

Approximate Implementations of Pure Random Search in the Presence of Noise^{*}

DAVID L.J. ALEXANDER¹, DAVID W. BULGER², JAMES M. CALVIN³,
H. EDWIN ROMEIJN⁴ and RYAN L. SHERRIFF⁵

¹College of Sciences, Massey University, Wellington, New Zealand (e-mail: David.Alexander.3@uni.massey.ac.nz)

²Department of Statistics, Macquarie University, NSW 2109, Australia (e-mail: dbulger@efs.mq.edu.au)

³Department of Computer and Information Science, New Jersey Institute of Technology, New York, NJ, U.S.A. (e-mail: calvin@homer.njit.edu)

⁴Department of Industrial and Systems Engineering, University of Florida, Gainesville, FL, U.S.A. (e-mail: romeijn@ise.ufl.edu)

⁵Institute of Information Sciences and Technology, Massey University, Palmerston North, New Zealand (e-mail: flick@infogen.net.nz)

Abstract. We discuss the noisy optimisation problem, in which function evaluations are subject to random noise. Adaptation of pure random search to noisy optimisation by repeated sampling is considered. We introduce and exploit an *improving bias condition* on noise-affected pure random search algorithms. Two such algorithms are considered; we show that one requires infinite expected work to proceed, while the other is practical.

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

1. Introduction

This paper reports the results of the Noisy Optimisation group at the Stochastic Global Optimisation 2001 Workshop in Hanmer, New Zealand in June, 2001. The problem under consideration is as outlined in the pre-workshop discussion paper also appearing in this issue. Recall that we wish to

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

where $S \subseteq \mathbb{R}^n$ is a measurable space and $g : S \rightarrow \mathbb{R}$ is a measurable function. However, we can only *approximately* observe $g(x)$, which we take to effectively mean that successive approximate observations of $g(x)$ are in fact independent observations from some probability distribution with

^{*}Supported by a Bright Future Scholarship administered by the Foundation for Research, Science and Technology, New Zealand.

mean $g(x)$. As a default, we will often assume that this distribution is in fact the normal distribution with variance σ_x^2 . We will not require g to have a minimum; rather, an algorithm will be considered to terminate upon sampling a function value sufficiently small. We will refer to this problem as the *noisy global optimisation problem*.

In contrast with the wide attention that global optimisation of deterministic functions has received over the last decades, the additional problems that arise when functions cannot be evaluated exactly, but result, for example, from a simulation, have received relatively little attention. Extensions of the *pure random search (PRS)* algorithm for traditional global optimisation to the problem of solving noisy global optimisation problems can be found in Gurin [9] and Devroye [7]. Recently, several simulated annealing algorithms have been proposed for solving the noisy global optimisation problem in the case where the feasible region is discrete (see, e.g., [1]), and an interesting algorithm with a convergence result for the continuous case appears in Baumert and Smith [3].

In Section 2 we will introduce a new class of algorithms that generalise the PRS algorithm, and a stochastic bound condition we call the *improving bias property*. After deriving a probabilistic asymptotic convergence guarantee, we use the class of generalised PRS algorithms meeting this condition as our basic framework for suggesting, in Section 3, several potential implementations.

2. Algorithm

The most elementary stochastic algorithm for solving noiseless global optimisation problems is PRS (see, for instance, Brooks [5] or Boender and Romeijn [4]). In this algorithm, we repeatedly sample points in the feasible region S from a *sampling distribution* δ , and keep track of the best (with respect to the objective function value) point sampled:

Pure random search (PRS)

Step 0. Set $k = 0$. Generate $x_0 \sim \delta$ and set $y_0 = g(x_0)$.

Step 1. Generate a point $z \sim \delta$.

Step 2. If $g(z) < y_k$, then set $y_{k+1} = g(z)$ and $x_{k+1} = z$. Otherwise, set $y_{k+1} = y_k$ and $x_{k+1} = x_k$.

Step 3. Increment k and return to Step 1.

Devroye [6] has shown that the sequence of function values $\{Y_n; n = 0, 1, 2, \dots\}$ converges with probability one to the global minimum (or infimum) of g over S if δ is an absolutely continuous probability distribution. Note that the sampling distribution δ is often chosen to be the uniform distribution on S .

Extensions of this algorithm to the solution of noisy global optimisation problems can be found in Gurin [9] and Devroye [7]. These papers generalise PRS directly, using the quite intuitive approach of *estimating* the value of $g(x)$ as the sample mean of many observations thereof. In Section 3, we will obtain an algorithm that is similar in spirit to this idea as one of the possible implementations of our class of algorithms.

Note that each iteration of the PRS algorithm must decide whether the newly observed point is an improvement over the best point observed before. This clearly is a trivial matter in noiseless global optimisation. In the case of noisy global optimisation, however, it is uncertain whether the previous “best” point or the newly observed one is best. Accordingly we make the following generalisation:

Generalised pure random search with acceptance rule A (GPRS(A))

Step 0. Set $k = 0$. Generate $x_0 \sim \delta$.

Step 1. Generate a point $z \sim \delta$.

Step 2. With probability $A(k, x_k, g(x_k), z, g(z))$, set $x_{k+1} = z$. Otherwise, set $x_{k+1} = x_k$.

Step 3. Increment k and return to Step 1.

The acceptance rule A is a function mapping $\{0, 1, \dots\} \times S \times \mathbb{R} \times S \times \mathbb{R}$ to $[0, 1]$. It encapsulates the influence of the iteration counter and the locations and values of the incumbent and candidate points on whether the candidate is accepted. The acceptance rule’s probabilistic nature describes randomness due to observational noise (as well as any deliberate randomness in the algorithm).

As a simple example, suppose that the observational noise is normal with variance $\sigma_x^2 = 1$ for each $x \in S$. Consider the following algorithm:

Example algorithm

Step 0. Set $k = 0$. Generate $x_0 \sim \delta$.

Step 1. Generate a point $z \sim \delta$.

Step 2. Let y_x and y_z be noisy evaluations of $g(x_k)$ and $g(z)$.

Step 3. If $y_z < y_x$, set $x_{k+1} = z$. Otherwise, set $x_{k+1} = x_k$.

Step 4. Increment k and return to Step 1.

This algorithm falls within the GPRS class: its acceptance rule A is given by

$$A(k, x, g_x, z, g_z) = \Phi\left(\frac{g_x - g_z}{\sqrt{2}}\right) \quad (1)$$

for any integer k , domain points x, z and real numbers g_x, g_z , where Φ denotes the standard normal cumulative distribution function. Also it is readily implementable, whereas the problem’s observational noise prevents the implementation of standard PRS.

Let us take the opportunity to dispel one possible point of confusion. Step 2 in GPRS does *not* require us to determine $g(z)$ and $g(x_k)$, evaluate

$A(k, x_k, g(x_k), z, g(z))$, and use this probability to choose between the candidate and the incumbent. Indeed, that would be impossible, due to the observational noise. Rather, the acceptance rule A is probabilistic due to our lack of precise knowledge of g . The example algorithm above illustrates this; there is no explicit randomisation, but the observational noise leads to the probabilistic acceptance rule (1).

While the example algorithm is an approximation to PRS which can be implemented in the noisy case, it does not converge to the global optimum; it is apparent that, though correct acceptance-rejection decisions will dominate the algorithm's behaviour, it will never stop making occasional incorrect decisions. In an effort to describe stochastic algorithms for noisy optimisation which *do* converge to the global optimum, we now define the "improving bias" property.

Let $\{a_k : \mathbb{R} \rightarrow [0, 1]\}$ be a sequence of nondecreasing *acceptance probability bound* functions converging pointwise on $\mathbb{R} \setminus \{0\}$ to the Heaviside step function H (i.e., the function that takes the value 0 on the set of negative reals, and the value 1 on the set of nonnegative reals). If, for every k, x, z and g with $g(x) < g(z)$ we have

$$A(k, x, g(x), z, g(z)) \leq a_k(g(x) - g(z)),$$

and for every k, x, z and g with $g(x) > g(z)$ we have

$$A(k, x, g(x), z, g(z)) \geq a_k(g(x) - g(z)),$$

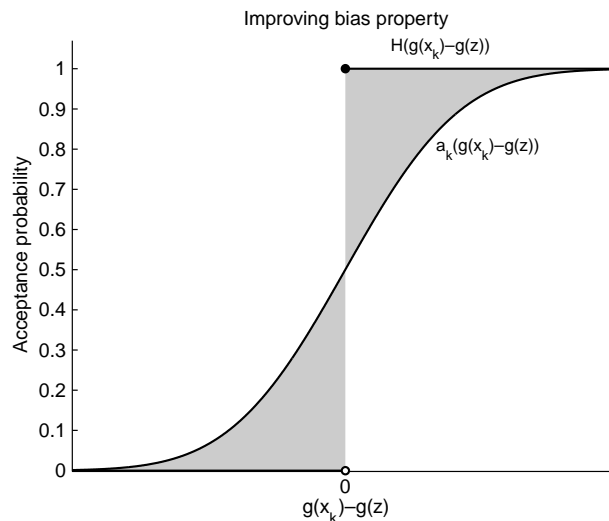


Figure 1. The acceptance probability at the k th step of a practical algorithm may depend on the candidate z and the incumbent x , as well as the value of the objective function at these points. The improving bias property restricts the acceptance probability to the shaded region. Since the Heaviside step function H represents the correct decision (which can only be guessed in noisy optimisation), the function a_k can be thought of as a lower bound on the accuracy of the acceptance-rejection decision.

then we say that A satisfies the *improving bias property with parameters* $\{a_k\}$.

This property bounds the accuracy of the acceptance–rejection decision. PRS makes the decision with perfect accuracy, and its acceptance probability is given by $H(g(x_k) - g(z))$. The improving bias property bounds the acceptance probability between a_k and the ideal probability H . This is illustrated in Figure 1.

The following theorem demonstrates the significance of the improving bias property to the asymptotic behaviour of GPRS.

THEOREM 2.1. *If A satisfies the improving bias property, then the $GPRS(A)$ algorithm converges in the sense that, if*

$$B_y \equiv \{x \in S : g(x) \leq y\}$$

has positive measure under δ , then $\pi_y^k \equiv \mathbb{P}[X_k \in B_y] \rightarrow 1$ as $k \rightarrow \infty$.

Proof. Let p denote the cumulative distribution function of $g(X)$, where $X \sim \delta$, that is, let

$$p(t) = \delta(\{x \in S : g(x) \leq t\})$$

for all real t . Let z denote the candidate point at the $(k + 1)$ st pass through the loop, so that x_{k+1} is probabilistically set to equal either x_k or z . Since x_k and z are independent,

$$\begin{aligned} \pi_y^{k+1} &= \mathbb{P}(g(X_{k+1}) \leq y | g(X_k) \leq y \text{ and } g(Z) \leq y) \mathbb{P}[g(X_k) \leq y] \mathbb{P}[g(Z) \leq y] \\ &\quad + \mathbb{P}(g(X_{k+1}) \leq y | g(X_k) \leq y \text{ and } g(Z) > y) \mathbb{P}[g(X_k) \leq y] \mathbb{P}[g(Z) > y] \\ &\quad + \mathbb{P}(g(X_{k+1}) \leq y | g(X_k) > y \text{ and } g(Z) \leq y) \mathbb{P}[g(X_k) > y] \mathbb{P}[g(Z) \leq y] \\ &\quad + \mathbb{P}(g(X_{k+1}) \leq y | g(X_k) > y \text{ and } g(Z) > y) \mathbb{P}[g(X_k) > y] \mathbb{P}[g(Z) > y] \\ &\geq \pi_y^k p(y) + \pi_y^k (1 - p(y)) \mathbb{E}(1 - a_k(g(X_k) - g(Z)) | g(X_k) \leq y \text{ and } g(Z) > y) \\ &\quad + (1 - \pi_y^k) p(y) \mathbb{E}(a_k(g(X_k) - g(Z)) | g(X_k) > y \text{ and } g(Z) \leq y) \\ &\geq \pi_y^k p(y) + \pi_y^k (1 - p(y)) \mathbb{E}(1 - a_k(y - g(Z)) | g(X_k) \leq y \text{ and } g(Z) > y) \\ &\quad + (1 - \pi_y^k) p(y) \mathbb{E}(a_k(y - g(Z)) | g(X_k) > y \text{ and } g(Z) \leq y) \\ &= \pi_y^k p(y) + \pi_y^k (1 - p(y)) \mathbb{E}(1 - a_k(y - g(Z)) | g(Z) > y) \\ &\quad + (1 - \pi_y^k) p(y) \mathbb{E}(a_k(y - g(Z)) | g(Z) \leq y). \end{aligned}$$

Now

$$\mathbb{E}(a_k(y - g(Z)) | g(Z) \leq y) = \frac{1}{1 - p(y)} \int_y^\infty a_k(y - c) dp(c) \rightarrow 1 \text{ as } k \rightarrow \infty$$

by the Dominated Convergence Theorem (see, e.g., [2]); thus there exists K_1 such that

$$E(a_k(y - g(Z)) | g(Z) \leq y) \geq \frac{1}{2}, \quad (3)$$

whenever $k \geq K_1$. Similarly

$$E(1 - a_k(y - g(Z)) | g(Z) > y) \rightarrow 1,$$

thus since, for small positive ϵ , $p(y)\epsilon/[2(2 - \epsilon)(1 - p(y))]$ is positive, there exists some K_2 such that

$$E(1 - a_k(y - g(Z)) | g(Z) > y) \geq 1 - \frac{p(y)\epsilon}{2(2 - \epsilon)(1 - p(y))}, \quad (4)$$

whenever $k \geq K_2$.

By substituting the bounds (3) and (4) into (2), we see that for any small positive ϵ , we can choose $K = \max\{K_1, K_2\}$ such that

$$\begin{aligned} \pi_y^{k+1} &\geq \pi_y^k p(y) + \pi_y^k (1 - p(y)) \left(1 - \frac{p(y)\epsilon}{2(2 - \epsilon)(1 - p(y))} \right) + \frac{1}{2} (1 - \pi_y^k) p(y) \\ &= \left(\frac{2 - p(y) - \epsilon}{2 - \epsilon} \right) \pi_y^k + \frac{p(y)}{2}, \end{aligned}$$

whenever $k \geq K$. Reiterating this inequality gives

$$\begin{aligned} \pi_y^{k+l} &\geq \left(\frac{2 - p(y) - \epsilon}{2 - \epsilon} \right)^l \pi_y^k + \frac{p(y) \left(1 - \left(\frac{2 - p(y) - \epsilon}{2 - \epsilon} \right)^l \right)}{2 \left(1 - \left(\frac{2 - p(y) - \epsilon}{2 - \epsilon} \right) \right)} \\ &\geq \left(1 - \frac{\epsilon}{2} \right) \left(1 - \left(\frac{2 - p(y) - \epsilon}{2 - \epsilon} \right)^l \right), \end{aligned}$$

whenever $k \geq K$ and $l \geq 0$. In particular, whenever

$$l \geq \frac{\ln \frac{\epsilon}{2}}{\ln \left(\frac{2 - p(y) - \epsilon}{2 - \epsilon} \right)},$$

we will have

$$1 - \left(\frac{2 - p(y) - \epsilon}{2 - \epsilon} \right)^l \geq 1 - \frac{\epsilon}{2},$$

and thus

$$\pi_y^{k+l} \geq \left(1 - \frac{\epsilon}{2} \right)^2 > 1 - \epsilon.$$

Since π_y^{k+l} is a probability and thus bounded above by 1, it follows that $\pi_y^k \rightarrow 1$ as $k \rightarrow \infty$. \square

3. Implementing GPRS with Improving Bias

In this section, we consider the implementation of GPRS algorithms with the improving bias property. Step 2 of the GPRS algorithm requires that

we (fallibly) decide which of the current point x_k and the candidate point z has the better value of g , and the improving bias property requires the accuracy of these decisions to improve as the algorithm progresses. As described in Section 1, the decisions must be based only on noisy observations of g . Clearly, more accuracy requires a greater number of noisy evaluations of g , and any satisfactory algorithm will strike a compromise between the accuracy of the decisions and the number of evaluations each decision requires.

This section presents an informal discussion of three possible approaches. Firstly we investigate the possibility of uniform convergence of the acceptance probability bounds a_k , that is, a decision method whose error rate converges to zero at a fixed rate, regardless of the magnitude of the improvement or deterioration entailed in acceptance of the candidate point. Secondly we look at the possibility of making a number of noisy observations of $g(x_k)$ and $g(z)$ which grows with k in a predetermined sequence. Lastly, we consider an intermediate approach suggested by the theory of sequential analysis.

3.1. UNIFORM CONVERGENCE OF a_k

The method addressed in this subsection is as follows. Let α_k be a sequence of positive numbers, less than $1/2$ and tending to 0.

Increasing confidence acceptance–rejection search (ICARS)

Step 0. Set $k = 0$. Generate $x_0 \sim \delta$.

Step 1. Generate a point $z \sim \delta$.

Step 2. Conduct a sequential test to choose between the hypotheses $H_<:g(z) < g(x_k)$ and $H_>:g(z) > g(x_k)$, gathering as many samples of $g(z)$ and $g(x_k)$ as necessary to provide a result with confidence level α_k .

Step 3. If the sample mean estimating $g(z)$ is less than the sample mean estimating $g(x_k)$, then set $x_{k+1} = z$. Otherwise, set $x_{k+1} = x_k$.

Step 4. Increment k and return to Step 1.

ICARS is a GPRS algorithm satisfying the improving bias property, with

$$a_k(y) = \begin{cases} \alpha_k & \text{for } y < 0, \\ 1 - \alpha_k & \text{for } y \geq 0, \end{cases}$$

due to the sequential test in Step 2. This test is not fully specified by the algorithm description. Many sequential experiments could be designed to fit the parameters given (see, e.g., [8, 10] for the theory of sequential analysis). In this case, all such experimental designs are on the same poor footing; although intuitively attractive, ICARS has a fundamental flaw, as shown in the following theorem. Assume the distribution of $g(z)$ is absolutely continuous where $z \sim \delta$. Also, for simplicity, assume that the noise magnitude is fixed at a known value $\sigma_x = \sigma$ for all $x \in S$.

THEOREM 3.1. *The expected number of function evaluations required by ICARS for each pass through the loop (Steps 1–4) is infinite.*

Proof. Fix the value of the loop counter k . Let x_k and z be two fixed points in S . Now each noisy evaluation of $g(x_k)$ and $g(z)$ can be considered as a noisy observation of $d = g(z) - g(x_k)$; thus it will be a Gaussian variable with unknown mean $\mu = g(z) - g(x_k)$ and known variance $2\sigma^2$.

Let $D(\mu, 2\sigma^2, \alpha_k)$ represent the expected number of samples of independent random variables, normally distributed with mean $\mu \neq 0$ and variance $2\sigma^2$, required to choose between $\mu < 0$ and $\mu > 0$ at confidence level α_k , under the sequential test selected for Step 2. The distribution of $g(x_k)$ is absolutely continuous (this can be seen by induction, starting with the absolute continuity of the distribution of $g(z)$). Find an interval (q, r) on which the distribution of $g(x_k)$ is minorised by a positive constant. The distribution of $g(z)$ on this same interval must also be minorised by a positive constant. Their joint distribution can now be viewed as a nontrivial mixture of the distribution in which $g(z)$ and $g(x_k)$ are independently and uniformly distributed on this interval, and a remainder distribution. In the doubly uniform case, the expected number of passes through the inner loop during the k th pass through the outer loop is

$$\int_q^r \int_q^r \frac{D(g_x - g_z, 2\sigma^2, \alpha_k)}{(r - q)^2} dg_z dg_x.$$

This integral diverges, provided that for all σ and α_k , $D(\mu, 2\sigma^2, \alpha_k)$ is bounded below by a multiple of $1/|\mu|$ where $\mu \neq 0$; the remainder of the proof establishes this bound.

We use a stochastic coupling argument. Without loss of generality fix $\mu > 0$, and for simplicity assume henceforth $2\sigma^2 = 1$. Let f represent the density function of the standard normal distribution. Define

$$\hat{f}(x) = \begin{cases} f(x - \mu), & x < 0, \\ f(x + \mu), & x \geq 0. \end{cases}$$

Clearly $\int_{-\infty}^{\infty} \hat{f}(x) dx = 2(1 - \Phi(\mu))$, and moreover

$$\int_{-\infty}^{\infty} \hat{f}(x + \mu) - f(x) dx = \int_{-\infty}^{\infty} \hat{f}(x - \mu) - f(x) dx = \Phi(\mu) - \frac{1}{2},$$

so that by normalising we can create three probability density functions,

$$f_A(x) = \frac{f(x + \mu) - \hat{f}(x)}{\Phi(\mu) - 1/2}, f_B(x) = \frac{\hat{f}(x)}{2(1 - \Phi(\mu))}, f_C(x) = \frac{f(x - \mu) - \hat{f}(x)}{\Phi(\mu) - 1/2}.$$

Now suppose the random variables Y and Y' are constructed as follows. With probability $2(1 - \Phi(\mu))$, we let Y be a random variable with distribution f_B , and set $Y' = Y$. Otherwise, we let Y and Y' be independent random variables with distributions f_C and f_A . Then the marginal distributions of Y

and Y' are normal with variance 1 and means μ and $-\mu$, respectively, but with probability $2(1 - \Phi(\mu)) > 1 - \sqrt{2/\pi}\mu$ we have $Y' = Y$.

Now let $\{(Y_j, Y'_j): j = 1, 2, \dots\}$ be an IID sequence of ordered pairs distributed like (Y, Y') in the previous paragraph. Define the stopping time J as the first index j for which $Y_j \neq Y'_j$. Then J is geometrically distributed with

$$E(J) \geq \sqrt{\pi}/(\sqrt{2}\mu). \quad (5)$$

Note that $\{Y_j\}$ is an IID sequence of normal random variables with mean μ and variance 1, and $\{Y'_j\}$ is similar but with mean $-\mu$.

Let T be the event that, given the sequence $\{Y_j\}$, the sequential test will terminate at some step $j < J$. This can be partitioned into two events $T_<$ and $T_>$ according to whether the test, upon terminating, infers that $\mu < 0$ or that $\mu > 0$. Now $T_<$ is a subset of the event that both sequences $\{Y_j\}$ and $\{Y'_j\}$ will lead to the decision that $\mu < 0$; this in turn is a subset of the event that the sequence $\{Y_j\}$ will lead to the decision that $\mu < 0$, and this event can have probability no greater than α_k . Thus $P(T_<) \leq \alpha_k$, and similarly $P(T_>) \leq \alpha_k$.

It follows that, with probability at least $1 - 2\alpha_k$, the event T does not occur. In this case, the sample size required from the sequence $\{Y_j\}$ in order to choose between the hypotheses $\mu < 0$ and $\mu > 0$ is greater than J . (In the event T , of course, the sample size is bounded below by 1.) Recalling (5), we have a lower bound for $D(\mu, 2\sigma^2, \alpha_k)$ (which is the expectation of this sample size) of

$$2\alpha_k + (1 - 2\alpha_k)\sqrt{\pi}/(\sqrt{2}\mu),$$

which is greater than $(1 - 2\alpha_k)\sqrt{\pi/2}$ divided by $|\mu|$, as required. \square

We have established the impracticality of the ICARS algorithm; its expected computational effort per iteration is infinite. Therefore a practical implementation of GPRS with improving bias will need to be formulated along slightly different lines.

3.2. FIXED SCHEDULE OF EVALUATION COUNT GROWTH

The previous approach amounted to testing each candidate point to a predetermined confidence level, without regard for the number of samples required. In this section we consider the other extreme, which we will call *Increasing sample size acceptance-rejection search*: we make a predetermined number of function evaluations, without regard to the confidence level.

Suppose $\{N_k: k = 0, 1, 2, \dots\}$ is a sequence of natural numbers, and that the decision between retaining x_k or adopting z is made on the basis of N_k noisy evaluations of each of $g(x_k)$ and $g(z)$. Denoting the corresponding sample means by M_{x_k} and M_z , the best decision is clearly to choose whichever of x_k or z yields the smaller sample mean.

Increasing sample-size acceptance–rejection search (ISSARS)**Step 0.** Set $k = 0$. Generate $x_0 \sim \delta$.**Step 1.** Generate a point $z \sim \delta$.**Step 2.** Take samples of g of size N_k at x_k and z , and compute the means M_{x_k} and M_z .**Step 3.** If $M_z < M_{x_k}$, then set $x_{k+1} = z$. Otherwise, set $x_{k+1} = x_k$.**Step 4.** Increment k and return to Step 1.

Whereas ICARS in Section 3.1 needs, in expectation, an infinite number of function evaluations in each iteration, ISSARS proceeds in such a way that any desired number of passes through the algorithm loop can be achieved with a finite cost, known in advance.

THEOREM 3.2. *If the error variances σ_x^2 are uniformly bounded, the sequence of points generated by ISSARS converges to the global optimum of (P) with probability one.*

Proof. Assume $\sigma_x \leq \sigma$, say, for all $x \in S$. Since the sample means are normally distributed, their difference is normally distributed as well, with mean $g(z) - g(x_k)$ and variance $\sqrt{\sigma_x^2 + \sigma_z^2}$. The random variable

$$\frac{(M_{x_k} - M_z) - (g(x_k) - g(z))}{\sqrt{(\sigma_{x_k}^2 + \sigma_z^2)/N_k}},$$

then has the standard normal distribution. The probability of accepting the candidate point z is equal to

$$\begin{aligned} \mathbb{P}[M_z < M_{x_k}] &= \mathbb{P}\left[\frac{(M_{x_k} - M_z) - (g(x_k) - g(z))}{\sqrt{(\sigma_{x_k}^2 + \sigma_z^2)/N_k}} > -\frac{g(x_k) - g(z)}{\sqrt{(\sigma_{x_k}^2 + \sigma_z^2)/N_k}}\right] \\ &= \Phi\left(\frac{g(x_k) - g(z)}{\sqrt{(\sigma_{x_k}^2 + \sigma_z^2)/N_k}}\right) \end{aligned}$$

(where Φ again denotes the standard normal cumulative distribution function). Thus the acceptance probability is bounded from below by

$$a_k(g(x_k) - g(z)) = \Phi\left(\frac{g(x_k) - g(z)}{\sigma\sqrt{2/N_k}}\right),$$

whenever $g(x_k) > g(z)$ and from above by the same expression whenever $g(x_k) < g(z)$. Since $N_k \rightarrow \infty$, $a_k(y) \rightarrow H(y)$, in satisfaction of the improving bias property. The result now follows by Theorem 2.1. \square

ISSARS is stricter, in terms of confidence required, for larger jumps up or down, as is appropriate. Since our aim is to achieve optimal or nearly-

optimal function values, it is more important to decide correctly whether a proposed transition is an improvement or a deterioration if that improvement or deterioration is large than if it is small. Thus the kind of nonuniformity shown in the convergence of $\Phi\left(|g(x_k) - g(z)| / (\sigma\sqrt{2/N_k})\right)$ to 0, specifically, that it is faster for larger $|g(x_k) - g(z)|$, is appropriate. There is, however, no reason to believe that the amount of extra strictness it places on larger jumps is optimal. A good understanding of this approach would require much more work.

3.3. SEQUENTIAL-ANALYTIC APPROACH

Calculation of an optimal schedule of evaluation count growth is likely to be difficult or impossible, but a hybrid method somewhere in between the methods proposed in the previous two subsections seems worth study. As we have seen, an unswerving determination to know, with a predetermined level of confidence, whether the current candidate point is an improvement or a deterioration may lead us to make infinitely many function evaluations in expectation. On the other hand, using a predetermined number of function evaluations may waste time confirming what is already almost certain, in cases where the current point and the candidate point vary greatly in objective function value.

The field of sequential analysis is concerned with hypothesis tests in which the number of data collected is not known in advance; data are collected until a decision can satisfactorily be made. A common example of a sequential hypothesis test is the *sequential probability ratio test* (SPRT). In this test, data are gathered until the ratio of the likelihoods of observing the collected sample under two opposing hypotheses strays outside a continuation interval. Then, the hypothesis according to which the sample is more likely is accepted. Sample size optimality results are known for the SPRT and other standard sequential hypothesis tests [8]. A deeper study of improving bias GPRS might view each pass through the algorithm loop as a sequential analysis problem.

4. Conclusion

This article has considered generalisations of the pure random search algorithm for global optimisation under the influence of observational noise. Such algorithms can be analysed via the parametrised improving bias condition. The method of Section 3.1, while intuitively reasonable, was shown to require infinite expected effort for each pass through its main loop. On the other hand, Section 3.2 discussed a method with prescribed effort at each iteration, and proved its convergence to the optimum.

Acknowledgement

This paper is the result of work conducted at the SGO 2001 Workshop in Hanmer, New Zealand, June 2001, funded in part by the Marsden Fund of the Royal Society of New Zealand.

References

1. Alrefaie, M.H. and Andradóttir, S. (1996), A simulated annealing algorithm with constant temperature for discrete stochastic optimization. *Management Science* 45(5), 748–764.
2. Ash, R.B. and Doléans-Dade, C.A. (2000), *Probability and Measure Theory*. Academic Press, San Diego, California.
3. Baumert, S. and Smith, R.L. (2001), Pure Random Search for Noisy Objective Functions. University of Michigan Technical Report 01–03.
4. Boender, C.G.E. and Romeijn, H.E. (1995), Stochastic methods. In: Horst, R. and Pardalos, P.M. (eds.), *Handbook of Global Optimization*, pp. 829–869. Kluwer Academic Publishers, Dordrecht, The Netherlands.
5. Brooks, S.H. (1958), A discussion of random methods for seeking maxima. *Operations Research* 6, 244–251.
6. Devroye, L. (1978), Progressive global random search of continuous functions. *Mathematical Programming* 15, 330–342.
7. Devroye, L.P. (1976), On the convergence of statistical search. *IEEE Transactions on Systems, Man, and Cybernetics* 6(1), 46–56.
8. Ghosh, B.K. and Sen, P.K. (eds.) (1991), *Handbook of Sequential Analysis*, Dekker, New York.
9. Gurin, L.S. (1966), Random search in the presence of noise. *Engineering Cybernetics* 4(3), 252–260.
10. Siegmund, D. (1985), *Sequential Analysis*. Springer-Verlag, New York.