

Stochastic Approximation on Discrete Sets Using Simultaneous Difference Approximations

Stacy D. Hill, László Gerencsér, and Zsuzsanna Vágó

Abstract—A stochastic approximation method for optimizing a class of discrete functions is considered. The procedure is a version of the Simultaneous Perturbation Stochastic Approximation (SPSA) method that has been modified to obtain a stochastic optimization method for cost functions defined on a grid of points in Euclidean p -space having integer components. We discuss the algorithm and examine its convergence properties.

I. INTRODUCTION

THE simultaneous perturbation stochastic approximation (SPSA) method [1] is a tool for solving continuous optimization problems in which the cost function is differentiable but analytically unavailable or difficult to compute. The method is essentially a randomized version of the Kiefer-Wolfowitz method in which the gradient is estimated at each iteration from two measurements of the cost function. SPSA in the continuous setting is particularly efficient in problems of high dimension and where the cost function estimates is obtained through expensive simulations. The convergence properties of the algorithm have been established in a series of papers ([2], [3], [4], [5]).

The present paper discusses a version of SPSA for discrete optimization. The problem is to minimize a cost function that is defined on a subset of points in R^p with integer coordinates. It is assumed that only noisy measurements of the function are available and that the exact form of the function is analytically unavailable or difficult to obtain. An ordinal optimization method for finding the minimum was introduced in [6]. The method discussed here, which was introduced in [7], relies on simultaneous perturbation difference approximations.

This work was partially supported by the Johns Hopkins University Applied Physics Laboratory IR&D Program and the National Research Foundation of Hungary.

S. D. Hill is with the Johns Hopkins University Applied Physics Laboratory, Laurel, MD 20723 USA (e-mail: stacy.hill@jhuapl.edu)

L. Gerencsér, is with the Computer and Automation Research Institute, Hungarian Academy of Sciences, 1111 Budapest XI. Kende u. 13-17, HUNGARY (email: gerencser@szttaki.hu)

Z. Vágó is with the Computer and Automation Research Institute, Hungarian Academy of Sciences, 1111 Budapest XI. Kende u. 13-17, HUNGARY and also with and also with Pázmány Péter Catholic University, Budapest, HUNGARY (email: vago@itk.ppke.hu)

The main motivation for the algorithm is a class of discrete resource allocation problems ([8], [9]), which arise in a variety of applications that include, for example, the problems of distributing search effort to detect a target, allocating buffers in a queueing network, and scheduling data transmission in a communication network.

II. NOTATION AND PROBLEM FORMULATION

Let Z denote the set of integers and consider the grid Z^p of points in Z^p with integer coordinates. For $x', x'' \in Z^p$, we adopt the notation $x' \leq x''$ if and only if $x'_i \leq x''_i$ for $i = 1, \dots, p$, where x'_i and x''_i denote the coordinates of x' and x'' .

Consider a real-valued function $L: Z^p \rightarrow R$. The function is not assumed to be explicitly known, but noisy measurements of it are available:

$$y_n(\theta) = L(\theta) + \varepsilon_n(\theta) \quad (1)$$

where $\{\varepsilon_n(\theta)\}$ is a zero-mean stochastic process. The sequence $\varepsilon_n(\theta)$ is not necessarily independent; however, sufficient conditions are imposed to ensure that the $y_n(\theta)$'s are integrable. We assume also that L is bounded below. The problem is to minimize L using only the measurements y_n . The constrained version of this problem assumes that $L: \Theta \rightarrow R$, where the subset Θ of Z^p is a discrete rectangle or hypercube in Z^p , i.e., for some $a, b \in Z^p$, the point x belongs to Θ if and only if $a \leq x \leq b$.

Similar to [6], we restrict our attention to cost functions that satisfy a certain integer convexity condition. For the case $p = 1$, the function $L: Z \rightarrow R$ satisfies the inequality

$$L(\theta+1) - L(\theta) \geq L(\theta) - L(\theta-1) \quad (2)$$

or, equivalently,

$$2L(\theta) \leq L(\theta+1) + L(\theta-1) \quad (3)$$

for each $\theta \in Z$. The latter inequality is the discrete analogue of mid-convexity. If strict inequality holds, then L is also said to be strictly convex. Analogous to the continuous case, the problem of minimizing L reduces to the problem of finding its stationary values, i.e., any point $\theta' \in Z$ such that

$$L(\theta' \pm 1) \geq L(\theta') \quad (4)$$

or, equivalently,

$$L(\theta' + 1) - L(\theta') \geq 0 \geq L(\theta') - L(\theta' - 1) \quad (5)$$

If L is strictly convex its stationary point is unique.

The notion of integer convexity can be extended to Z^p as follows (see, e.g., [10], [11]). For $x', x'' \in Z^p$, $x' \leq x''$ if and only if $x'_i \leq x''_i$ for $i = 1, \dots, p$, where x'_i and x''_i denote the coordinates of x' and x'' . For $x \in Z^p$, let $\lfloor x \rfloor$ and $\lceil x \rceil$ denote the vectors obtained by rounding down and rounding up, respectively, the components of x to the nearest integers. The discrete neighborhood $N(x) \subseteq Z^p$ about $x \in Z^p$, is the set of points

$$N(x) = \{\theta \in Z^p : \lfloor \theta \rfloor \leq x \leq \lceil \theta \rceil\}$$

which is simply the smallest hypercube in Z^p about x . A real-valued function L on Z^p is integrally or discretely convex if for any $\theta', \theta'' \in Z^p$ and scalar λ in the interval $[0, 1]$

$$\min_{\theta \in N(\lambda\theta' + (1-\lambda)\theta'')} L(\theta) \leq \lambda L(\theta') + (1-\lambda)L(\theta''). \quad (6)$$

Observe that this condition implies (3) since for any $\theta \in Z^p$, $N(\frac{1}{2}(\theta+1) + \frac{1}{2}(\theta-1)) = \{\theta\}$.

A discretely convex function L defined on Z^p can be extended to a convex function L^* defined on all of R^p . The extension is continuous and piecewise linear ([11]). If L is strictly convex, then so is its continuous extension. For the case $p = 1$, the extension L^* is obtained by linearly interpolating L between points in Z .

The following is a consequence of (2) and (5).

Lemma 1: Assume that L is a strictly and discretely convex function on Z . The function

$$g(\theta) = L^*(\theta) - L^*(\theta - 1) \quad (7)$$

is continuous and strictly monotonic on R . Furthermore, if θ^* is a zero of g , then $\lfloor \theta^* \rfloor$ or $\lceil \theta^* \rceil$ minimizes L .

III. FINITE-DIFFERENCE BASED ALGORITHM

The difference function in (7) is not directly available, we must rely on noisy estimates \hat{g} to obtain θ^* . We can then find the minimum of the discrete function L by means of a stochastic approximation procedure based on the estimates \hat{g} of g . The approximation is obtained from the difference estimates $y(\theta) - y(\theta - 1)$ of $L(\theta) - L(\theta - 1)$.

To be more specific, consider the following stochastic approximation algorithm:

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a_k \left(y_k \left(\lfloor \hat{\theta}_k \rfloor \right) - y_k \left(\lfloor \hat{\theta}_k - 1 \rfloor \right) \right), \quad \hat{\theta}_1 \in Z^p \quad (8)$$

where the sequence $\{a_n\}$ satisfies the standard conditions, i.e. $a_k > 0$, $\sum a_k^2 < \infty$ and $\sum a_k = \infty$.

Proposition 1: Assume that L is a discretely and strictly convex function on Z and is bounded below. Suppose also that $L^2(\theta) + E(\varepsilon^2(\theta)) \leq O(1 + \theta^2)$. The sequence in (8) converges almost surely to the minimum θ^* of L^* .

Proof: The proof follows straightforwardly from Theorem 1 of [12] as a result of Lemma 1.

In [6], the cost function $L : Z^p \rightarrow R$ is assumed to be discretely convex and separable, i.e.,

$$L(\theta) = \sum_{i=1}^p L_i(\theta_i) \quad (9)$$

where each L_i is a discretely convex function on Z . If L is separable, then a necessary and sufficient condition for θ to be a minimum is that $L_i(\theta_i \pm 1) \geq L_i(\theta_i)$ for $i = 1, \dots, p$ (see [10]). In other words, a separable convex function achieves its minimum at the point whose components correspond to the stationary points of the L_i 's. If each L_i is strictly convex then the global minimum is unique and any local minimum is also a global minimum.

The minimization of separable convex functions on Z^p can be handled in a manner similar to that for the scalar case. Consider the vector-valued function \hat{g} with i -th component \hat{g}_i given by

$$\hat{g}_i(\theta) = y(\theta) - y(\theta_1, \dots, \theta_i - 1, \dots, \theta_p).$$

Thus, \hat{g}_i is an estimate of the i -th difference

$g_i(\theta) = L_i(\theta_i) - L_i(\theta_i - 1)$ in (9), which is strictly monotonic on Z . These difference estimates are analogous to finite difference estimates of gradients in continuous optimization.

We have the following multivariate extension of Proposition 1.

Proposition 2: Assume that L is a discretely and strictly convex separable function on Z^p and is bounded below. Assume also that $\sum_{i=1}^p L_i^2(\theta_i) + E(\varepsilon^2(\theta)) \leq O(1 + \theta^2)$. Then (8) converges almost surely to the unique global minimum θ^* of the continuous extension L^* of L .

Proof: Let g^* denote the function obtained by extending each component of g to a continuous convex functions on R . We need only check that the function $-g^*(x)$ satisfies the conditions in Theorem 4 of [12]. The main condition that must be verified (the others hold by assumption) is that $-g^*(x)^T(x - \theta) > 0$, for some $\theta \in R^p$ and all $x \in R^p$, $x \neq \theta$. Since the L_i 's are strictly convex, this inequality follows from (2) applied to each term in (9), when $\theta = \theta^*$.

IV. THE SPSA METHOD

The SPSA method is based on simultaneous perturbation estimates of the gradient. In the discrete case, differences replace the gradient. To estimate the differences of $L(\theta)$ we use simultaneous random perturbations. At each iteration k of the algorithm, we take a random perturbation vector $\Delta_k = (\Delta_{k1}, \dots, \Delta_{kp})^T$, where the Δ_{ki} 's form an i.i.d. sequence of Bernoulli random variables taking the values ± 1 . The perturbations are assumed to be independent of the measurement noise process. For cost functions defined on R^p , the difference estimate at iteration k is obtained by evaluating $y_k(\cdot)$ at two values:

$$y_k^+(\theta) = L(\theta + c_k \Delta_k) + \varepsilon_{2k-1}(\theta + c_k \Delta_k),$$

$$y_k^-(\theta) = L(\theta - c_k \Delta_k) + \varepsilon_{2k}(\theta - c_k \Delta_k).$$

where $c_k > 0$. The i -th component of the difference estimate H is

$$H_i(k, \theta) = (y_k^+(\theta) - y_k^-(\theta)) / 2c_k \Delta_{ki} \quad (10).$$

Consider the sequence

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a_k H(k+1, \lfloor \hat{\theta}_k \rfloor). \quad (11)$$

with an initial estimate $\hat{\theta}_1 \in Z^p$.

The discrete, fixed gain version of this algorithm, i.e. $a_k \equiv a$, $c_k \equiv 1$, where $a > 0$, was introduced in [7]. In that algorithm, the difference estimate in (11) was replaced by its truncation $\lfloor H \rfloor$ and the iterates were constrained to lie in the set Z^p .

Consider the following discrete optimization algorithm, which was introduced in [13]

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a_k H^*(k+1, \hat{\theta}_k) \quad (12)$$

where the components of H^* are obtained from an approximation to L^* based on noise-corrupted measurements $y(\theta)$ of L . In this version, $\{a_k\}$ and $\{c_k\}$ satisfy the standard conditions for a Kiefer-Wolfowitz type algorithm. (Also, assumptions on the perturbations Δ_k can be relaxed, i.e., they need not be Bernoulli random variables.) The sequence $\hat{\theta}_k$ in (12) provides an estimate of the minimum of the extension L^* .

Proposition 3: Assume the conditions of Proposition 2. Assume also that the components of $\Delta_k = (\Delta_{k1}, \dots, \Delta_{kp})^T$ are bounded, i.i.d., symmetrically distributed about zero and satisfy $E|\Delta_{ki}^{-1}| < \infty$. The sequence in (12) converges almost surely to the stationary value of L^* .

The proof of this result relies on the notion of subgradients. A subgradient of L^* at θ is any $\xi \in R^p$ such that $L^*(\theta + h) - L^*(\theta) \geq \xi^T h$ for all $h \in R^p$. Since L^* is a convex function and is continuous at θ , the set of subgradients of L^* at θ is a nonempty compact set. ([14]).

Proof: The result follows from [15] (see Proposition 2) or [16] (see Theorem 5.6.2) if we can show that H^* is an approximate subgradient of L^* in the following sense. For each $k \geq 1$ and $\varepsilon > 0$, there is a subgradient of L^* at θ , denoted $\xi_k^T(\theta, \Delta_k)$, such that

$$\begin{aligned} \frac{y_k^+(\theta) - y_k^-(\theta)}{2c_k} &= H_i^*(k, \theta) \Delta_{ki} \\ &= \xi_k^T(\theta, \Delta_k) \Delta_k + \varepsilon. \end{aligned}$$

Lemma 1 of [15] derives this approximation under the assumption that the Δ_{ki} 's are supported on a finite discrete set. Since L^* is piecewise linear, this restriction can be relaxed. The subgradient of L^* at θ can then be chosen so that it is independent of Δ_k . In other words, for each θ ,

there is a subgradient $\xi(\theta)$ of L^* at θ such that

$$E\{H(k, \theta)\} = \xi(\theta) + o(c_k)/c_k.$$

The conclusion now follows from this and Proposition 2 of [15] or Theorem 5.6.2 of [16].

REFERENCES

- [1] J. C. Spall. "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," IEEE Trans. on Automat. Contr., vol. 37, pp. 332--341, 1992.
- [2] H. F. Chen, T. E. Duncan, and B. Pasik-Duncan. "A stochastic approximation algorithm with random differences," J. Gertler, J. B. Cruz, and M. Peshkin, editors, Proceedings of the 13th Triennial IFAC World Congress, pages 493--496, 1996.
- [3] L. Gerencsér. "On fixed gain recursive estimation processes," J. of Mathematical Systems, Estimation and Control, vol. 6, pp. 355--358, 1996.
- [4] L. Gerencsér. "Rate of convergence of moments for a simultaneous perturbation stochastic approximation method for function minimization," IEEE Trans. on Automat. Contr., vol. 44, pp. 894--906, 1999.
- [5] L. Gerencsér. "SPSA with state-dependent noise—a tool for direct adaptive control," Proceedings of the Conference on Decision and Control, CDC 37, 1998.
- [6] C. G. Cassandras, L. Dai, and C. G. Panayiotou. "Ordinal optimization for a class of deterministic and stochastic discrete resource allocation problems," IEEE Trans. Auto. Contr., vol. 43(7): pp. 881--900, 1998.
- [7] L. Gerencsér, S. D. Hill, and Z. Vágó. "Optimization over discrete sets via SPSA," Proceedings of the Conference on Decision and Control, CDC 38, 1999.
- [8] T. Ibaraki and N. Katoh. Resource Allocation Problems: Algorithmic Approaches. MIT Press, 1988.
- [9] J.E. Whitney, II, S. D. Hill, and L. I. Solomon. "Constrained optimization over discrete sets via SPSA with application to non-separable resource allocation," Proceedings of the 2001 Winter Simulation Conference, pp. 313-317.
- [10] B.L. Miller, "On Minimizing Nonseparable Functions Defined on the Integers with an Inventory Application, SIAM Journal on Applied Mathematics, vol. 21, pp. 166-185.
- [11] P. Favati and F. Tardella, "Convexity in nonlinear integer programming," Ricerca Operativa, vol. 53, pp. 3-34, 1990.
- [12] V. Dupac and U. Herkenrath, "Stochastic approximation on a discrete set and the multi-armed bandit problem," Communications in Statistics--Sequential Analysis, vol 1, pp. 1025, 1982.
- [13] Hill, S. D., L. Gerencsér and Z. Vágó, "Stochastic Approximation on Discrete Sets Using Simultaneous Perturbation Difference Approximations," Proc. of the 2003 Conf. on Information Science and Systems, The Johns Hopkins University, March 12-14, 2003
- [14] R. T. Rockafellar, Convex Analysis. Princeton, NJ: Princeton Univ. Press, 1970.
- [15] Y. He, M. C. Fu, and S. I. Marcus, "Convergence of Simultaneous Perturbation Stochastic Approximation for Nondifferentiable Optimization," IEEE Trans. on Auto. Contr., vol. 48, pp. 1459-1463.
- [16] H. J. Kushner and G. G. Yin, Stochastic Approximation Algorithms and Applications. New York: Springer-Verlag, 1997.
- [17] J. N. Eagle and J. R. Yee. "An optimal branch-and-bound procedure for the constrained path, moving target search problem," Operations Research, vol. 8, 1990.