New Developments and Applications in Experimental Design
IMS Lecture Notes – Monograph Series (1998) Volume 34

SEQUENTIAL DESIGN AND ACTIVE CONTROL

BY ROLAND GAUTIER AND LUC PRONZATO

Université de Nice

In nonlinear situations, optimal experimental conditions generally depend upon unknown parameters to be estimated from the data collected during the experiments. A natural approach then consists in designing the experiments sequentially, that is, alternating estimation and design phases. Each design phase can be considered as a control action applied on the system. Sequential design thus corresponds to adaptive control, with, e.g., the precision of estimation as objective. Even for a purely static system, designing the experiments sequentially introduces a feedback of information, which induces dynamics into the design procedure. Several sequential schemes, corresponding to different control policies, are considered. The optimal one corresponds to closed–loop control and is the solution of a stochastic dynamic–programming problem, which is extremely difficult to solve even in very simple cases. Suboptimal strategies are thus proposed. Examples for nonlinear regression models are presented.

1. Introduction. In nonlinear situations the optimal experimental conditions generally depend upon parameters to be estimated during the experiments. Sequential design provides a method of circumventing this issue: after each observation, or each sample of n observations, the parameters are estimated, and this information is used to design the experiment for the next observation(s). One can refer to Chernoff (1972) for a monograph on sequential analysis, including the construction of stopping rules. Since stopping rules will not be discussed here, one may prefer to call the designs considered "adaptive". We shall, however, call them "sequential" since this seems to be common practice. The design phases can then be considered as control actions applied on the system, the control objective being for instance the precision of estimation. The dependence of the kth experiment on previous observations introduces a feedback of information, and thus induces dynamics into the design procedure, even in cases where the process under study is purely static. In this sense, one can think about the kth design step as the design of the experiment to be performed at *time* k, and sequential design can be considered as an adaptive–control problem. For instance, the

Received September 1997; revised December 1997.
AMS 1991 subject classifications. Primary 62L05, 62K05; secondary 49L20, 49N35.
Key words and phrases. Active control, adaptive control, certainty equivalence, closed–loop control, dynamic programming, open–loop feedback, optimal design.

usual non–Bayesian myopic approach of sequential design can be considered as *Forced Certainty Equivalence* control [see Pronzato, Walter and Kulcsár (1993), Runggaldier (1993)]. Relations between experimental design and control are briefly described in Section 2. Section 3.1 discusses convergence problems for batch–sequential design, whereas Section 3.2 concerns fully–sequential design. The case where the number of observations is fixed is considered in Section 4. The optimal strategy corresponds to closed–loop control, and is the solution of a stochastic dynamic–programming problem [see Bellman (1957)], which is extremely difficult to solve even in very simple situations. Examples with two–stage sampling can be found in Zacks (1977), Pronzato, Walter and Kulcsár (1993), Kulcsár, Pronzato and Walter (1994). Different suboptimal strategies, passive and active, are presented. In conclusion, Section 5 gives some perspectives.

2. Experimental design and control. Consider a system on which N observations $y_1, \ldots, y_N$ are to be performed, with

$$y_k = \eta(\bar{\theta}, \xi_k) + \epsilon_k \,, \tag{2.1}$$

where $\bar{\theta} \in \Theta$ is the true value of the model parameters, with $\Theta \subset I\!R^p$ a compact set, ξ_k denotes the experimental conditions for the k-th observation and is assumed to belong to a compact set, $\{\epsilon_k\}$ is an i.i.d. sequence of normal variables $\mathcal{N}(0, \sigma^2)$, with σ known, and $\eta(\theta, \xi)$ is the model response for the value θ of the model parameters and experimental conditions ξ. The response may be nonlinear in θ, and it will then be assumed two times continuously differentiable with respect to θ. Least–squares (LS) estimation is considered, which coincides here with maximum–likelihood estimation, and we shall denote

$$\hat{\theta}^k = \hat{\theta}_{LS}(y_1^k) = \arg\min_{\theta \in \Theta} \sum_{i=1}^{k} [y_i - \eta(\theta, \xi_i)]^2 \,.$$

Here and in what follows x_i^j denotes the list of variables $(x_i, x_{i+1}, \ldots, x_j)$.

The design problem is to choose the experimental conditions that yield the highest possible accuracy for the parameter estimators $\hat{\theta}^N$. Under suitable conditions [Jennrich (1969)], for large values of N, $\hat{\theta}^N$ approximately follows the normal distribution $\mathcal{N}(\bar{\theta}, \mathbf{M}^{-1}(\bar{\theta}, \xi_1^N))$, where

$$\mathbf{M}(\theta, \xi_1^N) = \frac{1}{\sigma^2} \sum_{i=1}^{N} \frac{\partial \eta(\theta, \xi_i)}{\partial \theta} \frac{\partial \eta(\theta, \xi_i)}{\partial \theta^T} \tag{2.2}$$

is the Fisher information matrix. We shall thus use $\Phi(\mathbf{M})$ as a criterion for experimental design (to be maximized) [see, e.g., Fedorov (1972)]. Special attention will be devoted to D–optimality, for which $\Phi(\cdot) = \det(\cdot)$. Note that a more accurate approximation of the distribution of $\hat{\theta}^N$ may be used to define more accurate design criteria [see Pázman and Pronzato (1992), Pronzato and Pázman (1994)].

There are two situations where design and control are strongly related: (i) the experimental conditions correspond to the input of a *dynamical system*, (ii) the experiment is designed *sequentially*. This paper essentially concerns case (ii). Case (i) can

be illustrated by the simple example of model with Finite Impulse Response (FIR), which is classical in control engineering, signal processing, time–series analysis, etc.

EXAMPLE 1. Consider the situation where the model response $\eta(\theta, k)$ associated with the k-th measurement is a linear combination of design variables $(\xi_{k-p+1}, \ldots, \xi_k)$; i.e.,

$$\eta(\theta, k) = \sum_{i=0}^{p-1} \theta_i \xi_{k-i}\,, \tag{2.3}$$

and $\xi_i = 0$ for $i \leq 0$. D–optimal design then consists in choosing the input sequence $\{\xi_k\}$ that maximizes the determinant

$$\det \left[\sum_{k=1}^{N} (\xi_k, \xi_{k-1}, \ldots, \xi_{k-p+1})^T (\xi_k, \xi_{k-1}, \ldots, \xi_{k-p+1}) \right],$$

with e.g. a power constraint of the form $\sum_{k \leq N} \xi_k^2 \leq P_{\max}$, which defines the admissible design region. The determination of the optimal input sequence is then a nonlinear control problem.

Since $\eta(\theta, k)$ given by (2.3) is linear in θ, the optimal design does not depend on θ and sequential design is not needed. Consider now a trajectory–following problem, where $\{\xi_k\}$ should minimize

$$\sum_{k=1}^{N} [\eta(\bar{\theta}, k) - \mathrm{ref}(k)]^2 + \sum_{k=1}^{N} \xi_k^2\,,$$

with ref$(\cdot)$ a given reference trajectory to be followed by the system and $\eta(\theta, k)$ still given by (2.3). The optimal input–design then depends on $\bar{\theta}$, which is unknown and can be estimated at time k by $\hat{\theta}^k$. This type of problem will be considered in Section 4. Another situation where the optimal design depends on the value of $\bar{\theta}$ is when the dynamical system contains an autoregressive part [see, e.g., Goodwin and Payne (1977), Zarrop (1979), Walter and Pronzato (1997)]. The determination of the optimal input sequence for the estimation of θ is then again a nonlinear control problem, but it may be useful to design the experiment sequentially.

Sequential design is strongly related to control even if the model has no dynamical aspects. In the rest of the paper we shall focuss on static regression models such as (2.1), where $\eta(\theta, \xi_k)$ is independent of $\xi_{k-1}, \xi_{k-2}, \ldots$ and is nonlinear in θ, so that $\mathbf{M}(\theta, \xi_1^N)$ depends on θ. Classical fully–sequential design is then as shown in Figure 1, with $\hat{\theta}^0$ some prior guess for θ.

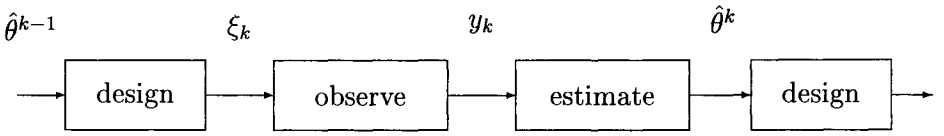

FIG. 1. *Fully–sequential design.*

Since ξ_k depends on y_1^{k-1}, designing ξ_k corresponds to controlling a dynamical system. One may have two different objectives in mind: first, one may wish that $\xi_1^N = (\xi_1, \ldots, \xi_N)$ tend to the optimal design for $\bar{\theta}$ when N tends to infinity, this is considered in Section 3; second, for a fixed value of N, one may wish to maximize $\Phi[\mathbf{M}(\bar{\theta}, \xi_1^N)]$, with $\bar{\theta}$ unknown. This corresponds to adaptive control with finite horizon and is considered in Section 4.

REMARK 1. We shall still use $\mathbf{M}(\theta, \xi_1^N)$, given by (2.2), to characterize the precision of the estimation, although it does not correspond to the observed or expected information matrix [Ford, Titterington and Wu (1985), Wu (1985)].

3. Convergence issues.

3.1. *Batch–sequential design.* Classical batch–sequential D–optimal design is as shown in Figure 2, where $\breve{\theta}^0$ is some prior guess for θ, and where $\Xi_D^k = \xi_{(k-1)n+1}^{kn}$ is the exact D–optimal design of size $n \geq p =$ dim θ for the estimator $\breve{\theta}^{k-1}$, that is

$$\Xi_D^k = \Xi_D(\breve{\theta}^{k-1}) = \arg\max_{\Xi} \det \mathbf{M}(\breve{\theta}^{k-1}, \Xi)\,.$$

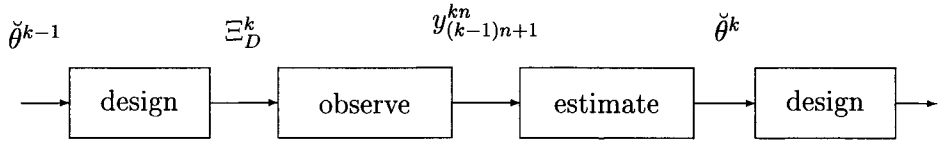

FIG. 2. *Batch sequential D–optimal design.*

We shall consider three alternatives for $\breve{\theta}^k$, to show the importance of this choice. Let $\tilde{\theta}^k$ denote the LS estimator based on the n observations associated with Ξ_D^k, that is

$$\tilde{\theta}^k = \arg\min_{\theta \in \mathbb{R}^p} \sum_{i=(k-1)n+1}^{kn} [y_i - \eta(\theta, \xi_i)]^2\,. \tag{3.1}$$

(i) A first choice, obviously not recommended, is $\breve{\theta}^k = \tilde{\theta}^k$: the design Ξ_D^k then does not converge since $\tilde{\theta}^k$ is based on n observations only, with n fixed.

(ii) A second choice, often used in practice, is

$$\breve{\theta}^k = \frac{1}{k} \sum_{i=1}^{k} \tilde{\theta}^i\,. \tag{3.2}$$

Let $\pi(\cdot|\bar{\theta}, \Xi)$ denote the density of the LS estimator for the design Ξ and true parameters $\bar{\theta}$. The estimator $\tilde{\theta}^k$ has the density $\pi[\hat{\theta}|\bar{\theta}, \Xi_D(\breve{\theta}^{k-1})]$, and one can show [Pronzato,

Walter and Kulcsár (1993)] that sequential design corresponds to a Robbins–Monro stochastic–approximation procedure for computing the value θ^* that satisfies

$$\theta^* = \int \hat{\theta}\, \pi[\hat{\theta}|\bar{\theta}, \Xi_D(\theta^*)]\, d\hat{\theta}.$$

Under rather general assumptions, $\breve{\theta}^k$ tends to θ^* and Ξ^k tends to $\Xi_D(\theta^*)$ when k tends to infinity. However, in nonlinear regression $\tilde{\theta}^k$ is not a sufficient statistic: each $\tilde{\theta}^k$ is biased, and this bias does not vanish asymptotically since $\tilde{\theta}^k$ is based on n observations only. Therefore the bias of $\breve{\theta}^k$ does not tend to zero, and θ^* differs from $\bar{\theta}$. A similar problem is encountered when the successive experiments are performed on different processes, or individuals, sampled from a population [Pronzato, Walter and Kulcsár (1993)], see D'Argenio (1981) for an example in pharmacokinetics.

EXAMPLE 2. Consider the case of a simple kinetic model (e.g. in chemical engineering or pharmacokinetics) with elimination rate constant θ, where $\eta(\theta, \xi) = \exp(-\theta\xi)$, $n = p = 1$, $\Xi_D(\theta) = 1/\theta$ and $\bar{\theta} = 0.5$. The measurement error ϵ_k for the experiment ξ_k has a normal density $\mathcal{N}(0, \sigma^2)$, truncated to values $\epsilon_k > -\exp(-\bar{\theta}\xi_k)$, so that $y_k > 0$ for all k. Figure 3 gives θ^* as a function of σ. For small values of σ the bias is small, and $\breve{\theta}^k$ converges to θ^* close to $\bar{\theta}$, whereas $\breve{\theta}^k$ converges to a value far from $\bar{\theta}$ when σ is large.

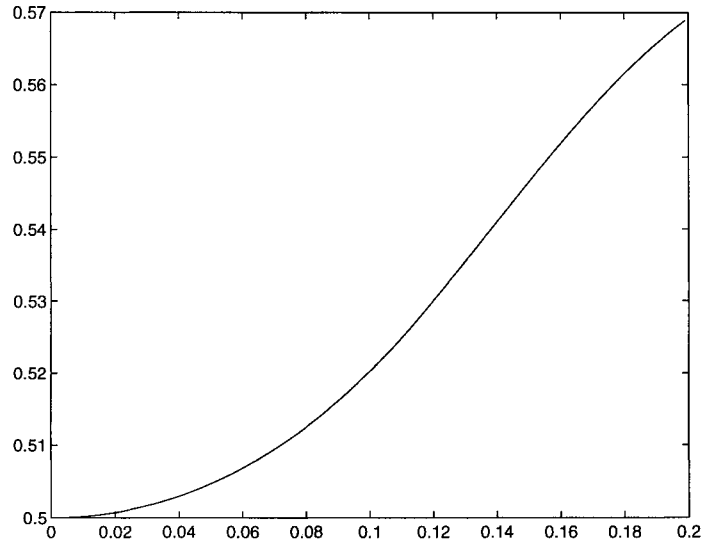

FIG. 3. *Bias in batch sequential design when (3.1,3.2) is used (θ^* as a function of σ, with $\bar{\theta} = 0.5$).*

(iii) A third choice corresponds to $\breve{\theta}^k = \hat{\theta}^{kn}$, that is the LS estimator based on all the observations collected. Since each Ξ_D^k is non–degenerate, $\breve{\theta}^k$ converges to $\bar{\theta}$ and Ξ_D^k converges to $\Xi_D(\bar{\theta})$ when k tends to infinity.

3.2. *Fully–sequential design.* Fully–sequential designs correspond to batch–sequential designs with $n = 1$ (only the case $p \geq 2$ is thus of interest). Consider the case of a

fully–sequential D–optimal design, see Figure 1, where for $k \geq p-1$ the design points are sequentially chosen as

$$\xi_{k+1} = \arg\max_{\xi}[\det \mathbf{M}(\hat{\theta}^k, [\xi_1^k, \xi])] = \arg\max_{\xi}[d(\hat{\theta}^k, \xi_1^k, \xi)]\,, \tag{3.3}$$

where

$$d(\theta, \xi_i^j, \xi) = \frac{\partial \eta(\theta,\xi)}{\partial \theta^T} \mathbf{M}^{-1}(\theta, \xi_i^j) \frac{\partial \eta(\theta,\xi)}{\partial \theta}\,. \tag{3.4}$$

This one–step ahead procedure corresponds to an adaptive version of Wynn's algorithm [Wynn (1970)] for the construction of a D–optimal design (approximate theory). Convergence results for such sequential schemes are still very partial. For instance, convergence is proved in Ford and Silvey (1980), Müller and Pötscher (1992) for estimation of the turning point of a quadratic regression: $\eta(\theta, \xi_k) = \theta_1 \xi_k + \theta_2 (\xi_k)^2 + \epsilon_k$, with the design criterion (to be minimized) $\Phi[\mathbf{M}(\xi_1^N)] = \mathbf{c}^T(\theta)\mathbf{M}^{-1}(\xi_1^N)\mathbf{c}(\theta)$, where $\mathbf{c}(\theta) = (1, -\theta_1/\theta_2)^T$ [$\Phi[\mathbf{M}(\xi_1^N)]$ is thus proportional to the asymptotic variance of $-\theta_1/(2\theta_2)$]. Convergence of fully–sequential designs for location and location–scale models is considered in Ying and Wu (1997).

To the best of our knowledge, no general result is available about the convergence of (3.3) in nonlinear regression but the following obvious one: if there is convergence to a non–degenerate design, then the parameter estimators converge to the true value $\bar{\theta}$, and the design obtained is thus D–optimal for $\bar{\theta}$. Note, however, that this convergence issue is not as crucial as it might seem. Indeed, one can always construct a non–degenerate design first, and then successively use its support points every M iterations of (3.3). The result above then applies, and the design which is constructed can be made arbitrarily close to the D–optimal design for $\bar{\theta}$ by choosing M large enough.

4. Sequential design for fixed N. All the design policies that have been considered so far substitute the current estimator $\hat{\theta}^k$ for the unknown value of $\bar{\theta}$, and the experiment is designed optimally for this estimator. This corresponds to *Forced Certainty Equivalence* (*FCE*), see Runggaldier (1993). In order to define new design policies, we first write optimal sequential design as a stochastic–control problem.

4.1. *Optimal sequential design and stochastic control.* Let the number of observations be fixed to N, and let $\mathcal{I}^0$ denote the prior information: θ has the density $\pi^0(\cdot)$ and $\{\epsilon_k\}$ is i.i.d. with the density $\pi_\epsilon(\cdot)$. We shall adopt a Bayesian approach (for design, *not for estimation*), and maximize $\mathrm{E}\{\Phi[\mathbf{M}(\theta, \xi_1^N)]|\mathcal{I}^0\}$ with respect to ξ_1^N. Note that ξ_k may depend on $\mathcal{I}^{k-1}$ given by $\mathcal{I}^{k-1} = \{\mathcal{I}^0, \xi_1^{k-1}, y_1^{k-1}\}$, which corresponds to the information available at time k. The expectation $\mathrm{E}\{\cdot\}$ is thus with respect to all random variables involved (θ and $\{\epsilon_k\}$). The optimal strategy is then the solution of the following stochastic dynamic–programming problem:

$$\begin{aligned} &\max_{\xi_1}[\mathrm{E}_{y_1}\{\max_{\xi_2}[\mathrm{E}_{y_2}\{\ldots \\ &\max_{\xi_{N-1}}[\mathrm{E}_{y_{N-1}}\{\max_{\xi_N}[\mathrm{E}_\theta\{\Phi[\mathbf{M}(\theta, \xi_1^N)]|\mathcal{I}^{N-1}\}]|\mathcal{I}^{N-2}\}]\ldots|\mathcal{I}^1\}]|\mathcal{I}^0\}]\,. \end{aligned} \tag{4.1}$$

It can be solved analytically in very special situations [for instance LQG control, see, e.g., Whittle (1982)], and its numerical solution is feasible only for N very small [$N = 2$

in Zacks (1977), Pronzato, Walter and Kulcsár (1993), Kulcsár, Pronzato and Walter (1994)]. Therefore, suboptimal solutions must be considered. Most of them use an approximation of the future posterior density $\pi(\theta|\mathcal{I}^{N-1})$ in (4.1). All the strategies to be considered are as follows. Suppose we are at time j, where $\mathcal{I}^{j-1}$ is available. We shall design ξ_j^N, then apply ξ_j, increment j by 1 and repeat.

4.2. Suboptimal passive strategies.

FORCED CERTAINTY EQUIVALENCE. *FCE* corresponds to *sequential local design*, and amounts to substituting the Dirac function $\delta(\theta - \hat{\theta}^{j-1})$ for $\pi(\theta|\mathcal{I}^{N-1})$ in (4.1), with $\hat{\theta}^{j-1}$ some estimate of θ, e.g. $\hat{\theta}^{j-1} = \arg\max_{\theta \in \mathbb{R}^p} \pi(\theta|\mathcal{I}^{j-1})$. Then $(\xi_j^N)_{FCE}$ is obtained by maximizing $\Phi[\mathbf{M}(\hat{\theta}^{j-1}, [\xi_1^{j-1}, \xi_j^N])]$, with ξ_1^{j-1} fixed to $(\xi_1^{j-1})_{FCE}$.

OPEN–LOOP FEEDBACK. *OLF* corresponds to *sequential Bayesian* (or *average–optimal* design), and amounts to substituting the current posterior density $\pi(\theta|\mathcal{I}^{j-1})$ for $\pi(\theta|\mathcal{I}^{N-1})$ in (4.1). Then $(\xi_j^N)_{OLF}$ is obtained by maximizing $\hat{\mathrm{E}}_\theta^j\{\Phi[\mathbf{M}(\theta, [\xi_1^{j-1}, \xi_j^N])]\}$, where $\xi_1^{j-1} = (\xi_1^{j-1})_{OLF}$ and $\hat{\mathrm{E}}_\theta^j\{\cdot\}$ denotes the expectation for $\pi(\theta|\mathcal{I}^{j-1})$.

OPEN–LOOP. The *OL* policy involves no feedback of information and is thus not sequential. In this case, $\pi(\theta|\mathcal{I}^{N-1})$ is simply replaced by the prior density $\pi(\theta|\mathcal{I}^0)$.

A control policy that yields better results than *OL* is called *quasi–adaptive* [Witsenhausen (1966)]. It is shown in Bertsekas (1976) that *OLF* is quasi–adaptive and in Thau and Witsenhausen (1966) that *FCE* is not.

When $\pi^0(\cdot)$ is the normal distribution $\mathcal{N}(\hat{\theta}^0, \Sigma_0)$ and $\pi_\epsilon(\cdot)$ is normal $\mathcal{N}(0, \sigma^2)$, $\pi(\theta|\mathcal{I}^{j-1})$ in *OLF* can be approximated by $\mathcal{N}(\hat{\theta}^{j-1}, \Sigma_{j-1}(\hat{\theta}^{j-1}))$, where

$$\Sigma_k(\theta) = [\Sigma_0^{-1} + \mathbf{M}(\theta, \xi_1^k)]^{-1} . \tag{4.2}$$

FCE is independent of σ^2. Since it ignores uncertainty on θ, it is also independent of Σ_0.

EXAMPLE 3. Consider the nonlinear regression model given by $\eta(\theta, \xi) = \theta_1 \exp(-\theta_2 \xi)$, with a normal prior $\mathcal{N}(\hat{\theta}^0, \Sigma_0)$ for θ, $\hat{\theta}^0 = (1, 0.5)^T$, $\Sigma_0 = \mathrm{diag}(0.04, 0.01)$, i.i.d. errors $\mathcal{N}(0, 0.015)$ and $N = 4$. We compare the *OLF* and *FCE* strategies for $\Phi(\cdot) = \det(\cdot)$ by performing M repetitions of the sequential experiments, as explained in the Appendix. Figure 4 gives the evolution of $t_{M-1,0.05}$ and $\rho_{OLF-FCE}(M)$ given by (A.1) as functions of M. We can conclude from the figure that *OLF* performs significantly better than *FCE*.

In *FCE* and *OLF* at any time j all support points in ξ_j^N play the same role. The sequential character of the experiment is thus ignored: the fact that at a future time $k > j$, $\mathcal{I}^{k-1}$ will differ from $\mathcal{I}^{j-1}$ is not taken into account. In other words, *FCE* and *OLF* use feedback but are *passively adaptive*: decisions are taken as if no other observation would take place in the future. This is the reason why these policies are generally implemented in a fully–sequential form, with N not fixed *a priori*. In the case of D-optimality, this gives the following one–step ahead myopic procedures:

$$\xi_{FCE}^j = \arg\max_\xi d(\hat{\theta}^{j-1}, (\xi_1^{j-1})_{FCE}, \xi) ,$$

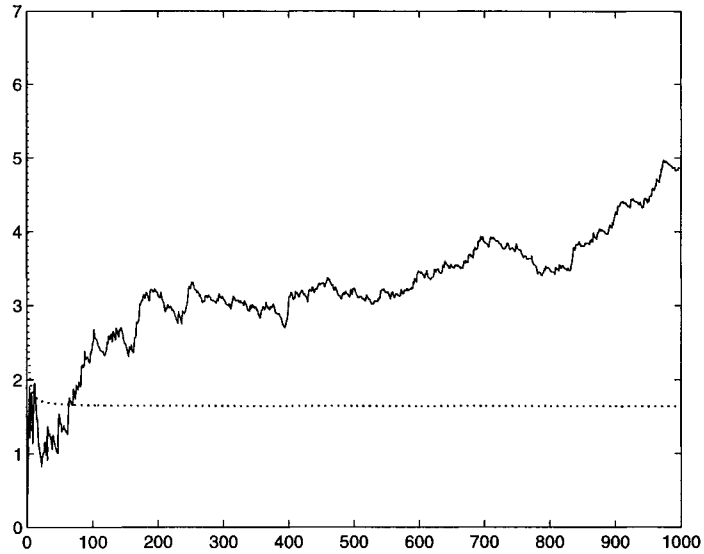

FIG. 4. *Evolution of* $\rho_{OLF-FCE}(M)$ *(full line) and* $t_{M-1,0.05}$ *(dotted line) as functions of* M.

$$\xi_{OLF}^{j} = \arg\max_{\xi} E_{\theta}\{d(\theta,(\xi_1^{j-1})_{OLF},\xi)|\mathcal{I}^{j-1}\}\,.$$

The opposite to this is the idea of *active control*: a control policy is said to be *actively adaptive* if it takes its influence on future uncertainty into account. New sequential strategies based on active control will now be presented.

4.3. *Suboptimal active strategies.* A first suboptimal active strategy, denoted $A1$ in what follows, consists in approximating at time j the future posterior $\pi(\theta|\mathcal{I}^{N-1})$ by the normal density $\mathcal{N}(\hat{\theta}^{j-1},\Sigma_{N-1}(\hat{\theta}^{j-1}))$, where $\Sigma_{N-1}(\hat{\theta}^{j-1})$ is given by (4.2) and does not depend on y_j^{N-1}. $(\xi_j^N)_{A1}$ is then obtained by maximizing $\hat{\mathrm{E}}_{\theta}^{j,N-1}\{\Phi[\mathbf{M}(\theta,[\xi_1^{j-1},\xi_j^N])]\}$, where $\xi_1^{j-1}=(\xi_1^{j-1})_{A1}$ and $\hat{\mathrm{E}}_{\theta}^{j,N-1}\{\cdot\}$ denotes the expectation for θ distributed

$$\mathcal{N}(\hat{\theta}^{j-1},\Sigma_{N-1}(\hat{\theta}^{j-1}))\,.$$

Note that when σ^2 tends to zero, this distribution tends to the Dirac function $\delta(\theta-\hat{\theta}^{j-1})$ and $A1$ tends to FCE.

Applying the same idea to an additive decomposition of the terminal cost $\Phi[\mathbf{M}(\theta,\xi_1^N)]$ we obtain another active strategy. This is illustrated in the case where $\Phi(\cdot)=\log\det(\cdot)$. We can write

$$\begin{aligned}\log\det[\mathbf{M}(\theta,\xi_1^N)] &= \log[\sigma^2+d(\theta,\xi_1^{N-1},\xi_N)]+\log[\sigma^2+d(\theta,\xi_1^{N-2},\xi_{N-1})]\\ &\quad+\cdots+\log[\sigma^2+d(\theta,\xi_1^{p},\xi_{p+1})]+\log\det[\mathbf{M}(\theta,\xi_1^{p})]\\ &\quad-(N-p)\log(\sigma^2)\,,\end{aligned}$$

where $d(\theta,\xi_i^j,\xi)$ is given by (3.4). The dynamic–programming formulation of the design

problem then becomes:

$$\max_{\xi_1}[\mathrm{E}_{y_1}\{\max_{\xi_2}[\mathrm{E}_{y_2}\{\dots\max_{\xi_p}[\mathrm{E}_{y_p}\{\mathrm{E}_\theta\{\log\det[\mathbf{M}(\theta,\xi_1^p)]|\mathcal{I}^p\}+$$
$$\max_{\xi_{p+1}}[\mathrm{E}_{y_{p+1}}\{\mathrm{E}_\theta\{\log[\sigma^2+d(\theta,\xi_1^p,\xi_{p+1})]|\mathcal{I}^{p+1}\}+\dots+$$
$$\max_{\xi_{N-1}}[\mathrm{E}_{y_{N-1}}\{\mathrm{E}_\theta\{\log[\sigma^2+d(\theta,\xi_1^{N-2},\xi_{N-1})]|\mathcal{I}^{N-1}\}+$$
$$\max_{\xi_N}[\mathrm{E}_\theta\{\log[\sigma^2+d(\theta,\xi_1^{N-1},\xi_N)]|\mathcal{I}^{N-1}\}]|\mathcal{I}^{N-2}\}]\dots|\mathcal{I}^p\}]|\mathcal{I}^{p-1}\}]\dots|\mathcal{I}^1\}]|\mathcal{I}^0\}]\,.$$

The approximations of future posterior densities used for FCE and OLF yield the same strategies as in the previous section, whereas approximating at time j the future posterior density $\pi(\theta|\mathcal{I}^k)$, $k\geq j$, by $\mathcal{N}(\hat\theta^{j-1},\Sigma_k(\hat\theta^{j-1}))$ gives a new strategy, denoted $A2$ in what follows. Since $\mathrm{E}_{y_k}\{\mathrm{E}_\theta\{f(\theta,\mathcal{I}^{k-1})|\mathcal{I}^k\}|\mathcal{I}^{k-1}\}=\mathrm{E}_\theta\{f(\theta,\mathcal{I}^{k-1})|\mathcal{I}^{k-1}\}$, $(\xi_j^N)_{A2}$ is obtained at time j by maximizing

$$\hat{\mathrm{E}}_\theta^{j,p-1}\{\log\det[\mathbf{M}(\theta,\xi_1^p)]\}+\hat{\mathrm{E}}_\theta^{j,p}\{\log[\sigma^2+d(\theta,\xi_1^p,\xi_{p+1})]\}+\dots+$$
$$\hat{\mathrm{E}}_\theta^{j,N-2}\{\log[\sigma^2+d(\theta,\xi_1^{N-2},\xi_{N-1})]\}+\hat{\mathrm{E}}_\theta^{j,N-1}\{\log[\sigma^2+d(\theta,\xi_1^{N-1},\xi_N)]\}\,,$$

where $\xi_1^{j-1}=(\xi_1^{j-1})_{A2}$ and $\hat{\mathrm{E}}_\theta^{j,k}\{\cdot\}$ denotes the expectation for $\mathcal{N}(\hat\theta^{j-1},\Sigma_k(\hat\theta^{j-1}))$.

The idea used in methods $A1$ and $A2$ has been successfully applied to standard control problems, such as regulation, trajectory following and target attainment, for which better performances than FCE or OLF were obtained. For instance, in the case of Example 1, one may wish to minimize $\mathrm{E}\{\sum_{k=1}^N[\eta(\theta,k)-\mathrm{ref}(k)]^2+\sum_{k=1}^N u_k^2\}$ [see Pronzato, Kulcsár and Walter (1996), Kulcsár, Pronzato and Walter (1996)]. The reason for good performance is that active policies force the input to excite the system, which helps in estimating θ. This is not so crucial in sequential design, since the input (i.e. the design) is always chosen so as to excite the system. We shall thus consider a third active strategy, denoted $A3$, which at time j takes the dependence of $\pi(\theta|\mathcal{I}^k)$ in y_j^k into account. The idea is as follows.

Assume that σ^2 tends to zero. Then at time j, if ξ_j^{j+p-1} is non–degenerate, for $k\geq p$ $\pi(\theta|\mathcal{I}^k)$ tends to the Dirac function $\delta(\theta-\hat\theta^{j+p-1})$. Therefore, the problem to solve becomes:

$$\max_{\xi_j}[\mathrm{E}_{y_j}\{\dots\max_{\xi_{j+p-1}}[\mathrm{E}_{y_{j+p-1}}\{\max_{\xi_{j+p}^N}[\Phi[\mathbf{M}(\hat\theta^{j+p-1},\xi_1^N)]]|\mathcal{I}^{j+p-2}\}]\dots|\mathcal{I}^{j-1}\}]\,. \tag{4.3}$$

REMARK 2. $A3$ tends to be optimal if $N=p$ when σ^2 tends to zero. This strategy is closely related to *p–measurement feedback control* [Curry (1969)], which has better performances than OLF [Bayard (1987)]. For this policy, at time j the next p steps are performed in closed loop, and the rest with OLF, that is $\mathrm{E}_\theta\{\Phi[\mathbf{M}(\theta,\xi_1^N)]|\mathcal{I}^{j+p-1}\}$ is substituted for $\Phi[\mathbf{M}(\hat\theta^{j+p-1},\xi_1^N)]$ in (4.3).

REMARK 3. Problem (4.3) is much simpler to solve than (4.1), but still involves p imbedded expectations and maximizations. To facilitate the calculations, at time j one can approximate $\hat\theta^k$, $k\geq j$, through the recursion:

$$\hat\theta^k=\hat\theta^{k-1}+\frac{\Sigma_{k-1}(\hat\theta^{j-1})\mathbf{s}(\hat\theta^{j-1},\xi_k)}{\left(\sigma^2+\mathbf{s}^T(\hat\theta^{j-1},\xi_k)\Sigma_{k-1}(\hat\theta^{j-1})\mathbf{s}(\hat\theta^{j-1},\xi_k)\right)^{1/2}}\omega_k\,,$$

with $\mathbf{s}(\theta,\xi) = \partial\eta(\theta,\xi)/\partial\theta$ and ω_k distributed $\mathcal{N}(0,1)$. The expectation with respect to y_k in (4.3) is then replaced by an expectation with respect to ω_k.

To reveal the active character of strategies $A1$, $A2$ and $A3$ we introduce constraints between design points, by considering the case where ξ is time. The sequential construction of the design then implies the time constraint $\xi_{k+1} \geq \xi_k$ for any k. Since all design points play the same role in passive strategies, they can be permuted without changing the value of the criterion. The time constraint has thus no other effect than restricting the design space to $[\xi_{j-1}, \infty)$ at step j. On the other hand, active strategies are affected by the time constraint. In $A1$ the influence of ξ_N is different from that of the other ξ_i's; in $A2$ the ξ_i's for $i \geq p$ play different roles in the criterion, whereas in $A3$, at step j, the roles of the ξ_i's with $j \leq i \leq j+p-1$ are different.

EXAMPLE 4. Consider the model $\eta(\theta,\xi) = \exp(-\theta\xi)$, with measurement errors $\{\epsilon_k\}$ i.i.d. $\mathcal{N}(0,\sigma^2)$ and a normal prior $\mathcal{N}(0.5, 0.01)$ for θ. We use $\Phi[\mathbf{M}(\theta,\xi_1^N)] = \det\mathbf{M}(\theta,\xi_1^N) = \sum_{i=1}^N \xi^2/\sigma^2 \exp(-2\theta\xi)$. This additive form makes it possible to use $A2$. At step 1, FCE and OLF can be calculated analytically: $(\xi_1)_{FCE} = 2$ and $(\xi_1)_{OLF} = 2.1922$. We consider the case $N = 2$, for which $A1$ and $A2$ are computed without numerical integration, and the computation of $A3$ only requires one numerical integration. We also compute the strategy which is optimal up to a linearization of the model response (to get a normal posterior for θ). Figure 5 gives the evolution of the first sampling time ξ_1 as a function of σ for the different strategies when there is no time constraint. The optimal strategy and $A2$ remain close to OLF, which is independent of the value of σ. $A1$ (resp. $A3$) tends to FCE (resp. OLF) when σ tends to zero. Figure 6 corresponds to the case where time constraints are present. FCE and OLF stay the same as previously, whereas the first sampling time for the optimal strategy, $A2$, and for $A3$ is shifted towards smaller values: the strategies are thus cautious, in the sense that performing the first observation too late would imply performing the second one even later, which might then yield bad performance if the value of $\bar{\theta}$ happens to be large. Note that in Figures 5 and 6, $A3$ tends to the optimal strategy when σ tends to zero, but quickly moves away from optimality as σ increases.

We compared the performances of the suboptimal strategies by performing 1000 repetitions of the sequential experiments, as explained in the Appendix, for $\sigma = 0.1225$ and $N = 4$. The results are as follows, where $\leq$ means that the difference was not significant at level 0.05 and $\approx$ means that the performances were similar:

without time constraint:	$FCE < A3 < A1 < OLF \approx A2$
with time constraint:	$A3 \leq OLF \leq FCE < A1 < A2$

These results confirm in particular that the performance of $A3$ may quickly deteriorate as σ increases.

Finally, Figure 7 presents the evolution of ξ_1 as a function of N for different strategies, when $\sigma = 0.02$ and time constraints are present. FCE and OLF do not depend on N and σ^2. The optimal strategy is not computed for $N > 2$ due to the complexity of (4.1). Note that, as intuition would suggest, active strategies tend to be more cautious as N increases. $A1$ remains close to FCE due to the small value of σ.

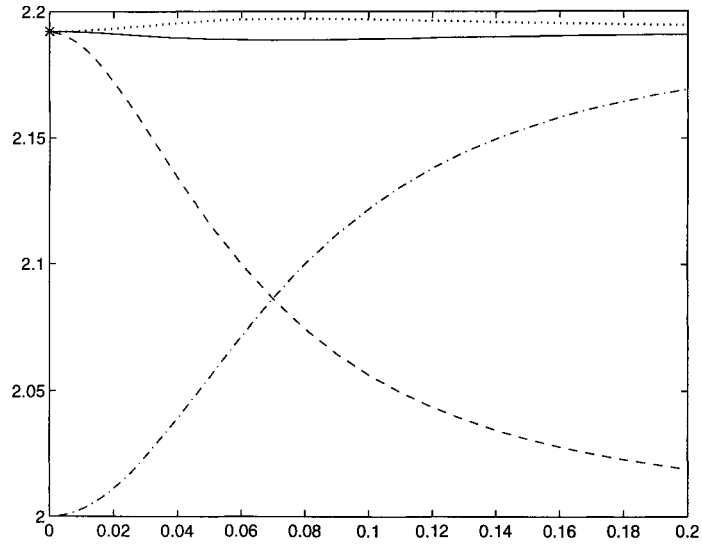

FIG. 5. ξ_1 *as a function of* σ *for different strategies without time constraint when* $N = 2$. *A*1 *is in dash–dotted line, A*2 *in dotted line, A*3 *in dashed line, the optimal strategy is in full line,* $(\xi_1)_{FCE} = 2$, $(\xi_1)_{OLF}$ *is indicated by a star.*

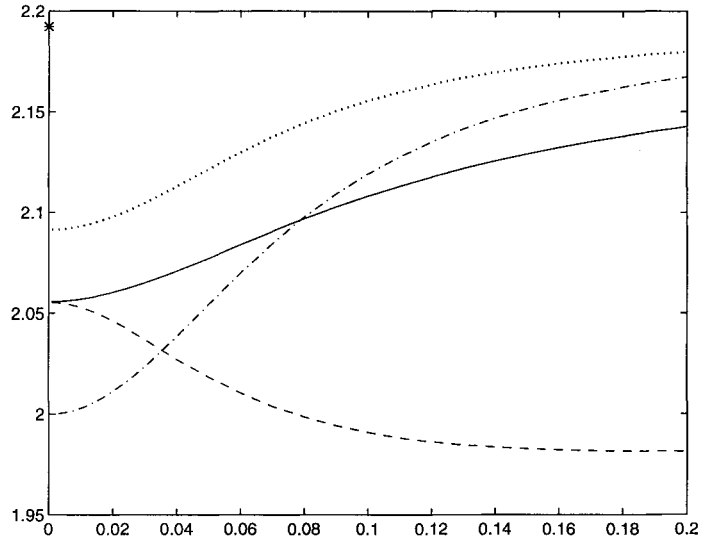

FIG. 6. ξ_1 *as a function of* σ *for different strategies with the time constraint* $\xi_2 \geq \xi_1$ *when* $N = 2$. *A*1 *is in dash–dotted line, A*2 *in dotted line, A*3 *in dashed line, the optimal strategy is in full line,* $(\xi_1)_{FCE} = 2$, $(\xi_1)_{OLF}$ *is indicated by a star.*

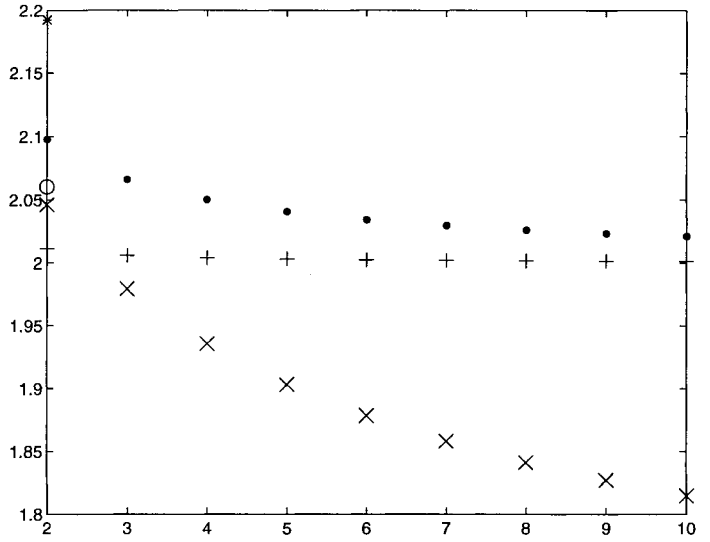

FIG. 7. ξ_1 *as a function of* N *for different strategies with the time constraint* $\xi_k \geq \xi_{k-1}$ *and* $\sigma = 0.02$. *(A1:* $+$, *A2:* $\bullet$, *A3:* $\times$, $(\xi_1)_{OLF}$*:* $*$, *optimal strategy for* $N = 2$*:* $\circ$, $(\xi_1)_{FCE} = 2$*).*

5. Conclusions. Designing experiments sequentially seems a very natural approach in nonlinear situations and is more and more used due to the increasing presence of computers on sites where the experiments are conducted.

Convergence is generally a difficult issue in sequential strategies, but we indicated why it is mainly of theoretical interest: one can always introduce periodically the support points of a nondegenerate design in the sequential procedure to ensure consistency of the estimator. Batch–sequential design is easier to analyse than fully–sequential design: convergence is generally achieved provided a consistent estimator is used, see Section 3.1.

Various methods have been considered for finite N, with a distinction between passive and active strategies. Passive strategies ignore the future. Open–loop feedback is preferable to forced certainty equivalence, but requires integrations with respect to the posterior density of θ. Active strategies take into account the fact that future observations will take place. The support points play different roles in the criteria since the information available changes with time. Active strategies $A1$ and $A2$ rely on a normal approximation of the posterior of θ which is independent of future observations. Integrations with respect to this approximated posteriors are required. Strategy $A3$ relies on the assumption that the variance σ^2 of the disturbances is small. Only one–dimensional integrals need to be computed, but integrations are imbedded within maximizations. Numerical simulations with $N = 2$ have shown that the strategy tends to be optimal when σ tends to zero. However, other suboptimal strategies perform better when σ is not negligible.

Further studies are required to compare these strategies, both from a theoretical

and practical point of view, and to study the convergence properties of the new active strategies proposed. Combining strategy $A2$ with $A3$, or with m–measurement feedback control, seems promising. Approximating integrals (e.g. using Laplace approximation) might be useful to reduce the computational task. Also, stochastic approximation methods might help in cases where expectations and maximizations are imbedded [see Bayard (1988)]. Finally, other criteria than $\Phi(\cdot) = \det(\cdot)$ might be considered. In particular, $\mathrm{E}\{\log\det(\cdot)\}$ is more justified than $\mathrm{E}\{\det(\cdot)\}$ [see Chaloner and Verdinelli (1995)].

Appendix. Suppose that sequential design strategies $S1$ and $S2$ have to be compared, with $\pi^0(\cdot)$ the prior distribution for θ, $\pi_\epsilon(\cdot)$ the distribution for the i.i.d. disturbances $\{\epsilon_k\}$ and N the number of observations.

We perform M independent repetitions of the sequential experiments. For each experiment, say the i-th one, a value $\bar{\theta}^i$ is generated according to $\pi^0(\cdot)$ and N independent disturbances $\{\epsilon_k\}^i_{k=1,\dots,N}$ are generated according to $\pi_\epsilon(\cdot)$. Let $(\xi_1^N)_1^i$ (resp. $(\xi_1^N)_2^i$) denote the design points produced by the strategy $S1$ (resp. $S2$). We compute the differences of performances between strategies S_1 and S_2 in the i-th experiment, that is

$$\Delta^i_{S1-S2} = \Phi[\mathbf{M}(\bar{\theta}^i, (\xi_1^N)_1^i)] - \Phi[\mathbf{M}(\bar{\theta}^i, (\xi_1^N)_2^i)]\,.$$

We thus obtain M independent realizations of Δ^i_{S1-S2}, and compute

$$\bar{\Delta}_{S1-S2}(M) = \frac{1}{M}\sum_{i=1}^{M}\Delta^i_{S1-S2}\,, \quad V_{S1-S2}(M) = \frac{1}{M-1}\sum_{i=1}^{M}(\Delta^i_{S1-S2} - \bar{\Delta}_{S1-S2})^2\,.$$

For large M, $\bar{\Delta}_{S1-S2}(M)$ tends to be normally distributed, and one can use the t-test to test if $E\{\Delta^i_{S1-S2}\} > 0$. We then say that $S1$ performs significantly better than $S2$ at level α if $\rho_{S1-S2}(M) > t_{M-1,\alpha}$, where $t_{M-1,\alpha}$ has probability α of being exceeded by a random variable with Student's t-distribution with $M-1$ degrees of freedom, and

$$\rho_{S1-S2}(M) = \sqrt{M}\frac{\bar{\Delta}_{S1-S2}(M)}{\sqrt{V_{S1-S2}(M)}}\,. \tag{A.1}$$

REFERENCES

BAYARD, D. S. (1987). Proof of the quasi-adaptivity for the m-measurement feedback class of stochastic control policies. *IEEE Trans. Autom. Cont.* **32** 447–451.

BAYARD, D. S. (1988). A forward method for optimal stochastic nonlinear and adaptive control. In *Proc. 27th Conf. on Decision and Control*, 280–285. IEEE, Piscataway.

BELLMAN, R. (1957). *Dynamic Programming.* Princeton University Press, Princeton.

BERTSEKAS, D. M. (1976). *Dynamic Programming and Stochastic Control.* Academic Press, New York.

CHALONER, K. AND VERDINELLI, I. (1995). Bayesian experimental design: a review. *Statist. Sci.* **10** 273–304.

CHERNOFF, H. (1972). *Sequential Analysis and Optimal Design.* SIAM, Philadelphia.

CURRY, R. E. (1969). A new algorithm for suboptimal stochastic control. *IEEE Trans. Autom. Cont.* **14** 533–536.

D'ARGENIO, D. Z. (1981). Optimal sampling times for pharmacokinetic experiments. *J. Pharm. Biopharm.* **9** 739–756.

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

FORD, I. AND SILVEY, S. D. (1980). A sequentially constructed design for estimating a nonlinear parametric function. *Biometrika* **67** 381–388.

FORD, I., TITTERINGTON, D. M. AND WU, C. F. J. (1985). Inference and sequential design. *Biometrika* **72** 545–551.

GOODWIN, G. C. AND PAYNE, R. L. (1977). *Dynamic System Identification: Experiment Design and Data Analysis.* Academic Press, New York.

JENNRICH, R. L. (1969). Asymptotic properties of nonlinear least squares estimation. *Ann. Math. Statist.* **40** 633–643.

KULCSÁR, C., PRONZATO, L. AND WALTER, L. (1994). Optimal experimental design and therapeutic drug monitoring. *Int. J. Biomed. Computing* **36** 95–101.

KULCSÁR, C., PRONZATO, L. AND WALTER, E. (1996). Dual control of linearly parameterised models via prediction of posterior densities. *European J. Control* **2** 135–143.

MÜLLER, W. G. AND PÖTSCHER, B. M. (1992). Batch sequential design for a nonlinear estimation problem. In *Model-Oriented Data Analysis II, Proceedings 2nd IIASA Workshop* (Fedorov, V. V., Müller, W. G. and Vuchkov, I. N., eds.), 76–87. Physica Verlag, Heidelberg.

PÁZMAN, A. AND PRONZATO, L. (1992). Nonlinear experimental design based on the distribution of estimators. *J. Statist. Plann. and Inf.* **33** 385–402.

PRONZATO, L., KULCSÁR, C. AND WALTER, E. (1996). An actively adaptive control policy for linear models. *IEEE Trans. Autom. Cont.*, **41** 855–858.

PRONZATO, L. AND PÁZMAN, A. (1994). Second-order approximation of the entropy in nonlinear least-squares estimation. *Kybernetika* **30** 187–198 (Erratum, **32**, p. 104, 1996).

PRONZATO, L., WALTER, E. AND KULCSÁR, C. (1993). A dynamical-system approach to sequential design. In *Model-Oriented Data Analysis III, Proceedings MODA3* (Müller, W. G., Wynn, H. P. and Zhigljavsky, A. A., eds.), 11–24. Physica Verlag, Heidelberg.

RUNGGALDIER, W. J. (1993). Concepts of optimality in stochastic control. In *Reliability and Decision Making* (Barlow, R. *et al.*, eds.), 101–114. Elsevier, Amsterdam.

THAU, F. E., AND WITSENHAUSEN, H. S. (1966). A comparison of closed-loop and open-loop optimum systems. *IEEE Trans. Autom. Cont.* **11** 619–621.

WALTER, E. AND PRONZATO, L. (1997). *Identification of Parametric Models from Experimental Data.* Springer, Heidelberg.

WHITTLE, P. (1982). *Optimization Over Time*, vol. 1. Wiley, Chichester.

WITSENHAUSEN, H. S. (1966). *Minimax Control of Uncertain Systems.* Massachusetts Institute of Technology, Cambridge, MA (doctoral dissertation).

WU, C. F. J. (1985). Asymptotic inference from sequential design in a nonlinear situation. *Biometrika* **72** 553–558.

WYNN, H. P. (1970). The sequential generation of D-optimum experimental designs. *Ann. of Math. Statist.* **41** 1655–1664.

YING, Z. AND WU, C. F. J. (1997). An asymptotic theory of sequential designs based on maximum likelihood recursions. *Statist. Sinica* **7** 75–91.

ZACKS, S. (1977). Problems and approaches in design of experiments for estimation and testing in nonlinear models. In *Multivariate Analysis IV* (Krishnaiah, P. ed.), 209–223. North Holland, Amsterdam.

ZARROP, M. B. (1979). *Optimal Experiment Design for Dynamic System Identification.* Springer, Heidelberg.

LABORATOIRE I3S,
CNRS/UPRES–A 6070, BÂT. 4,
250 RUE A. EINSTEIN,
06560 VALBONNE
FRANCE