*Proceedings of the 2002 Winter Simulation Conference*
*E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, eds.*

# CONFIDENCE REGIONS FOR STOCHASTIC APPROXIMATION ALGORITHMS

Ming-hua Hsieh

Department of Management Information Systems
National Chengchi University
Wenshan, Taipei 11623, TAIWAN

Peter W. Glynn

Department of Management Science and Engineering
Stanford University
Stanford, CA 94305, U.S.A.

## ABSTRACT

In principle, known central limit theorems for stochastic approximation schemes permit the simulationist to provide confidence regions for both the optimum and optimizer of a stochastic optimization problem that is solved by means of such algorithms. Unfortunately, the covariance structure of the limiting normal distribution depends in a complex way on the problem data. In particular, the covariance matrix depends not only on variance constants but also on even more statistically challenging parameters (e.g. the Hessian of the objective function at the optimizer). In this paper, we describe an approach to producing such confidence regions that avoids the necessity of having to explicitly estimate the covariance structure of the limiting normal distribution. This procedure offers an easy way for the simulationist to provide confidence regions in the stochastic optimization setting.

## 1 INTRODUCTION

Stochastic approximation algorithms are iterative procedures that permit simulationists to numerically optimize complex stochastic systems. This class of algorithms exhibits a convergence rate that is typically not faster than of order $c^{-1/2}$, where $c$ is the size of the computer time budget. Given this relatively slow convergence rate (as compared to most non-random iterative procedures), it is desirable, from practical standpoint, to assess the accuracy of the computed solution at the conclusion of the calculation. Given that stochastic approximation algorithms are driven by random numbers, the most natural means of assessing error is via a confidence interval (for use when the optimizing decision variable is scalar) or, more generally, via a confidence region (for use when the optimizing decision variable is vector-valued).

This paper explores the construction of confidence regions for such iterative algorithms. We focus on Robbins-Monro procedures in this paper, although much of the theory we present has analogs in the setting of Kiefer-Wolfowitz procedures. Section 2 describes a central limit theorem for the Robbins-Monro algorithm. This central limit theorem lies at the basis of the confidence region procedures we describe. Specifically, Section 2 discusses a confidence region procedure that requires consistent estimation of a certain covariance matrix, whereas Section 3 describes a more easily implemented "cancellation" procedure. In Section 4, we discuss some of our computational experience with the procedure introduced in Section 3. Finally, Section 5 offers some concluding remarks.

## 2 CENTRAL LIMIT THEOREM FOR STOCHASTIC APPROXIMATIONS

We start by describing the problem setting precisely. Suppose that our goal is to numerically maximize an objective function $\alpha(\theta)$ over a continuous decision parameter $\theta \in \Re^d$. If $\alpha(\cdot)$ is smooth, it is well known that maximization requires computing an appropriate root $\theta^*$ of the equation

$$\nabla\alpha(\theta^*) = 0 \tag{1}$$

where $\nabla\alpha(\theta)$ is the gradient of $\alpha(\cdot)$ at $\theta \in \Re^d$. (In this paper, we take $\nabla\alpha(\theta)$ to be a row vector.) If $\nabla\alpha(\theta)$ can be numerically evaluated, the root $\theta^*$ can often be efficiently computed by (deterministic) gradient-based Newton-type algorithms (see, e.g., Luenberger 1984; Gill, et al. 1981).

We focus on the case that $\nabla\alpha(\cdot)$ must be computed by Monte Carlo sampling. It is well known that simulation-based algorithms enjoy a much broader range of applicability than do methods requiring numerical evaluation of closed-form expressions for the expectations involved. Specifically, we assume existence of a family $(Z(\theta) : \theta \in \Re^d)$ of random vectors that act as unbiased estimators of the gradient, namely

$$EZ(\theta) = \nabla\alpha(\theta). \tag{2}$$

(Note that $Z(\theta)$ is also encoded as a row vector.) A number of different procedures have been proposed in the literature for obtaining such unbiased estimators of the gradient, including likelihood ratio methods (Glynn 1986, Glynn 1990), infinitesimal perturbation analysis (Glasserman 1991), Conditional Monte Carlo (Fu and Hu 1997), and the "push-out" approach (Rubinstein 1992).

For $\theta \in \Re^d$, let

$$F(dz;\theta) = P(Z(\theta) \in dz)$$

for $z \in \Re^d$. Given the existence of unbiased estimators for the gradient, we are naturally led to the consideration of a Robbins-Monro (R-M) algorithm for computing $\theta^*$. Such a R-M algorithm proceeds by first choosing an initial guess $\theta_0$ for the maximizer $\theta^*$, and subsequently iterates to $\theta_{n+1}$ from $\theta_n$ via the recursion

$$\theta_{n+1} = \theta_n + \frac{a}{n+1} Z_{n+1}(\theta_n) \tag{3}$$

for $a > 0$, where

$$P(Z_{n+1}(\theta_n) \in dz | \theta_0, Z_1(\theta_0), \cdots, Z_n(\theta_{n-1})) = F(dz;\theta_n)$$

for $z \in \Re^d$. (Again, we choose to encode the $\theta_n$'s as row vectors.)

We are now ready to describe one of many known central limit theorems (CLT's) for $\theta_n$; see p.147-150 of Nevel'son and Has'minskii (1973) for a complete proof. For $x \in \Re^d$ (encoded as a row vector), let $\|x\| = \sqrt{xx^T}$ be its Euclidian norm.

**Assumption A.** The sequence $(\theta_n : n \geq 0)$ satisfies the following conditions:

i) $\nabla\alpha(\theta) = (\theta - \theta^*)H + o(\|\theta - \theta^*\|)$ as $\theta \to \theta^*$;
ii) $aH + \frac{1}{2}I$ is symmetric and all its eigenvalues are negative;
iii) for all $\varepsilon > 0$, $\sup_{\varepsilon < \|\theta - \theta^*\| < \frac{1}{\varepsilon}} \nabla\alpha(\theta)(\theta - \theta^*)^T < 0$;
iv) $F(\cdot;\theta) \Rightarrow F(\cdot;\theta^*)$ as $\theta \to \theta^*$, where $\Rightarrow$ denotes weak convergence;
v) there exists $\varepsilon_0 > 0$ such that $(\|Z(\theta)\|^2 : \|\theta - \theta^*\| < \varepsilon_0)$ is a uniformly integrable family of random variables;
vi) there exists $k < \infty$ such that $E\|Z(\theta)\|^2 \leq k(1 + \|\theta\|^2)$.

Recall that for a square matrix $B$,

$$\exp(B) \equiv \sum_{n=0}^{\infty} \frac{B^n}{n!}$$

(which is guaranteed to always converge absolutely and be well-defined).

**Theorem 1.** *Suppose that $(\theta_n : n \geq 0)$ satisfies (2), (3), and assumption **A**. Then,*

$$n^{\frac{1}{2}}(\theta_n - \theta^*) \Rightarrow N(0, C)$$

*as $n \to \infty$, where $N(0, C)$ is a multivariate normal random vector with mean zero and covariance matrix $C$ given by*

$$C = a^2 \int_0^\infty \exp((aH + \frac{1}{2}I)u) E Z(\theta^*)^T Z(\theta^*) \cdot \exp((aH + \frac{1}{2}I)u) du.$$

Before proceeding further, let us briefly discuss assumption **A**. When $\alpha(\cdot)$ is twice continuously differentiable (as will typically be the case), i) is immediately satisfied at a stationary point $\theta^*$ and $H$ is just the Hessian of $\alpha$ evaluated at $\theta^*$. At an isolated maximizer $\theta^*$, $H$ must be negative definite and symmetric, so that $aH + \frac{1}{2}I$ will have negative eigenvalues for $a$ sufficiently large; ii) will then be satisfied automatically. Condition iii) is a sufficient condition that guarantees that $\theta^*$ is a global maximizer of $\alpha(\cdot)$. Condition iv) is a mild continuity hypothesis on the distribution of $Z(\theta)$, and condition v) is a technical integrability hypothesis that is satisfied in great generality. Finally, condition vi) is an assumption that controls the rate at which the variance of $\|Z(\theta)\|$ grows as $\|\theta\|$ tends to infinity.

Theorem 1 asserts that $\theta_n$ converges to $\theta^*$ at rate $n^{-\frac{1}{2}}$ in the number of iterations $n$. Furthermore, the error of $\theta_n$ is approximately normally distributed when $n$ is large. The latter observation suggests the possibility of constructing confidence regions for $\theta^*$ based on the above CLT.

**Proposition 1.** *Suppose that $(\theta_n : n \geq 0)$ satisfies the conditions of Theorem 1 and that $EZ(\theta^*)^T Z(\theta^*)$ is a non-singular matrix. If $C_n \Rightarrow C$ as $n \to \infty$, then*

$$n(\theta_n - \theta^*)C_n^{-1}(\theta_n - \theta^*)^T \Rightarrow \chi_d^2$$

*as $n \to \infty$, where $\chi_d^2$ is a chi-square random variable with $d$ degrees of freedom.*

**Proof** We start by showing that $C$ is non-singular. The matrix $C$ can alternatively be represented as the matrix solution to the equation

$$(aH + \frac{1}{2}I)C + C(aH + \frac{1}{2}I) = EZ(\theta^*)^T Z(\theta^*);$$

see p. 77-78 of Ljung, Pflug, and Walk (1992). Lyapunov's lemma (see p. 133 of Nevel'son and Has'minskii 1973) establishes that $C$ is positive definite and therefore non-singular. Since the matrix inverse functional is continuous in a (matrix) neighborhood of a non-singular matrix, it follows from the continuous mapping principle for weak convergence that $C_n^{-1} \Rightarrow C^{-1}$ as $n \to \infty$.

To continue the argument, we write the quadratic form $n(\theta_n - \theta^*)C_n^{-1}(\theta_n - \theta^*)^T$ as follows:

$$\begin{aligned} & n(\theta_n - \theta^*)C_n^{-1}(\theta_n - \theta^*)^T \\ = \; & n(\theta_n - \theta^*)C^{-1}(\theta_n - \theta^*)^T \\ & + n(\theta_n - \theta^*)(C_n^{-1} - C^{-1})(\theta_n - \theta^*)^T. \end{aligned} \tag{4}$$

The continuous mapping principle implies that the first term on the right hand side of (4) converges weakly to $N(0, C)C^{-1}N(0, C)^T$, which is easily seen to have a $\chi_d^2$ distribution. On the other hand,

$$\begin{aligned} & \|n(\theta_n - \theta^*)(C_n^{-1} - C^{-1})(\theta_n - \theta^*)^T\| \\ & \qquad \leq \|n^{\frac{1}{2}}(\theta_n - \theta^*)\|^2 \cdot \|C_n^{-1} - C^{-1}\|. \end{aligned} \tag{5}$$

The first factor on the right-hand side of (5) converges weakly, by the continuous mapping principle to $\|N(0, C)\|^2$, whereas the second factor converges weakly to zero. Hence, the product converges weakly to zero, establishing that the left-hand side of (4) does indeed converge weakly to $N(0, C)C^{-1}N(0, C)^T$, as desired. □

Thus, if $C_n$ is a consistent estimator for $C$, the region

$$\{\theta : n(\theta_n - \theta)C_n^{-1}(\theta_n - \theta)^T \leq z\}$$

is an approximate $100(1 - \delta)\%$ confidence region for $\theta^*$, provided that $z$ is selected so that $P(\chi_d^2 > z) = \delta$. The key to constructing confidence regions for $\theta^*$ is therefore the consistent estimation of $C$.

As is evident from the formula for $C$, even when $H$ and $EZ(\theta^*)^T Z(\theta^*)$ are known, the numerically evaluation of $C$ is non-trivial. Of course, in practice, both $H$ and $EZ(\theta^*)^T Z(\theta^*)$ are generally unknown and must themselves be estimated, substantially complicating the task. In the current setting, $EZ(\theta^*)^T Z(\theta^*)$ can be estimated in a straightforward fashion. In particular, let

$$\Sigma_n = \frac{1}{n}\sum_{i=1}^{n} Z_i(\theta_{i-1})^T Z_i(\theta_{i-1}). \tag{6}$$

**Proposition 2.** *Suppose that* $\sup\{E\|Z(\theta)\|^4 : \theta \in \Re^d\} < \infty$. *If, in addition,* $(\theta_n : n \geq 0)$ *satisfies the conditions of Theorem 1, then*

$$\Sigma_n \to EZ(\theta^*)^T Z(\theta^*) \quad \textit{a.s. as } n \to \infty.$$

**Proof.** Let $\gamma(\theta) = EZ(\theta)^T Z(\theta)$ and set $D_i = Z_i(\theta_{i-1})^T Z_i(\theta_{i-1}) - \gamma(\theta_{i-1})$. Because $\sup\{E\|Z(\theta)\|^4 : \theta \in \Re^d\} < \infty$, it follows that $(D_i : i \geq 1)$ is a sequence of martingale differences (adapted to $\theta_0, Z_1(\theta_0), Z_2(\theta_1), \ldots$) for which $\sup\{\mathrm{Var}(D_i) : i \geq 1\} < \infty$. Consequently, the martingale $(\sum_{i=1}^n D_i/i : n \geq 1)$ satisfies the conditions of the Martingale Convergence Theorem (see p. 468 of Billingsley 1995, or ch. 12 of Williams 1991), so that there exists a finite-valued random variable $M_\infty$ such that $\sum_{i=1}^n D_i/i \to M_\infty$ a.s. as $n \to \infty$. Kronecker's lemma (see p. 250 of Loeve 1977) then guarantees that $n^{-1}\sum_{i=1}^n D_i/i \to 0$ a.s. as $n \to \infty$. But

$$\Sigma_n = n^{-1}\sum_{i=1}^{n} D_i + n^{-1}\sum_{i=0}^{n-1}\gamma(\theta_i). \tag{7}$$

Conditions **A** iv) and v) ensure that $\gamma(\theta) \to \gamma(\theta^*)$ as $\theta \to \theta^*$. In addition, it is known that $\theta_n \to \theta^*$ a.s. as $n \to \infty$ under the conditions of Theorem 1; see p. 93 of Nevel'son and Has'minskii (1973). Hence,

$$n^{-1}\sum_{i=0}^{n-1}\gamma(\theta_i) \to \gamma(\theta^*) = EZ(\theta^*)^T Z(\theta^*) \quad \text{a.s.}$$

as $n \to \infty$. It follows from (7) that

$$\Sigma_n \to EZ(\theta^*)^T Z(\theta^*) \quad \text{a.s. as } n \to \infty.$$

□

The estimator defined by (6) can be modified (and possibly improved) by weighting more recent observations more heavily, rather than weighting all observations equally.

In the setting of Theorem 1, the greater challenge in constructing a consistent estimator $C_n$ is the estimation of the matrix $H$. The $(i, j)$'th entry of the matrix $H$ is given by

$$H_{ij} = \frac{\partial^2 \alpha(\theta)}{\partial\theta_i \partial\theta_j}\,|_{\theta=\theta^*}\,.$$

A number of the gradient estimation algorithms mentioned earlier, while capable of producing unbiased estimates of first order partial derivatives, do not easily extend to construction of unbiased estimators for second order partial derivatives. For example, infinitesimal perturbation analysis typically estimates only first derivatives consistently. One notable exception is the likelihood ratio method, for which second derivative estimators are easily constructed.

In general, however, estimation of the second order partial derivatives requires either a non-trivial extension of the methods available for estimation of first order partial derivatives, or use of finite difference estimators in which first differences of first order derivative estimators are utilized. In either case, the need to consistently estimate the $d(d+1)/2$ distinct elements of $H$ makes constructing a consistent estimator $C_n$ for $C$ a highly non-trivial exercise.

## 3 A GENERAL APPROACH TO CONSTRUCTION OF CONFIDENCE REGIONS

Given the difficulties associated with consistent estimation of the covariance matrix $C$, we seek an alternative method. To place our alternative methodology in its appropriate context, we describe the idea in a more general setting.

Suppose that we wish to compute a parameter $\alpha \in \Re^d$ for which there exists an estimator $\alpha_n$ satisfying a CLT. In particular, assume:

**Assumption B.** There exists a symmetric positive definite matrix $\Gamma$ such that $n^{1/2}(\alpha_n - \alpha) \Rightarrow N(0, \Gamma)$ as $n \to \infty$.

We wish to construct confidence regions for $\alpha$ without requiring existence of a consistent estimator for $\Gamma$.

The idea is to compute $\alpha$ by simulating $m$ independent replications of the estimator $\alpha_n$ (thereby consuming roughly $mn$ computer time units). Let $\alpha_{1n}, \alpha_{2n}, \ldots, \alpha_{mn}$ be the resulting $m$ copies of the random vector $\alpha_n$. The natural estimator for $\alpha$ is then $\alpha(m,n) = \frac{1}{m}\sum_{i=1}^{m} \alpha_{in}$. The sample covariance matrix $V(m,n)$ is given by

$$V(m,n) = \frac{1}{m-1}\sum_{i=1}^{m}(\alpha_{in} - \alpha(m,n))^T(\alpha_{in} - \alpha(m,n)).$$

**Theorem 2.** *Under condition **B**,*

*a)* $\sqrt{mn}(\alpha(m,n) - \alpha) \Rightarrow N(0,\Gamma)$ *as* $n \to \infty$*;*

*b) if* $m \geq d+1$*, then*

$$m(\alpha(m,n) - \alpha)^T V(m,n)^{-1}(\alpha(m,n) - \alpha) \Rightarrow \frac{d(m-1)}{m-d} F_{(d,m-d)}$$

*as* $n \to \infty$*, where* $F_{(d,m-d)}$ *is an* $F$ *distributed random variable with* $(d, m-d)$ *degrees of freedom.*

**Proof.** Part a) of Theorem 2 is a simple consequence of the continuous mapping principle. For part b) of Theorem 2, we note that if $N_1(0,\Gamma), \ldots, N_m(0,\Gamma)$ are $m$ independent and identically distributed random vectors with distribution $N(0,\Gamma)$, then the random matrix

$$\frac{1}{m-1}\sum_{i=1}^{m}\left(N_i(0,\Gamma) - \frac{1}{m}\sum_{j=1}^{m} N_j(0,\Gamma)\right)^T \cdot \left(N_i(0,\Gamma) - \frac{1}{m}\sum_{j=1}^{m} N_j(0,\Gamma)\right)$$

is almost surely non-singular (since $m \geq d+1$); see p.208 of Searle (1982).

Because the matrix inverse functional is continuous in a neighborhood of any non-singular matrix, the following function $g : \Re^{mn} \to \Re$ is almost surely continuous in a neighborhood of $(N_1(0,\Gamma), \ldots, N_m(0,\Gamma))$:

$$g(x_1, \ldots, x_m) = \frac{1}{m}\sum_{j=1}^{m} x_j \cdot \left(\frac{1}{m-1}\sum_{l=1}^{m}(x_l - \frac{1}{m}\sum_{j=1}^{m} x_j)^T(x_l - \frac{1}{m}\sum_{j=1}^{m} x_j)\right)^{-1} \cdot \frac{1}{m}\sum_{j=1}^{m} x_j^T.$$

The continuous mapping principle then yields the second result. □

Theorem 2 applies immediately to the construction of confidence regions for stochastic approximations. Suppose that we independently replicate the stochastic approximation $m$ times, thereby yielding $m$ independent copies $\theta_{1n}, \theta_{2n}, \ldots, \theta_{mn}$ of the random vector $\theta_n$. Put

$$\theta(m,n) = \frac{1}{m}\sum_{i=1}^{m}\theta_{in},$$

$$\tilde{V}(m,n) = \frac{1}{m-1}\sum_{i=1}^{m}(\theta_{in} - \theta(m,n))^T(\theta_{in} - \theta(m,n))$$

and set

$$\Lambda_{mn}(z) = \left\{\theta : (\theta(m,n) - \theta)\tilde{V}(m,n)^{-1}(\theta(m,n) - \theta)^T \leq \frac{zd(m-1)}{m-d}\right\}.$$

**Proposition 3.** *Assume the conditions of Theorem 1, and suppose that* $EZ(\theta^*)^T Z(\theta^*)$ *is non-singular. If* $m \geq d+1$*, then*

$$P(\theta^* \in \Lambda_{mn}(z)) \to 1 - \delta$$

*as* $n \to \infty$*, where* $z$ *is selected so that* $P(F_{(d,m-d)} \leq z) = 1 - \delta$.

The proof is an easy application of Theorem 2. Proposition 3 asserts that $\Lambda_{mn}(z)$ is (for large $n$) an approximate $100(1-\delta)\%$ confidence region for $\theta^*$. In contrast to the confidence region approach of Section 2 (in which $C$ was consistently estimated), this procedure can be easily implemented. Note that the current procedure is essentially a "cancellation procedure", in the terminology of Glynn and Iglehart (1990).

One additional advantage of the current approach is that it requires the simulationist to run $m$ independent replications of the stochastic approximation, providing $m$ independent

estimators of the maximizer $\theta^*$. These multiple independent re-starts of the stochastic approximation procedure permit the simulationist to test the question of whether the data collected is consistent with $\alpha(\cdot)$ having a unique local maximizer. This question is of great importance from an applied standpoint, since global maximization is typically the simulationist's objective. We shall pursue this important issue in an expanded version of this paper.

## 4 COMPUTATIONAL RESULTS

In this section, we test the confidence region procedure proposed in Section 3 on two different examples. The first involves a scalar decision parameter, whereas the second example concerns a two-dimensional vector of decision parameters.

### 4.1 Model 1

The first model is an M/M/1 queue with arrival rate $\lambda = 1$ and mean service time

$$\theta \in \Theta = [0.05, 0.95].$$

Therefore, the service time $V$ has density $f_\theta(x) = \exp(-x/\theta)/\theta$. The goal is to minimize

$$\alpha(\theta) = w(\theta) + \frac{1}{\theta}, \tag{8}$$

where $w(\theta)$ is the mean steady-state sojourn time per customer in the system associated with parameter $\theta$. It can be shown that $w(\theta) = \theta/(1-\theta)$. Thus the optimal value $\theta^* = 0.5$. Figure 1 shows the shape of $\alpha(\theta)$ on the interval $[0.05, 0.95]$. This model appeared in L'Ecuyer, Giroux, and Glynn (1994).

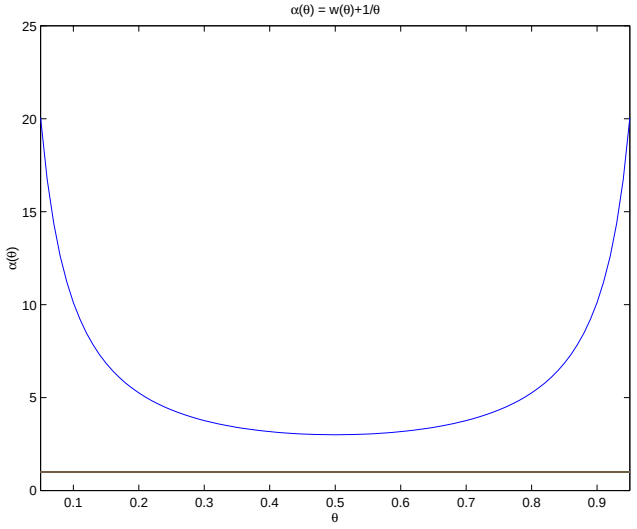


Figure 1: Objective Function of Model 1

We implemented the Robbins-Monro algorithm using likelihood ratio technique to estimate the derivative; see L'Ecuyer and Glynn (1994) for detail on constructing the likelihood ratio derivative estimator for this model. The initial point $\theta_0$ is uniformly chosen from the interval. The step parameter $a$ in (3) is set to 0.1, and the simulation budget at iteration $n$ is set to be proportional to $\sqrt{n}$ (such simulation time allocation has been shown to be empirically efficient in L'Ecuyer, Giroux, and Glynn (1994).)

Table 1 summarizes the numerical results of the simulation runs. The confidence level is 95%, and the 95% confidence interval for the coverage probability is estimated from 100 replications. The coverage probabilities are reasonably close to the nominal level of 95%, and exhibit some degree of convergence to their nominal level as $n$ grows larger. (See, in particular, the results reported for $m = 5$.) Although $n$ seems not have significant affect on the coverage accuracy, long $n$ does narrow the confidence interval, i.e., the variation of $\theta_i$ is smaller.

Table 1: Numerical Summary of Model 1. The Ideal Coverage Probability Is 0.95

| | $m = 3$ | $m = 5$ |
|---|---|---|
| $n$ | CI for coverage prob. | CI for coverage prob. |
| 64 | $0.87 \pm 0.066$ | $0.54 \pm 0.098$ |
| 128 | $0.86 \pm 0.068$ | $0.66 \pm 0.093$ |
| 256 | $0.90 \pm 0.059$ | $0.82 \pm 0.076$ |
| 512 | $0.89 \pm 0.062$ | $0.74 \pm 0.086$ |
| 1024 | $0.86 \pm 0.068$ | $0.80 \pm 0.079$ |
| 2048 | $0.86 \pm 0.068$ | $0.87 \pm 0.066$ |

### 4.2 Model 2

The second model is an M/M/1 queue with arrival rate $\lambda \in [1, 2.5]$ and service rate $\mu \in [3.5, 6]$. Thus, the service time $V$ has density $f_\mu(x) = \mu \exp(-\mu x)$ and the inter-arrival time $U$ has density $f_\lambda(x) = \lambda \exp(-\lambda x)$. The goal is to minimize

$$\alpha(\lambda, \mu) = w(\lambda, \mu) + \frac{1}{\lambda} + \frac{\mu}{4}, \tag{9}$$

where $w(\lambda, \mu)$ is the mean steady-state customer sojourn time for the system associated with arrival rate $\lambda$ and and service rate $\mu$. Again, it can be shown $w(\lambda, \mu) = 1/(\mu - \lambda)$, so that the optimal value $(\lambda^*, \mu^*) = (2, 4)$. The shape of $\alpha(\lambda, \mu)$ is displayed in Figure 2.

Again, we implement the Robbins-Monro algorithm using the likelihood ratio gradient estimation algorithm. The initial $(\lambda_0, \mu_0)$ is chosen uniformly from the rectangle $([1, 2.5], [3.5, 6])$. The step parameter $a$ in (3) is set to 0.5, and the simulation budget at iteration $n$ is set to be proportional to $\sqrt{n}$.

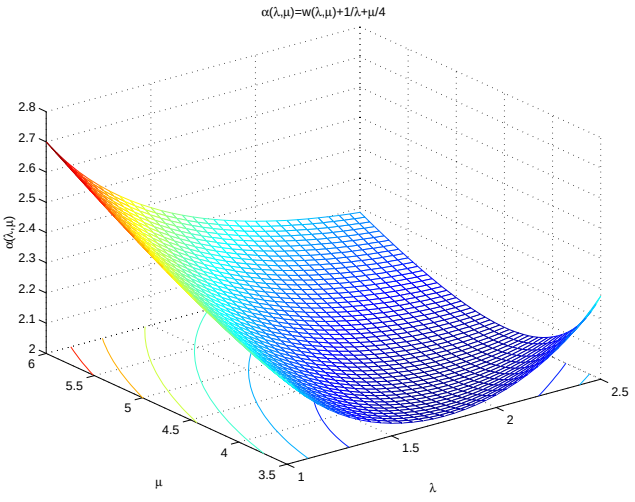


Figure 2: Objective Function of Model 2

Table 2 summarizes the numerical results of the simulation runs. Again, the confidence level is 95%, and the 95% confidence interval for coverage probability is estimated from 100 replications. Again, the empirical evidence is in reasonably close agreement with the theory developed in Section 3.

Table 2: Numerical Summary of Model 2. The Ideal Coverage Probability is 0.95

| | $m = 3$ | $m = 5$ |
|---|---|---|
| $n$ | CI for coverage prob. | CI for coverage prob. |
| 64 | $0.97 \pm 0.034$ | $0.89 \pm 0.062$ |
| 128 | $0.93 \pm 0.050$ | $0.89 \pm 0.062$ |
| 256 | $0.95 \pm 0.043$ | $0.88 \pm 0.064$ |
| 512 | $0.91 \pm 0.056$ | $0.93 \pm 0.050$ |
| 1024 | $0.96 \pm 0.039$ | $0.90 \pm 0.059$ |
| 2048 | $0.96 \pm 0.039$ | $0.91 \pm 0.056$ |

The simulation results from these two models suggest that the procedure described in Section 3 is a pragmatic approach for constructing confidence intervals/regions for stochastic approximation algorithms.

## 5 CONCLUDING REMARKS

In this paper, we have offered an approach to the construction of confidence regions for stochastic approximation algorithms of Robbins-Monro type. A principal advantage of our proposed methodology is the ease with which it can be implemented from a practical standpoint. An additional feature of our proposed method is that its requirement to simulate multiple independent replications of the stochastic approximation procedure offers the simulationist an opportunity to test the question of whether the stochastic approximation has indeed converged to a global optimizer (as opposed to convergence to a local optimizer). We intend to pursue this question at greater depth in future work.

In future work, we hope to study the following issues concerning extensions of the methodology developed in this paper as well as extensions of other methods for simulation optimization:

1. Extend the theory and body of computational experience to the class of Kiefer-Wolfowitz algorithms;
2. Extend the theory and body of computational experience to the averaging algorithms proposed by Polyak and Juditsky (1992), as well as Kushner and Yin (1997);
3. Develop theory and algorithms to address the situation in which the gain constant $a$ is chosen (inadvertently) so small that the matrix $aH + (1/2)I$ has possible positive eigenvalues.

## REFERENCES

Billingsley, P. 1995. *Probability and Measure*. Third Edition. John Wiley & Sons, New York.

Fu, M. C., and J. Q. Hu. 1997. *Conditional Monte Carlo: Gradient estimation and optimization applications*. Kluwer Academic Publisher, Boston.

Gill, P. E., W. Murray, and M.H. Wright. 1981. *Practical Optimization*, Academic Press, London and New York.

Glasserman, P. 1991. *Gradient estimation via perturbation analysis*. Boston: Kluwer.

Glynn, P. W. 1986. Stochastic approximation for Monte Carlo optimization. In *Proceedings of the 1986 Winter Simulation Conference*, 356–365.

Glynn, P. W. 1990. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM* 33:75–84.

Glynn, P. W. and D. L. Iglehart 1990. Simulation output analysis using standardized time series. *Mathematics of Operations Research* 15:1–16.

Kushner, H. J., and G. G. Yin. 1997. *Stochastic approximation algorithms and applications*. New York: Springer-Verlag.

L'Ecuyer, P., and P. W. Glynn. 1994. Stochastic optimization by simulation: Convergence proofs for the GI/G/1 queue in steady-state. *Management Science* 40:1562–1578.

L'Ecuyer, P., N. Giroux, and P. W. Glynn 1994. Stochastic optimization by simulation: Numerical Experiments with the M/M/1 queue in steady-state. *Management Science* 40:1245–1261.

Ljung, L., G. Ch. Pflug, and H. Walk. 1992. *Stochastic approximation and optimization of random systems*. Basel: Birkhäuser Verlag.

Loeve, M. 1977. *Probability Theory*, Vol. 1, 4th edition, Springer-Verlag.

Luenberger, D. G. 1984. *Introduction to Linear and Nonlinear Programming*, Addison-Wesley, Menlo Park, California.

Nevel'son, M. B., and R. Z. Has'minskii. 1973. *Stochastic Approximation and Recursive Estimation*. American mathematical Society, Providence, Rhode Island.

Polyak, B. T., and A. B. Juditsky. 1992. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization* 30:838–855.

Rubinstein, R. Y. 1992. Sensitivity analysis of discrete event systems by the "push-out" method. *Annals of Operations Research* 39:229–237.

Searle, S.R. 1982. *Matrix Algebra Useful for Statistics*. John Wiley, New York.

Williams, D. 1991. *Probability with Martingales*. Cambridge Mathematical Textbooks. Cambridge University Press, Cambridge.

## AUTHOR BIOGRAPHIES

**MING-HUA HSIEH** is an assistant Professor of the Department of Management Information Systems of National Chengchi University, Taiwan. From 1997 to 1999, he was a software designer at Hewlett Packard company, California. He is a member of INFORMS. His research interests include simulation methodology and financial engineering. His e-mail address is <mhsieh@mis.nccu.edu.tw>.

**PETER W. GLYNN** received his Ph.D. from Stanford University, after which he joined the faculty of the Department of Industrial Engineering at the University of Wisconsin- Madison. In 1987, he returned to Stanford, where he is the Thomas Ford Professor of Engineering in the Department of Management Science and Engineering. His research interests include discrete-event simulation, computational probability, queueing, and general theory for stochastic systems. His e-mail address is <glynn@leland.stanford.edu>.