

Faster Convergence of Stochastic Accelerated Gradient Descent under Interpolation

Aaron Mishkin^{1*}, Mert Pilanci² and Mark Schmidt^{3,4}

¹Department of Computer Science, Stanford University.

²Department of Electrical Engineering, Stanford University.

³Department of Computer Science, University of British Columbia.

⁴Canada CIFAR AI Chair, Amii.

*Corresponding author(s). E-mail(s): amishkin@cs.stanford.edu;

Contributing authors: pilanci@stanford.edu; schmidtm@cs.ubc.ca;

Abstract

We prove new convergence rates for a generalized version of stochastic Nesterov acceleration under interpolation conditions. Unlike previous analyses, our approach accelerates any stochastic gradient method which makes sufficient progress in expectation. The proof, which proceeds using the estimating sequences framework, applies to both convex and strongly convex functions and is easily specialized to accelerated SGD under the strong growth condition. In this special case, our analysis reduces the dependence on the strong growth constant from ρ to $\sqrt{\rho}$ as compared to prior work. This improvement is comparable to a square-root of the condition number in the worst case and address criticism that guarantees for stochastic acceleration could be worse than those for SGD.

1 Introduction

A continuing trend in machine learning is the adoption of powerful prediction models which can exactly fit, or *interpolate*, their training data (Zhang et al., 2017). Methods such as over-parameterized neural networks (Zhang and Yin, 2013; Belkin et al., 2019a), kernel machines (Belkin et al., 2019b), and boosting (Schapire et al., 1997) have all been shown to achieve zero training loss in practice. This phenomena is particularly prevalent in modern deep learning, where interpolation is conjectured to be key to both optimization (Liu et al., 2022; Oymak and Soltanolkotabi, 2019) and generalization (Belkin, 2021).

Recent experimental and theoretical evidence shows stochastic gradient descent (SGD) matches the fast convergence rates of deterministic gradient methods up to problem-dependent constants when training interpolating models (Arora et al., 2018; Ma et al., 2018; Zou and Gu, 2019). With additional assumptions, interpolation also implies the strong (Polyak, 1987) and weak (Bassily et al., 2018; Vaswani et al., 2019) growth conditions, which bound the second moment of the stochastic gradients. Under strong/weak growth, variance-reduced algorithms typically exhibit slower convergence than stochastic gradient methods despite using more computation or memory (Defazio and Bottou, 2019; Ma et al., 2018), perhaps because these conditions already imply a form of “automatic variance reduction” (Liu et al., 2022). A combination of interpolation and growth conditions has been used to prove fast convergence rates for SGD with line-search (Vaswani et al., 2019), with the stochastic Polyak step-size (Loizou et al., 2020; Berrada et al., 2020), for mirror descent (D’Orazio et al., 2021), and for model-based methods (Asi and Duchi, 2019).

While these results show interpolation is sufficient to break the $\Omega(\epsilon^{-4})$ complexity barrier for computing stationary points of smooth, convex functions with stochastic, first-order oracles (Arjevani et al., 2019), significantly less work has been done to obtain the accelerated rates possible in the deterministic setting (Nemirovsky and Nesterov, 1985). Vaswani et al. (2019) analyze a stochastic version of Nesterov’s accelerated gradient method (AGD) (Nesterov, 1983) under the strong growth condition, but their bounds have a linear dependence on the strong growth constant and can be slower than SGD (Liu and Belkin, 2020). In contrast, Liu and Belkin (2020) propose a modified version of stochastic AGD and extend the statistical condition number approach of Jain et al. (2018) to the interpolation setting. However, their results apply primarily to quadratics and are not accelerated for general convex functions.

In this work, we apply the estimating sequences analysis developed by Nesterov (2004) to the interpolation setting. Our approach hinges on a simple, in-expectation progress guarantee for SGD, which we prove is a sufficient condition for generic acceleration of stochastic algorithms. This proof technique is completely different from that used by Vaswani et al. (2019) and yields an improved dependence on the strong growth constant. In the worst-case, the improvement is proportional to the square-root of the condition number and guarantees stochastic AGD is always at least as fast as SGD. In what follows, all proofs are deferred to the corresponding section of the appendix.

1.1 Additional Related Work

A large literature on stochastic optimization under interpolation rapidly developed following the seminal work by Bassily et al. (2018) and Vaswani et al. (2019). For instance, Xiao et al. (2022) analyze Frank-Wolfe under interpolation, Vaswani et al. (2020) prove fast convergence for Adagrad-type methods (Duchi et al., 2011), and Meng et al. (2020) show fast rates for sub-sampled Newton method under interpolation. Interpolation has also been used to study last-iterate convergence of SGD (Varre et al., 2021) and subgradient methods (Fang et al., 2021).

Interpolation is a key sufficient condition for growth conditions, which are a set of general assumptions controlling the magnitude of the noise in stochastic gradients. Strong growth was introduced by Polyak (1987, Section 4.2.5), who called the condition *relative random noise* and used it to prove q -linear convergence of SGD. Solodov (1998) and Tseng (1998) later used a variation of strong growth to analyze incremental gradient methods; their variation was also used by Schmidt and Le Roux (2013) to prove linear convergence of SGD with a constant step-size for strongly-convex functions. Vaswani et al. (2019) returned to the original definition given by Polyak (1987) and used it to analyze stochastic AGD, as we do in this paper.

Many authors have tried to accelerate stochastic gradient methods, including under a variety of growth conditions. By accelerate, we mean improve the dependence on the condition number or improve the order of convergence (i.e. from $O(1/k)$ to $O(1/k^2)$) as compared to SGD. For example, Schmidt et al. (2011) establish orders of growth on the gradient noise which still permit stochastic proximal-gradient methods to be accelerated. In contrast, d’Aspremont (2008) and Devolder et al. (2014) assume bounded, deterministic gradient errors to derive convergence rates, while Cohen et al. (2018) develop a noise-resistant acceleration scheme. Most recently, Chen et al. (2020) analyze stochastic AGD under expected smoothness, but their rate only holds under interpolation when the strong growth constant is less than two.

As discussed above, Liu and Belkin (2020) extend the approach of Jain et al. (2018) to the interpolation setting. Their assumptions imply strong growth and the analysis is limited to least-squares problems, although similar rates have been obtained for continuized AGD applied to kernel least-squares (Even et al., 2021). Valls et al. (2022) take a different view and extend the work by Vaswani et al. (2019) to constrained optimization. Unfortunately, none of these algorithms are necessarily accelerated and both Assran and Rabbat (2020) and Liu and Belkin (2020) prove that stochastic AGD may not obtain accelerated rates of convergence even under interpolation. We address this criticism and make detailed comparisons between convergence rates in Section 4.

2 Assumptions

Consider the unconstrained minimization of a smooth, convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. We assume f has at least one minimizer w^* and denote the optimal set by \mathcal{W}^* . At each iteration k , we sample a stochastic gradient $\nabla f(w_k, z_k)$ such that,

$$\mathbb{E}_{z_k} [\nabla f(w_k, z_k)] = \nabla f(w_k),$$

meaning the stochastic gradient is unbiased. We assume that z_k, z_j are independent for $k \neq j$, but they do not need to have the same distribution. The stochastic gradients satisfy the strong growth condition when there exists $\rho \geq 1$ for which

$$\mathbb{E}_{z_k} [\|\nabla f(w_k, z_k)\|_2^2] \leq \rho \|\nabla f(w_k)\|_2^2, \quad (1)$$

holds given any sequence $\{w_k, z_k\}$. We say that interpolation is satisfied if

$$\nabla f(w) = 0 \implies \nabla f(w, z_k) = 0,$$

for all z_k . That is, stationarity of f implies stationarity of the stochastic gradients. Although the strong growth condition implies interpolation, we will see that the converse requires further assumptions on f and on the stochastic gradients.

We assume that f is L -smooth, meaning ∇f is L -Lipschitz continuous. L -smoothness of f implies the following quadratic upper-bound for all $w, u \in \mathbb{R}^d$ (Nesterov, 2004):

$$f(u) \leq f(w) + \langle \nabla f(w), u - w \rangle + \frac{L}{2} \|u - w\|_2^2. \quad (2)$$

Similarly, we will sometimes require the stochastic gradients to be L_{\max} -individually smooth, meaning they satisfy

$$f(u, z_k) \leq f(w, z_k) + \langle \nabla f(w, z_k), u - w \rangle + \frac{L_{\max}}{2} \|u - w\|_2^2, \quad (3)$$

almost surely for all k . We also assume that f is μ -strongly convex, by which we mean

$$f(u) \geq f(w) + \langle \nabla f(w), u - w \rangle + \frac{\mu}{2} \|u - w\|_2^2, \quad (4)$$

holds for all $w, u \in \mathbb{R}^d$ and some $\mu \geq 0$. When $\mu = 0$, strong convexity reduces to convexity of f . If f is strongly convex with $\mu > 0$, the stochastic gradients are L_{\max} -individually smooth, and interpolation holds, then the strong growth constant is bounded as $\rho \leq L_{\max}/\mu$ (see Lemma 13). Recalling that the ratio $\kappa = L/\mu$ is the condition number of f , we see that the strong-growth constant is bounded by a quantity like the condition number of the worst-conditioned stochastic function.¹

3 Convergence of Stochastic AGD

Our analysis in this section builds on the estimating sequences approach developed by Nesterov (1988). However, we consider variable step-sizes η_k and allow the iterates w_k to be set using a general update scheme. The procedure takes $\gamma_0 > 0$ as input and

¹Note that $f(u, z_k)$ is typically not strongly convex, so this analogy is not formal.

uses the following updates, which we call *generalized stochastic AGD*:

$$\begin{aligned}
\alpha_k^2 &= \eta_k(1 - \alpha_k)\gamma_k + \eta_k\alpha_k\mu \\
\gamma_{k+1} &= (1 - \alpha_k)\gamma_k + \alpha_k\mu \\
y_k &= \frac{1}{\gamma_k + \alpha_k\mu} [\alpha_k\gamma_kv_k + \gamma_{k+1}w_k], \\
w_{k+1} &= m(\eta_k, y_k, \nabla f(y_k, z_k)) \\
v_{k+1} &= \frac{1}{\gamma_{k+1}} [(1 - \alpha_k)\gamma_kv_k + \alpha_k\mu y_k - \alpha_k\nabla f(y_k)].
\end{aligned} \tag{5}$$

where m is an as-yet unspecified update for the “primal sequence” w_k and $v_0 = w_0$ (which implies $y_0 = w_0$). Note that the step-size η_k is required at the start of the iteration to compute α_k . As a result, local search methods like the Armijo line-search (Armijo, 1966) must re-evaluate γ_{k+1} , y_k , and w_{k+1} for each candidate step-size.

Choosing m to be one step of SGD with step-size η_k yields the familiar updates of the standard version of stochastic AGD (see Nesterov (2004, Eq. 2.2.20)):

$$\begin{aligned}
w_{k+1} &= w_k - \eta_k\nabla f(y_k, z_k) \\
\alpha_{k+1}^2 &= (1 - \alpha_{k+1})\alpha_k^2\frac{\eta_{k+1}}{\eta_k} + \eta_{k+1}\alpha_{k+1}\mu \\
y_{k+1} &= w_{k+1} + \frac{\alpha_k(1 - \alpha_k)}{\alpha_k^2 + \alpha_{k+1}}(w_{k+1} - w_k).
\end{aligned} \tag{6}$$

Our approach hinges on the fact that, under strong growth, stochastic gradient updates for w_k give a similar per-step progress condition as deterministic gradient descent. Since descent in the primal step is the only link between w_k and the “dual sequence” y_k , generalized stochastic AGD with any primal update obtaining similar per-iteration progress can be analyzed in the same fashion as stochastic AGD. This allows us to derive a fast convergence rate for the general scheme in Eq. (5).

We start by deriving the progress condition for SGD. It is straightforward to prove the following bound using L -smoothness, the definition of w_{k+1} , and strong growth:

Lemma 1. *Suppose f is L -smooth, the strong growth condition holds, and η_k is independent of z_k . Then the stochastic gradient step in Eq. (6) makes progress as,*

$$\mathbb{E}_{z_k} [f(w_{k+1})] \leq f(y_k) - \eta_k\left(1 - \frac{\eta_k\rho L}{2}\right)\|\nabla f(y_k)\|_2^2. \tag{7}$$

Substituting any fixed step-size $\eta_k \leq 1/\rho L$ into Eq. (7) gives the following equation, which we call the *expected progress condition*:

$$\mathbb{E}_{z_k} [f(w_{k+1})] \leq f(y_k) - \frac{\eta_k}{2}\|\nabla f(y_k)\|_2^2, \tag{8}$$

which is equivalent to the progress made by gradient descent with step-size $\eta_k \leq 1/L$ up to a factor of ρ (Bertsekas, 1997). In order to make use of the expected progress condition, we now introduce the estimating sequences framework.

Definition 2 (Estimating Sequences). *Two sequences λ_k, ϕ_k are estimating sequences if the following hold almost surely: (i) $\lambda_k \geq 0$ (ii) $\lim \lambda_k = 0$, and (iii) for all $w \in \mathbb{R}^d$,*

$$\phi_k(w) \leq (1 - \lambda_k)f(w) + \lambda_k\phi_0(w). \quad (9)$$

Unlike the standard definition, we permit λ_k and ϕ_k to depend on z_0, \dots, z_{k-1} , making them random variables. Typically the first function ϕ_0 is deterministic and chosen to satisfy $\phi_0(w) \geq f(w)$ for all w near w_0 . Since Eq. (9) guarantees $\lim_k \phi_k(w) \leq f(w)$, ϕ_k can be interpreted as a sequence of relaxing upper-bounds, where the rate of relaxation is controlled by λ_k . The next condition captures when ϕ_k is a good local model of f .

Definition 3. *The local upper-bound property holds in expectation if*

$$\mathbb{E}_{z_0, \dots, z_{k-1}} [f(w_k)] \leq \mathbb{E}_{z_0, \dots, z_{k-1}} \left[\inf_u \phi_k(u) \right]. \quad (10)$$

In what follows, we use \mathbb{E} without subscripts to denote the total expectation with respect to z_0, \dots, z_{k-1} . That is, all randomness in the procedure up to iteration k . If ϕ_k maintains the local upper-bound property in expectation, then Eq. (9) guarantees

$$\begin{aligned} \mathbb{E}[f(w_k)] &\leq \mathbb{E} \left[\inf_u \phi_k(u) \right] \leq \mathbb{E}[\phi_k(w^*)] \leq \mathbb{E}[(1 - \lambda_k)f(w^*) + \lambda_k\phi_0(w^*)] \\ &\implies \mathbb{E}[f(w_k)] - f(w^*) \leq \mathbb{E}[\lambda_k(\phi_0(w^*) - f(w^*))], \end{aligned} \quad (11)$$

which shows that generalized stochastic AGD converges in expectation at a rate controlled by λ_k . As a result, the bulk of our analysis focuses on establishing the local upper-bound property for a suitable choice of estimating sequences.

Following Nesterov (2004), choose $\lambda_0 = 1$, $\phi_0(w) = f(w_0) + \frac{\gamma_0}{2} \|w - w_0\|_2^2$, and

$$\begin{aligned} \lambda_{k+1} &= (1 - \alpha_k)\lambda_k \\ \phi_{k+1}(w) &= (1 - \alpha_k)\phi_k(w) + \alpha_k \left(f(y_k) + \langle \nabla f(y_k), w - y_k \rangle + \frac{\mu}{2} \|w - y_k\|_2^2 \right). \end{aligned} \quad (12)$$

The initial curvature γ_0 is an input parameter; differentiating shows that v_{k+1} is actually the minimizer of ϕ_{k+1} , while $\nabla^2 \phi_{k+1} = \gamma_{k+1}I$ (Lemma 14). Thus, the auxiliary sequences γ_k, v_k can be viewed as arising from our choice of local model. The next lemma proves these are valid estimating sequences when the step-size sequence is well-behaved. In what follows, we use the convention $1/0 = \infty$ to cover the case of non-strongly convex functions.

Lemma 4. *Assume f is μ -strongly convex with $\mu \geq 0$ and $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely. Then λ_k and ϕ_k given in Eq. (9) are estimating sequences.*

The parameter $\eta_{\min} > 0$ is required for λ_k to decrease sufficiently fast, while the upper-bound $\eta_k \leq 1/\mu$ is only necessary when $\mu > 0$. In this case, it guarantees $\lambda_k \geq 0$. This choice of estimating sequences also satisfies the local error-bound property in expectation when $m(\eta_k, y_k, \nabla f(y_k, z_k))$ matches the progress of fixed step-size SGD.

Proposition 5. *If f is L -smooth and μ -strongly convex with $\mu \geq 0$, $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely for every $k \in \mathbb{N}$, and the primal update $m(\eta_k, y_k, \nabla f(y_k, z_k))$ satisfies the sufficient progress condition (Eq. (8)), then ϕ_k has the local upper-bound property in expectation. That is, for every $k \in \mathbb{N}$,*

$$\mathbb{E}[f(w_k)] \leq \mathbb{E}\left[\inf_u \phi_k(u)\right].$$

Proposition 5 is our main theoretical contribution and immediately leads to two accelerated convergence rates for the generalized stochastic AGD scheme.

Theorem 6. *Suppose f is L -smooth and μ -strongly convex with $\mu > 0$, $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely for every $k \in \mathbb{N}$, and the primal update $m(\eta_k, y_k, \nabla f(y_k, z_k))$ satisfies the sufficient progress condition (Eq. (8)). If $\gamma_0 = \mu$, then generalized stochastic AGD has the following rate of convergence:*

$$\begin{aligned} \mathbb{E}[f(w_{k+1})] - f(w^*) &\leq \mathbb{E}\left[\prod_{i=0}^k (1 - \sqrt{\eta_k \mu})\right] \left[f(w_0) - f(w^*) + \frac{\mu}{2}\|w_0 - w^*\|_2^2\right] \\ &\leq (1 - \sqrt{\eta_{\min} \mu})^{k+1} \left[f(w_0) - f(w^*) + \frac{\mu}{2}\|w_0 - w^*\|_2^2\right]. \end{aligned} \quad (13)$$

This linear rate of convergence requires knowledge of the strong convexity constant in order to set γ_0 . However, we can still obtain a $O(1/k^2)$ rate without knowing μ so long as the smoothness constant can be estimated. The following theorem is a generalization of [Nesterov \(2004\)](#) to stochastic optimization under interpolation.

Theorem 7. *Suppose f is L -smooth and μ -strongly convex with $\mu \geq 0$, $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely for every $k \in \mathbb{N}$, and the primal update $m(\eta_k, y_k, \nabla f(y_k, z_k))$ satisfies the sufficient progress condition (Eq. (8)). If $\gamma_0 \in (\mu, 3/\eta_{\min})$, then generalized stochastic AGD has the following rate of convergence:*

$$\mathbb{E}[f(w_{k+1})] - f(w^*) \leq \frac{4}{\eta_{\min}(\gamma_0 - \mu)(k+1)^2} \left[f(w_0) - f(w^*) + \frac{\gamma_0}{2}\|w_0 - w^*\|_2^2\right]. \quad (14)$$

3.1 Specializations

Theorems 6 and 7 provide accelerated guarantees for any stochastic primal update $m(\eta_k, y_k, \nabla f(y_k, z_k))$ satisfying the sufficient progress condition. Assuming strong growth holds, we may specialize m to fixed step-size SGD with $\eta_k = 1/\rho L$ (sufficient progress is satisfied according to Lemma 1). This yields the standard version of stochastic AGD analyzed by [Vaswani et al. \(2019\)](#). However, instantiating our convergence guarantees shows a faster rate for stochastic AGD with an improved dependence on ρ .

Corollary 8. *If f is L -smooth and μ -strongly convex with $\mu > 0$, strong growth holds, and $\eta_k = 1/\rho L$, then stochastic AGD with $\gamma_0 = \mu$ converges as,*

$$\mathbb{E}[f(w_{k+1})] - f(w^*) \leq \left(1 - \sqrt{\frac{\mu}{\rho L}}\right)^{k+1} \left[f(w_0) - f(w^*) + \frac{\mu}{2}\|w_0 - w^*\|_2^2\right]. \quad (15)$$

Alternatively, if $\mu \geq 0$ and $\gamma_0 \in (\mu, 3\rho L)$, then stochastic AGD satisfies,

$$f(w_{k+1}) - f(w^*) \leq \frac{4\rho L}{(\gamma_0 - \mu)(k+1)^2} \left[f(w_0) - f(w^*) + \frac{\gamma_0}{2}\|w_0 - w^*\|_2^2\right]. \quad (16)$$

Corollary 8 shows that SGD can be accelerated to obtain the same convergence rate as deterministic AGD up to a factor of ρ . We emphasize that some dependence on ρ cannot be avoided; generic acceleration of SGD is not possible, even in the interpolation setting (Assran and Rabbat, 2020), so convergence rates must incorporate some measure of hardness due to stochasticity. In the next section, we compare our convergence guarantees against other results from the literature and give simple conditions under which acceleration is achieved.

An advantage of our analysis is that it also extends to more complex methods, such as SGD with full matrix preconditioning,

$$w_{k+1} = w_k - \eta_k D_k^{-1} \nabla f(w_k, z_k), \quad (17)$$

where $D_k \in \mathbb{R}^{d \times d}$ is a positive-definite matrix and η_k is a step-size sequence. We say that matrix strong growth holds in the norm $\|x\|_{D_k}^2 = x^\top D_k x$ with constant ρ_{D_k} if

$$\mathbb{E}_{z_k} \left[\|\nabla f(w, z_k)\|_{D_k^{-1}}^2 \right] \leq \rho_{D_k} \|\nabla f(w)\|_{D_k}^2. \quad (18)$$

If interpolation holds, $\nabla f(u, z_k)$ are $L_{\max}^{D_k}$ -individually smooth in $\|\cdot\|_{D_k}$, and f is μ_{D_k} -strongly convex in $\|\cdot\|_{D_k}$, then $\rho_{D_k} \leq L_{\max}^{D_k}/\mu_{D_k}$ (see Lemma 16) and the following convergence rate is obtained by combining Lemma 17 with our main convergence theorems for generalized stochastic AGD.

Corollary 9. *Assume f is μ -strongly convex and $0 \prec D_k \preceq I$ for every $k \in \mathbb{N}$. Suppose f is L_{D_k} -smooth, ρ_{D_k} matrix strong growth holds, and $\eta_k = 1/\rho_{D_k} L_{D_k}$. Let $C_\infty = \sup_k \{\rho_{D_k} L_{D_k}\}$. If $\gamma_0 = \mu > 0$, then stochastic preconditioned AGD satisfies,*

$$\mathbb{E}[f(w_{k+1})] - f(w^*) \leq \prod_{i=0}^k \left(1 - \sqrt{\frac{\mu}{\rho_{D_i} L_{D_i}}}\right) \left[f(w_0) - f(w^*) + \frac{\mu}{2}\|w_0 - w^*\|_2^2\right]. \quad (19)$$

Alternatively, if $\mu \geq 0$ and $\gamma_0 \in (\mu, 3C_\infty)$, then we obtain,

$$f(w_{k+1}) - f(w^*) \leq \frac{4C_\infty}{(\gamma_0 - \mu)(k+1)^2} \left[f(w_0) - f(w^*) + \frac{\gamma_0}{2}\|w_0 - w^*\|_2^2\right]. \quad (20)$$

Assumptions	SGD	S-AGD (Ours)	S-AGD (VSB)	MaSS
Strongly Convex	$O\left(\frac{L_{\max}}{\mu} \log\left(\frac{1}{\epsilon}\right)\right)$	$O\left(\sqrt{\frac{\rho L}{\mu}} \log\left(\frac{1}{\epsilon}\right)\right)$	$O\left(\rho \sqrt{\frac{L}{\mu}} \log\left(\frac{1}{\epsilon}\right)\right)$	$\sqrt{\kappa \tilde{\kappa}} \log\left(\frac{1}{\epsilon}\right)$
Convex	$O\left(\frac{L_{\max}}{\epsilon}\right)$	$O\left(\sqrt{\frac{\rho L}{\epsilon}}\right)$	$O\left(\rho \sqrt{\frac{L}{\epsilon}}\right)$	N/A

Table 1 Comparison of iteration complexities for stochastic acceleration schemes under strong growth and individual smoothness. VSB indicates Vaswani et al. (2019) and MaSS is the modified stochastic AGD iteration proposed by Liu and Belkin (2020). The strongly-convex rate for MaSS applies only to quadratics; although MaSS has a convergence guarantee for convex functions, we omit it here because it relies on a hard-to-interpret assumption and is not accelerated.

Compared to standard stochastic AGD, preconditioning allows us to measure stochasticity in the norm induced by D_k . This is advantageous when $\rho_{D_k} L_{D_k} \leq \rho L$, meaning f is smoother and/or the stochastic gradients are better conditioned in $\|\cdot\|_{D_k}$ than in the standard Euclidean norm. In such a setting, our theory suggests preconditioning is a simple way to further speed-up accelerated stochastic optimization.

4 Comparison to Existing Rates

Now we compare our rates to those existing in the literature. Throughout this section, we assume that f is individually smooth, interpolation holds, and the strong growth condition is satisfied. Recall that the strong growth constant is bounded above as $\rho \leq L_{\max}/\mu$ under these conditions. This worst-case bound on ρ is critical to understanding when stochastic AGD does or does not accelerate.

Before proceeding, we introduce the notion of statistical condition number proposed by Jain et al. (2018) and used by Liu and Belkin (2020) to analyze their modified version of stochastic AGD (called MaSS) in the least-squares setting. Let \mathcal{P} be a probability distribution over (x, y) and define the least squares objective as

$$f_{\text{ls}}(w) = \mathbb{E}_{(x,y) \sim \mathcal{P}} \left[\frac{1}{2} (x^\top w - y)^2 \right]. \quad (21)$$

Define the stochastic functions and gradients to be $f_{\text{ls}}(w, z_k) = \frac{1}{2} (x_{z_k}^\top w - y_{z_k})^2$ and $\nabla f_{\text{ls}}(w, z_k) = x_{z_k} (x_{z_k}^\top w - y_{z_k})$, where $(x_{z_k}, y_{z_k}) \sim \mathcal{P}$. These stochastic gradients are L_{\max} -individually smooth with $L_{\max} = \sup \{\|x_{z_k}\|_2^2 : (x_{z_k}, y_{z_k}) \in \text{Supp}(\mathcal{P})\}$. Assuming we may interchange expectation and differentiation, the Hessian is $H = \mathbb{E}_x [xx^\top]$ and the condition number κ and statistical condition number $\tilde{\kappa}$ are defined as,

$$\kappa = \inf \{t/\mu : \mathbb{E}_x [\|x\|_2^2 (xx^\top)] \preceq tH\}, \quad \tilde{\kappa} = \inf \{t : \mathbb{E}_x [\|x\|_{H^{-1}}^2 (xx^\top)] \preceq tH\}. \quad (22)$$

It is straightforward to prove $\kappa, \tilde{\kappa} \leq L_{\max}/\mu$, similar to the strong growth constant.

Table 1 compares our iteration complexities for stochastic AGD to the complexity of SGD under interpolation, the analysis of stochastic AGD by Vaswani et al. (2019), and the complexity of MaSS. Unlike Vaswani et al. (2019), who use strong growth to show both the optimality gap and distance to a minimizer decrease in expectation at each

iteration, our approach only requires the sufficient progress condition. This allows us to shrink the dependence on the strong growth constant from ρ to $\sqrt{\rho}$, which — since $\rho \leq L_{\max}/\mu$ — can be larger than $\sqrt{\kappa}$ in the worst case. Substituting this into the complexity bound shows stochastic AGD requires $O((\sqrt{LL_{\max}}/\mu) \log(1/\epsilon))$ iterations to reach ϵ -sub-optimality. That is, stochastic AGD is always at least as fast as SGD and faster when $L_{\max} \gg L$.

Our convergence rate for stochastic AGD also improves over that for SGD under the strong growth condition (Schmidt and Le Roux, 2013). The improvement is by a factor $\sqrt{\rho/\mu}$, indicating that acceleration actually shrinks the dependence on the noise level. This is quite different from results in the general stochastic setting, where accelerated methods are typically more sensitive to noise (Honorio, 2012). For example, Schmidt et al. (2011) show that the noise level must decrease faster for accelerated methods to converge compared to (proximal) SGD when interpolation does not hold. We conclude that interpolation seems to be key when proving fast rates for stochastic AGD.

Comparing our results against MaSS is more difficult due to the dependence on $\tilde{\kappa}$. To understand the difference in convergence rates, we consider two finite-sum example problems. In what follows, let e_1, \dots, e_n be the standard basis for \mathbb{R}^n .

Example 10 ($L_{\max} \gg L$). Consider the least-squares problem setting in Eq. (21) and choose $y = 0$, $x \sim \text{Uniform}(e_1, \dots, e_n)$. A short calculation shows $L_{\max} = 1$, $\rho = n$, $L = \mu = 1/n$, and $\kappa = \tilde{\kappa} = n$. As a result, stochastic AGD and MaSS have the following complexity bounds:

$$\text{S-AGD: } O\left(\sqrt{n} \log\left(\frac{1}{\epsilon}\right)\right) \quad \text{vs} \quad \text{MaSS: } O\left(n \log\left(\frac{1}{\epsilon}\right)\right) \quad (23)$$

As expected, stochastic AGD accelerates due to the gap between the smoothness and individual smoothness constants. In comparison, the complexity bound for MaSS is not accelerated and only matches that for SGD. The next example considers the opposite setting, where $L \approx L_{\max}$ and we do not expect stochastic AGD to be faster than SGD.

Example 11 ($L_{\max} \approx L$). Consider the least-squares problem setting in Eq. (21). Let $y = 0$ and x be distributed as follows: $P(x = e_1) = 1 - 1/n$ and $P(x = e_2) = 1/n$. It is straightforward to show that $L_{\max} = 1$, $\mu = 1/n$, and $L = (n - 1)/n$, while $\rho = n$ and $\tilde{\kappa} = \kappa = n$. As a result, the complexity estimates for stochastic AGD and MaSS are,

$$\text{S-AGD: } O\left(\sqrt{n(n-1)} \log\left(\frac{1}{\epsilon}\right)\right) \quad \text{vs} \quad \text{MaSS: } O\left(n \log\left(\frac{1}{\epsilon}\right)\right). \quad (24)$$

As $n \rightarrow \infty$, stochastic AGD attains the same complexity as SGD and is not accelerated. In comparison, the guarantee for MaSS always matches SGD and is slower than stochastic AGD for every finite n . We conclude that while both methods are restricted by lower bounds on stochastic acceleration, AGD can accelerate on some simple problems where MaSS fails.

5 Conclusion

We derive new convergence rates for a generalized version of stochastic Nesterov acceleration. Our approach extends the estimating sequences framework to the stochastic setting and shows that any update scheme making sufficient progress in expectation can be accelerated. As this sufficient progress condition is satisfied by SGD under the strong growth condition, our proof immediately specializes to give fast rates for stochastic AGD. Compared to previous work, our convergence bounds improve the dependence on the strong growth constant from ρ to $\sqrt{\rho}$. This improvement can be larger than the square-root of the condition number, shows stochastic AGD is at least as fast as SGD, and explains the strong empirical performance of stochastic acceleration shown by Vaswani et al. (2019). We also leverage our generalized algorithm to prove convergence guarantees for stochastic AGD with preconditioning. In particular, we show that preconditioning further speeds-up accelerated SGD when the stochastic gradients are small in the matrix norm induced by the preconditioner.

In addition to these results, the utility of our theoretical approach is further demonstrated by recent literature. Our core result for stochastic AGD (Proposition 5) was previously made available in a master’s thesis (Mishkin, 2020) and the proof technique has since been leveraged to give optimal bounds for stochastic acceleration in the general setting (Vaswani et al., 2022). Yet, several questions remain unanswered. For example, the convergence of stochastic AGD under relaxed conditions, like weak growth or with a stochastic line-search (Vaswani et al., 2019), has not been proved. And while our generalized AGD scheme also suggests accelerating methods like stochastic proximal-point, establishing the expected progress condition appears difficult and new insights may be required. We leave these questions to future work.

Acknowledgments. We would like to thank Frederik Kunstner, Victor Sanchez-Portella, and Sharan Vaswani for many insightful discussions.

Funding. Aaron Mishkin was supported by NSF GRF Grant No. DGE-1656518 and by NSERC PGS D Grant No. PGSD3-547242-2020. Mert Pilanci was supported by the NSF under Grant ECCS-2037304 and Grant DMS-2134248, by an NSF CAREER Award under Grant CCF-2236829, by the U.S. Army Research Office Early Career Award under Grant W911NF-21-1-0242, and by the Stanford Precourt Institute. Mark Schmidt was partially supported by the Canada CIFAR AI Chair Program and NSERC Discovery Grant No. RGPIN-2022-03669.

Data Availability. This paper does not rely any data, public or private. Derivations for all theoretical results are included in the appendix.

References

- Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires rethinking generalization. In: 5th International Conference on Learning Representations, ICLR 2017. OpenReview.net (2017)
- Zhang, H., Yin, W.: Gradient methods for convex minimization: Better rates under weaker conditions. arXiv preprint arXiv:1303.4645 (2013)
- Belkin, M., Hsu, D., Ma, S., Mandal, S.: Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences* **116**(32), 15849–15854 (2019)
- Belkin, M., Rakhlin, A., Tsybakov, A.B.: Does data interpolation contradict statistical optimality? In: Chaudhuri, K., Sugiyama, M. (eds.) *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019. Proceedings of Machine Learning Research*, vol. 89, pp. 1611–1619. PMLR (2019)
- Schapire, R.E., Freund, Y., Barlett, P., Lee, W.S.: Boosting the margin: A new explanation for the effectiveness of voting methods. In: Fisher, D.H. (ed.) *Proceedings of the Fourteenth International Conference on Machine Learning (ICML 1997)*, pp. 322–330. Morgan Kaufmann (1997)
- Liu, C., Zhu, L., Belkin, M.: Loss landscapes and optimization in over-parameterized non-linear systems and neural networks. *Applied and Computational Harmonic Analysis* **59**, 85–116 (2022)
- Oymak, S., Soltanolkotabi, M.: Overparameterized nonlinear learning: Gradient descent takes the shortest path? In: Chaudhuri, K., Salakhutdinov, R. (eds.) *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA. Proceedings of Machine Learning Research*, vol. 97, pp. 4951–4960. PMLR (2019)
- Belkin, M.: Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation. *Acta Numer.* **30**, 203–248 (2021)
- Arora, S., Cohen, N., Hazan, E.: On the optimization of deep networks: Implicit acceleration by overparameterization. In: Dy, J.G., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning, ICML 2018. Proceedings of Machine Learning Research*, vol. 80, pp. 244–253. PMLR (2018)
- Ma, S., Bassily, R., Belkin, M.: The power of interpolation: Understanding the effectiveness of SGD in modern over-parametrized learning. In: Dy, J.G., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning, ICML 2018. Proceedings of Machine Learning Research*, vol. 80, pp. 3331–3340. PMLR (2018)
- Zou, D., Gu, Q.: An improved analysis of training over-parameterized deep neural

- networks. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32: NeurIPS 2019*, pp. 2053–2062 (2019)
- Polyak, B.T.: *Introduction to optimization* (1987)
- Bassily, R., Belkin, M., Ma, S.: On exponential convergence of SGD in non-convex over-parametrized learning. *arXiv preprint arXiv:1811.02564* (2018)
- Vaswani, S., Mishkin, A., Laradji, I.H., Schmidt, M., Gidel, G., Lacoste-Julien, S.: Painless stochastic gradient: Interpolation, line-search, and convergence rates. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32: NeurIPS 2019*, pp. 3727–3740 (2019)
- Defazio, A., Bottou, L.: On the ineffectiveness of variance reduced optimization for deep learning. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32: NeurIPS 2019*, pp. 1753–1763 (2019)
- Loizou, N., Vaswani, S., Laradji, I., Lacoste-Julien, S.: Stochastic Polyak step-size for SGD: An adaptive learning rate for fast convergence. *arXiv preprint arXiv:2002.10542* (2020)
- Berrada, L., Zisserman, A., Kumar, M.P.: Training neural networks for and by interpolation. In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event. Proceedings of Machine Learning Research*, vol. 119, pp. 799–809. PMLR (2020)
- D'Orazio, R., Loizou, N., Laradji, I.H., Mitliagkas, I.: Stochastic mirror descent: Convergence analysis and adaptive variants via the mirror stochastic polyak stepsize. *CoRR* **abs/2110.15412** (2021)
- Asi, H., Duchi, J.C.: Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization* **29**(3), 2257–2290 (2019)
- Arjevani, Y., Carmon, Y., Duchi, J.C., Foster, D.J., Srebro, N., Woodworth, B.: Lower bounds for non-convex stochastic optimization. *arXiv preprint arXiv:1912.02365* (2019)
- Nemirovsky, A.S., Nesterov, Y.E.: Optimal methods of smooth convex minimization. *USSR Computational Mathematics and Mathematical Physics* **25**(2), 21–30 (1985)
- Vaswani, S., Bach, F., Schmidt, M.W.: Fast and faster convergence of SGD for over-parameterized models and an accelerated perceptron. In: Chaudhuri, K., Sugiyama, M. (eds.) *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019. Proceedings of Machine Learning Research*, vol. 89, pp. 1195–1204.

PMLR (2019)

- Nesterov, Y.: A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In: Doklady an USSR, vol. 269, pp. 543–547 (1983)
- Liu, C., Belkin, M.: Accelerating SGD with momentum for over-parameterized learning. In: 8th International Conference on Learning Representations, ICLR 2020. OpenReview.net (2020)
- Jain, P., Kakade, S.M., Kidambi, R., Netrapalli, P., Sidford, A.: Accelerating stochastic gradient descent for least squares regression. In: Bubeck, S., Perchet, V., Rigollet, P. (eds.) Conference On Learning Theory, COLT 2018. Proceedings of Machine Learning Research, vol. 75, pp. 545–604. PMLR (2018)
- Nesterov, Y.E.: Introductory Lectures on Convex Optimization - A Basic Course. Applied Optimization, vol. 87. Springer (2004)
- Xiao, T., Balasubramanian, K., Ghadimi, S.: Improved complexities for stochastic conditional gradient methods under interpolation-like conditions. Oper. Res. Lett. **50**(2), 184–189 (2022)
- Vaswani, S., Kunstner, F., Laradji, I., Meng, S.Y., Schmidt, M., Lacoste-Julien, S.: Adaptive gradient methods converge faster with over-parameterization (and you can do a line-search). arXiv preprint arXiv:2006.06835 (2020)
- Duchi, J.C., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. J. Mach. Learn. Res. **12**, 2121–2159 (2011)
- Meng, S.Y., Vaswani, S., Laradji, I.H., Schmidt, M., Lacoste-Julien, S.: Fast and furious convergence: Stochastic second order methods under interpolation. In: Chiappa, S., Calandra, R. (eds.) The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020. Proceedings of Machine Learning Research, vol. 108, pp. 1375–1386. PMLR (2020)
- Varre, A.V., Pillaud-Vivien, L., Flammarion, N.: Last iterate convergence of SGD for least-squares in the interpolation regime. In: Ranzato, M., Beygelzimer, A., Dauphin, Y.N., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, Virtual, pp. 21581–21591 (2021)
- Fang, H., Fan, Z., Friedlander, M.P.: Fast convergence of stochastic subgradient method under interpolation. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net (2021)
- Solodov, M.V.: Incremental gradient algorithms with stepsizes bounded away from zero. Comp. Opt. and Appl. **11**(1), 23–35 (1998)

- Tseng, P.: An incremental gradient(-projection) method with momentum term and adaptive stepsize rule. *SIAM Journal on Optimization* **8**(2), 506–531 (1998)
- Schmidt, M., Le Roux, N.: Fast convergence of stochastic gradient descent under a strong growth condition. *arXiv preprint arXiv:1308.6370* (2013)
- Schmidt, M., Le Roux, N., Bach, F.R.: Convergence rates of inexact proximal-gradient methods for convex optimization. In: Shawe-Taylor, J., Zemel, R.S., Bartlett, P.L., Pereira, F.C.N., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 24: NeurIPS 2011*, pp. 1458–1466 (2011)
- d’Aspremont, A.: Smooth optimization with approximate gradient. *SIAM J. Optim.* **19**(3), 1171–1183 (2008)
- Devolder, O., Glineur, F., Nesterov, Y.: First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming* **146**(1-2), 37–75 (2014)
- Cohen, M., Diakonikolas, J., Orecchia, L.: On acceleration with noise-corrupted gradients. In: Dy, J.G., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning, ICML 2018. Proceedings of Machine Learning Research*, vol. 80, pp. 1018–1027. PMLR (2018)
- Chen, Y.-L., Na, S., Kolar, M.: Convergence analysis of accelerated stochastic gradient descent under the growth condition. *arXiv preprint arXiv:2006.06782* (2020)
- Even, M., Berthier, R., Bach, F.R., Flammarion, N., Gaillard, P., Hendriks, H., Mas-soulié, L., Taylor, A.B.: A continuized view on nesterov acceleration for stochastic gradient descent and randomized gossip. *CoRR* **abs/2106.07644** (2021)
- Valls, V., Wang, S., Jiang, Y., Tassiulas, L.: Accelerated convex optimization with stochastic gradients: Generalizing the strong-growth condition. *arXiv preprint arXiv:2207.11833* (2022)
- Assran, M., Rabbat, M.: On the convergence of Nesterov’s accelerated gradient method in stochastic settings. *arXiv preprint arXiv:2002.12414* (2020)
- Nesterov, Y.: On an approach to the construction of optimal methods of minimization of smooth convex functions. *Ekonomika i Mateaticheskie Metody* **24**(3), 509–517 (1988)
- Armijo, L.: Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific Journal of Mathematics* **16**(1), 1–3 (1966)
- Bertsekas, D.P.: Nonlinear programming. *Journal of the Operational Research Society* **48**(3), 334–334 (1997)
- Honorio, J.: Convergence rates of biased stochastic optimization for learning sparse

Ising models. In: Proceedings of the 29th International Conference on Machine Learning, ICML 2012. icml.cc / Omnipress (2012)

Mishkin, A.: Interpolation, growth conditions, and stochastic gradient descent. PhD thesis, University of British Columbia (2020)

Vaswani, S., Dubois-Taine, B., Babanezhad, R.: Towards noise-adaptive, problem-adaptive (accelerated) stochastic gradient descent. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., Sabato, S. (eds.) International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA. Proceedings of Machine Learning Research, vol. 162, pp. 22015–22059. PMLR (2022)

Appendix A Assumptions: Proofs

Lemma 12. *Suppose f is convex and L -smooth, the stochastic gradients $\{\nabla f(w_k, z_k)\}$ are L_{\max} individually smooth, and interpolation holds. Then the weak growth condition holds with constant $\alpha \leq L_{\max}/L$.*

Proof. First, recall that the weak growth condition (Vaswani et al., 2019) is given by

$$\mathbb{E}_{z_k} [\|\nabla f(w_k, z_k)\|_2^2] \leq 2\alpha L (f(w_k) - f(w^*)). \quad (\text{A1})$$

Now, starting from L_{\max} individual-smoothness,

$$f(u, z_k) \leq f(w, z_k) + \langle \nabla f(w, z_k), u - w \rangle + \frac{L_{\max}}{2} \|u - w\|^2,$$

and choosing $u = w - \frac{1}{L_{\max}} \nabla f(w, z_k)$, we obtain

$$\begin{aligned} f(u, z_k) &\leq f(w, z_k) - \frac{1}{L_{\max}} \langle \nabla f(w, z_k), \nabla f(w, z_k) \rangle + \frac{L_{\max}}{2L_{\max}^2} \|\nabla f(w, z_k)\|^2 \\ &= f(w, z_k) - \frac{1}{2L_{\max}} \|\nabla f(w, z_k)\|^2. \end{aligned}$$

Noting that $f(u, z_k) \geq f(w^*, z_k)$ by convexity of f and interpolation and taking expectations with respect to z_k gives the following:

$$\begin{aligned} f(w^*, z_k) &\leq f(w, z_k) - \frac{1}{2L_{\max}} \|\nabla f(w, z_k)\|^2 \\ \implies \mathbb{E}_{z_k} [f(w^*, z_k)] &\leq \mathbb{E}_{z_k} [f(w, z_k)] - \frac{1}{2L_{\max}} \mathbb{E}_{z_k} [\|\nabla f(w, z_k)\|^2] \\ \implies f(w^*) &\leq f(w) - \frac{1}{2L_{\max}} \mathbb{E}_{z_k} [\|\nabla f(w, z_k)\|^2]. \end{aligned}$$

Re-arranging this final equation gives the desired result as follows:

$$\begin{aligned} \mathbb{E}_{z_k} [\|\nabla f(w, z_k)\|^2] &\leq 2L_{\max} (f(w) - f(w^*)) \\ &= 2 \left(\frac{L_{\max}}{L} \right) L (f(w) - f(w^*)). \end{aligned}$$

We conclude that weak growth holds with $\alpha \leq \frac{L_{\max}}{L}$. \square

Lemma 13. *Suppose f is L -smooth and μ -strongly convex, the stochastic gradients $\{\nabla f(w_k, z_k)\}$ are L_{\max} individually smooth, and interpolation holds. Then strong growth holds with constant $\rho \leq L_{\max}/\mu$.*

Proof. Lemma 12 implies that f satisfies the weak growth condition with parameter

$$\alpha \leq \frac{L_{\max}}{L}.$$

Vaswani et al. (2019, Proposition 1) now implies that f satisfies strong growth with parameter

$$\rho \leq \frac{\alpha L}{\mu} \leq \frac{L_{\max}}{\mu}.$$

This concludes the proof. \square

Appendix B Convergence of Stochastic AGD: Proofs

Lemma 1. *Suppose f is L -smooth, the strong growth condition holds, and η_k is independent of z_k . Then the stochastic gradient step in Eq. (6) makes progress as,*

$$\mathbb{E}_{z_k} [f(w_{k+1})] \leq f(y_k) - \eta_k \left(1 - \frac{\eta_k \rho L}{2}\right) \|\nabla f(y_k)\|_2^2. \quad (7)$$

Proof. The proof is a modification of the standard descent lemma. Starting from L -smoothness of f , we obtain

$$\begin{aligned} f(w_{k+1}) &\leq f(y_k) + \langle \nabla f(y_k), w_{k+1} - y_k \rangle + \frac{L}{2} \|w_{k+1} - y_k\|_2^2 \\ &= f(y_k) - \eta_k \langle \nabla f(y_k), \nabla f(y_k, z_k) \rangle + \frac{\eta_k^2 L}{2} \|\nabla f(y_k, z_k)\|_2^2 \end{aligned}$$

Taking expectations with respect to z_k and using the strong growth condition,

$$\begin{aligned} \mathbb{E}_{z_k} [f(w_{k+1})] &\leq f(y_k) - \eta_k \langle \nabla f(y_k), \mathbb{E}_{z_k} [\nabla f(y_k, z_k)] \rangle + \frac{\eta_k^2 L}{2} \mathbb{E}_{z_k} [\|\nabla f(y_k, z_k)\|_2^2] \\ &\leq f(y_k) - \eta_k \|\nabla f(y_k)\|_2^2 + \frac{\eta_k^2 L \rho}{2} \|\nabla f(y_k)\|_2^2 \\ &= f(y_k) - \eta_k \left(1 - \frac{\eta_k L \rho}{2}\right) \|\nabla f(y_k)\|_2^2. \end{aligned}$$

\square

Lemma 14. *The ϕ_k sequence in Eq. (12) satisfies the following canonical form:*

$$\phi_{k+1}(w) = \phi_{k+1}^* + \frac{\gamma_{k+1}}{2} \|w - v_{k+1}\|_2^2, \quad (\text{B2})$$

where the curvature γ_{k+1} , minimizer v_{k+1} , and minimum value ϕ_{k+1}^* are given as follows:

$$\begin{aligned}\gamma_{k+1} &= (1 - \alpha_k)\gamma_k + \alpha_k\mu \\ v_{k+1} &= \frac{1}{\gamma_{k+1}} [(1 - \alpha_k)\gamma_k v_k + \alpha_k\mu y_k - \alpha_k \nabla f(y_k)] \\ \phi_{k+1}^* &= (1 - \alpha_k)\phi_k^* + \alpha_k f(y_k) - \frac{\alpha_k^2}{2\gamma_{k+1}} \|\nabla f(y_k)\|_2^2 \\ &\quad + \frac{\alpha_k(1 - \alpha_k)\gamma_k}{\gamma_{k+1} \left(\frac{\mu}{2} \|y_k - v_k\|_2^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right)}.\end{aligned}\tag{B3}$$

Furthermore, the relationship between γ_{k+1} and α_k is the following:

$$\gamma_{k+1} = \alpha_k^2 / \eta_k.\tag{B4}$$

Proof. The canonical form for ϕ_{k+1} follows directly from [Nesterov \(2004, Lemma 2.2.3\)](#). To see the relationship between γ_{k+1} and α_k , re-arrange the update for α_{k+1} in Eq. (5) to obtain,

$$\frac{\alpha_k^2}{\eta_k} = (1 - \alpha_k)\gamma_k + \alpha_k\mu.\tag{B5}$$

By comparison to the update for γ_k , we deduce that $\gamma_{k+1} = \alpha_k^2 / \eta_k$. \square

Lemma 15. Assume $\alpha_k \in (0, 1)$ and $\eta_{\min} \leq \eta_k \leq 1/\mu$ almost surely for all $k \in \mathbb{N}$. If $\mu > 0$ and $\gamma_0 = \mu$, then

$$\lambda_k \leq \prod_{i=0}^{k-1} (1 - \sqrt{\eta_i \mu}).\tag{B6}$$

Alternately, if $\gamma_0 \in (\mu, \mu + 3/\eta_{\min})$, we obtain

$$\lambda_k \leq \frac{4}{\eta_{\min}(\gamma_0 - \mu)(k+1)^2}.\tag{B7}$$

Proof. Case 1: $\gamma_0 = \mu > 0$. Then $\gamma_k = \mu$ for all k and

$$\begin{aligned}\alpha_k^2 &= (1 - \alpha_k)\eta_k\mu + \alpha_k\eta_k\mu \\ &= \eta_k\mu.\end{aligned}$$

Thus, we deduce that

$$\lambda_k = \prod_{i=0}^{k-1} (1 - \sqrt{\eta_i \mu}),$$

and if $\eta_k > \eta_{\min}$, then $\alpha_k \geq \sqrt{\eta_{\min} \mu}$ and

$$\lambda_k \leq (1 - \sqrt{\eta_{\min} \mu})^k.$$

Case 2: $\gamma_0 \in (\mu, 3L + \mu)$. Using the update rule for γ_k in Lemma 14, we find

$$\gamma_{k+1} - \mu = (1 - \alpha_k)\gamma_k + (\alpha_k - 1)\mu = (1 - \alpha_k)(\gamma_k - \mu)$$

Recurring on this equality implies

$$\gamma_{k+1} = (\gamma_0 - \mu) \prod_{i=0}^k (1 - \alpha_i) = \lambda_{k+1}(\gamma_0 - \mu).$$

Similarly, using $\lambda_{k+1} = (1 - \alpha_k)\lambda_k$ and $\alpha_k^2/\gamma_{k+1} = \eta_k$ yields

$$\begin{aligned} 1 - \frac{\lambda_{k+1}}{\lambda_k} &= \alpha_k = (\gamma_{k+1}\eta_k)^{1/2} \\ &= (\eta_k\mu + \eta_k\lambda_{k+1}(\gamma_0 - \mu))^{1/2} \\ \implies \frac{1}{\lambda_{k+1}} - \frac{1}{\lambda_k} &= \frac{1}{\lambda_{k+1}^{1/2}} \left[\frac{\eta_k\mu}{\lambda_{k+1}} + \eta_k(\gamma_0 - \mu) \right]^{1/2} \\ &\geq \frac{1}{\lambda_{k+1}^{1/2}} \left[\frac{\eta_{\min}\mu}{\lambda_{k+1}} + \eta_{\min}(\gamma_0 - \mu) \right]^{1/2}. \end{aligned}$$

Finally, this implies

$$\begin{aligned} \frac{2}{\lambda_{k+1}^{1/2}} \left(\frac{1}{\lambda_{k+1}^{1/2}} - \frac{1}{\lambda_k^{1/2}} \right) &\geq \left(\frac{1}{\lambda_{k+1}^{1/2}} - \frac{1}{\lambda_k^{1/2}} \right) \left(\frac{1}{\lambda_{k+1}^{1/2}} + \frac{1}{\lambda_k^{1/2}} \right) \\ &\geq \frac{1}{\lambda_{k+1}^{1/2}} \left[\frac{\eta_{\min}\mu}{\lambda_{k+1}} + \eta_{\min}(\gamma_0 - \mu) \right]^{1/2}. \end{aligned}$$

Moreover, this bound holds uniformly for all $k \in \mathbb{N}$. We have now exactly reached Eq. 2.2.11 of Nesterov (2004, Lemma 2.2.4) with L replaced by η_{\min} . Applying that Lemma with this modification, we obtain

$$\lambda_k \leq \frac{4}{\eta_{\min}(\gamma_0 - \mu)(k+1)^2},$$

which completes the proof. □

Lemma 4. *Assume f is μ -strongly convex with $\mu \geq 0$ and $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely. Then λ_k and ϕ_k given in Eq. (9) are estimating sequences.*

Proof. Recalling the update for α_k , we obtain,

$$\alpha_k^2 = (1 - \alpha_k)\gamma_k\eta_k + \alpha_k\eta_k\mu.$$

Define $\hat{L}_k = 1/\eta_k$ to obtain the following quadratic formula,

$$\hat{L}_k \alpha_k^2 + (\gamma_k - \mu) \alpha_k - \gamma_k = 0.$$

Using the quadratic equation, we find

$$\alpha_k = \frac{\mu - \gamma_k \pm \sqrt{(\mu - \gamma_k)^2 + 4\hat{L}_k \gamma_k}}{2\hat{L}_k}.$$

As a result, $\alpha_k > 0$ if and only if

$$(\mu - \gamma_k) + \left((\mu - \gamma_k)^2 + 4\hat{L}_k \gamma_k \right)^{1/2} > 0.$$

If $\mu \geq \gamma_k$, then this holds trivially. Otherwise, we require,

$$(\mu - \gamma_k)^2 + 4\hat{L}_k \gamma_k > (\mu - \gamma_k)^2,$$

which holds if and only if $\eta_k, \gamma_k > 0$. Similarly, $\alpha_k < 1$ if and only if

$$4\hat{L}_k^2 + 4\hat{L}_k(\gamma_k - \mu) + (\mu - \gamma_k)^2 > (\mu - \gamma_k)^2 + 4\hat{L}_k \gamma_k \iff \eta_k < \frac{1}{\mu}.$$

Since this condition holds by assumption, we have $\alpha_k \in (0, 1)$ for all k .

Recall $\lambda_0 = 1$ and $\lambda_{k+1} = (1 - \alpha_k)\lambda_k$. Since $\alpha_k \in (0, 1)$, $\lambda_k \geq 0$ holds by induction. It remains to show that λ_k tends to zero. Invoking Lemma 15 establishes this result almost surely.

Finally, we must show,

$$\phi_k(w) \leq (1 - \lambda_k)f(w) + \lambda_k \phi_0(w).$$

We proceed by induction. Since $\lambda_0 = 1$, we immediately obtain

$$\phi_0(w) = (1 - \lambda_0)f(w) + \lambda_0 \phi_0(w).$$

Now assume the condition holds at ϕ_k ; by the construction of ϕ_{k+1} , we have

$$\begin{aligned} \phi_{k+1}(w) &= (1 - \alpha_k)\phi_k(w) + \alpha_k \left[f(y_k) + \langle \nabla f(y_k), w - y_k \rangle + \frac{\mu}{2} \|w - y_k\| \right] \\ &\leq (1 - \alpha_k)\phi_k(w) + \alpha_k f(w) \\ &\leq (1 - \alpha_k) [(1 - \lambda_k)f(w) + \lambda_k \phi_0(w)] + \alpha_k f(w) \\ &= \lambda_{k+1} \phi_0(w) - \lambda_{k+1} f(w) + (1 - \alpha_k)f(w) + \alpha_k f(w) \\ &= \lambda_{k+1} \phi_0(w) + (1 - \lambda_{k+1})f(w). \end{aligned}$$

This completes the proof. □

Proposition 5. *If f is L -smooth and μ -strongly convex with $\mu \geq 0$, $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely for every $k \in \mathbb{N}$, and the primal update $m(\eta_k, y_k, \nabla f(y_k, z_k))$ satisfies the sufficient progress condition (Eq. (8)), then ϕ_k has the local upper-bound property in expectation. That is, for every $k \in \mathbb{N}$,*

$$\mathbb{E}[f(w_k)] \leq \mathbb{E}\left[\inf_u \phi_k(u)\right].$$

Proof. The choice of $\phi_0^* = f(x_0)$ ensures $\inf \phi_0(w) = f(w_0)$ deterministically, which is the base case for induction. The inductive assumption is $\mathbb{E}[\inf \phi_k(w)] \geq \mathbb{E}[f(w_k)]$; let us use this to show

$$\mathbb{E}\left[\inf_w \phi_{k+1}(w)\right] \geq \mathbb{E}[f(w_{k+1})].$$

Lemma 14 implies that the explicit form of the minimizer $\inf \phi_{k+1}(w) = \phi_{k+1}^*$ is

$$\begin{aligned} \phi_{k+1}^* &= (1 - \alpha_k)\phi_k^* + \alpha_k f(y_k) - \frac{\alpha_k^2}{2\gamma_{k+1}} \|\nabla f(y_k)\|^2 \\ &\quad + \frac{\alpha_k(1 - \alpha_k)\gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right) \end{aligned}$$

Taking expectations with respect to z_0, \dots, z_k and using linearity of expectation:

$$\begin{aligned} \mathbb{E}[\phi_{k+1}^*] &= \mathbb{E}[(1 - \alpha_k)\phi_k^*] + \mathbb{E}\left[\alpha_k f(y_k) - \frac{\alpha_k^2}{2\gamma_{k+1}} \|\nabla f(y_k)\|^2\right] \\ &\quad + \mathbb{E}\left[\frac{\alpha_k(1 - \alpha_k)\gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right)\right] \\ &\geq \mathbb{E}[(1 - \alpha_k)f(w_k)] + \mathbb{E}\left[\alpha_k f(y_k) - \frac{\alpha_k^2}{2\gamma_{k+1}} \|\nabla f(y_k)\|^2\right] \\ &\quad + \mathbb{E}\left[\frac{\alpha_k(1 - \alpha_k)\gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right)\right], \end{aligned}$$

where the inequality follows from the inductive assumption. Convexity of f implies $f(w_k) \geq f(y_k) + \langle \nabla f(y_k), w_k - y_k \rangle$. Recalling $\frac{\alpha_k^2}{\gamma_{k+1}} = \eta_k$ from Lemma 14 allows us to obtain,

$$\begin{aligned} \mathbb{E}[\phi_{k+1}^*] &\geq \mathbb{E}[(1 - \alpha_k)(f(y_k) + \langle \nabla f(y_k), w_k - y_k \rangle)] + \mathbb{E}\left[\alpha_k f(y_k) - \frac{\eta_k}{2} \|\nabla f(y_k)\|^2\right] \\ &\quad + \mathbb{E}\left[\frac{\alpha_k(1 - \alpha_k)\gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right)\right] \\ &= \mathbb{E}[f(y_k)] + \mathbb{E}[(1 - \alpha_k) \langle \nabla f(y_k), w_k - y_k \rangle] - \mathbb{E}\left[\frac{\eta_k}{2} \|\nabla f(y_k)\|^2\right] \\ &\quad + \mathbb{E}\left[\frac{\alpha_k(1 - \alpha_k)\gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right)\right] \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E} \left[f(y_k) - \frac{\eta_k}{2} \|\nabla f(y_k)\|^2 \right] + \mathbb{E} \left[(1 - \alpha_k) \left(\langle \nabla f(y_k), w_k - y_k \rangle \right. \right. \\
&\quad \left. \left. + \frac{\alpha_k \gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right) \right) \right]
\end{aligned}$$

The sufficient progress condition (Eq. (8)) now implies

$$\begin{aligned}
\mathbb{E}[\phi_{k+1}^*] &\geq \mathbb{E}[f(w_{k+1})] + \mathbb{E} \left[(1 - \alpha_k) \left(\langle \nabla f(y_k), w_k - y_k \rangle \right. \right. \\
&\quad \left. \left. + \frac{\alpha_k \gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right) \right) \right].
\end{aligned}$$

The remainder of the proof is largely unchanged from the deterministic case. The definition of y_k gives $w_k - y_k = \frac{\alpha_k \gamma_k}{\gamma_k + \alpha_k \mu} (w_k - v_k)$, which we use to obtain

$$\begin{aligned}
\mathbb{E}[\phi_{k+1}^*] &\geq \mathbb{E}[f(w_{k+1})] + \mathbb{E} \left[(1 - \alpha_k) \left(\frac{\alpha_k \gamma_k}{\gamma_k + \alpha_k \mu} \langle \nabla f(y_k), w_k - v_k \rangle \right. \right. \\
&\quad \left. \left. + \frac{\alpha_k \gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right) \right) \right]
\end{aligned}$$

Noting that $v_k - y_k = \frac{\gamma_{k+1}}{\gamma_k + \alpha_k \mu} (v_k - w_k)$ gives

$$\begin{aligned}
\mathbb{E}[\phi_{k+1}^*] &\geq \mathbb{E}[f(w_{k+1})] + \mathbb{E} \left[(1 - \alpha_k) \left(\frac{\alpha_k \gamma_k}{\gamma_k + \alpha_k \mu} \langle \nabla f(y_k), w_k - v_k \rangle \right. \right. \\
&\quad \left. \left. + \frac{\alpha_k \gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \langle \nabla f(y_k), v_k - y_k \rangle \right) \right) \right] \\
&= \mathbb{E}[f(w_{k+1})] + \mathbb{E} \left[(1 - \alpha_k) \left(\frac{\alpha_k \gamma_k}{\gamma_k + \alpha_k \mu} \langle \nabla f(y_k), w_k - v_k \rangle \right. \right. \\
&\quad \left. \left. + \frac{\alpha_k \gamma_k}{\gamma_{k+1}} \left(\frac{\mu}{2} \|y_k - v_k\|^2 + \frac{\gamma_{k+1}}{\gamma_k + \alpha_k \mu} \langle \nabla f(y_k), v_k - w_k \rangle \right) \right) \right] \\
&= \mathbb{E}[f(w_{k+1})] + \mathbb{E} \left[\frac{\mu \alpha_k (1 - \alpha_k) \gamma_k}{2 \gamma_{k+1}} \|y_k - v_k\|^2 \right] \\
&\geq \mathbb{E}[f(w_{k+1})].
\end{aligned}$$

since $\frac{\mu \alpha_k (1 - \alpha_k) \gamma_k}{2 \gamma_{k+1}} \geq 0$. We conclude that $\mathbb{E}[\inf \phi_k(w)] \geq \mathbb{E}[f(w_k)]$ holds for all $k \in \mathbb{N}$ by induction. \square

Theorem 6. *Suppose f is L -smooth and μ -strongly convex with $\mu > 0$, $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely for every $k \in \mathbb{N}$, and the primal update $m(\eta_k, y_k, \nabla f(y_k, z_k))$ satisfies the sufficient progress condition (Eq. (8)). If $\gamma_0 = \mu$, then generalized stochastic AGD*

has the following rate of convergence:

$$\begin{aligned}\mathbb{E}[f(w_{k+1})] - f(w^*) &\leq \mathbb{E}\left[\prod_{i=0}^k (1 - \sqrt{\eta_k \mu})\right] \left[f(w_0) - f(w^*) + \frac{\mu}{2}\|w_0 - w^*\|_2^2\right] \\ &\leq (1 - \sqrt{\eta_{\min} \mu})^{k+1} \left[f(w_0) - f(w^*) + \frac{\mu}{2}\|w_0 - w^*\|_2^2\right].\end{aligned}\quad (13)$$

Proof. Under the conditions of the theorem, Proposition 5 holds and Eq. (11) implies

$$\mathbb{E}[f(w_k)] - f(w^*) \leq \mathbb{E}[\lambda_k(\phi_0(w^*) - f(w^*))]$$

Since $\gamma_0 = \mu > 0$, Lemma 15 implies $\lambda_k = \prod_{i=0}^{k-1} (1 - \sqrt{\eta_i \mu})$ and we obtain

$$\begin{aligned}\mathbb{E}[f(w_k)] - f(w^*) &\leq \mathbb{E}\left[\prod_{i=0}^{k-1} (1 - \sqrt{\eta_i \mu})\right] (\phi_0(w^*) - f(w^*)) \\ &= \mathbb{E}\left[\prod_{i=0}^{k-1} (1 - \sqrt{\eta_i \mu})\right] \left[f(w_0) - f(w^*) + \frac{\mu}{2}\|w_0 - w^*\|_2^2\right],\end{aligned}$$

which completes the proof. \square

Theorem 7. *Suppose f is L -smooth and μ -strongly convex with $\mu \geq 0$, $\eta_{\min} \leq \eta_k < 1/\mu$ almost surely for every $k \in \mathbb{N}$, and the primal update $m(\eta_k, y_k, \nabla f(y_k, z_k))$ satisfies the sufficient progress condition (Eq. (8)). If $\gamma_0 \in (\mu, 3/\eta_{\min})$, then generalized stochastic AGD has the following rate of convergence:*

$$\mathbb{E}[f(w_{k+1})] - f(w^*) \leq \frac{4}{\eta_{\min}(\gamma_0 - \mu)(k+1)^2} \left[f(w_0) - f(w^*) + \frac{\gamma_0}{2}\|w_0 - w^*\|_2^2\right]. \quad (14)$$

Proof. Under the conditions of the theorem, Proposition 5 holds and Eq. (11) implies

$$\mathbb{E}[f(w_k)] - f(w^*) \leq \mathbb{E}[\lambda_k(\phi_0(w^*) - f(w^*))]$$

Since $\gamma_0 \in (\mu, \mu + 3/\eta_{\min})$, Lemma 15 implies $\lambda_k \leq 4/\eta_{\min}(\gamma_0 - \mu)(k+1)^2$ and we obtain

$$\begin{aligned}\mathbb{E}[f(w_k)] - f(w^*) &\leq \frac{4}{\eta_{\min}(\gamma_0 - \mu)(k+1)^2} (\phi_0(w^*) - f(w^*)) \\ &= \frac{4}{\eta_{\min}(\gamma_0 - \mu)(k+1)^2} \left[f(w_0) - f(w^*) + \frac{\gamma_0}{2}\|w_0 - w^*\|_2^2\right],\end{aligned}$$

which completes the proof. \square

B.1 Specializations: Proofs

Lemma 16. *Let $D \in \mathbb{R}^{d \times d}$ be independent of z_k . Suppose f is convex and L -smooth, the stochastic gradients $\nabla f(w_k, z_k)$ are L_{\max}^D individually smooth with respect to the matrix norm $\|\cdot\|_D$, and interpolation holds. If f is also μ_D -strongly convex with respect to $\|\cdot\|_D$, then*

$$\mathbb{E}_{z_k} [\|\nabla f(w, z_k)\|_{D^{-1}}^2] \leq \frac{L_{\max}^D}{\mu_D} \|\nabla f(w)\|_{D^{-1}}^2. \quad (\text{B8})$$

That is, strong growth holds in $\|\cdot\|_D$ with constant $\rho_D \leq \frac{L_{\max}^D}{\mu_D}$.

Proof. Starting from L_{\max}^D individual-smoothness,

$$f(u, z_k) \leq f(w, z_k) + \langle \nabla f(w, z_k), u - w \rangle + \frac{L_{\max}^D}{2} \|u - w\|_D^2,$$

and choosing $u = w - \frac{1}{L_{\max}^D} D^{-1} \nabla f(w, z_k)$, we obtain

$$\begin{aligned} f(u, z_k) &\leq f(w, z_k) - \frac{1}{L_{\max}^D} \langle \nabla f(w, z_k), D^{-1} \nabla f(w, z_k) \rangle + \frac{1}{2L_{\max}^D} \|D^{-1} \nabla f(w, z_k)\|_D^2 \\ &= f(w, z_k) - \frac{1}{2L_{\max}^D} \|\nabla f(w, z_k)\|_{D^{-1}}^2. \end{aligned}$$

Noting that $f(u, z_k) \geq f(w^*, z_k)$ by convexity of f and interpolation and taking expectations with respect to z_k gives the following:

$$\begin{aligned} f(w^*, z_k) &\leq f(w, z_k) - \frac{1}{2L_{\max}^D} \|\nabla f(w, z_k)\|_{D^{-1}}^2 \\ \implies \mathbb{E}_{z_k} [f(w^*, z_k)] &\leq \mathbb{E}_{z_k} [f(w, z_k)] - \frac{1}{2L_{\max}^D} \mathbb{E}_{z_k} [\|\nabla f(w, z_k)\|_{D^{-1}}^2] \\ \implies f(w^*) &\leq f(w) - \frac{1}{2L_{\max}^D} \mathbb{E}_{z_k} [\|\nabla f(w, z_k)\|_{D^{-1}}^2]. \end{aligned}$$

Re-arranging this final equation and using μ_D -strong convexity gives the desired result,

$$\begin{aligned} \mathbb{E}_{z_k} [\|\nabla f(w, z_k)\|_{D^{-1}}^2] &\leq 2L_{\max}^D (f(w) - f(w^*)) \\ &\leq \frac{L_{\max}^D}{\mu_D} \|\nabla f(w)\|_{D^{-1}}^2. \end{aligned}$$

□

Lemma 17. *Let $D \in \mathbb{R}^{d \times d}$. Assume the f is both L_D -smooth and satisfies ρ_D strong growth in the matrix norm $\|\cdot\|_D$. If $0 \prec D \leq I$ and $\eta_k \leq \frac{1}{\rho_D L_D}$, then preconditioned SGD satisfies,*

$$\mathbb{E}_{z_k} [f(w_{k+1})] \leq f(y_k) - \frac{\eta_k}{2} \|\nabla f(y_k)\|^2.$$

Proof. Starting from smoothness in $\|\cdot\|_D$,

$$\begin{aligned} f(w_{k+1}) &\leq f(y_k) + \langle \nabla f(y_k), w_{k+1} - y_k \rangle + \frac{L_D}{2} \|w_{k+1} - y_k\|_D^2 \\ &\leq f(y_k) - \eta_k \langle \nabla f(y_k), D^{-1} \nabla f(y_k, z_k) \rangle + \frac{\eta_k^2 L_D}{2} \|\nabla f(y_k, z_k)\|_{D^{-1}}^2. \end{aligned}$$

Taking expectations with respect to z_k ,

$$\begin{aligned} \implies \mathbb{E}_{z_k} [f(w_{k+1})] &\leq f(y_k) - \eta_k \langle \nabla f(y_k), D_k^{-1} \nabla f(y_k) \rangle + \frac{\eta_k^2 L_D}{2} \mathbb{E}_{z_k} [\|\nabla f(y_k, z_k)\|_{D^{-1}}^2] \\ &\leq f(y_k) - \eta_k \langle \nabla f(y_k), D_k^{-1} \nabla f(y_k) \rangle + \frac{\eta_k^2 \rho_D L_D}{2} \|\nabla f(y_k)\|_{D^{-1}}^2 \\ &= f(y_k) - \eta_k \left(1 - \frac{\eta_k \rho_D L_D}{2}\right) \|\nabla f(y_k)\|_{D^{-1}}^2 \\ &\leq f(y_k) - \frac{\eta_k}{2} \|\nabla f(y_k)\|_{D^{-1}}^2 \\ &\leq f(y_k) - \frac{\eta_k}{2} \|\nabla f(y_k)\|_2^2. \end{aligned}$$

□

Appendix C Comparison to Existing Rates: Proofs

Example 10 ($L_{\max} \gg L$). Consider the least-squares problem setting in Eq. (21) and choose $y = 0$, $x \sim \text{Uniform}(e_1, \dots, e_n)$. A short calculation shows $L_{\max} = 1$, $\rho = n$, $L = \mu = 1/n$, and $\kappa = \tilde{\kappa} = n$. As a result, stochastic AGD and MaSS have the following complexity bounds:

$$\text{S-AGD: } O\left(\sqrt{n} \log\left(\frac{1}{\epsilon}\right)\right) \quad \text{vs} \quad \text{MaSS: } O\left(n \log\left(\frac{1}{\epsilon}\right)\right) \quad (23)$$

Proof. It is easy to see that $L_{\max} = 1$. Taking expectations, we find that

$$f_{\text{ls}}(w) = \frac{1}{2n} \|w\|_2^2,$$

which implies $L = \mu = 1/n$. It is also straightforward to compute ρ :

$$\mathbb{E} [\|\nabla f_{\text{ls}}(w, z_k)\|_2^2] = \frac{1}{n} \sum_{i=1}^n \|e_i(e_i^\top w)\|_2^2 = \frac{1}{n} \sum_{i=1}^n w_i^2 = \frac{1}{n} \|w\|_2^2 = n \|\nabla f_{\text{ls}}(w)\|_2^2,$$

which implies $\rho = n$. Note that this is tight with the bound $\rho \leq L_{\max}/\mu$.

Now we compute the values of κ and $\tilde{\kappa}$. Since the Hessian satisfies $H = I/n$, it is straightforward to see that

$$\begin{aligned}\mathbb{E}_x [\|x\|_2^2(xx^\top)] &= \frac{1}{n} \sum_{i=1}^n \|e_i\|_2^2 e_i e_i^\top = H \\ \mathbb{E}_x [\|x\|_{H^{-1}}^2(xx^\top)] &= \frac{1}{n} \sum_{i=1}^n \|e_i\|_{H^{-1}}^2 e_i e_i^\top = \sum_{i=1}^n \|e_i\|_2^2 e_i e_i^\top = nH,\end{aligned}$$

which implies that $\kappa = \tilde{\kappa} = n$. Substituting these values into the complexity bounds for stochastic AGD and MaSS completes the example. \square

Example 11 ($L_{\max} \approx L$). Consider the least-squares problem setting in Eq. (21). Let $y = 0$ and x be distributed as follows: $P(x = e_1) = 1 - 1/n$ and $P(x = e_2) = 1/n$. It is straightforward to show that $L_{\max} = 1$, $\mu = 1/n$, and $L = (n - 1)/n$, while $\rho = n$ and $\tilde{\kappa} = \kappa = n$. As a result, the complexity estimates for stochastic AGD and MaSS are,

$$\text{S-AGD: } O\left(\sqrt{n(n-1)} \log\left(\frac{1}{\epsilon}\right)\right) \quad \text{vs} \quad \text{MaSS: } O\left(n \log\left(\frac{1}{\epsilon}\right)\right). \quad (24)$$

Proof. Again, it is easy to see that $L_{\max} = 1$. Taking expectations, we find that

$$f_{\text{ls}}(w) = \frac{n-1}{2n} w_1^2 + \frac{1}{2n} w_2^2,$$

which implies $L = \frac{n-1}{n}$ and $\mu = \frac{1}{n}$. The strong growth constant is given by

$$\begin{aligned}\mathbb{E} [\|\nabla f(w, z_k)\|_2^2] &= \frac{n-1}{n} w_1^2 + \frac{1}{n} w_2^2 \\ &= n \left(\frac{n-1}{n^2} w_1^2 + \frac{1}{n^2} w_2^2 \right) \\ &\leq n \left(\frac{(n-1)^2}{n^2} w_1^2 + \frac{1}{n^2} w_2^2 \right) \\ &= n \|\nabla f(w)\|_2^2,\end{aligned}$$

which implies $\rho \leq n$. It's easy to see that this is tight by taking $w_1 = 0$ and $w_2 \neq 0$.

Now we compute the values of κ and $\tilde{\kappa}$. The Hessian is given by

$$H = \begin{bmatrix} \frac{n-1}{n} & 0 \\ 0 & \frac{1}{n} \end{bmatrix},$$

and thus

$$\mathbb{E}_x [\|x\|_2^2(xx^\top)] = \frac{n-1}{n} e_1 e_1^\top + \frac{1}{n} e_2 e_2^\top = H$$

$$\mathbb{E}_x [\|x\|_{H^{-1}}^2(xx^\top)] = e_1 e_1^\top + e_2 e_2^\top = I \preceq nH,$$

which implies $\tilde{\kappa} = \kappa = n$. Substituting these values into the complexity bounds for stochastic AGD and MaSS completes the example. \square