The case of adaptive D -optimal design is considered in (Pronzato, 2009b) (notice that it also covers the situation considered by Dragalin and Fedorov (2006); Dragalin et al. (2008), which can be formulated as a standard D -optimal design problem).

In the case of adaptive penalized D -optimal design, the design point after N observations is taken as

$$x_{N+1} = \arg \max_{x \in \mathcal{X}} \left\{ \text{trace}[\mu(x, \hat{\theta}^N) \mathbf{M}^{-1}(\xi_N, \hat{\theta}^N)] - \lambda_N \phi(x, \hat{\theta}^N) \right\}. \quad (8)$$

Following an approach similar to that in (Pronzato, 2009b), one can show that when \mathcal{X} is finite, λ_N is the optimal Lagrange coefficient for (2) with the estimated value $\hat{\theta}^N$ substituted for θ , and under standard regularity assumptions, this procedure is asymptotically “optimal” in the sense that the estimated value of the parameters (by least-squares in a nonlinear regression model or by the maximum-likelihood in Bernoulli trials) converges a.s. to its true value $\bar{\theta}$ and the information matrix tends a.s. to the penalized D -optimal matrix at $\bar{\theta}$, see Pronzato (2009a). Also, the estimator is asymptotically normal, with variance-covariance matrix given by the inverse of the usual information matrix, similarly to the non-adaptive case.

The strong consistency of $\hat{\theta}^N$ is preserved when λ_N is taken as a control parameter that tends to infinity not too fast (more slowly than $N/\log \log N$). Letting λ_N tend to infinity permit to focus more and more on cost minimization and to obtain design measures that converge weakly to the delta measure at $x^* = \arg \min_{x \in \mathcal{X}} \phi(x, \bar{\theta})$ (and all design points tend to concentrate around x^* for suitable penalty functions). In dose-finding problems, it means that for suitable penalty functions, when the weight given to the cost for bad treatment increases with the number of patients enrolled, the doses allocated converge to the OSD while the parameters are still estimated consistently.

Numerical simulations with the example of Section 2 (trials on 36 patients with the cost function $\phi(x, \theta) = \pi_{10}^{-1}(x, \theta)$) indicate much better performance in terms of information gained (precision of the estimation of θ in the trial) for (8) with λ_N adapted to $\hat{\theta}^N$ than for the up-and-down rule of Ivanova (2003), defined by

$$x_{N+1} = \begin{cases} \max\{x^{(i_N-1)}, x^{(1)}\} & \text{if } Z_N = 1, \\ x^{(i_N)} & \text{if } Y_N = 1 \text{ and } Z_N = 0, \\ \min\{x^{(i_N+1)}, x^{(11)}\} & \text{if } Y_N = 0 \text{ and } Z_N = 0. \end{cases} \quad (9)$$

Here, the index $i_N \in \{1, \dots, 11\}$ is defined by $x^{(i_N)} = x_N$ and (Y_N, Z_N) denotes the observation for x_N . This up and down rule is also considered by Dragalin and Fedorov (2006) (see also Kpamegan and Flournoy (2001, p. 221)). Simulations also indicate that the number of times the OSD is estimated

correctly after the trial is much larger for (8) with λ_N adapted to $\hat{\theta}^N$ than for (9). This performance comes with a prize, and the cost $\Phi(\xi, \theta)$ is (slightly) higher for (8) than for (9). Considering longer trials (240 patients) with λ_N increasing with N in (8) permits to outperform (9) both in terms of precision of the estimation of θ (and location of the OSD) and cost: as λ_N increases, the design points generated by (8) tend to concentrate around the OSD.

4 Conclusions

The approach used in (Dragalin and Fedorov, 2006; Dragalin et al., 2008) makes a clear compromise between the efficient treatment of individuals in the trial (by preventing the use of doses with low efficacy or high toxicity) and the precise estimation of the model parameters (accompanied with measures of statistical accuracy), to be used for making efficient decisions for future treatments. As such, it has a great potential in combining traditional phase I and phase II clinical trials into a single one, thereby accelerating the drug development process.

We have shown that a different formulation of the problem permits to introduce some flexibility in setting the compromise between the information gained (in terms of precision of parameter estimation) and the cost of the experiment (in terms of poor success for the patients enrolled in the trial). We have shown in particular that, for suitable penalty functions, by increasing the weight set on the penalty one guarantees that all doses in the experiment have a small cost (and concentrate around the optimal dose when this one is unique). This permits the avoidance of extreme doses generally suggested by optimal design for parameter estimation. Further developments and numerical studies are required to define suitable rules for selecting cost functions and for choosing the value (or the sequence of values) for the penalty coefficients λ_N .

Bibliography

- Chaudhuri, P., Mykland, P., 1993. Nonlinear experiments: optimal design and inference based likelihood. *Journal of the American Statistical Association* 88 (422), 538–546.
- Cook, D., Fedorov, V., 1995. Constrained optimization of experimental design (invited discussion paper). *Statistics* 26, 129–178.
- Cook, D., Wong, W., 1994. On the equivalence between constrained and compound optimal designs. *Journal of the American Statistical Association* 89 (426), 687–692.
- Dragalin, V., Fedorov, V., 2006. Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 136, 1800–1823.
- Dragalin, V., Fedorov, V., Wu, Y., 2008. Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 138, 352–373.
- Fedorov, V., Hackl, P., 1997. *Model-Oriented Design of Experiments*. Springer, Berlin.
- Ivanova, A., 2003. A new dose-finding design for bivariate outcomes. *Biometrics* 59, 1001–1007.
- Kpamegan, E., Flournoy, N., 2001. An optimizing up-and-down design. In: Atkinson, A., Bogacka, B., Zhigljavsky, A. (Eds.), *Optimum Design 2000*. Kluwer, Dordrecht, Ch. 19, pp. 211–224.
- Mikulecká, J., 1983. On a hybrid experimental design. *Kybernetika* 19 (1), 1–14.
- Prinzato, L., 2009a. Asymptotic properties of nonlinear least squares estimates in stochastic regression models over a finite design space. Application to self-tuning optimisation. In: *Proc. 15th IFAC Symposium on System Identification*, Saint-Malo, France.
- Prinzato, L., 2009b. One-step ahead adaptive D -optimal design on a finite design space is asymptotically optimal. *Metrika* (to appear).