Adaptive design in regression and control

(iterated least squares/adaptive stochastic approximation/nonlinear regression/control theory/optimal dosage estimation)

T. L. LAI AND HERBERT ROBBINS

Department of Mathematical Statistics, Columbia University, New York, New York 10027

Contributed by Herbert Robbins, November 18, 1977

ABSTRACT When $y = M(x) + \epsilon$, where M may be nonlinear, adaptive regression designs of the levels $x_1, x_2, \dots$ at which $y_1, y_2, \dots$ are observed lead to asymptotically efficient estimates of the value θ of x for which $M(\theta)$ is equal to any desired value y^* . More importantly, these designs also make the "cost" of the observations, defined at the n th stage to be $\sum_1^n (x_i - \theta)^2$, to be of the order of $\log n$ instead of n , an obvious advantage in medical and other applications.

We shall consider a general, not necessarily linear, regression model

$$y_i = M(x_i) + \epsilon_i \quad (i = 1, 2, \dots), \quad [1]$$

where the errors ϵ_i are independent, identically distributed random variables with mean 0 and variance σ^2 . The adaptive designs for choosing the levels of x at which y is to be observed are motivated by applications of the following nature. Suppose that x_i is the dosage level of some drug and that y_i is the response value of the i th patient who turns up for treatment. Suppose also that an optimal response value y^* is known. If the function M were known, then the dosage level x could be kept at the value θ such that $M(\theta) = y^*$. But if M is unknown, how can the dosage levels $x_1, x_2, \dots$ be chosen so as to approach the unknown θ as rapidly as possible?

We take $\sum_1^n (x_i - \theta)^2$ to be the (cumulative) cost at stage n . To fix ideas, consider the linear model $M(x) = \alpha + \beta x$ with $\beta > 0$ and α, β unknown, and suppose without loss of generality that the optimal response value y^* is 0. Then $\theta = -\alpha/\beta$, and if we assign the values $x_1, x_2, \dots$ in advance and estimate θ by $\hat{\theta}_n = -\hat{\alpha}_n/\hat{\beta}_n$ after observing $x_1, y_1, \dots, x_n, y_n$, where $\hat{\alpha}_n, \hat{\beta}_n$ are the usual least squares estimates of α, β , under mild regularity assumptions it can be shown that for large n , $\hat{\theta}_n$ is approximately normal with mean θ and variance

$$\frac{\sigma^2}{n\beta^2} \left[1 + \frac{(\bar{x}_n - \theta)^2}{\sum_1^n (x_i - \bar{x}_n)^2/n} \right] \quad [2]$$

where the notation $\bar{x}_n$ for the arithmetic mean of any n quantities $x_1, \dots, x_n$ will be used throughout the sequel. In ignorance of θ , to make the square bracket in Eq. 2 close to its minimal value 1 we would have to choose the x_i so as to make $\sum_1^n (x_i - \bar{x}_n)^2$, and hence *a fortiori* $\sum_1^n (x_i - \theta)^2 = \text{cost}$, of order larger than n . The results of this paper show that, by a suitable adaptive choice of $x_1, \dots, x_n$, the unknown θ can be estimated with an asymptotic variance of $\sigma^2/n\beta^2$ but with a cost of order only $(\sigma^2/\beta^2) \log n$. Thus, the apparent dilemma of having to choose between a good final estimate of θ (of interest perhaps to future

patients) and a small cost $\sum_1^n (x_i - \theta)^2$ (of interest to the current patients) is happily resolved for large n .

In an adaptive design, the choice of each level x_i depends on all the previous data values $x_1, y_1, \dots, x_{i-1}, y_{i-1}$. In the linear model $M(x) = \alpha + \beta x$ with desired mean response 0 and $\theta = -\alpha/\beta$, the general model 1 can be written as

$$y_i = \beta(x_i - \theta) + \epsilon_i. \quad [3]$$

A valuable hint toward the construction of an adaptive design in the general case comes from considering the special situation that arises when the true value of β in Eq. 3 is known. Then at stage n the least squares estimate of θ becomes

$$\theta_n^* = \bar{x}_n - (\bar{y}_n/\beta) = \theta - (\bar{\epsilon}_n/\beta). \quad [4]$$

In this case, *no matter how the levels $x_1, \dots, x_n$ are chosen*, whether fixed in advance or by an adaptive design, Eq. 4 shows that $E(\theta_n^* - \theta)^2 = \sigma^2/(n\beta^2)$ and

$$\sqrt{n}(\theta_n^* - \theta) \rightarrow N(0, \sigma^2/\beta^2) \text{ in distribution as } n \rightarrow \infty. \quad [5]$$

Now, to minimize the cost, we want each x_n to be as close as possible to θ , and since our estimate of θ at stage n is θ_n^* , it is natural to choose the next level x_{n+1} to be precisely θ_n^* . This amounts to the following adaptive design. Let x_1 (the initial best guess of θ) be a random variable with finite second moment, and define inductively for $n = 1, 2, \dots$

$$x_{n+1} = \theta_n^* = \bar{x}_n - (\bar{y}_n/\beta). \quad [6]$$

Then Eqs. 4 and 5 still hold, while from Eqs. 4 and 6

$$\begin{aligned} E \left(\sum_1^n (x_i - \theta)^2 \right) &= E(x_1 - \theta)^2 + \sum_1^{n-1} E(\theta_i^* - \theta)^2 \\ &= E(x_1 - \theta)^2 + \frac{\sigma^2}{\beta^2} \left[1 + \frac{1}{2} + \dots + \frac{1}{n-1} \right] \\ &= (\sigma^2/\beta^2) \log n + O(1). \end{aligned} \quad [7]$$

The desirable properties 5 and 7 of the adaptive design 6 have been obtained under the restrictive assumption of a linear model 3 with β known. The following theorem shows that similar properties still hold for 6 when the regression function M is not linear but the slope $M'(\theta)$ at the unique root of the equation $M(x) = 0$ is known to be some positive constant β .

THEOREM 1. Let β, θ , and the Borel function M be such that

$$M(\theta) = 0, M'(\theta) = \beta > 0 \quad [8]$$

$$\inf_{\delta \leq |x - \theta| \leq 1/\delta} [M(x)(x - \theta)] > 0 \text{ for all } 0 < \delta < 1 \quad [9]$$

$$|M(x)| \leq c|x| + d \text{ for some } c, d > 0 \text{ and all } x. \quad [10]$$

The costs of publication of this article were defrayed in part by the payment of page charges. This article must therefore be hereby marked "advertisement" in accordance with 18 U. S. C. §1734 solely to indicate this fact.

Let x_1 be any random variable independent of $\epsilon_1, \epsilon_2, \dots$. Define inductively y_i by 1 and x_{n+1} by 6. Then as $n \rightarrow \infty$

$$\sqrt{n}(x_n - \theta) \rightarrow N(0, \sigma^2/\beta^2) \text{ in distribution} \quad [11]$$

$P(\lim x_n = \theta) = 1$, and in fact

$$P\left(\lim \sqrt{\frac{n}{2 \log \log n}}(x_n - \theta) = \sigma/\beta\right) = 1 \quad [12]$$

$$P\left(\lim \left(\sum_1^n (x_i - \theta)^2 / \log n\right) = \sigma^2/\beta^2\right) = 1. \quad [13]$$

Since in practice the value $M'(\theta) = \beta$ will rarely be known, it is natural to try replacing 6 by

$$x_{n+1} = \bar{x}_n - (\bar{y}_n/b_n), \quad [14]$$

where $b_n = b_n(x_1, y_1, \dots, x_n, y_n)$ is some estimate of β based on the data observed by stage n . If b_n is a consistent estimate of β (no matter how inefficient), we may hope that the asymptotic properties 11, 12, and 13 of Theorem 1 will still hold. An obvious choice for b_n is the usual least squares estimate

$$\hat{\beta}_n = \frac{\sum_1^n (x_i - \bar{x}_n)(y_i - \bar{y}_n)}{\sum_1^n (x_i - \bar{x}_n)^2} \quad [15]$$

at least in the linear case $M(x) = \beta(x - \theta)$. But when the x_i are sequentially determined random variables it is not evident that $\hat{\beta}_n$ will in fact be a consistent estimate of β . [Even for the case of a fixed design $x_1, x_2, \dots$ the strong consistency of $\hat{\beta}_n$ under the sole condition that $\sum_1^n (x_i - \bar{x}_n)^2 \rightarrow \infty$ has only recently been established (1).] However, a modified version of $\hat{\beta}_n$, asymptotically equal to it but somewhat more stable in its initial stages, has the desired properties. This is the content of

THEOREM 2. Suppose that, in addition to 8, 9, and 10

$$\lim |M(x)/x| > 0 \text{ as } |x| \rightarrow \infty, \quad [16]$$

and that positive constants b and B are known such that $b \leq \beta = M'(\theta) \leq B$. Define inductively y_i by 1 and x_{n+1} by 14, where the b_n are defined as follows. Let b_1 be any constant between b and B , and let c_n be any sequence of positive constants such that

$$\lim c_n > 0, \quad \overline{\lim} c_n < 1. \quad [17]$$

Define $s_n^2 = \sum_1^n (y_i - \bar{y}_n)^2/n$, and for $n \geq 2$ if

$$\sum_1^n (x_i - \bar{x}_n)^2 \leq (c_n s_n^2 / B^2) \cdot \log n \quad [18]$$

set $b_n = b_{n-1}$; otherwise set b_n equal to $\hat{\beta}_n$ truncated below by b and above by B . Then 11, 12, and 13 still hold and with

probability 1 as $n \rightarrow \infty$,

$$\lim b_n = \beta, \quad \lim s_n^2 = \sigma^2. \quad [19]$$

If, furthermore, for some $r > 0$

$$M(x) = \beta(x - \theta) + O(|x - \theta|^{1+r}) \text{ as } x \rightarrow \theta \quad [20]$$

then

$$(\log n)^{1/2}(b_n - \beta) \rightarrow N(0, \beta^2) \text{ in distribution as } n \rightarrow \infty. \quad [21]$$

We call the design of Theorem 2 iterated least squares modified by using b_n instead of $\hat{\beta}_n$ in 14. We now discuss a different design which we shall call adaptive stochastic approximation. Going back to 6, it is easy to show that, for any constant $b > 0$, the two sequences

$$x_{n+1} = \bar{x}_n - (\bar{y}_n/b), \quad x_{n+1} = x_n - (y_n/nb) \quad [22]$$

are identical. This suggests using, instead of 14, the design

$$x_{n+1} = x_n - (y_n/nb_n) \quad [23]$$

where again b_n is an estimate of β . Properties of 23 are given by

THEOREM 3. Suppose that 8, 9, and 10 hold and that positive constants b and B are known such that $b \leq \beta = M'(\theta) \leq B$. Define inductively y_i by 1 and x_{n+1} by 23, where b_1 is any constant between b and B and for $n \geq 2$, b_n is defined as follows:

(i) For b_n defined as in Theorem 2, 11, 12, 13, and 19 still hold, and if 20 also holds then so does 21.

(ii) Alternatively, simply put b_n equal to $\hat{\beta}_n$ truncated below by b and above by B [if $\sum_1^n (x_i - \bar{x}_n)^2 > 0$]; otherwise put $b_n = b_{n-1}$. Then the conclusions of (i) remain valid.

Design 23 with $b_n = b$, some positive constant, is the original Robbins-Monro stochastic approximation procedure and has an extensive literature. In Theorem 3 we have taken b_n to be a random variable, and (ii) says that for $b_n = \hat{\beta}_n$ truncated below and above by known positive bounds on β , procedure 23 makes b_n a strongly consistent estimator of β and that the conclusions of Theorem 1 still hold for 23.

Proofs of Theorems 1, 2, and 3 and simulation studies of several adaptive designs will be given elsewhere, along with a treatment of the case in which prior lower and upper bounds for β are not known; the latter case can be handled by a modification of b_n in 23.

This research was supported by the National Institute of General Medical Sciences, the National Science Foundation, and the Office of Naval Research.

1. Lai, T. L. & Robbins, H. (1977) Proc. Natl. Acad. Sci. USA 74, 2667-2669.