Stochastic Newton methods with enhanced Hessian estimation

Danda Sai Koti Reddy

M.Sc(engg) thesis defence

Research Advisor : Prof. Shalabh Bhatnagar

Stochastic Systems Lab,
Department of Computer Science and Automation,
Indian Institute of Science, Bangalore

Overview

Simulation Optimization

Random Directions Stochastic Approximation (RDSA) + Improved Hessian Estimation

Improved Hessian Estimation for Simultaneous Perturbation Stochastic Approximation with 3 Simulations (SPSA-3)

Numerical Results

Simulation Optimization

Optimization under Uncertainity

Energy Demand Management

- Consumer demand, energy generation are uncertain
- Objective is to minimize absolute difference

Optimization under Uncertainity

Energy Demand Management

- Consumer demand, energy generation are uncertain
- Objective is to minimize absolute difference

Traffic Signal Control

- Optimal order to switch traffic lights
- Objective is to minimize waiting time

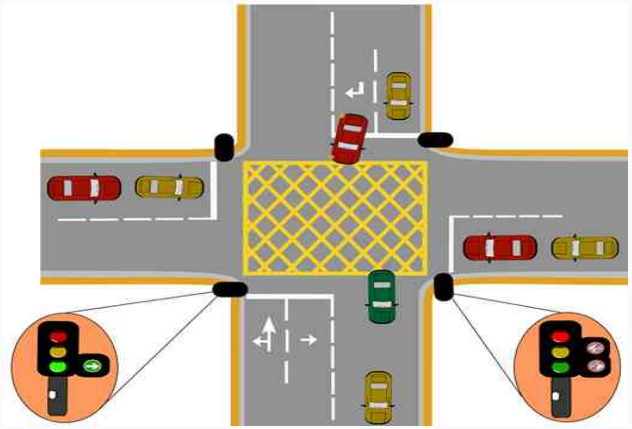

Basic optimization problem

To find θ^* that minimizes the objective function $\mathrm{f}(\theta)$:

$$\theta^* = \arg\min_{\theta \in \Theta} \mathrm{f}(\theta) \tag{1}$$

- $\mathrm{f}\colon \mathbb{R}^{\mathrm{N}} \to \mathbb{R}$ is called the objective function
- θ is tunable N-dimensional parameter
- $\Theta \subseteq \mathbb{R}^{\mathrm{N}}$ is the feasible region in which θ takes values

Classification of optimization problems

Deterministic optimization problem

- Complete information about objective function f
- First and higher order derivatives
- Feasible region

Classification of optimization problems

Deterministic optimization problem

- Complete information about objective function f
- First and higher order derivatives
- Feasible region

Stochastic optimization problem

- We have little knowledge on the structure of f
- f cannot be obtained directly
- $f(\theta) \equiv E_{\xi}[h(\theta, \xi)]$, where ξ comprises the randomness in the system

Difficult to find θ^* only on the basis of noisy samples

Stochastic optimization via simulation

Stochastic optimization deals with highly nonlinear and high dimensional systems. The challenges with these problems are:

- Too complex to solve analytically
- Many simplifying assumptions are required

Stochastic optimization via simulation

Stochastic optimization deals with highly nonlinear and high dimensional systems. The challenges with these problems are:

- Too complex to solve analytically
- Many simplifying assumptions are required

A good alternative of modelling and analysis is "Simulation"

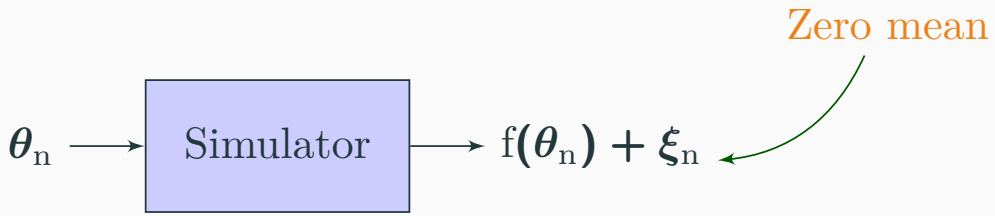

Figure 1: Simulation optimization

Stochastic analog of gradient descent

$$\theta_{n+1} = \Gamma_{\Theta}\left[\theta_n - a_n \widehat{\nabla} f(\theta_n)\right] \tag{2}$$

- $\widehat{\nabla} f(\theta_n)$ is a noisy estimate of the gradient $\nabla f(\theta_n)$, and it should satisfy $E\left[\widehat{\nabla} f(\theta_n)\right] - \nabla f(\theta_n) \to 0$
- $\{a_n\}$ are pre-determined step-sizes satisfying:

$$\sum_{n=1}^{\infty} a_n = \infty, \quad \sum_{n=1}^{\infty} a_n^2 < \infty$$

- Γ_{Θ} denotes the projection of a point onto Θ

2-slide summary

Related second-order methods

(Spall 2000)[1]	Second-order SPSA (2SPSA)	4 simulations/iteration
(Spall 2009)[2]	2SPSA + feedback	4 simulations/iteration
(Prashanth L.A. et al 2016)[3]	Second-order RDSA (2RDSA)	3 simulations/iteration
(S. Bhatnagar et al 2015)[4]	Second-order SPSA-3 (2SPSA-3)	3 simulations/iteration

[1] J. C. Spall (2000), "Adaptive stochastic approximation by the simultaneous perturbation method," IEEE TAC.

[2] J. C. Spall (2009), "Feedback and weighting mechanisms for improving Jacobian estimates in the adaptive simultaneous perturbation algorithm," IEEE TAC.

[3] Prashanth L. A, Shalabh Bhatnagar, Michael Fu, Steve Marcus (2016), "Adaptive system optimization using random directions stochastic approximation," IEEE TAC.

[4] S. Bhatnagar, Prashanth L. A (2015) ,"Simultaneous perturbation Newton algorithms for simulation optimization," JOTA.

2-slide summary

Our work

- We propose generalised RDSA algorithm[1] + feedback and weighting mechanisms for improving Hessian estimate[2]

- We propose feedback and weighting mechanisms for improving Hessian estimate of 2SPSA-3 algorithm[1]

---

[1] Under preparation - "https://github.com/dsai1215/asgoodasitgets/tree/master/Journal".

[2] D. Sai Koti Reddy, Prashanth L.A, Shalabh Bhatnagar (2016), "Improved Hessian estimation for adaptive random directions stochastic approximation," IEEE CDC 2016.

Random Directions Stochastic Approximation (RDSA) + Improved Hessian Estimation

Our algorithm

- Matrix projection
- Gradient estimate

$$\theta_{n+1} = \theta_n - a_n \Upsilon (\overline{H}_n)^{-1} \widehat{\nabla} f(\theta_n) \tag{3}$$

Our algorithm

- Matrix projection
- Gradient estimate

$$\theta_{n+1} = \theta_n - a_n \, \Upsilon \, (\overline{H}_n)^{-1} \, \widehat{\nabla} f(\theta_n) \tag{3}$$

$$\overline{H}_n = (1 - b_n)\overline{H}_{n-1} + b_n (\widehat{H}_n - \widehat{\psi}_n) \tag{4}$$

- Optimal step-sizes
- Hessian estimate
- Feedback term

Overall flow of 2RDSA-IH

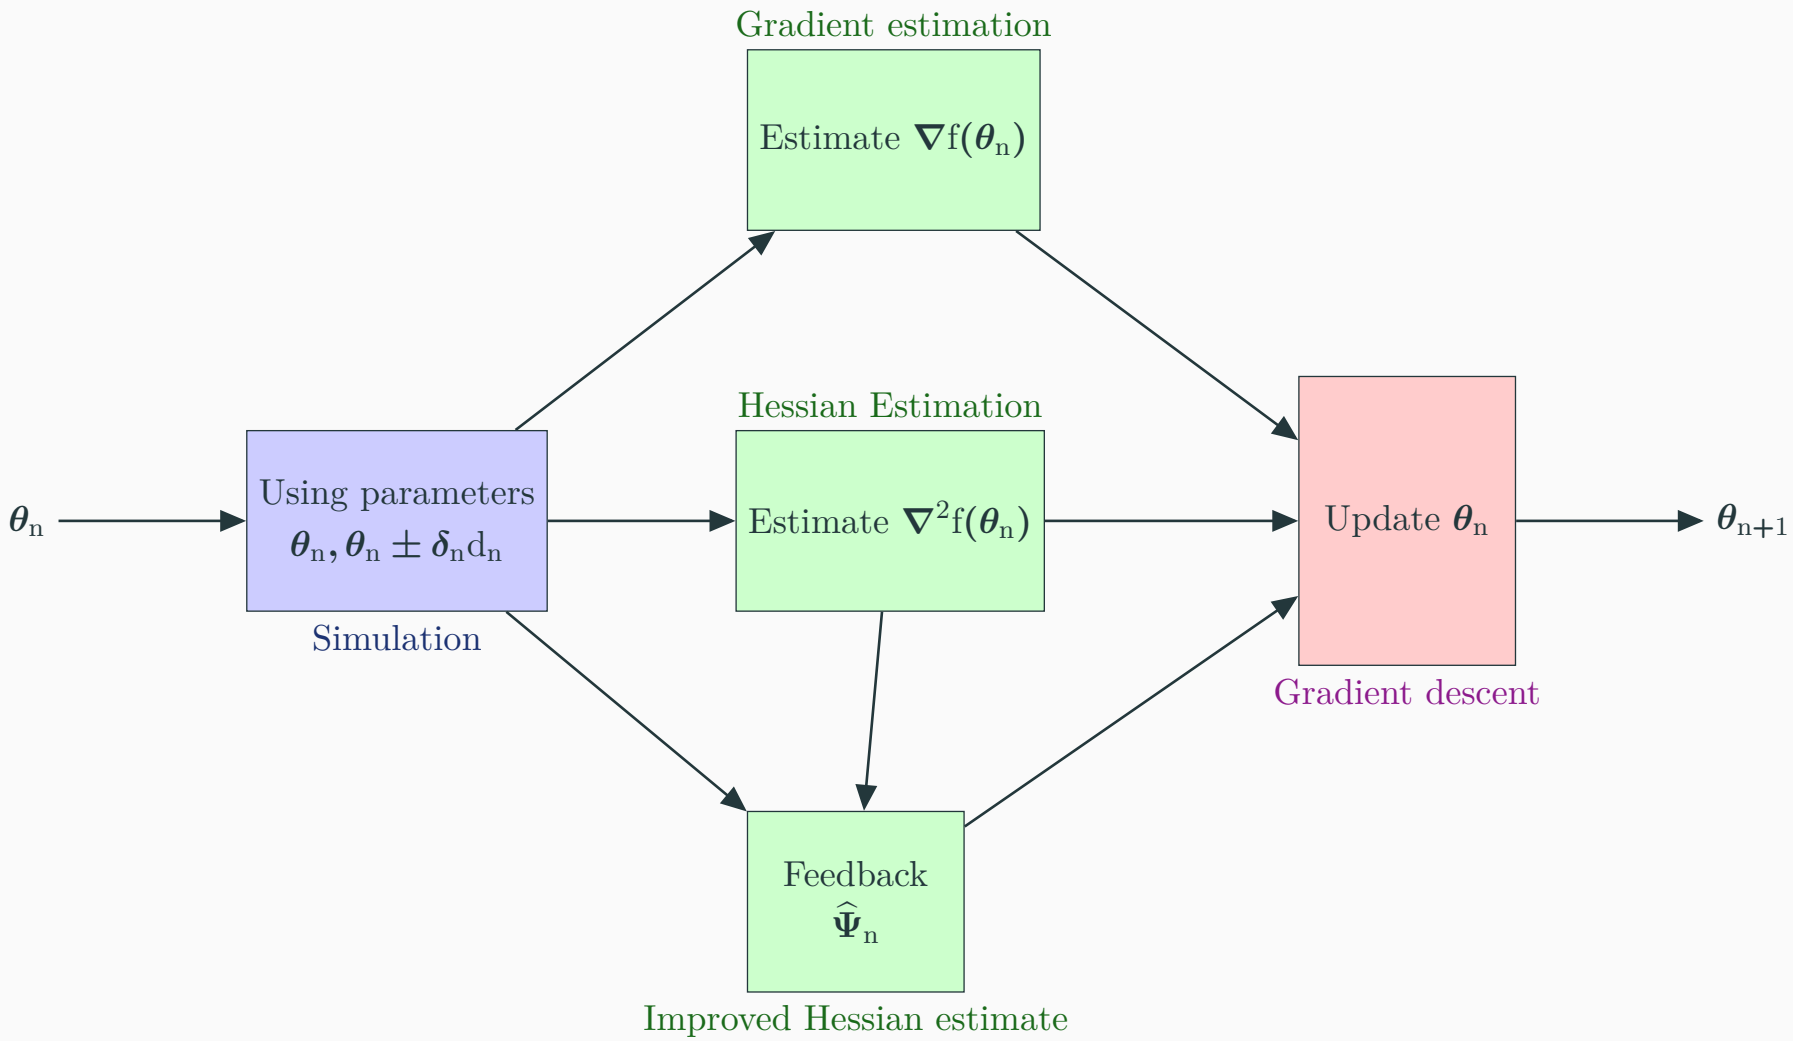

RDSA gradient estimate

Function measurements

$$\mathrm{y}_\mathrm{n}^{+} = \mathrm{f}(\ \theta_\mathrm{n} + \delta_\mathrm{n}\mathrm{d}_\mathrm{n}\ ) + \xi_\mathrm{n}^{+}, \quad \mathrm{y}_\mathrm{n}^{-} = \mathrm{f}(\ \theta_\mathrm{n} - \delta_\mathrm{n}\mathrm{d}_\mathrm{n}\ ) + \xi_\mathrm{n}^{-}$$

RDSA gradient estimate

Function measurements

$$y_n^+ = f(\theta_n + \delta_n d_n) + \xi_n^+, \quad y_n^- = f(\theta_n - \delta_n d_n) + \xi_n^-$$

Gradient estimate

$$\widehat{\nabla} f(\theta_n) = \frac{1}{\lambda} d_n \left[\frac{y_n^+ - y_n^-}{2\delta_n} \right] \tag{5}$$

Where $d_n = (d_n^1, \ldots, d_n^N)^T$, and $\lambda = \mathbb{E}(d_n^i)^2$

2RDSA Hessian estimate

Function measurements

$$y_n^+ = f(\theta_n + \delta_n d_n) + \xi_n^+, \quad y_n^- = f(\theta_n - \delta_n d_n) + \xi_n^-, \quad y_n = f(\theta_n) + \xi_n$$

2RDSA Hessian estimate

Function measurements

$$y_n^+ = f(\theta_n + \delta_n d_n) + \xi_n^+, \quad y_n^- = f(\theta_n - \delta_n d_n) + \xi_n^-, \quad y_n = f(\theta_n) + \xi_n$$

Hessian estimate $\widehat{H}_n$

$$\begin{aligned}
\widehat{H}_n &= M_n \left(\frac{y_n^+ + y_n^- - 2y_n}{\delta_n^2} \right) \\
&= M_n \left[\left(\frac{f(\theta_n + \delta_n d_n) + f(\theta_n - \delta_n d_n) - 2f(\theta_n)}{\delta_n^2} \right) \right. \\
&\qquad \left. + \left(\frac{\xi_n^+ + \xi_n^- - 2\xi_n}{\delta_n^2} \right) \right] \\
&= M_n \left(d_n^T \nabla^2 f(\theta_n) d_n + O(\delta_n^2) + \left(\frac{\xi_n^+ + \xi_n^- - 2\xi_n}{\delta_n^2} \right) \right) \qquad (6)
\end{aligned}$$

Want to recover $\nabla^2 f(\theta_n)$ from this

Zero-mean

How to choose M_n?

$$M_n = \begin{bmatrix} \frac{1}{\kappa}\left((d_n^1)^2 - \lambda\right) & \cdots & \frac{1}{2\lambda^2} d_n^1 d_n^N \\ \frac{1}{2\lambda^2} d_n^2 d_n^1 & \cdots & \frac{1}{2\lambda^2} d_n^2 d_n^N \\ \cdots & \cdots & \cdots \\ \frac{1}{2\lambda^2} d_n^N d_n^1 & \cdots & \frac{1}{\kappa}\left((d_n^N)^2 - \lambda\right) \end{bmatrix} \tag{7}$$

where $\lambda = \mathbb{E}(d_n^i)^2$, $\tau = E(d_n^i)^4$, and $\kappa = \left(\tau - \lambda^2\right)$ for any $i = 1, \ldots, N$

Zero-mean feedback term

Zero-mean term

Mean of the Hessian estimate

$$\mathbb{E}\left[\widehat{\mathrm{H}}_{\mathrm{n}}\middle|\mathcal{F}_{\mathrm{n}}\right] = \nabla^2 \mathrm{f}(\theta_{\mathrm{n}}) + \mathbb{E}\left[\Psi_{\mathrm{n}}(\nabla^2 \mathrm{f}(\theta_{\mathrm{n}}))\middle|\mathcal{F}_{\mathrm{n}}\right] + \mathrm{O}(\delta_{\mathrm{n}}^2) + \mathbb{E}\left[\left(\frac{\xi_{\mathrm{n}}^{+} + \xi_{\mathrm{n}}^{-} - 2\xi_{\mathrm{n}}}{\delta_{\mathrm{n}}^2}\right)\middle|\mathcal{F}_{\mathrm{n}}\right] \tag{8}$$

Zero-mean

[1] For any matrix P, $[\mathrm{P}]_{\mathrm{D}}$ refers to a matrix that retains only the diagonal entries of P and replaces all the remaining entries with zero

[2] $[\mathrm{P}]_{\mathrm{N}}$ to refer to a matrix that retains only the off-diagonal entries of P, while replaces all the diagonal entries with zero

Zero-mean feedback term

Zero-mean term

Mean of the Hessian estimate

$$\mathbb{E}\left[\widehat{\mathrm{H}}_{\mathrm{n}}\middle|\mathcal{F}_{\mathrm{n}}\right] = \nabla^2 \mathrm{f}(\theta_{\mathrm{n}}) + \mathbb{E}\left[\Psi_{\mathrm{n}}(\nabla^2 \mathrm{f}(\theta_{\mathrm{n}}))\middle|\mathcal{F}_{\mathrm{n}}\right] + \mathrm{O}(\delta_{\mathrm{n}}^2) + \mathbb{E}\left[\left(\frac{\xi_{\mathrm{n}}^{+} + \xi_{\mathrm{n}}^{-} - 2\xi_{\mathrm{n}}}{\delta_{\mathrm{n}}^2}\right)\middle|\mathcal{F}_{\mathrm{n}}\right] \tag{8}$$

Zero-mean

Feedback term

$$\Psi_{\mathrm{n}}(\mathrm{H}) = [\mathrm{M}_{\mathrm{n}}]_{\mathrm{D}}\left(\mathrm{d}_{\mathrm{n}}^{\mathrm{T}}\,[\mathrm{H}]_{\mathrm{N}}\,\mathrm{d}_{\mathrm{n}}\right) + [\mathrm{M}_{\mathrm{n}}]_{\mathrm{N}}\left(\mathrm{d}_{\mathrm{n}}^{\mathrm{T}}\,[\mathrm{H}]_{\mathrm{D}}\,\mathrm{d}_{\mathrm{n}}\right) \tag{9}$$

[1] For any matrix P, $[\mathrm{P}]_{\mathrm{D}}$ refers to a matrix that retains only the diagonal entries of P and replaces all the remaining entries with zero

[2] $[\mathrm{P}]_{\mathrm{N}}$ to refer to a matrix that retains only the off-diagonal entries of P, while replaces all the diagonal entries with zero

Zero-mean feedback term

Problem

Feedback term is function of current Hessian $\nabla^2 \mathrm{f}$

Zero-mean feedback term

Problem

Feedback term is function of current Hessian $\nabla^2 \mathrm{f}$

Solution

Use $\overline{\mathrm{H}}_{\mathrm{n}-1}$ as a proxy for $\nabla^2 \mathrm{f}$

$$\widehat{\Psi}_{\mathrm{n}} = \Psi_{\mathrm{n}}(\,\overline{\mathrm{H}}_{\mathrm{n}-1}\,) \tag{10}$$

Step-size optimization

Recall the Hessian recursion, $\overline{\mathrm{H}}_{\mathrm{n}} = (1 - \mathrm{b}_{\mathrm{n}})\overline{\mathrm{H}}_{\mathrm{n}-1} + \mathrm{b}_{\mathrm{n}}(\widehat{\mathrm{H}}_{\mathrm{n}} - \widehat{\Psi}_{\mathrm{n}})$

Step-size optimization

Recall the Hessian recursion, $\overline{H}_n = (1 - b_n)\overline{H}_{n-1} + b_n(\widehat{H}_n - \widehat{\Psi}_n)$

Rewriting the Hessian recursion

$$\overline{H}_n = \sum_{i=0}^{n} \tilde{b}_i(\widehat{H}_i - \widehat{\Psi}_i) \tag{11}$$

Step-size optimization

Recall the Hessian recursion, $\overline{H}_n = (1 - b_n)\overline{H}_{n-1} + b_n(\widehat{H}_n - \widehat{\Psi}_n)$

Rewriting the Hessian recursion

$$\overline{H}_n = \sum_{i=0}^{n} \tilde{b}_i(\widehat{H}_i - \widehat{\Psi}_i) \tag{11}$$

Optimization problem for weights

$$\min_{\{\tilde{b}_i\}} \sum_{i=0}^{n} (\tilde{b}_i)^2 \delta_i^{-4}, \text{ subject to} \tag{12}$$

$$\tilde{b}_i \geq 0 \ \forall i \text{ and } \sum_{i=0}^{n} \tilde{b}_i = 1 \tag{13}$$

Step-size optimization

Above optimization problem solution

$$\tilde{b}_i^* = \delta_i^4 / \sum_{j=0}^{n} \delta_j^4, i = 1, \ldots, n \tag{14}$$

[1] Step-size optimization is a relatively straightforward migration from Spall 2009

Step-size optimization

Above optimization problem solution

$$\tilde{b}_i^* = \delta_i^4 / \sum_{j=0}^{n} \delta_j^4, i = 1, \ldots, n \tag{14}$$

Optimal weights for original Hessian recursion

$$b_i = \delta_i^4 / \sum_{j=0}^{i} \delta_j^4 \tag{15}$$

[1] Step-size optimization is a relatively straightforward migration from Spall 2009

Improved Hessian Estimation for Simultaneous Perturbation Stochastic Approximation with 3 Simulations (SPSA-3)

SPSA gradient estimate[1]

Function measurements

$$y_n^+ = f(\theta_n + \delta_n \Delta_n) + \xi_n^+, \quad y_n^- = f(\theta_n - \delta_n \Delta_n) + \xi_n^-$$

Gradient estimate

$$\widehat{\nabla} f(\theta_n) = \begin{pmatrix} \frac{f(\theta_n + \delta_n \Delta_n) - f(\theta_n - \delta_n \Delta_n)}{2\delta_n \Delta_{n1}} + \frac{\xi_n^+ - \xi_n^-}{2\delta_n \Delta_{n1}} \\ \vdots \\ \frac{f(\theta_n + \delta_n \Delta_n) - f(\theta_n - \delta_n \Delta_n)}{2\delta_n \Delta_{nN}} + \frac{\xi_n^+ - \xi_n^-}{2\delta_n \Delta_{nN}} \end{pmatrix} \tag{16}$$

[1] J. C. Spall (1992), "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," IEEE TAC.

SPSA-3 Hessian estimate[1]

Function measurements

$$y_n^{++} = f(\theta_n + \delta_n \Delta_n + \delta_n \widehat{\Delta}_n) + \xi_n^{++}, \qquad y_n = f(\theta_n) + \xi_n,$$
$$y_n^{--} = f(\theta_n - \delta_n \Delta_n - \delta_n \widehat{\Delta}_n) + \xi_n^{--}$$

Hessian estimate $\widehat{H}_n$

$$\left(\widehat{H}_n\right)_{ij} = \left(\frac{y_n^{++} + y_n^{--} - 2y_n}{2\delta_n^2 \Delta_n^{(i)} \widehat{\Delta}_n^{(j)}}\right) \tag{17}$$

[1] S. Bhatnagar, Prashanth L. A (2015) ,"Simultaneous perturbation Newton algorithms for simulation optimization," JOTA.

Zero-mean feedback term

Simplified Hessian estimate

$$\widehat{\mathrm{H}}_{\mathrm{n}} = \nabla^2 \mathrm{f}(\theta_{\mathrm{n}}) + \Psi_{\mathrm{n}}(\nabla^2 \mathrm{f}(\theta_{\mathrm{n}})) + \mathrm{O}(\delta_{\mathrm{n}}^2) + \mathrm{O}(\delta_{\mathrm{n}}^{-2}) \tag{18}$$

Feedback term

$$\Psi_{\mathrm{n}}(\mathrm{H}) = \frac{1}{2}\mathrm{P}_{\mathrm{n}}\left[\Delta_{\mathrm{n}}^{\mathrm{T}}\mathrm{H}\Delta_{\mathrm{n}} + \widehat{\Delta}_{\mathrm{n}}^{\mathrm{T}}\mathrm{H}\widehat{\Delta}_{\mathrm{n}}\right] + \widehat{\mathrm{N}}_{\mathrm{n}}^{\mathrm{T}}\mathrm{H}\mathrm{N}_{\mathrm{n}} + \widehat{\mathrm{N}}_{\mathrm{n}}^{\mathrm{T}}\mathrm{H} + \mathrm{H}\mathrm{N}_{\mathrm{n}} \tag{19}$$

Where,

$\mathrm{P}_{\mathrm{n}} = [1./\Delta_{\mathrm{n}}][1./\Delta_{\mathrm{n}}]^{\mathrm{T}},\ \mathrm{N}_{\mathrm{n}} = \Delta_{\mathrm{n}}[1./\Delta_{\mathrm{n}}]^{\mathrm{T}} - \mathrm{I}_{\mathrm{N}},\ \widehat{\mathrm{N}}_{\mathrm{n}} = \widehat{\Delta}_{\mathrm{n}}[1./\widehat{\Delta}_{\mathrm{n}}]^{\mathrm{T}} - \mathrm{I}_{\mathrm{N}}$

Lemma

(Bias in Hessian estimate) Under assumptions similar to those for 2SPSA and 2RDSA, we have a.s. that[1], for $i, j = 1, \ldots, N$,

$$\left| \mathbb{E}\left[\widehat{H}_n(i,j) \middle| \mathcal{F}_n \right] - \nabla^2_{ij} f(\theta_n) \right| = O(\delta_n^2) \tag{20}$$

Theorem

(Strong Convergence of Hessian) Under assumptions similar to those for 2SPSA and 2RDSA, we have that

$$\theta_n \to \theta^*, \overline{H}_n \to \nabla^2 f(\theta^*) \text{ a.s. as } n \to \infty$$

[1] Here $\widehat{H}_n(i,j)$ and $\nabla^2_{ij} f(\cdot)$ denote the (i,j)th entry in the Hessian estimate $\widehat{H}_n$ and the true Hessian $\nabla^2 f(\cdot)$, respectively

Convergence analysis

Recall the Hessian recursion, $\overline{\mathrm{H}}_{\mathrm{n}} = (1 - \mathrm{b}_{\mathrm{n}})\overline{\mathrm{H}}_{\mathrm{n-1}} + \mathrm{b}_{\mathrm{n}}(\widehat{\mathrm{H}}_{\mathrm{n}} - \widehat{\Psi}_{\mathrm{n}})$

Convergence analysis

Recall the Hessian recursion, $\overline{\mathrm{H}}_{\mathrm{n}} = (1 - \mathrm{b}_{\mathrm{n}})\overline{\mathrm{H}}_{\mathrm{n-1}} + \mathrm{b}_{\mathrm{n}}(\widehat{\mathrm{H}}_{\mathrm{n}} - \widehat{\Psi}_{\mathrm{n}})$

Rewriting Hessian recursion as SA scheme

$$\begin{aligned}\overline{\mathrm{H}}_{\mathrm{n}} &= \overline{\mathrm{H}}_{\mathrm{n-1}} - \mathrm{b}_{\mathrm{n}}(\overline{\mathrm{H}}_{\mathrm{n-1}} - \widehat{\mathrm{H}}_{\mathrm{n}} + \widehat{\Psi}_{\mathrm{n}}) \\ &= \overline{\mathrm{H}}_{\mathrm{n-1}} - \mathrm{b}_{\mathrm{n}}(\overline{\mathrm{H}}_{\mathrm{n-1}} - \mathrm{H}^{*} + \widehat{\Psi}_{\mathrm{n}} - \Psi_{\mathrm{n}}(\mathrm{H}^{*})) \qquad (21)\end{aligned}$$

Recall the Hessian recursion, $\overline{\mathrm{H}}_{\mathrm{n}} = (1 - \mathrm{b}_{\mathrm{n}})\overline{\mathrm{H}}_{\mathrm{n}-1} + \mathrm{b}_{\mathrm{n}}(\widehat{\mathrm{H}}_{\mathrm{n}} - \widehat{\Psi}_{\mathrm{n}})$

Rewriting Hessian recursion as SA scheme

$$\begin{aligned}\overline{\mathrm{H}}_{\mathrm{n}} &= \overline{\mathrm{H}}_{\mathrm{n}-1} - \mathrm{b}_{\mathrm{n}}(\overline{\mathrm{H}}_{\mathrm{n}-1} - \widehat{\mathrm{H}}_{\mathrm{n}} + \widehat{\Psi}_{\mathrm{n}}) \\ &= \overline{\mathrm{H}}_{\mathrm{n}-1} - \mathrm{b}_{\mathrm{n}}(\overline{\mathrm{H}}_{\mathrm{n}-1} - \mathrm{H}^{*} + \widehat{\Psi}_{\mathrm{n}} - \Psi_{\mathrm{n}}(\mathrm{H}^{*})) \qquad (21)\end{aligned}$$

Theorem

(2SPSA-3-IH Quadratic case - Convergence rate) Let $\mathrm{b}_{\mathrm{n}} = \mathrm{b}_0/\mathrm{n}^{\mathrm{r}}$, where $1/2 < \mathrm{r} < 1$ and $0 < \mathrm{b}_0 \leq 1$, $\mathrm{H}^{*} = \nabla^2 \mathrm{f}(\theta^{*})$ and $\Lambda_{\mathrm{k}} = \overline{\mathrm{H}}_{\mathrm{k}} - \mathrm{H}^{*}$. Under noise-free setting, we have

$$\mathrm{trace}[\mathbb{E}(\Lambda_{\mathrm{n}}^{\mathrm{T}}\Lambda_{\mathrm{n}})] = \mathrm{O}(\mathrm{e}^{-2\mathrm{b}_0\mathrm{n}^{1-\mathrm{r}}/1-\mathrm{r}}) \qquad (22)$$

Numerical Results

Distribution for $d_n^i, i = 1, \ldots, N$

Asymmetric Bernoulli distribution[1]

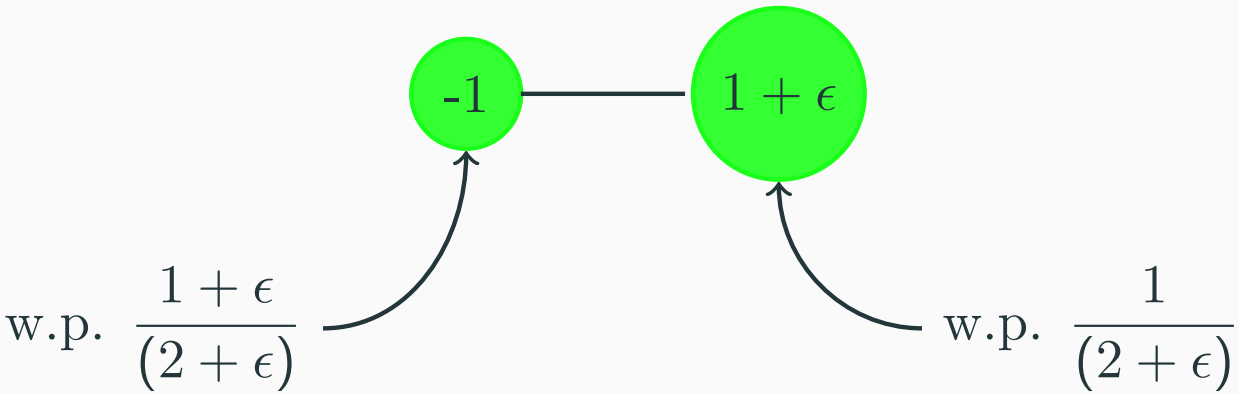

Uniform perturbations[1]

$$d_n^i = \text{Unif}[-\eta, \eta], \eta > 0 \tag{23}$$

[1] Prashanth L. A, Shalabh Bhatnagar, Michael Fu, Steve Marcus (2016), "Adaptive system optimization using random directions stochastic approximation," IEEE TAC.

Numerical Results

Quadratic loss

$$\mathrm{f}(\theta) = \theta^{\mathrm{T}} \mathrm{A} \theta + \mathrm{b}^{\mathrm{T}} \theta \tag{24}$$

Fourth-order loss

$$\mathrm{f}(\theta) = \theta^{\mathrm{T}} \mathrm{A}^{\mathrm{T}} \mathrm{A} \theta + 0.1 \sum_{\mathrm{j}=1}^{\mathrm{N}} (\mathrm{A}\theta)_{\mathrm{j}}^{3} + 0.01 \sum_{\mathrm{j}=1}^{\mathrm{N}} (\mathrm{A}\theta)_{\mathrm{j}}^{4} \tag{25}$$

Additive Noise : $[\theta^{\mathrm{T}}, 1]\mathrm{Z}$, where $\mathrm{Z} \approx \mathcal{N}(0, \sigma^2 \mathrm{I}_{\mathrm{N}+1 \times \mathrm{N}+1})$

[1] The implementation is available at https://github.com/prashla/RDSA/archive/master.zip

Numerical Results

Normalized MSE (NMSE)

$$\left\|\theta_{\mathrm{n_{end}}} - \theta^*\right\|^2 / \left\|\theta_0 - \theta^*\right\|^2 \tag{26}$$

Normalized loss

$$\mathrm{f}(\theta_{\mathrm{n_{end}}})/\mathrm{f}(\theta_0) \tag{27}$$

Table 1: Normalized loss values for fourth-order objective (25) with noise: simulation budget = 10,000 and standard error from 500 replications shown after ±

Noise parameter $\sigma = 0.1$		
	Regular	Improved Hessian estimation
2SPSA	0.132 ± 0.0267	0.104 ± 0.0355
2SPSA-3	0.0951 ± 0.0031	0.0594 ± 0.0014
2RDSA-Unif	0.115 ± 0.0214	0.0271 ± 0.0538
2RDSA-AsymBer	0.0471 ± 0.021	0.0099 ± 0.0014

[1] Observation 1: Schemes with improved Hessian estimation performs better than their respective regular schemes

[2] Observation 2: 2RDSA-IH-AsymBer is performing the best overall

Table 2: NMSE values for quadratic objective (24) with noise: simulation budget = 10,000 and standard error from 500 replications shown after ±

Noise parameter $\sigma = 0.1$		
	Regular	Improved Hessian estimation
2SPSA	0.9491 ± 0.0131	0.5495 ± 0.0217
2SPSA-3	0.8378 ± 0.0179	0.1045 ± 0.0005
2RDSA-Unif	1.0073 ± 0.0140	0.1953 ± 0.0095
2RDSA-AsymBer	0.1667 ± 0.0095	0.0324 ± 0.0007

[1] Observation 1: Schemes with improved Hessian estimation performs better than their respective regular schemes

[2] Observation 2: 2RDSA-IH-AsymBer is performing the best overall

Conclusions

- Proposed generalised RDSA algorithm + Improved Hessian estimation scheme
- Improved Hessian estimation scheme for 2SPSA-3 algorithm
- 2RDSA-IH, 2SPSA-3 requires only 75% of the simulation cost per-iteration for 2SPSA, 2SPSA-IH

Future work

- To improve rate of convergence of first-order methods by incorporating ideas of momentum descent, Hessian-free methods, and conjugate methods
- To develop stochastic approximation versions of quasi-Newton schemes given by the Broyden family
- To explore applications of these methods to the design of reinforcement learning algorithms
- Deterministic perturbations for Second-order methods

Thank You