Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

SPSA ALGORITHMS WITH MEASUREMENT REUSE

Mohammed Shahid Abdulla
Shalabh Bhatnagar

Department of Computer Science and Automation
Indian Institute of Science
Bangalore 560012, INDIA

ABSTRACT

Four algorithms, all variants of Simultaneous Perturbation Stochastic Approximation (SPSA), are proposed. The original one-measurement SPSA uses an estimate of the gradient of objective function L containing an additional bias term not seen in two-measurement SPSA. As a result, the asymptotic covariance matrix of the iterate convergence process has a bias term. We propose a one-measurement algorithm that eliminates this bias, and has asymptotic convergence properties making for easier comparison with the two-measurement SPSA. The algorithm, under certain conditions, outperforms both forms of SPSA with the only overhead being the storage of a single measurement. We also propose a similar algorithm that uses perturbations obtained from normalized Hadamard matrices. The convergence w.p. 1 of both algorithms is established. We extend measurement reuse to design two second-order SPSA algorithms and sketch the convergence analysis. Finally, we present simulation results on an illustrative minimization problem.

1 INTRODUCTION

Simultaneous Perturbation Stochastic Approximation (SPSA) is an efficient parameter optimization method that operates under the constraint that only noisy measurements of the objective function L are available at each parameter iterate θ_k. First proposed in Spall (1992), it involves making only two measurements of L at each update epoch k that are obtained by perturbing θ_k along random directions. A plethora of applications and enhancements of this technique can be found at Spall (2001). A variant of SPSA that reduces the number of function measurements made at each iteration k from two to one and establishes the conditions under which a net lower number of observations suffice to attain the same Mean Square Error (MSE) is provided in Spall (1997). However, an impediment in rapid convergence to θ^* is that the algorithm constructs a gradient estimate of L at θ_k that contains an additional error term over the scheme in Spall (1992) and that contributes heavily to the bias in the estimate. A solution to this problem was proposed in Bhatnagar et al. (2003) in the simulation optimization setting, where the perturbation to θ_k in the one-simulation case is generated in a deterministic manner. While this algorithm performs considerably better in practice, the asymptotic convergence properties in the setting of Spall (1992) and Spall (1997) were derived in Xiong, Wang, and Fu (2002) and found to be on par with those of one-measurement SPSA.

In this work, we first propose two first-order algorithms: one using randomly generated perturbations (cf. Section 2) and the other using deterministic perturbations (cf. Section 4). We show convergence w.p. 1 for both the algorithms. For the first algorithm, we also derive the asymptotic convergence properties and compare these with Spall (1992) (cf. Section 3). Further, we design two second-order algorithms based on the measurement-storage concept in Section 5. A numerical example is used to justify our findings (cf. Section 6).

The general structure of gradient descent algorithms is as follows. Suppose $\theta_k \stackrel{\triangle}{=} (\theta_{k,1},...,\theta_{k,p})^T$ where $\theta_{k,i}$, $1 \leq i \leq p$, are the p components of parameter θ_k. Let $G_{k,l}$ be an estimate of the $l-$th partial derivative of the cost $L(\theta_k)$, $l \in \{1,2,...,p\}$. Then,

$$\theta_{k+1,l} = \theta_{k,l} - a_k G_{k,l}(\theta_k), 1 \leq l \leq p, k \geq 0, \qquad (1)$$

where $\{a_k\}$ is a step-size sequence. In the following, we refer to the one-measurement form of SPSA as SPSA2-1R and the two-measurement form as SPSA2-2R following the convention of Bhatnagar et al. (2003). In such a convention, the 'R' refers to perturbations which are randomly obtained, in contrast to deterministic perturbations in Section 4. The trailing '1' in SPSA2-1R refers to the fact that at each iteration, the algorithm makes one measurement. The leading '2' stands for a variant of the algorithm that makes parameter updates at every epoch, in contrast to algorithms like SPSA1-2R which update the parameter after an (increasingly large) number of epochs. The current parameter estimate θ_k is perturbed with a vector $\Delta_k = (\Delta_{k,1}, ..., \Delta_{k,p})^T$ to produce $\theta_k^+ = \theta_k + c_k\Delta_k$, where c_k is a small step-size parameter that satisfies Assumption 1 (below) together with the step-size parameter a_k in (1). The gradient estimates $G_{k,l}(\theta_k)$ used in SPSA2-1R are:

$$\begin{aligned} G_{k,l}(\theta_k) &= \tfrac{L(\theta_k^+)+\varepsilon_k^+}{c_k\Delta_{k,l}} \\ &= \tfrac{L(\theta_k)}{c_k\Delta_{k,l}} + g_l(\theta_k) + \sum_{i=1,i\neq l}^{p} g_i(\theta_k)\tfrac{\Delta_{k,i}}{\Delta_{k,l}} \\ &+ \tfrac{c_k^2\Delta_k^T H(\theta_k)\Delta_k}{2c_k\Delta_{k,l}} + \tfrac{\varepsilon_k^+}{c_k\Delta_{k,l}} \\ &+ \tfrac{c_k^3 L^{(3)}(\bar{\theta}_k)\Delta_k\otimes\Delta_k\otimes\Delta_k}{6c_k\Delta_{k,l}}. \end{aligned} \quad (2)$$

We assume here that L is twice continuously differentiable with bounded third derivative. Note that $H(\theta_k)$ is the Hessian evaluated at θ_k and $L^{(3)}(\bar{\theta}_k)\Delta_k \otimes \Delta_k \otimes \Delta_k = \Delta_k^T(L^{(3)}(\bar{\theta}_k)\Delta_k)\Delta_k$ where $\bar{\theta}_k = \theta_k + \lambda_k c_k \Delta_k$ for some $0 \le \lambda_k \le 1$ and $L^{(3)}$ is the third derivative of objective function $L(\cdot)$, where $\otimes$ denotes the Kronecker product. Also, ε_k^+ corresponds to additive observation noise. Thus, $G_{k,l}$ is a random variable, which we assume is measurable with respect to (w.r.t.) the $\sigma-$algebra $\mathcal{F}_k = \sigma(\theta_i, \Delta_i, 0 \le i \le k-1, \theta_k)$. In contrast, $g_l(\theta_k)$ is the $l-$th component of the derivative of $L(\cdot)$ at θ_k. The gradient of $L(\cdot)$ at θ_k is now defined as $g(\theta_k) = (g_l(\theta_k),\ 1 \le l \le p)^T$. Although not the current object of study, we observe that the estimate of SPSA2-2R needs two measurements of $L(\cdot)$ about θ_k:

$$G_{k,l}(\theta_k) = \tfrac{L(\theta_k^+)+\varepsilon_k^+ - L(\theta_k^-)-\varepsilon_k^-}{2c_k\Delta_{k,l}}.$$

Here $G_{k,l}$ uses function measurements at both θ_k^+ and $\theta_k^- = \theta_k - c_k\Delta_k$ and the measurement noise values at these points are ε_k^+ and ε_k^-, respectively.

We retain all assumptions of Spall (1997), most of which are carried over from Spall (1992). As in Spall (1997), the key assumption requires the measurement noise ε_k^+ to be mean 0: $E(\varepsilon_k^+|\theta_k, \Delta_k) = 0,\ \forall k \ge 1$, and $var(\varepsilon_k^+) \to \sigma_\varepsilon^2$, where σ_ε^2 is some finite constant. The step-size sequences used are of the form $a_k = ak^{-\alpha}$ and $c_k = ck^{-\gamma}$, respectively, where $k \ge 1$, $a, c > 0$, are given constants and with constraints on $0 < \gamma, \alpha \le 1$ such that the following assumption holds

Assumption 1 $\sum_k a_k = \infty$ *and* $\sum_k \frac{a_k^2}{c_k^2} < \infty$.

2 ALGORITHM SPSA2-1UR

The proposed algorithm also has a similar structure as SPSA2-1R and we call this algorithm SPSA2-1UR, the alphabet 'U' indicating 'unbiased'. We utilize the noisy measurement already made at θ_{k-1}^+, the storage of which results in unit space complexity.

Algorithm 1 (SPSA2-1UR)

$$\theta_{k+1,l} := \theta_{k,l} - a_k \tfrac{L(\theta_k^+)+\varepsilon_k^+ - L(\theta_{k-1}^+)-\varepsilon_{k-1}^+}{c_k\Delta_{k,l}},$$

where $k \ge 0$, $1 \le l \le p$. We have in the above,

$$\begin{aligned} G_{k,l}(\theta_k) &= \tfrac{L(\theta_k)-L(\theta_{k-1})}{c_k\Delta_{k,l}} + g_l(\theta_k) \\ &+ \sum_{i=1,i\neq l}^{p} g_i(\theta_k)\tfrac{\Delta_{k,i}}{\Delta_{k,l}} \\ &- \sum_{i=1}^{p} \tfrac{c_{k-1}}{c_k} g_i(\theta_{k-1})\tfrac{\Delta_{k-1,i}}{\Delta_{k,l}} \\ &+ \tfrac{c_k^2\Delta_k^T H(\theta_k)\Delta_k - c_{k-1}^2\Delta_{k-1}^T H(\theta_{k-1})\Delta_{k-1}}{2c_k\Delta_{k,l}} \\ &+ \tfrac{c_k^3 L^{(3)}(\bar{\theta}_k)\Delta_k\otimes\Delta_k\otimes\Delta_k}{6c_k\Delta_{k,l}} \\ &- \tfrac{c_{k-1}^3 L^{(3)}(\bar{\theta}_{k-1})\Delta_{k-1}\otimes\Delta_{k-1}\otimes\Delta_{k-1}}{6c_k\Delta_{k,l}} \\ &+ \tfrac{\varepsilon_k^+ - \varepsilon_{k-1}^+}{c_k\Delta_{k,l}}. \end{aligned}$$

Using a similar analysis as in Proposition 2 of (Spall 1992), we identify below the order of convergence of the bias to 0:

Lemma 1 *Suppose for each $k > K$ for some $K < \infty$, $\{\Delta_{k,i}\}$ are i.i.d., symmetrically distributed about 0, with $\Delta_{k,i}$ independent of θ_j, ε_j^+, $1 \le j < k$. Further let $|\Delta_{k,i}| \le \beta_0$ a.s., $E|\Delta_{k,i}^{-1}| \le \beta_1$, and L be thrice continuously differentiable with $|L_{i_1,i_2,i_3}^{(3)}| \le \beta_2$, $\forall i_1, i_2, i_3 \in \{1, 2, ..., p\}$, for some constants β_0, β_1, and β_2. Then, $E\{G_{k,l}(\theta_k)|\mathcal{F}_k\} = g_l(\theta_k) + O(c_k^2)$, $1 \le l \le p$, a.s.*

Proof: The bias vector represented by $b_k(\theta) = (b_{k,1}(\theta), b_{k,2}(\theta), \ldots, b_{k,p}(\theta))^T$ is defined as:

$$b_k(\theta_k) = E\{G_k(\theta_k) - g(\theta_k)|\theta_k\}, \quad (3)$$

where $G_k(\theta_k) = (G_{k,1}(\theta_k), G_{k,2}(\theta_k), ..., G_{k,p}(\theta_k))^T$. Due to the mean-zero assumption on ε_k^+ and Δ_k^{-1} w.r.t. $\mathcal{F}_k$, we have

$$E\{\tfrac{\varepsilon_k^+ - \varepsilon_{k-1}^+}{\Delta_{k,l}}|\mathcal{F}_k\} = 0.$$

It is crucial here to note that, despite the previous equality, $E\{\varepsilon_k^+ - \varepsilon_{k-1}^+|\mathcal{F}_k, \Delta_k\} \neq 0$. Further, using the properties of independence, symmetry and finite inverse moments of perturbation vector elements (i.e., $\Delta_{k,l}$, $\Delta_{k,i}$, and $\Delta_{k-1,i}$), observe that terms on the RHS of (3) have zero mean, with the exception of

$$b_{k,l}(\theta_k) = E\left\{\frac{c_k^3 L^{(3)}(\bar{\theta}_k)\Delta_k \otimes \Delta_k \otimes \Delta_k}{6 c_k \Delta_{k,l}} | \mathcal{F}_k\right\}.$$

Observe that the bias term here is the same as for SPSA2-1R. The claim is now obtained by the arguments following Equation (3.1) in Lemma 1 of Spall (1992). In particular, note that

$$\begin{aligned} |b_{k,l}(\theta_k)| &\leq \frac{\beta_2 c_k^2}{6} \sum_{i_1}\sum_{i_2}\sum_{i_3} E\left|\frac{\Delta_{k,i_1}\Delta_{k,i_2}\Delta_{k,i_3}}{\Delta_{k,l}}\right| \\ &\leq \frac{\beta_2 c_k^2}{6}\left([p^3-(p-1)^3]\beta_0^2 + (p-1)^3\beta_1\beta_0^3\right) \end{aligned}$$

□

The relation in Lemma 1 is of value in establishing a form of asymptotic normality of the scaled iterate convergence process, see Section 3. Note that Lemma 1 will not hold if normalized Hadamard matrix-based $\{\pm 1\}^p$−valued perturbations Δ_k that were first introduced in Bhatnagar et al. (2003) (Section 4 below explains this deterministic perturbation method in some detail). This is because there is no assurance that the term $\sum_{i=1}^p \frac{c_{k-1}}{c_k} g_i(\theta_{k-1})\frac{\Delta_{k-1,i}}{\Delta_{k,l}}$ will average to 0 as $k\to\infty$, unlike the previous term $\sum_{i=1,i\neq l}^p g_i(\theta_k)\frac{\Delta_{k,i}}{\Delta_{k,l}}$. In such a case, a different method for unbiasing that does not use the immediate past measurement, in the spirit of Section 4 later, would be appropriate. A consequence of the a.s. convergence of the bias $b_k(\theta_k)$ is the strong convergence of the iterates θ_k to a local minimum θ^*. We now state Assumption A2 of Spall (1992) (that is also applicable to the setting of Spall (1997)):

Assumption 2 *$\exists \alpha_0,\alpha_1,\alpha_2 > 0$ and $\forall k$, $E{\varepsilon_k^+}^2 \leq \alpha_0$, $EL^2(\theta_k^+) \leq \alpha_1$, and $\Delta_{k,l}^{-2} \leq \alpha_2$ a.s., for $1 \leq l \leq p$.*

While this does not entail any difference, observe that we use $\Delta_{k,l}^{-2} \leq \alpha_2$ a.s. instead of the original $E\{\Delta_{k,l}^{-2}\} \leq \alpha_2$ in Spall (1992).

Lemma 2 *Under assumptions of Spall (1997), as $k\to\infty$: $\theta_k \to \theta^*$ a.s.*

Proof: Follows almost verbatim as Proposition 1 of Spall (1992). The only modifications are due to a different error process e_k, defined as $e_k(\theta_k) = G_k(\theta_k) - E(G_k(\theta_k)|\theta_k)$. We can thus rewrite recursion (1) as: $\theta_{k+1} = \theta_k - a_k(g(\theta_k) + b_k(\theta_k) + e_k(\theta_k))$. The claim is obtained if the following conditions are satisfied:

(a) $\|b_k(\theta_k)\| < \infty$, $\forall k$ and $b_k(\theta_k)\to 0$ a.s.

(b) $\lim_{k\to\infty} P(\sup_{m\geq k}\|\sum_{i=k}^m a_i e_i(\theta_i)\| \geq \eta) = 0$, for any $\eta > 0$.

where $\|\cdot\|$ represents the Euclidean norm in parameter space $\mathcal{R}^p$. Lemma 1 establishes (a) whilst for (b), notice that $\{\sum_{i=k}^m a_i e_i\}_{m\geq k}$ is a martingale sequence (since $E(e_{i+1}|\mathcal{F}_i) = 0$) and the martingale inequality gives:

$$P(\sup_{m\geq k}\|\sum_{i=k}^m a_i e_i(\theta_i)\| \geq \eta) \leq \eta^{-2} E\|\sum_{i=k}^\infty a_i e_i\|^2.$$

This upper bound equals $\eta^{-2}\sum_{i=k}^\infty a_i^2 E\|e_i\|^2$ since $E(e_i^T e_j) = E(e_i^T E(e_j|\theta_j)) = 0$, $\forall j \geq i+1$.

Further, for $1 \leq l \leq p$ using Hölder's inequality:

$$\begin{aligned} &E\left(G_{i,l}^2(\theta_i)\right) \\ &\leq E(L(\theta_i^+) - L(\theta_{i-1}^+) + \varepsilon_i^+ - \varepsilon_{i-1}^+)^2 \cdot \frac{\|\Delta_{i,l}^{-2}(\omega)\|_\infty}{c_i^2} \\ &\leq 2(\alpha_1+\alpha_0)\alpha_2 c_i^{-2}. \end{aligned}$$

Due to the mean-zero property of $e_{i,l}(\theta_i)$, we have $E\left(G_{i,l}^2(\theta_i)\right) = \left(g_l(\theta_i) + b_{i,l}(\theta_i)\right)^2 + E(e_{i,l}^2(\theta_i))$, thus having $E(e_{i,l}^2(\theta_i)) \leq E(G_{i,l}^2(\theta_i))$, and resulting in $E\|e_i\|^2 \leq 2p(\alpha_1+\alpha_0)\alpha_2 c_i^{-2}$. The square summability of $\frac{a_k}{c_k}$, from Assumption 1, now establishes (b). □

3 ASYMPTOTIC NORMALITY AND COMPARISON

The results obtained so far aid us in establishing the asymptotic normality of a scaled iterate convergence process. We show that

$$k^{\frac{\beta}{2}}(\theta_k - \theta^*) \xrightarrow{D} N(\mu, P\tilde{M}_1 P^T)$$

as $k\to\infty$ where the indicated convergence is in distribution, $\beta = \alpha - 2\gamma > 0$ (given $3\gamma - \frac{\alpha}{2} \geq 0$), and the mean μ is the same as in SPSA2-2R (Spall 1992, Proposition 2) and SPSA2-1R. The orthogonal matrix P above satisfies $P^T aH(\theta^*)P = \mathrm{Diag}(\{\lambda_l\}_{l=1}^p)$, $\lambda_1,\ldots,\lambda_p$ being the p eigenvalues of $aH(\theta^*)$. Unlike Spall (1997), $\tilde{M}_1$ above does not have an $L^2(\theta^*)$ bias; however, it is scaled by a factor of 2. This factor arises due to the use of the additional noisy measurement $L(\theta_{k-1}^+) + \varepsilon_k^+$ in (3). In particular,

$$\tilde{M}_1 = 2a^2c^{-2}\rho^2\sigma_\varepsilon^2 \mathrm{Diag}(\{(2\lambda_l - \beta_+)^{-1}\}_{l=1}^p),$$

where $\beta_+ = \beta$ if $\alpha = 1$ and 0 otherwise, and $E\Delta_{k,l}^{-2} \to \rho^2$. However, as confirmed by Spall (2005), the M_1 in (5) of Spall (1997) should be

$$M_1 = a^2c^{-2}\rho^2(\sigma_\varepsilon^2 + L^2(\theta^*))\mathrm{Diag}(\{(2\lambda_l - \beta_+)^{-1}\}_{l=1}^p), \tag{4}$$

and not $a^2c^{-2}\rho^2(\sigma_\varepsilon^2\mathrm{Diag}(\{(2\lambda_l-\beta_+)^{-1}\}_{l=1}^p) + L^2(\theta^*)I)$ as printed (see Appendix 1 of Abdulla and Bhatnagar 2006 for derivation of M_1). Similarly, Appendix 2 there establishes the form of $\tilde{M}_1$ and mean μ that we claim.

We compare the proposed SPSA2-1UR with SPSA2-2R in the number of measurements of cost function L using variables $\tilde{n}_1$ and n_2, respectively. As in Spall (1997), we consider the case $\alpha = 1$ and $\gamma = \frac{1}{6}$ (giving $\beta = \frac{2}{3}$) and $E((\varepsilon_k^+ - \varepsilon_k^-)^2 | \theta_k, \Delta_k) = 2\sigma_\varepsilon^2$, resulting in $M_2 = \frac{1}{2} a^2 c^{-2} \rho^2 \sigma_\varepsilon^2 \mathrm{Diag}(\{(2\lambda_l - \beta_+)^{-1}\}_{l=1}^p)$ and $\tilde{M}_1 = 4M_2$. This gives us:

$$\frac{\tilde{n}_1}{n_2} \to \frac{1}{2}\left(\frac{4\mathrm{tr}PM_2P^T + \mu^T\mu}{\mathrm{tr}PM_2P^T + \mu^T\mu}\right)^{\frac{3}{2}}, \qquad (5)$$

where tr stands for the trace of the matrix. The ratio above depends upon quantity μ and to achieve $\tilde{n}_1 < n_2$ we need that

$$\mu^T\mu > \left(\frac{4 - 2^{\frac{2}{3}}}{2^{\frac{2}{3}} - 1}\right)\mathrm{tr}PM_2P^T \approx 4.11\mathrm{tr}PM_2P^T.$$

We use n_1 to denote the number of measurements of L made by SPSA2-1R. In the special case $L(\theta^*) = 0$, it is shown in Equation (8) of Spall (1997) that $n_1 < n_2$ when $\mu^T\mu > 0.7024\mathrm{tr}PM_2P^T$. While our result does not compare favorably, the advantage is that (5) holds for all values of $L(\theta^*)$.

The comparison with SPSA2-1R yields an interesting rule of thumb. Using D_λ to represent the diagonal matrix $\mathrm{Diag}(\{2\lambda_l - \beta_+\}_{l=1}^p)$, we have:

$$\begin{aligned}\frac{\tilde{n}_1}{n_1} &\to \left(\frac{\mathrm{tr}P\tilde{M}_1P^T + \mu^T\mu}{\mathrm{tr}PM_1P^T + \mu^T\mu}\right)^{\frac{3}{2}} \\ &= \left(\frac{2a^2c^{-2}\rho^2\sigma_\varepsilon^2 \mathrm{tr}PD_\lambda P^T + \mu^T\mu}{a^2c^{-2}\rho^2(\sigma_\varepsilon^2 + L^2(\theta^*))\mathrm{tr}PD_\lambda P^T + \mu^T\mu}\right)^{\frac{3}{2}}.\end{aligned}$$

Irrespective of μ, D_λ, and P (quantities that may require substantial knowledge of the system), it suffices to have $L^2(\theta^*) > \sigma_\varepsilon^2$ to ensure that $\frac{\tilde{n}_1}{n_1} \le 1$. The experimental results in Section 6 provide verification of these claims.

4 ALGORITHM SPSA2-1UH

We now propose a fast convergence algorithm by modifying SPSA2-1H from §3 of Bhatnagar et al. (2003). The key departure in SPSA2-1H from gradient estimate (2) of SPSA2-1R is that perturbation vectors Δ_k are now deterministically obtained from normalized Hadamard matrices. The kind of matrices considered are the following: Let H_2 be a 2×2 matrix with elements $H_2(1,1) = H_2(1,2) = H_2(2,1) = 1$ and $H_2(2,2) = -1$. Likewise for any $q > 1$, let the block matrices $H_{2^q}(1,1)$, $H_{2^q}(1,2)$, and $H_{2^q}(2,1)$ equal $H_{2^{q-1}}$. Also, let $H_{2^q}(2,2) = -H_{2^{q-1}}$. For a parameter of dimension p, the dimension of the Hadamard matrix needed is 2^q where $q = \lceil \log_2(p+1) \rceil$. Next, p columns from the above matrix H_q are arbitrarily chosen from the $q-1$ columns that remain after the first column is removed. The latter column is removed as it does not satisfy a key property of the perturbation sequence. Each row of the resulting $q \times p$ matrix $\hat{H}$ is now used for the perturbation vector Δ_k in a cyclic manner, i.e., $\Delta_k = \hat{H}(k\%q + 1)$, where $\%$ indicates the modulo operator. Though not shown here, the convergence of SPSA2-1H can be shown as a special case of Proposition 2.5 of Xiong, Wang, and Fu (2002). The proposed algorithm, which we call SPSA2-1UH, has two steps:

Algorithm 2 (SPSA2-1UH) 1. *For* $k \ge 0$, $1 \le l \le p$,

$$\theta_{k+1,l} := \theta_{k,l} - a_k \frac{L(\theta_k^+) + \varepsilon_k^+ - \bar{L}_k}{c_k \Delta_{k,l}}.$$

2. *If* $k\%q = 0$, $\bar{L}_k := L(\theta_k^+) + \varepsilon_k^+$ *else* $\bar{L}_{k+1} := \bar{L}_k$.

In the above, $\bar{L}_k$ changes only periodically in epochs of size q and the algorithm has a unit space requirement. Given index k, define $\bar{k} = \max\{m : m < k, m\%q = 0\}$. For SPSA2-1UH, (2) is now modified to:

$$\begin{aligned}G_{k,l}(\theta_k) &= \frac{L(\theta_k) - L(\theta_{\bar{k}})}{c_k \Delta_{k,l}} + g_l(\theta_k) \\ &\quad + \sum_{i=1, i \ne l}^{p} \frac{\Delta_{k,i}}{\Delta_{k,l}} g_i(\theta_k) - \sum_{i=1}^{p} \frac{c_{\bar{k}}}{c_k}\frac{\Delta_{\bar{k},i}}{\Delta_{k,l}} g_i(\theta_{\bar{k}}) \\ &\quad + O(c_k) + \frac{\varepsilon_k^+ - \varepsilon_{\bar{k}}^+}{c_k \Delta_{k,l}},\end{aligned}$$

where $O(c_k)$ contains higher order terms. Since $\Delta_{\bar{k}} = \underline{1}$, we have $\frac{\Delta_{\bar{k},i}}{\Delta_{k,l}} = \frac{1}{\Delta_{k,l}}$, $\forall 1 \le i, l \le p$ and $\forall k$. Therefore, it can be shown as in Lemma 3.5 of Bhatnagar et al. (2003) that the fourth term above averages to 0 over q steps as $k \to \infty$, thus settling the problem posed in §2. In passing, we also note that step 2 can be written as $k\%q = m$ for any given m for $0 \le m \le q-1$.

4.1 Convergence Analysis

We can now formally establish convergence w.p. 1 of θ_k. The original SPSA2-1H algorithm can be expanded as follows

$$\begin{aligned}\theta_{k+1} &= \theta_k - a_k \Delta_k^{-1} \Delta_k^T g(\theta_k + \lambda_k c_k \Delta_k) \\ &\quad - \frac{a_k}{c_k} L(\theta_k) \Delta_k^{-1} - \frac{a_k}{c_k} \varepsilon_k^+ \Delta_k^{-1}, \qquad (6)\end{aligned}$$

where $0 \le \lambda_k \le 1$. Here Δ_k^{-1} is the vector $\Delta_k^{-1} = (\frac{1}{\Delta_{k,1}}, \ldots, \frac{1}{\Delta_{k,p}})^T$. This recursion is now presented in the manner of Equation 6 of Xiong, Wang, and Fu (2002), with r_k, d_k and e_k^+ there replaced by Δ_k^{-1}, Δ_k and ε_k^+, respectively:

$$\begin{aligned}\theta_{k+1} &= \theta_k - a_k g(\theta_k) \\ &\quad - a_k \Delta_k^{-1} \Delta_k^T \{g(\theta_k + \lambda_k c_k \Delta_k) - g(\theta_k)\} \\ &\quad - a_k \{\Delta_k^{-1} \Delta_k^T - I\} g(\theta_k) - \frac{a_k}{c_k} L(\theta_k) \Delta_k^{-1}\end{aligned}$$

$$-\tfrac{a_k}{c_k}\varepsilon_k^+\Delta_k^{-1}.$$

In the manner of (6), the SPSA2-1UH recursion is written as:

$$\theta_{k+1} = \theta_k - a_k\Delta_k^{-1}\Delta_k^T g(\theta_k+\lambda_k c_k\Delta_k) \\ - \tfrac{a_k}{c_k}(L(\theta_k)-L(\theta_{\bar{k}}^+))\Delta_k^{-1} - \tfrac{a_k}{c_k}(\varepsilon_k^+ - \varepsilon_{\bar{k}}^+)\Delta_k^{-1}$$

which can be expanded as

$$\begin{aligned}\theta_{k+1} = \theta_k &- a_k g(\theta_k) \\ &- a_k\Delta_k^{-1}\Delta_k^T\{g(\theta_k+\lambda_k c_k\Delta_k)-g(\theta_k)\} \\ &- a_k\{\Delta_k^{-1}\Delta_k^T - I\}g(\theta_k) \\ &+ a_k\Delta_k^{-1}\Delta_{\bar{k}}^T\{g(\theta_{\bar{k}}+\lambda_{\bar{k}}c_{\bar{k}}\Delta_{\bar{k}})-g(\theta_{\bar{k}})\} \\ &+ a_k\Delta_k^{-1}\Delta_{\bar{k}}^T g(\theta_{\bar{k}}) \\ &- \tfrac{a_k}{c_k}\left(L(\theta_k)-L(\theta_{\bar{k}})-\varepsilon_k^+ + \varepsilon_{\bar{k}}^+\right)\Delta_k^{-1}.\end{aligned}$$

However, we need to make a non-restrictive assumption:

Assumption 3 *The function g (cf. A1 of Xiong, Wang, and Fu 2002) is uniformly continuous.*

Theorem 1 *Under Assumptions from Spall (1997) and 3, Algorithm 2 produces iterates θ_k where $\theta_k \to \theta^*$ w.p. 1.*

Proof: We first show that terms in (7) and (7) are error terms in the nature of $e_i(\theta_i)$ in condition (b) of Lemma 2. In particular, we show that these satisfy the conditions (B1) and (B4) in Xiong, Wang, and Fu (2002). We reproduce these two conditions for clarity:

(B1) For some $T>0$,

$$\lim_{n\to\infty}\left(\sup_{n\le k\le m(n,T)} \|\sum_{i=n}^{k} a_i e_i\|\right) = 0,$$

where $m(n,T) \stackrel{\Delta}{=} \max\{k : a_n+\ldots+a_k \le T\}$.

(B4) There exist sequences $\{e_{1,n}\}$ and $\{e_{2,n}\}$ with $e_n = e_{1,n}+e_{2,n}$ for all n such that $\sum_{k=1}^{n} a_k e_{1,k}$ converges, and $\lim_{n\to\infty} e_{2,n} = 0$.

Observe that due to $\lim_{k\to\infty} c_{\bar{k}} = 0$ and the uniform continuity of g, $\Delta_k^{-1}\Delta_{\bar{k}}^T\{g(\theta_{\bar{k}}+\lambda_{\bar{k}}c_{\bar{k}}\Delta_{\bar{k}})-g(\theta_{\bar{k}})\}$ satisfies (B4). Since $\lim_{k\to\infty}\theta_{\bar{k}} - \theta_{\overline{k+1}} = 0$, $\Delta_k^{-1}\Delta_{\bar{k}}^T g(\theta_{\bar{k}})$ satisfies (B1). This is shown by applying Lemma 2.2 of Xiong, Wang, and Fu (2002) with the substitution $\{x_n\}$ where $x_n = 1\ \forall n\ge 1$, $\{\Delta_n^{-1}\Delta_{\bar{n}}\}$ and $\{g(\theta_{\bar{n}})\}$ for $\{c_n\}$, $\{r_n\}$, and $\{e_n\}$, respectively. We have $\forall k$,

$$\tfrac{|(L(\theta_k)-L(\theta_{\bar{k}}))-(L(\theta_{k+1})-L(\theta_{\overline{k+1}}))|}{c_k} \le \\ \tfrac{|L(\theta_k)-L(\theta_{\bar{k}})|}{c_k} I_{k\%q=0} + \tfrac{|L(\theta_k)-L(\theta_{k+1})|}{c_k}.$$

We consider the first term on the RHS, the second follows similarly:

$$\begin{aligned}\tfrac{|L(\theta_k)-L(\theta_{\bar{k}})|}{c_k} I_{k\%q=0} &\le \tfrac{M_0}{c_k}\sum_{m=\bar{k}}^{k-1}\|\theta_{m+1}-\theta_m\| I_{k\%q=0} \\ &\quad + \tfrac{M_0}{c_k}\sum_{m=\bar{k}}^{k-1}\left(M_1\tfrac{a_m}{c_m} + M_2\tfrac{a_m}{c_m}|\varepsilon_m^+|\right) I_{k\%q=0} \\ &\le \left(\tfrac{M_0 M_1 q}{c_k}\tfrac{a_{\bar{k}}}{c_{\bar{k}}} + \tfrac{M_0 M_2}{c_k}\tfrac{a_{\bar{k}}}{c_{\bar{k}}}\sum_{m=\bar{k}}^{k-1}|\varepsilon_m^+|\right) I_{k\%q=0},\end{aligned}$$

where M_0, M_1, and M_2 represent appropriate bounds. The summability of $\{\frac{a_k a_{\bar{k}}}{c_k c_{\bar{k}}}\}$ is obtained using Assumption 1 — implying that the LHS satisfies (B1). This fact is used when we apply Lemma 2.2 of Xiong, Wang, and Fu (2002) again (with $\{\Delta_n^{-1}\}$, $\{L(\theta_n)-L(\theta_{\bar{n}})\}$ replacing $\{r_n\}$ and $\{e_n\}$, respectively, and $\{c_n\}$ as is) to see that $\frac{L(\theta_k)-L(\theta_{\bar{k}})}{c_k}\Delta_k^{-1}$ satisfies (B1). We now consider the last term, i.e., $\frac{\varepsilon_k^+ - \varepsilon_{\bar{k}}^+}{c_k}\Delta_k^{-1}$. However, now the noise term $\bar{\varepsilon}_{k,l} = \frac{\varepsilon_k^+ - \varepsilon_{\bar{k}}^+}{\Delta_{k,l}}$ is not mean 0 w.r.t. $\mathcal{F}_k$ but letting $\tilde{k} = k+q$, $\forall k$ we see that $E(\bar{\varepsilon}_{\tilde{k},l}|\mathcal{F}_k) = 0$. This results in $\left\{\sum_{k=\tilde{n}}^{m}\frac{a_k}{c_k}\bar{\varepsilon}_k\right\}_{m\ge\tilde{n}}$ being a martingale sequence w.r.t. $\mathcal{F}_n$, where we again utilize the inequality

$$P\left(\sup_{m\ge\tilde{n}}\left\|\sum_{k=\tilde{n}}^{m}\tfrac{a_k}{c_k}\bar{\varepsilon}_k\right\| \ge \eta\right) \quad\le\quad \eta^{-2}\sum_{k=\tilde{n}}^{\infty}\left(\tfrac{a_k}{c_k}\right)^2 E\|\bar{\varepsilon}_k\|^2,$$

the LHS modified to obtain

$$P\left(\sup_{m\ge n}\left\|\sum_{k=n}^{m}\tfrac{a_k}{c_k}\bar{\varepsilon}_k\right\| \ge \eta\right) \le \eta^{-2}\sum_{k=\tilde{n}}^{\infty}\left(\tfrac{a_k}{c_k}\right)^2 E\|\bar{\varepsilon}_k\|^2 \\ + P\left(\sup_{\tilde{n}>m\ge n}\left\|\sum_{k=n}^{m}\tfrac{a_k}{c_k}\bar{\varepsilon}_k\right\| \ge \eta\right).$$

The square summability of $\frac{a_k}{c_k}$ and boundedness of $\bar{\varepsilon}_k$ result in quantities on the RHS vanishing as $n\to\infty$. The proof of Proposition 2.3 in Xiong, Wang, and Fu (2002) handles the terms in the RHS of (7), thus resulting in the claim. □

5 SECOND-ORDER ALGORITHMS

We now propose two second order SPSA algorithms, both re-use noisy function measurements. The first algorithm — called 2SPSA-3UR since it is a modification of 2SPSA of Spall (2000) — makes three measurements in the vicinity of each iterate θ_k and re-uses the current gradient estimate $G_k(\theta_k)$ to estimate the Hessian matrix $H_k(\theta_k)$ at θ_k. The second algorithm 2SPSA-2UR makes two measurements at θ_k and reuses the value $L(\theta_{k-1}^+)$ in the Hessian matrix estimation. A third algorithm, 2SPSA-1UR, makes a single measurement per iteration and is described in Abdulla and Bhatnagar (2006). Second-order SPSA algorithms, which are stochastic analogs of the Newton-Raphson algorithm, are also proposed in Spall (2000) and Bhatnagar (2005). The two algorithms that we propose are modifications of 3SA and 2SA of Bhatnagar (2005) although differing in a few details.

- Unlike the 2SPSA in Spall (2000), all three algorithms 2SPSA-nUR $n = 1,2,3$ use an additional a_k-like step-size sequence $\{b_k\}$ (not to be confused with the bias term $b_k(\theta_k)$ in Lemma 1) in the recursion to compute H_k. Such an additional step-size $\{b_k\}$ is employed in all the four second-order SPSA algorithms described in Bhatnagar (2005). The property of b_k relative to a_k is the well-known 'two-timescale' property: $\sum_k b_k = \infty$, $\sum_k b_k^2 < \infty$ and $a_k = o(b_k)$.
- Similar to 2SPSA, we employ an auxiliary perturbation sequence $\{\tilde{\Delta}_k\}$ with the same properties as the original $\{\Delta_k\}$, although independently generated. There is also an associated scaling parameter $\{\tilde{c}_k\}$. We will also require an analog of Assumption 1: replace the pair (a_k, c_k) in Assumption 1 with the pairs $(a_k, \tilde{c}_k)$, $(b_k, \tilde{c}_k)$, (b_k, c_k).
- We use the 'unbiasing' concept by storing past or current measurements of L and gradient estimate G. However, unlike the unit storage overhead in SPSA2-1UR and SPSA2-1UH, this retention of the current estimate of gradient G arguably costs $O(p)$ in storage. Second-order algorithms of Spall (2000) and Bhatnagar (2005) do implicitly assume memory to store, manipulate and multiply Hessian estimates H_k — which are $O(p^2)$ data structures.

5.1 2SPSA-3UR

As used in Bhatnagar (2005), the function Γ used below maps from the set of general $p \times p$ matrices to the set of positive definite matrices. There are many possible candidates for such a Γ, as explained in §II-D of Spall (2000) where the notation f_k is used.

$$\begin{aligned} \theta_{k+1} &= \theta_k - a_k H_k^{-1} G_k(\theta_k) && (7) \\ H_k &= \Gamma(\bar{H}_k) \\ \bar{H}_k &= \bar{H}_{k-1} + b_{k-1}(\hat{H}_k - \bar{H}_{k-1}) \end{aligned}$$

where

$$\hat{H}_k = \frac{1}{2}\left[\frac{\delta G_k^T}{c_k \Delta_k} + \left(\frac{\delta G_k^T}{c_k \Delta_k}\right)^T\right]$$

$$\delta G_k = G_k^1(\theta_k^+) - G_k(\theta_k).$$

Note the re-use of the current gradient estimate $G_k(\theta_k)$ in the second recursion above. This estimate is computed as in the algorithm SPSA2-2R. In addition to θ_k^+ and θ_k^-, we now employ the shorthand notation $\theta_k^{++} = \theta_k + c_k\Delta_k + \tilde{c}_k\tilde{\Delta}_k$. Similarly, we denote the measurement noise incurred at θ_k^{++} as ε_k^{++}. The terms used above are:

$$\begin{aligned} G_k^1(\theta_k^+) &= \frac{\tilde{\Delta}_k^{-1}}{\tilde{c}_k}\left(L(\theta_k^{++}) + \varepsilon_k^{++} - L(\theta_k^+) - \varepsilon_k^+\right), \text{ and} \\ G_k(\theta_k) &= \frac{\Delta_k^{-1}}{2c_k}\left(L(\theta_k^+) + \varepsilon_k^+ - L(\theta_k^-) - \varepsilon_k^-\right). \end{aligned}$$

Appendix 3 of Abdulla and Bhatnagar (2006) contains the derivation regarding $E(\hat{H}_k|\mathcal{F}_k) = H(\theta_k) + O(c_k)$. The convergence analysis of $\theta_k \to \theta^*$ proceeds as in Bhatnagar (2005), outlines of which we explain here. We construct a time-axis using the step-size b_k: assume that $t(n) = \sum_{m=0}^n b_m$ and define a function $H(\cdot)$ as $H(t(k)) = H_k$ with linear interpolation between $[t(k), t(k+1))$. Similarly define function $\theta(\cdot)$ by setting $\theta(t(k)) = \theta_k$ and linear interpolation on the interval $[t(k), t(k+1)]$. Let $T > 0$ be a scalar and define a sequence $\{T_k\}$ as $T_0 = 0$ and $T_k = \min\{t(m)|t(m) \geq T_{k-1} + T\}$. Then $T_k = t(m_k)$ for some m_k and $T_k - T_{k-1} = T$. Now define functions $\bar{H}(\cdot)$ and $\bar{\theta}(\cdot)$ as $\bar{H}(T_k) = H(t_{m_k}) = H_k$ and $\bar{\theta}(T_k) = \theta(t_{m_k}) = \theta_k$, and for $t \in [T_k, T_{k+1}]$, the evolution is according to the system of ODEs:

$$\begin{aligned} \dot{\bar{H}}_{i,j}(t) &= \bigtriangledown_{i,j}^2 L(\bar{\theta}(t)) - \bar{H}_{i,j}(t) \\ \dot{\bar{\theta}}(t) &= 0, \end{aligned}$$

where $\bigtriangledown_{i,j}^2$ indicates $\frac{\partial^2 L(\bar{\theta})}{\partial\theta_i \partial\theta_j}$. One can now show as in Lemma A.8 of Bhatnagar (2005) that

$$\lim_{k\to\infty} \sup_{t\in[T_k, T_{k+1}]} \|H(t) - \bar{H}(t)\| = 0$$

and

$$\lim_{k\to\infty} \sup_{t\in[T_k, T_{k+1}]} \|\theta(t) - \bar{\theta}(t)\| = 0.$$

Recursion (7) can now be shown to asymptotically track the trajectories of the ODE $\dot{\theta}(t) = -H^{-1}(\theta(t)) \bigtriangledown L(\theta(t))$ in a similar manner as above on the slower timescale $\{a_k\}$ (cf. Theorem 3.1 of Bhatnagar (2005)).

5.2 2SPSA-2UR

The proposed algorithm can be understood in terms of the gradient-free four-measurement algorithm 2SPSA of Spall (2000). In Footnote 6 of that article, the SPSA2-1R analog of 2SPSA was not considered due to the inherent variability of both estimates G_k and H_k if the one-measurement form of SPSA were to be used. We employ the technique of the proposed SPSA2-1UR to overcome this hurdle, in the process reducing the number of function measurements required from 4 to 2. The family of recursions (7) is retained with the differences that

$$\begin{aligned} G_k(\theta_k) &= \frac{\tilde{\Delta}_k^{-1}}{\tilde{c}_k}\left(L(\theta_k^{++})+\varepsilon_k^{++}-L(\theta_k^{+})-\varepsilon_k^{+}\right), \\ \hat{H}_k &= \frac{1}{2}\left[\frac{\delta G_k^T}{c_k\Delta_k}+\left(\frac{\delta G_k^T}{c_k\Delta_k}\right)^T\right], \text{ and} \\ \delta G_k &= G_k^1(\theta_k^+)-\tilde{G}_k^1(\theta_k), \end{aligned}$$

followed by a correction of the diagonal terms in $\hat{H}_k$:

$$\hat{H}_k(i,i) := \hat{H}_k(i,i)+\frac{L(\theta_k^+)+\varepsilon_k^+-L(\theta_{k-1}^+)-\varepsilon_{k-1}^+}{c_k^2} \tag{8}$$

where the measurement at θ_{k-1}^+ is reused. This correction assumes that Δ_k and $\tilde{\Delta}_k$ are both Bernoulli distributed over $\{+c,-c\}$ for some $c>0$, although a similar corrective term can be derived for other classes of perturbations. Appendix 4 of Abdulla and Bhatnagar (2006) derives the steps leading to this correction. In the above,

$$\tilde{G}_k^1(\theta_k) = \frac{L(\theta_k^+)+\varepsilon_k^+-L(\theta_{k-1}^+)-\varepsilon_{k-1}^+}{c_k}\Delta_k^{-1}$$

and $G_k^1(\theta_k^+)$ is as in 2SPSA-3UR. Note that $\tilde{G}_k^1(\theta_k)$ is precisely the gradient estimate in the algorithm SPSA2-1UR of Section 2. Also, $G_k(\theta_k)=G_k^1(\theta_k^+)$ indicating that a re-use of the gradient estimate G_k is being made to compute the Hessian estimate $\hat{H}_k$. Here, G_k is computed using a one-sided difference just as in 2SA of Bhatnagar (2005). Such an estimate still uses two measurements, yet is different from the one-measurement form of G_k as in SPSA2-1R or the unbiased G_k of SPSA2-1UR proposed in Section 2.

In place of a detailed convergence analysis, we provide an outline:

$$\begin{aligned} & E(G_k^1(\theta_k^+)-\tilde{G}_k^1(\theta_k)|\theta_k,\Delta_k) \\ &= E(G_k^1(\theta_k^+)|\theta_k,\Delta_k)-E(\tilde{G}_k^1(\theta_k)|\theta_k,\Delta_k) \\ &= g(\theta_k^+)-\frac{L(\theta_k^+)-L(\theta_{k-1}^+)-\varepsilon_{k-1}^+}{c_k}\Delta_k^{-1}. \end{aligned}$$

Also,

$$E\left(\frac{g(\theta_k^+)-\frac{L(\theta_k^+)-L(\theta_{k-1}^+)-\varepsilon_{k-1}^+}{c_k}\Delta_k^{-1}}{c_k}(\Delta_k^{-1})^T|\mathcal{F}_k\right) = H(\theta_k)+O(c_k),$$

the proof being in Appendix 4 of Abdulla and Bhatnagar (2006). Here $H(\theta_k)$ is the Hessian at θ_k while the error term corresponds to a matrix with an induced norm bounded above by $O(c_k)$. We write this as:

$$E\left(\frac{\delta G_{k,i}}{c_k\Delta_{k,j}}|\mathcal{F}_k\right) = H_{i,j}(\theta_k)+O(c_k), 1\le i,j\le p.$$

The convergence analysis uses the ODE technique of 2SPSA-3UR, and since G_k is the same as algorithm 2SA of Bhatnagar (2005), convergence of θ_k is assured using Theorem 3.3 of Bhatnagar (2005). The convergence can also be obtained in a manner similar to that of Theorems 1a and 2a of Spall (2000). Note that Spall (2000) uses the step-size $b_{k+1}=\frac{1}{k+1}$. Our algorithm is applicable for more general step-sizes as long as the requirement $a_k=o(b_k)$ is met.

6 NUMERICAL EXAMPLE

We first compare algorithm SPSA2-2R of Spall (1992) with the proposed SPSA2-1UR using the setting of Spall (1997). In particular, the objective function used is

$$L_b(\theta) = b+\theta^T\theta+0.1\sum_{i=1}^{5}\theta_i^3+0.01\sum_{i=1}^{5}\theta_i^4,$$

with $\theta^*=0$ and $L_b(\theta^*)=0$ for all b. We keep $b=0$ for comparison with SPSA2-2R and change to 0.1 for comparison with SPSA2-1R. We use $a=c=1$, $\alpha=6\gamma=1$ and $\theta_0=0.1\underline{1}$ (i.e., the vector with 0.1 in all its components) in all the experiments. Assume that ε_k^+ are i.i.d., mean-zero, Gaussian random variables with variance σ_ε^2. The formula for asymptotic normality derived previously lets us consider two cases for the observation noise:

1. $\sigma_\varepsilon=0.1$ where $\frac{\tilde{n}_1}{n_2}\to 1.30$, and
2. $\sigma_\varepsilon=0.07$ where $\frac{\tilde{n}_1}{n_2}\to 0.93$, respectively.

Each run of the SPSA2-2R algorithm is for 2000 iterations, thus making 4000 observations of the objective function. Table 1 summarizes the results, the mean square error (MSE) obtained being over 100 runs of each algorithm. The MSE values for SPSA2-2R are less when compared to SPSA2-1UR, the proportion being 0.93 and 0.92, respectively for the two cases. However, this ratio improves if we use the SPSA2-1UH algorithm, which we compare with the analogous SPSA2-2H algorithm in Table 2.

Table 1: Mean Square Error and No. of Iterations

	$\sigma_\varepsilon=0.1$		$\sigma_\varepsilon=0.07$	
Algorithm	MSE	Iter.	MSE	Iter.
SPSA2-2R	0.0135	2000	0.0130	2000
SPSA2-1UR	0.0145	5200	0.0144	3600

While we have no asymptotic normality results for SPSA2-1UH, the performance obtained is better than that of SPSA2-1UR. We also observe the performance of SPSA2-1UR vis-a-vis SPSA2-1R in Table 3. Possibly due to the larger number of iterations required to achieve asymptotic normality, the MSE is always higher. A notable change in the behaviour of SPSA2-1R is the higher MSE when $b = 0.1$. This is due to the $L^2(\theta^*)$ bias term in (6). Note that we use $\sigma_e = 0.1$ in both the above comparisons.

We compare the second-order algorithms on the same setting. For algorithms 2SPSA-3UR and 2SPSA-2UR, we use $\tilde{\Delta}_{k,i} \in \{+1,-1\}$ while the step-size $\tilde{c}_k$ was the same as c_k, with $b_k = \frac{1}{k^{0.55}}$. We used a similar projection operator $\Gamma(\cdot)$ as in the experiments of Bhatnagar (2005), i.e., choose the diagonal elements $\bar{H}_k(i,i)$, $1 \leq i \leq p$ of the Hessian estimate and then truncate to interval $[0.1, 10.0]$. This upper bound of 10 on $H_k(i,i)$ was justified since typically two-timescale algorithms are known to perform better with an additional averaging on the faster timescale, where $L >> 1$ measurements are made. Since recourse to multiple measurements is ruled out in this setting, we chose to prune the fluctuations in the diagonal terms $\bar{H}_k(i,i)$.

We compare 2SPSA-3UR with the four-measurement 2SPSA of Spall (2000) to obtain the results in Table 4. We run both algorithms in such a manner that the number of function evaluations is the same: 4000. The convergence of the bias (of $\hat{H}_k$) in 2SPSA-3UR is $O(c_k)$, resulting in problems establishing any asymptotic normality results. As a consequence, there is no clear set of parameters for which 2SPSA-3UR would outperform 2SPSA. This slower order of convergence may also be responsible for the poor performance of the algorithm. The experiments indicate the disconnect between finite-time performance of the second-order algorithms vis-a-vis the robust convergence behaviour expected from a Newton-Raphson method. We chose this numerical setting to compare the proposed algorithms with those in the literature. The work Zhu and Spall (2002) explores both finite-time performance and a computationally-efficient second-order SPSA algorithm. The difference with Zhu and Spall (2002) would lie in choosing the Γ operator of (7). This is an issue also identified in Bhatnagar (2005), from where we chose the 3SA and 2SA algorithms for modification. Table 5 compares performance of 2SPSA-2UR w.r.t. 2SA of Bhatnagar (2005). The algorithms are more or less on par with each other.

Table 2: Mean Square Error and No. of Iterations

	$\sigma_\varepsilon = 0.1$		$\sigma_\varepsilon = 0.07$	
Algorithm	MSE	Iter.	MSE	Iter.
SPSA2-2H	0.0133	2000	0.0127	2000
SPSA2-1UH	0.0109	5200	0.0109	3600

Table 3: Mean Square Error and No. of Iterations

	$b = 0$		$b = 0.1$	
Algorithm	MSE	Iter.	MSE	Iter.
SPSA2-1R	0.0443	4000	0.0492	4000
SPSA2-1UR	0.0132	6000	0.0147	4000

Table 4: Comparison of 2SPSA-3UR

	$\sigma_\varepsilon = 0.1$		$\sigma_\varepsilon = 0.07$	
Algorithm	MSE	Iter.	MSE	Iter.
2SPSA	0.037	1000	0.039	1000
2SPSA-3UR	0.078	1333	0.073	1333

Table 5: Comparison of 2SPSA-2UR

	$\sigma_\varepsilon = 0.1$		$\sigma_\varepsilon = 0.07$	
Algorithm	MSE	Iter.	MSE	Iter.
2SA	0.076	2000	0.077	2000
2SPSA-2UR	0.072	2000	0.078	2000

7 FUTURE DIRECTIONS

The asymptotic convergence properties of SPSA2-1H have been theoretically shown to be on par with SPSA2-1R in Proposition 2.5 of Xiong, Wang, and Fu (2002). Yet, it is unclear why SPSA2-1H performs better in practice and this represents an avenue for future investigation. Also of interest is the possibility of reducing the scale factor 2 in the asymptotic covariance matrix $\tilde{M}_1$ using an average of past measurements $L(\theta_{k-j})$, $j > 1$. Whether online function regression mechanisms will serve as a 'critic' to speed up SPSA gradient descent by yielding an approximation of the objective function remains to be seen. Such an arrangement would place the resulting algorithm in-between the accepted forms of 'gradient-free' and 'gradient-based' methods. Further, in line with the asymptotic normality results of both first and second order SPSA algorithms, work such as Konda and Tsitsiklis (2004) that identifies rate of convergence of two-timescale recursions should be useful.

ACKNOWLEDGMENTS

This work was supported in part by grant number SR/S3/EE/43/2002-SERC-Engg from the Department of Science and Technology, Government of India. We thank Dr. James C. Spall of Applied Physics Laboratory, Johns Hopkins University for his kind help in identifying the correct expression for the covariance matrix M_1 of the algorithm SPSA2-1R.

REFERENCES

Abdulla, M., and S. Bhatnagar. 2006. SPSA algorithms with measurement reuse. Technical Report 2006–8, Department of CSA, Indian Institute of Science. URL: `<http://archive.csa.iisc.ernet.in/TR/2006/8/>`.

Bhatnagar, S. 2005. Adaptive multivariate three–timescale stochastic approximation algorithms for simulation based optimization. *ACM Transactions on Modeling and Computer Simulation* 15 (1): 74–107.

Bhatnagar, S., M. Fu, S. Marcus, and I.-J. Wang. 2003. Two–timescale simultaneous perturbation stochastic approximation using deterministic perturbation sequences. *ACM Transactions on Modeling and Computer Simulation* 13 (4): 180–209.

Konda, V. R., and J. N. Tsitsiklis. 2004. Convergence rate of linear two-time-scale stochastic approximation. *Annals Of Applied Probability* 14 (2): 796–819.

Spall, J. C. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control* 37 (1): 332–341.

Spall, J. C. 1997. A one–measurement form of simultaneous perturbation stochastic approximation. *Automatica* 33 (1): 109–112.

Spall, J. C. 2000. Adaptive stochastic approximation by the simultaneous perturbation method. *IEEE Transactions on Automatic Control* 45 (10): 1839–1853.

Spall, J. C. 2001. Selected references on theory, applications, and numerical Analysis. URL <http://www.jhuapl.edu/SPSA/Pages/References-List_Ref.htm>, accessed Aug 2006.

Spall, J. C. 2005. Personal communication. Dated 5 October.

Xiong, X., I.-J. Wang, and M. C. Fu. 2002. Randomized-direction stochastic approximation algorithms using deterministic sequences. In *Proceedings of the 2002 Winter Simulation Conference*, 285–291.

Zhu, X., and J. C. Spall. 2002. A modified second-order SPSA optimization algorithm for finite samples. *International Journal of Adaptive Control and Signal Processing* 16:397–409.

AUTHOR BIOGRAPHIES

MOHAMMED SHAHID ABDULLA is a Ph.D. student in the Department Of Computer Science and Automation of the Indian Institute Of Science, Bangalore, India. His e-mail address is <shahid@csa.iisc.ernet.in>.

SHALABH BHATNAGAR is an Assistant Professor in the Department Of Computer Science and Automation of the Indian Institute Of Science, Bangalore, India. His e-mail address is <shalabh@csa.iisc.ernet.in>, and his web page is <http://drona.csa.iisc.ernet.in/~shalabh/>.