

Simultaneous perturbation method for processing magnetospheric images

Daniel C. Chin

The Johns Hopkins University Applied
Physics Laboratory
11100 Johns Hopkins Road
Laurel, Maryland 20723-6099
E-mail: daniel.chin@jhuapl.edu

Abstract. Extracting a multivariate nonlinear physical model from a set of satellite images is considered as a multivariate nonlinear regression problem. Multiple local solutions often prevent gradient type algorithms from obtaining global optimal solutions. A method of solving this problem is presented based on the simultaneous perturbation stochastic approximation (SPSA) algorithm. The method is applied to a problem of estimating the distribution of energetic ion populations in the magnetosphere from global images of the magnetosphere. The approach uses multiple objective functions: single image errors and the summation of square image errors. The algorithm is demonstrated on simulated energetic-neutral atom (ENA) images. Within a reasonable number of function evaluations, the process converges and reconstructs the images with a mean square error less than or equal to 0.1% of the original image. Also, the SPSA method is compared with results obtained from simulated annealing (SAN) in a single objective function setting. In the comparison study, SPSA has a 3:1 advantage over SAN in both accuracy and efficiency measures. © 1999 Society of Photo-Optical Instrumentation Engineers. [S0091-3286(99)00304-9]

Subject terms: optimization; image processing; system identification; model estimation; simultaneous perturbation stochastic approximation; data fusion.

Paper 980077 received Mar. 5, 1998; revised manuscript received Aug. 21, 1998, and Oct. 23, 1998; accepted for publication Oct. 28, 1998.

1 Introduction

This paper presents an image-processing algorithm for extracting a physical model from a system that either has multiple local solutions or whose structure is complicated, such as when using the global magnetospheric images formed by energetic-neutral atoms (ENA) to estimate magnetospheric energetic ion distributions. Roelof¹ was the first to demonstrate the ENA imaging technique. Chase and Roelof² discussed the extraction of the hypothesized ion-intensity model from the ENA sensor onboard a satellite. The problems in using the ENA emission images are the Poisson counting fluctuations in the image taking, the expensive time-consuming line-of-sight integration in the measurement evaluations, and the variability of ion distributions. This paper presents a modified simultaneous perturbation stochastic approximation (SPSA) algorithm for the inversion process of the satellite data in a single global image.

The inversion process using a finite amount of data in a complicated system often runs into multiple solution problems. There are many optimization techniques available for estimating the hypothesized model, such as Bayesian estimation,^{3,4} maximum likelihood estimation,⁵ and steepest descent and Newton-Raphson.⁶ They all require detailed information of the system, either for computing gradients or for forming statistical distributions. These algorithms are hard to use when their objective functions have multiple local solutions. Other estimation techniques rely only on observations, such as Kiefer/Wolfowitz stochastic approximation⁷ (KWSA), simulated annealing^{8,9} (SAN),

and other linear search type algorithms, as in Ref. 2. Only KWSA and SAN have been proven to demonstrate convergence using noisy measurements. Both KWSA and SAN have been shown in general forms to work in finding the global optimal solution.

The SPSA algorithm¹⁰ is one of the KWSA types of algorithms that has been shown to require less data to converge than the other currently available KWSA algorithms.¹¹ Chin¹² (and Styblinski and Tang¹³) concluded that SPSA requires fewer function evaluations than SAN. This paper also compares SPSA with SAN directly. The comparison study shows SPSA is about three times more efficient than SAN; in an average of 10 runs, the level of accuracy at 300 function evaluations for SAN is about the level of accuracy at 100 function evaluations for SPSA. Some other successful examples of SPSA applications are in detecting faults in a power plant,¹⁴ learning rules for a neurocontroller,¹⁵ and estimating an electric current conductivity map to discriminate buried objects.¹⁶

When applying SPSA to the inversion process in the global magnetospheric image setting, the original algorithm must be modified. The original algorithm was designed for asymptotic convergence and required an open-ended sequential data set. Because of the time-varying ion intensity in the magnetosphere, the field changes over time, so the inversion process is forced to use a single global image. Some of the nice convergence properties are lost, such as the strong convergence in the original theory, where the SPSA iteration converges to the truth asymptotically; but in the magnetospheric global image setting, the SPSA itera-

tion will converge to a solution that is represented by the single global image. The process used here estimates a model of the system by matching a finite quantity of sensor data with the values computed from an estimated model; the process then iteratively modifies the model values in an attempt to obtain a better fit. The data noise (including sensor background noise and random fluctuations) tends to average out for the original SPSA algorithm, but that is not the case for the finite data stream of a single global image.

In dealing with the multiple-root problem, Chin¹² shows that SPSA with a certain stepped gain sequence (function dependent) can be used to find a global solution. (There may be more than one representation for a single global image.) The simulation studies presented in this paper show that the SPSA iteration converges to the solution within a reasonable number of iterations; there are many local solutions in the system (the final estimates are trapped in different solutions for other estimation methods).

This paper also discusses a new approach for SPSA that uses multiple objective functions to find the common optimal solution for all functions. The purpose here is to fuse the magnetospheric image pixels (counts for each line of sight direction) for a global fit and still to match each individual count (pixel) locally. A single objective function lumps all the components together in a single mathematical formula that fits the counts (pixels) globally. In the single objective function approach, the parameters to which the objective function is most sensitive tend to dominate the estimation process. A separation of the single objective function into multiple functions or an addition type function may desensitize the domination and get a better overall fit. The example given here shows that the multiple-objective function algorithm achieves the lowest level of the total of measurement residual errors.

2 The Global Magnetospheric Image Problem

The magnetosphere exists in a region of space that surrounds the earth out to 100,000 km above the earth's surface. Its field lines, which pass through the earth's surface, fill the space. The global magnetospheric image process is used to estimate the energetic ion-populations trapped within the field. A global magnetospheric image is represented in pixel units, each of which is an accumulation of ENA (weighted by the exoatmospheric hydrogen density) along the line of sight extended from the ENA camera (actually a particle counter). The number of pixels needed to represent a global image depends on the resolution of the sensors; for a sensor equipped with a 4×4 deg pixel resolution, it takes 4050 pixels over 4π sr (180 deg in latitude and 360 deg in longitude) to represent a global image. The actual numbers are slightly less because the top and bottom latitude regions have smaller surface area and use fewer pixels. If an ion-intensity model is known, the image value can be obtained by integrating along the line of sight the ion intensity times the product of the hydrogen atom intensity and the cross section of the exchange process by which an ion strips an electron from a hydrogen atom to become a fast neutral atom detected at the ENA camera. Some highly complex nonparametric simulations are available for the global magnetospheric ion intensities (e.g., the Rice convection model and the magnetosphere specification model¹⁷ (MSM) and the 3-D ring current decay model¹⁸). The ob-

ject of this paper is to provide an estimation technique that attempts to invert this process. That is, given the hydrogen density and cross section and ENA counts over each 4×4 deg angular sector (or whatever the instrument's resolution) over a full 4π sr, one would like to estimate the ion density of the magnetosphere.

2.1 Chase and Roelof Model

The field lines of the magnetosphere leave and enter the earth at the magnetic north and south poles (offset by about 12 deg from the geodetic poles) and approximate a dipole field. Due to dominance of the magnetic field over near earth plasma, a simplified distribution of ring current ions as a function of only two variables can be used. Quoting from Chase and Roelof,² "the two variables are L (constant along a dipole field line defined in spherical coordinates by $r = aL \cos^2 \Lambda$, where $a = 1R_E$ and Λ is the geomagnetic latitude) and the azimuthal angle ϕ measured anticlockwise from the sunward direction." Here R_E is one earth radius. To modulate the ion intensity as a function of ϕ , a second-order harmonic expansion is used: $F_\phi = k_1[1 - \cos(\phi - \phi_1)] + k_2[1 - \cos 2(\phi - \phi_2)]$. To modulate the intensity as a function of L , Chase and Roelof define a piecewise function F_L , which is parameterized on the five parameters $(L_0, L_1, L_2, \delta L_1, \delta L_2)$:

$$F_L = \begin{cases} (L - L_1)^2 / 2\delta L_1^2 & L < L_{11} \\ (L - L_{11})/L_0 + (1/2)(\delta L_1/L_0)^2 & L_{11} \leq L \leq L_{22} \\ (L - L_2)^2 / 2\delta L_2^2 + (L_2 - L_1)/L_0 \\ \quad + (1/2)(\delta L_2/L_0)^2 - (1/2)(\delta L_1/L_0)^2 & L > L_{22}, \end{cases}$$

where $L_{11} = L_1 + \delta L_1^2/L_0$ and $L_{22} = L_2 + \delta L_2^2/L_0$, and $(L_0, L_1, L_2, \delta L_1, \delta L_2)$ are also functions of ϕ that are represented by five different harmonic expansions. Then, the ion density $j_{\text{ion}}(\phi, L)$ in the magnetosphere can be written as

$$j_{\text{ion}} = j_0 \exp(-F_\phi - F_L),$$

where j_0 is a constant. An ENA image is an image each of whose pixels represents the number of fast neutral atoms detected over each angular sector defined by the resolution of the ENA camera. The expected intensity of these counts is a function of the ion intensity along the line of sight, but the actual realized count in each angular sector (represented by a single pixel) is a random sample from a Poisson distribution with intensity parameter given as a function of the ion intensity. If $\lambda(j_{\text{ion}})$ is the ENA intensity of the Poisson distribution Π along a particular line of sight, then an ENA image pixel represents $\sum \Pi[\lambda(j_{\text{ion}})]$ over all lines of sight in the cone subtended at the camera by the spherical angle determined by the resolution of the camera. But mathematically $\sum \Pi[\lambda(j_{\text{ion}})] = \Pi[\sum \lambda(j_{\text{ion}})]$, so that one can first compute the expected intensity of counts for each particular angular sector subtended by the camera with its resolution, and then sample only once per angular sector from the Poisson distribution. This is what is done in the simulations of this paper. The simulations are inherently random in that even with the same ion-intensity model, two simulations of

an ENA image will differ due to the randomness of the Poisson distribution that characterizes the ENA counts along a given line of sight. For the purposes of the estimation of the ion intensity from an ENA image, the randomness of the counts acts as a kind of measurement noise. Figure 1 of Ref. 2 shows the degrading effect of Poisson counting statistics on the ENA images. Chase and Roelof² also show that the images from their ion-intensity model match with the simulated images from the Rice model within a very small variation when the Poisson counting statistics are ignored.

3 Multiple-Objective Function SPSA

Let θ represent the ion-intensity model parameters. For the i 'th pixel ($i = 1, \dots, n$), let $j_{\text{ENA},i}$ be the magnetospheric image sensor measurement for pixel i and let $F_i(\theta)$ be the associated line-of-sight integration value predicted by the ion-intensity model. If the background noise level of the sensor is represented as σ , the usual SPSA algorithm optimizes a single loss (objective) function P defined as:

$$P(\theta) = \sum_{i=1}^{\ell} [F_i(\theta)j_{\text{ENA},i}]^2,$$

where ℓ is the total number of pixels in the global image whose values are greater than the background noise level σ .

In contrast, the multiple-objective function SPSA algorithm seeks to optimize *each* of the normalized functions

$$P_i(\theta) = k_i [F_i(\theta) - j_{\text{ENA},i}]^2,$$

where $k_i = 1/j_{\text{ENA},i}, j_{\text{ENA},i} > \sigma$, for each pixel $i = 1, \dots, \ell$ whose measurement count is greater than σ . For each iteration of the usual SPSA algorithm, it substitutes a sequence of ℓ SPSA estimates, updating θ based on the gradient information of each function $P_i(\theta)$ in turn. The normalizing factor $k_i = 1/j_{\text{ENA},i}, j_{\text{ENA},i} > \sigma$, is meant to keep the information from those pixels with the largest measurements from dominating the updates, i.e., it is meant to enable the algorithm to do its best to fit *all* the observations, not just the large ones.

To simplify the discussion, the original SPSA algorithm is discussed first, then the multiple-objective function SPSA is explained from the original algorithm. Letting $g(\bullet)$ denote the gradient of $P(\theta)$ with respect to θ , $\hat{\theta}_k$ denote the estimate for θ at the k 'th iteration, and $\hat{g}_k(\hat{\theta}_k)$ denote the SPSA approximated gradient at $\hat{\theta}_k$, the SPSA algorithm has the form

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a_k \hat{g}_k(\hat{\theta}_k),$$

where the gain sequence $\{a_k\}$ satisfies certain well-known stochastic approximation conditions.¹¹ Let $\Delta \in R^p$ be a vector of p mutually independent mean-zero random variables $\{\Delta_1, \Delta_2, \dots, \Delta_p\}$ satisfying conditions given in Section III of Spall,¹¹ where p is the number of estimated parameters. Subject to the important conditions in Ref. 11, the user has full control over Δ . The recommended choice for the dis-

tribution of Δ_i is Bernoulli (± 1). (Gaussian and uniform are not allowed.) Each iteration, SPSA uses two function evaluations to approximate the gradient, $\hat{g}_k(\bullet)$. In particular, at design levels $\hat{\theta}_k \pm c_k \Delta_k$, with c_k a positive scalar, let

$$y_k^{(+)} = P(\hat{\theta}_k + c_k \Delta_k),$$

$$y_k^{(-)} = P(\hat{\theta}_k - c_k \Delta_k).$$

Then the SPSA estimate of $g(\bullet)$ at the k 'th iteration is

$$\hat{g}_k(\hat{\theta}_k) = \begin{bmatrix} \frac{y_k^{(+)} - y_k^{(-)}}{2c_k \Delta_1} \\ \cdot \\ \cdot \\ \cdot \\ \frac{y_k^{(+)} - y_k^{(-)}}{2c_k \Delta_p} \end{bmatrix}.$$

The sequence $\{c_k\}$, for $k = 1, 2, \dots$, should satisfy the convergence conditions as they are stated in Ref. 11.

For the multiple-objective function SPSA algorithm, Δ_k, c_k, a_k are selected for each k , as earlier, and an approximate gradient $\hat{g}_{k,i}$ of each function P_i is formed just as for P in the preceding, using the same Δ_k, c_k for each P_i . Then, holding Δ_k, c_k, a_k constant, one updates θ sequentially for $i = 1, \dots, \ell$ to yield at the k 'th major iteration and i 'th minor iteration (for function P_i)

$$\hat{\theta}_{k+1,i} = \hat{\theta}_{k+1,i-1} - a_k \hat{g}_{k,i}(\hat{\theta}_{k,i}),$$

$$\text{for } i = 1, \dots, \ell \text{ and } \hat{\theta}_{k+1,0} = \hat{\theta}_{k,\ell}.$$

4 Accelerated Procedure

This paper uses the SPSA iterative procedure in stages. Each stage consists of a set of runs with identical conditions except for the seed for the random number generator. In the stochastic approximation algorithm, the final estimates of a run are different from the other final estimates according to the initial selection of the seed of the random number generator. It is very difficult to determine whether the iterative process approaches a global solution or a local solution, especially if the initial values of model parameters are very different from the solution. Also, it is a waste to let a run proceed when the iteration has taken some wrong turns, which is normal in optimization iterations. Therefore, it is a good idea to restart a new stage with an averaged estimate from the last estimates in the previous stage. Furthermore, a selective average can be done between stages. In the simulation study, the initial values for the model parameters in the second stage are the average of four runs from the first stage for which the errors of the measurement residuals in the last iterations are less than 0.3 in value. (It is just a natural separation among the runs, no particular meaning.)

To avoid ambiguities, constraints have been put on the scale and trigonometric parameters of the harmonic expression determining the functions F_ϕ and F_L . The scale pa-

Table 1 SPSA results.

Elements of Objective Function	SPSA Initial Values	Simulation Values	Second Stage Estimates	Final Estimates
F_ϕ	(2.30, 0.1, 43 deg)	(2.50, 0.5, 29 deg)	(2.46, 0.5, 42 deg)	(2.45, 0.5, 41 deg)
L_0	(0.20, 0.1, 43 deg)	(0.22, 0.5, 19 deg)	(0.21, 0.3, 44 deg)	(0.21, 0.4, 43 deg)
L_1	(2.00, 0.1, 43 deg)	(2.50, 0.4, 66 deg)	(2.52, 0.3, 53 deg)	(2.51, 0.3, 56 deg)
L_2	(4.50, 0.1, 43 deg)	(5.00, 0.8, 29 deg)	(5.07, 0.4, 41 deg)	(5.05, 0.4, 41 deg)
δL_1	(0.20, 0.1, 43 deg)	(0.14, 0.2, 27 deg)	(0.13, 0.3, 32 deg)	(0.14, 0.3, 33 deg)
δL_2	(0.10, 0.1, 43 deg)	(0.11, 0.8, 46 deg)	(0.12, 0.5, 40 deg)	(0.11, 0.5, 41 deg)

rameters are limited to positive values; and the trigonometric parameters are restricted to the interval $[-\pi, \pi]$. If an estimate for these constrained parameters exceeds its range, the value of the estimate is reset to the initial value of that stage. There are other choices for determining the reset values depending on the purpose to which the procedure is applied.

Because the Chase and Roelof model used for the simulation has only piecewise continuous gradients, the perturbed values $\hat{\theta}_k \pm c_k \Delta_k$ used to compute the SPSA approximate gradient $\hat{g}_k(\hat{\theta}_k)$ may lie on either side of a gradient discontinuity. In this situation, the approximated gradients are not representative of either piece and cause the measurement residual error to jump. If the jump makes the measurement residual error small, it is beneficial. If the jump goes the other direction, it is not beneficial. The simulation study here ignores such updates if the updated measurement residual error is greater than 110% of the previous value.

5 Simulation Study

This study is performed using measurements simulated from a simple Chase and Roelof model; first-order harmonic expansions in ϕ being used for all elements (F_ϕ , L_0 , L_1 , L_2 , δL_1 , δL_2). The first-order expansions consist of three parameters, constants, scalars, and phase angles. Let k , s , and φ represent these three parameters, respectively, and let E be an arbitrary element from the six just listed. Then the harmonic expansion is $E = k - s \cos(\phi - \varphi)$. The entries in Table 1 show the three parameters for each of F_ϕ , L_0 , L_1 , L_2 , δL_1 , δL_2 inside parentheses separated with commas.

The values used to generate the images are listed in the "simulation values" column in Table 1. The setting of the study is based on an earlier mission with a 4×12 deg pixel image and a maximum of 1000 counts per pixel. These values were used to define the response function of the ENA camera used in this study. The SPSA algorithm was then applied to estimate the model parameters using a set of initial values, as shown in the entries of Table 1 in the SPSA initial values column, to start the multiple-objective SPSA iterations. The initial values are inside the physically expected ranges. For example, the constant part of L_2 is expected to be between 4.00 and 6.00; thus, a value of 4.50 was given. The studies show the convergence of the estimation is independent of initial values. Many local solutions exist between the initial values and the final estimates.

There are three stages in the simulation study. The averages of the final estimates of the runs from the second and third stages are listed in the second stage estimates and final estimates columns of Table 1. Each stage consists of a total of 10 runs; each run is an SPSA process with identical setting, except the seeds used for the pseudorandom number generator are different. Each run consists of 150 iterations; each iteration uses four function evaluations (a total of 600 function evaluations); a function evaluation is defined by one round of computations of all pixels in the image of the magnetosphere field. As mentioned, the initial values of the SPSA runs in stage 1 are the values shown in Table 1; the initial values for the runs in the subsequent stages are the average of the final estimates of the runs in the previous stages.

The initial gain constants a and c used in SPSA were set to 0.005 and 0.001 at stage 1 for each run. Following the practical guidelines from Spall,¹⁹ the reduction rate α for a is 0.602 and γ for c is 0.101. The reduction for a is applied continuously as regular SPSA, but the reduction for c is applied once every 30 iterations. The gain constants for the following runs in stages 2 and 3 are a continuation of stage 1 runs.

An obvious question is how well the multiple-objective SPSA inversion process estimates its parameters. Table 1 shows differences in parameters between the simulation values column and the estimates. These differences are suspected to come primarily from the random fluctuations in the image pixels, which act as a kind of measurement noise; the multiple-objective SPSA iterations converge to a solution that represents a model for the ion intensity that would mimic the single global image without any addition of random Poisson fluctuations. To verify this, another simulated global image was generated under the same setting without the Poisson statistics in the counts and a number of estimation runs were made using this data. The total estimated parameter errors based on simulation values as truth converged below the 10^{-9} error level for all the runs. The total estimated parameter error is the summation over the 18 parameters of Table 1 of the difference squares between each individual estimate and its target value.

Theoretically the SPSA iterations asymptotically converge under some conditions stated in Ref. 11 and the differences between estimates of the SPSA algorithm and the true values are asymptotically normal at each iteration.¹¹ If the conditions of the multiple-objective SPSA runs satisfy the same convergence conditions of the SPSA algorithm, the multiple-objective SPSA iterations here will asymptoti-

cally converge, and the errors of the estimates will also be asymptotically normal. There are variations among the estimates of the multiple-objective SPSA runs at each iteration, the standard deviation of the variations is larger at the beginning iterations; the standard deviation becomes smaller when there are more iterations. Among three stages of multiple-objective SPSA runs, stage 1 reduces the estimated parameter error quickly, stage 2 finalizes the estimates, and stage 3 is just for confirmation. In fact, the values of the estimated model change very slightly between stages 2 and 3. The deviations of the final estimates between runs within the stage have narrowed for stage 3 in comparison with the deviation for stage 2. The maximum deviation in stage 3 for a single parameter is less than 0.3% of parameter value (0.003 in parameter units), which is negligible in most applications of the model. In practical applications, a single run can be used for stages 2 and 3.

The success in extracting the distribution from the ENA images can be judged by deviations of the estimated parameters from the simulated parameters for the ion-intensity model (the final estimates and simulation values columns in Table 1) for the sensor resolution of 4×12 deg pixels. The values in Table 1 show that the constants of the element parameters can be recovered pretty well. The scale factors and the phase angles of the element parameters are shifted. The simulation studies show that the estimation errors due to the estimation method are small in comparison with the instrumentation errors after a sufficient number of estimation iterations. To improve the accuracy, one must decrease the measurement errors, for example, by increasing the sensor resolution that reduces the random fluctuations, improving the model structure to reduce the sensitivities to the random fluctuation in the data, or introducing replicated data points to average the random fluctuations.

6 Comparison of SPSA with SAN

A well-known method for a global optimization problem is SAN. The question of how well SPSA compares with SAN is of interest given that SAN has been widely applied in optimization. Because SAN appears not to have been applied to multiple objective function problems, the comparison is based only on the results from the single-objective function—sum square measurement-residual error. For the single objective function, neither SPSA nor SAN shows any sign of convergence within 300 function evaluations if the same initial values are used as in stage 1 of the multiple-objective SPSA study. When the initial values are close to the solutions (using the initial values of stage 2), the single objective function SAN does show signs of convergence after tuning. Therefore, this study compares these two algorithms starting from a set of parameters close to the simulation values, comparable to the initial values of stage 2.

A total of 10 runs was studied. Figure 1 shows that at the 300th function evaluation, the level of mean sum square error is about 3:1 in favor of SPSA. It also shows a 3:1 ratio in the number of function evaluations required to achieve the same level of accuracy. (The total error level at the 300th function evaluation for SAN is about equal to the total error level at the 100th function evaluation for SPSA.) Because the results are in favor of SPSA, there is no further development for SAN using multiple objective functions.




Fig. 1 Convergence of SPSA versus SAN.

Also, with SAN it is inherently difficult to determine the acceptance of the updated parameter values at each SAN step for multiple objective functions. The averaged absolute mean total measurement-residual error for multiple-objective SPSA is also much lower than for the single objective SPSA or SAN (compare with Fig. 1).

7 Conclusions

This paper applies the SPSA algorithm for two types of objective functions, pixel-wide multiple objective functions and a total summation function. The comparison of sensor residuals for simulated annealing and for single- and multiple-objective SPSA algorithms indicates that the multiple-objective function SPSA algorithm gives the lowest total measurement residual errors. Also, the data noise level and resolution in the simulation study are adequate to estimate the constant parameters in the Chase and Roelof model²; some of the other parameters are not so well estimated due primarily to the finite sample with random fluctuations in the data counts.

When the multiple-objective SPSA algorithm is applied to a real satellite mission for the global magnetospheric image problem, there will be other error sources affecting the final estimates. Reductions of those errors are not within the scope of this paper. The paper shows that multiple-objective SPSA could be used to reduce the estimation errors and to isolate the other error sources. This may help the satellite mission planning and data analysis work for the global magnetospheric image problem.

This paper also compares the efficiency and accuracy levels for SPSA and SAN. In a single objective function setting, the results of the comparison were 3:1 in favor of SPSA in both efficiency and accuracy at the given measures. An additional advantage in using SPSA is that it can be extended to an infinite data setting. Instead of reiterating on a fixed global image, SPSA could iterate on consecutive sets of image counts. This changes the objective functions into noisy functions, which work well for SPSA. The ion intensity in the magnetosphere changes over time; the underlying model of the consecutive image counts also changes from image to image. SPSA with gain-sequence manipulation, as demonstrated in Ref. 20, could adapt to the changes in the model. (This is different from restarting the iteration from the final estimates in the previous data images.)

A second-order version of SPSA presented in Ref. 21 can be used to speed up the convergence after the estimates of the parameters are isolated within the region containing


the global optimization solution. When second-order SPSA is used, the constraints of the estimates for regular SPSA should also apply for the second order. This paper uses only the "standard" (first-order) SPSA approach and therefore omits the details of the second-order SPSA.

Acknowledgments

The problem was suggested by Drs. C. J. Chase and E. C. Roelof who provided the line-of-sight integration algorithm for the simulated ENA images in this study. This author appreciates Dr. D. Payne for his help in editing this paper. This paper was jointly supported by the Johns Hopkins University Applied Physics Laboratory Internal Research and Development program and the National Aeronautics and Space Administration Space Research program, Grant No. NAGW-2691.

References

1. E. C. Roelof, "Energetic neutral atom image of a storm—time ring current," *Geophys. Rev. Lett.* **14**, 652–655 (1987).
2. C. J. Chase and E. C. Roelof, "Extracting evolving structures from global magnetospheric images via model fitting and video visualization," *Johns Hopkins APL Tech. Dig.* **16**(2), 111–122 (1995).
3. J. C. Spall, *Bayesian Analysis of Time Series and Dynamic Models*, Dekker (1988).
4. P. K. Venkatesh, M. H. Cohen, R. W. Carr, and A. M. Dean, "Bayesian method for global optimization," *Phys. Rev. E* **55**(5), 6219–6232 (1997).
5. A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. R. Stat. Soc. Ser. B* **39**, 1–38 (1977).
6. D. R. Haley, J. P. Garner, and W. S. Levine, "Efficient maximum likelihood identification of a positive semi-definite covariance of initial population statistics," in *Proc. 1984 Am. Control Conf.*, pp. 1085–1089 (1984).
7. M. Bazaraa and C. M. Shetty, *Nonlinear Programming*, Chapter 8, Wiley, New York (1979).
8. J. Kiefer and J. Wolfowitz, "Stochastic estimation of a regression function," *Ann. Math. Stat.* **23**, 462–466 (1952).
9. S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, "Optimization by simulated annealing," *Science* **220**(4598), 671–680 (1983).
10. J. R. Ray and R. W. Harris, "Simulated annealing in the microcanonical ensemble," *Phys. Rev. E* **55**(5, Pt. a), 5270–5274 (1997).
11. J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Trans. Autom. Control* **27**, 332–341 (1992).
12. D. C. Chin, "Comparative study of stochastic algorithms for system optimization based on gradient approximations," *IEEE Trans. Syst. Man Cybern.* **27**(2), 244–249 (1997).
13. D. C. Chin, "A more efficient global optimization algorithm based on Styblinski and Tang," *Neural Networks* **7**, 573–574 (1994).
14. M. A. Styblinski and T. S. Tang, "Experiments in nonconvex optimization: stochastic approximation with function smoothing and simulated annealing," *Neural Networks* **3**, 467–483 (1990).
15. A. Alessandri and T. Parisini, "Nonlinear modeling of complex large-scale plants using neural networks and stochastic approximation," *IEEE Trans. Syst. Man Cybern.* **750**–757 (1997).
16. D. C. Chin and R. Srinivasan, "Electrical conductivity object locator: location of small objects buried at shallow depths," in *UXO Forum'97 Conf. Proc.*, pp. 50–57 (1997).
17. R. A. Wolf, R. W. Spino, and F. J. Rich, "Extension of the Rice convection model into the high-latitude ionosphere," *J. Atmos. Terr. Phys.* **53**, 817–829 (1991).
18. M. Fok, T. E. Moore, J. U. Kozyra, G. C. Ho, and D. C. Hamilton, "Three-dimensional ring current decay model," *J. Geophys. Res.* **100**(A6), 9619–9632 (1995).
19. J. C. Spall, "Implementation of the simultaneous perturbation algorithm for stochastic optimization," *IEEE Trans. Aerospace Electron. Syst.* (in press).
20. J. C. Spall and J. A. Cristion, "Nonlinear adaptive control using neural networks: estimation based on a smoothed form of simultaneous perturbation gradient approximation," *Stat. Sin.* **4**, 1–27 (1994).
21. J. C. Spall, "Adaptive stochastic approximation by the simultaneous perturbation method," in *Proc. 36th IEEE Conf. on Decision and Control* (in press).



Daniel C. Chin received his BS degree in mathematics from Chung Yuan University, Taiwan, in 1966 and his MS degree in mathematics from Northern Illinois University in 1970. Mr. Chin is a member of the senior professional staff of the Johns Hopkins University Applied Physics Laboratory and is experienced in stochastic approximation, Bayesian analysis, statistical estimation and simulation, and image data processing. His accomplishments include the identification of buried objects (metallic or nonmetallic) with induced electric current, control strategies and simulation studies for systemwide traffic-adaptive control, and an optimal representation of the Kalman filter process and its square root formulation. Mr. Chin is a member of the IEEE society, Sigma Xi, and the team that won the Hart Prize for most outstanding IR&D project at the Johns Hopkins University Applied Physics Lab in 1990.