

ON EFFICIENCIES OF STOCHASTIC OPTIMIZATION PROCEDURES UNDER IMPORTANCE SAMPLING

Henry Lam

Department of Industrial Engineering
and Operations Research
Columbia University
New York, NY 10027, USA

Guangxin Jiang

School of Management
Shanghai University
Shanghai, 200444, CHINA

Michael C. Fu

Robert H. Smith School of Business and Institute for Systems Research
University of Maryland
College Park, MD, 20742, USA

ABSTRACT

We study and compare the efficiencies of stochastic approximation (SA) and sample average approximation (SAA) for stochastic optimization problems when the decision variables are inside the generating probability distributions or when importance sampling can be applied (as in the case of simulation optimization with Monte Carlo samples). We explain how SA is statistically more efficient than SAA in such contexts, a behavior different from conventional situations where SAA is usually held as the statistically optimal procedure. We support our claim with a theoretical minimax framework and a weak duality argument. We also demonstrate our theoretical findings with some simple simulation examples.

1 INTRODUCTION

This paper compares the efficiencies of two commonly used methods in stochastic optimization, namely stochastic approximation (SA) and sample average approximation (SAA). We are particularly interested in the situation where one is allowed to control the probability distribution that drives the stochasticity in the problem. This occurs, for instance, when the decision variable lies inside the distribution, or when we consider applying importance sampling (IS) to sample the stochasticity. These situations arise quite often in simulation optimization where the objective function is an expectation that can be evaluated from Monte Carlo samples under some parametric distribution.

When the distribution driving the stochasticity is fixed, or when the Monte Carlo samples, or data, are given in advance, SAA is usually held as the statistically optimal procedure. By this we mean the following asymptotic property. Suppose the objective function is smooth enough and there is a unique optimal solution that lies in the interior of the feasible region, then, under some mild conditions, both the estimated optimal solutions of SA and SAA satisfy respective central limit theorems. Under these theorems, the asymptotic variance of the solution from SAA is at most that of SA, regardless of what stepsize it takes (when the solution is multi-dimensional, this comparison is based on a matrix inequality between the covariance matrices of the solutions or an inequality between the variances of the attained objective values; when the solution is single-dimensional, the meaning of this comparison is obvious).

In this paper, we explore how the above conclusion should be modified when one has the flexibility to change the distribution used in driving the stochasticity. Because SA is a sequential scheme that updates the decision variable and generates Monte Carlo samples “on the fly”, the distribution can be chosen sequentially based on the current best solution guess. On the other hand, SAA uses the samples at the beginning to form a deterministic problem. Thus, any flexibility in choosing the distribution only benefits the initial setup of the problem, which cannot meaningfully take into account any updated information about the solution. To rigorize the above intuition, we will use a minimax framework and show concretely that SA always has a better asymptotic variance than SAA, under an idealized setting that allows us to use a weak duality argument. We focus on the case of single-dimensional decision in this paper, which is simpler and illustrates our claim transparently. We will also discuss similar implications with finite-sample bounds, and substantiate our claims with some simple numerical examples.

We close this introduction with a brief literature review. SA is an analog of gradient descent in nonlinear programming to the stochastic setting, which produces iterates that converge to the root of a stochastic equation or an optimum of a stochastic optimization problem. When there exists an unbiased gradient estimator, one could use the Robbins-Monro (R-M) algorithm (Robbins and Monro 1951), which we focus on in this paper. When such an estimator is unavailable, one could resort to finite-difference schemes, which give rise to the Kiefer-Wolfowitz (K-W) algorithm (Kiefer and Wolfowitz 1952). Variants of these algorithms have been substantially studied, including notably the simultaneous perturbation stochastic approximation (Spall 1992) and Gaussian smoothing (Nesterov and Spokoiny 2017) for efficient biased gradient estimation, averaging schemes (Polyak and Juditsky 1992; Nemirovski et al. 2009) and scaled-and-shifted schemes (Broadie et al. 2011) for enhancing algorithmic efficiency, and other recent variants such as the accelerated SA (Ghadimi and Lan 2013) and the secant-tangent averaged SA (Chau and Fu 2018). More details of SA can be found in Bottou (2010), Chau and Fu (2015), Ghadimi and Lan (2015). On the other hand, SAA replaces an expected value objective function with a sample average from Monte Carlo samples. This formulation then allows the use of deterministic optimization routines for an approximate solution. Theoretical properties of SAA, including consistency, asymptotic normality and optimality gaps, can be found in Kleywegt et al. (2002), Mak et al. (1999), and Shapiro and Nemirovski (2005). More details are in Shapiro (2003), Kim et al. (2015) and the references therein. Due to its simplicity, SAA has been applied in many fields such as finance (Alexander et al. 2006; Jiang and Fu 2016), vehicle routing (Verweij et al. 2003) and inventory management (Xu and Zhang 2009). Lastly, our comparison in this paper is on situations where one can apply IS, which is a common variance reduction technique in simulation via change-of-measure; general discussion can be found in Asmussen and Glynn (2007), Glasserman (2013).

The remainder of this paper is as follows. Section 2 compares SA and SAA when applying IS and presents our main asymptotic results. Section 3 discusses the implications in finite-time analysis. Section 4 presents two simulation examples. Section 5 concludes the paper.

2 ASYMPTOTIC COMPARISONS

Our discussion focuses on the stochastic optimization problem

$$\min_{\theta \in \Theta} f(\theta) = \mathbb{E}_{\theta}[F(\theta, \mathbf{X})], \quad (1)$$

where $\Theta \subset \mathfrak{R}$ is a feasible region, and $F(\theta, \mathbf{X})$ is the performance function. Note that we focus on a single-dimensional decision space here for ease of discussion. The expectation $\mathbb{E}_{\theta}[\cdot]$ is under a distribution $P(\theta, \cdot)$ that generates the random variable \mathbf{X} . Generally, we allow the decision variable θ to affect both the performance function and the distribution; the next subsections will focus on specific cases.

We assume that an unbiased gradient estimator is available (whose form varies in different cases). Moreover, we assume that there exists a unique optimal solution $\theta^* \in \Theta$ for (1) that is in the interior of Θ , f is twice continuously differentiable at θ^* , and the gradient estimator has finite variance. This allows us to obtain central limit theorems for the estimated solutions for both SA and SAA (Asmussen and Glynn

2007, Chapters 3 and 8), as we will describe in more detail. Furthermore, for reason that will become obvious, we assume $f''(\theta^*) = 1$ for simplicity.

In the subsequent discussion, we assume that the Monte Carlo simulation or data acquisition is the dominant cost and ignore the effort from the optimization routine in SAA. As will be explained, we will use the asymptotic variance from the central limit theorems as a measurement of the statistical performance.

2.1 θ in the Performance Function Only

We first look at the case where θ appears only in F , the performance function, but not in the distribution. So, we can use $P(\cdot)$ to denote the distribution of \mathbf{X} . Suppose that $F(\theta, \mathbf{X})$ is differentiable in θ , and $G(\theta, \mathbf{X})$ is the derivative. Suppose also that an interchange of derivative and expectation is allowed, so that $\mathbb{E}[G(\theta, \mathbf{X})] = f'(\theta)$.

To use SA, we will use the unbiased gradient estimator $G(\theta, \mathbf{X})$ and apply the R-M algorithm by generating a sequence of iterates $\{\theta_k\}$ via the recursion

$$\theta_{k+1} = \Pi_{\Theta}(\theta_k - a_k G(\theta_k, \mathbf{X}_k)), \quad (2)$$

where Π_{Θ} is the projection of θ into the feasible region Θ if $\theta \notin \Theta$, and a_k is an appropriate stepsize, e.g., a/k for some constant $a > 0$.

Under the assumptions above, we have a central limit theorem that $\sqrt{k}(\theta_k - \theta^*) \Rightarrow N(0, \sigma^2)$. When the stepsize a_k is chosen optimally, i.e., when it is set to a/k where $a = 1/f''(\theta^*)$, or when the Polyak-Ruppert averaging is used, then $\sigma^2 = \text{Var}(G(\theta^*, \mathbf{X}))/f''(\theta^*)^2 = \text{Var}(G(\theta^*, \mathbf{X}))$ (as we assume $f''(\theta^*) = 1$). This σ^2 cannot be improved using iteration of the form (2).

On the other hand, to use SAA, we generate i.i.d. copies $\mathbf{X}_1, \dots, \mathbf{X}_k$ from $P(\cdot)$ and solve

$$\min_{\theta \in \Theta} \left\{ f_k(\theta) = \frac{1}{k} \sum_{i=1}^k F(\theta, \mathbf{X}_i) \right\} \quad (3)$$

by applying a deterministic optimization routine. Suppose this routine generates an exact solution θ_k for the sample problem. Then, under the same assumptions as above, we have a central limit theorem that $\sqrt{k}(\theta_k - \theta^*) \Rightarrow N(0, \sigma^2)$, where $\sigma^2 = \text{Var}(G(\theta^*, \mathbf{X}))/f''(\theta^*)^2 = \text{Var}(G(\theta^*, \mathbf{X}))$. Note that this σ^2 is the same as the optimal σ^2 obtained in SA. In this sense, SAA always achieves the optimal statistical performance compared to all SA schemes in the form (2), and is thus viewed as the superior procedure statistically. This is particularly the case when the samples $\mathbf{X}_1, \dots, \mathbf{X}_k$ are given in advance, which occurs when they constitute a real data set for instance.

In simulation-based contexts, however, one can plausibly change the distribution $P(\cdot)$ in generating \mathbf{X} and apply IS in any part of the algorithms. Suppose $P(\cdot)$ lies in a class of distributions, denoted $\{P(\alpha, \cdot) : \alpha \in \Gamma\}$, and in particular say $P(\cdot) = P(\alpha_0, \cdot)$. For convenience, we abuse notations to use P also as the density or the mass function of \mathbf{X} . At each iteration k of SA, we can choose some $P(\alpha_k, \cdot)$ to generate the new \mathbf{X}_k . To retain unbiasedness, the gradient estimator at iteration k becomes

$$G(\theta_k, \mathbf{X}_k) \frac{P(\alpha_0, \mathbf{X}_k)}{P(\alpha_k, \mathbf{X}_k)}.$$

Suppose we can optimize the choice of α_k at each step. In other words, let $\alpha^*(\theta) \in \Gamma$ denote the optimal IS parameter given θ in terms of variance, i.e.,

$$\alpha^*(\theta) = \arg \min_{\alpha \in \Gamma} \text{Var}_{\alpha} \left(G(\theta, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha, \mathbf{X})} \right) \quad (4)$$

where $\text{Var}_{\alpha}(\cdot)$ denotes the variance under $P(\alpha, \cdot)$, and we pick $\alpha_k = \alpha^*(\theta_k)$ at iteration k to obtain an unbiased gradient

$$G(\theta_k, \mathbf{X}_k) \frac{P(\alpha_0, \mathbf{X}_k)}{P(\alpha^*(\theta_k), \mathbf{X}_k)}. \quad (5)$$

For convenience, we call this scheme SA-IS. A central limit theorem for θ_k obtained from SA-IS holds as before, with an asymptotic variance given by

$$\text{Var}_{\alpha^*(\theta^*)} \left(G(\theta^*, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha^*(\theta^*), \mathbf{X})} \right).$$

This is equivalent to

$$\min_{\alpha \in \Gamma} \text{Var}_{\alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha, \mathbf{X})} \right).$$

Since θ^* is generally unknown in advance, we take a worst-case view and use the minimax criterion that is in line with decision theory. In particular, the worst-case asymptotic variance for SA-IS among all possibilities of θ^* is

$$\max_{\theta^* \in \Theta} \min_{\alpha \in \Gamma} \text{Var}_{\alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha, \mathbf{X})} \right). \quad (6)$$

On the other hand, if we want to apply IS in SAA, we must apply it at the beginning when we form the sample problem (unless we expand the SAA method to allow for iterations). Say we use the distribution $P(\alpha_1, \cdot)$ to generate $\mathbf{X}_1, \dots, \mathbf{X}_k$. Then the SAA formulation becomes

$$\min_{\theta \in \Theta} \frac{1}{k} \sum_{i=1}^k F(\theta, \mathbf{X}_i) \frac{P(\alpha_0, \mathbf{X}_i)}{P(\alpha_1, \mathbf{X}_i)}. \quad (7)$$

A central limit theorem again holds for the obtained solution θ_k , with an asymptotic variance

$$\text{Var}_{\alpha_1} \left(G(\theta^*, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha_1, \mathbf{X})} \right).$$

Note that again we do not know the value θ^* . Suppose we want to optimize the worst-case asymptotic variance, we will find α_1 that solves

$$\min_{\alpha_1 \in \Gamma} \max_{\theta^* \in \Theta} \text{Var}_{\alpha_1} \left(G(\theta^*, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha_1, \mathbf{X})} \right). \quad (8)$$

Expression (8) is the best worst-case asymptotic variance for SAA, assuming α_1 can be picked suitably. For convenience, we call this scheme SAA-IS.

Comparing (6) and (8), the worst-case asymptotic variances of SAA-IS and SA-IS respectively, we see that (6) is at most (8), by using a standard weak duality (e.g., Boyd and Vandenberghe (2004)) that interchanges the max and the min. We summarize this as:

Theorem 1. *Suppose that the decision variable θ only appears in the performance function in (1), and all assumptions discussed above hold. SA-IS that uses (5) as the gradient, and SAA-IS that uses formulation (7), each has an asymptotic variance of the estimated optimal solution given by (6) and (8), respectively. The worst-case asymptotic variance for SA-IS is at most that of SAA-IS, that is,*

$$\max_{\theta^* \in \Theta} \min_{\alpha \in \Gamma} \text{Var}_{\alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha, \mathbf{X})} \right) \leq \min_{\alpha \in \Gamma} \max_{\theta^* \in \Theta} \text{Var}_{\alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\alpha_0, \mathbf{X})}{P(\alpha, \mathbf{X})} \right).$$

This result shows that SA can outperform SAA when the distribution can be changed. This presumes, however, that the parameter in SA can be tuned optimally, for example by Polyak-Ruppert averaging and by solving (4) at each θ_k . The result can be intuited by the adaptivity of IS at each step of SA, whereas IS can only be applied a priori in SAA. The adaptivity allows efficient digestion of information about θ^* , which in turn leads to a better worst-case performance of SA.

2.2 θ in the Distribution Only

In this subsection, we assume that the decision variable θ is only in the distribution of \mathbf{X} . Thus, we can write the performance function as $F(\mathbf{X})$. Suppose that we have an unbiased gradient estimator $G(\theta, \mathbf{X})$ obtained by, e.g., the likelihood ratio or the score function method (Glynn 1990; Reiman and Weiss 1989; Fu 2006), so that $G(\theta, \mathbf{X}) = F(\mathbf{X})s(\theta, \mathbf{X})$ for some score function $s(\theta, \mathbf{X})$.

We consider SA and the R-M algorithm. The standard way is to generate \mathbf{X}_k , at iteration k and given the current iterate θ_k , under the distribution $P(\theta_k, \cdot)$. The asymptotic variance of the estimated optimal solution is given by $\text{Var}_{\theta^*}(G(\theta^*, \mathbf{X}))$, and thus the worst-case asymptotic variance is

$$\max_{\theta^* \in \Theta} \text{Var}_{\theta^*}(G(\theta^*, \mathbf{X})). \quad (9)$$

For SAA, we have to first pick a particular choice of θ to generate $\mathbf{X}_1, \dots, \mathbf{X}_k$ to form the sample problem. This necessitates the use of IS. Suppose we pick θ_1 . This gives

$$\min_{\theta \in \Theta} \frac{1}{k} \sum_{i=1}^k F(\mathbf{X}_i) \frac{P(\theta, \mathbf{X}_i)}{P(\theta_1, \mathbf{X}_i)}, \quad (10)$$

and the asymptotic variance of the obtained optimal solution is

$$\text{Var}_{\theta_1} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \mathbf{X})}{P(\theta_1, \mathbf{X})} \right).$$

The best worst-case asymptotic variance is given by

$$\min_{\theta \in \Theta} \max_{\theta^* \in \Theta} \text{Var}_{\theta} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \mathbf{X})}{P(\theta, \mathbf{X})} \right). \quad (11)$$

Under this setting, it is unclear whether the worst-case asymptotic variance (9) or (11) is better. In other words, if the decision variable is inside the distribution, then SAA may no longer be statistically superior to SA (this reasoning precludes the use of more general IS). For a rough comparison, let us focus on the second moment, i.e., we compare

$$\max_{\theta^* \in \Theta} \mathbb{E}_{\theta^*} [G(\theta^*, \mathbf{X})^2] \quad (12)$$

and

$$\min_{\theta \in \Theta} \max_{\theta^* \in \Theta} \mathbb{E}_{\theta} \left[\left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \mathbf{X})}{P(\theta, \mathbf{X})} \right)^2 \right]. \quad (13)$$

We re-express (13) under the same probability measure as (12), and obtain

$$\min_{\theta \in \Theta} \max_{\theta^* \in \Theta} \mathbb{E}_{\theta^*} \left[G(\theta^*, \mathbf{X})^2 \frac{P(\theta^*, \mathbf{X})}{P(\theta, \mathbf{X})} \right]. \quad (14)$$

We approximate (14) by choosing a proper θ to make (12) and (14) comparable, as in the following proposition that is straightforward to verify:

Proposition 1. *Suppose that the decision variable θ only appears in the distribution in (1). If we choose the sampling distribution in SAA as $P(\theta_1, \mathbf{X})$, where θ_1 gives rise to the worst-case likelihood ratio*

$$\max_{\theta^* \in \Theta} \frac{P(\theta^*, \mathbf{X})}{P(\theta_1, \mathbf{X})} \leq K,$$

for some constant $K > 0$, then the asymptotic second moment of SAA is at most that of SA (without using IS), up to a constant factor K , i.e.,

$$\max_{\theta^* \in \Theta} \mathbb{E}_{\theta^*} \left[G(\theta^*, \mathbf{X})^2 \frac{P(\theta^*, \mathbf{X})}{P(\theta_1, \mathbf{X})} \right] \leq K \max_{\theta^* \in \Theta} \mathbb{E}_{\theta^*} [G(\theta^*, \mathbf{X})^2].$$

Next, we consider the use of IS in SA. To make the discussion more general, suppose $P(\theta, \cdot)$ lies in a bigger class of distributions, denoted $\{P(\theta, \alpha, \cdot) : \theta \in \Theta, \alpha \in \Gamma\}$, where in particular $P(\theta, \cdot) = P(\theta, \alpha_0, \cdot)$. Thus, at every iteration k , we can choose some $P(\tilde{\theta}_k, \tilde{\alpha}_k, \cdot)$ to generate the variable \mathbf{X}_k . To retain unbiasedness, the gradient estimator is

$$G(\theta_k, \mathbf{X}_k) \frac{P(\theta_k, \alpha_0, \mathbf{X}_k)}{P(\tilde{\theta}_k, \tilde{\alpha}_k, \mathbf{X}_k)}.$$

Given a θ , denote $(\tilde{\theta}^*(\theta), \tilde{\alpha}^*(\theta))$ as the minimizer of the variance of the associated unbiased gradient estimator, i.e.,

$$(\tilde{\theta}^*(\theta), \tilde{\alpha}^*(\theta)) = \arg \min_{\theta_1 \in \Theta, \alpha_1 \in \Gamma} \text{Var}_{\theta_1, \alpha_1} \left(G(\theta, \mathbf{X}) \frac{P(\theta, \alpha_0, \mathbf{X})}{P(\theta_1, \alpha_1, \mathbf{X})} \right)$$

where $\text{Var}_{\theta_1, \alpha_1}(\cdot)$ denotes the variance under $P(\theta_1, \alpha_1, \cdot)$. We choose $(\tilde{\theta}_k, \tilde{\alpha}_k) = (\tilde{\theta}^*(\theta_k), \tilde{\alpha}^*(\theta_k))$, so that the gradient estimator becomes

$$G(\theta_k, \mathbf{X}) \frac{P(\theta_k, \alpha_0, \mathbf{X}_k)}{P(\tilde{\theta}^*(\theta_k), \tilde{\alpha}^*(\theta_k), \mathbf{X}_k)}, \quad (15)$$

where \mathbf{X}_k is generated from $P(\tilde{\theta}^*(\theta_k), \tilde{\alpha}^*(\theta_k), \cdot)$. This gives an asymptotic variance

$$\text{Var}_{\tilde{\theta}^*(\theta^*), \tilde{\alpha}^*(\theta^*)} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \alpha_0, \mathbf{X})}{P(\tilde{\theta}^*(\theta^*), \tilde{\alpha}^*(\theta^*), \mathbf{X})} \right),$$

which is equivalent to

$$\min_{\theta \in \Theta, \alpha \in \Gamma} \text{Var}_{\theta, \alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \alpha_0, \mathbf{X})}{P(\theta, \alpha, \mathbf{X})} \right).$$

So the worst-case asymptotic variance is given by

$$\max_{\theta^* \in \Theta} \min_{\theta \in \Theta, \alpha \in \Gamma} \text{Var}_{\theta, \alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \alpha_0, \mathbf{X})}{P(\theta, \alpha, \mathbf{X})} \right). \quad (16)$$

On the other hand, the SAA problem within the same IS class has the form

$$\min_{\theta \in \Theta} \frac{1}{k} \sum_{i=1}^k F(\mathbf{X}_i) \frac{P(\theta, \alpha_0, \mathbf{X}_i)}{P(\theta_1, \alpha_1, \mathbf{X}_i)}, \quad (17)$$

where $\mathbf{X}_1, \dots, \mathbf{X}_k$ are generated from $P(\theta_1, \alpha_1, \cdot)$, and the best worst-case asymptotic variance is given by

$$\min_{\theta \in \Theta, \alpha \in \Gamma} \max_{\theta^* \in \Theta} \text{Var}_{\theta, \alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \alpha_0, \mathbf{X})}{P(\theta, \alpha, \mathbf{X})} \right). \quad (18)$$

Comparing (16) and (18), we see that, similar to Theorem 1, SA-IS is preferable. We summarize the result as follows:

Theorem 2. *Suppose that the decision variable θ only appears in the distribution in (1). SA-IS using a gradient estimator (15) and SAA-IS using formulation (17) have worst-case asymptotic variances of the respective estimated optimal solutions given by (16) and (18). The worst-case asymptotic variance of SA-IS is at most that of SAA-IS, i.e.,*

$$\max_{\theta^* \in \Theta} \min_{\theta \in \Theta, \alpha \in \Gamma} \text{Var}_{\theta, \alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \alpha_0, \mathbf{X})}{P(\theta, \alpha, \mathbf{X})} \right) \leq \min_{\theta \in \Theta, \alpha \in \Gamma} \max_{\theta^* \in \Theta} \text{Var}_{\theta, \alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \alpha_0, \mathbf{X})}{P(\theta, \alpha, \mathbf{X})} \right). \quad (19)$$

Like Theorem 1, the above conclusion assumes that the parameters in SA-IS can be optimally tuned. Note that even when the α parameter is absent, i.e., there is only the freedom to choose the decision variable θ , SA-IS is still better under this setting. This, in particular, indicates that the “standard way” to generate \mathbf{X}_k in SA discussed at the beginning of this subsection is suboptimal. This is summarized as:

Corollary 1. *Suppose that the decision variable θ only appears in the distribution in (1), and we are restricted to the class of distributions $\{P(\theta, \cdot) : \theta \in \Theta\}$ to generate \mathbf{X} . SA-IS using a gradient estimator*

$$G(\theta_k, \mathbf{X}) \frac{P(\theta_k, \mathbf{X}_k)}{P(\tilde{\theta}^*(\theta_k), \mathbf{X}_k)},$$

where

$$\tilde{\theta}^*(\theta) = \arg \min_{\theta_1 \in \Theta} \text{Var}_{\theta_1} \left(G(\theta, \mathbf{X}) \frac{P(\theta, \mathbf{X})}{P(\theta_1, \mathbf{X})} \right),$$

has a worst-case asymptotic variance of the estimated optimal solution given by

$$\max_{\theta^* \in \Theta} \min_{\theta \in \Theta} \text{Var}_{\theta, \alpha} \left(G(\theta^*, \mathbf{X}) \frac{P(\theta^*, \mathbf{X})}{P(\theta, \mathbf{X})} \right).$$

This quantity is at most the asymptotic variance of the standard SA that is given by (9), and also the worst-case asymptotic variance of SAA-IS formulated in (10) that is given by (11).

2.3 θ in Both the Performance and the Distribution

The above discussion largely holds when the decision variable θ is in both the performance function and the distribution. In this case, we assume an unbiased gradient estimator $G(\theta, \mathbf{X})$ can be derived via a combination of the pathwise differentiation and the likelihood ratio / score function method, i.e.,

$$G(\theta, \mathbf{X}) = \tilde{G}(\theta, \mathbf{X}) + F(\theta, \mathbf{X})s(\theta, \mathbf{X}), \tag{20}$$

where $\tilde{G}(\theta, \mathbf{X})$ is the pathwise differentiation of $F(\theta, \mathbf{X})$. We suppose that $P(\theta, \cdot)$ is in the class $\{P(\theta, \alpha, \cdot) : \theta \in \Theta, \alpha \in \Gamma\}$, where in particular $P(\theta, \cdot) = P(\theta, \alpha_0, \cdot)$, as in Section 2.2. Then, using the notations there, we can derive SA-IS with a gradient estimator the same as (15) and SAA-IS using formulation

$$\min_{\theta \in \Theta} \frac{1}{k} \sum_{i=1}^k F(\theta, \mathbf{X}_i) \frac{P(\theta, \alpha_0, \mathbf{X}_i)}{P(\theta_1, \alpha_1, \mathbf{X}_i)}. \tag{21}$$

We have the following result that is similar to Theorems 1 and 2:

Theorem 3. *Suppose that the parameter of interest θ appears in both the performance function and the distribution. SA-IS with gradient estimator given by (20) and SAA-IS formulation given by (21) have respective worst-case asymptotic variances of the estimated optimal solutions satisfying (16) and (18). Moreover, (19) holds.*

Remark 1. The estimator (20) can be extended to the case of discontinuous performance functions by using the generalized likelihood ratio method of Peng et al. (2018).

3 FINITE-TIME COMPARISON

This section discusses the finite-time performances of SAA-IS and SA-IS. Here we consider only the case that θ is in the performance function, and there exists an unbiased gradient estimator $G(\theta, \mathbf{X})$ obtainable from pathwise differentiation.

We consider the R-M algorithm (2) in the SA context. Specifically, let $a_k = a/k$, where a is a proper constant. We assume that $f(\theta)$ is strongly convex, i.e., there is a constant $c > 0$ such that

$$f(\theta) \geq f(\theta') + (\theta - \theta')f'(\theta') + \frac{1}{2}c(\theta - \theta')^2, \quad \forall \theta, \theta' \in \Theta.$$

Then, Nemirovski et al. (2009) gives the following upper bound for the optimality gap of the obtained solution θ_k in terms of L_2 -distance:

$$\mathbb{E} [(\theta_k - \theta^*)^2] \leq \frac{\max \{a^2 M^2 / (2ca - 1), (\theta_1 - \theta^*)^2\}}{k}, \quad (22)$$

where M is a positive number such that

$$\mathbb{E} [G(\theta, \mathbf{X})^2] \leq M^2, \quad \forall \theta \in \Theta. \quad (23)$$

For SAA given by (3), we assume that F is Lipschitz continuous, i.e., there exists a positive constant C such that

$$|F(\theta, \mathbf{X}) - F(\theta', \mathbf{X})| \leq C, \quad \text{for all } \theta, \theta' \in \Theta, \text{ and for all } \mathbf{X}.$$

Then, by Theorem 1 in Shapiro and Nemirovski (2005), the obtained SAA solution θ_k satisfies

$$\Pr \{|\theta_k - \theta^*| > \varepsilon\} \leq C_1 D \exp \left(-\frac{C_2 \varepsilon^2 k}{M^2 D^2} \right),$$

where M is as defined in (23), D is the diameter of Θ , and C_1 and C_2 are positive constants. Thus, a rather crude bound for the optimality gap measured by L_2 -distance can be obtained as

$$\begin{aligned} \mathbb{E} [(\theta_k - \theta^*)^2] &= \int_0^\infty \Pr \{(\theta_k - \theta^*)^2 > \delta\} d\delta \\ &\leq C_1 D \int_0^\infty \exp \left(-\frac{C_2 \delta k}{M^2 D^2} \right) d\delta \\ &= \frac{C_1 M^2 D^3}{C_2 k}. \end{aligned} \quad (24)$$

Note that the optimality gap bounds for both SA and SAA, namely (22) and (24), are both $O(k^{-1})$.

When we apply IS using the class of distributions $\{P(\alpha, \mathbf{X}) : \alpha \in \Gamma\}$, where $P(\cdot) = P(\alpha_0, \cdot)$, introduced in Section 2.1, we attempt to find an optimal IS parameter $\alpha^*(\theta)$ that satisfies (4) and use the gradient (5). Thus, in the above analysis, M can be replaced by a better constant $M_* \leq M$ such that

$$\mathbb{E}_{\alpha^*(\theta)} \left[\left(G(\theta, X) \frac{P(\alpha_0, X)}{P(\alpha^*(\theta), X)} \right)^2 \right] \leq M_*^2, \quad \forall \theta \in \Theta$$

or equivalently

$$\max_{\theta \in \Theta} \min_{\alpha \in \Gamma} \mathbb{E}_\alpha \left[\left(G(\theta, X) \frac{P(\alpha_0, X)}{P(\alpha, X)} \right)^2 \right] \leq M_*^2. \quad (25)$$

In SAA, applying IS can also change the constant M in the bound to a better constant $M^* \leq M$, but in the form

$$\min_{\alpha \in \Gamma} \max_{\theta \in \Theta} \mathbb{E}_\alpha \left[\left(G(\theta, X) \frac{P(\alpha_0, X)}{P(\alpha, X)} \right)^2 \right] \leq M^{*2} \quad (26)$$

since we must apply the IS and choose an α at the beginning. As in the asymptotic analysis in Section 2, the new constant for SA-IS in (25) is better than that for SA-SAA in (26).

4 SIMULATION EXAMPLES

In this section, we provide two simulation examples as illustration. In the first example, the decision variable θ is not in the distribution function. Let $X \sim \mathcal{N}(\mu, \sigma^2)$. Consider the optimization problem

$$\min_{\theta \in \Theta} \mathbb{E} \left[(X^2 + \theta - C)^2 \right], \quad (27)$$

where C is a constant. Specifically, let $\Theta = [-3, 1]$, $\mu = 1$, $\sigma = 1$, and $C = 1$. The true optimal solution is $\theta^* = C - \mathbb{E}[X^2] = -1$. We compare SA, SAA, SA-IS, and SAA-IS with different sample sizes. We use the root mean square error (RMSE) of the estimated optimal solution to measure the performance of each approach (which is a similar criterion as the variance asymptotically), where the RMSE is estimated from 50 replications. To compare SA and SAA methods fairly, we generate one sample at each iteration in SA and SA-IS, so the total number of Monte Carlo samples are the same in SA and SAA methods. We use the stepsize constant $a = 0.4$ in SA and SA-IS. In SA-IS, we choose the optimal $\tilde{\theta}^*(\theta_k)$ at each step by minimizing the corresponding variance via numerical integration and optimization. In SAA-IS, we uniformly discretize Θ and find the IS parameter that gives the smallest RMSE under the given sample size (note that this assumes knowledge of the optimal solution θ^* , which is impossible in practice and gives an edge to SAA, but this helps illustrate our conclusions). The results are shown in Table 1. We see that SAA performs better than SA across the settings. Better performances are detected generally when IS is applied to both methods (but for SAA-IS, the performances are sometimes slightly worse, possibly due to statistical noises in estimating the RMSE). Importantly, when IS is applied, SA-IS appears superior to both SAA and SAA-IS, which is consistent with our theoretical results in Section 2.1. In particular, the RMSE point estimates of SA-IS are generally lower than SAA-IS, though the differences are not statistically significant. However, as described above, we have assumed extra knowledge in implementing SAA-IS, and we foresee the differences will be more obvious when using SAA-IS without any a priori information.

Table 1: Comparing RMSE of SA, SA-IS, SAA, and SAA-IS in the first example, based on 50 simulation replications. The corresponding standard errors are provided in parentheses.

n	100	200	500	1000	2000	5000	10000	20000
SA	0.28 (0.05)	0.18 (0.03)	0.14 (0.02)	0.091 (0.02)	0.065 (0.01)	0.048 (0.007)	0.035 (0.004)	0.028 (0.004)
SA-IS	0.22 (0.02)	0.14 (0.01)	0.10 (0.009)	0.063 (0.007)	0.050 (0.005)	0.032 (0.003)	0.021 (0.002)	0.015 (0.001)
SAA	0.21 (0.03)	0.15 (0.02)	0.11 (0.01)	0.077 (0.01)	0.059 (0.009)	0.035 (0.006)	0.026 (0.004)	0.017 (0.003)
SAA-IS	0.21 (0.03)	0.16 (0.02)	0.11 (0.01)	0.065 (0.01)	0.051 (0.008)	0.032 (0.004)	0.024 (0.004)	0.016 (0.003)

Next, we consider the second example where the decision variable θ is in both the distribution function and the performance function. Let $X \sim \mathcal{N}((\theta - 1)^2, \sigma^2)$, and we consider the same objective function as in (27), but $\Theta = [-2, 2]$. We consider two cases: (i) $\sigma = 1$, $C = 1$; (ii) $\sigma = 2$, $C = -2$. The true optimal solutions are $\theta^* = 0.6085$ and $\theta^* = 0.5169$, respectively. In SA and SA-IS, we choose the stepsize constant $a = 0.5$ for Case (i) and $a = 0.7$ for Case (ii). As in the first example, we choose the optimal $\tilde{\theta}^*(\theta_k)$ at each step via numerical integration and optimization in SA-IS. For SAA, note that, as discussed in Section 2.2, it cannot be applied directly here and necessitates the use of IS at the beginning. For this, we consider two choices of the IS parameter. First is to pick 0 as an arbitrary choice (we label it as ‘‘SAA-IS0’’). Second, we uniformly discretize Θ and find the IS parameter that gives the smallest RMSE, like in the first example. The comparisons among SA, SA-IS, SAA-IS0, and SAA-IS are shown in Figure 1. We see that an arbitrary choice of IS parameter for SAA does not perform as well as SA; this can be by chance and

another choice may give better results, but there is no general guarantee. Applying IS on both methods generally improves performances (but, like in the first example, the performances of SAA-IS are worse than SAA-IS0 sometimes, which could be due to estimation noises for the RMSE). SA-IS appears superior to both SAA-IS0 and SAA-IS in this example, consistent with the findings in Sections 2.2 and 2.3.

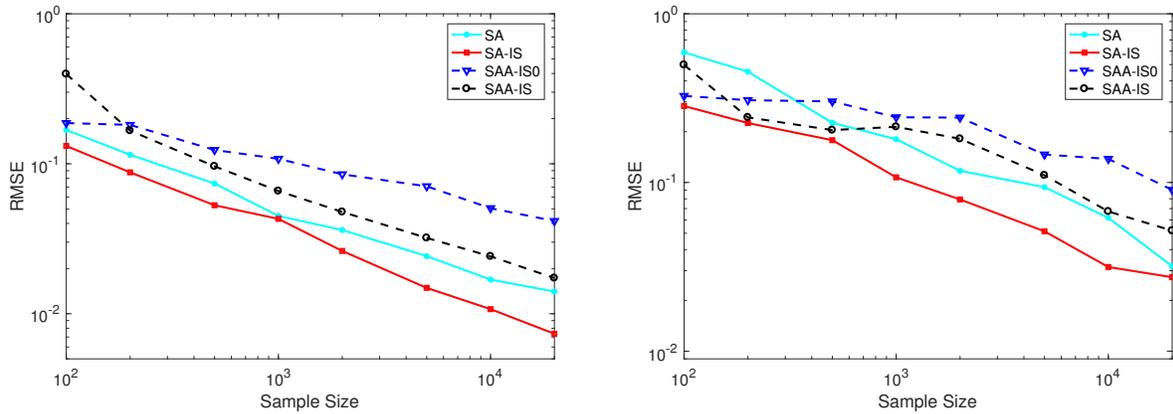


Figure 1: Comparing RMSE among SA, SA-IS, SAA-IS0, and SAA-IS in the second example. Left panel: $\sigma = 1, C = 1$. Right panel: $\sigma = 2, C = -2$.

Lastly, we give further details on how we find the IS parameters in SAA-IS. The procedure is similar for both examples, and here we discuss only the second one. We discretize $\Theta = [-2, 2]$ into 81 points $\theta = -2, -1.95, -1.90, \dots, 1.90, 1.95, 2$, and estimate the RMSE on each of these. Here we replicate 40 times to estimate the RMSE. From these values (shown in Figure 2), we select the IS parameters $\theta_b^* = 0.7$ for Case (i) and $\theta_b^* = 1.35$ for Case (ii), respectively, with the smallest RMSEs.

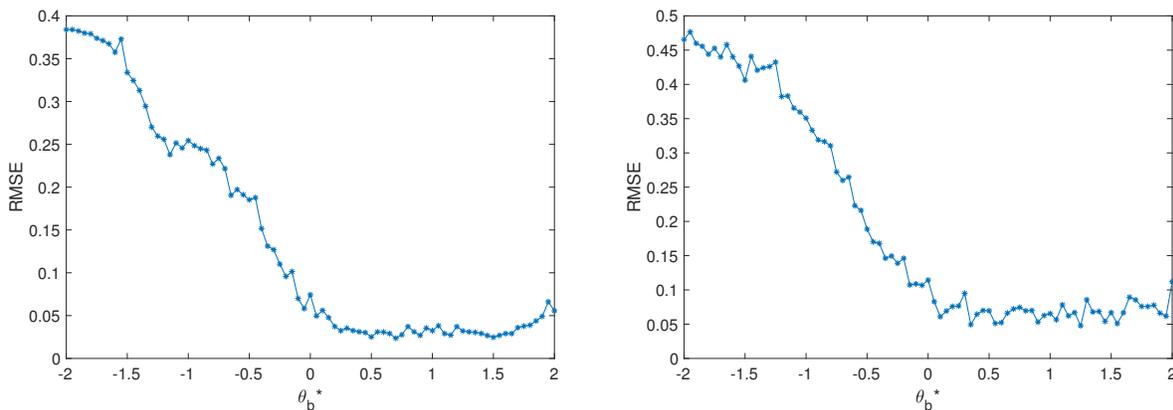


Figure 2: RMSE under different IS parameters in SAA with sample size 5000.

5 CONCLUSION AND DISCUSSION

We have compared the use of SA and SAA for stochastic optimization problems, particularly in the situations where the decision variable is in the distribution generating the stochasticity or when IS can be applied. In an idealized setting where the IS parameter can be chosen optimally, we show that SA performs better than SAA in terms of asymptotic variance. This modifies the conventional conclusion that SAA is statistically optimal when no IS can be applied. The derivation of our theoretical comparison uses a

worst-case framework over all possibilities of the unknown optimal solution, and a minimax weak duality argument. We further discuss similar observations in the finite-time analysis where the improvement in gradient estimation depicted in the performance bound in SA is bigger than in SAA. Finally, we present simulation examples to substantiate our theoretical findings. In future research, we will consider multi-dimensional decision variables, which are considerably more challenging, as the criterion for choosing the optimal IS is less clear. We will also investigate practical methods to tune IS schemes in suitable contexts that can show significant improvements in the efficiencies of the stochastic optimization algorithms, and apply the methods to realistic examples.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (NSF) under Grants CMMI-1434419, CMMI-1542020, CMMI-1523453, CAREER CMMI-1653339/1834710, and Shanghai Young Eastern Scholar Program [N.60-D129-18-202].

REFERENCES

- Alexander, S., T. F. Coleman, and Y. Li. 2006. “Minimizing CVaR and VaR for a Portfolio of Derivatives”. *Journal of Banking and Finance* 30(2):583–605.
- Asmussen, S., and P. W. Glynn. 2007. *Stochastic Simulation - Algorithms and Analysis*. New York: Springer.
- Bottou, L. 2010. “Large-Scale Machine Learning with Stochastic Gradient Descent”. In *Proceedings of COMPSTAT’2010*, edited by Y. Lechevallier and G. Saporta, 177–187. Berlin, Germany: Physica-Verlag.
- Boyd, S., and L. Vandenberghe. 2004. *Convex Optimization*. Cambridge, UK: Cambridge University Press.
- Broadie, M., D. Cicek, and A. Zeevi. 2011. “General Bounds and Finite-Time Improvement for the Kiefer-Wolfowitz Stochastic Approximation Algorithm”. *Operations Research* 59(5):1211–1224.
- Chau, M., and M. C. Fu. 2015. “An Overview of Stochastic Approximation”. In *Handbooks of Simulation Optimization*, edited by M. C. Fu, Chapter 6, 149–178. New York: Springer.
- Chau, M., and M. C. Fu. 2018. “Multivariate Stochastic Approximation Using a Secant-Tangents Averaged (STAR) Gradient”. *Operations Research* submitted.
- Fu, M. C. 2006. “Gradient Estimation”. In *Handbooks in Operations Research and Management Science*, edited by S. G. Henderson and B. L. Nelson, Volume 13, Chapter 19, 575–616. Amsterdam: Springer.
- Ghadimi, S., and G. Lan. 2013. “Optimal Stochastic Approximation Algorithms for Strongly Convex Stochastic Composite Optimization II: Shrinking Procedures and Optimal Algorithms”. *SIAM Journal on Optimization* 23(4):2061–2089.
- Ghadimi, S., and G. Lan. 2015. “Stochastic Approximation Methods and Their Finite-Time Convergence Properties”. In *Handbooks of Simulation Optimization*, edited by M. C. Fu, Chapter 7, 179–206. New York: Springer.
- Glasserman, P. 2013. *Monte Carlo Methods in Financial Engineering*. New York: Springer Science & Business Media.
- Glynn, P. W. 1990. “Likelihood Ratio Gradient Estimation for Stochastic Systems”. *Communications of the ACM* 33(10):75–84.
- Jiang, G., and M. C. Fu. 2016. “On Sample Average Approximation Algorithms for Determining the Optimal Importance Sampling Parameters in Pricing Financial Derivatives on Lévy Processes”. *Operations Research Letters* 44(1):44–49.
- Kiefer, K., and J. Wolfowitz. 1952. “Stochastic Estimation of the Maximum of a Regression Function”. *The Annals of Mathematical Statistics* 23(3):462–466.
- Kim, S., R. Pasupathy, and S. G. Henderson. 2015. “A Guide to Sample-Average Approximation”. In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 8, 207–241. New York: Springer.
- Kleywegt, A. J., A. Shapiro, and T. Homem-de Mello. 2002. “The Sample Average Approximation Method for Stochastic Discrete Optimization”. *SIAM Journal on Optimization* 12(2):479–502.

- Mak, W.-K., D. P. Morton, and R. K. Wood. 1999. "Monte Carlo Bounding Techniques for Determining Solution Quality in Stochastic Programs". *Operations Research Letters* 24(1-2):47–56.
- Nemirovski, A., A. Juditsky, G. Lan, and A. Shapiro. 2009. "Robust Stochastic Approximation Approach to Stochastic Programming". *SIAM Journal on Optimization* 19(4):1574–1609.
- Nesterov, Y., and V. Spokoiny. 2017. "Random Gradient-Free Minimization of Convex Functions". *Foundations of Computational Mathematics* 17(2):527–566.
- Peng, Y., M. C. Fu, J.-Q. Hu, and B. Heidergott. 2018. "A New Unbiased Stochastic Derivative Estimator for Discontinuous Sample Performances with Structural Parameters". *Operations Research* 66(2):487–499.
- Polyak, B. T., and A. B. Juditsky. 1992. "Acceleration of Stochastic Approximation by Averaging". *SIAM Journal on Control and Optimization* 30(4):838–855.
- Reiman, M. I., and A. Weiss. 1989. "Sensitivity Analysis for Simulations via Likelihood Ratios". *Operations Research* 37(5):830–844.
- Robbins, H., and S. Monro. 1951. "A Stochastic Approximation Method". *The Annals of Mathematical Statistics* 22(3):400–407.
- Shapiro, A. 2003. "Monte Carlo Sampling Methods". In *Handbooks in Operations Research and Management Science*, edited by A. Ruszczyński and A. Shapiro, Volume 10, Chapter 6, 353–425. Amsterdam: Springer.
- Shapiro, A., and A. Nemirovski. 2005. "On Complexity of Stochastic Programming Problems". In *Continuous Optimization: Current Trends and Application*, edited by V. Jeyakumar and A. M. Rubinov, 111–146. New York: Springer.
- Spall, J. C. 1992. "Multivariate Stochastic Approximation Using a Simultaneous Perturbation Gradient Approximation". *IEEE Transactions on Automatic Control* 37(3):332–341.
- Verweij, B., S. Ahmed, A. J. Kleywegt, G. Nemhauser, and A. Shapiro. 2003. "The Sample Average Approximation Method Applied to Stochastic Routing Problems: A Computational Study". *Computational Optimization and Applications*. 24(2-3):289–333.
- Xu, H., and D. Zhang. 2009. "Smooth Sample Average Approximation of Stationary Points in Nonsmooth Stochastic Optimization and Applications". *Mathematical Programming* 119(2):371–401.

AUTHOR BIOGRAPHIES

HENRY LAM is an Associate Professor in the Department of Industrial Engineering and Operations Research at Columbia University. He received his Ph.D. degree in statistics at Harvard University in 2011, and was on the faculty of Boston University and the University of Michigan before joining Columbia in 2017. His research focuses on Monte Carlo simulation, risk and uncertainty quantification, and stochastic optimization. His email address is khl2114@columbia.edu.

GUANGXIN JIANG is an Assistant Professor in the School of Management, at Shanghai University. He received his Ph.D. degree in applied mathematics at Tongji University, and was a postdoctoral fellow in the Department of Economics and Finance, at City University of Hong Kong. His research interests lie in simulation methodology, analytics, machine learning, and optimization. His email address is gxjiang@shu.edu.cn.

MICHAEL C. FU holds the Smith Chair of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research, A. James Clark School of Engineering, University of Maryland, College Park. He has a Ph.D. in applied math from Harvard and degrees in math and EECS from MIT. He served as WSC2011 Program Chair, NSF Operations Research Program Director, *Management Science* Stochastic Models and Simulation Department Editor, and *Operations Research* Simulation Area Editor. He is a Fellow of **INFORMS** and **IEEE**. His email address is mfu@umd.edu.