Journal of Global Optimization (2005) 31: 599–600
DOI 10.1007/s10898-004-9969-x
© Springer 2005

Optimising Noisy Objective Functions

DAVID W. BULGER[1] and H. EDWIN ROMEIJN[2]
[1]*Statistics Department, Macquarie University, NSW 2109, Australia* (*e-mail: dbulger@efs.mq.edu.au*)
[2]*Department of Industrial and Systems Engineering, University of Florida, Gainesville, Florida, U.S.A.* (*e-mail: romeijn@ise.ufl.edu*)

Abstract. This discussion paper considers the use of stochastic algorithms for solving global optimisation problems in which function evaluations are subject to random noise. An idea is outlined for discussion at the forthcoming Stochastic Global Optimisation 2001 workshop in Hanmer in June; we propose that a noisy version of pure random search be studied.

Key words: Global optimisation, Noisy objective function, Pure random search

Generally in the study of global optimisation methods, it is assumed that we have access to a *black box* which will, for a fixed cost per sample, output the exact value of an objective function at the sample point. Often, however, in practical applications, our observations of the objective function are subject to random noise. Some examples of this might be

- a physical experiment in which only some of the environmental variables are controllable,
- an industrial problem in, for instance, agriculture, where again we have incomplete knowledge of and control over the system,
- the performance of a complex stochastic system, the value of which is obtained via a simulation of the system.

For simplicity, we will assume that the observational noise is normally distributed, although we will allow the variance to be unknown and to vary across the domain. Thus we wish to

$$\text{minimise } g(x), \text{ subject to } x \in S,$$

where S is a measurable space and $g : S \to \mathbb{R}$ is a measurable function, assuming that we can only observe $g(x)$ by sampling independent values from a normal distribution with mean $g(x)$ and error variance σ_x^2.

Observational noise in the objective function complicates the already difficult global optimisation problem. The pure random search algorithm provides some of the most elementary analysis of stochastic optimisation methods in the case of a noiseless objective function. One step toward an analogous analysis for the noisy case appears in Baumert and Smith [1], where it is established that, for a continuous objective function, if the sequence of sample points is everywhere dense in the domain, and if we estimate the true objective function value at each point by averaging noisy sample values in a neighbourhood whose radius shrinks toward 0 as the algorithm progresses, then the point with the most favourable estimate will converge to the true optimum with probability 1. It is clear that the everywhere dense sampling condition is indispensible, since a finite number of samples in one area of the domain will have a positive probability of giving an arbitrarily inaccurate portrayal of the true objective function in that area.

But assuming the sequence of sample points to be dense – assuming for simplicity, in fact, that they are independent and identically distributed – is only part of the story. Rather than keeping track of every sample point and value arising during algorithm execution, we might aim for something simpler in terms of memory usage and stochastic analysis. Accordingly, we propose during the forthcoming workshop to formulate and analyse a variation of noisy pure random search which is Markovian in the sequence traced by the current favourite domain point.

Reference

1. Baumert, S. and Smith, R.L. (2001), *Pure Random Search for Noisy Objective Functions*, University of Michigan Technical Report 01-03.