

Adaptive gradient descent without descent

Yura Malitsky*

Konstantin Mishchenko[†]

Abstract

We present a strikingly simple proof that two rules are sufficient to automate gradient descent: 1) don't increase the stepsize too fast and 2) don't overstep the local curvature. No need for functional values, no line search, no information about the function except for the gradients. By following these rules, you get a method adaptive to the local geometry, with convergence guarantees depending only on smoothness in a neighborhood of a solution. Given that the problem is convex, our method will converge even if the global smoothness constant is infinity. As an illustration, it can minimize arbitrary continuously twice-differentiable convex function. We examine its performance on a range of convex and nonconvex problems, including matrix factorization and training of ResNet-18.

1 Introduction

Since the early days of optimization it was evident that there is a need for algorithms that are as independent from the user as possible. First-order methods have proven to be versatile and efficient in a wide range of applications, but one drawback has been present all that time: the stepsize. Despite some certain success stories, line search procedures and adaptive online methods have not removed the need to manually tune the optimization parameters. Even in smooth convex optimization, which is often believed to be much simpler than the nonconvex counterpart, robust rules for stepsize selection have been elusive. The purpose of this work is to remedy this deficiency.

The problem formulation that we consider is the basic unconstrained optimization problem

$$\min_{x \in \mathbb{R}^d} f(x), \quad (1)$$

where $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is a differentiable function. Throughout the paper we assume that (1) has a solution and we denote its optimal value by f_* .

The simplest and most known approach to this problem is the gradient descent method (GD), whose origin can be traced back to Cauchy [7, 20]. Although it is probably the oldest optimization method, it continues to play a central role in modern algorithmic theory and applications. Its definition can be written in a mere one line,

$$x^{k+1} = x^k - \lambda \nabla f(x^k), \quad k \geq 0, \quad (2)$$

where $x^0 \in \mathbb{R}^d$ is arbitrary and $\lambda > 0$. Under assumptions that f is convex and L -smooth¹, that is

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|, \quad \forall x, y, \quad (3)$$

one can show that GD with $\lambda \in (0, \frac{2}{L})$ converges to an optimal solution [30]. Moreover, with $\lambda = \frac{1}{L}$ the convergence rate [10] is

$$f(x^k) - f_* \leq \frac{L\|x^0 - x^*\|^2}{2(2k+1)}, \quad (4)$$

*Laboratory for Information and Inference Systems, EPFL, Lausanne, Switzerland, yurii.malitskyi@epfl.ch

[†]KAUST, Thuwal, Saudi Arabia, konstantin.mishchenko@kaust.edu.sa

¹Alternatively, we will say that ∇f is L -Lipschitz.

where x^* is any solution of (1). Note that this bound is not improvable [10].

We identify four important challenges that limit the applications of gradient descent even in the convex case:

1. GD is not general: many functions do not satisfy (3) globally.
2. GD is not a free lunch: one needs to guess λ , potentially trying many values before a success.
3. GD is not robust: failing to provide $\lambda < \frac{2}{L}$ may lead to divergence.
4. GD is slow: even if L is finite, it might be arbitrarily larger than local smoothness.

Certain ways to address some of the issues above already exist in the literature. They include line search, adaptive Polyak’s stepsize, mirror descent, dual preconditioning and stepsize estimation for subgradient methods. We discuss them one by one below, in a process reminiscent of cutting off Hydra’s limbs: if one issue is fixed, two others take its place.

If our goal is to decrease the function value robustly, without user input and by a significant amount, why don’t we just estimate the best stepsize at each iteration? This idea has been around for decades under the name line search (or backtracking) and, indeed, it remains the most practical and generic solution to tackle the aforementioned issues. This direction of research started from the seminal works of Goldstein [13] and Armijo [1] and continues to attract attention, see [5, 35] and references therein.

In general, at each iteration the line search executes another subroutine with additional evaluations of ∇f and/or f until some condition is met, potentially exploding the cost of a single iteration. Naturally, this makes complexity bounds no longer informative.

At the same time, the famous Polyak’s stepsize [31] stands out as a very fast alternative to gradient descent. Furthermore, it does not depend on the global smoothness constant and uses the current gradient to estimate the geometry. The formula might look deceitfully simple, $\lambda_k = \frac{f(x^k) - f_*}{\|\nabla f(x^k)\|^2}$, but there is a catch: it is rarely possible to know f_* . This method, again, requires the user to guess f_* . What is more, with λ it was fine to underestimate it by a factor of 10, but the guess for f_* must be tight, otherwise it has to be reestimated later [15].

Seemingly no issue is present in the Barzilai-Borwein [3] stepsize. Motivated by the quasi-Newton schemes, the authors of [3] suggested using steps

$$\lambda_k = \frac{\langle x^k - x^{k-1}, \nabla f(x^k) - \nabla f(x^{k-1}) \rangle}{\|\nabla f(x^k) - \nabla f(x^{k-1})\|^2}.$$

Alas, the convergence results regarding this choice of λ_k are very limited, only the case with a quadratic f is well understood [9, 32], otherwise one needs to add again a line search [33], see also the counter-example for constrained minimization in [8].

Other more interesting ways to deal with non-Lipschitzness of ∇f use the problem structure. The first method, pioneered recently in [4], shows that the mirror descent method, which is another extension of GD, can be used with a fixed stepsize, whenever f satisfies a certain generalization of (3). In the other paper [21], the authors proposed the dual preconditioning method—also a refined version of GD. Similarly to the former technique, it also goes beyond the standard smoothness assumption of f , but in a different way. Unfortunately, these two simple and elegant approaches cannot resolve all issues yet. First, not many functions fulfill respective generalized conditions. And second, both methods still get us back to the problem of not knowing the allowed range of stepsizes.

A whole branch of optimization considers adaptive extensions of GD that deal with functions whose (sub)gradients are bounded. Probably the earliest work in that direction was written by Shor [36]. He showed that the method

$$x^{k+1} = x^k - \lambda_k \frac{g^k}{\|g^k\|},$$

where $g^k \in \partial f(x^k)$ is a subgradient, converges for properly chosen sequences (λ_k) , see, e.g., Section 3.2.3 in [26]. Moreover, λ_k requires no knowledge about the function whatsoever.

Similar methods that work in online setting such as Adagrad [11, 23] received a lot of attention in recent years and remain an active topic of research [41]. Methods similar to Adagrad—Adam [18, 34], RMSprop [40] and Adadelta [43]—remain state-of-the-art for training neural networks. The corresponding objective is usually neither smooth nor convex, and the theory often assumes Lipschitzness of the function rather than of the gradients. Therefore, this direction of research is mostly orthogonal to ours, although we do compare with some of these methods in our ResNet-18 experiment.

We also note that without momentum Adam and RMSprop reduce to signSGD [6], which is known to be non-convergent for arbitrary stepsizes on a simple quadratic problem [16]. This issue seems to have something to do with pathological dependence on gradients noise, see the bounds in [6, 42].

In a close relation to ours is the recent work [22], where there was proposed an adaptive golden ratio algorithm for monotone variational inequalities. As it solves a more general problem, it does not exploit the structure of (1) and, as most variational inequality methods, has a more conservative update. Although the method estimates the smoothness, it still requires an upper bound on the stepsize as input.

Contribution. We propose a new version of GD that at no cost resolves all aforementioned issues. The idea is simple, and it is surprising that it has not been yet discovered. In each iteration we choose λ_k as a certain approximation of the inverse local Lipschitz constant. With such a choice, we prove that convexity and local smoothness of f and sufficient for convergence of iterates with the complexity $\mathcal{O}(1/k)$ for $f(x^k) - f_*$ in the worst case.

Discussion. Let us now briefly discuss why we believe that proofs based on monotonicity and global smoothness lead to slower methods.

Gradient descent is by far not a recent method, so there have been obtained optimal rates of convergence. However, we argue that adaptive methods require rethinking optimality of the stepsizes. Take as an example a simple quadratic problem, $f(x, y) = \frac{1}{2}x^2 + \frac{\delta}{2}y^2$, where $\delta \ll 1$. Clearly, the smoothness constant of this problem is equal to $L = 1$ and the strong convexity one is $\mu = \delta$. If we run GD from an arbitrary point (x^0, y^0) with the “optimal” stepsize $\lambda = \frac{1}{L} = 1$, then one iteration of GD gives us $(x^1, y^1) = (0, (1 - \delta)y^0)$, and similarly $(x^k, y^k) = (0, (1 - \delta)^k y^0)$. Evidently for δ small enough we might converge to the solution $(0, 0)$ quite long. Instead GD would converge in two iterations, if it adjusts its step after the first iteration to $\lambda = \frac{1}{\delta}$.

Nevertheless, all existing analyses of the gradient descent with L -smooth f use stepsizes bounded by $2/L$. In addition, functional analysis gives

$$f(x^{k+1}) \leq f(x^k) - \lambda \left(1 - \frac{\lambda L}{2}\right) \|\nabla f(x^k)\|^2,$$

from which $1/L$ can be seen as the “optimal” stepsize. Alternatively, we can assume that f is μ -strongly convex, and the analysis in norms gives

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &\leq \left(1 - 2\frac{\lambda L \mu}{L + \mu}\right) \|x^k - x^*\|^2 \\ &\quad - \lambda \left(\frac{2}{L + \mu} - \lambda\right) \|\nabla f(x^k)\|^2, \end{aligned}$$

whence the “optimal” step is $\frac{2}{L + \mu}$.

Finally, line search procedures use some certain type of monotonicity, for instance ensuring that $f(x^{k+1}) \leq f(x^k) - c \|\nabla f(x^k)\|^2$ for some $c > 0$. We break with this tradition and merely ask for convergence in the end.

Algorithm 1 Adaptive gradient descent

```
1:  $x^0 \in \mathbb{R}^d$ ,  $\lambda_0 > 0$ ,  $\theta_0 = +\infty$ ,  $x^1 = x^0 - \lambda_0 \nabla f(x^0)$ 
2: for  $k = 1, 2, \dots$  do
3:    $\lambda_k = \min \left\{ \sqrt{1 + \theta_{k-1} \lambda_{k-1}}, \frac{\|x^k - x^{k-1}\|}{2\|\nabla f(x^k) - \nabla f(x^{k-1})\|} \right\}$ 
4:    $x^{k+1} = x^k - \lambda_k \nabla f(x^k)$ 
5:    $\theta_k = \frac{\lambda_k}{\lambda_{k-1}}$ 
6: end for
```

2 Main part

2.1 Local smoothness of f

Recall that the mapping is locally Lipschitz if it is Lipschitz over any bounded set. It is natural to ask whether some interesting functions are smooth locally, but not globally.

It turns out there is no shortage of examples, most prominently among highly nonlinear functions. In \mathbb{R} , they include $x \mapsto \exp(x)$, $\log(x)$, $\tan(x)$, x^p , for $p > 1$, etc. More generally, they include any twice differentiable f , since $\nabla^2 f(x)$, as a continuous mapping, is bounded over any bounded set \mathcal{C} . In this case, we have that ∇f is Lipschitz on \mathcal{C} , due to the mean value inequality

$$\|\nabla f(x) - \nabla f(y)\| \leq \max_{z \in \mathcal{C}} \|\nabla^2 f(z)\| \|x - y\|, \quad \forall x, y \in \mathcal{C}.$$

Algorithm 1 we propose is just a slight modification of the GD. The quick explanation why local Lipschitzness of ∇f does not cause us any problems, unlike to most other methods, lies in the way we prove its convergence. Whenever the stepsize λ_k satisfies two inequalities²

$$\begin{cases} \lambda_k^2 & \leq (1 + \theta_{k-1})\lambda_{k-1}^2, \\ \lambda_k & \leq \frac{\|x^k - x^{k-1}\|}{2\|\nabla f(x^k) - \nabla f(x^{k-1})\|}, \end{cases}$$

independently of the properties of f (apart from convexity), we can show that the iterates (x^k) remain bounded. Here and everywhere else we use the convention $1/0 = +\infty$, so if $\nabla f(x^k) - \nabla f(x^{k-1}) = 0$, the second inequality can be ignored.

In the introduction we stated that we are interested in algorithms as independent from the user as possible. However, Algorithm 1 still needs x^0 and λ_0 as input. This is not an issue, as one can simply fix $x^0 = 0$ and $\lambda_0 = 10^{-10}$. Equipped with a tiny λ_0 , we ensure that x^1 will be close enough to x^0 and likely will give a good estimate for λ_1 . Otherwise, this has no influence on the further steps.

2.2 Analysis without descent

It is now time to show our main contribution, the new analysis technique. The tools that we are going to use are the well-known Cauchy-Schwarz and convexity inequalities. In addition, our methods are related to potential functions [38], which is a powerful tool for producing tight bounds for GD.

Another divergence from the common practice is that our main lemma includes not only x^{k+1} and x^k , but also x^{k-1} . This can be seen as a two-step analysis, while the majority of optimization methods have one-step bounds. However, as we want to adapt to the local

²It can be shown that instead of the second condition it is enough to ask for $\lambda_k^2 \leq \frac{\|x^k - x^{k-1}\|^2}{[3\|\nabla f(x^k)\|^2 - 4\langle \nabla f(x^k), \nabla f(x^{k-1}) \rangle]_+}$, where $[a]_+ \stackrel{\text{def}}{=} \max\{0, a\}$, but we prefer the option written in the main text for its simplicity.

geometry of our objective, it is rather natural to have two terms to capture the change in the gradients.

Now, it is time to derive a characteristic inequality for a specific Lyapunov energy.

Lemma 1. Let x^* be any solution of (1). Then for (x^k) generated by Algorithm 1 it holds

$$\begin{aligned} \|x^{k+1} - x^*\|^2 + \frac{1}{2}\|x^{k+1} - x^k\|^2 + 2\lambda_k(1 + \theta_k)(f(x^k) - f_*) \\ \leq \|x^k - x^*\|^2 + \frac{1}{2}\|x^k - x^{k-1}\|^2 + 2\lambda_k\theta_k(f(x^{k-1}) - f_*). \end{aligned} \quad (5)$$

Proof. Let $k \geq 1$. We start from the standard way of analyzing GD:

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &= \|x^k - x^*\|^2 + 2\langle x^{k+1} - x^k, x^k - x^* \rangle + \|x^{k+1} - x^k\|^2 \\ &= \|x^k - x^*\|^2 - 2\lambda_k\langle \nabla f(x^k), x^k - x^* \rangle + \|x^{k+1} - x^k\|^2. \end{aligned}$$

As usually, we can bound the scalar product by convexity of f :

$$-2\lambda_k\langle \nabla f(x^k), x^k - x^* \rangle \leq -2\lambda_k(f(x^k) - f_*), \quad (6)$$

which gives us

$$\|x^{k+1} - x^*\|^2 \leq \|x^k - x^*\|^2 - 2\lambda_k(f(x^k) - f_*) + \|x^{k+1} - x^k\|^2. \quad (7)$$

These two steps have been repeated thousands of times, but now we continue in a completely different manner. We have precisely one “bad” term in (7), which is $\|x^{k+1} - x^k\|^2$. The reason it appears lies in the fact that gradients change from point to point, so we will bound it using the difference of gradients:

$$\begin{aligned} \|x^{k+1} - x^k\|^2 &= 2\|x^{k+1} - x^k\|^2 - \|x^{k+1} - x^k\|^2 \\ &= -2\lambda_k\langle \nabla f(x^k), x^{k+1} - x^k \rangle - \|x^{k+1} - x^k\|^2 \\ &= 2\lambda_k\langle \nabla f(x^k) - \nabla f(x^{k-1}), x^k - x^{k+1} \rangle \\ &\quad + 2\lambda_k\langle \nabla f(x^{k-1}), x^k - x^{k+1} \rangle - \|x^{k+1} - x^k\|^2. \end{aligned} \quad (8)$$

Let us estimate the first two terms in the right-hand side above. First, definition of λ_k , followed by Cauchy-Schwarz and Young’s inequalities, yields

$$\begin{aligned} 2\lambda_k\langle \nabla f(x^k) - \nabla f(x^{k-1}), x^k - x^{k+1} \rangle &\leq 2\lambda_k\|\nabla f(x^k) - \nabla f(x^{k-1})\|\|x^k - x^{k+1}\| \\ &\leq \|x^k - x^{k-1}\|\|x^k - x^{k+1}\| \\ &\leq \frac{1}{2}\|x^k - x^{k-1}\|^2 + \frac{1}{2}\|x^{k+1} - x^k\|^2. \end{aligned} \quad (9)$$

Second, by convexity of f ,

$$\begin{aligned} 2\lambda_k\langle \nabla f(x^{k-1}), x^k - x^{k+1} \rangle &= \frac{2\lambda_k}{\lambda_{k-1}}\langle x^{k-1} - x^k, x^k - x^{k+1} \rangle \\ &= 2\lambda_k\theta_k\langle x^{k-1} - x^k, \nabla f(x^k) \rangle \leq 2\lambda_k\theta_k(f(x^{k-1}) - f(x^k)). \end{aligned} \quad (10)$$

Plugging (9) and (10) in (8), we obtain

$$\|x^{k+1} - x^k\|^2 \leq \frac{1}{2}\|x^k - x^{k-1}\|^2 - \frac{1}{2}\|x^{k+1} - x^k\|^2 + 2\lambda_k\theta_k(f(x^{k-1}) - f(x^k)).$$

Finally, using the produced estimate for $\|x^{k+1} - x^k\|^2$ in (7), we deduce the desired inequality (5). \square

The above lemma already might give a good hint why our method works. From inequality (5) together with condition $\lambda_k^2 \leq (1 + \theta_{k-1})\lambda_{k-1}^2$, we obtain that the Lyapunov energy—the left-hand side of (5)—is decreasing. This gives us boundedness of (x^k) , which is often the key ingredient for proving convergence. In the next theorem we formally state our result.

Theorem 1. *Suppose that $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is convex with locally Lipschitz gradient ∇f . Then (x^k) generated by Algorithm 1 converges to a solution of (1) and we have that*

$$f(\hat{x}^k) - f_* \leq \frac{D}{2S_k} = \mathcal{O}\left(\frac{1}{k}\right),$$

where $\hat{x}^k \stackrel{\text{def}}{=} \frac{\sum_{i=1}^k \lambda_i x^i + \lambda_1 \theta_1 x^1}{S_k}$ and $S_k \stackrel{\text{def}}{=} \sum_{i=1}^k \lambda_i + \lambda_1 \theta_1$ and D is a constant that explicitly depends on the initial data and the solution set.

Our proof will consist of two parts. The first one is a straightforward application of Lemma 1, from which we derive boundedness of (x^k) and complexity result. Due to its conciseness, we provide it directly after this remark. In the second part, we prove that the whole sequences (x^k) converges to a solution. Surprisingly, this part is a bit more technical than expected, and thus we postpone it to the appendix.

Proof. (Boundedness and complexity result).

Fix any x^* from the solution set of eq. (1). Telescoping inequality (5), we deduce

$$\begin{aligned} & \|x^{k+1} - x^*\|^2 + \frac{1}{2} \|x^{k+1} - x^k\|^2 + 2\lambda_k(1 + \theta_k)(f(x^k) - f_*) \\ & \quad + 2 \sum_{i=1}^{k-1} [\lambda_i(1 + \theta_i) - \lambda_{i+1}\theta_{i+1}](f(x^i) - f_*) \\ & \leq \|x^1 - x^*\|^2 + \frac{1}{2} \|x^1 - x^0\|^2 + 2\lambda_1\theta_1[f(x^0) - f_*] \stackrel{\text{def}}{=} D. \end{aligned} \tag{11}$$

Note that by definition of λ_k , the second line above is always nonnegative. Thus, the sequence (x^k) is bounded. Since ∇f is locally Lipschitz, it is Lipschitz continuous on bounded sets. It means that for the bounded set³ $\mathcal{C} = \overline{\text{conv}}\{x^*, x^0, x^1, \dots\}$ there exists $L > 0$ such that

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \quad \forall x, y \in \mathcal{C}.$$

Clearly, $\lambda_1 = \frac{\|x^1 - x^0\|}{2\|\nabla f(x^1) - \nabla f(x^0)\|} \geq \frac{1}{2L}$, thus, by induction one can prove that $\lambda_k \geq \frac{1}{2L}$, in other words, the sequence (λ_k) is separated from zero.

Now we want to apply the Jensen inequality for the sum of all terms $f(x^i) - f_*$ in the left-hand side of (11). Notice, that the total sum of coefficients at these terms is

$$\lambda_k(1 + \theta_k) + \sum_{i=1}^{k-1} [\lambda_i(1 + \theta_i) - \lambda_{i+1}\theta_{i+1}] = \sum_{i=1}^k \lambda_i + \lambda_1\theta_1 = S_k$$

Thus, by Jensen's inequality,

$$\frac{D}{2} \geq \frac{\text{LHS of (11)}}{2} \geq S_k(f(\hat{x}^k) - f_*),$$

where $\hat{x}^k = \frac{\sum_{i=1}^k \lambda_i x^i + \lambda_1 \theta_1 x^1}{S_k}$. And the first part of the proof is complete. \square

As we have shown that $\lambda_i \geq \frac{1}{2L}$ for all i , we have a theoretical upper bound $f(\hat{x}^k) - f_* \leq \frac{DL}{k}$. Note that in practice however, (λ_k) might be much larger than a pessimistic bound $\frac{1}{2L}$, which we observe in our experiments together with a faster convergence.

³Obviously, the closed convex-hull of a bounded set remains bounded.

Algorithm 2 Adaptive GD (L is known)

```
1:  $x^0 \in \mathbb{R}^d$ ,  $\lambda_0 = \frac{1}{L}$ ,  $\theta_0 = +\infty$ ,  $x^1 = x^0 - \lambda_0 \nabla f(x^0)$ 
2: for  $k = 1, 2, \dots$  do
3:    $L_k = \frac{\|\nabla f(x^k) - \nabla f(x^{k-1})\|}{\|x^k - x^{k-1}\|}$ 
4:    $\lambda_k = \min \left\{ \sqrt{1 + \theta_{k-1}} \lambda_{k-1}, \frac{1}{\lambda_{k-1} L^2} + \frac{1}{2L_k} \right\}$ 
5:    $x^{k+1} = x^k - \lambda_k \nabla f(x^k)$ 
6:    $\theta_k = \frac{\lambda_k}{\lambda_{k-1}}$ 
7: end for
```

2.3 f is L -smooth

The proof demonstrated above is novel even when f is L -smooth and we will use a fixed stepsize. A simple distinguished feature of it is that we did not use the descent lemma. This has, however, a drawback, as our analysis only allows one to take $\lambda \leq \frac{1}{2L}$ instead of $\lambda < \frac{2}{L}$. The natural question is then whether one can improve our algorithm and/or its analysis providing that f is L -smooth and L is known.

In this case we propose Algorithm 2, which is just a slight modification of the previous one. Interestingly, Lemma 1 remains valid and we get the next theorem (the proofs are in the appendix).

Theorem 2. *Let f be convex and L -smooth. Then for (x^k) generated by Algorithm 2 inequality (5) holds. As a corollary, it holds for some ergodic vector \hat{x}^k that $f(\hat{x}^k) - f_* = \mathcal{O}\left(\frac{1}{k}\right)$.*

Clearly, analysis of Algorithm 2 now allows one to take any fixed step $\lambda \leq \frac{3}{2L}$.

2.4 f is μ -strongly convex

Since one of our goals is to make optimization easy to use, we believe that a good method should have state-of-the-art guarantees in various scenarios. For strongly convex functions, this means that we want to see linear convergence, which is not covered by normalized GD or online methods. In section 2.1 we have shown that Algorithm 1 matches the $\mathcal{O}(1/\varepsilon)$ complexity of GD on convex problems. Now we show that it also matches $\mathcal{O}\left(\frac{L}{\mu} \log \frac{1}{\varepsilon}\right)$ complexity of GD when f is μ -strongly convex.

For proof simplicity, instead of using bound $\lambda_k \leq \sqrt{1 + \theta_{k-1}} \lambda_{k-1}$ as in line 3 of Algorithm 1 we will use a more conservative bound $\lambda_k \leq \sqrt{1 + \frac{\theta_{k-1}}{2}} \lambda_{k-1}$ (otherwise the derivation would be too technical). It is clear that with such a change Theorem 1 still holds true, so the sequence is bounded and we can rely on local smoothness and local strong convexity.

Theorem 3. *Suppose that $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is locally strongly convex and ∇f is locally Lipschitz. Then (x^k) generated by Algorithm 1 (with the modification mentioned above) converges to the solution x^* of (1). The complexity to get $\|x^k - x^*\|^2 \leq \varepsilon$ is $\mathcal{O}\left(\kappa \log \frac{1}{\varepsilon}\right)$, where $\kappa = \frac{L}{\mu}$ and L, μ are the smoothness and strong convexity constants of f on the set $\mathcal{C} = \overline{\text{conv}}\{x^*, x^0, x^1, \dots\}$.*

We want to highlight that in our rate κ depends on the local Lipschitz and strong convexity constants L and μ , which is meaningful even when these properties are not satisfied globally. Similarly, if f is globally smooth and strongly convex, our rate is still faster as it depends on the smaller local constant.

3 Heuristics

In this section we describe several extensions of our method. We do not have full theory for them, but believe that they are of interest in applications.

Algorithm 3 Adaptive accelerated gradient descent

```
1:  $x^0 \in \mathbb{R}^d$ ,  $\lambda_0 > 0$ ,  $\Lambda_0 > 0$ ,  $\theta_0 = \Theta_0 = +\infty$ ,  $y^1 = x^1 = x^0 - \lambda_0 \nabla f(x^0)$ 
2: for  $k = 1, 2, \dots$  do
3:    $\lambda_k = \min \left\{ \sqrt{1 + \theta_{k-1} \lambda_{k-1}}, \frac{\|x^k - x^{k-1}\|}{2\|\nabla f(x^k) - \nabla f(x^{k-1})\|} \right\}$ 
4:    $\Lambda_k = \min \left\{ \sqrt{1 + \Theta_{k-1} \Lambda_{k-1}}, \frac{\|\nabla f(x^k) - \nabla f(x^{k-1})\|}{2\|x^k - x^{k-1}\|} \right\}$ 
5:    $\beta_k = \frac{\sqrt{1/\lambda_k} - \sqrt{\Lambda_k}}{\sqrt{1/\lambda_k} + \sqrt{\Lambda_k}}$ 
6:    $y^{k+1} = x^k - \lambda_k \nabla f(x^k)$ 
7:    $x^{k+1} = y^{k+1} + \beta_k (y^{k+1} - y^k)$ 
8:    $\theta_k = \frac{\lambda_k}{\lambda_{k-1}}$ ,  $\Theta_k = \frac{\Lambda_k}{\Lambda_{k-1}}$ 
9: end for
```

3.1 Acceleration

Suppose that f is μ -strongly convex. One version of the accelerated gradient method proposed by Nesterov [26] is

$$\begin{aligned} y^{k+1} &= x^k - \frac{1}{L} \nabla f(x^k), \\ x^{k+1} &= y^{k+1} + \beta (y^{k+1} - y^k), \end{aligned}$$

where $\beta = \frac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}}$. We know that in the gradient descent L can be efficiently estimated using the scheme

$$\lambda_k = \min \left\{ \sqrt{1 + \theta_{k-1} \lambda_{k-1}}, \frac{\|x^k - x^{k-1}\|}{2\|\nabla f(x^k) - \nabla f(x^{k-1})\|} \right\}.$$

What about the strong convexity constant μ ? We know that it equals to the inverse smoothness constant of the conjugate $f^*(y) \stackrel{\text{def}}{=} \sup_x \{ \langle x, y \rangle - f(x) \}$. Thus, it is tempting to estimate this inverse constant just as we estimated inverse smoothness of f , i.e., by formula

$$\Lambda_k = \min \left\{ \sqrt{1 + \Theta_{k-1} \Lambda_{k-1}}, \frac{\|p^k - p^{k-1}\|}{2\|\nabla f^*(p^k) - \nabla f^*(p^{k-1})\|} \right\}$$

where p^k and p^{k-1} are some elements of the dual space and $\Theta_k = \frac{\Lambda_k}{\Lambda_{k-1}}$. A natural choice then is $p^k = \nabla f(x^k)$ since it is an element of the dual space that we use. What is its value? It is well known that $\nabla f^*(\nabla f(x)) = x$, so we come up with the update rule

$$\Lambda_k = \min \left\{ \sqrt{1 + \Theta_{k-1} \Lambda_{k-1}}, \frac{\|\nabla f(x^k) - \nabla f(x^{k-1})\|}{2\|x^k - x^{k-1}\|} \right\},$$

and hence we can estimate β by $\beta_k = \frac{\sqrt{1/\lambda_k} - \sqrt{\Lambda_k}}{\sqrt{1/\lambda_k} + \sqrt{\Lambda_k}}$.

We summarize our arguments in Algorithm 3. Unfortunately, we do not have any theoretical guarantees for it. Estimating strong convexity parameter μ is very important in practice. Most common approaches rely on restarting technique, proposed by Nesterov [25], see also [12] and references therein. Unlike Algorithm 3, these works have theoretical guarantees, however these methods by itself are more complicated and still require tuning of other unknown parameters.

Algorithm 4 Adaptive SGD

- 1: $x^0 \in \mathbb{R}^d$, $\lambda_0 > 0$, $\theta_0 = +\infty$, ξ^0 , $x^1 = x^0 - \lambda_0 \nabla f_{\xi^0}(x^0)$, $\alpha > 0$
 - 2: **for** $k = 1, 2, \dots$ **do**
 - 3: Sample ξ^k and optionally ζ^k
 - 4: Option I (biased): $L_k = \frac{\|\nabla f_{\xi^k}(x^k) - \nabla f_{\xi^k}(x^{k-1})\|}{\|x^k - x^{k-1}\|}$
 - 5: Option II (unbiased): $L_k = \frac{\|\nabla f_{\zeta^k}(x^k) - \nabla f_{\zeta^k}(x^{k-1})\|}{\|x^k - x^{k-1}\|}$
 - 6: $\lambda_k = \min\left\{\sqrt{1 + \theta_{k-1}}\lambda_{k-1}, \frac{1}{\alpha L_k}\right\}$
 - 7: $x^{k+1} = x^k - \lambda_k \nabla f_{\xi^k}(x^k)$
 - 8: $\theta_k = \frac{\lambda_k}{\lambda_{k-1}}$
 - 9: **end for**
-

3.2 Uniting our steps with stochastic gradients

Here we would like to discuss applications of our method to the problem

$$\min_x \mathbb{E} [f_\xi(x)],$$

where f_ξ is almost surely L -smooth and μ -strongly convex. Assume that at each iteration we get sample ξ^k to make a stochastic gradient step,

$$x^{k+1} = x^k - \lambda_k \nabla f_{\xi^k}(x^k).$$

Then, we have two ways of incorporating our stepsize into SGD: by using ξ^k in the estimate of λ_k , or using an extra sample ζ^k . If we use ξ^k , then we estimate $L_k = \frac{\|\nabla f_{\xi^k}(x^k) - \nabla f_{\xi^k}(x^{k-1})\|}{\|x^k - x^{k-1}\|}$, and compute $\lambda_k = \min\left\{(1 + \theta_k)\lambda_{k-1}, \frac{1}{\alpha L_k}\right\}$ and similarly we do if we used ζ^k . Clearly, the update with ξ^k will not be unbiased, although it is more intuitive to estimate smoothness on the same function that we use in the update. Surprisingly, the option with a biased estimate performed much better in our experiments. Full algorithm formulation is given in Algorithm 4.

The theorem below provides convergence guarantees for both cases, but with different assumptions. Unfortunately, we cannot guarantee similarly to Theorem 1 that one can use $\alpha = 2$.

Theorem 4. *Let f_ξ be L -smooth and μ -strongly convex almost surely. Assuming $\alpha \geq 2\kappa$ and estimating L_k with ∇f_{ζ^k} , the complexity to get $\mathbb{E} [\|x^k - x^*\|^2] \leq \varepsilon$ is not worse than $\mathcal{O}\left(\frac{\kappa^2}{\varepsilon} \log \frac{\kappa}{\varepsilon}\right)$. Furthermore, if the model is overparameterized, i.e., $\nabla f_\xi(x^*) = 0$ almost surely, then one can estimate L_k with ∇f_{ξ^k} and the complexity is $\mathcal{O}\left(\kappa^2 \log \frac{1}{\varepsilon}\right)$.*

Note that in both cases we match the known dependency on ε up to logarithmic terms, but we get an extra κ as the price for adaptive estimation of the stepsize.

3.3 Decreasing stepsizes in SGD

Another potential application of our techniques is estimation of decreasing stepsizes in SGD. Namely, the best known rates for SGD [37], which are optimal up to some constants, are obtained using λ_k that decreases as $\mathcal{O}\left(\frac{1}{L+\mu k}\right)$. For instance, [37] proposes to use stepsize $\lambda_k = \frac{2}{L+\mu[k-k_0]_+}$, where $[k-k_0]_+ \stackrel{\text{def}}{=} \max\{0, k-k_0\}$, $k_0 \stackrel{\text{def}}{=} \lceil K/2 \rceil$ and K is the total number of iterations. This requires estimates of both smoothness and strong convexity, which can be borrow from the previous discussion. We leave rigorous proof of such schemes for future work.

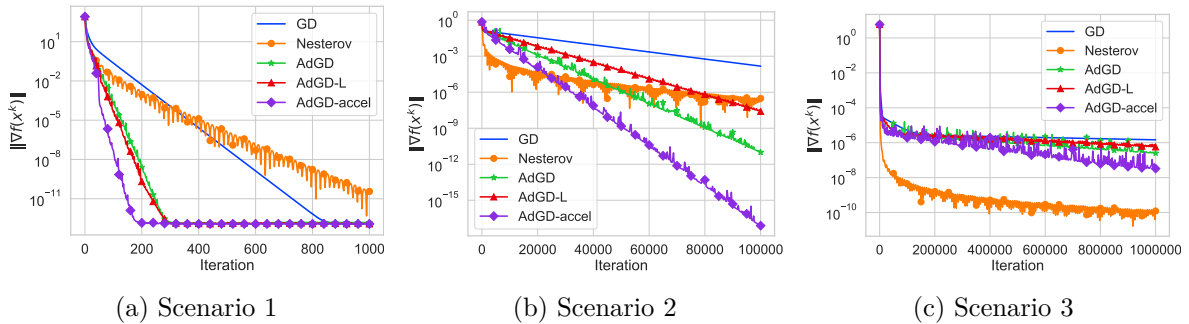


Figure 1: Results for the quadratic problem.




Figure 2: Results for the logistic regression problem. Note that ‘AdGD’ and ‘AdGD-L’ are almost indistinguishable as the global value of L is quite large compared to the local value, which can be seen from the plot in middle.

4 Experiments

In the experiments⁴, we compare our approach with the two most related methods: GD and Nesterov’s acceleration of GD for convex functions [24] (AGD). We always tune the stepsize for GD and AGD unless L is known. If no finite L exists, we still run the methods and observe their convergence even though it is not backed by theory.

Quadratic minimization. First, we consider probably the simplest possible optimization problem—minimization of the quadratic $f(x) = \frac{1}{2}\langle Ax, x \rangle$, for $A \in \mathbb{R}^{n \times n}$ in different settings. We run experiments for three scenarios: 1) $A = B^\top B$ with a Gaussian matrix $B = \mathbb{R}^{\frac{n}{2} \times n}$; 2) $A = B^{-1}$, where $B_{ij} = 0.99^{|i-j|}$; 3) $A = H^\top H$ with the Hilbert matrix $H \in \mathbb{R}^{n \times n}$: $H_{ij} = \frac{1}{i+j-1}$. In the experiments we used $n = 100$ and random initialization. The results are presented in Figure 1. This problem illustrates well that the “convergence rates are in the eye of a practitioner”.

Logistic regression. The ℓ_2 -regularized logistic regression objective is given by $\frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-b_i a_i^\top x)) + \frac{\ell_2}{2} \|x\|^2$, where n is the number of observations and $(a_i, b_i) \in \mathbb{R}^{d \times 1}$, $i = 1, \dots, n$ are the observations. We use ‘mushrooms’ and ‘covtype’ datasets to run the experiments. The results are provided in Figure 2. We choose the amount of ℓ_2 regularization proportionally to $\frac{1}{n}$ as often done in practice. Since we have closed-form expressions to estimate L and we know μ , we use them directly in GD and its acceleration.

⁴The code for the experiments can be found at https://github.com/yimalitsky/adaptive_gd




Figure 3: Results for matrix factorization. The objective is neither convex nor smooth.




Figure 4: Results for training ResNet-18 on Cifar10. Labels for AdGD correspond to how λ_k was estimated.

Matrix factorization. Given a matrix $A \in \mathbb{R}^{m \times n}$ and $r < \min\{m, n\}$, we want to solve $\min_{X=[U,V]} f(X) = f(U, V) = \frac{1}{2} \|UV^\top - A\|_F^2$ for $U \in \mathbb{R}^{m \times r}$ and $V \in \mathbb{R}^{n \times r}$. This is useful when we want to approximate A by some low rank matrix. It is a nonconvex problem, moreover, the gradient ∇f is not globally Lipschitz. With some tuning one still can apply GD and Nesterov’s accelerated GD, but—and we want to emphasize it—it was not a trivial thing to find the steps in practice. The steps we have chosen were in some sense optimal, i.e., GD and AGD did not converge if we doubled the steps. Interestingly, for both these methods their “optimal” steps were quite different. In contrast, our methods do not require any tuning, so even in this regard they are much more practical. For the experiments we used Movilens 100K dataset [14] with more than million entries and several values of $r = 10, 20, 30$. All algorithms were initialized at the same point, chosen randomly. The results are presented in Figure 3.

Neural network. We use standard ResNet-18 architecture implemented in Pytorch⁵ [29] and train it to classify images from the Cifar10 dataset [19] with cross-entropy loss. We use batch size 128 for all methods. For Adam and Adadelta, we took the default parameters, and learning rate 0.1 for SGD.

For our method, we observed that $\frac{1}{L_k}$ works better than $\frac{1}{2L_k}$. In addition, we ran it with $\sqrt{1 + \theta_k}$, $\sqrt{1 + \frac{\theta_k}{2}}$ and $\sqrt{1 + \frac{\theta_k}{10}}$ in the other factor and the latter performed the best. For reference, we provide the result for the theoretical estimate as well. We did not test any other values than the specified combinations, so we imagine that some other variation can perform better. We used the variant of SGD $x^{k+1} = x^k - \lambda_k \nabla f_{\xi^k}(x^k)$ with λ_k computed using ξ^k as well. See Figure 4 for the results.

⁵Code taken from <https://github.com/kuangliu/pytorch-cifar/blob/master/models/resnet.py>




Figure 5: Results for the non-smooth subproblem from cubic regularization.

Cubic regularization. In cubic regularization of Newton method [27], at each iteration we need to minimize $f(x) = g^\top x + \frac{1}{2}x^\top Hx + \frac{M}{6}\|x\|^3$, where $g \in \mathbb{R}^d$, $H \in \mathbb{R}^{d \times d}$ and $M > 0$ are given. This objective is smooth only locally due to the cubic term, which is our motivation to consider it. g and H were the gradient and the Hessian of the logistic loss with the Covtype dataset, evaluated at $x = 0 \in \mathbb{R}^d$. Although the values of $M = 10, 20, 100$ led to similar results, they also required different numbers of iterations, so we present the corresponding results in Figure 5.

5 Perspectives

We briefly provide a few directions which we personally consider to be important and challenging.

1. **Nonconvex case.** A great challenge for us is to understand theoretical guarantees of the proposed method in the nonconvex settings. We are not aware of any generic first-order method for nonconvex optimization that does not rely on the descent lemma (or its generalization), see, e.g., [2].
2. **Performance estimation.** In our experiments we often observed much better performance of Algorithm 1, than GD or AGD. However, the theoretical rate we can show coincides with that of GD. The challenge here is to bridge this gap and we hope that the approach pioneered by [10] and further developed in [17, 38, 39] has a potential to do that.
3. **Composite minimization.** In classical first-order methods the transition from smooth to composite minimization [25] is rather straightforward. Unfortunately, the proposed proof of Algorithm 1 does not seem to provide any route for generalization and we hope there is some way of resolving this issue.
4. **Stochastic optimization.** The derived bounds for the stochastic case are not satisfactory and have suboptimal dependency on κ . However, it is not clear to us whether one can extend the techniques from the deterministic analysis to improve the rate.
5. **Heuristics.** Finally, we want to have some solid ground in understanding the performance of the proposed heuristics.

Acknowledgments. Yura Malitsky wishes to thank Roman Cheplyaka for his interest in optimization that partly inspired the current work.

References

- [1] L. ARMIJO, *Minimization of functions having lipschitz continuous first partial derivatives.*, Pacific Journal of Mathematics, 16 (1966), pp. 1–3.
- [2] H. ATTOUCH, J. BOLTE, AND B. F. SVAITER, *Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized Gauss–Seidel methods*, Mathematical Programming, 137 (2013), pp. 91–129.
- [3] J. BARZILAI AND J. M. BORWEIN, *Two-point step size gradient methods*, IMA Journal of Numerical Analysis, 8 (1988), pp. 141–148.
- [4] H. H. BAUSCHKE, J. BOLTE, AND M. TEBoulLE, *A descent lemma beyond Lipschitz gradient continuity: first-order methods revisited and applications*, Mathematics of Operations Research, 42 (2016), pp. 330–348.
- [5] J. Y. BELLO CRUZ AND T. T. NGHIA, *On the convergence of the forward–backward splitting method with linesearches*, Optimization Methods and Software, 31 (2016), pp. 1209–1238.
- [6] J. BERNSTEIN, Y.-X. WANG, K. AZIZZADENESHELI, AND A. ANANDKUMAR, *signSGD: Compressed optimisation for non-convex problems*, in International Conference on Machine Learning, 2018, pp. 559–568.
- [7] A. CAUCHY, *Méthode générale pour la résolution des systemes d’équations simultanées*, Comp. Rend. Sci. Paris, 25 (1847), pp. 536–538.
- [8] Y.-H. DAI AND R. FLETCHER, *Projected Barzilai-Borwein methods for large-scale box-constrained quadratic programming*, Numerische Mathematik, 100 (2005), pp. 21–47.
- [9] Y.-H. DAI AND L.-Z. LIAO, *R-linear convergence of the barzilai and borwein gradient method*, IMA Journal of Numerical Analysis, 22 (2002), pp. 1–10.
- [10] Y. DRORI AND M. TEBoulLE, *Performance of first-order methods for smooth convex minimization: a novel approach*, Mathematical Programming, 145 (2014), pp. 451–482.
- [11] J. DUCHI, E. HAZAN, AND Y. SINGER, *Adaptive subgradient methods for online learning and stochastic optimization*, Journal of Machine Learning Research, 12 (2011), pp. 2121–2159.
- [12] O. FERCOQ AND Z. QU, *Adaptive restart of accelerated gradient methods under local quadratic growth condition*, arXiv preprint arXiv:1709.02300, (2017).
- [13] A. GOLDSTEIN, *Cauchy’s method of minimization*, Numerische Mathematik, 4 (1962), pp. 146–150.
- [14] F. M. HARPER AND J. A. KONSTAN, *The movielens datasets: History and context*, ACM transactions on interactive intelligent systems (tiis), 5 (2016), p. 19.
- [15] E. HAZAN AND S. KAKADE, *Revisiting the Polyak step size*, arXiv preprint arXiv:1905.00313, (2019).
- [16] S. P. KARIMIREDDY, Q. REBJOCK, S. STICH, AND M. JAGGI, *Error feedback fixes signsgd and other gradient compression schemes*, in International Conference on Machine Learning, 2019, pp. 3252–3261.
- [17] D. KIM AND J. A. FESSLER, *Optimized first-order methods for smooth convex minimization*, Mathematical programming, 159 (2016), pp. 81–107.

- [18] D. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, International Conference on Learning Representations, (2014).
- [19] A. KRIZHEVSKY, G. HINTON, ET AL., *Learning multiple layers of features from tiny images*, tech. rep., Citeseer, 2009.
- [20] C. LEMARÉCHAL, *Cauchy and the gradient method*, Doc Math Extra, 251 (2012), p. 254.
- [21] C. J. MADDISON, D. PAULIN, Y. W. TEH, AND A. DOUCET, *Dual space preconditioning for gradient descent*, arXiv preprint arXiv:1902.02257, (2019).
- [22] Y. MALITSKY, *Golden ratio algorithms for variational inequalities*, Mathematical Programming, (2019).
- [23] H. B. MCMAHAN AND M. STREETER, *Adaptive bound optimization for online convex optimization*, arXiv preprint arXiv:1002.4908, (2010).
- [24] Y. NESTEROV, *A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$* , Doklady AN SSSR, 269 (1983), pp. 543–547.
- [25] Y. NESTEROV, *Gradient methods for minimizing composite functions*, Mathematical Programming, 140 (2013), pp. 125–161.
- [26] Y. NESTEROV, *Introductory lectures on convex optimization: A basic course*, vol. 87, Springer Science & Business Media, 2013.
- [27] Y. NESTEROV AND B. T. POLYAK, *Cubic regularization of Newton method and its global performance*, Mathematical Programming, 108 (2006), pp. 177–205.
- [28] L. NGUYEN, P. H. NGUYEN, M. DIJK, P. RICHTÁRIK, K. SCHEINBERG, AND M. TAKAC, *SGD and Hogwild! convergence without the bounded gradients assumption*, in International Conference on Machine Learning, 2018, pp. 3747–3755.
- [29] A. PASZKE, S. GROSS, S. CHINTALA, G. CHANAN, E. YANG, Z. DEVITO, Z. LIN, A. DESMAISON, L. ANTIGA, AND A. LERER, *Automatic differentiation in pytorch*, (2017).
- [30] B. T. POLYAK, *Gradient methods for minimizing functionals*, Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki, 3 (1963), pp. 643–653.
- [31] ———, *Minimization of nonsmooth functionals*, Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki, 9 (1969), pp. 509–521.
- [32] M. RAYDAN, *On the Barzilai and Borwein choice of steplength for the gradient method*, IMA Journal of Numerical Analysis, 13 (1993), pp. 321–326.
- [33] ———, *The Barzilai and Borwein gradient method for the large scale unconstrained minimization problem*, SIAM Journal on Optimization, 7 (1997), pp. 26–33.
- [34] S. J. REDDI, S. KALE, AND S. KUMAR, *On the convergence of adam and beyond*, in International Conference on Learning Representations, 2018.
- [35] S. SALZO, *The variable metric forward-backward splitting algorithm under mild differentiability assumptions*, SIAM Journal on Optimization, 27 (2017), pp. 2153–2181.
- [36] N. SHOR, *An application of the method of gradient descent to the solution of the network transportation problem*, Materialy Naucnovo Seminara po Teoret i Priklad. Voprosam Kibernet. i Issted. Operacii, Nuyni Sov. po Kibernet, Akad. Nauk Ukrain. SSSR, vyp, 1 (1962), pp. 9–17.

- [37] S. U. STICH, *Unified optimal analysis of the (stochastic) gradient method*, arXiv preprint arXiv:1907.04232, (2019).
- [38] A. TAYLOR AND F. BACH, *Stochastic first-order methods: non-asymptotic and computer-aided analyses via potential functions*, in Proceedings of the Thirty-Second Conference on Learning Theory, A. Beygelzimer and D. Hsu, eds., vol. 99 of Proceedings of Machine Learning Research, Phoenix, USA, 25–28 Jun 2019, PMLR, pp. 2934–2992.
- [39] A. B. TAYLOR, J. M. HENDRICKX, AND F. GLINEUR, *Smooth strongly convex interpolation and exact worst-case performance of first-order methods*, Mathematical Programming, 161 (2017), pp. 307–345.
- [40] T. TIELEMAN AND G. HINTON, *Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude*. COURSERA: Neural Networks for Machine Learning, 2012.
- [41] R. WARD, X. WU, AND L. BOTTOU, *AdaGrad stepsizes: Sharp convergence over nonconvex landscapes*, in Proceedings of the 36th International Conference on Machine Learning, K. Chaudhuri and R. Salakhutdinov, eds., vol. 97 of Proceedings of Machine Learning Research, Long Beach, California, USA, 09–15 Jun 2019, PMLR, pp. 6677–6686.
- [42] M. ZAHEER, S. REDDI, D. SACHAN, S. KALE, AND S. KUMAR, *Adaptive methods for nonconvex optimization*, in Advances in Neural Information Processing Systems, 2018, pp. 9793–9803.
- [43] M. D. ZEILER, *Adadelta: an adaptive learning rate method*, arXiv preprint arXiv:1212.5701, (2012).

Appendix:

A Missing proofs

Lemma 2 (Theorem 2.1.5 in [26]). Let \mathcal{C} be a closed convex set in \mathbb{R}^d . If $f: \mathcal{C} \rightarrow \mathbb{R}$ is convex and L -smooth, then $\forall x, y \in \mathcal{C}$ it holds

$$f(x) - f(y) - \langle \nabla f(y), x - y \rangle \geq \frac{1}{2L} \|\nabla f(x) - \nabla f(y)\|^2. \quad (12)$$

We also need some variation of the Opial lemma.

Lemma 3. Let (x^k) and (a_k) be two sequences in \mathbb{R}^d and \mathbb{R}_+ respectively. Suppose that (x^k) is a bounded sequence whose cluster points belong to $S \subset \mathbb{R}^d$ and it also holds that

$$\|x^{k+1} - x\|^2 + a_{k+1} \leq \|x^k - x\|^2 + a_k, \quad \forall x \in S. \quad (13)$$

Then (x^k) converges to some element in S .

Proof. Let \bar{x}_1, \bar{x}_2 be any cluster points of (x^k) . Thus, there exist two subsequences (x^{k_i}) and (x^{k_j}) such that $x^{k_i} \rightarrow \bar{x}_1$ and $x^{k_j} \rightarrow \bar{x}_2$. Obviously, $\lim_{k \rightarrow \infty} \|x^k - x\|^2 + a_k$ exists for any $x \in S$. Let $x = \bar{x}_1$. This yields

$$\begin{aligned} \lim_{k \rightarrow \infty} \|x^k - \bar{x}_1\|^2 + a_k &= \lim_{i \rightarrow \infty} \|x^{k_i} - \bar{x}_1\|^2 + a_{k_i} = \lim_{i \rightarrow \infty} a_{k_i} \\ &= \lim_{j \rightarrow \infty} \|x^{k_j} - \bar{x}_1\|^2 + a_{k_j} = \|\bar{x}_2 - \bar{x}_1\|^2 + \lim_{j \rightarrow \infty} a_{k_j}. \end{aligned}$$

Hence, $\lim_{i \rightarrow \infty} a_{k_i} = \lim_{j \rightarrow \infty} a_{k_j} + \|\bar{x}_1 - \bar{x}_2\|^2$. Doing the same with $x = \bar{x}_2$ instead of $x = \bar{x}_1$, yields $\lim_{j \rightarrow \infty} a_{k_j} = \lim_{i \rightarrow \infty} a_{k_i} + \|\bar{x}_1 - \bar{x}_2\|^2$. Hence, we obtain that $\bar{x}_1 = \bar{x}_2$, which finishes the proof. \square

Proof of Theorem 1. (Convergence of (x^k)).

Note that in the first part we have already proved that (x^k) is bounded and that ∇f is L -Lipschitz on $\mathcal{C} = \overline{\text{conv}}\{x^*, x^0, x^1, \dots\}$. Invoking Lemma 1, we deduce that

$$\lambda_k(f(x^*) - f(x^k)) \geq \lambda_k \langle \nabla f(x^k), x - x^k \rangle + \frac{\lambda_k}{2L} \|\nabla f(x^k)\|^2. \quad (14)$$

This indicates that instead of using inequality (6) in the proof of Lemma 1, we could use a better estimate (14). However, we want to emphasize that we did not assume that ∇f is globally Lipschitz, but rather obtained Lipschitzness on \mathcal{C} as an artifact of our analysis. Clearly, in the end this improvement gives us an additional term $\frac{\lambda_k}{L} \|\nabla f(x^k)\|^2$ in the left-hand side of (5). Thus, now telescoping (5), one obtains that $\sum_{i=1}^k \frac{\lambda_k}{2L} \|\nabla f(x^k)\|^2 \leq D$. As $\lambda_k \geq \frac{1}{2L}$, one has that $\nabla f(x^k) \rightarrow 0$. Now we might conclude that all cluster points of (x^k) are solutions of (1).

Let S be the solution set of (1) and $a_k = \frac{1}{2} \|x^k - x^{k-1}\|^2 + 2\lambda_k \theta_k (f(x^{k-1}) - f_*)$. We want to finish the proof applying Lemma 3. To this end, notice that inequality (5) yields (13), since $\lambda_{k+1} \theta_{k+1} \leq (1 + \theta_k) \lambda_k$. This completes the proof. \square

Proof of Lemma 2. Proceeding similarly as in Lemma 1, we have

$$\|x^{k+1} - x^*\|^2 = \|x^k - x^*\|^2 - 2\lambda_k \langle \nabla f(x^k), x^k - x^* \rangle + \|x^{k+1} - x^k\|^2. \quad (15)$$

By convexity of f and Lemma 2,

$$\begin{aligned} 2\lambda_k \langle \nabla f(x^k), x^* - x^k \rangle &\stackrel{(12)}{\leq} 2\lambda_k (f(x^*) - f(x^k)) - \frac{1}{2L} \|\nabla f(x^k)\|^2 \\ &= 2\lambda_k (f_* - f(x^k)) - \frac{1}{\lambda_k L} \|x^{k+1} - x^k\|^2. \end{aligned} \quad (16)$$

As in (8), we have

$$\|x^{k+1} - x^k\|^2 = 2\lambda_k \langle \nabla f(x^k) - f(x^{k-1}), x^k - x^{k+1} \rangle + 2\lambda_k \langle f(x^{k-1}), x^k - x^{k+1} \rangle - \|x^{k+1} - x^k\|^2. \quad (17)$$

Again, instead of using merely convexity of f , we combine it with Lemma 2. This gives

$$\begin{aligned} 2\lambda_k \langle \nabla f(x^{k-1}), x^k - x^{k+1} \rangle &= \frac{2\lambda_k}{\lambda_{k-1}} \langle x^{k-1} - x^k, x^k - x^{k+1} \rangle \\ &= 2\lambda_k \theta_k \langle x^{k-1} - x^k, \nabla f(x^k) \rangle \\ &\stackrel{(12)}{\leq} 2\lambda_k \theta_k (f(x^{k-1}) - f(x^k)) - \frac{\lambda_k \theta_k}{L} \|\nabla f(x^k) - \nabla f(x^{k-1})\|^2. \end{aligned} \quad (18)$$

Since now we have two additional terms $\frac{1}{\lambda_k L} \|x^{k+1} - x^k\|^2$ and $\frac{\lambda_k \theta_k}{L} \|\nabla f(x^k) - \nabla f(x^{k-1})\|^2$, we can do better than (9). But first we need a simple, yet a bit tedious fact. By our choice of λ_k , in every iteration $\lambda_k \leq \frac{1}{\lambda_{k-1} L^2} + \frac{1}{2L_k}$ with $L_k = \frac{\|\nabla f(x^k) - \nabla f(x^{k-1})\|}{\|x^k - x^{k-1}\|}$. We want to show that it implies

$$2 \left(\lambda_k - \frac{\sqrt{\theta_k}}{L} \right) \leq \frac{1}{L_k}, \quad (19)$$

which is equivalent to $\lambda_k - \frac{\sqrt{\lambda_k}}{\sqrt{\lambda_{k-1} L}} - \frac{1}{2L_k} \leq 0$. Nonnegative solutions of the quadratic inequality $t^2 - \frac{t}{\sqrt{\lambda_{k-1} L}} - \frac{1}{2L_k} \leq 0$ are

$$0 \leq t \leq \frac{1}{2\sqrt{\lambda_{k-1} L}} + \frac{1}{2} \sqrt{\frac{1}{\lambda_{k-1} L^2} + \frac{2}{L_k}} = \frac{1}{2\sqrt{\lambda_{k-1} L}} \left(1 + \sqrt{1 + \frac{2\lambda_{k-1} L^2}{L_k}} \right).$$

Let us prove that $\sqrt{\frac{1}{\lambda_{k-1} L^2} + \frac{1}{2L_k}}$ falls into this segment and, hence, $\sqrt{\lambda_k}$ does as well. Using a simple inequality $4 + a \leq (1 + \sqrt{1 + a})^2$, for $a > 0$, we obtain

$$\frac{1}{\lambda_{k-1} L^2} + \frac{1}{2L_k} = \frac{1}{4\lambda_{k-1} L^2} \left(4 + \frac{2\lambda_{k-1} L^2}{L_k} \right) \leq \frac{1}{4\lambda_{k-1} L^2} \left(1 + \sqrt{1 + \frac{2\lambda_{k-1} L^2}{L_k}} \right)^2.$$

This confirms that (19) is true. Thus, by Cauchy-Schwarz and Young's inequalities, one has

$$\begin{aligned} &2\lambda_k \langle \nabla f(x^k) - \nabla f(x^{k-1}), x^k - x^{k+1} \rangle \leq 2\lambda_k \|\nabla f(x^k) - \nabla f(x^{k-1})\| \|x^k - x^{k+1}\| \\ &= 2 \left(\lambda_k - \frac{\sqrt{\theta_k}}{L} \right) \|\nabla f(x^k) - \nabla f(x^{k-1})\| \|x^k - x^{k+1}\| + \frac{2\sqrt{\theta_k}}{L} \|\nabla f(x^k) - \nabla f(x^{k-1})\| \|x^k - x^{k+1}\| \\ &\stackrel{(19)}{\leq} \frac{1}{L_k} \|\nabla f(x^k) - \nabla f(x^{k-1})\| \|x^k - x^{k+1}\| + \frac{\lambda_k \theta_k}{L} \|\nabla f(x^k) - \nabla f(x^{k-1})\|^2 + \frac{1}{\lambda_k L} \|x^{k+1} - x^k\|^2 \\ &= \|x^k - x^{k-1}\| \|x^k - x^{k+1}\| + \frac{\lambda_k \theta_k}{L} \|\nabla f(x^k) - \nabla f(x^{k-1})\|^2 + \frac{1}{\lambda_k L} \|x^{k+1} - x^k\|^2 \\ &\leq \frac{1}{2} \|x^k - x^{k-1}\|^2 + \frac{1}{2} \|x^{k+1} - x^k\|^2 + \frac{\lambda_k \theta_k}{L} \|\nabla f(x^k) - \nabla f(x^{k-1})\|^2 + \frac{1}{\lambda_k L} \|x^{k+1} - x^k\|^2. \end{aligned} \quad (20)$$

Combining everything together, we obtain the statement of the lemma. \square

Proof of Theorem 3. First of all, we note that using stricter inequality $\lambda_k \leq \sqrt{1 + \frac{\theta_{k-1}}{2}} \lambda_{k-1}$ does not change the statement of Theorem 1. Hence, $x^k \rightarrow x^*$ and there exist $\mu, L > 0$ such that f is μ -strongly convex and ∇f is L -Lipschitz on $\mathcal{C} = \overline{\text{conv}}\{x^*, x^0, x^1, \dots\}$. Secondly, due to strong convexity, $\|\nabla f(x^k) - \nabla f(x^{k-1})\| \geq \mu \|x^k - x^{k-1}\|$, and hence $\lambda_k \leq \frac{1}{2\mu}$ for $k \geq 1$.

Now we tighten some steps in the analysis to improve bound (6). By strong convexity,

$$\lambda_k \langle \nabla f(x^k), x^* - x^k \rangle \leq \lambda_k (f(x^*) - f(x^k)) - \lambda_k \frac{\mu}{2} \|x^* - x^k\|^2.$$

By L -smoothness and bound $\lambda_k \leq \frac{1}{2\mu}$,

$$\begin{aligned} \lambda_k \langle \nabla f(x^k), x^* - x^k \rangle &\leq \lambda_k (f(x^*) - f(x^k)) - \lambda_k \frac{1}{2L} \|\nabla f(x^k)\|^2 \\ &= \lambda_k (f_* - f(x^k)) - \frac{1}{2L\lambda_k} \|x^{k+1} - x^k\|^2 \\ &\leq \lambda_k (f_* - f(x^k)) - \frac{\mu}{L} \|x^{k+1} - x^k\|^2. \end{aligned}$$

Together, these two bounds give us

$$\lambda_k \langle \nabla f(x^k), x^* - x^k \rangle \leq \lambda_k (f_* - f(x^k)) - \lambda_k \frac{\mu}{4} \|x^k - x^*\|^2 - \frac{\mu}{2L} \|x^{k+1} - x^k\|^2.$$

We keep inequality (10) and the rest of the proof as is. Then the strengthened analog of (5) will be

$$\begin{aligned} &\|x^{k+1} - x^*\|^2 + \frac{1}{2} \left(1 + \frac{2\mu}{L}\right) \|x^{k+1} - x^k\|^2 + 2\lambda_k (1 + \theta_k) (f(x^k) - f_*) \\ &\leq \left(1 - \frac{\lambda_k \mu}{2}\right) \|x^k - x^*\|^2 + \frac{1}{2} \|x^k - x^{k-1}\|^2 + 2\lambda_k \theta_k (f(x^{k-1}) - f_*) \\ &\leq \left(1 - \frac{\lambda_k \mu}{2}\right) \|x^k - x^*\|^2 + \frac{1}{2} \|x^k - x^{k-1}\|^2 + 2\lambda_{k-1} \left(1 + \frac{\theta_{k-1}}{2}\right) (f(x^{k-1}) - f_*), \end{aligned} \quad (21)$$

where in the last inequality we used our new condition on λ_k . Under the new update we have contraction in every term: $1 - \frac{\lambda_k \mu}{2}$ in the first, $\frac{1}{1+2\mu/L} = 1 - \frac{2\mu}{L+2\mu}$ in the second and $\frac{1+\theta_{k-1}/2}{1+\theta_{k-1}} = 1 - \frac{\theta_{k-1}}{2(1+\theta_{k-1})}$ in the last one.

To further bound the last contraction, recall that $\lambda_k \in \left[\frac{1}{2L}, \frac{1}{2\mu}\right]$ for $k \geq 1$. Therefore, $\theta_k = \frac{\lambda_k}{\lambda_{k-1}} \geq \frac{1}{\kappa}$ for any $k > 1$, where $\kappa \stackrel{\text{def}}{=} \frac{L}{\mu}$. Since the function $\theta \mapsto \frac{\theta}{1+\theta}$ monotonically increases with $\theta > 0$, this implies $\frac{\theta_{k-1}}{2(1+\theta_{k-1})} \geq \frac{1}{2(\kappa+1)}$ when $k > 2$. Thus, for the full energy Ψ^{k+1} (the left-hand side of (21)) we have

$$\Psi^{k+1} \leq \left(1 - \min \left\{ \frac{\lambda_k \mu}{2}, \frac{1}{2(\kappa+1)}, \frac{2\mu}{L+2\mu} \right\}\right) \Psi^k.$$

Using simple bounds $\frac{\lambda_k \mu}{2} \geq \frac{1}{4\kappa}$, $\frac{2\mu}{L+2\mu} = \frac{2}{\kappa+2} \geq \frac{1}{4\kappa}$, and $\frac{1}{2(\kappa+1)} \geq \frac{1}{4\kappa}$, we obtain $\Psi^{k+1} \leq \left(1 - \frac{1}{4\kappa}\right) \Psi^k$ for $k > 2$. This gives $\mathcal{O}\left(\kappa \log \frac{1}{\varepsilon}\right)$ convergence rate. \square

B Stochastic analysis

B.1 Different samples

Consider the following version of SGD, in which we have two samples at each iteration, ξ^k and ζ^k to compute

$$\begin{aligned} \lambda_k &= \min \left\{ \sqrt{1 + \theta_k} \lambda_{k-1}, \frac{\|x^k - x^{k-1}\|}{\alpha \|\nabla f_{\zeta^k}(x^k) - \nabla f_{\zeta^k}(x^{k-1})\|} \right\}, \\ x^{k+1} &= x^k - \lambda_k \nabla f_{\xi^k}(x^k). \end{aligned}$$

As before, we assume that $\theta_0 = +\infty$, so $\lambda_1 = \frac{\|x^1 - x^0\|}{\alpha \|\nabla f_{\zeta^1}(x^1) - \nabla f_{\zeta^1}(x^0)\|}$.

Lemma 4. Let f_ξ be L -smooth μ -strongly convex almost surely. It holds for λ_k produced by the rule above

$$\frac{1}{L\alpha} \leq \lambda_k \leq \frac{1}{\mu\alpha} \quad \text{a.s.} \quad (22)$$

Proof. Let us start with the upper bound. Strong convexity of f_{ξ^k} implies that $\|x - y\| \leq \frac{1}{\mu} \|\nabla f_{\xi^k}(x) - \nabla f_{\xi^k}(y)\|$ for any x, y . Therefore, $\lambda_k \leq \min \{\sqrt{1 + \theta_k} \lambda_{k-1}, 1/(\mu\alpha)\} \leq 1/(\mu\alpha)$ a.s. On the other hand, L -smoothness gives $\lambda_k \geq \min \{\sqrt{1 + \theta_k} \lambda_{k-1}, 1/(L\alpha)\} \geq \min \{\lambda_{k-1}, 1/(L\alpha)\}$ a.s. Recursing this inequality, we obtain the stated lower bound. \square

Proposition 1 (Lemma 1 in [28]). Denote $\sigma^2 \stackrel{\text{def}}{=} \mathbb{E} [\|\nabla f_\xi(x^*)\|^2]$ and assume f to be almost surely L -smooth and convex. Then it holds for any x

$$\mathbb{E} [\|\nabla f_\xi(x)\|^2] \leq 4L(f(x) - f_*) + 2\sigma^2. \quad (23)$$

Another fact that we will use is a strong convexity bound, which states for any x, y

$$\langle \nabla f(x), x - y \rangle \geq \frac{\mu}{2} \|x - y\|^2 + f(x) - f(y). \quad (24)$$

Theorem 5. Let f_ξ be L -smooth and μ -strongly convex almost surely. If we choose some $\alpha \geq 2\frac{L}{\mu}$, then

$$\mathbb{E} [\|x^k - x^*\|^2] \leq \exp\left(-k \frac{\mu}{L\alpha}\right) C_0 + \frac{\sigma^2}{\mu^2\alpha},$$

where $C_0 \stackrel{\text{def}}{=} 2(1 + 2\lambda_0^2 L^2) \|x^0 - x^*\|^2 + 4\lambda_0^2 \sigma^2$ and $\sigma^2 \stackrel{\text{def}}{=} \mathbb{E} [\|\nabla f_\xi(x^*)\|^2]$.

Proof. Under our assumptions on $\alpha \geq 2\frac{L}{\mu}$, we have $\lambda_k \leq \frac{1}{\mu\alpha} \leq \frac{1}{2L}$. Since λ_k is independent of ξ^k , we have $\mathbb{E} [\lambda_k \nabla f_{\xi^k}(x^k)] = \mathbb{E} [\lambda_k] \mathbb{E} [\nabla f(x^k)]$ and

$$\begin{aligned} \mathbb{E} [\|x^{k+1} - x^*\|^2] &= \mathbb{E} [\|x^k - x^*\|^2] - 2\mathbb{E} [\lambda_k \langle \nabla f(x^k), x^k - x^* \rangle] + \mathbb{E} [\lambda_k^2] \mathbb{E} [\|\nabla f_{\xi^k}(x^k)\|^2] \\ &\stackrel{(24)}{\leq} \mathbb{E} [(1 - \lambda_k \mu) \|x^k - x^*\|^2] - 2\mathbb{E} [\lambda_k (f(x^k) - f_*)] + \mathbb{E} [\lambda_k^2] \mathbb{E} [\|\nabla f_{\xi^k}(x^k)\|^2] \\ &\stackrel{(23)}{\leq} \mathbb{E} [(1 - \lambda_k \mu) \|x^k - x^*\|^2] - 2\mathbb{E} [\lambda_k \underbrace{(1 - 2\lambda_k L)}_{\geq 0} (f(x^k) - f_*)] + \mathbb{E} [\lambda_k^2] \sigma^2 \\ &\stackrel{(22)}{\leq} \mathbb{E} [1 - \lambda_k \mu] \mathbb{E} [\|x^k - x^*\|^2] + \frac{\mathbb{E} [\lambda_k] \sigma^2}{\mu\alpha}. \end{aligned}$$

Therefore, if we subtract $\frac{\sigma^2}{\mu^2\alpha}$ from both sides, we obtain

$$\mathbb{E} \left[\|x^{k+1} - x^*\|^2 - \frac{\sigma^2}{\mu^2\alpha} \right] \leq \mathbb{E} [1 - \lambda_k \mu] \mathbb{E} \left[\|x^k - x^*\|^2 - \frac{\sigma^2}{\mu^2\alpha} \right].$$

If $\mathbb{E} [\|x^k - x^*\|^2] \leq \frac{\sigma^2}{\mu^2\alpha}$ for some k , it follows that $\mathbb{E} [\|x^t - x^*\|^2] \leq \frac{\sigma^2}{\mu^2\alpha}$ for any $t \geq k$. Otherwise, we can recurse the produced bound to obtain

$$\mathbb{E} [\|x^{k+1} - x^*\|^2] \leq \prod_{t=1}^k \mathbb{E} [1 - \lambda_t \mu] \|x^1 - x^*\|^2 + \frac{\sigma^2}{\mu^2\alpha}.$$

By inequality $1 - x \leq e^{-x}$, we have $\prod_{t=1}^k \mathbb{E}[1 - \lambda_t \mu] \leq \exp\left(-\mu \sum_{t=0}^k \lambda_t\right)$. In addition, recall that in accordance with (22) we have $\lambda_k \geq \frac{1}{L\alpha}$. Thus,

$$\mathbb{E}\left[\|x^{k+1} - x^*\|^2\right] \leq \exp\left(-k \frac{\mu}{L\alpha}\right) \mathbb{E}\left[\|x^1 - x^*\|^2\right] + \frac{\sigma^2}{\mu^2 \alpha}.$$

It remains to mention that

$$\mathbb{E}\left[\|x^1 - x^*\|^2\right] \leq 2\|x^0 - x^*\| + 2\lambda_0^2 \mathbb{E}\left[\|\nabla f_{\xi^0}(x^0)\|^2\right] \stackrel{(23)}{\leq} 2\|x^0 - x^*\| + 2\lambda_0^2 \left(2L^2\|x^0 - x^*\|^2 + 2\sigma^2\right).$$

□

This gives the following corollary.

Corollary 1. Choose $\alpha = 2\gamma \frac{L}{\mu}$ with $\gamma \geq 1$. Then, to achieve $\mathbb{E}\left[\|x^k - x^*\|^2\right] = \mathcal{O}\left(\varepsilon + \frac{\sigma^2}{\gamma}\right)$ we need only $k = \mathcal{O}\left(\gamma \frac{L^2}{\mu^2} \log \frac{1+\lambda_0^2}{\varepsilon}\right)$ iterations. If we choose γ proportionally to $\frac{1}{\varepsilon}$, it implies $\mathcal{O}\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ complexity.

B.2 Same sample: overparameterized models

Assume additionally that the model is overparameterized, i.e., $\nabla f_{\xi}(x^*) = 0$ with probability one. In that case, we can prove that one can use the same stochastic sample to compute the stepsize and to move the iterate. The update becomes

$$\lambda_k = \min\left\{\sqrt{1 + \theta_k} \lambda_{k-1}, \frac{\|x^k - x^{k-1}\|}{\alpha \|\nabla f_{\xi^k}(x^k) - \nabla f_{\xi^k}(x^{k-1})\|}\right\},$$

$$x^{k+1} = x^k - \lambda_k \nabla f_{\xi^k}(x^k).$$

Theorem 6. Let f_{ξ} be L -smooth, μ -strongly convex and satisfy $\nabla f_{\xi}(x^*) = 0$ with probability one. If we choose $\alpha \geq \frac{L}{\mu}$, then

$$\mathbb{E}\left[\|x^k - x^*\|^2\right] \leq \exp\left(-k \frac{\mu}{L\alpha}\right) C_0,$$

where $C_0 \stackrel{\text{def}}{=} 2(1 + \lambda_0^2 L^2) \|x^0 - x^*\|^2$.

Proof. Now λ_k depends on ξ^k , so we do not have an unbiased update anymore. However, under the new assumption, $\nabla f_{\xi^k}(x^*) = 0$, so we can write

$$\langle \nabla f_{\xi^k}(x^k), x^k - x^* \rangle \stackrel{(24)}{\geq} \frac{\mu}{2} \|x^k - x^*\|^2 + f_{\xi^k}(x^k) - f_{\xi^k}(x^*).$$

In addition, L -smoothness and convexity of f_{ξ^k} gives

$$\|\nabla f_{\xi^k}(x^k)\|^2 \leq 2L(f_{\xi^k}(x^k) - f_{\xi^k}(x^*)).$$

Since our choice of α implies $\lambda_k \leq \frac{1}{L}$, we conclude that

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &= \|x^k - x^*\|^2 - 2\lambda_k \langle \nabla f_{\xi^k}(x^k), x^k - x^* \rangle + \lambda_k^2 \|\nabla f_{\xi^k}(x^k)\|^2 \\ &\leq (1 - \lambda_k \mu) \|x^k - x^*\|^2 - 2\lambda_k (1 - \lambda_k L) (f_{\xi^k}(x^k) - f_{\xi^k}(x^*)) \\ &\leq (1 - \lambda_k \mu) \|x^k - x^*\|^2. \end{aligned}$$

Furthermore, as $\|\nabla f_{\xi^0}(x^0)\| = \|\nabla f_{\xi^0}(x^0) - \nabla f_{\xi^0}(x^*)\| \leq L\|x^0 - x^*\|$, we also get a better bound on $\mathbb{E}\left[\|x^1 - x^*\|^2\right]$, namely

$$\mathbb{E}\left[\|x^1 - x^*\|^2\right] \leq 2\|x^0 - x^*\| + 2\lambda_0^2 \mathbb{E}\left[\|\nabla f_{\xi^0}(x^0)\|^2\right] \leq 2(1 + \lambda_0^2 L^2) \|x^0 - x^*\|.$$

□