Stochastic Reweighted Gradient Descent

Ayoub El Hanchi [1] **David A. Stephens** [2] **Chris J. Maddison** [1]

Abstract

Importance sampling is a promising strategy for improving the convergence rate of stochastic gradient methods. It is typically used to precondition the optimization problem, but it can also be used to reduce the variance of the gradient estimator. Unfortunately, this latter point of view has yet to lead to practical methods that provably improve the asymptotic error of stochastic gradient methods. In this work, we propose stochastic reweighted gradient descent (SRG), a stochastic gradient method based solely on importance sampling that can reduce the variance of the gradient estimator and improve on the asymptotic error of stochastic gradient descent (SGD) in the strongly convex and smooth case. We show that SRG can be extended to combine the benefits of both importance-sampling-based preconditioning and variance reduction. When compared to SGD, the resulting algorithm can simultaneously reduce the condition number and the asymptotic error, both by up to a factor equal to the number of component functions. We demonstrate improved convergence in practice on regularized logistic regression problems.

1. Introduction

Unconstrained optimization of finite-sum objectives is a core algorithmic problem in machine learning. The prototypical way of solving such problems is by viewing them through the lens of stochastic optimization, where the source of stochasticity resides in the choice of the index in the sum. Stochastic gradient descent (SGD) (Robbins & Monro, 1951) remains the standard algorithm for this class of problems.

A natural way to improve on SGD is by considering importance sampling schemes. This idea is not new and dates back to (Needell et al., 2014) who uses importance sampling as a preconditioning technique. They propose sampling the indices with probabilities proportional to the smoothness constants of the corresponding component functions, and show that this sampling scheme provably reduces the condition number of the problem.

In another line of work, variance-reduced methods were found to achieve linear convergence in the strongly convex and smooth case (Roux et al., 2012; Schmidt et al., 2017). Many of these methods rely on control variates to reduce the variance of the gradient estimator used by SGD (Johnson & Zhang, 2013; Defazio et al., 2014). While very successful, the applicability of these methods is limited by the large memory overhead that they introduce, or the periodic full-gradient recomputation that they require (Defazio & Bottou, 2019). Despite strong progress in this research area, an importance-sampling-based analogue to these algorithms, which is potentially free from these drawbacks, has yet to emerge.

In this work, we propose such an analogue. We introduce stochastic reweighted gradient (SRG), an importance-sampling-based stochastic optimization algorithm that reduces the variance of the gradient estimator. Similar to SGD, SRG requires a single gradient oracle call per iteration, and only requires $O(n)$ additional memory, and $O(\log n)$ additional floating point operations per iteration. We analyze the convergence rate of SRG in the strongly-convex and smooth case, and show that it can provably improve the asymptotic error of SGD. Finally, we show how our importance sampling strategy can be combined with smoothness-based importance sampling, and prove that the resulting algorithm simultaneously performs variance reduction and preconditioning. We demonstrate improved convergence in practice on ℓ_2-regularized logistic regression problems.

2. Preliminaries

We consider the finite-sum optimization problem:

$$\min_{x \in \mathbb{R}^d} \left\{ F(x) := \frac{1}{n} \sum_{i=1}^{n} f_i(x) \right\} \tag{1}$$

where F is μ-strongly convex for $\mu > 0$, and for all $i \in [n]$, f_i is convex and L_i-smooth for $L_i > 0$. Note that

[1]University of Toronto and Vector Institute [2]McGill University. Correspondence to: Ayoub El Hanchi <aelhan@cs.toronto.edu>.

Proceedings of the 39^{th} International Conference on Machine Learning, Baltimore, Maryland, USA, PMLR 162, 2022.

by strong-convexity, F has a unique minimizer $x^* \in \mathbb{R}^d$. We define the maximum $L_{\max} := \max_{i \in [n]} L_i$ and average $\overline{L} := n^{-1} \sum_{j=1}^{n} L_i$ smoothness constants. Similarly, we define the maximum $\kappa_{\max} := L_{\max}/\mu$ and average $\overline{\kappa} := \overline{L}/\mu$ condition numbers.

The classical way of solving (1) is by viewing it as a stochastic optimization problem where the randomness comes from the choice of the index $i \in [n]$. Starting from some $x_0 \in \mathbb{R}^d$, and for an iteration number $k \in \mathbb{N}$, stochastic gradient descent (SGD) performs the following update:

$$x_{k+1} = x_k - \alpha_k \nabla f_{i_k}(x_k)$$

for a step size $\alpha_k > 0$ and a random index i_k drawn uniformly from $[n]$.

The idea behind importance sampling for SGD is to instead sample the index i_k according to a chosen distribution p_k on $[n]$, and to perform the update (Needell et al., 2014):

$$x_{k+1} = x_k - \alpha_k \frac{1}{n p_k^{i_k}} \nabla f_{i_k}(x_k) \tag{2}$$

where p_k^i is the i^{th} component of the probability vector p_k. It is immediate to verify that the importance sampling estimator of the gradient is unbiased as long as $p_k > 0$.

The question that we address in this paper is how to design a sequence $\{p_k\}_{k=0}^{\infty}$ that produces more efficient gradient estimators than the ones produced by uniform sampling. One way to design such a sequence is by adopting a greedy strategy: at each iteration k we choose p_k to minimize the conditional variance of the gradient estimator, which is given by, up to an additive constant:

$$\sigma^2(x_k, p) = \frac{1}{n^2} \sum_{i=1}^{n} \frac{1}{p^i} \left\| \nabla f_i(x_k) \right\|_2^2, \tag{3}$$

This conditional variance is minimized at (Zhao & Zhang, 2015):

$$\arg\min_p \sigma^2(x_k, p) = \left(\frac{\left\| \nabla f_i(x_k) \right\|_2}{\sum_{j=1}^{n} \left\| \nabla f_j(x_k) \right\|_2} \right)_{i=1}^{n} \tag{4}$$

Ideally, we would like to set p_k to this minimizer. However, this requires knowledge of the gradient norms $(\left\| \nabla f_i(x_k) \right\|_2)_{i=1}^{n}$, which, in general, requires n gradient evaluations per iteration, making this approach intractable.

3. Algorithm

In this section, we show how to design a tractable sequence of importance sampling distributions for SGD that approximate the conditional-variance-minimizing distributions (4). First we construct efficient approximations of the conditional variances (3). We then state a simple bound on their approximation errors and use it to motivate our choice of importance sampling distributions.

Algorithm 1 SRG

Parameters: step sizes $(\alpha_k)_{k=0}^{\infty} > 0$, mixture coefficients $(\theta_k)_{k=0}^{\infty} \in (0, 1]$
Initialization: $x_0 \in \mathbb{R}^d$, $(\left\| g_0^i \right\|_2)_{i=1}^{n} \in \mathbb{R}^n$
for $k = 0, 1, 2, \ldots$ **do**
  $p_k = (1 - \theta_k) q_k + \theta_k / n$ $\{q_k$ is defined in (6)$\}$
  $b_k \sim \text{Bernoulli}(\theta_k)$
  if $b_k = 1$ **then** $i_k \sim 1/n$ **else** $i_k \sim q_k$
  $x_{k+1} = x_k - \alpha_k \frac{1}{n p_k^{i_k}} \nabla f_{i_k}(x_k)$
  $\left\| g_{k+1}^i \right\|_2 = \begin{cases} \left\| \nabla f_i(x_k) \right\|_2 & \text{if } b_k = 1 \text{ and } i_k = i \\ \left\| g_k^i \right\|_2 & \text{otherwise} \end{cases}$
end for

To approximate the conditional variances, we follow the strategy of certain variance-reduced methods (Roux et al., 2012; Schmidt et al., 2017; Defazio et al., 2014). These methods maintain a table $(g_k^i)_{i=1}^{n}$ that tracks the component gradients $(\nabla f_i(x_k))_{i=1}^{n}$ which is updated at each iteration at the index i_k used to update the iterates x_k (2). Our method instead maintains an array of gradient *norms*, from which we construct an approximation of the conditional variance (3) of the gradient estimator:

$$\tilde{\sigma}^2(x_k, p) := \frac{1}{n^2} \sum_{i=1}^{n} \frac{1}{p^i} \left\| g_k^i \right\|_2^2 \tag{5}$$

this approximation is minimized at:

$$q_k = \arg\min_p \tilde{\sigma}^2(x_k, p) = \left(\frac{\left\| g_k^i \right\|_2}{\sum_{j=1}^{n} \left\| g_k^j \right\|_2} \right)_{i=1}^{n} \tag{6}$$

We do not directly use q_k as an importance sampling distribution, because this approximation may be poor. In particular, we have the following bound that relates the conditional variance (3) to the approximation (5):

$$\sigma^2(x_k, p) \leq \frac{2}{n^2} \sum_{i=1}^{n} \frac{\left\| \nabla f_i(x_k) - g_k^i \right\|_2^2}{p^i} + 2\tilde{\sigma}^2(x_k, p) \tag{7}$$

Recall that our goal is to pick p_k that minimizes $\sigma^2(x_k, p)$. q_k minimizes the second term on the right-hand side of 7, but we must ensure that both terms are small. Two conditions are needed to keep the first term small. The first is to control the terms $\left\| \nabla f_i(x_k) - g_k^i \right\|_2^2$, which we can do by making sure that the historical gradients g_k^i are frequently updated. The second is to ensure that the probabilities p_k^i

are lower bounded. We achieve both of these properties, as well as approximately minimize the right-hand side of (7) by mixing q_k with the uniform distribution over $[n]$. For a given mixture coefficient $\theta_k \in (0, 1]$, this yields our final importance sampling distribution:

$$p_k = (1 - \theta_k) q_k + \frac{\theta_k}{n} \tag{8}$$

Our analysis in section 4 clarifies the role of the sequence of mixture coefficients $(\theta_k)_{k=0}^{\infty}$, and relates it to both the step size sequence $(\alpha_k)_{k=0}^{\infty}$ and the asymptotic error of SRG in the constant step size setting.

Curiously, our analysis requires performing the update of the array $(\left\|g_k^i\right\|_2)_{i=1}^n$ only when the index i_k is drawn from the uniform mixture component. It is not clear to us whether this constraint is an artifact of the analysis or a property of the algorithm. We discuss this further after the statement of Lemma 4.1 in section 4. Pseudocode for our method is given in Algorithm 1.

We briefly discuss how SRG can be efficiently implemented. To sample from q_k, we store $(\left\|g_k^i\right\|_2)_{i=1}^n$ in a binary indexed tree using $O(n)$ memory, along with the normalizing constant $\lambda_k = \sum_{i=1}^n \left\|g_k^i\right\|_2$. The binary indexed tree can be updated and maintained in $O(\log n)$ operations, while the normalizing constant can be maintained in constant time. We can then sample from q_k using inverse transform sampling: we multiply a uniform random variable $u \in [0, 1)$ by λ_k, then return the largest index i_k whose corresponding prefix sum $\sum_{i=1}^{i_k} \left\|g_k^i\right\|_2$ is less than $\lambda_k u$. This procedure also requires only $O(\log n)$ operations. The total overhead of SRG when compared to SGD is therefore $O(n)$ memory and $O(\log n)$ floating point operations per iteration, both of which are dimension-free.

4. Theory

In this section, we analyze the convergence rate of SRG, and show that it can achieve a better asymptotic error than SGD. Two key constants are helpful in contrasting the asymptotic errors of SRG and SGD. Recall the definition of the conditional variance $\sigma^2(x_k, p_k)$ in (3), and define:

$$\sigma^2 := \sigma^2(x^*, 1/n) = \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(x^*)\|_2^2$$

$$\sigma_*^2 := \min_p \sigma^2(x^*, p) = \frac{1}{n^2} \left(\sum_{i=1}^n \|\nabla f_i(x^*)\|_2 \right)^2$$

It is well known that the asymptotic error of SGD in the strongly-convex and smooth case depends linearly on σ^2 (Needell et al., 2014). We here show that SRG reduces this to a linear dependence on σ_*^2, which can be up to n times smaller.

To study the convergence rate of SRG, we use the following Lyapunov function, which is similar to the one used to study the convergence rate of SAGA (Hofmann et al., 2015; Defazio, 2016):

$$T^k := \frac{\alpha_k}{\theta_k} \frac{a}{L_{\max}} \sum_{i=1}^n \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + \|x_k - x^*\|_2^2 \tag{9}$$

for a constant $a > 0$ that we set during the analysis. The proofs of this section can be found in Appendix B.

4.1. Intermediate lemmas

Before proceeding with the main result, let us first state two intermediate lemmas. The first studies the evolution of $(g_k^i)_{i=1}^n$ from one iteration to the next.

Lemma 4.1. *Let $k \in \mathbb{N}$ and suppose that $(g_k^i)_{i=1}^n$ evolves as in Algorithm 1. Taking expectation with respect to (b_k, i_k), conditional on $(b_t, i_t)_{t=0}^{k-1}$, we have:*

$$\mathbb{E}\left[\sum_{i=1}^n \left\|g_{k+1}^i - \nabla f_i(x^*)\right\|_2^2\right] \leq 2\theta_k L_{max} \left[F(x_k) - F(x^*)\right] + \left(1 - \frac{\theta_k}{n}\right) \sum_{i=1}^n \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2$$

The use of the Bernoulli random variable b_k in Algorithm 1 to monitor the update of $(g_k^i)_{i=1}^n$ is necessary for Lemma 4.1 to hold. In particular, without the use of b_k, the elements of $(g_k^i)_{i=1}^n$ may have different probabilities of being updated. In that case, the first term of Lemma 4.1 becomes a weighted average of terms which is not easily bounded. When the importance sampling distribution does not depend on the iteration number, this issue can be fixed with a slight modification of the Lyapunov function (9) (Schmidt et al., 2015). In our case however, we are dealing with time-varying importance sampling distributions, and this approach fails. Instead, we condition the update of $(g_k^i)_{i=1}^n$ on the Bernoulli random variable b_k such that the probability of updating g_k^i is fixed to θ_k/n for all $i \in [n]$.

The second lemma is a bound on the conditional variance of the gradient estimator used by SRG. We intentionally leave as many free parameters as possible in the bound and optimize over them in the main result of section 4.2.

Lemma 4.2. *Let $k \in \mathbb{N}$ and assume that $\theta_k \in (0, 1/2]$. Taking expectation with respect to (b_k, i_k), conditional on $(b_t, i_t)_{t=0}^{k-1}$, we have, for all $\beta, \gamma, \delta, \eta > 0$:*

$$\mathbb{E}_{i_k \sim p_k}\left[\left\|\frac{1}{np_k^{i_k}} \nabla f_{i_k}(x_k)\right\|_2^2\right] \leq \frac{2D_1 L_{max}}{\theta_k} \left[F(x_k) - F^*\right] + \frac{D_2}{\theta_k n} \sum_{i=1}^n \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + D_3(1 + 2\theta_k)\sigma_*^2$$

where D_1, D_2 and D_3 are given by:

$$D_1 := (1 + \beta + \gamma)$$
$$D_2 := (1 + \beta^{-1} + \delta) + (1 + \gamma^{-1} + \delta^{-1})(1 + \eta)$$
$$D_3 := (1 + \gamma^{-1} + \delta^{-1})(1 + \eta^{-1})$$

4.2. Main result

Our main result is a bound on the evolution of the Lyapunov function T^k along the steps of SRG.

Theorem 4.3. *Suppose that $(x_k, (g_k^i)_{i=1}^n)$ evolves according to Algorithm 1. Further, assume that for all $k \in \mathbb{N}$: (i) α_k/θ_k is non-increasing. (ii) $\theta_k \in (0, 1/2]$. (iii) $\alpha_k \leq \theta_k/12L_{max}$. Then:*

$$\mathbb{E}\left[T^{k+1}\right] \leq (1 - \rho_k)\mathbb{E}\left[T^k\right] + (1 + 2\theta_k)6\alpha_k^2\sigma_*^2$$

for all $k \in \mathbb{N}$, and where:

$$\rho_k := \min\left\{\frac{\theta_k}{12n}, \alpha_k\mu\right\}$$

The constants ($1/12$ in the bound on the step size and in ρ_k, 6 in front of the σ_*^2 term) in this theorem are optimized under the following constraints. First, the parameterized form of the above bound shows that the largest allowable step size is given by $c_2\theta_k/L_{\text{max}}$. Using it we get:

$$\rho_k = \min\left\{c_1\frac{\theta_k}{n}, c_2\frac{\theta_k}{\kappa_{\text{max}}}\right\}$$

for some constants $c_1, c_2 > 0$. As we do not know a priori the relative magnitudes of n and κ_{max}, and since c_1 and c_2 are inversely proportional, we impose the constraint $c_1 = c_2$. Similarly, our parameterized bound gives an asymptotic error of the form $(1+2\theta_k)c_3\alpha_k^2\sigma_*^2$ for a constant c_3. We chose to impose the constraint $c_3 = 6$. c_1, c_2 and c_3 are all functions of the free parameters of Lemma 4.2 and the constant a of the Lyapunov function (9). Numerically maximizing c_2 (and therefore the largest allowable step size) subject to these two constraints ($c_3 = 6$ and $c_1 = c_2$) with respect to these parameters yields the result in Theorem 4.3.

To give the reader an idea of the sensitivity of the result to the choice of c_3, note that setting $c_3 = 2$ yields $c_2 \approx 1/20$, whereas taking $c_3 \gg 1$ yields $c_2 \approx 1/10$. We have attempted to obtain the best bound possible on the largest step size allowable, but the rather small prefactor $c_2 = 1/12$ seems to be an inevitable consequence of the multiple (but as far as we can tell necessary) uses of Young's inequality in our analysis. Our experiments suggest that the dependence of the largest step size allowable on the mixture coefficient θ is real, but that the prefactor $c_2 = 1/12$ may be an artifact of the analysis.

For SRG with constant mixture coefficient and step size, its convergence rate and complexity are given by the following corollary of Theorem 4.3:

Corollary 4.4. *Suppose that $(x_k, (g_k^i)_{i=1}^n)$ evolves according to Algorithm 1 with a constant mixture coefficient $\theta_k = \theta \in (0, 1/2]$ and a constant step size $\alpha_k = \alpha \leq \theta/12L_{max}$. Then for any $k \in \mathbb{N}$:*

$$\mathbb{E}\left[T^k\right] \leq (1 - \rho)^k T^0 + (1 + 2\theta)\frac{6\alpha^2\sigma_*^2}{\rho}$$

where $\rho = \rho_k$ is as defined in Theorem 4.3. For any $\varepsilon > 0$ and $\theta \in (0, 1/2]$, choosing:

$$\alpha = \min\left\{\frac{\theta}{12L_{max}}, \frac{\varepsilon\mu}{(1+2\theta)12\sigma_*^2}, \sqrt{\frac{\theta}{1+2\theta}\frac{\varepsilon}{144n\sigma_*^2}}\right\}$$

and:

$$k \geq \max\left\{\frac{12n}{\theta}, \frac{1}{\alpha\mu}\right\}\log\left(\frac{2T^0}{\varepsilon}\right)$$

guarantees $\mathbb{E}\left[\|x_k - x^*\|_2^2\right] \leq \varepsilon$

Comparing the convergence rate of SRG in Corollary 4.4 with the standard result for SGD (Needell et al., 2014), we see that they are of similar form. When $\rho = \alpha\mu$, the bound of Corollary 4.4 is better asymptotically. Indeed, in this case, and as $k \to \infty$, the iterates of SRG stay within a ball of radius $O(\sqrt{\sigma_*^2})$ of the minimizer, while those of SGD stay withing a ball of radius $O(\sqrt{\sigma^2})$. The equality $\rho = \alpha\mu$ holds when $\alpha \leq 1/n\mu$, which is true for all allowable step sizes if the problem is dominated by its maximum condition number $\kappa_{\text{max}} \geq n$, and for smaller step sizes otherwise.

In terms of complexity, we have the following comparison. Up to constants, the complexity of SRG with a constant mixture coefficient and step size is of the form:

$$O\left(n + \sqrt{\frac{n\sigma_*^2}{\mu^2\varepsilon}} + \kappa_{\text{max}} + \frac{\sigma_*^2}{\mu^2\varepsilon}\right)\log\left(\frac{1}{\varepsilon}\right)$$

We compare this to the complexity of SGD with constant step size (Needell et al., 2014):

$$O\left(\kappa_{\text{max}} + \frac{\sigma^2}{\mu^2\varepsilon}\right)\log\left(\frac{1}{\varepsilon}\right)$$

In the high accuracy regime, the ε^{-1} terms dominate the complexities of SRG and SGD. In this case, SRG enjoys a better complexity than SGD since $\sigma_*^2 \leq \sigma^2$.

Algorithm 2 SRG+

Parameters: step sizes $(\alpha_k)_{k=0}^{\infty} > 0$, mixture coefficients $(\theta_k)_{k=0}^{\infty} \in (0, 1]$
Initialization: $x_0 \in \mathbb{R}^d$, $(\left\|g_0^i\right\|_2)_{i=1}^n \in \mathbb{R}^n$
for $k = 0, 1, 2, \ldots$ **do**
  $p_k = (1 - \theta_k)q_k + \theta_k v$
  $\{q_k$ is given by (6), v is given by (12)$\}$
  $b_k \sim \text{Bernoulli}(\theta_k)$
  if $b_k = 1$ **then** $(i_k, j_k) \sim \pi$ **else** $i_k \sim q_k$
  $\{\pi$ maximally couples $(v, 1/n)\}$
  $x_{k+1} = x_k - \alpha_k \frac{1}{np_k^{i_k}} \nabla f_{i_k}(x_k)$
  $\left\|g_{k+1}^j\right\|_2 = \begin{cases} \|\nabla f_j(x_k)\|_2 & \text{if } b_k = 1 \text{ and } j = j_k \\ \left\|g_k^j\right\|_2 & \text{otherwise} \end{cases}$
end for

5. Extension

In this section, we extend SRG to combine its variance reduction capacity with the preconditioning ability of smoothness-based importance sampling.

A straightforward way to generalize the argument given in the derivation of SRG in section 3 is to consider the following bound on the conditional variance (3) of the gradient estimator, which can be derived from Young's inequality and the L_i-smoothness of each f_i:

$$\sigma^2(x_k, p) \leq \frac{3}{n^2} \sum_{i=1}^{n} \frac{L_i}{p^i} \langle \nabla f_i(x_k) - \nabla f_i(x^*), x_k - x^* \rangle + \frac{3}{n^2} \sum_{i=1}^{n} \frac{1}{p^i} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + 3\tilde{\sigma}^2(x_k, p) \quad (10)$$

While the motivating bound (7) seems more intuitive, because we think of g_k^i as tracking $\nabla f_i(x_k)$, it turns out that this second bound better captures the evolution of g_k^i in relation to $\nabla f_i(x_k)$. At a high-level, this is because g_k^i tracks $\nabla f_i(x_k)$ indirectly: both hover around $\nabla f_i(x^*)$ as k gets large.

Similar to our approach in section 3, our goal is to pick p_k that minimizes the right-hand side of (10). We know that q_k (6) minimizes the third term, but we need to make sure that the first two are also small. To minimize the first term, knowing nothing about the relative sizes of the inner product terms, it makes sense to have probabilities proportional to the smoothness constants [1]. On the other hand, to keep the second term small, we need to ensure that the historical gradients $(g_k^i)_{i=1}^n$ are frequently updated, which we can do by imposing a uniform lower bound on the probabilities. These considerations motivate us to consider the following distributions:

$$p_k' = (1 - \eta_k - \theta_k)q_k + \eta_k v + \frac{\theta_k}{n} \quad (11)$$

for positive mixture coefficients (θ_k, η_k) satisfying $\theta_k + \eta_k \in (0, 1]$, and where v is given by:

$$v = \left(\frac{L_i}{n\overline{L}}\right)_{i=1}^n \quad (12)$$

Using these probabilities, we are able to show that the resulting algorithm does indeed achieve both variance reduction and preconditioning. However, in the worst case, its complexity is twice as much as what we would expect from simply replacing $L_{\max}$ with $\overline{L}$ in Corollary 4.4. Intuitively, this is because the probability assigned to the uniform component in (8) needs to be split between the uniform and the smoothness-based components in (11).

5.1. Carefully decoupling the updates of $(g_k^i)_{i=1}^n$ and x_k

Here we show how to design our algorithm such that this additional factor of 2 (described above) in the complexity is reduced to:

$$1 + \|v - 1/n\|_{TV} \leq 2 - 1/n$$

which is usually much closer to 1 in practical settings where the smoothness constants are roughly on the same scale. Following (Schmidt et al., 2015), our method is based on the observation that we can decouple the index used to update the historical gradients and the index used to update the iterates, which we refer to as j_k and i_k, respectively. Intuitively, to minimize (10) we would ideally want j_k to be uniformly distributed and i_k to be distributed according to $p_k = (1 - \theta_k)q_k + \theta_k v$. Unfortunately, sampling (j_k, i_k) independently with these marginals does not address our issue, because we would still require an average of approximately two gradient evaluations per iteration.

Our main observation is that we can use any coupling between (j_k, i_k), because the evolution of the expectation of the Lyapunov function (9) only depends on their marginal distributions and not on their joint. In particular, we obtain the same bound on the evolution of $\mathbb{E}\left[T^k\right]$ regardless of how i_k and j_k are coupled. It therefore makes sense to pick the coupling that maximizes the probability that $i_k = j_k$, since this minimizes the number of gradient evaluations required per iteration. Such couplings are known as maximal couplings in the literature, and can easily be computed for discrete random variables (see, e.g., Algorithm 5 in Biswas et al. (2019)). With a maximal coupling,

[1] Based on this argument alone, one would want them to be proportional to the square root of the smoothness constants. This however does not lead to a nice averaging of the inner product terms, which is important for reasons related to the strong-convexity of F but not of the component functions f_i.

the expected number of gradient evaluations per iteration becomes $1 + \|v - 1/n\|_{TV}$. Using one such coupling leads to SRG+ described in Algorithm 2.

Let us briefly discuss the implementation of SRG+. Sampling from π, a maximal coupling of v and the uniform distribution, requires forming three probability vectors (see Algorithm 5 in Biswas et al. (2019) for details). As both v and $1/n$ are constant throughout the optimization process, we can form these vectors at the beginning of the algorithm along with their partial sums for a total initial cost of $O(n)$ operations. We can then sample from π in $O(\log n)$ time using binary search on the partial sums at each iteration. The overhead of SRG+ is therefore the same as that of SRG.

5.2. Analysis

The analysis of SRG+ is similar to that of SRG. In particular, the iterates of SRG+ also obey a slight modification of Theorem 4.3 where the bound on the largest allowable step size is loosened to $\alpha_k \leq \theta_k / 12\overline{L}$. We refer the reader to Appendix C for more details. Due to this improvement, we get that the complexity of SRG+ with a constant mixture coefficient and step size is given by:

$$O\left(n + \sqrt{\frac{n\sigma_*^2}{\mu\varepsilon}} + \overline{\kappa} + \frac{\sigma_*^2}{\mu^2\varepsilon}\right) \log\left(\frac{1}{\varepsilon}\right)$$

This shows that SRG+ performs both variance reduction as shown by the dependence of the complexity on σ_*^2 instead of σ^2 and preconditioning as shown by the dependence on $\overline{\kappa}$ instead of κ_{max}.

6. Related work

At a high-level, three lines of work exist that study the use of importance sampling with SGD. The first one considers fixed importance sampling distributions based on the constants of the problem (Needell et al., 2014; Zhao & Zhang, 2015; Gower et al., 2019), and shows that such a strategy leads to improved conditioning of the problem. The second considers adaptive importance sampling and similar to our work targets the variance of the gradient estimator, but generally fails at providing strong convergence rate guarantees under standard assumptions (Papa et al., 2015; Gopal, 2016; Alain et al., 2016; Canevet et al., 2016; Yuan et al., 2016; Stich et al., 2017; Katharopoulos & Fleuret, 2018; Johnson & Guestrin, 2018; Liu et al., 2021). The third line of work frames the problem as an online learning problem with bandit feedback and provides guarantees on the regret of the proposed distributions in terms of the variance of the resulting gradient estimators (Namkoong et al., 2017; Salehi et al., 2017; Borsos et al., 2018; 2019; El Hanchi & Stephens, 2020). Our main method is very closely related to the one proposed by (Papa et al., 2015); the main result of their analysis however is asymptotic in nature. (Liu et al., 2021) provide a non-asymptotic analysis of the algorithm proposed by (Papa et al., 2015), but under strong and non-standard assumptions. To the best of our knowledge, our work is the first to provide non-asymptotic guarantees on the suboptimality of the iterates for variance-reducing importance sampling under standard technical assumptions.

7. Experiments

In this section, we empirically verify our two main claims: (i) SRG performs variance reduction which can improve the asymptotic error of SGD. (ii) SRG+ performs both variance reduction and preconditioning, and can both reduce the asymptotic error of SGD and allow the use of larger step sizes. We start with controlled synthetic experiments that provide direct support for our claims. We then compare SRG to other baseline optimizers on ℓ_2-regularized logistic regression problems. In all experiments, we ran each algorithm 10 times and present the averaged result.

7.1. Synthetic experiments

For our first experiment, we consider the following toy problem. We let $x \in \mathbb{R}$ and $f_i(x) = \frac{1}{2}(x - a_i)^2$ where $a_i = 0$ for $i \in [n-1]$ and $a_n = 1$. In this case, $x^* = 1/n$, $\sigma^2 \approx 1/n$, and $\sigma_*^2 \approx 4/n^2$. We consider five instantiations of this problem with $n \in \{8, 16, 32, 64, 128\}$, yielding ratios σ^2/σ_*^2 approximately in $\{2, 4, 8, 16, 32\}$. For each instantiation, we ran SGD and SRG until they reached stationarity and recorded their asymptotic errors $\lim_{k\to\infty} \mathbb{E}\left[\|x_k - x^*\|_2^2\right]$, which we denote by Δ_{SGD} and Δ_{SRG} respectively. We experimented with three different step sizes, two of them allowed by Corollary 4.4, and one larger one for which we can only prove that SRG has a similar convergence guarantee as SGD. For SRG, we used the mixture coefficient $\theta = 1/2$.

We plot $\Delta_{SGD}/\Delta_{SRG}$ against σ^2/σ_*^2 in Figure 1 (left) for each of the three step sizes, from which we see that the relationship between the two ratios is linear, and very close to identity. From an asymptotic error point of view, these results support our theory in that the improvement is seen to be directly proportional to the ratio σ^2/σ_*^2. On the other hand, the constant $6(1 + 2\theta)/\rho$ in Corollary 4.4 would suggest that the proportionality constant is much smaller than 1, particularly when n is large, but this is not what we observe in practice. This could be because the first term of the Lyapunov function (9) is quite large at stationarity, or because the constants in our bound are not sharp due to the multiple uses of Young's inequality. This latter possibility is further supported by the fact that we see a similar behaviour for SRG for step sizes larger than the ones allowed by our theory. We have consistently made these two observations

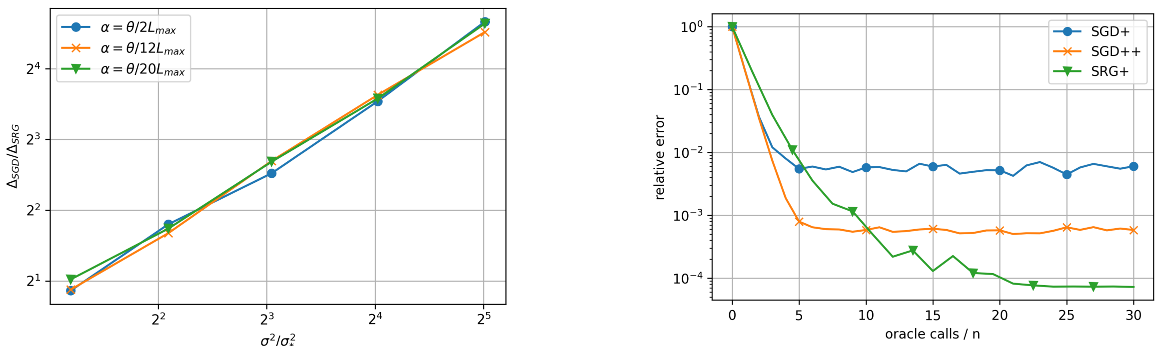

Figure 1: Left: Linear dependence of $\Delta_{SGD}/\Delta_{SRG}$, the ratio of asymptotic errors of SGD and SRG, on σ^2/σ_*^2, the ratio of the uniform and optimal variances at the minimum. Right: SRG+ achieves smaller asymptotic error than both SGD with L_i-sampling (SGD+) and SGD with partially biased sampling (SGD++), while using large $O(1/\overline{L})$ step sizes just like SGD+ and SGD++.

in other experiments. It is however unclear to us how our analysis can be improved to match these observations.

For our second experiment, we considered the following problem. We let $x \in \mathbb{R}$ and $f_i(x) = \frac{1}{2}L_i(x-a_i)^2$, and fixed $n = 20$. Similar to the first experiment, we take $a_i = 0$ for $i \in [n-1]$ and $a_n = 1$. We then set $L_1 = n-1$, $L_n = 1/n$, and $L_i = \frac{n-1}{n(n-2)}$, so that $\overline{L} = 1, L_{\max} = n-1$. In this case, we get $x^* = 1/n^2$, $\sigma_+^2 := \sigma^2(x^*, v) \approx 1/n$, $\sigma^2 \approx 2/n^2$, and $\sigma_*^2 \approx 4/n^3$. We then ran SGD with L_i-sampling (SGD+), SGD with partially biased sampling (SGD++, Needell et al. (2014)), and SRG+, all with the largest allowable step size $\alpha = \theta/12\overline{L}$ from our theory for SRG+, and we use the mixture coefficient $\theta = 1/2$ for both SRG+ and SGD++. We have obtained very similar results when using the larger step size $\alpha = \theta/2\overline{L}$, which is the maximum allowable for SGD+.

We plot the relative error $\mathbb{E}\left[\|x_k - x^*\|_2^2\right] / \|x_0 - x^*\|_2^2$ for each algorithm against the number of gradient oracle calls it makes in Figure 1 (right). We see that all three algorithms are able to converge even when using the large $O(1/\overline{L})$ step sizes. The theory of SGD+ allows for larger step sizes compared to SGD, but the asymptotic error of the algorithm depends on σ_+^2 (Needell et al., 2014), which can be up to n times bigger than σ^2. To solve this problem, Needell et al. (2014) proposed partially biased sampling (SGD++) which mixes the smoothness-based distribution with uniform sampling, allowing the use of larger step sizes just like SGD+, but preserving the asymptotic error of SGD. SRG+ further improves on SGD++, and reduces the asymptotic error, making it proportional to σ_*^2.

7.2. ℓ_2-regularized logistic regression

For our last experiment, we test SRG on ℓ_2-regularized logistic regression problems. In this case, the functions f_i are given by:

$$f_i(x) := \log\left(1 + \exp\left(-y_i a_i^T x\right)\right) + \frac{\mu}{2}\|x\|_2^2$$

where $y_i \in \{-1, 1\}$ is the label of data point $a_i \in \mathbb{R}^d$. Each f_i is convex and $L_i = 0.25\|a_i\|_2^2 + \mu$ smooth. Their average F is also μ-strongly convex. As is standard, we select $\mu = 1/n$.

We experiment with SRG on four datasets from LIBSVM (Chang & Lin, 2011): *ijcnn1, w8a, mushrooms* and *phishing*. For each dataset, and to be able to efficiently run our experiments, we randomly select a subset of the data of size $n = 1000$. As the datasets are normalized, we have that $\|a_i\|_2 = 1$ for all $i \in [n]$. This makes $L_i = L = 0.25 + \mu$, which reduces SGD+ to SGD, and SRG+ to SRG.

We tested the performance of SRG against three baselines, the first of which is standard SGD. The second is SGD with the optimal conditional-variance-minimizing probabilities $p_k^i \propto \|\nabla f_i(x_k)\|_2$, which allows us to compare SRG with the best possible variance-reducing importance sampling scheme. The last baseline is SGD with random shuffling, which is also known to improve the asymptotic error of SGD (Mishchenko et al., 2020).[2] We evaluate the performance of the algorithms by tracking the average relative error $\mathbb{E}\left[\|x_k - x^*\|_2^2\right] / \|x_0 - x^*\|_2^2$. We used the mixture coefficient $\theta = 1/2$ for SRG, and used the same step size $\alpha = \theta/2L$ for all algorithms.

The results of this experiment are shown in Figure 2. We observe that SRG consistently outperforms SGD on all datasets, and that it closely matches the performance of SGD with the variance-minimizing distributions, which it tries to approximate. We also see that SRG outperforms

[2] under the additional assumption that each f_i is also strongly convex

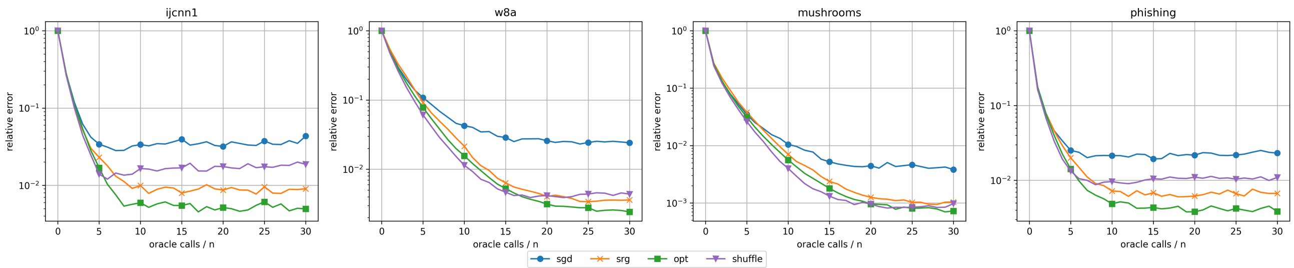

Figure 2: Comparison of the evolution of the average relative error $\|x_k - x^*\|_2^2 \,/\, \|x_0 - x^*\|_2^2$ for different optimizers on ℓ_2-regularized logistic regression problems using the datasets *ijcnn1, w8a, mushrooms, phishing*. We compare SRG (orange) with SGD (blue), SGD with random shuffling (purple), and SGD with the optimal variance-minimizing distributions at each iteration (green).

SGD with random shuffling on two datasets, and is competitive with it on the remaining two.

8. Conclusion

We introduced SRG, a new importance-sampling based stochastic optimization algorithm for finite-sum problems that reduces the variance of the gradient estimator. We analyzed its convergence rate in the strongly convex and smooth case, and showed that it can improve on the asymptotic error of SGD. We also introduced SRG+, an extension of SRG which simultaneously performs variance reduction and preconditioning through importance sampling. We expect our algorithms to be most useful in the medium accuracy regime, where the required accuracy is higher than the one achieved by SGD, but low enough that the overhead of classical variance reduced methods becomes significant. Finally, an interesting future direction would be to explore non-greedy strategies for the design of importance sampling distributions for SGD that not only minimize the variance of the current gradient estimator, but also take into account the variance of the gradient estimators of subsequent iterations.

References

Alain, G., Lamb, A., Sankar, C., Courville, A., and Bengio, Y. Variance Reduction in SGD by Distributed Importance Sampling. *arXiv:1511.06481 [cs, stat]*, April 2016. arXiv: 1511.06481.

Biswas, N., Jacob, P. E., and Vanetti, P. Estimating Convergence of Markov chains with L-Lag Couplings. *Advances in Neural Information Processing Systems*, 32, 2019.

Borsos, Z., Krause, A., and Levy, K. Y. Online Variance Reduction for Stochastic Optimization. In *Conference On Learning Theory*, pp. 324–357. PMLR, July 2018. ISSN: 2640-3498.

Borsos, Z., Curi, S., Levy, K. Y., and Krause, A. Online Variance Reduction with Mixtures. In *International Conference on Machine Learning*, pp. 705–714. PMLR, May 2019. ISSN: 2640-3498.

Canevet, O., Jose, C., and Fleuret, F. Importance Sampling Tree for Large-scale Empirical Expectation. In *International Conference on Machine Learning*, pp. 1454–1462. PMLR, June 2016. ISSN: 1938-7228.

Chang, C.-C. and Lin, C.-J. Libsvm: A library for support vector machines. *ACM Trans. Intell. Syst. Technol.*, 2(3), May 2011.

Defazio, A. A Simple Practical Accelerated Method for Finite Sums. *Advances in Neural Information Processing Systems*, 29:676–684, 2016.

Defazio, A. and Bottou, L. On the Ineffectiveness of Variance Reduced Optimization for Deep Learning. *Advances in Neural Information Processing Systems*, 32: 1755–1765, 2019.

Defazio, A., Bach, F., and Lacoste-Julien, S. SAGA: A Fast Incremental Gradient Method With Support for Non-Strongly Convex Composite Objectives. *Advances in*

Neural Information Processing Systems, 27:1646–1654, 2014.

El Hanchi, A. and Stephens, D. Adaptive Importance Sampling for Finite-Sum Optimization and Sampling with Decreasing Step-Sizes. *Advances in Neural Information Processing Systems*, 33, 2020.

Gopal, S. Adaptive Sampling for SGD by Exploiting Side Information. In *International Conference on Machine Learning*, pp. 364–372. PMLR, June 2016. ISSN: 1938-7228.

Gower, R. M., Loizou, N., Qian, X., Sailanbayev, A., Shulgin, E., and Richtárik, P. SGD: General Analysis and Improved Rates. In *International Conference on Machine Learning*, pp. 5200–5209. PMLR, May 2019. ISSN: 2640-3498.

Hofmann, T., Lucchi, A., Lacoste-Julien, S., and McWilliams, B. Variance Reduced Stochastic Gradient Descent with Neighbors. *Advances in Neural Information Processing Systems*, 28:2305–2313, 2015.

Johnson, R. and Zhang, T. Accelerating Stochastic Gradient Descent using Predictive Variance Reduction. *Advances in Neural Information Processing Systems*, 26:315–323, 2013.

Johnson, T. B. and Guestrin, C. Training Deep Models Faster with Robust, Approximate Importance Sampling. *Advances in Neural Information Processing Systems*, 31: 7265–7275, 2018.

Katharopoulos, A. and Fleuret, F. Not All Samples Are Created Equal: Deep Learning with Importance Sampling. In *International Conference on Machine Learning*, pp. 2525–2534. PMLR, July 2018. ISSN: 2640-3498.

Liu, H., Wang, X., Li, J., and So, A. M.-C. Low-Cost Lipschitz-Independent Adaptive Importance Sampling of Stochastic Gradients. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 2150–2157, Milan, Italy, January 2021. IEEE.

Mishchenko, K., Khaled Ragab Bayoumi, A., and Richtarik, P. Random Reshuffling: Simple Analysis with Vast Improvements. *Advances in Neural Information Processing Systems*, 33, 2020.

Namkoong, H., Sinha, A., Yadlowsky, S., and Duchi, J. C. Adaptive Sampling Probabilities for Non-Smooth Optimization. In *International Conference on Machine Learning*, pp. 2574–2583. PMLR, July 2017. ISSN: 2640-3498.

Needell, D., Ward, R., and Srebro, N. Stochastic Gradient Descent, Weighted Sampling, and the Randomized Kaczmarz algorithm. *Advances in Neural Information Processing Systems*, 27:1017–1025, 2014.

Nesterov, Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Applied Optimization. Springer US, 2004.

Papa, G., Bianchi, P., and Clémençon, S. Adaptive Sampling for Incremental Optimization Using Stochastic Gradient Descent. In Chaudhuri, K., GENTILE, C., and Zilles, S. (eds.), *Algorithmic Learning Theory*, Lecture Notes in Computer Science, pp. 317–331, Cham, 2015. Springer International Publishing.

Robbins, H. and Monro, S. A Stochastic Approximation Method. *Annals of Mathematical Statistics*, 22(3):400–407, September 1951. Publisher: Institute of Mathematical Statistics.

Roux, N., Schmidt, M., and Bach, F. A Stochastic Gradient Method with an Exponential Convergence _rate for Finite Training Sets. *Advances in Neural Information Processing Systems*, 25:2663–2671, 2012.

Salehi, F., Celis, L. E., and Thiran, P. Stochastic Optimization with Bandit Sampling. *arXiv preprint arXiv:1708.02544*, August 2017.

Schmidt, M., Babanezhad, R., Ahmed, M., Defazio, A., Clifton, A., and Sarkar, A. Non-Uniform Stochastic Average Gradient Method for Training Conditional Random Fields. In *Artificial Intelligence and Statistics*, pp. 819–828. PMLR, February 2015. ISSN: 1938-7228.

Schmidt, M., Le Roux, N., and Bach, F. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1):83–112, March 2017.

Stich, S. U., Raj, A., and Jaggi, M. Safe Adaptive Importance Sampling. *Advances in Neural Information Processing Systems*, 30:4381–4391, 2017.

Yuan, K., Ying, B., Vlaski, S., and Sayed, A. H. Stochastic gradient descent with finite samples sizes. In *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6. IEEE, 2016.

Zhao, P. and Zhang, T. Stochastic Optimization with Importance Sampling for Regularized Loss Minimization. In *International Conference on Machine Learning*, pp. 1–9. PMLR, June 2015. ISSN: 1938-7228.

Supplementary material

A. Preliminaries

Lemma A.1. *For all $i \in [n]$, assume that f_i is convex and L_i-smooth. Then we have:*

$$\frac{1}{n}\sum_{i=1}^{n}\|\nabla f_i(x)-\nabla f_i(x^*)\|_2^2 \le 2L_{max}\left[F(x)-F(x^*)\right]$$
$$\frac{1}{n}\sum_{i=1}^{n}\frac{\overline{L}}{L_i}\|\nabla f_i(x)-\nabla f_i(x^*)\|_2^2 \le 2\overline{L}\left[F(x)-F(x^*)\right]$$

Proof. By the convexity and smoothness assumptions on the functions $\{f_i\}_{i=1}^n$, and Theorem 2.1.5 in (Nesterov, 2004) we have for all $i \in [n]$:

$$f_i(x) \ge f_i(x^*) + \langle \nabla f_i(x^*), x - x^*\rangle + \frac{1}{2L_i}\|\nabla f_i(x)-\nabla f_i(x^*)\|_2^2$$

Rearranging gives:

$$\|\nabla f_i(x)-\nabla f_i(x^*)\|_2^2 \le 2L_i\left[f_i(x)-f_i(x^*)-\langle \nabla f_i(x^*), x-x^*\rangle\right]$$

Using $L_i \le L_{\text{max}}$, averaging over $i \in [n]$, and noticing that $\frac{1}{n}\sum_{i=1}^n \nabla f_i(x^*) = \nabla F(x^*) = 0$, we get the first inequality. For the second, we first divide both sides by $\overline{L}/L_i$, then average over $i \in [n]$ and notice that $\frac{1}{n}\sum_{i=1}^n \nabla f_i(x^*) = \nabla F(x^*) = 0$ to get the inequality. □

B. Analysis of SRG

We study the convergence rate of SRG by studying the one-step evolution of the Lyapunov function:

$$T^k = T(x_k, (g_k^i)_{i=1}^n) := \frac{\alpha_k}{\theta_k}\frac{a}{L_{\text{max}}}\sum_{i=1}^n \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + \|x_k - x^*\|_2^2$$

We start by proving Lemma 4.1 which studies the evolution of the first term of T^k.

Lemma 4.1. *Let $k \in \mathbb{N}$ and suppose that $(g_k^i)_{i=1}^n$ evolves as in Algorithm 1. Taking expectation with respect to (b_k, i_k), conditional on $(b_t, i_t)_{t=0}^{k-1}$, we have:*

$$\mathbb{E}\left[\sum_{i=1}^n \left\|g_{k+1}^i - \nabla f_i(x^*)\right\|_2^2\right] \le 2\theta_k L_{max}\left[F(x_k)-F(x^*)\right]$$
$$+\left(1-\frac{\theta_k}{n}\right)\sum_{i=1}^n \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2$$

Proof. Taking expectation with respect to (b_k, i_k) conditional on $(b_t, i_t)_{t=1}^{k-1}$ we have:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{i=1}^{n} \left\|g_{k+1}^{i} - \nabla f_i(x^*)\right\|_2^2\right] \\
&= \mathbb{P}(b_k = 0) \sum_{j=1}^{n} \mathbb{P}(i_k = j \mid b_k = 0) \left(\sum_{i=1}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2\right) \\
&\quad + \mathbb{P}(b_k = 1) \sum_{j=1}^{n} \mathbb{P}(i_k = j \mid b_k = 1) \left(\sum_{\substack{i=1 \\ i \neq j}}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + \left\|\nabla f_j(x_k) - \nabla f_j(x^*)\right\|_2^2\right) \\
&= (1-\theta_k) \sum_{j=1}^{n} q_k^j \left(\sum_{i=1}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2\right) \\
&\quad + \theta_k \sum_{j=1}^{n} \frac{1}{n} \left(\sum_{i=1}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 - \left\|g_k^j - \nabla f_j(x^*)\right\|_2^2 + \left\|\nabla f_j(x_k) - \nabla f_j(x^*)\right\|_2^2\right) \\
&= \left(1 - \frac{\theta_k}{n}\right) \left(\sum_{i=1}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2\right) + \frac{\theta_k}{n} \sum_{i=1}^{n} \left\|\nabla f_i(x_k) - \nabla f_i(x^*)\right\|_2^2 \\
&\leq \left(1 - \frac{\theta_k}{n}\right) \left(\sum_{i=1}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2\right) + 2\theta_k L_{\max} \left[F(x_k) - F(x^*)\right]
\end{aligned}$$

where the first and second equalities follow from the update of $(g_k^i)_{i=1}^n$ in Algorithm 1, and the last inequality follows from the first inequality of Lemma A.1. $\square$

The evolution of the second term of T^k depends on the second moment of the gradient estimator. We prove Lemma 4.2 which gives us a bound on it.

Lemma 4.2. *Let $k \in \mathbb{N}$ and assume that $\theta_k \in (0, 1/2]$. Taking expectation with respect to (b_k, i_k), conditional on $(b_t, i_t)_{t=0}^{k-1}$, we have, for all $\beta, \gamma, \delta, \eta > 0$:*

$$\begin{aligned}
\mathop{\mathbb{E}}_{i_k \sim p_k}\left[\left\|\frac{1}{np_k^{i_k}} \nabla f_{i_k}(x_k)\right\|_2^2\right] &\leq \frac{2D_1 L_{max}}{\theta_k}\left[F(x_k) - F^*\right] \\
&\quad + \frac{D_2}{\theta_k n} \sum_{i=1}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + D_3(1+2\theta_k)\sigma_*^2
\end{aligned}$$

where D_1, D_2 and D_3 are given by:

$$\begin{aligned}
D_1 &\coloneqq (1 + \beta + \gamma) \\
D_2 &\coloneqq (1 + \beta^{-1} + \delta) + (1 + \gamma^{-1} + \delta^{-1})(1 + \eta) \\
D_3 &\coloneqq (1 + \gamma^{-1} + \delta^{-1})(1 + \eta^{-1})
\end{aligned}$$

Proof. Taking expectation with respect to (i_k, b_k) conditional on $(i_t, b_t)_{t=1}^{k-1}$ we have:

$$\mathbb{E}\left[\left\|\frac{1}{np_k^{i_k}} \nabla f_{i_k}(x_k)\right\|_2^2\right] = \frac{1}{n^2} \sum_{i=1}^{n} \frac{1}{p_k^i} \left\|\nabla f_i(x_k)\right\|_2^2$$

Now by Young's inequality (Peter-Paul inequality):

$$\begin{aligned}
\frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|\nabla f_i(x_k)\right\|_2^2 &= \frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|\nabla f_i(x_k)-\nabla f_i(x^*)+\nabla f_i(x^*)-g_k^i+g_k^i\right\|_2^2 \\
&\leq (1+\beta+\gamma)\frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|\nabla f_i(x_k)-\nabla f_i(x^*)\right\|_2^2 \\
&\quad + (1+\beta^{-1}+\delta)\frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|g_k^i-\nabla f_i(x^*)\right\|_2^2 \\
&\quad + (1+\gamma^{-1}+\delta^{-1})\frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|g_k^i\right\|_2^2
\end{aligned}$$

Let us bound each of the three terms. The first can be bound using $p_k \geq \theta_k/n$ and the first inequality of Lemma A.1:

$$\frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|\nabla f_i(x_k)-\nabla f_i(x^*)\right\|_2^2 \leq \frac{1}{\theta_k n}\sum_{i=1}^{n}\left\|\nabla f_i(x_k)-\nabla f_i(x^*)\right\|_2^2 \leq \frac{2L_{\text{max}}}{\theta_k}\left[F(x_k)-F(x^*)\right]$$

The second term is easily bound using $p_k \geq \theta_k/n$:

$$\frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|g_k^i-\nabla f_i(x^*)\right\|_2^2 \leq \frac{1}{\theta_k n}\sum_{i=1}^{n}\left\|g_k^i-\nabla f_i(x^*)\right\|_2^2$$

Finally, the third term can be bounded by:

$$\begin{aligned}
&\frac{1}{n^2}\sum_{i=1}^{n}\frac{1}{p_k^i}\left\|g_k^i\right\|_2^2 \\
&\leq \frac{1}{1-\theta_k}\frac{1}{n^2}\left(\sum_{i=1}^{n}\left\|g_k^i\right\|_2\right)^2 \\
&\leq (1+2\theta_k)\frac{1}{n^2}\left(\sum_{i=1}^{n}\left\|g_k^i-\nabla f_i(x^*)\right\|_2+\sum_{i=1}^{n}\left\|\nabla f_i(x^*)\right\|_2\right)^2 \\
&\leq (1+2\theta_k)(1+\eta)\frac{1}{n^2}\left(\sum_{i=1}^{n}\left\|g_k^i-\nabla f_i(x^*)\right\|_2\right)^2 + (1+2\theta_k)(1+\eta^{-1})\frac{1}{n^2}\left(\sum_{i=1}^{n}\left\|\nabla f_i(x^*)\right\|_2\right)^2 \\
&\leq \theta_k(1+2\theta_k)(1+\eta)\frac{1}{\theta_k n}\sum_{i=1}^{n}\left\|g_k^i-\nabla f_i(x^*)\right\|_2^2 + (1+2\theta_k)(1+\eta^{-1})\sigma_*^2 \\
&\leq (1+\eta)\frac{1}{\theta_k n}\sum_{i=1}^{n}\left\|g_k^i-\nabla f_i(x^*)\right\|_2^2 + (1+2\theta_k)(1+\eta^{-1})\sigma_*^2
\end{aligned}$$

Where the second line follows from the inequality $p_k \geq (1-\theta_k)q_k$ and the definition of q_k in (6), the third from the triangle inequality and the inequality $1/(1-\theta_k) < 1+2\theta_k$ which holds since $\theta_k \in (0, 1/2]$ by hypothesis. The fourth line follows by Young's inequality (Peter-Paul inequality). The fifth line follows from the definition of σ_*^2 and an application of Cauchy-Schwarz inequality: let v be the vector whose i^{th} component is $\left\|g_k^i-\nabla f_i(x^*)\right\|_2$ and let 1_n be the vector of ones. Then:

$$\left(\sum_{i=1}^{n}\left\|g_k^i-\nabla f_i(x^*)\right\|_2\right)^2 = |\langle 1_n, v\rangle|^2 \leq \|1_n\|_2^2\,\|v\|_2^2 = n\sum_{i=1}^{n}\left\|g_k^i-\nabla f_i(x^*)\right\|_2^2$$

Finally, line six follows from the inequality $\theta_k(1+2\theta_k) \leq 1$ since $\theta_k \in (0, 1/2]$ by hypothesis. The result of the lemma then follows after defining:

$$
\begin{aligned}
D_1 &:= (1 + \beta + \gamma) \\
D_2 &:= (1 + \beta^{-1} + \delta) + (1 + \gamma^{-1} + \delta^{-1})(1 + \eta) \\
D_3 &:= (1 + \gamma^{-1} + \delta^{-1})(1 + \eta^{-1})
\end{aligned}
$$

□

Combining these two lemmas, we arrive at our main result, which is a per-step bound on the evolution of T^k under the dynamics of SRG.

Theorem 4.3. *Suppose that $(x_k, (g_k^i)_{i=1}^n)$ evolves according to Algorithm 1. Further, assume that for all $k \in \mathbb{N}$: (i) α_k/θ_k is non-increasing. (ii) $\theta_k \in (0, 1/2]$. (iii) $\alpha_k \leq \theta_k/12L_{max}$. Then:*

$$
\mathbb{E}\left[T^{k+1}\right] \leq (1 - \rho_k)\mathbb{E}\left[T^k\right] + (1 + 2\theta_k)6\alpha_k^2\sigma_*^2
$$

for all $k \in \mathbb{N}$, and where:

$$
\rho_k := \min\left\{\frac{\theta_k}{12n}, \alpha_k\mu\right\}
$$

Proof. All the expectations in this proof are with respect to (b_k, i_k) and are conditional on $(b_t, i_t)_{t=0}^{k-1}$. Since α_k/θ_k is non-increasing, Lemma 4.1 immediately gives us a bound on the first term of $\mathbb{E}\left[T^{k+1}\right]$. The second term of $\mathbb{E}\left[T^{k+1}\right]$ expands as:

$$
\begin{aligned}
\mathbb{E}\left[\|x_{k+1} - x^*\|_2^2\right] &= \mathbb{E}\left[\left\|x_k - \alpha_k \frac{1}{np_k^{i_k}}\nabla f_{i_k}(x_k) - x^*\right\|_2^2\right] \\
&= \|x_k - x^*\|_2^2 - 2\alpha_k\left\langle \mathbb{E}\left[\frac{1}{np_k^{i_k}}\nabla f_{i_k}(x_k)\right], x_k - x^*\right\rangle + \alpha_k^2\mathbb{E}\left[\left\|\frac{1}{np_k^{i_k}}\nabla f_{i_k}(x_k)\right\|_2^2\right] \\
&= \|x_k - x^*\|_2^2 - 2\alpha_k\langle\nabla F(x_k), x_k - x^*\rangle + \alpha_k^2\mathbb{E}\left[\left\|\frac{1}{np_k^{i_k}}\nabla f_{i_k}(x_k)\right\|_2^2\right] \\
&\leq (1 - \alpha_k\mu)\|x_k - x^*\|_2 - 2\alpha_k\left[F(x_k) - F(x^*)\right] + \alpha_k^2\mathbb{E}\left[\left\|\frac{1}{np_k^{i_k}}\nabla f_{i_k}(x_k)\right\|_2^2\right]
\end{aligned}
$$

where in the last line we use the μ-strong-convexity of F to bound the inner product term. Since we are assuming $\theta_k \in (0, 1/2]$, we can apply Lemma 4.2 to bound the last term above. Combining the resulting bound with the one on the first term of the Lyapunov function we get:

$$
\begin{aligned}
\mathbb{E}\left[T^{k+1}\right] \leq &\left(1 - \frac{\theta_k}{n} + \frac{D_2\alpha_k L_{\max}}{na}\right)\frac{\alpha_k}{\theta_k}\frac{a}{L_{\max}}\sum_{i=1}^n\left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + (1 - \alpha_k\mu)\|x_k - x^*\|_2^2 \\
&+ D_3(1 + 2\theta_k)\alpha_k^2\sigma_*^2 \\
&+ 2\alpha_k\left(\frac{D_1\alpha_k L_{\max}}{\theta_k} + a - 1\right)\left[F(x_k) - F(x^*)\right]
\end{aligned}
$$

To ensure that the last parenthesis is not positive, we need:

$$
\alpha_k \leq \frac{(1 - a)}{D_1}\frac{\theta_k}{L_{\max}} \tag{13}
$$

Assuming α_k satisfies this condition, and replacing in the first parenthesis we get:

$$
1 - \frac{\theta_k}{n} + \frac{D_2\alpha_k L_{\max}}{na} \leq 1 - \left(1 - \frac{1 - a}{a}\frac{D_2}{D_1}\right)\frac{\theta_k}{n}
$$

It remains to choose the parameters $\beta, \gamma, \delta, \eta > 0$, and the parameter $a > 0$ so as to maximize the largest allowable step size in (13). First however, note that we have the constraint $a < 1$ so that the step size can be allowed to be positive in (13). Furthermore, since we do not know the relative magnitudes of n and κ_{max}, we decide to impose the equality:

$$1 - \frac{1-a}{a}\frac{D_2}{D_1} = \frac{1-a}{D_1}$$

Finally, we choose to impose $D_3 = 6$. These considerations lead us to the following constrained optimization problem:

$$\begin{aligned} \max_{a,\beta,\gamma,\delta,\eta} \quad & \frac{(1-a)}{D_1} \\ \text{subject to:} \quad & 1 - \frac{1-a}{a}\frac{D_2}{D_1} = \frac{1-a}{D_1} \\ & D_3 = 6 \\ & 0 < a < 1 \\ & \beta, \gamma, \delta, \eta > 0 \end{aligned}$$

which we solve numerically to find the feasible point:

$$a = 0.677,\ \beta = 1.047,\ \gamma = 1.666,\ \delta = 1.591,\ \eta = 0.591$$

With this choice of parameters, we get:

$$\mathbb{E}\left[T^{k+1}\right] \leq (1-\rho_k)\, T^k + (1+2\theta_k) 6\alpha_k^2 \sigma_*^2$$

under the condition:

$$\alpha_k \leq \frac{1}{12}\frac{\theta_k}{L_{\text{max}}}$$

which ensures that (13) holds. □

From this result, we can derive the convergence rate of SRG when it is used with a constant mixture coefficient and step size.

Corollary 4.4. *Suppose that $(x_k, (g_k^i)_{i=1}^n)$ evolves according to Algorithm 1 with a constant mixture coefficient $\theta_k = \theta \in (0, 1/2]$ and a constant step size $\alpha_k = \alpha \leq \theta/12L_{max}$. Then for any $k \in \mathbb{N}$:*

$$\mathbb{E}\left[T^k\right] \leq (1-\rho)^k\, T^0 + (1+2\theta)\frac{6\alpha^2\sigma_*^2}{\rho}$$

where $\rho = \rho_k$ is as defined in Theorem 4.3. For any $\varepsilon > 0$ and $\theta \in (0, 1/2]$, choosing:

$$\alpha = \min\left\{\frac{\theta}{12L_{max}}, \frac{\varepsilon\mu}{(1+2\theta)12\sigma_*^2}, \sqrt{\frac{\theta}{1+2\theta}\frac{\varepsilon}{144n\sigma_*^2}}\right\}$$

and:

$$k \geq \max\left\{\frac{12n}{\theta}, \frac{1}{\alpha\mu}\right\}\log\left(\frac{2T^0}{\varepsilon}\right)$$

guarantees $\mathbb{E}\left[\|x_k - x^*\|_2^2\right] \leq \varepsilon$

Proof. Under the assumptions of the corollary, the conditions of Theorem 4.3 are satisfied for all $k \in \mathbb{N}$. Starting from $\mathbb{E}\left[T^k\right]$ and repeatedly applying Theorem 4.3 we get:

$$\begin{aligned} \mathbb{E}\left[T^k\right] &\leq (1-\rho)^k T^0 + (1+2\theta)6\alpha^2\sigma_*^2 \sum_{t=0}^{k-1}(1-\rho)^t \\ &\leq (1-\rho)^k T^0 + (1+2\theta)6\alpha^2\sigma_*^2 \sum_{t=0}^{\infty}(1-\rho)^t \\ &= (1-\rho)^k T^0 + (1+2\theta)\frac{6\alpha^2\sigma_*^2}{\rho} \end{aligned}$$

This proves the first part. For the complexity part, let $\varepsilon > 0$. We choose the step size α so that:

$$(1+2\theta)\frac{6\alpha^2\sigma_*^2}{\rho} \leq \frac{\varepsilon}{2}$$

We solve this inequality for α using the definition of ρ. Similarly, we need to ensure that the first term satisfies:

$$(1-\rho)^k T^0 \leq \frac{\varepsilon}{2}$$

using the inequality $(1-\rho) \leq \exp(-\rho)$, taking logarithms of both sides, and solving for k we get the desired lower bound. □

Using Theorem 4.3, one can similarly derive the convergence rate of SRG when it is used with decreasing step sizes. In such a case, the mixture coefficient can also be decreased, provided that the ratio α_k/θ_k remains non-increasing as required by Theorem 4.3. We do not pursue this further here.

C. Analysis of SRG+

The analysis of SRG+ is very similar to that of SRG. The Lyapunov function needs to be modified to make use of index-specific constants:

$$T^k = T(x_k, (g_k^i)_{i=1}^n) := \frac{\alpha_k}{\theta_k}\sum_{i=1}^n \frac{a}{L_i}\left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + \|x_k - x^*\|_2^2$$

Similar to Lemma 4.1, the following lemma studies the evolution of the first term of this Lyapunov function.

Lemma C.1. *Let $k \in \mathbb{N}$. Suppose $(g_k^i)_{i=1}^n$ evolves as in Algorithm 2. Taking expectation with respect to (b_k, j_k, i_k), conditional on $(b_t, j_t, i_t)_{t=0}^{k-1}$, we have:*

$$\mathbb{E}\left[\sum_{i=1}^n \frac{1}{L_i}\left\|g_{k+1}^i - \nabla f_i(x^*)\right\|_2^2\right] \leq \left(1 - \frac{\theta_k}{n}\right)\sum_{i=1}^n \frac{1}{L_i}\left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + 2\theta_k\left[F(x_k) - F(x^*)\right]$$

Proof. The proof is identical to that of Lemma 4.1, with the only difference being in the use of the second inequality of Lemma A.1 instead of the first. □

Similar to Lemma 4.2, the following lemma bounds the second moment of the gradient estimator of SRG+.

Lemma C.2. *Let $k \in \mathbb{N}$ and assume that $\theta_k \in (0, 1/2]$. Taking expectation with respect to (b_k, j_k, i_k), conditional on $(b_t, j_t, i_t)_{t=0}^{k-1}$, we have, for all $\beta, \gamma, \delta, \eta > 0$:*

$$\underset{i_k \sim p_k}{\mathbb{E}}\left[\left\|\frac{1}{np_k^{i_k}}\nabla f_{i_k}(x_k)\right\|_2^2\right] \leq \frac{2D_1\overline{L}}{\theta_k}\left[F(x_k) - F^*\right] + \frac{D_2}{\theta_k n}\sum_{i=1}^n \frac{\overline{L}}{L_i}\left\|g_k^i - \nabla f_i(x^*)\right\|_2^2 + D_3(1+2\theta_k)\sigma_*^2$$

where:

$$\begin{aligned}
D_1 &:= (1+\beta+\gamma) \\
D_2 &:= (1+\beta^{-1}+\delta) + (1+\gamma^{-1}+\delta^{-1})(1+\eta) \\
D_3 &:= (1+\gamma^{-1}+\delta^{-1})(1+\eta^{-1})
\end{aligned}$$

Proof. The proof is again very similar to that of Lemma 4.2. In particular, we use the same decomposition in three terms. The first term is bounded using the second inequality of Lemma A.1 instead of the first. The third term is also bounded using the same arguments except in the use of the Cauchy-Schwarz inequality. In particular, we use instead the following bound:

$$\left(\sum_{i=1}^{n} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2\right)^2 = |\langle u, v\rangle| \leq \|u\|_2^2 \|v\|_2^2 = n \sum_{i=1}^{n} \frac{\overline{L}}{L_i} \left\|g_k^i - \nabla f_i(x^*)\right\|_2^2$$

where u is the vector whose i^{th} component is $\sqrt{L_i}$ while v is the vector whose i^{th} component is $\frac{1}{\sqrt{L_i}} \left\|g_k^i - \nabla f_i(x^*)\right\|_2$. $\square$

Combining Lemmas C.1 and C.2, and using identical arguments to the ones used to prove Theorem 4.3, we arrive at Theorem 4.3 for the iterates of SRG+ with the loosening of the condition $\alpha_k \leq \theta_k/12L_{\text{max}}$ to $\alpha_k \leq \theta_k/12\overline{L}$. By extension, Corollary 4.4 also holds for SRG+, with every occurrence of L_{max} replaced by $\overline{L}$.