

3-20-2014

Stochastic Gradient Descent, Weighted Sampling, and the Randomized Kaczmarz Algorithm

Deanna Needell
Claremont McKenna College

Nathan Srebro

Rachel Ward

Recommended Citation

Needell, D., Srebro, N., Ward, R., "Stochastic Gradient Descent, Weighted Sampling, and the Randomized Kaczmarz algorithm", Submitted, arXiv preprint arXiv:1310.5715, 2014.

This Article - preprint is brought to you for free and open access by the CMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

STOCHASTIC GRADIENT DESCENT, WEIGHTED SAMPLING, AND THE RANDOMIZED KACZMARZ ALGORITHM

DEANNA NEEDELL, NATHAN SREBRO, AND RACHEL WARD

ABSTRACT. We obtain an improved finite-sample guarantee on the linear convergence of stochastic gradient descent for smooth and strongly convex objectives, improving from a quadratic dependence on the conditioning $(L/\mu)^2$ (where L is a bound on the smoothness and μ on the strong convexity) to a linear dependence on L/μ . Furthermore, we show how reweighting the sampling distribution (i.e. importance sampling) is necessary in order to further improve convergence, and obtain a linear dependence in the average smoothness, dominating previous results. We also discuss importance sampling for SGD more broadly and show how it can improve convergence also in other scenarios.

Our results are based on a connection we make between SGD and the *randomized Kaczmarz algorithm*, which allows us to transfer ideas between the separate bodies of literature studying each of the two methods. In particular, we recast the randomized Kaczmarz algorithm as an instance of SGD, and apply our results to prove its exponential convergence, but to the solution of a weighted least squares problem rather than the original least squares problem. We then present a modified Kaczmarz algorithm with partially biased sampling which does converge to the original least squares solution with the same exponential convergence rate.

Keywords. distribution reweighting, importance sampling, Kaczmarz method, stochastic gradient descent

1. INTRODUCTION

This paper connects two algorithms which until now have remained remarkably disjoint in the literature: the randomized Kaczmarz algorithm for solving linear systems and the stochastic gradient descent (SGD) method for optimizing a convex objective using unbiased gradient estimates. The connection enables us to make contributions by borrowing from each body of literature to the other. In particular, it helps us highlight the role of weighted sampling for SGD and obtain a tighter guarantee on the linear convergence regime of SGD.

Recall that stochastic gradient descent is a method for minimizing a convex objective $F(\mathbf{x})$ based on access to unbiased stochastic gradient estimates, i.e. to an estimate \mathbf{g} for the gradient at a given point \mathbf{x} , such that $\mathbb{E}[\mathbf{g}] = \nabla F(\mathbf{x})$. Viewing $F(\mathbf{x})$ as an expectation $F(\mathbf{x}) = \mathbb{E}_i[f_i(\mathbf{x})]$, the unbiased gradient estimate can be obtained by drawing i and using its gradient: $\mathbf{g} = \nabla f_i(\mathbf{x})$. SGD originated as “Stochastic Approximation” in the pioneering work of Robbins and Monroe [32], and has recently received renewed attention for confronting very large scale problems, especially in the context of machine learning [4, 33, 24, 2]. Classical analysis of SGD shows a polynomial rate on the sub-optimality of the objective value, $F(\mathbf{x}_k) - F(\mathbf{x}_*)$, namely $1/\sqrt{k}$ for non-smooth objectives, and $1/k$ for smooth, or non-smooth but strongly convex objectives. Such convergence

Date: March 21, 2014.

DN: Department of Mathematical Sciences, Claremont McKenna College (dneedell@cmc.edu). NS: Toyota Technological Institute at Chicago (nati@ttic.edu). RW: Department of Mathematics, University of Texas at Austin (rward@math.utexas.edu).

can be ensured even if the iterates \mathbf{x}_k do not necessarily converge to a unique optimum \mathbf{x}_* , as might be the case if $F(\mathbf{x})$ is not strongly convex. Here we focus on the strongly convex case, where the optimum is unique, and on convergence of the iterates \mathbf{x}_k to the optimum \mathbf{x}_* .

Bach and Moulines [1] recently provided a non-asymptotic bound on the convergence of the iterates in strongly convex SGD, improving on previous results of this kind [19, Section 2.2][3, Section 3.2][34][25]. In particular, they showed that if each $f_i(\mathbf{x})$ is smooth and if \mathbf{x}_* is a minimizer of (almost) all $f_i(\mathbf{x})$, i.e. $\mathbb{P}_i(\nabla f_i(\mathbf{x}_*) = 0) = 1$, then $\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|$ goes to zero exponentially, rather than polynomially, in k . That is, reaching a desired accuracy of $\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|^2 \leq \varepsilon$ requires a number of steps that scales only logarithmically in $1/\varepsilon$. Bach and Moulines’s bound on the required number of iterations further depends on the average *squared* conditioning number $\mathbb{E}[(L_i/\mu)^2]$, where L_i is the Lipschitz constant of $\nabla f_i(\mathbf{x})$ (i.e. $f_i(\mathbf{x})$ are “ L -smooth”), and $F(\mathbf{x})$ is μ -strongly convex. If \mathbf{x}_* is not an exact minimizer of each $f_i(\mathbf{x})$, the bound degrades gracefully as a function of $\sigma^2 = \mathbb{E}\|\nabla f_i(\mathbf{x})\|^2$, and includes an unavoidable term that behaves as σ^2/k .

In a seemingly independent line of research, the *Kaczmarz method* was proposed as an iterative method for solving (usually overdetermined) systems of linear equations [16]. The simplicity of the method makes it useful in a wide array of applications ranging from computer tomography to digital signal processing [13, 20, 15]. Recently, Strohmer and Vershynin [36] proposed a variant of the Kaczmarz method using a random selection method which select rows with probability proportional to their squared norm, and showed that using this selection strategy, a desired accuracy of ε can be reached in the noiseless setting in a number of steps that scales like $\log(1/\varepsilon)$ and *linearly* in the condition number.

1.1. Contributions. Inspired by the analysis of Strohmer and Vershynin and Bach and Moulines, we prove convergence results for stochastic gradient descent as well as for SGD variants where gradient estimates are chosen based on a *weighted sampling distribution*, highlighting the role of importance sampling in SGD.

We first show (Corollary 2.2 in Section 2) that without perturbing the sampling distribution, we can obtain a linear dependence on the *uniform conditioning* ($\sup L_i/\mu$), but it is *not possible* to obtain a linear dependence on the *average conditioning* $\mathbb{E}[L_i/\mu]$. This is a **quadratic improvement over the previous results [1] in regimes where the components have similar Lipschitz constants**.

We then turn to importance sampling, using a weighted sampling distribution. We show that weighting components proportionally to their Lipschitz constants L_i , as is essentially done by Strohmer and Vershynin, can reduce the dependence on the conditioning to a linear dependence on the average conditioning $\mathbb{E}[L_i/\mu]$. However, this comes at an increased dependence on the residual σ^2 . But, we show that by only *partially biasing* the sampling towards L_i , we can enjoy the best of both worlds, obtaining a linear dependence on the average conditioning $\mathbb{E}[L_i/\mu]$, without amplifying the dependence on the residual. Thus, **using importance sampling, we obtain a guarantee dominating, and improving over the previous best-known results [1]** (Corollary 3.1 in Section 2).

In Section 4, we consider the benefits of reweighted SGD also in other scenarios and regimes. We show how also for smooth but not-strongly-convex objectives, **importance sampling can improve a dependence on a uniform bound over smoothness, ($\sup L_i$), to a dependence on the average smoothness $\mathbb{E}[L_i]$** —such an improvement is not possible without importance sampling. For non-smooth objectives, we show that importance sampling can eliminate a dependence on the

variance in the Lipschitz constants of the components. In parallel work we recently became aware of, Zhao and Zhang [40] also consider importance sampling for non-smooth objectives, including composite objectives, suggesting the same reweighting as we obtain here.

Finally, in Section 5, we turn to the Kaczmarz algorithm, explain how it is an instantiation of SGD, and how using partially biased sampling improves known guarantees in this context as well. We show that the randomized Kaczmarz method with uniform i.i.d. row selection can be recast as an instance of preconditioned Stochastic Gradient Descent acting on a re-weighted least squares problem and through this connection, provide exponential convergence rates for this algorithm. We also consider the Kaczmarz algorithm corresponding to SGD with hybrid row selection strategy which shares the exponential convergence rate of Strohmer and Vershynin [36] while also sharing a small error residual term of the SGD algorithm. This presents a clear tradeoff between convergence rate and the convergence residual, not present in other results for the method.

2. SGD FOR STRONGLY CONVEX SMOOTH OPTIMIZATION

We consider the problem of minimizing a smooth convex function,

$$(2.1) \quad \mathbf{x}_\star = \underset{\mathbf{x}}{\operatorname{argmin}} F(\mathbf{x})$$

where $F(\mathbf{x})$ is of the form $F(\mathbf{x}) = \mathbb{E}_{i \sim \mathcal{D}} f_i(\mathbf{x})$ for smooth functionals $f_i : \mathcal{H} \rightarrow \mathcal{R}$ over $\mathcal{H} = \mathcal{R}^d$ endowed with the standard Euclidean norm $\|\cdot\|_2$, or over a Hilbert space \mathcal{H} with the norm $\|\cdot\|_2$. Here i is drawn from some *source distribution* \mathcal{D} over an arbitrary probability space. Throughout this manuscript, unless explicitly specified otherwise, expectations will be with respect to indices drawn from the source distribution \mathcal{D} . That is, we write $\mathbb{E} f_i(\mathbf{x}) = \mathbb{E}_{i \sim \mathcal{D}} f_i(\mathbf{x})$. We also denote by σ^2 the ‘‘residual’’ quantity at the minimum,

$$\sigma^2 = \mathbb{E} \|\nabla f_i(\mathbf{x}_\star)\|_2^2.$$

We will instate the following assumptions on the function F :

- (1) Each f_i is continuously differentiable and the gradient function ∇f_i has Lipschitz constant L_i ; that is, $\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\|_2 \leq L_i \|\mathbf{x} - \mathbf{y}\|_2$ for all vectors \mathbf{x} and \mathbf{y} .
- (2) F has strong convexity parameter μ ; that is, $\langle \mathbf{x} - \mathbf{y}, \nabla F(\mathbf{x}) - \nabla F(\mathbf{y}) \rangle \leq \mu \|\mathbf{x} - \mathbf{y}\|_2^2$ for all vectors \mathbf{x} and \mathbf{y} .

We denote $\sup L$ the supremum of the support of L_i , i.e. the smallest L such that $L_i \leq L$ a.s., and similarly denote $\inf L$ the infimum. We denote the average Lipschitz constant as $\bar{L} = \mathbb{E} L_i$.

A unbiased gradient estimate for $F(\mathbf{x})$ can be obtained by drawing $i \sim \mathcal{D}$ and using $\nabla f_i(\mathbf{x})$ as the estimate. The SGD updates with (fixed) step size γ based on these gradient estimates are then given by:

$$(2.2) \quad \mathbf{x}_{k+1} \leftarrow \mathbf{x}_k - \gamma \nabla f_{i_k}(\mathbf{x}_k)$$

where $\{i_k\}$ are drawn i.i.d. from \mathcal{D} . We are interested in the distance $\|\mathbf{x}_k - \mathbf{x}_\star\|_2^2$ of the iterates from the unique minimum, and denote the initial distance by $\varepsilon_0 = \|\mathbf{x}_0 - \mathbf{x}_\star\|_2^2$.

Bach and Moulines [1, Theorem 1] considered this setting¹ and established that

$$(2.3) \quad k = 2 \log(\varepsilon/\varepsilon_0) \left(\frac{\mathbb{E}L_i^2}{\mu^2} + \frac{\sigma^2}{\mu^2\varepsilon} \right)$$

SGD iterations of the form (2.2), with an appropriate step-size, are sufficient to ensure $\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \varepsilon$, where the expectations is over the random sampling. As long as $\sigma^2 = 0$, i.e. the same minimizer \mathbf{x}_* minimizes all components $f_i(\mathbf{x})$ (though of course it need not be a unique minimizer of any of them), this yields linear convergence to \mathbf{x}_* , with a graceful degradation as $\sigma^2 > 0$. However, in the linear convergence regime, the number of required iterations scales with the expected *squared* conditioning $\mathbb{E}L_i^2/\mu^2$. In this paper, we reduce this quadratic dependence to a linear dependence. We begin with a guarantee ensuring linear dependence, though with a dependence on $\sup L/\mu$ rather than $\mathbb{E}L_i/\mu$:

Theorem 2.1. *Let each f_i be convex where ∇f_i has Lipschitz constant L_i , with $L_i \leq \sup L$ a.s., and let $F(\mathbf{x}) = \mathbb{E}f_i(\mathbf{x})$ be μ -strongly convex. Set $\sigma^2 = \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|_2^2$, where $\mathbf{x}_* = \operatorname{argmin}_{\mathbf{x}} F(\mathbf{x})$. Suppose that $\gamma < \frac{1}{\sup L}$. Then the SGD iterates given by (2.2) satisfy:*

$$(2.4) \quad \mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \left[1 - 2\gamma\mu(1 - \gamma \sup L) \right]^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + \frac{\gamma\sigma^2}{\mu(1 - \gamma \sup L)},$$

where the expectation is with respect to the sampling of $\{i_k\}$.

If we are given a desired tolerance, $\|\mathbf{x} - \mathbf{x}_*\|_2^2 \leq \varepsilon$, and we know the Lipschitz constants and parameters of strong convexity, we may optimize the step-size γ , and obtain:

Corollary 2.2. *For any desired ε , using a step-size of*

$$\gamma = \frac{\mu\varepsilon}{2\varepsilon\mu \sup L + 2\sigma^2}$$

we have that after

$$(2.5) \quad k = 2 \log(\varepsilon_0/\varepsilon) \left(\frac{\sup L}{\mu} + \frac{\sigma^2}{\mu^2\varepsilon} \right)$$

SGD iterations, $\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \varepsilon$, where $\varepsilon_0 = \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2$ and where the expectation is with respect to the sampling of $\{i_k\}$.

Proof sketch. The crux of the improvement over Bach and Moulines is in a tighter recursive equation. Bach and Moulines rely on the recursion

$$\mathbb{E}\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2 \leq (1 - 2\gamma\mu + 2\gamma^2L_i^2) \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 + 2\gamma^2\sigma^2,$$

whereas we use the Co-Coercivity Lemma A.1, with which we can obtain the tighter recursion

$$\mathbb{E}\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2 \leq (1 - 2\gamma\mu + 2\gamma^2\mu L_i) \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 + 2\gamma^2\sigma^2,$$

where L_i is the Lipschitz constant of the component used in the current iterate. The significant difference is that one of the factors of L_i (an upper bound on the second derivative), in the third

¹Bach and Moulines's results are somewhat more general. Their Lipschitz requirement is a bit weaker and more complicated, but in terms of L_i yields (2.3). They also study the use of polynomial decaying step-sizes, but these do not lead to improved runtime if the target accuracy is known ahead of time.

term inside the parenthesis, is replaced by μ (a lower bound on the second derivative). A complete proof can be found in the appendix.

Comparison to results of Bach and Moulines. Our bound (2.5) replaces the dependence on the average *square* conditioning ($\mathbb{E}L_i^2/\mu^2$) with a linear dependence on the *uniform* conditioning ($\sup L/\mu$). When all Lipschitz constants L_i are of similar magnitude, this is a quadratic improvement in the number of required iterations. However, when different components f_i have widely different scaling, i.e. L_i are highly variable, the supremum might be larger than the average square conditioning.

Tightness. Considering the above, one might hope to obtain a linear dependence on the average conditioning $\bar{L}/\mu = \mathbb{E}L_i/\mu$. However, as the following example shows, this is not possible. Consider a uniform source distribution over $N + 1$ quadratics, with the first quadratic f_1 being $(\mathbf{x}[1] - b)^2$ and all others being $(\sqrt{\frac{1}{N}}\mathbf{x}[2])^2$, and $b = \pm 1$. It is clear that for any method one must examine f_1 in order to recover \mathbf{x} to within error less than one, but by uniformly sampling indices i , this takes $(N + 1)$ iterations in expectation. It is easy to verify that in this case, $\sup L_i = L_1 = N$, $\bar{L} = 2$, $\mathbb{E}L_i^2 = N$, and $\mu = 1$. For large N , a linear dependence on \bar{L}/μ would mean that a constant number of iterations suffice (as $\bar{L}/\mu \rightarrow 2$ as $N \rightarrow \infty$), but we just saw that *any* method that sampled i uniformly must consider at least $(N + 1)$ samples in expectation to get non-trivial error. Note that both $\sup L_i/\mu = N$ and $\mathbb{E}L_i^2/\mu^2 = N$ indeed correspond to the correct number of iterations required by SGD.

We therefore see that the choice between a dependence on the average *quadratic* conditioning $\mathbb{E}L_i^2/\mu^2$, or a linear dependence on the *uniform* conditioning $\sup L/\mu$, is unavoidable. A linear dependence on the average conditioning \bar{L}/μ is not possible with any method that samples from the source distribution \mathcal{D} . In the next Section, we will show how we *can* obtain a linear dependence on the average conditioning \bar{L}/μ , using *importance sampling*, i.e. by sampling from a modified distribution.

3. IMPORTANCE SAMPLING

We will now consider stochastic gradient descent, where gradient estimates are sampled from a *weighted distribution*.

3.1. Reweighting a Distribution. For a weight function $w(i)$ which assigns a non-negative weight $w(i) \geq 0$ to each index i , the weighted distribution $\mathcal{D}^{(w)}$ is defined as the distribution such that

$$\mathbb{P}_{\mathcal{D}^{(w)}}(I) \propto \mathbb{E}_{i \sim \mathcal{D}} [\mathbb{1}_I(i)w(i)],$$

where I is an event (subset of indices) and $\mathbb{1}_I(\cdot)$ its indicator function. For a discrete distribution \mathcal{D} with probability mass function $p(i)$ this corresponds to weighting the probabilities to obtain a new probability mass function:

$$p^{(w)}(i) \propto w(i)p(i).$$

Similarly, for a continuous distribution, this corresponds to multiplying the density by $w(i)$ and renormalizing.

One way to construct the weighted distribution $\mathcal{D}^{(w)}$, and sample from it, is through *rejection sampling*: sample $i \sim \mathcal{D}$, and accept with probability $w(i)/W$, for some $W \geq \sup_i w(i)$. Otherwise, reject and continue to re-sample until a suggestion i is accepted. The accepted samples are then distributed according to $\mathcal{D}^{(w)}$.

We use $\mathbb{E}^{(w)}[\cdot] = E_{i \sim \mathcal{D}^{(w)}}[\cdot]$ to denote an expectation where indices are sampled from the weighted distribution $\mathcal{D}^{(w)}$. An important property of such an expectation is that for any quantity $X(i)$ that depends on i :

$$(3.1) \quad \mathbb{E}^{(w)} \left[\frac{1}{w(i)} X(i) \right] = \mathbb{E} [w(i)] \cdot \mathbb{E} [X(i)],$$

where recall that the expectations on the r.h.s. are with respect to $i \sim \mathcal{D}$. In particular, when $\mathbb{E}[w(i)] = 1$, we have that $\mathbb{E}^{(w)} \left[\frac{1}{w(i)} X(i) \right] = \mathbb{E} X(i)$. In fact, we will consider only weights s.t. $\mathbb{E}[w(i)] = 1$, and refer to such weights as *normalized*.

3.2. Reweighted SGD. For any normalized weight function $w(i)$, we can weight each component f_i , defining:

$$(3.2) \quad f_i^{(w)}(\mathbf{x}) = \frac{1}{w(i)} f_i(\mathbf{x})$$

and obtain

$$(3.3) \quad F(\mathbf{x}) = \mathbb{E}^{(w)} [f_i^{(w)}(\mathbf{x})].$$

The representation (3.3) is an equivalent, and equally valid, stochastic representation of the objective $F(\mathbf{x})$, and we can just as well base SGD on this representation. In this case, at each iteration we sample $i \sim \mathcal{D}^{(w)}$ and then use $\nabla f_i^{(w)}(\mathbf{x}) = \frac{1}{w(i)} \nabla f_i(\mathbf{x})$ as an unbiased gradient estimate. SGD iterates based on the representation (3.3), which we will also refer to as w -weighted SGD, are then given by

$$(3.4) \quad \mathbf{x}_{k+1} \leftarrow \mathbf{x}_k - \frac{\gamma}{w(i_k)} \nabla f_{i_k}(\mathbf{x}_k)$$

where $\{i_k\}$ are drawn i.i.d. from $\mathcal{D}^{(w)}$.

The important observation here is that all SGD guarantees are equally valid for the w -weighted updates (3.4)—the objective is the same objective $F(w)$, the sub-optimality is the same, and the minimizer \mathbf{x}_* is the same. We do need, however, to calculate the relevant quantities controlling SGD convergence with respect to the modified components $f_i^{(w)}$ and the weighted distribution $\mathcal{D}^{(w)}$.

3.3. Strongly Convex Smooth Optimization using Weighted SGD. We now return to the analysis of strongly convex smooth optimization and investigate how re-weighting can yield a better guarantee. To do so, we must analyze the relevant quantities involved.

The Lipschitz constant $L_i^{(w)}$ of each component $f_i^{(w)}$ is now scaled, and we have, $L_i^{(w)} = \frac{1}{w(i)} L_i$. The supremum is given by:

$$(3.5) \quad \sup L_{(w)} = \sup_i L_i^{(w)} = \sup_i \frac{L_i}{w(i)}.$$

It is easy to verify that (3.5) is minimized by the weights

$$(3.6) \quad w(i) = \frac{L_i}{\bar{L}},$$

and that with this choice of weights

$$(3.7) \quad \sup L_{(w)} = \sup_i \frac{L_i}{L_i/\bar{L}} = \bar{L}.$$

Note that the average Lipschitz constant $\bar{L} = \mathbb{E}[L_i] = \mathbb{E}^{(w)}[L_i^{(w)}]$ is invariant under weightings.

Before applying Corollary 2.2, we must also calculate:

$$(3.8) \quad \begin{aligned} \sigma_{(w)}^2 &= \mathbb{E}^{(w)}[\|\nabla f_i^{(w)}(\mathbf{x}_*)\|_2^2] = \mathbb{E}^{(w)}\left[\frac{1}{w(i)^2} \|\nabla f_i(\mathbf{x}_*)\|_2^2\right] \\ &= \mathbb{E}\left[\frac{1}{w(i)} \|\nabla f_i(\mathbf{x}_*)\|_2^2\right] = \mathbb{E}\left[\frac{\bar{L}}{L_i} \|\nabla f_i(\mathbf{x}_*)\|_2^2\right] \leq \frac{\bar{L}}{\inf L} \sigma^2. \end{aligned}$$

Now, applying Corollary 2.2 to the w -weighted SGD iterates (3.4) with weights (3.6), we have that, with an appropriate stepsize,

$$(3.9) \quad \begin{aligned} k &= 2 \log(\varepsilon_0/\varepsilon) \left(\frac{\sup L_{(w)}}{\mu} + \frac{\sigma_{(w)}^2}{\mu^2 \varepsilon} \right) \\ &\leq 2 \log(\varepsilon_0/\varepsilon) \left(\frac{\bar{L}}{\mu} + \frac{\bar{L}}{\inf L} \cdot \frac{\sigma_{(w)}^2}{\mu^2 \varepsilon} \right) \end{aligned}$$

iterations are sufficient for $\mathbb{E}^{(w)}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \varepsilon$, where \mathbf{x}_* , μ and ε_0 are exactly as in Corollary 2.2.

3.4. Partially biased sampling. If $\sigma^2 = 0$, i.e. we are in the “realizable” situation, with true linear convergence, then we also have $\sigma_{(w)}^2 = 0$. In this case, we already obtain the desired guarantee: linear convergence with a linear dependence on the average conditioning \bar{L}/μ , strictly improving over Bach and Moulines. However, the inequality in (3.8) might be tight in the presence of components with very small L_i that contribute towards the residual error (as might well be the case for a small component). When $\sigma^2 > 0$, we therefor get a dissatisfying scaling of the second term, relative to Bach and Moulines, by a factor of $\bar{L}/\inf L$.

Fortunately, we can easily overcome this factor. To do so, consider sampling from a distribution which is a *mixture* of the original source distribution and its re-weighting using the weights (3.6). That is, sampling using the weights:

$$(3.10) \quad w(i) = \frac{1}{2} + \frac{1}{2} \cdot \frac{L_i}{\bar{L}}.$$

We refer to this as *partially biased sampling*. Using these weights, we have

$$(3.11) \quad \sup L_{(w)} = \sup_i \frac{1}{\frac{1}{2} + \frac{1}{2} \cdot \frac{L_i}{\bar{L}}} L_i \leq 2\bar{L}$$

and

$$(3.12) \quad \sigma_{(w)}^2 = \mathbb{E}\left[\frac{1}{\frac{1}{2} + \frac{1}{2} \cdot \frac{L_i}{\bar{L}}} \|\nabla f_i(\mathbf{x}_*)\|_2^2\right] \leq 2\sigma^2.$$

Plugging these into Corollary 2.2 we obtain:

Corollary 3.1. *Let each f_i be convex where ∇f_i has Lipschitz constant L_i and let $F(\mathbf{x}) = \mathbb{E}_{i \sim \mathcal{D}}[f_i(\mathbf{x})]$, where $F(\mathbf{x})$ is μ -strongly convex. Set $\sigma^2 = \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|_2^2$, where $\mathbf{x}_* = \operatorname{argmin}_{\mathbf{x}} F(\mathbf{x})$. For any desired ε , using a stepsize of*

$$\gamma = \frac{\mu\varepsilon}{4(\varepsilon\mu\bar{L} + \sigma^2)}$$

we have that after

$$(3.13) \quad k = 4 \log(\varepsilon_0/\varepsilon) \left(\frac{\bar{L}}{\mu} + \frac{\sigma^2}{\mu^2\varepsilon} \right)$$

iterations of w -weighted SGD (3.4) with weights specified by (3.10), $\mathbb{E}^{(w)}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \varepsilon$, where $\varepsilon_0 = \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2$ and $\bar{L} = \mathbb{E}L_i$.

We now obtain the desired linear scaling on \bar{L}/μ , without introducing any additional factor to the residual term, except for a constant factor of two. We thus obtain a result which dominates Bach and Moulines (up to a factor of 2) and substantially improves upon it (with a linear rather than quadratic dependence on the conditioning).

One might also ask whether the previous best known result (2.3) could be improved using weighted sampling. The relevant quantity to consider is the average square Lipschitz constant for the weighted representation: (3.3):

$$(3.14) \quad \overline{L_{(w)}^2} \doteq \mathbb{E}^{(w)} \left[\left(L_i^{(w)} \right)^2 \right] = \mathbb{E}^{(w)} \left[\frac{L_i^2}{w(i)^2} \right] = \mathbb{E} \left[\frac{L_i^2}{w(i)} \right].$$

Interestingly, this quantity is minimized by the same weights as $\sup L_{(w)}$, given by (3.6), and with these weights we have:

$$(3.15) \quad \overline{L_{(w)}^2} = \mathbb{E} \left[\frac{L_i^2}{L_i/\bar{L}} \right] = \bar{L} \mathbb{E}L_i = \bar{L}^2.$$

Again, we can use the partially biased weights give in (3.10), which yields $\overline{L_{(w)}^2} \leq 2\bar{L}^2$ and also ensures $\sigma_{(w)}^2 \leq 2\sigma^2$. In any case, we get a dependence on $\bar{L}^2 = (\mathbb{E}L_i)^2 \leq \mathbb{E}[L_i^2]$ instead of $\bar{L}^2 = \mathbb{E}[L_i^2]$, which is indeed an improvement. Thus, the Bach and Moulines guarantee is also improved by using biased sampling, and in particular the partially biased sampling specified by the weights (3.10). However, relying on Bach and Moulines we still have a quadratic dependence on $(\bar{L}/\mu)^2$, as opposed to the linear dependence we obtain in Corollary 3.1.

3.5. Implementing Importance Sampling. As discussed above, when the magnitudes of L_i are highly variable, importance sampling is necessary in order to obtain a dependence on the average, rather than worst-case, conditioning. In some applications, especially when the Lipschitz constants are known in advance or easily calculated or bounded, such importance sampling might be possible by directly sampling from $D^{(w)}$. This is the case, for example, in trigonometric approximation problems or linear systems which need to be solved repeatedly, or when the Lipschitz constant is easily computed from the data, and multiple passes over the data are needed anyway. We do acknowledge that in other regimes, when data is presented in an online fashion, or when we only have sampling access to the source distribution \mathcal{D} (or the implied distribution over gradient estimates), importance sampling might be difficult.

One option that could be considered, in light of the above results, is to use rejection sampling to simulate sampling from $\mathcal{D}^{(w)}$. For the weights (3.6), this can be done by accepting samples with probability proportional to $L_i / \sup L$. The overall probability of accepting a sample is then $\bar{L} / \sup L$, introducing an additional factor of $\sup L / \bar{L}$. This results in a sample complexity with a linear dependence on $\sup L$, as in Corollary 2.2 (for the weights (3.10), we can first accept with probability 1/2, and then if we do not accept, perform this procedure). Thus, if we are presented samples from \mathcal{D} , and the cost of obtaining the sample dominates the cost of taking the gradient step, we do not gain (but do not lose much either) from rejection sampling. We might still gain from rejection sampling if the cost of operating on a sample (calculating the actual gradient and taking a step according to it) dominates the cost of obtaining it and (a bound on) the Lipschitz constant.

3.6. A Family of Partially Biased Schemes. The choice of weights (3.10) corresponds to an equal mix of uniform and fully biased sampling. More generally, we could consider sampling according to any one of a family of weights which interpolate between uniform and fully biased sampling:

$$(3.16) \quad w^\lambda(i) = \lambda + (1 - \lambda) \frac{L_i}{\bar{L}}, \quad \lambda \in [0, 1].$$

To be concrete, we summarize below the a template algorithm for SGD with partially biased sampling:

Algorithm 3.1: Stochastic Gradient Descent with Partially Biased Sampling

Input:

- Initial estimate $\mathbf{x}_0 \in \mathbb{R}^d$
- Bias parameter $\lambda \in [0, 1]$
- Step size $\gamma > 0$
- Tolerance parameter $\delta > 0$
- Access to the source distribution \mathcal{D}
- If $\lambda < 1$: bounds on the Lipschitz constants L_i ; the weights $w^\lambda(i)$ derived from them (see eq. 3.16); and access to the weighted distribution $\mathcal{D}^{(\lambda)}$.

Output: Estimated solution $\hat{\mathbf{x}}$ to the problem $\min_{\mathbf{x}} F(\mathbf{x})$

$k \leftarrow 0$

repeat

$k \leftarrow k + 1$

 Draw an index $i \sim \mathcal{D}^{(\lambda)}$.

$\mathbf{x}_k \leftarrow \mathbf{x}_{k-1} - \frac{\gamma}{w^\lambda(i)} \nabla f_i(\mathbf{x}_{k-1})$

until $F(\mathbf{x}) \leq \delta$

$\hat{\mathbf{x}} \leftarrow \mathbf{x}_k$

For arbitrary $\lambda \in [0, 1]$, we have the bounds

$$\sup L_{(w^\lambda)} = \sup_i \frac{L_i}{\lambda + (1 - \lambda) \frac{L_i}{\bar{L}}} \leq \min \left(\frac{\bar{L}}{1 - \lambda}, \frac{\sup_i L_i}{\lambda} \right)$$

and

$$\sigma_{(w^\lambda)}^2 = \mathbb{E}\left[\frac{1}{\lambda + (1-\lambda)\frac{L_i}{\bar{L}}}\|\nabla f_i(\mathbf{x}_*)\|_2^2\right] \leq \max\left(\frac{1}{\lambda}, \frac{\bar{L}}{(1-\lambda)\inf_i L_i}\right)$$

Plugging these quantities into Corollary 2.2, we obtain:

Corollary 3.2. *Let each f_i be convex where ∇f_i has Lipschitz constant L_i and let $F(\mathbf{x}) = \mathbb{E}_{i \sim \mathcal{D}}[f_i(\mathbf{x})]$, where $F(\mathbf{x})$ is μ -strongly convex. Set $\sigma^2 = \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|_2^2$, where $\mathbf{x}_* = \operatorname{argmin}_{\mathbf{x}} F(\mathbf{x})$. For any desired ε , using a stepsize of*

$$\gamma = \frac{\mu\varepsilon}{2\varepsilon\mu \min\left(\frac{\bar{L}}{1-\lambda}, \frac{\sup_i L_i}{\lambda}\right) + 2 \max\left(\frac{1}{\lambda}, \frac{\bar{L}}{(1-\lambda)\inf_i L_i}\right)}$$

we have that after

$$k = 2 \log(\varepsilon_0/\varepsilon) \left(\frac{\min\left(\frac{\bar{L}}{1-\lambda}, \frac{\sup_i L_i}{\lambda}\right)}{\mu} + \frac{\max\left(\frac{1}{\lambda}, \frac{\bar{L}}{(1-\lambda)\inf_i L_i}\right)}{\mu^2\varepsilon} \right)$$

iterations of w -weighted SGD (3.4) with partially biased weights (3.16), $\mathbb{E}^{(w)}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \varepsilon$, where $\varepsilon_0 = \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2$ and $\bar{L} = \mathbb{E}L_i$.

It is important to note in this corollary that, **even if λ is close to 1, i.e. we add only a small amount of bias to the sampling**, we have immediately a linear dependence on the average conditioning \bar{L}/μ (multiplied by a factor of $\frac{1}{\lambda}$), rather than a dependence on $\sup L/\mu$.

4. IMPORTANCE SAMPLING FOR SGD IN OTHER SCENARIOS

In the previous Section, we considered SGD for smooth and strongly convex objectives, and were particularly interested in the regime where the residual σ^2 is low, and the linear convergence term is dominant. Weighted SGD could of course be relevant also in other scenarios, and we now briefly survey them, as well as relate them to our main scenario of interest.

4.1. Smooth, Not Strongly Convex. When each component f_i is convex, non-negative, and has an L_i -Lipschitz gradient, but the objective $F(\mathbf{x})$ is not necessarily strongly convex, then after

$$(4.1) \quad k = O\left(\frac{(\sup L)\|\mathbf{x}_*\|_2^2}{\varepsilon} \cdot \frac{F(\mathbf{x}_*) + \varepsilon}{\varepsilon}\right)$$

iterations of SGD with an appropriately chosen step-size we will have $F(\bar{\mathbf{x}}_k) \leq F(\mathbf{x}_*) + \varepsilon$, where $\bar{\mathbf{x}}_k$ is an appropriate averaging of the k iterates Srebro et al. [35]. The relevant quantity here determining the iteration complexity is again $\sup L$. Furthermore, Srebro et al. [35], relying on an example from Foygel and Srebro [11], point out that the dependence on the supremum is unavoidable and *cannot* be replaced with the average Lipschitz constant \bar{L} . That is, if we sample gradients according to the source distribution \mathcal{D} , we must have a linear dependence on $\sup L$.

The only quantity in the bound (4.1) that changes with a re-weighting is $\sup L$ —all other quantities ($\|\mathbf{x}_*\|_2^2$, $F(\mathbf{x}_*)$, and the sub-optimality ε) are invariant to re-weightings. We can therefore replace the dependence on $\sup L$ with a dependence on $\sup L_{(w)}$ by using a weighted SGD as in (3.4). As we already calculated, the optimal weights are given by (3.6), and using them we have

$\sup L_{(w)} = \bar{L}$. In this case, there is no need for partially biased sampling, and we obtain that with an appropriate step-size,

$$(4.2) \quad k = O\left(\frac{\bar{L}\|\mathbf{x}_*\|_2^2}{\varepsilon} \cdot \frac{F(\mathbf{x}_*) + \varepsilon}{\varepsilon}\right)$$

iterations of weighed SGD updates (3.4) using the weights (3.6) suffice.

We again see that using importance sampling allows us to reduce the dependence on $\sup L$, which is unavoidable without biased sampling, to a dependence on \bar{L} .

4.2. Non-Smooth Objectives. We now turn to non-smooth objectives, where the components f_i might not be smooth, but each component is G_i -Lipschitz. Roughly speaking, G_i is a bound on the first derivative (gradient) of f_i , while L_i is a bound on the second derivatives of f_i . Here, the performance of SGD depends on the second moment $\overline{G^2} = \mathbb{E}[G_i^2]$. The precise iteration complexity depends on whether the objective is strongly convex or whether \mathbf{x}_* is bounded, but in either case depends linearly on $\overline{G^2}$.

By using weighted SGD we can replace the linear dependence on $\overline{G^2}$ with a linear dependence on $\overline{G_{(w)}^2} = \mathbb{E}^{(w)}\left[(G_i^{(w)})^2\right]$, where $G_i^{(w)}$ is the Lipschitz constant of the scaled $f_i^{(w)}$ and is given by $G_i^{(w)} = G_i/w(i)$. Again, this follows directly from the standard SGD guarantees, where we consider the representation (3.3) and the unbiased gradient estimates $\nabla f_i^{(w)}(\mathbf{x})$.

We can calculate:

$$(4.3) \quad \overline{G_{(w)}^2} = \mathbb{E}^{(w)}\left[\frac{G_i^2}{w(i)^2}\right] = \mathbb{E}\left[\frac{G_i^2}{w(i)}\right]$$

which is minimized by the weights:

$$(4.4) \quad w(i) = \frac{G_i}{\overline{G}}$$

where $\overline{G} = \mathbb{E}G_i$. Using these weights we have $\overline{G_{(w)}^2} = \mathbb{E}[G_i^2] = \overline{G^2}$. Using importance sampling, we can thus reduce the linear dependence on $\overline{G^2}$ to a linear dependence on \overline{G} . It's helpful to recall that $\overline{G^2} = \overline{G}^2 + \text{Var}[G_i]$. What we save is therefore exactly the variance of the Lipschitz constants G_i .

In parallel work, Zhao and Zhang [40] also consider importance sampling for stochastic optimization for non-smooth objectives. Zhao and Zhang consider a more general setting, with a composite objective that is only partially linearized. But also there, the iteration complexity depends on the second moment of the gradient estimates, and the analysis performed above applies (Zhao and Zhang perform a specialized analysis instead).

4.3. Non-Realizable Regime. Returning to the smooth and strongly convex setting of Sections 2 and 3, let us consider more carefully the residual term $\sigma^2 = \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|_2^2$. This quantity definitively depends on the weighting, and in the analysis of Section 3.3, we avoided increasing it too much, introducing partial biasing for this purpose. However, if this is the dominant term, we might want to choose weights so as to minimize this term. The optimal weights here would be proportional to $\|\nabla f_i(\mathbf{x}_*)\|_2$. The problem is that we do not know the minimizer \mathbf{x}_* , and so cannot calculate these weights. Approaches which dynamically update the weights based on the current iterates as a surrogate for \mathbf{x}_* are possible, but beyond the scope of this paper.

An alternative approach is to bound $\|\nabla f_i(\mathbf{x}_*)\|_2 \leq G_i$ and so $\sigma^2 \leq \overline{G^2}$. Taking this bound, we are back to the same quantity as in the non-smooth case, and the optimal weights are proportional to G_i . Note that this is a different weighting than using weights proportional to L_i , which optimize the linear-convergence term as studied in Section 3.3.

To understand how weighting according to G_i and L_i are different, consider a generalized linear objective where $f_i(\mathbf{x}) = \phi_i(\langle \mathbf{z}_i, \mathbf{x} \rangle)$, and ϕ_i is a scalar function with $|\phi_i'| \leq G_\phi$ and $|\phi_i''| \leq L_\phi$. We have that $G_i \propto \|\mathbf{z}_i\|_2$ while $L_i \propto \|\mathbf{z}_i\|_2^2$. Weighting according to the Lipschitz constants of the gradients, i.e. the ‘‘smoothness’’ parameters, as in (3.6), versus weighting according to the Lipschitz constants of f_i as in (4.4), thus corresponds to weighting according to $\|\mathbf{z}_i\|_2^2$ versus $\|\mathbf{z}_i\|_2$, and are rather different. We can also calculate that weighing by $L_i \propto \|\mathbf{z}_i\|_2^2$ (i.e. following (3.6)), yields $\overline{G_{(w)}^2} = \overline{G^2} > \overline{G^2}$. That is, weights proportional to L_i yield a suboptimal gradient-dependent term (the same dependence as if no weighting at all was used). Conversely, using weights proportional to G_i , i.e. proportional to $\|\mathbf{z}_i\|_2$ yields $\sup L_{(w)} = (E[\sqrt{L_i}])\sqrt{\sup L}$ – a suboptimal dependence, though better than no weighting at all.

Again, as with partially biased sampling, we can weight by the average, $w(i) = \frac{1}{2} \cdot \frac{G_i}{G} + \frac{1}{2} \cdot \frac{L_i}{L}$ and ensure both terms are optimal up to a factor of two.

5. THE LEAST SQUARES CASE AND THE RANDOMIZED KACZMARZ METHOD

A special case of interest is the least squares problem, where

$$(5.1) \quad F(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^n (\langle \mathbf{a}_i, \mathbf{x} \rangle - b_i)^2 = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$$

with \mathbf{b} an n -dimensional vector, \mathbf{A} an $n \times d$ overdetermined full-rank matrix with rows \mathbf{a}_i , and $\mathbf{x}_* = \operatorname{argmin}_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$ is the least-squares solution. Writing the least squares problem (5.1) in the form (2.1), we see that the source distribution \mathcal{D} is uniform over $\{1, 2, \dots, n\}$, the components are $f_i = \frac{n}{2} (\langle \mathbf{a}_i, \mathbf{x} \rangle - b_i)^2$, the Lipschitz constants are $L_i = n \|\mathbf{a}_i\|_2^2$, the average Lipschitz constant is $\frac{1}{n} \sum_i L_i = \|\mathbf{A}\|_F^2$, the strong convexity parameter is $\mu = \frac{1}{\|(\mathbf{A}^T \mathbf{A})^{-1}\|_2}$, so that $K(\mathbf{A}) := \overline{L}/\mu = \|\mathbf{A}\|_F^2 \|(\mathbf{A}^T \mathbf{A})^{-1}\|_2$, and the residual is $\sigma^2 = n \sum_i \|\mathbf{a}_i\|_2^2 |\langle \mathbf{a}_i, \mathbf{x}_* \rangle - b_i|^2$.

The randomized *Kaczmarz method* [36, 6, 14, 20, 38, 7, 37, 12, 41, 21] for solving the least squares problem (5.1) begins with an arbitrary estimate \mathbf{x}_0 , and in the k th iteration selects a row $i = i(k)$ i.i.d. at random from the matrix \mathbf{A} and iterates by:

$$(5.2) \quad \mathbf{x}_{k+1} = \mathbf{x}_k + c \cdot \frac{b_i - \langle \mathbf{a}_i, \mathbf{x}_k \rangle}{\|\mathbf{a}_i\|_2^2} \mathbf{a}_i,$$

where the step size $c = 1$ in the standard method.

Strohmer and Vershynin provided the first non-asymptotic convergence rates, showing that drawing rows proportionally to $\|\mathbf{a}_i\|_2^2$ leads to provable exponential convergence in expectation [36]. Recent works use acceleration techniques to improve convergence rates [18, 27, 9, 31, 41, 10, 8, 5, 28, 29, 30, 23, 22].

However, one can easily verify that the iterates (5.2) are precisely weighted SGD iterates (3.4) with the fully biased weights (3.6).

The reduction of the quadratic dependence on the conditioning to a linear dependence in Theorem 2.1, as well as the use of biased sampling which we investigate here was motivated by

Strohmer and Vershynin’s analysis of the randomized Kaczmarz method. Indeed, applying Theorem 2.1 to the weighted SGD iterates (2.2) for (5.1) with the weights (3.6) and a stepsize of $\gamma = 1$ yields precisely the Strohmer and Vershynin [36] guarantee.

Understanding the randomized Kaczmarz method as SGD allows us also to obtain improved methods and results for the randomized Kaczmarz method:

Using Step-sizes. As shown by Strohmer and Vershynin [36] and extended by Needell [21], the randomized Kaczmarz method with weighted sampling exhibits exponential convergence, but only to within a radius, or *convergence horizon*, of the least-squares solution. This is because a stepsize of $\gamma = 1$ is used, and so the second term in (2.4) does not vanish. It has been shown [38, 7, 37, 12, 23] that changing the step size can allow for convergence inside of this convergence horizon, although non-asymptotic results have been difficult to obtain. Our results allow for finite-iteration guarantees with arbitrary step-sizes and can be immediately applied to this setting. Indeed, applying Theorem 2.1 with the weights (3.6) gives

Corollary 5.1. *Let \mathbf{A} be an $n \times d$ overdetermined full-rank matrix with rows \mathbf{a}_i . Set $\mathbf{e} = \mathbf{A}\mathbf{x}_* - \mathbf{b}$, where \mathbf{x}_* is the minimizer of the problem*

$$\mathbf{x}_* = \operatorname{argmin}_x \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2.$$

Suppose that $c < 1$. Set $a_{\min}^2 = \inf_i \|\mathbf{a}_i\|_2^2$, $a_{\max}^2 = \sup_i \|\mathbf{a}_i\|_2^2$. Then the expected error at the k th iteration of the Kaczmarz method described by (5.2) with row \mathbf{a}_i selected with probability $p_i = \|\mathbf{a}_i\|_2^2 / \|\mathbf{A}\|_F^2$ satisfies

$$(5.3) \quad \mathbb{E} \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \left[1 - \frac{2c(1-c)}{K(\mathbf{A})} \right]^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + \frac{c}{1-c} K(\mathbf{A}) r,$$

with $r = \frac{\sigma^2}{n \|\mathbf{A}\|_F^2 a_{\min}^2}$. The expectation is taken with respect to the weighted distribution over the rows.

When e.g. $c = \frac{1}{2}$, we recover the exponential rate of Strohmer and Vershynin [36] up to a factor of 2, and nearly the same convergence horizon. For arbitrary c , Corollary 5.1 implies a tradeoff between a smaller convergence horizon and a slower convergence rate.

Uniform Row Selection. The Kaczmarz variant of Strohmer and Vershynin [36] calls for weighted row sampling, and thus requires pre-computing all the row norms. Although certainly possible in some applications, in other cases this might be better avoided. Understanding the randomized Kaczmarz as SGD allows us to apply Theorem 2.1 also with uniform weights (i.e. to the unbiased SGD), and obtain a randomized Kaczmarz using uniform sampling, which converges to the least-squares solution and enjoys finite-iteration guarantees:

Corollary 5.2. *Let \mathbf{A} be an $n \times d$ overdetermined full-rank matrix with rows \mathbf{a}_i . Let \mathbf{D} be the diagonal matrix with terms $d_{j,j} = \|\mathbf{a}_i\|_2$, and consider the composite matrix $\mathbf{D}^{-1}\mathbf{A}$. Set $\mathbf{e}_w = \mathbf{D}^{-1}(\mathbf{A}\mathbf{x}_*^w - \mathbf{b})$, where \mathbf{x}_*^w is the minimizer of the weighted least squares problem*

$$\mathbf{x}_*^w = \operatorname{argmin}_x \frac{1}{2} \|\mathbf{D}^{-1}(\mathbf{A}\mathbf{x} - \mathbf{b})\|_2^2.$$

Suppose that $c < 1$. Then the expected error after k iterations of the Kaczmarz method described by (5.2) with uniform row selection satisfies

$$\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*^w\|_2^2 \leq \left[1 - \frac{2c(1-c)}{K(\mathbf{D}^{-1}\mathbf{A})}\right]^k \|\mathbf{x}_0 - \mathbf{x}_*^w\|_2^2 + \frac{c}{1-c} K(\mathbf{D}^{-1}\mathbf{A}) r_w,$$

where $r_w = \|\mathbf{e}_w\|_2^2/n$.

Note that the randomized Kaczmarz algorithm with uniform row selection converges exponentially to a *weighted* least-squares solution, to within arbitrary accuracy by choosing sufficiently small stepsize c . Thus, in general, the randomized Kaczmarz algorithms with uniform and biased row selection converge (up to a convergence horizon) towards different solutions.

Partially Biased Sampling. As in our SGD analysis, using the partially biased sampling weights is applicable also for the randomized Kaczmarz method. Applying Theorem 2.1 using weights (3.10) gives

Corollary 5.3 (Randomized Kaczmarz with partially biased sampling). *Let \mathbf{A} be an $n \times d$ overdetermined full-rank matrix with rows \mathbf{a}_i . Set $\mathbf{e} = \mathbf{A}\mathbf{x}_* - \mathbf{b}$, where \mathbf{x}_* is the minimizer of the problem*

$$\mathbf{x}_* = \operatorname{argmin}_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2.$$

Suppose $c < 1/2$. Then the iterate \mathbf{x}_k of the modified Kaczmarz method described by

$$(5.4) \quad \mathbf{x}_{k+1} = \mathbf{x}_k + 2c \cdot \frac{b_i - \langle \mathbf{a}_i, \mathbf{x}_k \rangle}{\|\mathbf{A}\|_F^2/n + \|\mathbf{a}_i\|_2^2} \mathbf{a}_i$$

with row \mathbf{a}_i selected with probability $p_i = \frac{1}{2} \cdot \frac{\|\mathbf{a}_i\|_2^2}{\|\mathbf{A}\|_F^2} + \frac{1}{2} \cdot \frac{1}{n}$ satisfies

$$(5.5) \quad \mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \left[1 - \frac{2c(1-2c)}{K(\mathbf{A})}\right]^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + \frac{cK(\mathbf{A})}{1-2c} \cdot \frac{2\sigma^2}{n\|\mathbf{A}\|_F^2}$$

The partially biased randomized Kaczmarz method described above (which does have modified update equation (5.4) compared to the standard update equation (5.2)) yields the same convergence rate as the fully biased randomized Kaczmarz method [36] (up to a factor of 2), but gives a better dependence on the residual error over the fully biased sampling, as the final term in (5.5) is smaller than the final term in (5.3).

6. NUMERICAL EXPERIMENTS

In this section we present some numerical results for the randomized Kaczmarz algorithm with partially biased sampling, that is, applying Algorithm 3.1 to the least squares problem $F(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$ (so $f_i(\mathbf{x}) = \frac{n}{2} (\langle \mathbf{a}_i, \mathbf{x} \rangle - b_i)^2$) and considering $\lambda \in [0, 1]$. Recall that $\lambda = 0$ corresponds to the randomized Kaczmarz algorithm of Strohmer and Vershynin with fully weighted sampling [36]. $\lambda = .5$ corresponds to the partially biased randomized Kaczmarz algorithm outlined in Corollary 5.3. We demonstrate how the behavior of the algorithm depends on λ , the conditioning of the system, and the residual error at the least squares solution. We focus on exploring the role of λ on the convergence rate of the algorithm for various types of matrices \mathbf{A} . We consider five types of systems, described below, each using a 1000×10 matrix \mathbf{A} . In each setting, we create a vector \mathbf{x} with standard normal entries. For the described matrix \mathbf{A} and residual \mathbf{e} , we create the

system $\mathbf{b} = \mathbf{A}\mathbf{x} + \mathbf{e}$ and run the randomized Kaczmarz method with various choices of λ . Each experiment consists of 100 independent trials and uses the optimal step size as in Corollary 3.2 with $\varepsilon = .1$; the plots show the average behavior over these trials. The settings below show the various types of behavior the Kaczmarz method can exhibit.

Case 1: Each row of the matrix \mathbf{A} has standard normal entries, except the last row which has normal entries with mean 0 and variance 10^2 . The residual vector \mathbf{e} has normal entries with mean 0 and variance 0.1^2 .

Case 2: Each row of the matrix \mathbf{A} has standard normal entries. The residual vector \mathbf{e} has normal entries with mean 0 and variance 0.1^2 .

Case 3: The j th row of \mathbf{A} has normal entries with mean 0 and variance j . The residual vector \mathbf{e} has normal entries with mean 0 and variance 20^2 .

Case 4: The j th row of \mathbf{A} has normal entries with mean 0 and variance j . The residual vector \mathbf{e} has normal entries with mean 0 and variance 10^2 .

Case 5: The j th row of \mathbf{A} has normal entries with mean 0 and variance j . The residual vector \mathbf{e} has normal entries with mean 0 and variance 0.1^2 .

Figure 1 shows the convergence behavior of the randomized Kaczmarz method in each of these five settings. As expected, when the rows of \mathbf{A} are far from normalized, as in Case 1, we see different behavior as λ varies from 0 to 1. Here, weighted sampling ($\lambda = 0$) significantly outperforms uniform sampling ($\lambda = 1$), and the trend is monotonic in λ . On the other hand, when the rows of \mathbf{A} are close to normalized, as in Case 2, the various λ give rise to similar convergence rates, as is expected. Out of the λ tested (we tested increments of 0.1 from 0 to 1), the choice $\lambda = 0.7$ gave the worst convergence rate, and again purely weighted sampling gives the best. Still, the worst-case convergence rate was not much worse, as opposed to the situation with uniform sampling in Case 1. Cases 3, 4, and 5 use matrices with varying row norms and cover “high”, “medium”, and “low” noise regimes, respectively. In the high noise regime (Case 3), we find that fully weighted sampling, $\lambda = 0$, is relatively very slow to converge, as the theory suggests, and hybrid sampling outperforms both weighted and uniform selection. In the medium noise regime (Case 4), hybrid sampling still outperforms both weighted and uniform selection. Again, this is not surprising, since hybrid sampling allows a balance between small convergence horizon (important with large residual norm) and convergence rate. As we decrease the noise level (as in Case 5), we see that again weighted sampling is preferred.

Figure 2 shows the number of iterations of the randomized Kaczmarz method needed to obtain a fixed approximation error. For the choice $\lambda = 1$ for Case 1, we cut off the number of iterations after 50,000, at which point the desired approximation error was still not attained. As seen also from Figure 1, Case 1 exhibits monotonic improvements as we scale λ . For Cases 2 and 5, the optimal choice is pure weighted sampling, whereas Cases 3 and 4 prefer intermediate values of λ .

7. SUMMARY AND OUTLOOK

We consider this paper as making three contributions: the improved dependence on the conditioning for smooth and strongly convex SGD, the discussion of importance sampling for SGD, and the connection between SGD and the randomized Kaczmarz method.

For simplicity, we only considered SGD iterates with a fixed step-size γ . This is enough for getting the optimal iteration complexity if the target accuracy ε is known in advance, which was

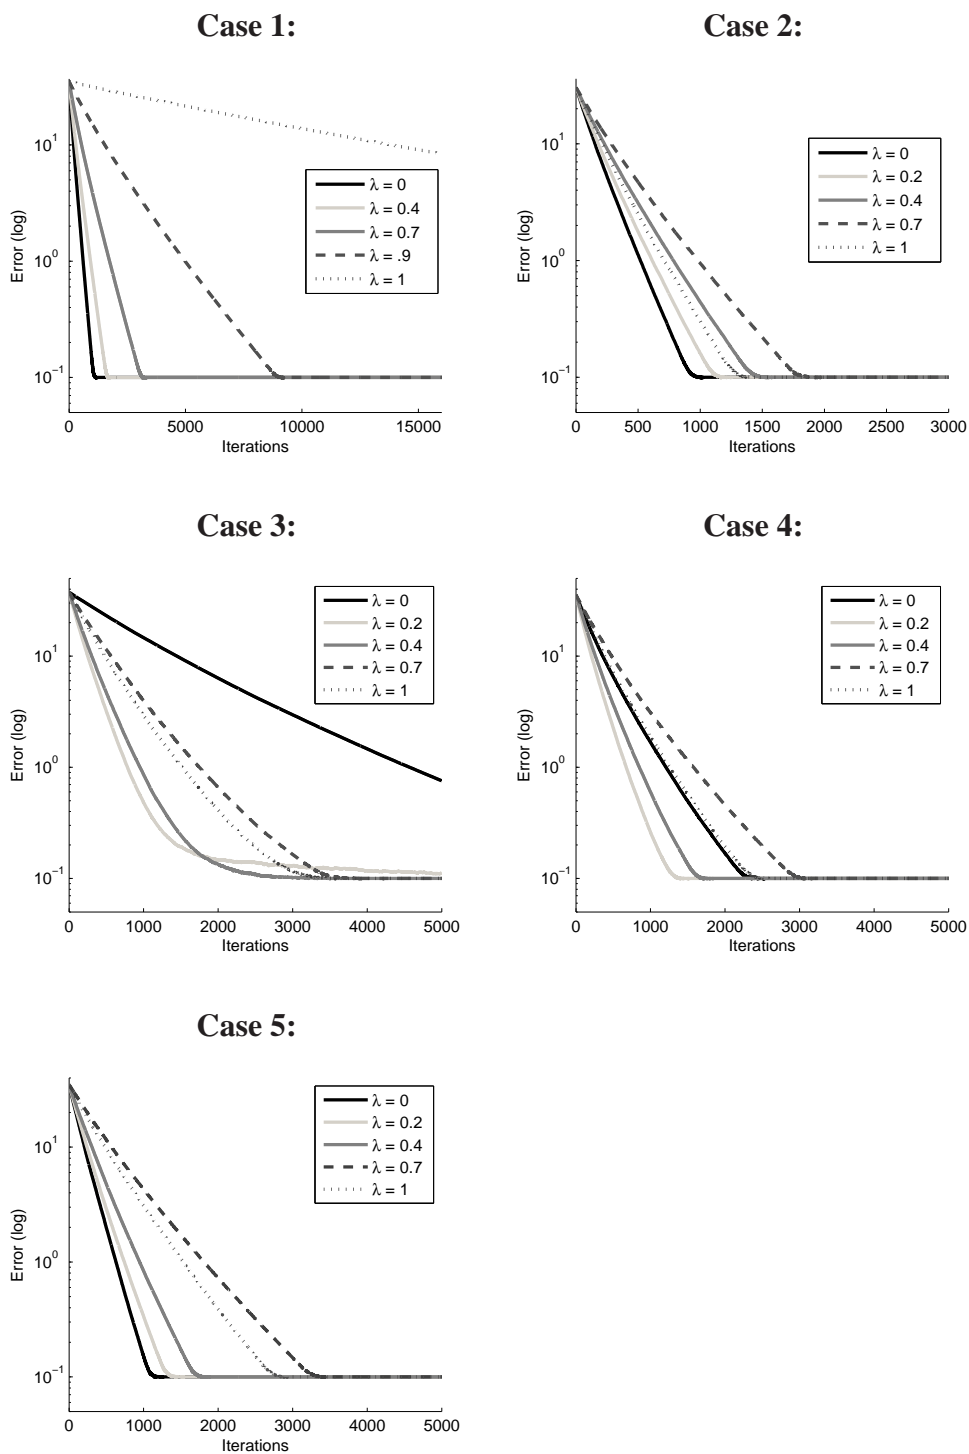


FIGURE 1. The convergence rates for the randomized Kaczmarz method with various choices of λ in the five settings described above. The vertical axis is in logarithmic scale and depicts the approximation error $\|x_k - x_*\|_2^2$ at iteration k (the horizontal axis).




FIGURE 2. Number of iterations k needed by the randomized Kaczmarz method with partially biased sampling, for various values of λ , to obtain approximation error $\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \varepsilon = 0.1$ in the five cases described above: Case 1 (blue with circle marker), Case 2 (red with square marker), Case 3 (black with triangle marker), Case 4 (green with x marker), and Case 5 (purple with star marker).

our approach in this paper. It is easy to adapt the analysis, using standard techniques, to incorporate decaying step-sizes, which are appropriate if we don't know ε in advance.

We suspect that the assumption of strong convexity can be weakened to *restricted strong convexity* [17, 39] without changing any of the results of this paper; we leave this analysis to future work.

Finally, our discussion of importance sampling is limited to a static reweighting of the sampling distribution. A more sophisticated approach would be to update the sampling distribution dynamically as the method progresses, and as we gain more information about the relative importance of components. Although such dynamic importance sampling is sometimes attempted heuristically, we are not aware of any rigorous analysis of such a dynamic biasing scheme.

ACKNOWLEDGEMENTS

DN was partially supported by a Simons Foundation Collaboration grant, NS was partially supported by a Google Research Award, and RW was supported in part by ONR Grant N00014-12-1-0743, an AFOSR Young Investigator Program Award, and an NSF CAREER award.

REFERENCES

- [1] F. Bach and E. Moulines. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [2] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.
- [3] L. Bottou and O. Bousquet. The tradeoffs of large-scale learning. *Optimization for Machine Learning*, page 351, 2011.

- [4] O. Bousquet and L. Bottou. The tradeoffs of large scale learning. In *Advances in Neural Information Processing Systems*, pages 161–168, 2007.
- [5] C. L. Byrne. *Applied iterative methods*. A K Peters Ltd., Wellesley, MA, 2008. ISBN 978-1-56881-342-4; 1-56881-271-X.
- [6] C. Cenker, H. G. Feichtinger, M. Mayer, H. Steier, and T. Strohmer. New variants of the POCS method using affine subspaces of finite codimension, with applications to irregular sampling. In *Proc. SPIE: Visual Communications and Image Processing*, pages 299–310, 1992.
- [7] Y. Censor, Paul PB Eggermont, and Dan Gordon. Strong underrelaxation in Kaczmarz’s method for inconsistent systems. *Numerische Mathematik*, 41(1):83–92, 1983.
- [8] P. P. B. Eggermont, G. T. Herman, and A. Lent. Iterative algorithms for large partitioned linear systems, with applications to image reconstruction. *Linear Algebra Appl.*, 40:37–67, 1981. ISSN 0024-3795. doi: 10.1016/0024-3795(81)90139-7. URL [http://dx.doi.org/10.1016/0024-3795\(81\)90139-7](http://dx.doi.org/10.1016/0024-3795(81)90139-7).
- [9] Y. C. Eldar and D. Needell. Acceleration of randomized Kaczmarz method via the Johnson-Lindenstrauss lemma. *Numer. Algorithms*, 58(2):163–177, 2011. ISSN 1017-1398. doi: 10.1007/s11075-011-9451-z. URL <http://dx.doi.org/10.1007/s11075-011-9451-z>.
- [10] T. Elfving. Block-iterative methods for consistent and inconsistent linear equations. *Numer. Math.*, 35(1):1–12, 1980. ISSN 0029-599X. doi: 10.1007/BF01396365. URL <http://dx.doi.org/10.1007/BF01396365>.
- [11] R. Foygel and N. Srebro. Concentration-based guarantees for low-rank matrix reconstruction. *24th Ann. Conf. Learning Theory (COLT)*, 2011.
- [12] M. Hanke and W. Niethammer. On the acceleration of Kaczmarz’s method for inconsistent linear systems. *Linear Algebra and its Applications*, 130:83–98, 1990.
- [13] G. T. Herman. *Fundamentals of computerized tomography: image reconstruction from projections*. Springer, 2009.
- [14] G.T. Herman and L.B. Meyer. Algebraic reconstruction techniques can be made computationally efficient. *IEEE Trans. Medical Imaging*, 12(3):600–609, 1993.
- [15] G. N Hounsfield. Computerized transverse axial scanning (tomography): Part 1. description of system. *British Journal of Radiology*, 46(552):1016–1022, 1973.
- [16] S. Kaczmarz. Angenäherte auflösung von systemen linearer gleichungen. *Bull. Int. Acad. Polon. Sci. Lett. Ser. A*, pages 335–357, 1937.
- [17] M. Lai and W. Yin. Augmented ℓ_1 and nuclear-norm models with a globally linearly convergent algorithm. *SIAM J. Imaging Sciences*, 2013.
- [18] Y. T. Lee and A. Sidford. Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems. submitted, 2013.
- [19] N. Murata. *A statistical study of on-line learning*. Cambridge University Press, Cambridge, UK, 1998.
- [20] F. Natterer. *The mathematics of computerized tomography*, volume 32 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001. ISBN 0-89871-493-1. doi: 10.1137/1.9780898719284. URL <http://dx.doi.org/10.1137/1.9780898719284>. Reprint of the 1986 original.

- [21] D. Needell. Randomized Kaczmarz solver for noisy linear systems. *BIT*, 50(2):395–403, 2010. ISSN 0006-3835. doi: 10.1007/s10543-010-0265-5. URL <http://dx.doi.org/10.1007/s10543-010-0265-5>.
- [22] D. Needell and J. A. Tropp. Paved with good intentions: Analysis of a randomized block kaczmarz method. *Linear Algebra and its Applications*, 2013.
- [23] D. Needell and R. Ward. Two-subspace projection method for coherent overdetermined linear systems. *Journal of Fourier Analysis and Applications*, 19(2):256–269, 2013.
- [24] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [25] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [26] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Kluwer, 2004.
- [27] C. Popa. Extensions of block-projections methods with relaxation parameters to inconsistent and rank-deficient least-squares problems. *BIT*, 38(1):151–176, 1998. ISSN 0006-3835. doi: 10.1007/BF02510922. URL <http://dx.doi.org/10.1007/BF02510922>.
- [28] C. Popa. Block-projections algorithms with blocks containing mutually orthogonal rows and columns. *BIT*, 39(2):323–338, 1999. ISSN 0006-3835. doi: 10.1023/A:1022398014630. URL <http://dx.doi.org/10.1023/A:1022398014630>.
- [29] C. Popa. A fast Kaczmarz-Kovarik algorithm for consistent least-squares problems. *Korean J. Comput. Appl. Math.*, 8(1):9–26, 2001. ISSN 1229-9502.
- [30] C. Popa. A Kaczmarz-Kovarik algorithm for symmetric ill-conditioned matrices. *An. Ştiinţ. Univ. Ovidius Constanţa Ser. Mat.*, 12(2):135–146, 2004. ISSN 1223-723X.
- [31] C. Popa, T. Preclik, H. Köstler, and U. Råde. On Kaczmarz's projection iteration as a direct solver for linear least squares problems. *Linear Algebra and Its Applications*, 436(2):389–404, 2012.
- [32] H. Robbins and S. Monroe. A stochastic approximation method. *Ann. Math. Statist.*, 22:400–407, 1951.
- [33] S. Shalev-Shwartz and N. Srebro. Svm optimization: inverse dependence on training set size. In *Proceedings of the 25th international conference on Machine learning*, pages 928–935, 2008.
- [34] O. Shamir and T. Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. *arXiv preprint arXiv:1212.1824*, 2012.
- [35] N. Srebro, K. Sridharan, and A. Tewari. Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems*, 2010.
- [36] T. Strohmer and R. Vershynin. A randomized Kaczmarz algorithm with exponential convergence. *J. Fourier Anal. Appl.*, 15(2):262–278, 2009. ISSN 1069-5869. doi: 10.1007/s00041-008-9030-4. URL <http://dx.doi.org/10.1007/s00041-008-9030-4>.
- [37] K. Tanabe. Projection method for solving a singular system of linear equations and its applications. *Numerische Mathematik*, 17(3):203–214, 1971.
- [38] T. M. Whitney and R. K. Meany. Two algorithms related to the method of steepest descent. *SIAM Journal on Numerical Analysis*, 4(1):109–118, 1967.
- [39] H. Zhang and W. Yin. Gradient methods for convex minimization: better rates under weaker conditions. CAM Report 13-17, UCLA, 2013.
- [40] P. Zhao and T. Zhang. Stochastic optimization with importance sampling. submitted, 2014.

[41] A. Zouzias and N. M. Freris. Randomized extended Kaczmarz for solving least-squares. *SIAM Journal on Matrix Analysis and Applications*, 2012.

APPENDIX A. PROOFS

Our main results utilize an elementary fact about smooth functions with Lipschitz continuous gradient, called the co-coercivity of the gradient. We state the lemma and recall its proof for completeness.

A.1. The Co-coercivity Lemma.

Lemma A.1 (Co-coercivity). *For a smooth function f whose gradient has Lipschitz constant L ,*

$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_2^2 \leq L \langle \mathbf{x} - \mathbf{y}, \nabla f(\mathbf{x}) - \nabla f(\mathbf{y}) \rangle.$$

Proof. Since ∇f has Lipschitz constant L , if \mathbf{x}_* is the minimizer of f , then [see e.g. 26, page 26] (A.1)

$$\frac{1}{2L} \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}_*)\|_2^2 = \frac{1}{2L} \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}_*)\|_2^2 + \langle \mathbf{x} - \mathbf{x}_*, \nabla f(\mathbf{x}_*) \rangle \leq f(\mathbf{x}) - f(\mathbf{x}_*);$$

Now define the convex functions

$$G(\mathbf{z}) = f(\mathbf{z}) - \langle \nabla f(\mathbf{x}), \mathbf{z} \rangle, \quad \text{and} \quad H(\mathbf{z}) = f(\mathbf{z}) - \langle \nabla f(\mathbf{y}), \mathbf{z} \rangle,$$

and observe that both have Lipschitz constants L and minimizers \mathbf{x} and \mathbf{y} , respectively. Applying (A.1) to these functions therefore gives that

$$G(\mathbf{x}) \leq G(\mathbf{y}) - \frac{1}{2L} \|\nabla G(\mathbf{y})\|_2^2, \quad \text{and} \quad H(\mathbf{y}) \leq H(\mathbf{x}) - \frac{1}{2L} \|\nabla H(\mathbf{y})\|_2^2.$$

By their definitions, this implies that

$$\begin{aligned} f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{x} \rangle &\leq f(\mathbf{y}) - \langle \nabla f(\mathbf{x}), \mathbf{y} \rangle - \frac{1}{2L} \|\nabla f(\mathbf{y}) - \nabla f(\mathbf{x})\|_2^2 \\ f(\mathbf{y}) - \langle \nabla f(\mathbf{y}), \mathbf{y} \rangle &\leq f(\mathbf{x}) - \langle \nabla f(\mathbf{y}), \mathbf{x} \rangle - \frac{1}{2L} \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_2^2. \end{aligned}$$

Adding these two inequalities and canceling terms yields the desired result. \square

A.2. Proof of Theorem 2.1. With the notation of Theorem 2.1, and where i is the random index chosen at iteration k , and $w = w_\lambda$, we have

$$\begin{aligned} \|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2 &= \|\mathbf{x}_k - \mathbf{x}_* - \gamma \nabla f_i(\mathbf{x}_k)\|_2^2 \\ &= \|(\mathbf{x}_k - \mathbf{x}_*) - \gamma(\nabla f_i(\mathbf{x}_k) - \nabla f_i(\mathbf{x}_*)) - \gamma \nabla f_i(\mathbf{x}_*)\|_2^2 \\ &= \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma \langle \mathbf{x}_k - \mathbf{x}_*, \nabla f_i(\mathbf{x}_k) \rangle + \gamma^2 \|\nabla f_i(\mathbf{x}_k) - \nabla f_i(\mathbf{x}_*) + \nabla f_i(\mathbf{x}_*)\|_2^2 \\ &\leq \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma \langle \mathbf{x}_k - \mathbf{x}_*, \nabla f_i(\mathbf{x}_k) \rangle + 2\gamma^2 \|\nabla f_i(\mathbf{x}_k) - \nabla f_i(\mathbf{x}_*)\|_2^2 + 2\gamma^2 \|\nabla f_i(\mathbf{x}_*)\|_2^2 \\ &\leq \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma \langle \mathbf{x}_k - \mathbf{x}_*, \nabla f_i(\mathbf{x}_k) \rangle \\ &\quad + 2\gamma^2 L_i \langle \mathbf{x}_k - \mathbf{x}_*, \nabla f_i(\mathbf{x}_k) - \nabla f_i(\mathbf{x}_*) \rangle + 2\gamma^2 \|\nabla f_i(\mathbf{x}_*)\|_2^2, \end{aligned}$$

where we have employed Jensen's inequality in the first inequality and the co-coercivity Lemma A.1 in the final line. We next take an expectation with respect to the choice of i . By assumption,

$i \sim \mathcal{D}$ such that $F(\mathbf{x}) = \mathbb{E}f_i(\mathbf{x})$ and $\sigma^2 = \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|^2$. Then $\mathbb{E}\nabla f_i(\mathbf{x}) = \nabla F(\mathbf{x})$, and we obtain:

$$\begin{aligned}
\mathbb{E}\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2 &\leq \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma \langle \mathbf{x}_k - \mathbf{x}_*, \nabla F(\mathbf{x}_k) \rangle + 2\gamma^2 \mathbb{E} [L_i \langle \mathbf{x}_k - \mathbf{x}_*, \nabla f_i(\mathbf{x}_k) - \nabla f_i(\mathbf{x}_*) \rangle] \\
&\quad + 2\gamma^2 \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|_2^2 \\
&\leq \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma \langle \mathbf{x}_k - \mathbf{x}_*, \nabla F(\mathbf{x}_k) \rangle + 2\gamma^2 \mathbb{E} [L_i \langle \mathbf{x}_k - \mathbf{x}_*, \nabla f_i(\mathbf{x}_k) - \nabla f_i(\mathbf{x}_*) \rangle] \\
&\quad + 2\gamma^2 \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|_2^2 \\
&\leq \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma \langle \mathbf{x}_k - \mathbf{x}_*, \nabla F(\mathbf{x}_k) \rangle + 2\gamma^2 \sup_i L_i \mathbb{E} \langle \mathbf{x}_k - \mathbf{x}_*, \nabla f_i(\mathbf{x}_k) - \nabla f_i(\mathbf{x}_*) \rangle \\
&\quad + 2\gamma^2 \mathbb{E}\|\nabla f_i(\mathbf{x}_*)\|_2^2 \\
&= \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma \langle \mathbf{x}_k - \mathbf{x}_*, \nabla F(\mathbf{x}_k) \rangle + 2\gamma^2 \sup L \langle \mathbf{x}_k - \mathbf{x}_*, \nabla F(\mathbf{x}_k) - \nabla F(\mathbf{x}_*) \rangle + 2\gamma^2 \sigma^2
\end{aligned}$$

We now utilize the strong convexity of $F(\mathbf{x})$ and obtain that

$$\begin{aligned}
&\leq \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 - 2\gamma\mu(1 - \gamma \sup L)\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 + 2\gamma^2 \sigma^2 \\
&= (1 - 2\gamma\mu(1 - \gamma \sup L))\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 + 2\gamma^2 \sigma^2
\end{aligned}$$

when $\gamma \leq \frac{1}{\sup L}$. Recursively applying this bound over the first k iterations yields the desired result,

$$\begin{aligned}
\mathbb{E}^{(w)}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 &\leq \left(1 - 2\gamma\mu(1 - \gamma \sup L)\right)^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + 2 \sum_{j=0}^{k-1} \left(1 - 2\gamma\mu(1 - \gamma \sup L)\right)^j \gamma^2 \sigma^2 \\
&= \left(1 - 2\gamma\mu(1 - \gamma \sup L)\right)^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + \frac{\gamma \sigma^2}{\mu(1 - \gamma \sup L)}.
\end{aligned}$$

A.3. Proof of Corollary 2.2.

Proof. Recall the main recursive step in the previous proof,

$$(A.2) \quad \mathbb{E}\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2 \leq (1 - 2\mu\gamma(1 - \gamma \sup L)) \|\mathbf{x}_k - \mathbf{x}_*\|_2^2 + 2\gamma^2 \sigma^2,$$

which is valid as long as $\gamma < \frac{1}{\sup L}$. The minimal value of the quadratic

$$F_\xi(\gamma) = (1 - 2\gamma\mu(1 - \gamma \sup L)) \xi + 2\sigma^2 \gamma^2$$

is achieved at

$$\gamma_\xi^* = \frac{\mu\xi}{2\xi\mu \sup L + 2\sigma^2},$$

and

$$(A.3) \quad F_\xi(\gamma_\xi^*) = \left(1 - \frac{\mu^2\xi}{2\mu \sup L\xi + 2\sigma^2}\right)\xi$$

Note that indeed $\gamma_\xi^* < \frac{1}{\sup L}$. Thus if we choose step-size $\gamma^* = \gamma_\varepsilon^*$,

$$\begin{aligned} \mathbb{E}\|\mathbf{x}_{k+1} - \mathbf{x}_\star\|_2^2 &\leq F_{\|\mathbf{x}_k - \mathbf{x}_\star\|_2^2}(\gamma^*) \\ &= \left(F_{\|\mathbf{x}_k - \mathbf{x}_\star\|_2^2}(\gamma^*) - F_\varepsilon(\gamma^*) \right) + F_\varepsilon(\gamma^*) \\ &\leq \left(1 - \frac{\mu^2 \varepsilon}{2\mu \varepsilon \sup L + 2\sigma^2} \right) \|\mathbf{x}_k - \mathbf{x}_\star\|_2^2. \end{aligned}$$

Iterating the expectation,

$$\mathbb{E}\|\mathbf{x}_{k+1} - \mathbf{x}_\star\|_2^2 \leq \left(1 - \frac{\mu^2 \varepsilon}{2\mu \varepsilon \sup L + 2\sigma^2} \right)^k \varepsilon_0.$$

It follows that if $\varepsilon \leq \mathbb{E}\|\mathbf{x}_{k+1} - \mathbf{x}_\star\|_2^2$, then

$$\begin{aligned} \log(\varepsilon/\varepsilon_0) &\leq k \log \left(1 - \frac{\mu^2 \varepsilon}{2\mu \varepsilon \sup L + 2\sigma^2} \right) \\ &\leq -k \left(\frac{\mu^2 \varepsilon}{2\mu \sup L \varepsilon + 2\sigma^2} \right) \end{aligned}$$

or, equivalently

$$\begin{aligned} k &\leq \log(\varepsilon_0/\varepsilon) \left(\frac{2\mu \sup L \varepsilon + 2\sigma^2}{\mu^2 \varepsilon} \right) \\ &= \log(\varepsilon_0/\varepsilon) \left(\frac{2 \sup L}{\mu} + \frac{2\sigma^2}{\mu^2 \varepsilon} \right). \end{aligned}$$

□