Stochastic Approximation on Discrete Sets Using Simultaneous Perturbation Difference Approximations

Stacy D. Hill
Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Road
Laurel, MD 20723-6099
stacy.hill@jhuapl.edu, (240) 228-4966

László Gerencsér and Zsuzsanna Vágó
Computer and Automation Institute
Hungarian Academy of Sciences
Kende 13-17, Budapest, 1111, HUNGARY

Abstract: **We consider a stochastic approximation method for optimizing a class of discrete functions. The procedure is a version of the Simultaneous Perturbation Stochastic Approximation (SPSA) method that has been modified to obtain an optimization method for cost functions defined on a grid of points in Euclidean *p*-space having integer components. We consider the convergence properties of discrete algorithm and discuss some related results on fixed gain SPSA. The application of the method to resource allocation is also briefly discussed.**

Keywords: Discrete optimization, resource allocation, stochastic approximation, convex, discretely convex

1. Introduction

The *simultaneous perturbation stochastic approximation* (SPSA) method [10] is a tool for solving continuous optimization problems in which the cost function is 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 only two measurements of the cost function. SPSA in the continuous setting is particularly efficient in problems of high dimension and where the cost function must be estimated through expensive simulations. The convergence properties of the algorithm have been established in a series of papers ([2], [5], [6], [10]).

The present paper discusses a modification of SPSA for discrete optimization. The problem is to minimize a cost function that is defined on a subset of points in $\mathbb{R}^p$ with integer coordinates. It is assumed that only noisy measurements of the cost function are available and that the exact form of the function is analytically unavailable or difficult to obtain. A method based on ordinal optimization was introduced in [1] for finding the cost function minimum. Alternatively, we consider a method based on stochastic approximation (SA), which may be applicable to a broader class of problems than those treated in [1]. In particular, the method discussed here is a version of the algorithm introduced in [7].

The main motivation for the algorithm is a class of discrete resource allocation problems, 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.

2. Problem Formulation

Let $\mathbb{Z}$ denote the set of integers and consider the grid $\mathbb{Z}^p$ of points in $\mathbb{R}^p$ with integer coordinates. Consider a real-valued function $L:\mathbb{Z}^p \to \mathbb{R}$. The function is not assumed to be explicitly known, but noisy measurements $y_n(\theta)$ of it are available:

$$y_n(\theta) = L(\theta) + \varepsilon_n(\theta), \quad \theta \in \mathbb{Z}^p, \qquad (2.1)$$

where $\{\varepsilon_n(\theta)\}$ is a sequence of zero-mean random variables. The sequence $\varepsilon_n(\theta)$ is not necessarily independent; however, sufficient conditions are imposed to ensure that the $y_n(\theta)$'s are integrable. Assume that L is bounded below. The problem is to minimize L using only the measurements y_n.

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

$$L(\theta+1) - L(\theta) \geq L(\theta) - L(\theta-1) \qquad (2.2)$$

or, equivalently,

$$2L(\theta) \le L(\theta+1) + L(\theta-1) \qquad (2.3)$$

for each $\theta \in \mathbb{Z}$. The latter inequality is the discrete analogue of mid-convexity. If strict inequality holds, then L is said to be *strictly convex*. Analogous to the continuous case, the problem of minimizing L reduces to the problem of finding a *stationary* value of L, i.e., a point $\theta^* \in \mathbb{Z}$ such that

$$L(\theta^* \pm 1) \ge L(\theta^*) \qquad (2.4)$$

or, equivalently,

$$L(\theta^*+1) - L(\theta^*) \ge 0 \ge L(\theta^*) - L(\theta^*-1) \qquad (2.5)$$

If L is strictly convex then the stationary point is unique.

The notion of integer convexity can be extended to $\mathbb{Z}^p$ as follows (see, e.g., [9], [12]). For $x', x'' \in \mathbb{R}^p$, $x' \le x''$ if and only if $x'_i \le x''_i$ for i = 1, …, p, where x'_i and x''_i denote the coordinates of x' and x''. For $x \in \mathbb{R}^p$, 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 $N(x) \subseteq \mathbb{Z}^p$ about $x \in \mathbb{R}^p$, is the set of points

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

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

$$\min_{\theta \in N(\lambda\theta' + (1-\lambda)\theta'')} L(\theta) \le \lambda L(\theta') + (1-\lambda) L(\theta''). \qquad (2.6)$$

Observe that this condition implies (2.3) since for any $\theta \in \mathbb{Z}$

$$N\left(\tfrac{1}{2}(\theta+1) + \tfrac{1}{2}(\theta-1)\right) = \{\theta\}.$$

A discretely convex function L defined on $\mathbb{Z}^p$ can be extended to a convex function L^* defined on all of $\mathbb{R}^p$. The extension is continuous and piecewise linear ([12]).

For the case p = 1, the extension L^* is obtained by linearly interpolating L between points in $\mathbb{Z}$. If L is strictly convex, then so is the continuous function L^*. Also, the function

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

is continuous and strictly monotonic. If θ' is a zero of g, then $\lfloor \theta' \rfloor$ or $\lceil \theta' \rceil$ minimizes L. Since g 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 an SA procedure based on the estimates $\hat{g}$ of g. The approximation $\hat{g}$ is obtained by linearly interpolating the difference estimates $y(\theta) - y(\theta-1)$ of $L(\theta) - L(\theta-1)$. The convergence of this procedure is readily established using results in [13], if we impose the additional condition on the cost function and noise:

$$L^2(\theta) + E\left(\varepsilon_n^2(\theta)\right) = O(1+\theta^2).$$

In [1], the cost function $L : \mathbb{Z}^p \to \mathbb{R}$ is assumed to be discretely convex and *separable*, i.e.,

$$L(\theta) = \sum_{i=1}^{p} L_i(\theta_i) \qquad (2.7)$$

where each L_i is a discretely convex function on $\mathbb{Z}$. If L is separable, then a necessary and sufficient condition for θ to be a minimum is that $L_i(\theta_i \pm 1) \ge L_i(\theta_i)$ for i = 1, …, p (see [9]). In other words, a separable convex function achieves its minimum at its stationary points. If each L_i is strictly convex then the global minimum is unique and any local minimum is also a global minimum.

Using results in [13], the minimization of separable convex functions on $\mathbb{Z}^p$ can be handled in a manner similar to that for the case p = 1. For this, consider the vector-valued function $h : \mathbb{Z}^p \to \mathbb{R}^p$ with i-th component h_i given by

$$h_i(\theta) = y(\theta) - y(\theta_1, \ldots, \theta_i - 1, \ldots, \theta_p).$$

Thus, $h_i(\theta)$ is an estimate of the i-th difference $L_i(\theta_i) - L_i(\theta_i - 1)$ in (2.7), which is strictly monotonic on $\mathbb{Z}$.

We seek to relax the separability assumption on the cost function, which is one motivation for considering the SPSA approach.

3. 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 $\Delta_k = (\Delta_{k1}, \ldots, \Delta_{kp})^T$, 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. In fixed gain SPSA, the step size of the perturbation is fixed at, say, some $c > 0$. To compute the difference estimate at iteration k, we evaluate $y_k(\cdot)$ at two values of θ:

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

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

The i-th component of the difference estimate is

$$H_i(k,\theta) = \frac{\left(y_k^+(\theta) - y_k^-(\theta)\right)}{2c_k\Delta_{ki}}.$$

Let $a_k > 0$ be a sequence such that $a_k \to 0$ and $\sum a_k = \infty$. Also, take $c_k \equiv 1$. Starting with an initial estimate $\hat{\theta}_1 \in \mathbb{Z}^p$, we recursively compute a sequence of estimates

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a_k H\left(k+1, \lfloor \hat{\theta}_k \rfloor \right) \qquad (3.1)$$

where the i-th component of H is H_i. Note that the iterates $\hat{\theta}_k$ are not constrained to lie in $\mathbb{Z}^p$. However, we are interested in the asymptotic behavior of the truncations $\lfloor \hat{\theta}_k \rfloor \in \mathbb{Z}^p$.

We also consider a version of (3.1)

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a_k H^*\left(k+1, \hat{\theta}_k \right) \qquad (3.2)$$

where the components of H^* are obtained from an approximation to L^* based on noise-corrupted measurements $y(\theta)$ of L. In this version, the sequence $\{ c_k \}$ satisfies the standard conditions for a Kiefer-Wolfowitz type algorithm. The sequence $\hat{\theta}_k$ in (3.2) provides an estimate of the minimum of the extension L^*.

Another version of (3.1), which was introduced in [7], is a fixed gain algorithm. For this, let $a > 0$ and consider the following form of the algorithm

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a\, H\left(k+1, \lfloor \hat{\theta}_k \rfloor \right). \qquad (3.3)$$

The assumed boundedness of the noise and assuming the stability of an associated ODE ensures that the sequence of estimates in (3.3) is bounded. The pathwise behavior of estimator process generated by fixed gain SPSA can be analyzed using the result of [4].

The analysis of (3.2) requires a different approach. The convergence can be studied using results in [14], which give conditions for the convergence of a Kiefer-Wolfowitz type SA algorithm when the cost function is continuous and convex. For example (see Corollary 3 of Theorem 2.3.5 in [14]):

Proposition: Let L be a continuous convex function on $\mathbb{R}^p$. Suppose that L is bounded below and that $|L(\theta)| \to \infty$ as $\|\theta\| \to \infty$. Assume the standard SA conditions on $\{ a_k \}$, $\{ c_k \}$, and the measurement noise process. Let $\{ \hat{\theta}_k \}$ be a sequence of estimates from a Kiefer-Wolfowitz SA algorithm and assume that $\{ \hat{\theta}_k \}$ is bounded with probability 1. Then $\{ \hat{\theta}_k \}$ converges with probability to the set of stationary points of L.

4. Resource Allocation

We consider the application of SPSA on grids to a class of multiple discrete resource allocation problems ([1], [8], [11]). The goal is to distribute a finite amount of resources of different types to finitely many classes of users, where the amount of resources that can be allocated to any user class is discrete.

There are n types of *resources*, where N_j denotes the number of resources of type j, $j = 1, \ldots, n$. These resources are to be allocated over M *user* classes. Let θ_{jk} denote the number of resources of type j that are allocated to user class k, and θ be the vector consisting of all the θ_{jk}'s. The allocation of resources to users in class k is denoted θ_k, thus $\theta_k = (\theta_{1k}, \ldots, \theta_{nk})$. For each allocation vector θ there is an associated cost function or performance index $L(\theta)$, which is the expected cost. The goal is to distribute the resources in such a way that cost is minimized:

$$\begin{aligned} &\text{minimize } L(\theta), \\ &\text{subject to } \sum_{k=1}^{M} \theta_{jk} = N_j,\ \theta_{jk} \geq 0,\ 1 \leq j \leq n \end{aligned} \tag{4.1}$$

where the θ_{jk}'s are nonnegative integers. The case of most interest, of course, is when the cost is observed with noise and the expected cost $L(\theta)$ is analytically unavailable. This problem includes many problems of practical interest including the problem of optimally distributing a search effort to locate a moving target whose position is unknown and time varying (cf. [3]) and the problem of scheduling time slots for the transmission of messages over nodes in a radio network (cf. [1]). The formulation in (4.1) is a generalization of the single resource allocation problem considered in [1] with separable cost.

5. References

[1] 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.

[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] 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.

[4] L. Gerencsér. "On fixed gain recursive estimation processes," J. of Mathematical Systems, Estimation and Control, vol. 6, pp. 355--358, 1996.

[5] 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.

[6] L. Gerencsér. "SPSA with state-dependent noise—a tool for direct adaptive control," *Proceedings of the Conference on Decision and Control*, *CDC 37*, 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] 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, 1971.

[10] J. C. Spall. "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Trans. on Automat. Contr.*, vol. 37, pp. 332--341, 1992.

[11] 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.

[12] P. Favati and F. Tardella, "Convexity in nonlinear integer programming," Ricerca Operativa, vol. 53, pp. 3-34, 1990.

[13] 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.

[14] Harold J. Kushner and Dean S. Clark, *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlag, 1978.

6. Acknowledgments

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