

 Open access • Journal Article • DOI:10.1007/BF02509828

Sequential experimental design and response optimisation — [Source link](#)

Luc Pronzato, E. Thierry

Institutions: University of Nice Sophia Antipolis

Published on: 01 Oct 2002 - Statistical Methods and Applications (Springer-Verlag)

Related papers:

- [Strongly consistent estimators of k-th order regression curves and rates of convergence](#)
- [Weak convergence of the weighted sequential empirical process of some long-range dependent data](#)
- [Further Characterizations of Design Optimality and Admissibility for Partial Parameter Estimation in Linear Regression](#)
- [A note on comparing the means of a multivariate normal population](#)
- [Finite memory algorithms for testing bernoulli random variables](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/sequential-experimental-design-and-response-optimisation-3v1pze1xt9>



HAL
open science

Sequential experimental design and response optimisation

Luc Pronzato, Eric Thierry

► **To cite this version:**

Luc Pronzato, Eric Thierry. Sequential experimental design and response optimisation. Statistical Methods and Applications, Springer Verlag, 2003, 11, pp.277-292. hal-01002348

HAL Id: hal-01002348

<https://hal.archives-ouvertes.fr/hal-01002348>

Submitted on 6 Jun 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sequential experimental design and response optimisation

Luc Pronzato, Éric Thierry

Laboratoire I3S, CNRS/Université de Nice-Sophia Antipolis, bât. Euclide
Les Algorithmes, 2000 route des Lucioles
BP 121, 06903 Sophia-Antipolis Cedex, France
e-mail: {pronzato,et}@i3s.unice.fr

Received: date / Revised version: date

Abstract We consider the situation where one wants to maximise a function $f(\theta, \mathbf{x})$ with respect to \mathbf{x} , with θ unknown and estimated from observations y_k . This may correspond to the case of a regression model, where one observes $y_k = f(\theta, \mathbf{x}_k) + \epsilon_k$, with ϵ_k some random error, or to the Bernoulli case where $y_k \in \{0, 1\}$, with $\Pr[y_k = 1 | \theta, \mathbf{x}_k] = f(\theta, \mathbf{x}_k)$. Special attention is given to sequences given by $\mathbf{x}_{k+1} = \arg \max_{\mathbf{x}} f(\hat{\theta}^k, \mathbf{x}) + \alpha_k d_k(\mathbf{x})$, with $\hat{\theta}^k$ an estimated value of θ obtained from $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_k, y_k)$ and $d_k(\mathbf{x})$ a penalty for poor estimation. Approximately optimal rules are suggested in the linear regression case with a finite horizon, where one wants to maximize $\sum_{i=1}^N w_i f(\theta, \mathbf{x}_i)$ with $\{w_i\}$ a weighting sequence. Various examples are presented, with a comparison with a Polya urn design and an up-and-down method for a binary response problem.

Key words Adaptive control – Bernoulli trials – binary response – dose-response – optimum design – parameter estimation – response optimisation – sequential design

1 Introduction

We consider an optimisation problem, where one wants to maximise a function $f(\theta, \mathbf{x})$ with respect to $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^q$, where θ is an unknown vector of parameters, $\theta \in \mathbb{R}^p$. Two cases will be considered. The first one corresponds to a regression problem, where one observes

$$y_k = f(\theta, \mathbf{x}_k) + \epsilon_k, \quad k = 1, 2, \dots \quad (1)$$

with ϵ_i an unobservable error such that $E\{\epsilon_k | \epsilon_1, \dots, \epsilon_{k-1}\} = 0$ and $E\{\epsilon_k^2 | \epsilon_1, \dots, \epsilon_{k-1}\} < \infty$, $k = 1, \dots$ almost surely (a.s.).

In the second case, when the design variable is \mathbf{x}_k , one observes a Bernoulli response:

$$y_k \in \{0, 1\} \quad \text{with} \quad \Pr[y_k = 1 | \theta, \mathbf{x}_k] = f(\theta, \mathbf{x}_k), \quad k = 1, 2, \dots \quad (2)$$

We shall denote \mathcal{F}_k the σ -field generated by y_1, \dots, y_k , $E\{\cdot | \mathcal{F}_k\}$ the corresponding posterior expectation and $E\{\cdot | \mathcal{F}_0\}$ the prior expectation, with a prior probability measure μ for θ . In the two cases above, the sequence of design points $\{\mathbf{x}_k\}$ and observations $\{y_k\}$ is used to estimate θ . If the issue were only to determine the value \mathbf{x}^* that maximises $f(\theta, \mathbf{x})$, one could resort to optimum-design theory for choosing an appropriate sequence of inputs $\{\mathbf{x}_k\}$, see, *e.g.*, [14, 27, 3, 4, 22, 24, 9]. However, here the objective is also to have each $f(\theta, \mathbf{x}_i)$ as large as possible. The sequence $\{\mathbf{x}_k\}$ must therefore fulfill two simultaneous (and generally opposite) objectives: (i) help to locate \mathbf{x}^* , (ii) be close to \mathbf{x}^* in order to maximise $f(\theta, \mathbf{x}_k)$, $k = 1, 2, \dots$. The problem is thus one of dual control, see the pioneering papers by Fel'dbaum [10–13]. Note that the \mathbf{x}_i 's will be chosen sequentially, contrary to the case considered in [26]; that is, \mathbf{x}_i is \mathcal{F}_{i-1} -measurable, $i = 1 \dots$.

A rather general formulation of the problem (see [15]) is:

$$\text{maximize} \quad E\left\{\sum_{i=1}^{\infty} w_i y_i | \mathcal{F}_0\right\} / \sum_{i=1}^{\infty} w_i \quad (3)$$

with respect to $\mathbf{x}_1, \mathbf{x}_2, \dots$, with $\{w_i\}$ a weighting (discount) sequence. For instance, the choice $w_i = 1$ for $i = N + 1$ and $w_i = 0$ otherwise corresponds to a pure design problem, where emphasis is put on the estimation of \mathbf{x}^* after N observations; $w_i = 1$, $i = 1, \dots, N$, $w_{N+1} = K$ and $w_i = 0$ for $i > N + 1$ corresponds to the case where the best guess for \mathbf{x}^* at step N is used for the next K steps; the finite horizon case with no discount corresponds to $w_i = 1$, $i = 1, \dots, N$, $w_i = 0$ for $i \geq N + 1$, *etc.*

In the two cases (1,2), $E\{y_i | \theta, \mathbf{x}_i\} = f(\theta, \mathbf{x}_i)$, so that (3) becomes

$$\text{maximize} \quad E\left\{\sum_{i=1}^{\infty} w_i f(\theta, \mathbf{x}_i) | \mathcal{F}_0\right\} / \sum_{i=1}^{\infty} w_i. \quad (4)$$

We shall mainly consider design sequences that are constructed as follows: at step k , \mathbf{x}_{k+1} maximizes the sum of the predicted value of f , that is, $f(\hat{\theta}^k, \mathbf{x})$ with $\hat{\theta}^k$ the current estimated value of θ , and a weighted penalty term $\alpha_k d_k(\mathbf{x})$, with $d_k(\mathbf{x})$, the penalty for poor estimation of θ , also a function of (\mathbf{x}_i, y_i) , $i = 1, \dots, k$:

$$\mathbf{x}_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\hat{\theta}^k, \mathbf{x}) + \alpha_k d_k(\mathbf{x}). \quad (5)$$

In Section 2 we give an asymptotic result (infinite horizon, no discount) obtained for the regression problem (1) when $f(\theta, \mathbf{x})$ is linear in θ and $d_k(\mathbf{x})$ is the variance function used in the construction of D -optimum designs. Other penalty functions related to L -optimum design are then suggested.

Section 3 is devoted to the finite horizon case: approximately optimal strategies are suggested, and a particular sequence of weights $\{\alpha_k\}$ and penalty function $d_k(\mathbf{x})$ for (5) are obtained through a series of approximations of the original problem (4). Extension to nonlinear problems is presented in Section 4, where the case of Bernoulli trials (2) is considered through an example. Finally, Section 5 concludes and points out some open issues and possible developments.

2 Asymptotic results for linear response optimisation

2.1 *D*-optimum penalty

Consider the case where the horizon is infinite ($N = \infty$) and there is no discount ($w_i = 1$ for any i), and assume that the regression function is linear in θ , that is, $f(\theta, \mathbf{x}) = \mathbf{r}^\top(\mathbf{x})\theta$. We use the penalty given by the variance function used in the sequential construction of *D*-optimum designs, see [29, 8],

$$d_k(\mathbf{x}) = d_k^D(\mathbf{x}) = \mathbf{r}^\top(\mathbf{x})\mathbf{M}_k^{-1}\mathbf{r}(\mathbf{x}), \quad (6)$$

with \mathbf{M}_k the design matrix

$$\mathbf{M}_k = \sum_{i=1}^k \mathbf{r}(\mathbf{x}_i)\mathbf{r}^\top(\mathbf{x}_i). \quad (7)$$

Assume that $\mathbf{r}(\mathbf{x})$ is continuous on \mathcal{X} compact, that the first K_0 regressors $\mathbf{r}(\mathbf{x}_1), \dots, \mathbf{r}(\mathbf{x}_{K_0})$ are such that \mathbf{M}_{K_0} is positive definite. Also assume for the moment that θ in (1) takes a deterministic (but unknown) value θ (we shall come back to Bayesian estimation later), and that $\mathbf{r}^\top(\mathbf{x})\theta$ has a unique global maximiser \mathbf{x}^* in \mathcal{X} ; that is:

$$\begin{aligned} &\text{for all } \beta > 0, \text{ there exists } \epsilon > 0 \text{ such that} \\ &\mathbf{r}^\top(\mathbf{x})\bar{\theta} + \epsilon > \mathbf{r}^\top(\mathbf{x}^*)\bar{\theta} \Rightarrow \|\mathbf{x} - \mathbf{x}^*\| < \beta. \end{aligned} \quad (8)$$

We estimate θ by least squares (LS):

$$\hat{\theta}^k = \arg \min_{\theta \in \Theta} \sum_{i=1}^k [y_i - \mathbf{r}^\top(\mathbf{x}_i)\theta]^2, \quad (9)$$

with Θ a compact subset of \mathbb{R}^p such that $\bar{\theta} \in \Theta$. The difficulty is that for a suitable weighting sequence $\{\alpha_k\}$, the sequence $\{\mathbf{x}_k\}$ accumulates at the value $\mathbf{x}^*(\theta)$ that maximises $f(\theta, \mathbf{x})$ for some θ , but, at the same time, when $p = \dim \theta > 1$, a sequence too concentrated gives a singular design matrix, and thus does not yield consistent estimates. The following theorem is proved in [23], using the results in [21] on almost sure convergence of LS estimates. We denote by ξ_k the empirical measure of design points generated by (5) and by $\xi_{\mathbf{x}}$ the discrete measure that puts weight 1 at the point \mathbf{x} .

Theorem 1 *Assume that sequence of weights $\{\alpha_k\}$ in (5) is such that (α_k/k) $\log \alpha_k$ decreases monotonically and $\alpha_k/(\log k)^{1+\delta}$ increases monotonically to ∞ for some $\delta > 0$. Then, the sequence $\{\mathbf{x}_k\}$ generated by (5) and (6) is such that $\hat{\theta}^k \rightarrow \bar{\theta}$, $(1/k) \sum_{i=1}^k \mathbf{r}^\top(\mathbf{x}_i)\bar{\theta} \rightarrow \mathbf{r}^\top(\mathbf{x}^*)\bar{\theta}$ and $\xi_k \xrightarrow{w} \xi_{\mathbf{x}^*}$ (in the sense of weak convergence of measures) a.s. as $k \rightarrow \infty$, with $\mathbf{x}^* = \mathbf{x}^*(\bar{\theta}) = \arg \max_{\mathbf{x} \in \mathcal{X}} \mathbf{r}^\top(\mathbf{x})\bar{\theta}$.*

Assume now that $\bar{\theta}$ has a prior μ . One can first note that from the Lebesgue dominated convergence theorem, Theorem 1 implies

$$E\{(1/k) \sum_{i=1}^k \mathbf{r}^\top(\mathbf{x}_i)\bar{\theta}\} \rightarrow E\{\mathbf{r}^\top[\mathbf{x}^*(\bar{\theta})]\bar{\theta}\}$$

when μ is supported on Θ compact. Also, using the Bayesian imbedding approach of [25,19], one can obtain a.s. convergence results under weaker conditions on $\{\alpha_k\}$ under the assumption that the errors ϵ_i are i.i.d. and Gaussian $\mathcal{N}(0, \sigma^2)$. Indeed, for $\Theta = \mathbb{R}^p$ and $k > K_0$ the LS estimator $\hat{\theta}^k$ then coincides with the Bayesian estimator $E\{\theta|\mathcal{F}_k\}$ for the prior μ given by $\mathcal{N}(\hat{\theta}^{K_0}, \sigma^2 \mathbf{M}_{K_0}^{-1})$ (\mathbf{M}_{K_0} is positive definite). Let Q denote the probability measure induced by $\{\epsilon_k\}$, and write $(\mu \times Q)$ -a.s. for a property almost sure in the sense of the product measure $\mu \times Q$. From the martingale convergence theorem, $\hat{\theta}^\infty = \lim_{k \rightarrow \infty} \hat{\theta}^k$ exists and is finite and $\mathbf{r}^\top(\mathbf{x})\hat{\theta}^\infty$ is bounded on \mathcal{X} , $(\mu \times Q)$ -a.s., and the posterior covariance matrix \mathbf{C}_k tends to some limit \mathbf{C}_∞ , $(\mu \times Q)$ -a.s. When the smallest eigenvalue of \mathbf{M}_k satisfies $\lambda_{\min}(\mathbf{M}_k) \rightarrow \infty$, $(\mu \times Q)$ -a.s., \mathbf{C}_∞ is the null matrix and $\hat{\theta}^k$ converges to $\bar{\theta}$, the value of θ that generates the observations, $(\mu \times Q)$ -a.s. A straightforward extension of Theorem 1 is then as follows.

Corollary 1 *Assume that the errors ϵ_i are i.i.d. $\mathcal{N}(0, \sigma^2)$ and that sequence of weights $\{\alpha_k\}$ in (5) is such that $\alpha_k \rightarrow \infty$ and $\alpha_k/k \rightarrow 0$. Then, the sequence $\{\mathbf{x}_k\}$ generated by (5) and (6) is such that $\hat{\theta}^k \rightarrow \bar{\theta}$, $(1/k) \sum_{i=1}^k \mathbf{r}^\top(\mathbf{x}_i)\bar{\theta} \rightarrow \mathbf{r}^\top(\mathbf{x}^*)\bar{\theta}$ and $\xi_k \xrightarrow{w} \xi_{\mathbf{x}^*}$ (in the sense of weak convergence of measures), $(\mu \times Q)$ -a.s., as $k \rightarrow \infty$.*

The condition on α_k is weaker in Corollary 1 than in Theorem 1, but note that there may be a singular set (with respect to the Lebesgue measure) for $\bar{\theta}$ for which $\hat{\theta}^k$ is not consistent.

Using the results of [17], the assumption of normality in Corollary 1 can be relaxed, provided (i) the errors ϵ_i are i.i.d. with an almost everywhere strictly positive density h with respect to the Lebesgue measure, such that h'' is continuous and $(\log h)'' < 0$, (ii) the prior measure μ is absolutely continuous with respect to the Lebesgue measure and (iii) the LS estimates $\hat{\theta}^k$ are replaced by $E\{\theta|\mathcal{F}_k\}$.

Both Theorem 1 and Corollary 1 indicate that a penalty of the form $\beta \det \mathbf{M}_{k+1} / \det \mathbf{M}_k$, see, e.g., [1], which corresponds to taking $\alpha_k = \alpha$ constant in (5), is not strong enough: the sequential design that is obtained will not guarantee convergence of the parameters to $\bar{\theta}$, and $\{\mathbf{x}_k\}$ may accumulate at a value $\hat{\mathbf{x}}$ different from $\mathbf{x}^*(\bar{\theta})$.

2.2 L -optimum penalty

Using an idea similar to previous section, one could use a penalty function related to L -optimum design; that is,

$$d_k(\mathbf{x}) = d_k^L(\mathbf{x}) = \mathbf{r}^\top(\mathbf{x})\mathbf{M}_k^{-1}\mathbf{H}\mathbf{M}_k^{-1}\mathbf{r}(\mathbf{x}), \quad (10)$$

with \mathbf{H} a non negative definite matrix. It is shown in [28] that when \mathbf{H} is positive definite and the sequence $\{\mathbf{x}_k\}$ is given by $\mathbf{x}_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} d_k^L(\mathbf{x})$, the design measure ξ_k converges to a L -optimum design measure ξ_L^* that minimises trace $\mathbf{H}\mathbf{I}^{-1}(\xi)$, where $\mathbf{I}(\xi) = \int_{\mathcal{X}} \mathbf{r}(\mathbf{x})\mathbf{r}^\top(\mathbf{x})\xi(d\mathbf{x})$. Further work is required to check if a property similar to Theorem 1 can be obtained in this case.

A case that has retained much attention is when the design objective corresponds to the estimation of the point \mathbf{x}^* where $f(\theta, \mathbf{x})$ achieves its maximum, see [24], especially when f is a quadratic function of a scalar variable x , see, *e.g.*, [3,9]. Assume that

$$f(\theta, x) = \theta_0 + \theta_1 x + \theta_2 x^2/2, \quad (11)$$

with $x \in \mathcal{X}$, a compact subset of the real line. One has $x^*(\theta) = -\theta_1/\theta_2$, and, in the case where the errors ϵ_i are i.i.d., the asymptotic variance of $x^*(\hat{\theta}^k)$, with $\hat{\theta}^k$ given by (9), is proportional to $\mathbf{c}^\top \mathbf{M}_k^{-1} \mathbf{c}$, with

$$\mathbf{c} = \mathbf{c}(\theta) = \frac{\partial x^*(\theta)}{\partial \theta} = (-1/\theta_2) (0 \ 1 \ x^*)^\top.$$

Choosing the \mathbf{x}_k 's to maximise the accuracy of the estimation of $\mathbf{x}^*(\theta)$ corresponds to c -optimum design, that is, to L -optimum design with \mathbf{H} given by the rank-one matrix $\mathbf{c}\mathbf{c}^\top$. Note that the dependence of \mathbf{c} in θ makes the problem nonlinear. A Bayesian approach is used in [3,6], based on a design criterion of the type $E\{\mathbf{c}^\top \mathbf{M}_k^{-1} \mathbf{c}\}$, where $E\{\cdot\}$ denotes the expectation with respect to θ for a given prior. Sequential approaches are considered in [14,22]. One can also refer to [2] for the use of c -optimal design in the context of Bayesian estimation and to [5] for a survey on Bayesian experimental design. Following (10), a penalty function related to c -optimal design is thus

$$d_k(\mathbf{x}) = d_k^c(\mathbf{x}) = [\mathbf{r}^\top(\mathbf{x})\mathbf{M}_k^{-1}\mathbf{c}(\hat{\theta}^k)]^2. \quad (12)$$

Another approach, used in [24], is to derive the design criterion from the construction of a Bayesian risk related to the maximisation of $f(\theta, \mathbf{x})$. Assume that θ has a normal prior $\mathcal{N}(\hat{\theta}^0, \sigma^2\Omega)$, that the errors ϵ_i are i.i.d. $\mathcal{N}(0, \sigma^2)$ and that the discount factors w_i satisfy $w_i = 1$ for $i = N + 1$ and $w_i = 0$ otherwise. When $\mathbf{x}_1, \dots, \mathbf{x}_N$ are all chosen at the same time, this leads to the following (non sequential) L_B -optimal design problem:

$$\text{minimise trace } \mathbf{H}(\hat{\theta}^0) (\mathbf{M}_k + \Omega^{-1})^{-1}$$

with

$$\mathbf{H}(\hat{\theta}) = \frac{\partial^2 f[\theta, \mathbf{x}^*(\theta)]}{\partial \theta \partial \theta^\top} \Big|_{\hat{\theta}}.$$

The matrix \mathbf{H} can easily be proved to be non negative definite when f is linear in θ , see [24], and

$$\mathbf{H}(\theta) = \frac{\partial \mathbf{r}(\mathbf{x})}{\partial \mathbf{x}^\top} \Big|_{\mathbf{x}^*(\theta)} \frac{\partial \mathbf{x}^*(\theta)}{\partial \theta^\top}, \quad (13)$$

which can be expressed analytically when f is quadratic in $\mathbf{x} \in \mathbb{R}^q$, $q \geq 1$. For instance, in the case where f is given by (11) (with $\theta_2 < 0$ in order to have a function concave in x), one gets

$$\mathbf{H}(\hat{\theta}^0) = (-1/\hat{\theta}_2^0) \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & \mathbf{x}^*(\hat{\theta}^0) \\ 0 & \mathbf{x}^*(\hat{\theta}^0) & [\mathbf{x}^*(\hat{\theta}^0)]^2 \end{pmatrix} = -\hat{\theta}_2^0 \mathbf{c}(\hat{\theta}^0) \mathbf{c}^\top(\hat{\theta}^0).$$

This suggests substitution of $\mathbf{H}(\tilde{\theta}^k)$ for \mathbf{H} in (10) in the case of sequential design, with $\tilde{\theta}^k$ a Bayesian estimator of θ . Again, in the infinite horizon case with no discount, the choice of a weighting sequence $\{\alpha_k\}$ in (5) ensuring convergence of $\mathbf{x}^*(\tilde{\theta}^k)$ to $\mathbf{x}^*(\bar{\theta})$ and of ξ_k to $\xi_{\mathbf{x}^*(\bar{\theta})}$ remains an open issue.

3 Linear response optimisation with finite horizon

Assume the the errors ϵ_k in (1) are i.i.d. $\mathcal{N}(0, \sigma^2)$. We shall use an expansion in σ^2 to obtain an approximate solution to the problem (4). We first need to prove the following Lemma.

Lemma 1 *Let $g(\mathbf{x})$ and $h(\mathbf{x})$ be two times continuously differentiable functions on \mathcal{X} , a compact set of \mathbb{R}^q . Assume that g has a unique global maximum at \mathbf{x}^* , an interior point of \mathcal{X} , with $\partial^2 g(\mathbf{x})/\partial \mathbf{x} \partial \mathbf{x}^\top \Big|_{\mathbf{x}^*}$ negative definite, and let $\hat{\mathbf{x}}$ denote the point where $f(\mathbf{x}) = g(\mathbf{x}) + u h(\mathbf{x})$ reaches its maximum in \mathcal{X} . Then, $\|\hat{\mathbf{x}} - \mathbf{x}^*\| = O(u)$ and $|f(\mathbf{x}^*) - f(\hat{\mathbf{x}})| = O(u^2)$, $u \rightarrow 0$.*

Proof For u small enough, $\hat{\mathbf{x}}$ is an interior point of \mathcal{X} so that

$$\begin{aligned} \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}^\top} \Big|_{\hat{\mathbf{x}}} = 0 &= \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}^\top} \Big|_{\mathbf{x}^*} + (\hat{\mathbf{x}} - \mathbf{x}^*)^\top \frac{\partial^2 g(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \Big|_{\mathbf{x}^*} + u \frac{\partial h(\mathbf{x})}{\partial \mathbf{x}^\top} \Big|_{\hat{\mathbf{x}}} + o(\|\hat{\mathbf{x}} - \mathbf{x}^*\|) \\ &= (\hat{\mathbf{x}} - \mathbf{x}^*)^\top \frac{\partial^2 g(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \Big|_{\mathbf{x}^*} + u \frac{\partial h(\mathbf{x})}{\partial \mathbf{x}^\top} \Big|_{\hat{\mathbf{x}}} + o(\|\hat{\mathbf{x}} - \mathbf{x}^*\|), \end{aligned}$$

and $\|\hat{\mathbf{x}} - \mathbf{x}^*\| = O(u)$. Therefore,

$$\begin{aligned} f(\mathbf{x}^*) &= f(\hat{\mathbf{x}}) + (\mathbf{x}^* - \hat{\mathbf{x}})^\top \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}} \Big|_{\hat{\mathbf{x}}} \\ &\quad + \frac{1}{2} (\mathbf{x}^* - \hat{\mathbf{x}})^\top \frac{\partial^2 f(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \Big|_{\hat{\mathbf{x}}} (\mathbf{x}^* - \hat{\mathbf{x}}) + o(\|\hat{\mathbf{x}} - \mathbf{x}^*\|^2) \\ &= f(\hat{\mathbf{x}}) + O(u^2). \end{aligned}$$

Next theorem shows how to construct an approximate solution to problem (4).

Theorem 2 Assume that $w_i > 0$, $i = 1, \dots, N$ and $w_i = 0$ otherwise; that f is linear in θ , that is, $f(\theta, \mathbf{x}) = \mathbf{r}^\top(\mathbf{x})\theta$, with \mathbf{r} two times continuously differentiable in \mathbf{x} ; that the errors ϵ_i are i.i.d. $\mathcal{N}(0, \sigma^2)$ and the prior distribution for θ is normal $\mathcal{N}(\hat{\theta}^0, \sigma^2\Omega)$. Denote $\mathbf{x}^*(\theta) = \arg \max_{\mathbf{x} \in \mathcal{X}} \mathbf{r}^\top(\mathbf{x})\theta$ and $\tilde{\theta}^k = E\{\theta | \mathcal{F}_k\}$, and assume that $f(\tilde{\theta}^j, \mathbf{x})$ has a unique global maximum at $\mathbf{x}^*(\tilde{\theta}^j)$ which lies in the interior of \mathcal{X} , with $\partial^2 f(\tilde{\theta}^j, \mathbf{x}) / \partial \mathbf{x} \partial \mathbf{x}^\top|_{\mathbf{x}^*(\tilde{\theta}^j)}$ negative definite, $j = 0, \dots, N - 2$. Define $j_{k+1}(\mathbf{x})$ as the expected optimal gain to go at step k when \mathbf{x} is applied:

$$j_{k+1}(\mathbf{x}) = E\{w_{k+1}f(\theta, \mathbf{x}) + \max_{\mathbf{z} \in \mathcal{X}} w_{k+2}j_{k+2}(\mathbf{z}) | \mathcal{F}_k\}, \quad k = 0, \dots, N - 2,$$

$$j_N(\mathbf{x}) = E\{w_N f(\theta, \mathbf{x}) | \mathcal{F}_{N-1}\}.$$

It satisfies

$$j_{k+1}(\mathbf{x}) = J_{k+1}(\mathbf{x}) + O(\sigma^4), \quad k = 0, \dots, N - 2, \quad (14)$$

where

$$J_{k+1}(\mathbf{x}) = (w_N + \dots + w_{k+2})$$

$$\times \left\{ \mathbf{r}^\top[\mathbf{x}^*(\tilde{\theta}^k)]\tilde{\theta}^k + \frac{\sigma^2}{2} \text{trace}[\mathbf{H}(\tilde{\theta}^k)(\Omega^{-1} + \mathbf{M}_k)^{-1}] \right\}$$

$$- \frac{\sigma^2}{2} \text{trace} \left\{ \mathbf{H}(\tilde{\theta}^k) \sum_{j=0}^{N-k-2} w_{k+j+2} [\Omega_{j,k}^{-1} + \mathbf{r}(\mathbf{x})\mathbf{r}^\top(\mathbf{x})^{-1}]^{-1} \right\}$$

$$+ w_{k+1} \mathbf{r}^\top(\mathbf{x})\tilde{\theta}^k, \quad (15)$$

with $\Omega_{j,k} = \left\{ \Omega^{-1} + \mathbf{M}_k + j\mathbf{r}[\mathbf{x}^*(\tilde{\theta}^k)]\mathbf{r}^\top[\mathbf{x}^*(\tilde{\theta}^k)] \right\}^{-1}$ and $\mathbf{M}_k, \mathbf{H}(\theta)$ respectively given by (7), (13). Moreover, the strategy defined by $\mathbf{x}_N = \mathbf{x}^*(\tilde{\theta}^{N-1})$ and $\mathbf{x}_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} J_{k+1}(\mathbf{x})$, $k = 0, \dots, N - 2$, satisfies

$$\mathbf{x}_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} \left\{ w_{k+1} \mathbf{r}^\top(\mathbf{x})\tilde{\theta}^k + \frac{\sigma^2}{2} \right.$$

$$\left. \times \sum_{j=0}^{N-k-2} w_{k+j+2} \frac{\mathbf{r}^\top(\mathbf{x})\Omega_{j,k}\mathbf{H}(\tilde{\theta}^k)\Omega_{j,k}\mathbf{r}(\mathbf{x})}{1 + \mathbf{r}^\top(\mathbf{x})\Omega_{j,k}\mathbf{r}(\mathbf{x})} \right\}, \quad k = 0, \dots, N - 2, \quad (16)$$

$$E \left\{ \sum_{i=k+1}^N w_i f(\theta, \mathbf{x}_i) | \mathcal{F}_k \right\} = J_{k+1}(\mathbf{x}_{k+1}) + O(\sigma^4), \quad k = 0, \dots, N - 2, \quad (17)$$

and

$$\|\hat{\mathbf{x}}_{k+1} - \mathbf{x}_{k+1}\| = O(\sigma^4), \quad k = 0, \dots, N - 2, \quad \hat{\mathbf{x}}_N = \mathbf{x}_N, \quad (18)$$

where $\hat{\mathbf{x}}_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} j_{k+1}(\mathbf{x})$ corresponds to the optimum strategy.

Proof Straightforward matrix manipulation shows that \mathbf{x}_{k+1} given by (16) maximises (15). We prove (14) and (17) by backward induction on k . For $k = N - 2$, we have

$$\begin{aligned} j_{N-1}(\mathbf{x}) &= w_N E\{\mathbf{r}^\top[\mathbf{x}^*(\tilde{\theta}^{N-1})]\theta|\mathcal{F}_{N-2}\} + w_{N-1}\mathbf{r}^\top(\mathbf{x})\tilde{\theta}^{N-2} \\ &= w_N E\{\mathbf{r}^\top[\mathbf{x}^*(\tilde{\theta}^{N-1})]\tilde{\theta}^{N-1}|\mathcal{F}_{N-2}\} + w_{N-1}\mathbf{r}^\top(\mathbf{x})\tilde{\theta}^{N-2}. \end{aligned}$$

When $\mathbf{x}^*(\tilde{\theta}^{N-2})$ is an interior point of \mathcal{X} , a second-order Taylor expansion around $\tilde{\theta}^{N-2}$ similar to that used in [24] gives

$$\begin{aligned} j_{N-1}(\mathbf{x}) &= w_N\{\mathbf{r}^\top[\mathbf{x}^*(\tilde{\theta}^{N-2})]\tilde{\theta}^{N-2} + \frac{\sigma^2}{2}\text{trace}[\mathbf{H}(\tilde{\theta}^{N-2})(\Omega^{-1} + \mathbf{M}_{N-2})^{-1}]\} \\ &\quad - w_N\frac{\sigma^2}{2}\text{trace}\{\mathbf{H}(\tilde{\theta}^{N-2})[\mathbf{r}(\mathbf{x})\mathbf{r}^\top(\mathbf{x}) + \Omega^{-1} + \mathbf{M}_{N-2}]^{-1}\} \\ &\quad + w_{N-1}\mathbf{r}^\top(\mathbf{x})\tilde{\theta}^{N-2} + O(\sigma^4), \end{aligned}$$

which proves (14) for $k = N - 2$. Since $E\{\sum_{i=N-1}^N w_i f(\theta, \mathbf{x}_i)|\mathcal{F}_{N-2}\} = j_{N-1}(\mathbf{x}_{N-1})$, it also proves (17) for $k = N - 2$.

Assume that (14) is true at step k . At step $k - 1$ we have $j_k(\mathbf{x}) = E\{w_k f(\theta, \mathbf{x}) + \max_{\mathbf{z} \in \mathcal{X}} j_{k+1}(\mathbf{z})|\mathcal{F}_{k-1}\}$. Using Lemma 1 with $u = \sigma^2$, we get

$$j_k(\mathbf{x}) = w_k \mathbf{r}^\top(\mathbf{x})\tilde{\theta}^{k-1} + E\{J_{k+1}[\mathbf{x}^*(\tilde{\theta}^k)]|\mathcal{F}_{k-1}\} + O(\sigma^4).$$

Grouping the terms $\mathbf{r}^\top[\mathbf{x}^*(\tilde{\theta}^k)]\tilde{\theta}^k$ and using a second-order Taylor expansion around $\tilde{\theta}^{k-1}$, we obtain

$$\begin{aligned} j_k(\mathbf{x}) &= (w_N + \dots + w_{k+2} + w_{k+1}) \left\{ \mathbf{r}^\top[\mathbf{x}^*(\tilde{\theta}^{k-1})]\tilde{\theta}^{k-1} \right. \\ &\quad \left. + \frac{\sigma^2}{2}\text{trace}[\mathbf{H}(\tilde{\theta}^{k-1})(\Omega^{-1} + \mathbf{M}_{k-1})^{-1}] \right\} \\ &\quad - (w_N + \dots + w_{k+2} + w_{k+1})\frac{\sigma^2}{2} \\ &\quad \times \text{trace}\{\mathbf{H}(\tilde{\theta}^{k-1})[\Omega^{-1} + \mathbf{M}_{k-1} + \mathbf{r}(\mathbf{x})\mathbf{r}^\top(\mathbf{x})]^{-1}\} \\ &\quad + (w_N + \dots + w_{k+2})\frac{\sigma^2}{2} \\ &\quad \times \text{trace}\{\mathbf{H}(\tilde{\theta}^{k-1})[\Omega^{-1} + \mathbf{M}_{k-1} + \mathbf{r}(\mathbf{x})\mathbf{r}^\top(\mathbf{x})]^{-1}\} \\ &\quad - \frac{\sigma^2}{2}\text{trace}\left\{ \mathbf{H}(\tilde{\theta}^{k-1}) \sum_{j=0}^{N-k-2} w_{k+j+2} [\Omega_{j+1,k-1}^{-1} + \mathbf{r}(\mathbf{x})\mathbf{r}^\top(\mathbf{x})]^{-1} \right\} \\ &\quad + w_k \mathbf{r}^\top(\mathbf{x})\tilde{\theta}^{k-1} + O(\sigma^4). \end{aligned}$$

A simplification of the terms on the last three lines gives $j_k(\mathbf{x}) = J_k(\mathbf{x}) + O(\sigma^4)$, which proves (14). Since \mathbf{x}_{k+1} maximises (15), similar arguments using (17) and Lemma 1 give $E\{\sum_{i=k}^N w_i f(\theta, \mathbf{x}_i)|\mathcal{F}_{k-1}\} = J_k(\mathbf{x}_k) + O(\sigma^4)$, and (17) is proved.

Finally, since \mathbf{x}_{k+1} maximises $J_{k+1}(\mathbf{x})$, (14) and Lemma 1 with $u = \sigma^4$ give (18).

Remark 1 The assumptions on $f(\tilde{\theta}^j, \mathbf{x})$ and $\mathbf{x}^*(\tilde{\theta}^j)$ for $j = 0, \dots, N-2$, in Theorem 2 are most often difficult, if not impossible, to check beforehand. Note, however, that it is always possible to apply the strategy and check the assumptions afterwards: if they are satisfied, the theorem applies and the strategy used is approximately optimal in the sense of the theorem.

The property stated in Theorem 2 suggests a simpler suboptimal strategy, as given in the following corollary.

Corollary 2 *Under the same conditions, and with the same notations, as in Theorem 2, the strategy defined by $\mathbf{x}'_N = \mathbf{x}^*(\tilde{\theta}^{N-1})$ and*

$$\mathbf{x}'_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} \left\{ w_{k+1} \mathbf{r}^\top(\mathbf{x}) \tilde{\theta}^k + \frac{\sigma^2}{2} \sum_{j=1}^{N-k-1} j w_{k+j+1} \frac{\mathbf{r}^\top(\mathbf{x}) \Omega_{0,k} \mathbf{H}(\tilde{\theta}^k) \Omega_{0,k} \mathbf{r}(\mathbf{x})}{1 + j \mathbf{r}^\top(\mathbf{x}) \Omega_{0,k} \mathbf{r}(\mathbf{x})} \right\}, \quad k = 0, \dots, N-2, \quad (19)$$

is approximately optimal in the following sense:

$$\max_{\mathbf{x} \in \mathcal{X}} j_{k+1}(\mathbf{x}) - E \left\{ \sum_{i=k+1}^N w_i f(\theta, \mathbf{x}'_i) \middle| \mathcal{F}_k \right\} = O(\sigma^4), \quad (20)$$

$$|\hat{\mathbf{x}}_{k+1} - \mathbf{x}'_{k+1}| = O(\sigma^2), \quad (21)$$

for $k = 0, \dots, N-1$.

Proof We first prove that $E \left\{ \sum_{i=k+1}^N w_i f(\theta, \mathbf{x}'_i) \middle| \mathcal{F}_k \right\} = J_{k+1}(\mathbf{x}'_{k+1}) + O(\sigma^4)$, for $k = N-2, \dots, 1$.

For $k = N-2$ we have again $E \left\{ \sum_{i=N-1}^N w_i f(\theta, \mathbf{x}'_i) \middle| \mathcal{F}_{N-2} \right\} = j_{N-1}(\mathbf{x}'_{N-1})$ which equals $J_{N-1}(\mathbf{x}'_{N-1}) + O(\sigma^4)$ from Theorem 2. Assume that the property is true at step k . At step $k-1$ we have

$$E \left\{ w_k f(\theta, \mathbf{x}) + \sum_{i=k+1}^N w_i f(\theta, \mathbf{x}'_i) \middle| \mathcal{F}_{k-1} \right\} = w_k \mathbf{r}^\top(\mathbf{x}) \tilde{\theta}^{k-1} + E \left\{ J_{k+1}(\mathbf{x}'_{k+1}) \middle| \mathcal{F}_{k-1} \right\} + O(\sigma^4),$$

and easy matrix manipulation using (19) and (15) shows that $\mathbf{x}'_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} [J_{k+1}(\mathbf{x}) + \sigma^2 h(\mathbf{x})]$, for some $h(\mathbf{x})$. Lemma 1 with $u = \sigma^2$ gives $J_{k+1}(\mathbf{x}'_{k+1}) = J_{k+1}[\mathbf{x}^*(\tilde{\theta}^k)] + O(\sigma^4)$, and similarly to the proof of Theorem 2, $E \left\{ w_k f(\theta, \mathbf{x}) + \sum_{i=k+1}^N w_i f(\theta, \mathbf{x}'_i) \middle| \mathcal{F}_{k-1} \right\} = J_k(\mathbf{x}) + O(\sigma^4)$.

Finally, since the optimal strategy $\hat{\mathbf{x}}_{k+1}$ maximises $j_{k+1}(\mathbf{x}) = J_{k+1}(\mathbf{x}) + O(\sigma^4)$, whereas \mathbf{x}'_{k+1} maximises a function that takes the form $J_{k+1}(\mathbf{x}) + \sigma^2 h(\mathbf{x})$, Lemma 1 gives (20-21).

Table 1 Empirical means and standard deviations (std) of $(1/N) \sum_{i=1}^N f(\bar{\theta}, x_i)$ for strategies S_1, \dots, S_4 in Example 1 (100 repetitions)

	mean	std
S_1	0.1788	0.0523
S_2	0.3774	0.0289
S_3	0.4234	0.0284
S_4	0.4767	0.0484

When $w_i = 1, i = 1, \dots, N$ and $w_i = 0$ otherwise, with N large, the strategy (19) can be further approximated by

$$\mathbf{x}''_{k+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} \left\{ \mathbf{r}^\top(\mathbf{x}) \tilde{\theta}^k + (N - k - 1) \frac{\sigma^2 \mathbf{r}^\top(\mathbf{x}) \Omega_{0,k} \mathbf{H}(\tilde{\theta}^k) \Omega_{0,k} \mathbf{r}(\mathbf{x})}{\mathbf{r}^\top(\mathbf{x}) \Omega_{0,k} \mathbf{r}(\mathbf{x})} \right\},$$

$$k = 0, \dots, N - 2, \quad (22)$$

which has the form (5).

Remark 2 It is clear from the proof of Corollary 2 that properties similar to (20-21) can be obtained for other strategies than (19). What makes this strategy attractive is that it amounts to substituting $\mathbf{r}(\mathbf{x})$ for $\mathbf{r}[\mathbf{x}(\tilde{\theta}^k)]$ in $J_{k+1}(\mathbf{x})$ to construct \mathbf{x}'_{k+1} . This makes the rule (19) suboptimal compared to (16) that directly maximises $J_{k+1}(\mathbf{x})$ (compare (21) to (18)), but, on the other hand, $\mathbf{r}(\mathbf{x}'_{k+1})$ can be expected to be close to $\mathbf{r}[\mathbf{x}(\tilde{\theta}^k)]$ and the loss can be expected to be small.

3.1 Example 1

We assume that $f(\theta, x)$ is given by (11), with $\mathcal{X} = [-1, 1]$, and that the observations y_k are generated by (1) for $\theta = \bar{\theta} = (0 \ 3.2 \ -8)^\top$ with the errors ϵ_k i.i.d. $\mathcal{N}(0, \sigma^2)$, $\sigma = 1$. The prior for θ is $\mathcal{N}(\hat{\theta}^0, \sigma^2 \Omega)$, with $\hat{\theta}^0 = (2 \ -4 \ -1)^\top$ and $\Omega = 10^6 \mathbf{I}_3$, with \mathbf{I}_3 the 3-dimensional identity matrix. Note that these numerical values give a prior guess for x^* at -4, whereas the true location is at 0.4.

We take $w_i = 1, i = 1, \dots, N$ and $w_i = 0$ otherwise, with $N = 100$, and compare four different strategies: S_1 corresponds to (5) with the penalty (6) and $\alpha_k = \sigma^2 (\log k)^2$, and S_2, S_3 and S_4 correspond respectively to (22), (19) and (16). Table 1 presents the results obtained for 100 independent repetitions of the experiment (for each experiment, the same values of observations errors ϵ_k are used for the four strategies). As expected, performances improve from S_1 to S_4 .

Figure 1 (resp. 2) presents a typical realization of the sequences $\{x_k\}$ (resp. $\{f(\bar{\theta}, x_k)\}$) generated by the four strategies. One can notice in Figure 1 that x_k converges to $x^* = -\bar{\theta}_1/\bar{\theta}_2 = 0.4$ for S_1 (see Theorem 1), which




Fig. 1 Sequences $\{x_k\}$ generated by strategies S_1 to S_4 (from left to right) in Example 1




Fig. 2 Sequences $\{f(\bar{\theta}, x_k)\}$ generated by strategies S_1 to S_4 (from left to right) in Example 1. The optimum value $f[\bar{\theta}, x^*(\bar{\theta})]$ is indicated by the dashed line

does not seem to be the case for the other strategies. However, the performance measured in terms of (4) is much better for the other strategies, in particular S_4 that makes a particularly good compromise between estimation and optimisation, see Figure 2 where the optimum value $f[\bar{\theta}, x^*(\bar{\theta})]$ is indicated by the dashed line.

4 Extension to nonlinear situations

4.1 Nonlinear regression

Whereas Theorem 1 relies on Corollary 3 of [21], deriving similar asymptotic results in the case of a nonlinear regression model would require using Theorem 2 of [20]. The conditions in [20] being more stringent than those in [21], one may conjecture that α_k should increase faster than any power of $\log k$ when $k \rightarrow \infty$, compare with Theorem 1. Using Bayesian estimation instead of LS would certainly require weaker conditions on α_k , as it was the case for Corollary 1 in comparison with Theorem 1, but deriving precise results on this is still an open issue.

In the case of finite horizon, the approaches of Section 3 can be used through a linearisation of the model response. Notice, however, that this implies that terms of order $O(\sigma^2)$ are neglected, so that the properties (17) and (18) are no longer true.

4.2 Example 2: Bernoulli trials

We consider the situation defined by (2), with

$$f(\theta, x) = \frac{\theta_1 + \theta_2}{\theta_2} [1 - \exp(-x/\theta_1)] \exp(-x/\theta_2),$$

with $\mathcal{X} = \{x_{\min}, x_{\min} + \delta, x_{\min} + 2\delta, \dots, x_{\max}\}$, where $x_{\min} = 2, x_{\max} = 100, \delta = 2$, which gives $\#\mathcal{X} = 50$. This may correspond to a dose-response problem, where efficiency is small at low doses x and toxic failures occur at large doses, see, *e.g.*, [18]. We assume that θ has a normal prior $\mathcal{N}(\hat{\theta}^0, \Omega)$, with $\hat{\theta}_0 = (70 \ 70)^\top$ and $\Omega = 10^3 \mathbf{I}_2$. The observations are generated with $\theta = \bar{\theta} = (3 \ 50)^\top$. Figure 3 presents the two probabilities $f(\theta, x)$ (full line) and $f(\hat{\theta}^0, x)$ (dashed line) as functions of x . The prior guess $x^*(\hat{\theta}^0)$ is at 48.52, whereas the true value is at $x^*(\bar{\theta}) = 8.615$.

We estimate θ by the maximum *a posteriori* estimator and denote

$$\hat{\theta}^k = \arg \max \sum_{i=1}^k \log \{y_i f(\theta, x_i) + (1 - y_i)[1 - f(\theta, x_i)]\} \\ - \frac{1}{2}(\theta - \hat{\theta}^0)^\top \Omega^{-1}(\theta - \hat{\theta}^0).$$

The expression of $f(\theta, x)$ gives $x^*(\theta) = -\theta_1 \log[\theta_1/(\theta_1 + \theta_2)]$ and

$$\mathbf{c}(\theta) = \frac{\partial x^*(\theta)}{\partial \theta} = \begin{pmatrix} -\theta_2/(\theta_1 + \theta_2) - \log[\theta_1/(\theta_1 + \theta_2)] \\ \theta_1/(\theta_1 + \theta_2) \end{pmatrix}.$$

The design matrix \mathbf{M}_k is replaced by $\mathbf{M}_k(\hat{\theta}^k)$, with

$$\mathbf{M}_k(\theta) = \sum_{i=1}^k \frac{\partial f(\theta, x_i)}{\partial \theta} \frac{\partial f(\theta, x_i)}{\partial \theta^\top} \frac{1}{f(\theta, x_i)[1 - f(\theta, x_i)]},$$




Fig. 3 Probabilities $f(\bar{\theta}, x)$ (full line) and $f(\hat{\theta}^0, x)$ (dashed line) as functions of x in Example 2

that would be the Fisher information matrix evaluated at θ if the design levels x_i were nonrandom constants. Note that it still has the form (7), with now

$$\mathbf{r}(x) = \mathbf{r}(\theta, x) = \frac{1}{\{f(\theta, x)[1 - f(\theta, x)]\}^{1/2}} \frac{\partial f(\theta, x)}{\partial \theta}.$$

We shall compare four strategies. S_1 and S_2 are defined by (5) with the weights $\alpha_k = k^{1/4}$, and penalty functions (6) and (12) respectively. S_3 corresponds to a randomized Polya urn design, see [7], and S_4 to an optimizing up-and-down design, see [18]. S_3 is initialized by putting one ball of each label in the urn, S_4 samples by pairs, at $(P_k - \delta/2, P_k + \delta/2)$ at step k , and we take P_1 at the median of \mathcal{X} .

Figure 4 (resp. 5) presents a typical realization of the sequences $\{x_k\}$ (resp. $\{f(\bar{\theta}, x_k)\}$) generated by the four strategies. On Figure 5, the optimum value $f[\bar{\theta}, x^*(\bar{\theta})]$ is indicated by the dotted line. Figure 6 presents histograms of the values of x_k generated by the four strategies.

Convergence to the maximum of $f(\theta, x)$ is faster for S_1 and S_2 than for S_3 and S_4 . The difference in performance is even stronger when δ decreases, that is, when the number of points in \mathcal{X} increases. Whereas the value of δ has no significant effect of S_1 and S_2 , decreasing δ makes the behaviour of S_3 more erratic and increases the time S_4 requires to move away from the initial values of x_k . This effect is already apparent on Figures 4 to 6. A possible compromise strategy would be to start with S_3 or S_4 and a large value of δ (that is, few point in \mathcal{X}), and then decrease δ and switch to S_1 or S_2 .




Fig. 4 Sequences $\{x_k\}$ generated by strategies S_1 to S_4 (from left to right) in Example 2




Fig. 5 Sequences $\{f(\bar{\theta}, x_k)\}$ generated by strategies S_1 to S_4 (from left to right) in Example 2. The optimum value $f[\bar{\theta}, x^*(\bar{\theta})]$ is indicated by the dotted line

5 Conclusions and further developments

We considered the problem of choosing a sequence of values \mathbf{x}_k for \mathbf{x} that maximize some function $f(\mathbf{x})$ in two different situations: $f(\mathbf{x}_k)$ is observed with errors (regression model), or binary responses are observed, the probability of a positive response at \mathbf{x}_k being $f(\mathbf{x}_k)$. Assuming that a parametric




Fig. 6 Histograms of the values of x_k generated by strategies S_1 to S_4 (from left to right) in Example 2

form is known for f , that is, $f(\mathbf{x}) = f(\theta, \mathbf{x})$ with θ unknown, different strategies have been suggested.

An approximately optimal strategy has been constructed in the case of a linear regression model with finite horizon. In more general situations, we considered sequences constructed according to the rule $\mathbf{x}_{k+1} = \arg \max_{\mathbf{x}} f(\hat{\theta}^k, \mathbf{x}) + \alpha_k d_k(\mathbf{x})$, with $\hat{\theta}^k$ an estimated value of θ obtained from $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_k, y_k)$ and $d_k(\mathbf{x})$ a penalty for poor estimation. This is easily implemented in the case of binary responses, and numerical examples illustrate that such strategies may outperform more traditional ones, like Polya urn designs and up-and-down methods. Such a result should not be a surprise, since methods based on a parametric form of f use more information than model-free approaches. On the other hand, their asymptotic behaviour is more difficult to study, due to the intricate connection between estimation of θ and optimisation. Only the linear regression case, with a penalty related to D -optimum design, seems to have been considered so far, see [23]. Extension to nonlinear regression problems is currently under study. In practice, it might be reasonable to start with a model-free approach, in order to gather enough information about the form of f , and then switch to a parametric approach of the type considered here. This raises many issues (choice of model structure, choice of switching time, robustness with respect to model mis-specification, *etc.*) which, together with comparisons with other approaches (e.g., the bandit design of [16]), deserve further study.

Acknowledgments. This work was presented at a meeting in Bertinoro on December 2000. Prof. A. Giovagnoli from the University of Bologna, who organised the meeting, is gratefully acknowledged for her invitation.

References

1. K.J. Aström and B. Wittenmark. *Adaptive Control*. Addison Wesley, 1989.
2. K. Chaloner. Optimal bayesian experimental design for linear models. *Annals of Statistics*, 12(1):283–300, 1984.
3. K. Chaloner. Bayesian design for estimating the turning point of a quadratic regression. *Commun. Statist.-Theory Meth.*, 18(4):1385–1400, 1989.
4. K. Chaloner and K. Larntz. Optimal Bayesian design applied to logistic regression experiments. *Journal of Statistical Planning and Inference*, 21:191–208, 1989.
5. K. Chaloner and I. Verdinelli. Bayesian experimental design: a review. *Statistical Science*, 10(3):273–304, 1995.
6. H. Dette. A note on Bayesian c - and D -optimal designs in nonlinear regression models. *Annals of Statistics*, 24(3):1225–1234, 1996.
7. S. Durham, N. Flournoy, and W. Li. A sequential design for maximizing the probability of a favorable response. *Can. J. Statist.*, 26:479–495, 1998.
8. V.V. Fedorov. *Theory of Optimal Experiments*. Academic Press, New York, 1972.
9. V.V. Fedorov and W.G. Müller. Another view on optimal design for estimating the point of extremum in quadratic regression. *Metrika*, 46:147–157, 1997.
10. A.A. Fel'dbaum. Dual control theory. I. *Avtomatika i Telemekhanika*, 21(9):1240–1249, 1960.
11. A.A. Fel'dbaum. Dual control theory. II. *Avtomatika i Telemekhanika*, 21(11):1453–1464, 1960.
12. A.A. Fel'dbaum. The theory of dual control. III. *Avtomatika i Telemekhanika*, 22(1):3–16, 1961.
13. A.A. Fel'dbaum. The theory of dual control. IV. *Avtomatika i Telemekhanika*, 22(2):129–142, 1961.
14. I. Ford and S.D. Silvey. A sequentially constructed design for estimating a nonlinear parametric function. *Biometrika*, 67(2):381–388, 1980.
15. J. Ginebra and M.K. Clayton. Response surface bandits. *Journal of Royal Statistical Society*, B57(4):771–784, 1995.
16. J. Hardwick and Q. Stout. Optimizing a unimodal response function for binary variables. In A. Atkinson, B. Bogacka, and A. Zhigljavsky, editors, *Optimum Design 2000*, chapter 18, pages 195–210. Kluwer, Dordrecht, 2001.
17. I. Hu. Strong consistency of Bayes estimates in stochastic regression models. *Journal of Multivariate Analysis*, 57:215–227, 1996.
18. E.E. Kpamegan and N. Flournoy. An optimizing up-and-down design. In *Optimum Design 2000*. Kluwer, 2000. (to appear).
19. P.R. Kumar. Convergence of adaptive control schemes using least-squares parameter estimates. *IEEE Transactions on Automatic Control*, 35(4):416–424, 1990.
20. T.L. Lai. Asymptotic properties of nonlinear least squares estimates in stochastic regression models. *Annals of Statistics*, 22(4):1917–1930, 1994.
21. T.L. Lai and C.Z. Wei. Least squares estimates in stochastic regression models with applications to identification and control of dynamic systems. *Annals of Statistics*, 10(1):154–166, 1982.
22. W.G. Müller and B.M. Pötscher. Batch sequential design for a nonlinear estimation problem. In V.V. Fedorov, W.G. Müller, and I.N. Vuchkov, editors, *Model-Oriented Data Analysis II, Proceedings 2nd IIASA Workshop, St Kyrík (Bulgaria), May 1990*, pages 76–87. Physica Verlag, Heidelberg, 1992.

23. L. Pronzato. Adaptive optimisation and D -optimum experimental design. *Annals of Statistics*, 28(6):1743–1761, 2000.
24. L. Pronzato and E. Walter. Experimental design for estimating the optimum point in a response surface. *Acta Applicandae Mathematicae*, 33:45–68, 1993.
25. J. Sternby. On consistency for the method of least squares using martingale theory. *IEEE Transactions on Automatic Control*, 22(3):346–352, 1977.
26. I. Verdinelli and J.K. Kadane. Bayesian designs for maximizing information and outcome. *J. Amer. Stat. Assoc.*, 87(418):510–515, 1992.
27. C.F.J. Wu. Optimal design for percentile estimation of a quantal response curve. In Y. Dodge, V. V. Fedorov, and H. P. Wynn, editors, *Optimal Design and Analysis of Experiments*, pages 213–223. North Holland, Amsterdam, 1988.
28. C.F.J. Wu and H.P. Wynn. The convergence of general step-length algorithms for regular optimum design criteria. *Annals of Statistics*, 6(6):1273–1285, 1978.
29. H.P. Wynn. The sequential generation of D -optimum experimental designs. *Annals of Math. Stat.*, 41:1655–1664, 1970.