Adaptive optimisation and D-optimum experimental design

Luc Pronzato

Laboratoire I3S, CNRS–UNSA, bât. Euclide,
Les Algorithmes, 2000 route des Lucioles,
Sophia Antipolis, 06410 Biot, France.
Tel: 33 4 92 94 27 08; Fax: 33 4 92 94 28 98;
email: pronzato@unice.fr

Abstract

We consider the situation where one has to maximise a function $\eta(\theta, \mathbf{x})$ with respect to $\mathbf{x} \in I\!\!R^q$, when θ is unknown and estimated by least squares through observations $y_k = \mathbf{f}^\top(\mathbf{x}_k)\theta + \epsilon_k$, with ϵ_k some random error. Classical applications are regulation and extremum control problems. The approach we adopt corresponds to maximizing the sum of the current estimated objective and a penalisation for poor estimation: $\mathbf{x}_{k+1}$ maximises $\eta(\hat{\theta}^k, \mathbf{x}) + (\alpha_k/k)\, d_k(\mathbf{x})$, with $\hat{\theta}^k$ the estimated value of θ at step k and d_k the penalisation function. Sufficient conditions for strong consistency of $\hat{\theta}^k$ and for almost sure convergence of $(1/k)\sum_{i=1}^k \eta(\theta, \mathbf{x}_i)$ to the maximum value of $\eta(\theta, \mathbf{x})$ are derived in the case where $d_k(\cdot)$ is the variance function used in the sequential construction of D-optimum designs. A classical sequential scheme from adaptive control is shown not to satisfy these conditions, and numerical simulations confirm that it indeed has convergence problems.

Keywords— adaptive control, convergence, LS estimation, D-optimum design, sequential design, strong consistency

1 Introduction

We consider an optimisation problem, where one wants to maximise a worth $\eta(\bar{\theta}, \mathbf{x})$ with respect to $\mathbf{x} \in \mathcal{X}$. The worth function $\eta(\bar{\theta}, \cdot)$ may be possibly multimodal on $\mathcal{X}$, a compact subset of $I\!\!R^q$. The value of $\bar{\theta}$ is unknown and estimated by Least–Squares (LS) from observations

$$y_k = \mathbf{f}^\top(\mathbf{x}_k)\bar{\theta} + \epsilon_k\,, \tag{1}$$