

ESCAPING SADDLE POINTS IN ZERO-ORDER OPTIMIZATION: TWO FUNCTION EVALUATIONS SUFFICE

Anonymous authors
Paper under double-blind review

ABSTRACT

Two-point zeroth order methods are important in many applications of zeroth-order optimization arising in robotics, wind farms, power systems, online optimization, and adversarial robustness to black-box attacks in deep neural networks, where the problem can be high-dimensional and/or time-varying. Furthermore, such problems may be nonconvex and contain saddle points. While existing works have shown that zeroth-order methods utilizing $\Omega(d)$ function evaluations per iteration (with d denoting the problem dimension) can escape saddle points efficiently, it remains an open question if zeroth-order methods based on two-point estimators can escape saddle points. In this paper, we show that by adding an appropriate isotropic perturbation at each iteration, a zeroth-order algorithm based on $2m$ (for any $1 \leq m \leq d$) function evaluations per iteration can not only find ϵ -second order stationary points polynomially fast, but do so using only $\tilde{O}(d/\epsilon^{2.5})$ function evaluations.

1 INTRODUCTION

Two-point (or in general $2m$ -point, where $1 \leq m < d$ with d being the problem dimension) estimators, which approximate the gradient using two (or $2m$) function evaluations per iteration, have been widely studied by researchers in the zeroth-order optimization literature, in convex (Nesterov and Spokoiny, 2017; Duchi et al., 2015; Shamir, 2017), nonconvex (Nesterov and Spokoiny, 2017), online (Shamir, 2017), as well as distributed settings (Tang et al., 2019). A key reason for doing so is that for applications of zeroth-order optimization arising in robotics (Li et al., 2022), wind farms (Tang et al., 2020a), power systems (Chen et al., 2020), online (time-varying) optimization (Shamir, 2017), learning-based control (Malik et al., 2019; Li et al., 2021), and improving adversarial robustness to black-box attacks in deep neural networks (Chen et al., 2017), it may be costly or impractical to wait for $\Omega(d)$ (where d denotes the problem dimension) function evaluations per iteration to make a step. This is especially true for high-dimensional and/or time-varying problems. See Appendix A for more discussion. However, despite the advantages of zeroth-order methods with two-point estimators, there has been a lack of existing work studying the ability of two-point estimators to escape saddle points in nonconvex optimization problems. Since nonconvex problems arise often in practice, it is crucial to know if two-point algorithms can efficiently escape saddle points of nonconvex functions and converge to second-order stationary points.

To motivate the challenges of escaping saddle points using two-point zeroth-order methods, we begin with a review of escaping saddle points using first-order methods. The problem of efficiently escaping saddle points in deterministic first-order optimization (with exact gradients) has been carefully studied in several earlier works (Jin et al., 2017; 2018). A key idea in these works is the injection of an isotropic perturbation whenever the gradient is small, facilitating escape from a saddle if a negative curvature direction exists even without actively identifying the direction. However, the analysis of efficient saddle point escape for stochastic gradient methods is often more complicated. In general, the behavior of the stochastic gradient near the saddle point can be difficult to characterize. Hence, strong concentration assumptions are typically made on the stochastic gradients being used, such as subGaussianity, boundedness of the variance or a bounded gradient estimator (Ge et al., 2015; Daneshmand et al., 2018; Xu et al., 2018; Fang et al., 2019; Roy et al., 2020; Vlaski and Sayed, 2021b), creating an analytical issue when such idealized assumptions fail to hold.

Indeed, though zeroth-order methods can be viewed as stochastic gradient methods, common zeroth order estimators, such as two-point estimators (Nesterov and Spokoiny, 2017), are not subGaussian, and can have unbounded variance. For instance, it can be shown that the variance of the two-point

estimator is on the order of $\Omega(d\|\nabla f(x)\|^2)$ (Nesterov and Spokoiny, 2017), with both a dependence on the problem dimension d as well as on the norm of the gradient, which can be unbounded. Due to non-subGaussianity and unboundedness, it is tricky to bound the effect of such zeroth-order estimators and establish tight concentration inequalities that facilitate its escape near saddle points. In addition, the large variance of the zeroth-order estimator is also an issue away from saddles when the gradient is large. While this is not an issue to show function improvement in expectation, as we discuss later, this becomes an issue when guaranteeing high probability bounds.

Due to these difficulties, previous works on escaping saddle points in zeroth-order optimization have exclusively focused on approaches requiring $\Omega(d)$ function evaluations per iteration to accurately estimate the gradient (Bai et al., 2020; Vlatakis-Gkaragkounis et al., 2019), or in some cases negative curvature directions (Zhang et al., 2022; Lucchi et al., 2021) or the Hessian itself (Balasubramanian and Ghadimi, 2022), reducing in a sense the zeroth-order problem back to a first-order one. However, as explained earlier, two-point or $2m$ -point zeroth-order algorithms are important for high-dimensional and/or time-varying problems in many applications areas. This raises an important question: **Can two-point zeroth-order methods escape saddle points and reach approximate second order stationary points efficiently?**

Our Contribution. In this work, we show that by adding an appropriate isotropic perturbation at each iteration, a zeroth-order algorithm based on *any* number m of pairs (m ranging from 1 to d) of function evaluations per iteration can not only find $(\epsilon, \sqrt{\epsilon})$ -second order stationary points (cf. the definition later in Definition 1) polynomially fast, but do so using only $\tilde{O}(\text{polylog}(\frac{1}{\epsilon})d/\epsilon^{2.5})$ function evaluations, with a probability of at least $1 - \delta$. In particular, this proves that using a single two-point zeroth-order estimator at each iteration (with appropriate perturbation) suffices to efficiently escape saddle points in zeroth-order optimization, with high probability. Moreover, for functions that are (ϵ, ψ) strict-saddle (see Definition 3 for a definition of strict saddle functions), our results become $\tilde{O}(\text{polylog}(\frac{1}{\epsilon})d/\psi\epsilon^2)$, which is a significant improvement when $\psi \gg \epsilon$; strict saddle functions have been identified as an important class of functions in nonconvex optimization, with several well-known examples such as tensor decomposition (Ge et al., 2015), dictionary learning and phase retrieval (Sun et al., 2015). A comparison of our results with existing zeroth-order and first-order methods is shown in Table 1. We also provide numerical results in Appendix G showing that our proposed two-point algorithm requires fewer total function evaluations to converge than zeroth order methods that use $2d$ function evaluations per iteration, for a nonconvex test function proposed in Du et al. (2017).

To overcome the theoretical challenges that were discussed earlier, we i) first show, via a careful analysis, that zeroth order methods can make function value improvement across iterates with large gradients with high probability, even when only a single two-point estimator (which can have significant variance at large gradients) is used per iteration. ii) Second, near saddle points, we overcome issues caused by the unbounded variance and non-subGaussianity of zeroth-order gradient estimators by developing new technical tools, including novel martingale concentration inequalities involving Gaussian vectors, to tightly bound such terms. In turn, this allows us to show that the noise emanating from the zeroth-order estimators will not overwhelm the effect of the additional isotropic perturbative noise, facilitating escape along negative curvature directions. To the best of our knowledge, both analyses are novel, and may be independent contributions on their own.

Related Work. Due to space considerations, we defer a full discussion of related work to Appendix A.

2 PROBLEM SETUP

We make the following assumptions on the class of functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ which we consider.

Assumption 1 (Properties of f). *We suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies the following properties:*

1. f is twice-differentiable and lower bounded, i.e. $f^* := \min_x f(x) > -\infty$.
2. f is L -Lipschitz, i.e. $\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \quad \forall x, y \in \mathbb{R}^d$
3. f is ρ -Hessian Lipschitz, i.e. $\|\nabla^2 f(x) - \nabla^2 f(y)\| \leq \rho\|x - y\| \quad \forall x, y \in \mathbb{R}^d$.

In our work, we focus on finding approximate second order stationary points, defined below.

Definition 1. A point $x \in \mathbb{R}^d$ is an (ϵ, φ) -second order stationary point if

$$\|\nabla f(x)\| < \epsilon, \quad \text{and} \quad \lambda_{\min}(\nabla^2 f(x)) > -\varphi.$$

		Iteration Complexity	Fun. evaluations per iter.
First-order	Jin et al. (2017) (deterministic)	$\tilde{O}\left(\frac{1}{\epsilon^2}\right)$	—
	Fang et al. (2019) (SGD)	$\tilde{O}\left(\frac{1}{\epsilon^{3.5}}\right)$	—
Zeroth-order	Bai et al. (2020)	$\tilde{O}\left(\frac{1}{\epsilon^2}\right)$	$\tilde{\Omega}\left(\frac{d^2}{\epsilon^8}\right)$
	Vlatakis-Gkaragkounis et al. (2019)	$\tilde{O}\left(\frac{1}{\epsilon^2}\right)$	$\tilde{\Omega}(d)$
	Balasubramanian and Ghadimi (2022)	$\tilde{O}\left(\frac{1}{\epsilon^{3/2}}\right)$	$\tilde{O}\left(\frac{d}{\epsilon^2} + \frac{d^4}{\epsilon}\right)$
	Lucchi et al. (2021) [†]	$\tilde{O}\left(\frac{1}{\epsilon^2}\right)$	$\tilde{O}\left(\frac{d}{\epsilon^{2/3}}\right)$
	Zhang et al. (2022)	$\tilde{O}\left(\frac{1}{\epsilon^2}\right)$	$\tilde{\Omega}(d)$
	Algorithm 1 (this paper) [‡]	$\tilde{O}\left(\frac{d}{\epsilon^2\psi m}\right)$	$2m$

Table 1: Selected comparison of convergence results to $(\epsilon, O(\sqrt{\epsilon})$ -second order stationary points in smooth, nonconvex functions; for [†], the convergence is to $(\epsilon, \epsilon^{2/3})$ -second order stationary points. For [‡], the term $\tilde{\psi}$ in the denominator is (i) ψ when the function f is (ϵ, ψ) -strict saddle for a $\psi > O(\sqrt{\epsilon})$ (see Definition 3 for a definition) and (ii) $O(\sqrt{\epsilon})$ if otherwise.

We define an (ϵ, φ) -approximate saddle point as follows.

Definition 2. A point $x \in \mathbb{R}^d$ is an (ϵ, φ) -approximate saddle point, if

$$\|\nabla f(x)\| < \epsilon, \quad \text{and} \quad \lambda_{\min}(\nabla^2 f(x)) \leq -\varphi.$$

Following past convention (Jin et al., 2019a), we will focus in particular on escaping $(\epsilon, \sqrt{\rho\epsilon})$ -saddle points. For notational simplicity, in following text, we refer to $(\epsilon, \sqrt{\rho\epsilon})$ -saddle points simply as ϵ -saddle points and $(\epsilon, \sqrt{\rho\epsilon})$ -second order stationary points as ϵ -second order stationary points. Beyond the definition of ϵ -approximate saddle points above, it is known that many nonconvex functions with saddle points, such as orthogonal tensor decomposition (Ge et al., 2015), phase retrieval and dictionary learning (Sun et al., 2015), satisfy what is known as a *strict saddle* condition (Ge et al., 2015). For the Hessians of the saddle points of such functions, there is always a strict negative eigenvalue whose magnitude is bounded from below. We provide a precise definition below.

Definition 3. A twice-differential function $f(x)$ is (ϵ, ψ) -strict saddle, if for any point x , at least one of the following is true:

$$\|\nabla f(x)\| \geq \epsilon \quad \text{or} \quad \lambda_{\min}(\nabla^2 f(x)) \leq -\psi$$

As a corollary of Theorem 1 (to be stated later), for functions f which are (ϵ, ψ) strict saddle, assuming that $\psi \geq \sqrt{\rho\epsilon}$, the sample complexity of our algorithm scales as $\tilde{\Omega}\left(\frac{d \max\{L^2, L\}(f(x_0) - f^*)}{m\epsilon^2\psi}\right)$, which scales as $\tilde{\Omega}\left(\frac{d}{m\epsilon^2}\right)$ when ψ is of size $\Omega(1)$. Thus, in this setting, for two-point estimators, where $m = 1$, the dependence on d and ϵ in our sample complexity (as measured by function evaluations) matches that achieved by the algorithms in Vlatakis-Gkaragkounis et al. (2019); Zhang et al. (2022), which have to use $2d$ function evaluations per iteration to estimate the gradient.

In our work, we consider the following batch symmetric two-point zeroth-order estimator.

Definition 4 ((Batch) two-point zeroth-order estimator with perturbation). We define a m -batch two-point zeroth order estimator as follows:

$$g_u^{(m)}(x) := \frac{1}{m} \sum_{i=1}^m \frac{f(x + uZ_i) - f(x - uZ_i)}{2u} Z_i, \quad (1)$$

where $Z_i \stackrel{i.i.d}{\sim} N(0, I)$, and $u > 0$ is a smoothing radius.

Such $2m$ zeroth-order gradient estimators have frequently been studied in zeroth-order optimization works (see e.g. Nesterov and Spokoiny (2017)). To facilitate efficient escape from saddle points, our proposed Algorithm 1 adds isotropic perturbation at each iteration.

Algorithm 1: Zeroth-order perturbed gradient descent (ZOPGD)

input : x_0 , horizon T , step-size η , smoothing radius u , perturbation radius r , batch size m

for step $t = 0, \dots, T_1$ **do**

Sample $Z^{(m)} = \{Z_{t,i}\}_{i=1}^m \sim N(0, I)$ to compute $g_u^{(m)}(x_t)$.

Update $x_{t+1} = x_t - \eta \left(g_u^{(m)}(x_t) + Y_t \right)$, where $Y_t \sim N(0, \frac{r^2}{d} I)$

We now state an informal version of our main result, and follow that with a few remarks.

Theorem 1 (Main result, informal version of Theorem 2). *Consider running Algorithm 1. Let \tilde{O} hide polylogarithmic terms in δ and other parameters. Suppose $\delta \in (0, 1/e]$. Suppose $\sqrt{\rho\epsilon} \leq \min\{1, L\}$ ¹, such that $\bar{\psi} \leq \min\{1, L\}$, where*

$$\bar{\psi} := \begin{cases} \min\{\psi, 1, L\} & \text{if } f(\cdot) \text{ is } (\epsilon, \psi)\text{-strict saddle for any } \psi > \sqrt{\rho\epsilon} \\ \sqrt{\rho\epsilon} & \text{otherwise.} \end{cases} \quad (2)$$

Suppose $u = \tilde{O}\left(\frac{\min\{\sqrt{\epsilon}, \sqrt{r}\}}{\sqrt{\rho d}}\right)$, $r = \tilde{O}(\epsilon)$, $\eta = \tilde{O}\left(\frac{m\bar{\psi}}{d \max\{L, L^2\}}\right)$. Then, in

$$\begin{aligned} T &= \tilde{\Omega}\left(\frac{(f(x_0) - f^*)}{\eta\epsilon^2} + \frac{\rho^2(f(x_0) - f^*)}{\eta\bar{\psi}^4}\right) \\ &= \tilde{\Omega}\left(\frac{d \max\{L, L^2\}(f(x_0) - f^*)}{m\bar{\psi}\epsilon^2} + \frac{d \max\{L, L^2\}\rho^2(f(x_0) - f^*)}{m\bar{\psi}^5}\right) \\ &= \tilde{\Omega}\left(\frac{d \max\{L, L^2\}\rho^2(f(x_0) - f^*)}{m\bar{\psi}\epsilon^2}\right) \end{aligned}$$

iterations (with each iteration using $2m$ function evaluations), with probability at least $1 - \delta$, at least half the iterates are ϵ -approximate second-order stationary points.

Remark 1. As the choice of η in Proposition 4 (Appendix D) and Theorem 2 (Appendix F) respectively imply, the $\tilde{\Omega}\left(\frac{f(x_0) - f^*}{\eta\epsilon^2}\right)$ term in the sample complexity comes from the large gradient iterations (Proposition 4), whereas the $\tilde{\Omega}\left(\frac{\rho^2(f(x_0) - f^*)}{\eta\bar{\psi}^4}\right)$ term comes from the escape saddle point phase.

Comparison to gradient-based methods. For first-order escape saddle point algorithms, standard perturbation-based methods (without acceleration) can find a $(\epsilon, O(\sqrt{\epsilon}))$ -second-order stationary point using $\tilde{O}(1/\epsilon^2)$ iterations for deterministic GD (Jin et al., 2019a), while for standard SGD the best-known rates are slower at $\tilde{O}(1/\epsilon^{3.5})$ (Fang et al., 2019). In contrast, our sample complexity (as measured by the total number of function evaluations) is $\tilde{O}\left(\frac{d}{\epsilon^2\bar{\psi}}\right)$, where $\bar{\psi}$ is defined in Eq. (2). The extra (linear) dependence on d is typical for zeroth-order algorithms (see e.g. Nesterov and Spokoiny (2017)); intuitively, gradient calculation for d -dimensional functions requires $O(d)$ calculations agnostically, so it makes sense that zeroth-order algorithms requires d times more iterations. For general non strict-saddle functions, our dependence on ϵ sits between that of the deterministic methods and SGD methods, and suggests the benefit of a specialized treatment of zeroth-order methods over considering them simply as a subclass of SGD methods. Moreover, for (ϵ, ψ) -strict-saddle functions where $\psi = \Omega(1)$, our sample complexity becomes $\tilde{O}\left(\frac{d}{\epsilon^2}\right)$, with an ϵ dependence that matches that of the best existing sample complexity for non-accelerated first-order escape saddle point methods Jin et al. (2017)

Comparison to existing zeroth-order methods. As Table 1 suggests, our sample complexity significantly outperforms that of Bai et al. (2020) and Balasubramanian and Ghadimi (2022), and also

¹In our paper, we focus on the case $\sqrt{\rho\epsilon} \leq L$; otherwise, by the L -Lipschitz assumption, $\lambda_{\min}(\nabla^2 f(x)) \geq -L$ for all $x \in \mathbb{R}^d$, which implies ϵ -first order stationary points are also ϵ -second order stationary points.

that in Lucchi et al. (2021), which is a random search method. The sample complexity in Vlatakis-Gkaragkounis et al. (2019); Zhang et al. (2022) outperform our method, with a function evaluation complexity of $\tilde{O}\left(\frac{d}{\epsilon^2}\right)$. However, for (ϵ, ψ) -strict-saddle functions where $\psi = \Omega(1)$, our sample complexity becomes $\tilde{O}\left(\frac{d}{\epsilon^2}\right)$, which matches the sample complexity in Vlatakis-Gkaragkounis et al. (2019); Zhang et al. (2022). Moreover, a key limitation of their methods is a requirement to use $\Omega(d)$ function evaluations to estimate the gradient at each iteration, which may not be practical in realistic applications when d is large. In contrast, our method supports any number of function evaluations at each iteration between 1 to d . Moreover, numerically, we found that for a test nonconvex function proposed in Du et al. (2017), our method (with two-point estimators) takes fewer function evaluations to escape saddle points and converge to the global minimum than the methods in Vlatakis-Gkaragkounis et al. (2019); Zhang et al. (2022).

3 PROOF STRATEGY AND KEY CHALLENGES IN THE ZERO-ORDER SETTING

Broadly speaking, our proof include two major parts, i) characterizing the progress made in iterations when the gradient is large (which we can define to be iterations t where $\|\nabla f(x_t)\| \geq \epsilon$) (Section 3.1), ii) and iterations when we are at an ϵ -approximate saddle point (where progress may be made along the negative eigendirection of the Hessian matrix) (Section 3.2). While the approach is similar to the first-order case (e.g. Jin et al. (2019a)), the zeroth-order setting brings forth several challenges as we explain later in the individual subsections. In the rest of this section, we explain these challenges, sketch out our high-level proof outlines, and provide statements of the main technical results. Note that due to the space limit, the full proof is provided in the Appendix.

3.1 SHOWING FUNCTION DECREASE WHEN GRADIENTS ARE LARGE

Challenge. Due to the noise in two-point (or $2m$ where m is a small constant) zeroth-order gradient, even when the gradient is large, it may not always be possible to make progress at each iteration, especially when $m < d$ is used in the gradient estimation equation in Eq. (1). While it is tempting to use an expectation-based argument to handle this issue, it is known that expectation-based function decrease arguments are insufficient for escape saddle point purposes (see e.g. Proposition 1 in Ziyin et al. (2021)). We tackle this issue by using high-probability arguments instead; we note that achieving these high-probability bounds is highly nontrivial due to the large variance of the two-point zeroth-order estimator (scaling with d times the squared norm of the gradient). Hence, any single iteration of the zeroth-order method may in fact lead to a function increase rather than decrease.

High-level proof outline. (i) We first characterize the function value change for our proposed algorithm (Lemma 1). (ii) Next, we tackle the issue of the possibility that the function value might increase for any given iteration. The key idea here is that across any small consecutive number of iterations, there will be one iteration where the zeroth-order estimator is sufficiently aligned with the gradient direction (Lemma 14 in Appendix D). (iii) Along with a series of other technical results in Appendix D, we then show that the function makes sufficient progress across the duration of the algorithm, with high probability (Proposition 1). To more concretely illustrate the key analytical challenge, we next introduce the following function decrease lemma, proved in Appendix D.

Lemma 1 (Function decrease for batch zeroth-order optimization). *Suppose at each time t , the algorithm performs the update step (with batch-size parameter $1 \leq m \leq d$)*

$$x_{t+1} = x_t - \eta \left(g_u^{(m)}(x_t) + Y_t \right),$$

where

$$g_u^{(m)}(x_t) = \frac{1}{m} \sum_{i=1}^m \frac{f(x_t + uZ_{t,i}) - f(x_t - uZ_{t,i})}{2u} Z_{t,i},$$

where each $Z_{t,i}$ is drawn i.i.d from $N(0, I)$, $u > 0$ is the smoothing radius, and $Y_t \sim N(0, \frac{r^2}{d}I)$ with $r > 0$ denoting the perturbation radius.

Then, there exist absolute constants $c_1 > 0, C_1 \geq 1$ such that, for any $T \in \mathbb{Z}^+$ and $T \geq \tau > 0$, $\alpha > 0$ and $\delta \in (0, 1/e]$, upon defining $\mathcal{H}_{0,\tau}(\delta)$ to be the event on which the inequality

$$f(x_\tau) - f(x_0) \leq -\frac{3\eta}{4} \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{\alpha} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2$$

$$\begin{aligned}
& + \tau\eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
& + \eta c_1 r^2 (\alpha + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2
\end{aligned} \tag{3}$$

is satisfied (where $\chi := \log(C_1 d m T / \delta)$), we have

$$\mathbb{P}(\mathcal{H}_{0,\tau}(\delta)) \geq 1 - \frac{(\tau + 4)\delta}{T}, \quad \mathbb{P}(\cap_{\tau=1}^{\tau'} \mathcal{H}_{0,\tau}(\delta)) \geq 1 - \frac{5\tau'\delta}{T}$$

for any $0 \leq \tau' \leq T$.

Our goal is to show that we can arrive at a contradiction $f(x_T) < \min_x f(x)$ when there is a large number of steps at which $\|\nabla f(x_t)\| \geq \epsilon$ (Proposition 1). As we can see from Eq. (3), this implies that we need to prove a lower bound of the form

$$\sum_{t=0}^{T-1} \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2 \geq \Omega \left(\frac{1}{\alpha} + \frac{c_1 L \eta \chi^3 d}{m} \right) \sum_{t=0}^{T-1} \|\nabla f(x_t)\|^2 \tag{4}$$

for some α which is not too large (an example would be picking α such that it only scales logarithmically in the problem parameters). However, it is tricky to prove such a lower-bound in the zeroth-order setting. In particular, for small batch-sizes m , $\frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2$ could be small even as $\|\nabla f(x_t)\|^2$ is large; this is because for each $i \in [m]$, $Z_{t,i}$ could have a negligible component in the $\nabla f(x_t)$ direction. This necessitates a more delicate analysis to prove a bound similar to Eq. (4). Due to space reasons, we defer our more detailed proof approach outline to Appendix D (see the discussion immediately following Lemma 1) The results in Appendix D culminates in the following result which limits the number of large-gradient.

Proposition 1 (Bound on number of iterates with large gradients, informal version of Proposition 4). *Let $\delta \in (0, 1/e]$ be arbitrary. Letting \tilde{O} hide polylogarithmic dependencies on δ (and other parameters), consider choosing u, r, η and T such that*

$$u = \tilde{O} \left(\frac{\sqrt{\epsilon}}{\sqrt{\rho d}} \right), \quad r = O(\epsilon), \quad \eta = \tilde{O} \left(\frac{m}{dL} \right), \quad T = \tilde{\Omega} \left(\frac{((f(x_0) - f^*) + \epsilon^2/L)}{\eta \epsilon^2} \right).$$

Then, with probability at least $1 - O(\delta)$, there are at most $T/4$ iterations for which $\|\nabla f(x_t)\| \geq \epsilon$.

3.2 MAKING PROGRESS NEAR SADDLE POINTS

Challenge. The noise in two-point zeroth-order estimators makes the analysis around ϵ -approximate saddle points challenging, because the concentration properties of the (non-subGaussian) noise are hard to characterize. Intuitively, a noisier estimator might facilitate easier escape from saddle point. However, without an appropriate concentration bound, the noise may behave in unpredictable ways, preventing escape from saddle regions. Previous analysis of saddle point escape using stochastic estimators typically requires these estimators to satisfy subGaussian properties (Jin et al., 2019a; Fang et al., 2019), which zeroth-order estimators do not satisfy.

High-level proof outline. (i) We first prove a technical result showing that the travelling distance of the iterates can be bounded in terms of the function value decrease (i.e., Improve or Localize, Lemma 2). (ii) Next, at any ϵ -saddle point, we consider a coupling argument and define two sequences running near-identical zeroth-order dynamics, differing only in the sign of their perturbative term along the minimum eigendirection of H , which denotes the Hessian of the saddle (Lemma 3). Using Lemma 2 in point (i), if we assume for contradiction that the two sequences both “get stuck” and make little function value progress, the dynamics of the difference between the two sequences will remain small as both sequences remain close to the saddle point. (iii) However, since the perturbation vectors of the two sequences differ in the (most) negative direction of H , the norm of the the difference of the two sequences will grow exponentially so long as a). the sequences remain close to the saddle point (and thus the Hessian has a negative curvature direction) and b). the effect of the zeroth-order stochastic noise can be controlled. This leads to a contradiction, implying that sufficient function decrease must have been made (Proposition 5 in Appendix E.3). (iv) To show that the zeroth-order stochastic noise can be controlled, we prove one technical result (Proposition 2), providing a

concentration bound for the product of (possibly unbounded) subGaussian random vectors that scales linearly with the dimension d . This enables us to control the effect of the zeroth-order noise near saddle points, and is essential in showing that the eventual sample complexity scales linearly with d .

We provide a more detailed proof sketch below, where we elaborate more on our analytical challenges and ideas. We first introduce an informal statement of a key technical result that bounds, with high probability, the travelling distance of the iterates in terms of the function value decrease.

Lemma 2 (Improve or Localize, informal version of Lemma 23). *Consider the perturbed zeroth-order update Algorithm 1. Let $\delta \in (0, 1/e]$ be arbitrary. Consider any $T_s = \tilde{\Omega}\left(\frac{1}{m} \log(1/\delta)\right)$, and any $t_0 \geq 0$. For any $F > 0$, suppose $f(x_{T_s+t_0}) - f(x_{t_0}) > -F$, i.e. $f(x_{t_0}) - f(x_{T_s+t_0}) < F$. Letting \tilde{O} hide polylogarithmic terms involving δ , suppose*

$$u = \tilde{O}\left(\frac{\min\{\sqrt{\epsilon}, \sqrt{r}\}}{\sqrt{\rho d}}\right), \quad r = \tilde{O}\left(\min\left\{\epsilon, \frac{F}{\eta T_s}\right\}\right), \quad \eta = \tilde{O}\left(\frac{m\sqrt{\rho\epsilon}}{dL}\right).$$

Then, with probability at least $1 - O\left(\frac{T_s\delta}{T}\right)$ (here $T \geq T_s$ denotes the total number of iterations), for each $\tau \in \{0, 1, \dots, T_s\}$, we have that

$$\|x_{t_0+\tau} - x_{t_0}\|^2 \leq \phi_{T_s}(\delta, F), \quad \text{where } \phi_{T_s}(\delta, F) = \tilde{O}\left(\max\left\{T_s, \frac{d}{m}\right\}\right) \eta F + \tilde{O}(\eta^2 \epsilon^2).$$

Intuitively, the above result shows that if little function value improvement has been made, then the algorithm's iterates have not moved much, such that it remains approximately in a saddle region if it started out in a saddle region. Next, Lemma 3 formally introduces the coupling we have mentioned, setting the stage for the rest of our arguments. For notational convenience, in this section, unless otherwise specified, we will assume that the initial iterate x_0 is an ϵ -saddle point.

Lemma 3. *Suppose x_0 is an ϵ -approximate saddle point. Without loss of generality, suppose that the minimum eigendirection of $H := \nabla^2 f(x_0)$ is the e_1 direction, and let γ to denote $-\lambda_{\min}(\nabla^2 f(x_0))$ (note $\gamma \geq \sqrt{\rho\epsilon}$). Consider the following coupling mechanism, where we run the zeroth-order gradient dynamics, starting with x_0 , with two isotropic noise sequences, Y_t and Y'_t respectively, where $(Y_t)_1 = -(Y'_t)_1$, and $(Y_t)_j = (Y'_t)_j$ for all other $j \neq 1$. Suppose that the sequence $\{Z_{t,i}\}_{t \in T, i \in [m]}$ is the same for both sequences. Let $\{x_t\}$ denote the sequence with the $\{Y_t\}$ noise sequence, and let the $\{x'_t\}$ denote the sequence with the $\{Y'_t\}$ noise sequence, where*

$$x'_{t+1} = x'_t - \eta \left(\frac{1}{m} \sum_{i=1}^m \left(Z_{t,i} Z_{t,i}^\top \nabla f(x'_t) + \frac{u}{2} Z_{t,i} Z_{t,i}^\top \tilde{H}'_{t,i} Z_{t,i} \right) + Y'_t \right), \quad x'_0 = x_0,$$

and $\tilde{H}'_{t,i} := \frac{H'_{t,i,+} - H'_{t,i,-}}{2}$, with $H'_{t,i,+} = \nabla^2 f(x'_t + \alpha'_{t,i,+} u Z'_i)$ for some $\alpha'_{t,i,+} \in [0, 1]$, and $H'_{t,i,-} = \nabla^2 f(x'_t - \alpha'_{t,i,-} u Z'_i)$ for some $\alpha'_{t,i,-} \in [0, 1]$. Then, for any $t \geq 0$,

$$\begin{aligned} \hat{x}_{t+1} := x_{t+1} - x'_{t+1} &= -\eta \underbrace{\sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}_{g_0}(\tau)}_{W_{g_0}(t+1)} - \eta \underbrace{\sum_{\tau=0}^t (I - \eta H)^{t-\tau} (\bar{H}_\tau - H) \hat{x}_\tau}_{W_H(t+1)} \\ &\quad - \eta \underbrace{\sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}_u(\tau)}_{W_u(t+1)} - \eta \underbrace{\sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{Y}_\tau}_{W_p(t+1)} \end{aligned}$$

where

$$\xi_{g_0}(t) = \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t), \quad \xi'_{g_0}(t) = \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x'_t), \quad \hat{\xi}_{g_0}(t) = \xi_{g_0}(t) - \xi'_{g_0}(t),$$

$$\xi_u(t) = \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i}, \quad \xi'_u(t) = \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i}^\top \tilde{H}'_{t,i} Z_{t,i}, \quad \hat{\xi}_u(t) = \xi_u(t) - \xi'_u(t),$$

$$\hat{Y}_t = Y_t - Y'_t, \quad \bar{H}_t = \int_0^1 \nabla^2 f(ax_t + (1-a)x'_t) da.$$

Our goal is to show that the dominating term in the evolution of the difference dynamics comes from the W_p term involving the additional perturbation. To this end, we need to bound the remaining terms, W_{g_0}, W_H, W_u . A key technical challenge is to find a precise concentration bound for the $W_{g_0}(t+1)$ term, where

$$W_{g_0}(t+1) = -\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \left(\frac{1}{m} \sum_{i=1}^m (Z_{\tau,i} Z_{\tau,i}^\top - I) (\nabla f(x_\tau) - \nabla f(x'_\tau)) \right).$$

For the simplicity of discussion, we assume for the time being that $m = 1$, and drop the i index in the subscript of $Z_{\tau,i}$. Since $\mathbb{E}[Z_\tau Z_\tau^\top] = I$, heuristically, assuming that $Z_\tau Z_\tau^\top - I$ satisfies “nice” concentration properties, utilizing the independence of the Z_τ ’s across time and the fact that $(I - \eta H) \preceq (1 + \eta\gamma)I$, we would like to show that with high probability,

$$\|W_{g_0}(t)\| \lesssim \eta \sqrt{\sum_{\tau=0}^{t-1} (1 + \eta\gamma)^{2(t-1-\tau)} \mathbb{E} \left[\|(Z_\tau Z_\tau - I) (\nabla f(x_\tau) - \nabla f(x'_\tau))\|^2 \mid \mathcal{F}_{\tau-1} \right]} \quad (5)$$

where $\mathcal{F}_{\tau-1}$ is a sigma-algebra containing all randomness up to and including iteration $\tau - 1$, such that x_τ and x'_τ are both in $\mathcal{F}_{\tau-1}$, but Z_τ is not. Then, assuming that Eq. (5) holds, since

$$\mathbb{E} \left[\|(Z_\tau Z_\tau - I) (\nabla f(x_\tau) - \nabla f(x'_\tau))\|^2 \mid \mathcal{F}_{\tau-1} \right] = O(d) \|\nabla f(x_\tau) - \nabla f(x'_\tau)\|^2,$$

it follows that

$$\|W_{g_0}(t)\| \leq \eta \sqrt{O(d) \sum_{\tau=0}^{t-1} (1 + \eta\gamma)^{2(t-1-\tau)} \|\nabla f(x_\tau) - \nabla f(x'_\tau)\|^2}$$

With this bound on $\|W_{g_0}(t)\|$, we eventually prove in Proposition 5 in Appendix E.3 that our algorithm escapes any ϵ -saddle point with constant probability and that the $O(d)$ term appearing in the square root term above will eventually lead to an $O(d)$ dependence in the sample complexity². We note that the $O(d)$ dimension dependence matches that of the best-known existing upper bound for finding first-order stationary points in smooth nonconvex zeroth-order optimization (Nesterov and Spokoiny, 2017), and has been conjectured to be the best possible dimension dependence for general smooth nonconvex zeroth-order optimization (Balasubramanian and Ghadimi, 2022).

Key technical challenge The key challenge in the above argument is to show that an equation in the form of Eq. (5) could in fact hold. At first glance, that an inequality such as Eq. (5) should hold is rather non-obvious — this is because while the variable $(Z_\tau Z_\tau - I) (\nabla f(x_\tau) - \nabla f(x'_\tau)) \mid \mathcal{F}_{\tau-1}$ is mean-zero, it is subExponential rather than subGaussian. In fact, even in the subGaussian case, given a sequence of random vectors $\mathbf{x}_0, \dots, \mathbf{x}_{t-1}$, such that each $\mathbb{E}[\mathbf{x}_\tau \mid \mathcal{F}_{\tau-1}] = 0$, and that each $\mathbf{x}_\tau \mid \mathcal{F}_{\tau-1}$ is norm-subGaussian with parameter $\sigma_\tau \in \mathcal{F}_{\tau-1}$ (which is an appropriate generalization of subGaussianity for vectors, proposed in Jin et al. (2019b)), proving a concentration inequality of the form $\left\| \sum_{\tau=0}^{t-1} \mathbf{x}_\tau \right\| \approx \tilde{O} \left(\sqrt{\sum_{\tau=0}^{t-1} \sigma_\tau^2} \right)$ is a very delicate matter. In our case, the analogue of \mathbf{x}_τ is $(I - \eta H)^{t-1-\tau} (Z_\tau Z_\tau - I) (\nabla f(x_\tau) - \nabla f(x'_\tau))$, while the analogue of σ_τ^2 is $(1 + \eta\gamma)^{2(t-1-\tau)} \mathbb{E} \left[\|(Z_\tau Z_\tau - I) (\nabla f(x_\tau) - \nabla f(x'_\tau))\|^2 \mid \mathcal{F}_{\tau-1} \right]$. Existing techniques (cf. Tropp et al. (2015); Jin et al. (2019b)) rely crucially on subGaussian properties that allow for each τ the moment-generating function $\mathbb{E}[e^{\theta \mathbf{Y}_\tau} \mid \mathcal{F}_{\tau-1}]$ to be defined for any fixed (and non-random) $\theta > 0$, where \mathbf{Y}_τ takes the form

$$\mathbf{Y}_\tau = \begin{bmatrix} 0 & \mathbf{x}_\tau^\top \\ \mathbf{x}_\tau & 0 \end{bmatrix},$$

such that $\mathbb{E}[\mathbf{Y}_\tau \mid \mathcal{F}_{\tau-1}] = 0$ (since $\mathbb{E}[\mathbf{x}_\tau \mid \mathcal{F}_{\tau-1}] = 0$), and the eigenvalues of \mathbf{Y}_τ are $\pm \|\mathbf{x}_\tau\|$. In the case when \mathbf{x}_τ is merely subExponential, the Moment Generating Function (MGF), $\mathbb{E}[e^{\theta \mathbf{Y}_\tau} \mid \mathcal{F}_{\tau-1}]$, will no longer be well-defined at any fixed (and non-random) $\theta > 0$. This poses a challenge in our setting, since \mathbf{x}_τ takes the form $(I - \eta H)^{t-1-\tau} (Z_\tau Z_\tau^\top - I) (\nabla f(x_\tau) - \nabla f(x'_\tau))$, which is subExponential rather than subGaussian. While it may be possible to force $(I - \eta H)^{t-1-\tau} (Z_\tau Z_\tau^\top - I) (\nabla f(x_\tau) - \nabla f(x'_\tau))$ to be sub-Gaussian, say by normalizing Z_τ to have norm \sqrt{d} (note any

²For general $1 \leq m \leq d$, there will also be an $O(1/m)$ dependence in the sample complexity.

bounded random vector is also subGaussian), such that $\|(Z_\tau Z_\tau^\top - I)g\|^2 \leq d^2 \|g\|^2$ for any vector $g \in \mathbb{R}^d$, a careful examination of the argument in Proposition 5 would show that this results in a $O(d^2)$ rather than $O(d)$ dependence in the sample complexity, incurring a heavy price on the overall sample complexity (extra factor of d) if d is large.

Our solution To overcome the issue, we build on the following observation: with high probability, for any vector $g \in \mathbb{R}^d$, $|Z_\tau^\top g|$ is bounded within some log factor of $\|g\|$. On the event $\{|Z_\tau^\top g| = \tilde{O}(\|g\|)\}$, the variable

$$(Z_\tau Z_\tau^\top - I)g = Z_\tau(Z_\tau^\top g) - g \approx Z_\tau \|g\| - g$$

behaves approximately like a subGaussian random vector since $Z_\tau \sim N(0, I_d)$. Based on this intuition, after some careful analysis, we can show that $(Z_\tau Z_\tau^\top - I)(\nabla f(x_\tau) - \nabla f(x'_\tau)) | \mathcal{F}_{\tau-1}$ is subGaussian on the event that $|Z_\tau^\top \nabla f(x_\tau)|$ is bounded within some log factor of $\|\nabla f(x_\tau)\|$, which happens with high probability. This then allows us to show that on this event, the corresponding MGF is well-defined for all fixed $\theta > 0$, enabling us to prove a concentration inequality of the form Eq. (5). This intuition is crystallized in the following proposition, which proves a more general bound than what we strictly need. For notational simplicity, we introduce the function $\text{lr}(x) := \log(x \log(x))$.

Proposition 2. *Let \mathcal{F}_t , $t \geq -1$ be a filtration. Let $(Z_t)_{t \geq 0}$ be a sequence of random vectors following the distribution $N(0, I)$ such that $Z_t \in \mathcal{F}_t$ and is independent of \mathcal{F}_{t-1} , and let $(v_t)_{t \geq 0}$ be a sequence of random vectors such that $v_t \in \mathcal{F}_{t-1}$. For each $\tau \geq 0$, let*

$$W_\tau = \sum_{t=0}^{\tau-1} M_t (Z_t Z_t^\top - I) v_t,$$

where each M_t is a deterministic matrix of appropriate dimension. Then, there exist some absolute constants $c', C > 0$ such that for any $\tau \in \mathbb{Z}^+$ and $\delta \in (0, 1/e]$, the following statements hold:

1. For any $\theta > 0$, with probability at least $1 - \delta$, we have

$$\|W_\tau\| \leq c' \theta \sum_{t=0}^{\tau-1} \|M_t\|_2^2 d (\text{lr}(C\tau/\delta))^2 \|v_t\|^2 + \frac{1}{\theta} \log(Cd\tau/\delta).$$

2. For any $B > b > 0$, with probability at least $1 - \delta$,

$$\text{either } \sum_{t=0}^{\tau-1} \|M_t\|_2^2 d (\text{lr}(C\tau/\delta))^2 \|v_t\|^2 \geq B$$

$$\text{or } \|W_\tau\| \leq c' \sqrt{\max \left\{ \sum_{t=0}^{\tau-1} \|M_t\|_2^2 d (\text{lr}(C\tau/\delta))^2 \|v_t\|^2, b \right\}} (\log(C\tau d/\delta) + \log(\log(B/b) + 1))$$

Moreover, as is clear from the bounds above, we may pick $C \geq 1$ such that $\log(\frac{C}{\delta}) \geq 1, \forall \delta \in (0, \frac{1}{e}]$.

With this result, along with a series of other technical results in Appendix E.3, we can show that the algorithm makes a function decrease of F with $\Omega(1)$ probability near an ϵ -saddle point (Proposition 5 in Appendix E.3). Armed with Proposition 5, as well as Proposition 1, the main result in Theorem 1 then follows. The complete detailed analysis can be found in Appendix E (Escaping saddle point) and Appendix F (main result).

4 CONCLUSION

In this paper, we proved that using two function evaluations per iteration suffices to escape saddle points and reach approximate second order stationary points efficiently in zeroth-order optimization. Along the way, we gave the first analysis of high-probability function change using two(or more)-point zeroth-order gradient estimators, as well as a novel concentration bound for sums of subExponential (but not subGaussian) vectors which are each the products of Gaussian vectors. These technical contributions may be of independent interest to researchers working in zeroth-order optimization as well as general stochastic optimization. There are a few limits of the current results which lead to several interesting future directions, such as extension to noisy function evaluations, as well as studying if some zeroth-order estimators such as asymmetric two-point estimators (Nesterov and Spokoiny, 2017) or single-point estimators (Flaxman et al., 2005) could actually escape saddle points without additional perturbation noise.

REFERENCES

- Adolphs, L., Daneshmand, H., Lucchi, A., and Hofmann, T. (2019). Local saddle point optimization: A curvature exploitation approach. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 486–495. PMLR.
- Antonakopoulos, K., Mertikopoulos, P., Piliouras, G., and Wang, X. (2022). Adagrad avoids saddle points. In *International Conference on Machine Learning*, pages 731–771. PMLR.
- Aydiukhin, D., Jin, C., and Yaroslavtsev, G. (2019). Escaping saddle points with inequality constraints via noisy sticky projected gradient descent. In *11th Annual Workshop on Optimization for Machine Learning*.
- Aydiukhin, D. and Yaroslavtsev, G. (2021). Escaping saddle points with compressed sgd. *Advances in Neural Information Processing Systems*, 34:10273–10284.
- Bai, Q., Agarwal, M., and Aggarwal, V. (2020). Escaping saddle points for zeroth-order non-convex optimization using estimated gradient descent. In *2020 54th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE.
- Balasubramanian, K. and Ghadimi, S. (2022). Zeroth-order nonconvex stochastic optimization: Handling constraints, high dimensionality, and saddle points. *Foundations of Computational Mathematics*, 22(1):35–76.
- Chen, P.-Y., Zhang, H., Sharma, Y., Yi, J., and Hsieh, C.-J. (2017). Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pages 15–26.
- Chen, Y., Bernstein, A., Devraj, A., and Meyn, S. (2020). Model-free primal-dual methods for network optimization with application to real-time optimal power flow. In *2020 American Control Conference (ACC)*, pages 3140–3147. IEEE.
- Crisciatiello, C. and Boumal, N. (2019). Efficiently escaping saddle points on manifolds. *Advances in Neural Information Processing Systems*, 32.
- Daneshmand, H., Kohler, J., Lucchi, A., and Hofmann, T. (2018). Escaping saddles with stochastic gradients. In *International Conference on Machine Learning*, pages 1155–1164. PMLR.
- Du, S. S., Jin, C., Lee, J. D., Jordan, M. I., Singh, A., and Póczos, B. (2017). Gradient descent can take exponential time to escape saddle points. In *Advances in neural information processing systems*, pages 1067–1077.
- Duchi, J. C., Jordan, M. I., Wainwright, M. J., and Wibisono, A. (2015). Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806.
- Fang, C., Lin, Z., and Zhang, T. (2019). Sharp analysis for nonconvex sgd escaping from saddle points. In *Conference on Learning Theory*, pages 1192–1234. PMLR.
- Flaxman, A. D., Kalai, A. T., and McMahan, H. B. (2005). Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 385–394.
- Ge, R., Huang, F., Jin, C., and Yuan, Y. (2015). Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pages 797–842.
- Ge, R., Li, Z., Wang, W., and Wang, X. (2019). Stabilized svrg: Simple variance reduction for nonconvex optimization. In *Conference on learning theory*, pages 1394–1448. PMLR.
- Han, A. and Gao, J. (2020). Escape saddle points faster on manifolds via perturbed riemannian stochastic recursive gradient. *arXiv preprint arXiv:2010.12191*.
- Huang, M. (2021). Escaping saddle points for nonsmooth weakly convex functions via perturbed proximal algorithms. *arXiv preprint arXiv:2102.02837*.

- Huang, M., Ji, K., Ma, S., and Lai, L. (2022). Efficiently escaping saddle points in bilevel optimization. *arXiv preprint arXiv:2202.03684*.
- Jin, C., Ge, R., Netrapalli, P., Kakade, S. M., and Jordan, M. I. (2017). How to escape saddle points efficiently. In *International Conference on Machine Learning*, pages 1724–1732. PMLR.
- Jin, C., Netrapalli, P., Ge, R., Kakade, S. M., and Jordan, M. I. (2019a). On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. *arXiv preprint arXiv:1902.04811*.
- Jin, C., Netrapalli, P., Ge, R., Kakade, S. M., and Jordan, M. I. (2019b). A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*.
- Jin, C., Netrapalli, P., and Jordan, M. I. (2018). Accelerated gradient descent escapes saddle points faster than gradient descent. In *Conference On Learning Theory*, pages 1042–1085. PMLR.
- Larson, J., Menickelly, M., and Wild, S. M. (2019). Derivative-free optimization methods. *Acta Numerica*, 28:287–404.
- Lee, J. D., Panageas, I., Piliouras, G., Simchowitz, M., Jordan, M. I., and Recht, B. (2019). First-order methods almost always avoid strict saddle points. *Mathematical programming*, 176(1):311–337.
- Li, J., Balasubramanian, K., and Ma, S. (2022). Stochastic zeroth-order riemannian derivative estimation and optimization. *Mathematics of Operations Research*.
- Li, Y., Tang, Y., Zhang, R., and Li, N. (2021). Distributed reinforcement learning for decentralized linear quadratic control: A derivative-free policy optimization approach. *IEEE Transactions on Automatic Control*.
- Li, Z. (2019). Ssrgd: Simple stochastic recursive gradient descent for escaping saddle points. *Advances in Neural Information Processing Systems*, 32.
- Liang, G., Tong, Q., Zhu, C., and Bi, J. (2021). Escaping saddle points with stochastically controlled stochastic gradient methods. *arXiv preprint arXiv:2103.04413*.
- Lucchi, A., Orvieto, A., and Solomou, A. (2021). On the second-order convergence properties of random search methods. *Advances in Neural Information Processing Systems*, 34:25633–25645.
- Malik, D., Pananjady, A., Bhatia, K., Khamaru, K., Bartlett, P., and Wainwright, M. (2019). Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. In *The 22nd international conference on artificial intelligence and statistics*, pages 2916–2925. PMLR.
- Mokhtari, A., Ozdaglar, A., and Jadbabaie, A. (2018). Escaping saddle points in constrained optimization. *Advances in Neural Information Processing Systems*, 31.
- Nesterov, Y. and Spokoiny, V. (2017). Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566.
- Panageas, I., Piliouras, G., and Wang, X. (2019). First-order methods almost always avoid saddle points: The case of vanishing step-sizes. *Advances in Neural Information Processing Systems*, 32.
- Reddi, S., Zaheer, M., Sra, S., Póczos, B., Bach, F., Salakhutdinov, R., and Smola, A. (2018). A generic approach for escaping saddle points. In *International conference on artificial intelligence and statistics*, pages 1233–1242. PMLR.
- Roy, A., Balasubramanian, K., Ghadimi, S., and Mohapatra, P. (2020). Escaping saddle-point faster under interpolation-like conditions. *Advances in Neural Information Processing Systems*, 33:12414–12425.
- Shamir, O. (2017). An optimal algorithm for bandit and zero-order convex optimization with two-point feedback. *The Journal of Machine Learning Research*, 18(1):1703–1713.
- Staib, M., Reddi, S., Kale, S., Kumar, S., and Sra, S. (2019). Escaping saddle points with adaptive gradient methods. In *International Conference on Machine Learning*, pages 5956–5965. PMLR.

- Sun, J., Qu, Q., and Wright, J. (2015). When are nonconvex problems not scary? *arXiv preprint arXiv:1510.06096*.
- Sun, T., Li, D., Quan, Z., Jiang, H., Li, S., and Dou, Y. (2019a). Heavy-ball algorithms always escape saddle points. *arXiv preprint arXiv:1907.09697*.
- Sun, Y., Flammarion, N., and Fazel, M. (2019b). Escaping from saddle points on riemannian manifolds. *Advances in Neural Information Processing Systems*, 32.
- Tang, H., Lian, X., Qiu, S., Yuan, L., Zhang, C., Zhang, T., and Liu, J. (2019). Deepsqueeze: Parallel stochastic gradient descent with double-pass error-compensated compression. *arXiv preprint arXiv:1907.07346*.
- Tang, Y., Ren, Z., and Li, N. (2020a). Zeroth-order feedback optimization for cooperative multi-agent systems. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 3649–3656. IEEE.
- Tang, Y., Zhang, J., and Li, N. (2020b). Distributed zero-order algorithms for nonconvex multiagent optimization. *IEEE Transactions on Control of Network Systems*, 8(1):269–281.
- Tropp, J. A. et al. (2015). An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230.
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.
- Vladimirova, M., Girard, S., Nguyen, H., and Arbel, J. (2020). Sub-weibull distributions: Generalizing sub-gaussian and sub-exponential properties to heavier tailed distributions. *Stat*, 9(1):e318.
- Vlaski, S. and Sayed, A. H. (2021a). Distributed learning in non-convex environments—part ii: Polynomial escape from saddle-points. *IEEE Transactions on Signal Processing*, 69:1257–1270.
- Vlaski, S. and Sayed, A. H. (2021b). Second-order guarantees of stochastic gradient descent in non-convex optimization. *IEEE Transactions on Automatic Control*.
- Vlatakis-Gkaragkounis, E.-V., Flokas, L., and Piliouras, G. (2019). Efficiently avoiding saddle points with zero order methods: No gradients required. *Advances in Neural Information Processing Systems*, 32.
- Wang, J.-K., Lin, C.-H., and Abernethy, J. (2021). Escaping saddle points faster with stochastic momentum. *arXiv preprint arXiv:2106.02985*.
- Xu, Y., Jin, R., and Yang, T. (2018). First-order stochastic algorithms for escaping from saddle points in almost linear time. *Advances in neural information processing systems*, 31.
- Zhang, C. and Li, T. (2021). Escape saddle points by a simple gradient-descent based algorithm. *Advances in Neural Information Processing Systems*, 34:8545–8556.
- Zhang, H., Xiong, H., and Gu, B. (2022). Zeroth-order negative curvature finding: Escaping saddle points without gradients. *arXiv preprint arXiv:2210.01496*.
- Ziyin, L., Li, B., Simon, J. B., and Ueda, M. (2021). Sgd with a constant large learning rate can converge to local maxima. *arXiv preprint arXiv:2107.11774*.

A RELATED WORK

Two-point methods in zeroth-order optimization. Two-point (or in general $2m$ -point, where $1 \leq m < d$ with d being the problem dimension) estimators, which approximate the gradient using two (or $2m$) function evaluations per iteration, have been widely studied by researchers in the zeroth-order optimization literature, in convex (Nesterov and Spokoiny, 2017; Duchi et al., 2015; Shamir, 2017), nonconvex (Nesterov and Spokoiny, 2017), online (Shamir, 2017), as well as distributed settings (Tang et al., 2019). A key reason for doing so is that for applications of zeroth-order optimization arising in robotics (Li et al., 2022), wind farms (Tang et al., 2020a), power systems (Chen et al., 2020), online (time-varying) optimization (Shamir, 2017), learning-based control (Malik et al., 2019; Li et al., 2021), and improving adversarial robustness to black-box attacks in deep neural networks (Chen et al., 2017), it may be costly or impractical to wait for $\Omega(d)$ (where d denotes the problem dimension) function evaluations per iteration to make a step. This is especially true for high-dimensional and/or time-varying problems. Indeed, for high-dimensional problems, two-point estimators can make swift progress even in the initial stage compared to $2d$ -point estimator, and can reach a higher-quality solution if computation is limited (Tang et al., 2020b; Chen et al., 2017). For instance, consider the work in (Chen et al., 2017), which studies the use of zeroth-order estimators to perform black-box attacks on deep neural networks, in order to identify (and then defend against) adversarial images that may lead to misclassification. In the paper, the authors employed two-point zeroth-order estimators, due to the high computational cost of using $2d$ function evaluations per iteration for hundreds of iterations (here d is the dimension of an image, which in this case is over 20000). The authors showed empirically that their two-point estimators worked well; however there over no accompanying theoretical results.

For online or time-varying environments, two-points estimators also often preferable. Since zeroth-order methods are often used in physical systems whose environment drifts or changes over time, this leads naturally to a *time-varying* or *online* optimization. For these problems, $2d$ -point estimators will not produce a good estimation because the underlying function can drift to a very different problem while waiting for the $2d$ function evaluations. Indeed, the fewer function evaluations an optimization procedure needs, the faster it can catch up with the time-varying environment. In fact, for online optimization, it has been shown that two points estimator is optimal for convex Lipschitz functions (Shamir, 2017). Thus, two-point estimators are a natural fit for time-varying online optimization problems.

Saddle point escape with access to deterministic gradient. While standard gradient descent can escape saddle points asymptotically (Lee et al., 2019; Panageas et al., 2019), it is known that standard gradient descent may take exponential time to escape saddle points (Du et al., 2017). Hence, when access to deterministic gradient is available, research has centered on escaping saddle points with adding perturbation (Jin et al., 2017), momentum/acceleration based methods (Jin et al., 2018; Sun et al., 2019a; Staib et al., 2019), or gradient-based robust Hessian power/curvature exploitation methods (Zhang and Li, 2021; Adolphs et al., 2019). In addition, there has also been work on escaping saddle points devoted to specific optimization settings, such as constrained optimization (Mokhtari et al., 2018; Avdiukhin et al., 2019), optimization of weakly convex functions (Huang, 2021), bilevel optimization (Huang et al., 2022), as well as on general manifolds (Sun et al., 2019b; Criscitiello and Boumal, 2019; Han and Gao, 2020).

Saddle point escape in stochastic gradient descent (SGD). In practice, only stochastic gradient estimators are available in many problems. While SGD may converge to local maxima in worst-case scenarios (Ziyin et al., 2021), under assumptions such as bounded variance or subGaussian noise, there have been many works that have studied the problem of saddle point escape in SGD (Ge et al., 2015; Daneshmand et al., 2018; Xu et al., 2018; Jin et al., 2019a; Vlaski and Sayed, 2021b). The best existing rate (without considering momentum/variance reduction techniques) appears to belong to that of Fang et al. (2019), which converges to ϵ -second order stationary points using $\tilde{O}(1/\epsilon^{3.5})$ stochastic gradients. While zeroth-order gradient estimators may also be viewed as stochastic gradients, they typically do not satisfy the bounded/subGaussian noise assumptions that are assumed in these works, making a direct comparison inappropriate. Escaping saddle point via momentum methods in SGD has also been studied (Wang et al., 2021; Antonakopoulos et al., 2022); while we do not consider incorporating momentum in our works, this may be interesting future work. A number of papers has also considered the specialized setting of escaping saddle points in nonconvex finite-sum optimization (Reddi et al., 2018; Liang et al., 2021), with many considering the case where variance-reduction is

used (Ge et al., 2019; Li, 2019). While the finite-sum problem is quite different from our problem, the variance reduction approach considered in these works may be a relevant future direction. The saddle point escape problem has also been studied in other specific settings such as compressed optimization (Avdiukhin and Yaroslavtsev, 2021), distributed optimization (Vlaski and Sayed, 2021a), or in the overparameterization case (Roy et al., 2020).

Saddle point escape with zeroth-order information. The problem of escaping saddle points in zeroth-order optimization has been studied less often, and we have already listed all known works comparable to our work in the introduction (Bai et al., 2020; Vlatakis-Gkaragkounis et al., 2019; Balasubramanian and Ghadimi, 2022); a more detailed comparison of these works with our results has been provided in the discussion following the statement of our main result Theorem 1. We would like to mention that Roy et al. (2020) also includes a convergence result of $\tilde{O}\left(\frac{d^{1.5}}{\epsilon^{4.5}}\right)$ for the case with noisy function evaluations, which is incomparable to our existing work which focuses on the case with exact function evaluation. In addition, Roy et al. (2020) also makes a subGaussian assumption on the estimator noise, which zeroth-order estimators in our paper do not satisfy. Nonetheless, considering the extension to noisy function evaluations will make for important future work.

Zeroth-order optimization. Our work rests on a line of research in zeroth-order optimization which focuses on constructing gradient estimators using zeroth-order function values (Flaxman et al., 2005; Duchi et al., 2015; Nesterov and Spokoiny, 2017; Shamir, 2017; Larson et al., 2019). As we have discussed, for smooth nonconvex functions, it is known that two-point zeroth-order estimators suffice to find first-order ϵ -stationary points using $\tilde{O}(d/\epsilon^2)$ function evaluations (Nesterov and Spokoiny, 2017). Our work studies the more complicated problem of reaching ϵ -second order stationary points, attaining a rate of $\tilde{O}(d/\epsilon^{2.5})$.

B PROOF ROADMAP

We begin by introducing several key concentration inequalities in Appendix C which we will frequently use in our proofs. We then describe in detail (and prove) the sequence of results that lead up to Proposition 4 in Appendix D, showing that there cannot be too many iterations with large gradients. Next, we describe the saddle point argument in detail, and prove Proposition 5 in Appendix E.3. Finally, we combine these results and prove our main result Theorem 2 (whose informal version is Theorem 1) in Appendix F

Throughout our proofs, absolute constants, as denoted by e.g. (c, c', C) , may change from line to line. However, within the same proof, for clarity, we try to index different constants differently. We assume $d \geq 2$ and $m \leq d$.

Notations. We shall denote the conditional expectation and conditional probability by $\mathbb{E}_{\mathcal{F}}[\cdot] = \mathbb{E}[\cdot | \mathcal{F}]$ and $\mathbb{P}_{\mathcal{F}}(\cdot) = \mathbb{P}(\cdot | \mathcal{F})$ where \mathcal{F} is a sigma-algebra.

C CONCENTRATION INEQUALITIES

This section serves to introduce several probabilistic results which will be useful for our main proofs in subsequent sections. We first introduce subGaussian, subExponential and norm-subGaussian random vectors in Appendix C.1. Next, in Appendix C.2, we provide concentration bounds for norm-subGaussian and subExponential random vectors. We then prove a novel concentration inequality involving products of subGaussian random vectors in Appendix C.3. We conclude by stating some concentration bounds for Appendix C.4 random variables.

C.1 SUBGAUSSIAN, SUBEXPONENTIAL AND NORM-SUBGAUSSIAN RANDOM VECTORS

We first define subGaussian and subExponential random vectors. A detailed reference for these concepts can be found in Vershynin (2018).

Definition 5 (subGaussian and subExponential random vectors). A random vector $\mathbf{x} \in \mathbb{R}^d$ is σ -subGaussian (SG(σ)), if there exists $\sigma > 0$ such that for any unit vector $g \in \mathbb{S}^{d-1}$,

$$\mathbb{E}[\exp(\lambda \langle g, \mathbf{x} - \mathbb{E}[\mathbf{x}] \rangle)] \leq \exp(\lambda^2 \sigma^2 / 2) \quad \forall \lambda \in \mathbb{R}.$$

Meanwhile, a random vector $\mathbf{x} \in \mathbb{R}^d$ is σ -subExponential (SE(σ)), if there exists $\sigma > 0$ such that for any unit vector $g \in \mathbb{S}^{d-1}$,

$$\mathbb{E}[\exp(\lambda\langle g, \mathbf{x} - \mathbb{E}[\mathbf{x}] \rangle)] \leq \exp(\lambda^2\sigma^2/2) \quad \forall |\lambda| \leq \frac{1}{\sigma}$$

An alternative concentration property for random vectors revolving around its norm, known as norm-subGaussianity (Jin et al., 2019b), is also relevant.

Definition 6 (norm-subGaussian random vectors). A random vector $\mathbf{x} \in \mathbb{R}^d$ is σ -norm-subGaussian (nSG(σ)), there exists $\sigma > 0$ such that

$$\mathbb{P}(\|\mathbf{x} - \mathbb{E}\mathbf{x}\| \geq s) \leq 2e^{-\frac{s^2}{2\sigma^2}} \quad \forall s \geq 0.$$

We recall the following result which provides several examples of nSG random vectors. In particular, it tells us a random vector $\mathbf{x} \in \mathbb{R}^d$ that is (σ/\sqrt{d}) -subGaussian is also σ -subGaussian.

Lemma 4 (Lemma 1 in Jin et al. (2019b)). *There exists absolute constant c such that the following random vectors are all nSG($c\sigma$).*

1. A bounded random vector $\mathbf{x} \in \mathbb{R}^d$ so that $\|\mathbf{x}\| \leq \sigma$.
2. A random vector $\mathbf{x} \in \mathbb{R}^d$, where $\mathbf{x} = \xi\mathbf{e}_1$ and the random variable $\xi \in \mathbb{R}$ is σ -subGaussian.
3. A random vector $\mathbf{x} \in \mathbb{R}^d$ that is (σ/\sqrt{d}) -subGaussian

In addition, if $\mathbf{x} \in \mathbb{R}^d$ is zero-mean nSG(σ), its component along a single direction is also subGaussian.

Lemma 5. *Suppose $\mathbf{x} \in \mathbb{R}^d$ is zero-mean nSG(σ). Then, for any fixed vector $\mathbf{v} \in \mathbb{R}^d$, $\langle \mathbf{v}, \mathbf{x} \rangle$ is zero-mean $\|\mathbf{v}\|\sigma$ -subGaussian.*

Proof. Without loss of generality, we assume that $\mathbf{v} \in \mathbb{S}^{d-1}$ is a unit vector. That $\langle \mathbf{v}, \mathbf{x} \rangle$ is zero-mean follows directly from \mathbf{x} being zero-mean and \mathbf{v} being fixed. Meanwhile, since $|\langle \mathbf{v}, \mathbf{x} \rangle| \leq \|\mathbf{v}\|\|\mathbf{x}\| = \|\mathbf{x}\|$, for any $s \geq 0$, it follows that

$$\mathbb{P}(|\langle \mathbf{v}, \mathbf{x} \rangle| \geq s) \leq \mathbb{P}(\|\mathbf{x}\| \geq s) \leq 2e^{-\frac{s^2}{2\sigma^2}},$$

where the last inequality follows from the fact that \mathbf{x} is zero-mean and also nSG(σ). Hence $\langle \mathbf{v}, \mathbf{x} \rangle$ is zero-mean SG(σ), as desired. \square

C.2 CONCENTRATION BOUNDS FOR NORM-SUBGAUSSIAN AND SUBEXPONENTIAL RANDOM VECTORS

We begin by giving some concentration bounds for norm-subGaussian random vectors. To do so, we introduce the following condition.

Condition 1. *Consider random vectors $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$, and corresponding filtrations \mathcal{F}_i generated by $(\mathbf{x}_1, \dots, \mathbf{x}_i)$. We assume $\mathbf{x}_i | \mathcal{F}_{i-1}$ is zero-mean, nSG(σ_i), with $\sigma_i \in \mathcal{F}_{i-1}$, i.e,*

$$\mathbb{E}[\mathbf{x}_i | \mathcal{F}_{i-1}] = 0,$$

and

$$\mathbb{P}(\|\mathbf{x}_i\| \geq s | \mathcal{F}_{i-1}) \leq 2e^{-\frac{s^2}{2\sigma_i^2}} \quad \forall s \geq 0,$$

where σ_i is a measurable function of $(\mathbf{x}_1, \dots, \mathbf{x}_{i-1})$ for each i .

For norm subGaussian random vectors satisfying Condition 1, we first have the following bound.

Lemma 6. *Suppose $(\mathbf{x}_1, \dots, \mathbf{x}_n) \in \mathbb{R}^d$ satisfy Condition 1, i.e. each $\mathbf{x}_i | \mathcal{F}_{i-1}$ is mean-zero, nSG(σ_i) with $\sigma_i \in \mathcal{F}_{i-1}$. Let $\{\mathbf{u}_i\}$ denote a sequence of random vectors such that $\mathbf{u}_i \in \mathcal{F}_{i-1}$ for every $i \in [n]$. Then, there exists an absolute constant c , such that for any $\delta \in (0, 1)$ and $\lambda > 0$, with probability at least $1 - \delta$,*

$$\sum_{i=1}^n \langle \mathbf{u}_i, \mathbf{x}_i \rangle \leq c\lambda \sum_{i=1}^n \|\mathbf{u}_i\|^2 \sigma_i^2 + \frac{1}{\lambda} \log(1/\delta).$$

Proof. We note that if \mathbf{x}_i is mean-zero and $\text{nSG}(\sigma_i)$, then by Lemma 5, $\langle \mathbf{u}_i, \mathbf{x}_i \rangle \mid \mathcal{F}_{i-1}$ is zero-mean and $\|\mathbf{u}_i\|_{\sigma_i}$ -subGaussian. The rest of the proof follows from the proof of Lemma 39 in Jin et al. (2019a) (key idea is exponentiate and then apply Markov's inequality). For completeness, we restate the proof here. Observe that for any i , since $\langle \mathbf{u}_i, \mathbf{x}_i \rangle$ is $\|\mathbf{u}_i\|_{\sigma_i}$ -subGaussian, for any $\lambda > 0$, we have that

$$\mathbb{E} [\exp(\lambda \langle \mathbf{u}_i, \mathbf{x}_i \rangle) \mid \mathcal{F}_{i-1}] \leq \exp(\lambda^2 \|\mathbf{u}_i\|^2 \sigma_i^2 / 2)$$

For any $\lambda > 0$ and $s \geq 0$, observe that

$$\begin{aligned} & \mathbb{P} \left(\sum_{i=1}^n \lambda \langle \mathbf{u}_i, \mathbf{x}_i \rangle - \lambda^2 \|\mathbf{u}_i\|^2 \sigma_i^2 / 2 \geq s \right) \\ &= \mathbb{P} \left(\exp \left(\lambda \sum_{i=1}^n \langle \mathbf{u}_i, \mathbf{x}_i \rangle - \lambda^2 \|\mathbf{u}_i\|^2 \sigma_i^2 / 2 \right) \geq \exp(\lambda s) \right) \\ &\leq \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^n \langle \mathbf{u}_i, \mathbf{x}_i \rangle - \lambda^2 \|\mathbf{u}_i\|^2 \sigma_i^2 / 2 \right) \right] \exp(-\lambda s) \\ &= \mathbb{E} \left[\mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^n \langle \mathbf{u}_i, \mathbf{x}_i \rangle - \lambda^2 \|\mathbf{u}_i\|^2 \sigma_i^2 / 2 \right) \mid \mathcal{F}_{n-1} \right] \right] \exp(-\lambda s) \\ &= \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{n-1} \langle \mathbf{u}_i, \mathbf{x}_i \rangle - \lambda^2 \|\mathbf{u}_i\|^2 \sigma_i^2 / 2 \right) \mathbb{E} \left[\exp \left(\lambda \langle \mathbf{u}_n, \mathbf{x}_n \rangle - \lambda^2 \|\mathbf{u}_n\|^2 \sigma_n^2 / 2 \right) \mid \mathcal{F}_{n-1} \right] \right] \exp(-\lambda s) \\ &\stackrel{(i)}{\leq} \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{n-1} \langle \mathbf{u}_i, \mathbf{x}_i \rangle - \lambda^2 \|\mathbf{u}_i\|^2 \sigma_i^2 / 2 \right) \right] \exp(-\lambda s) \leq \dots \leq \exp(-\lambda s) \end{aligned}$$

Above, (i) follows from the fact that $\langle \mathbf{u}_i, \mathbf{x}_i \rangle \mid \mathcal{F}_{i-1}$ is zero-mean and $\|\mathbf{u}_i\|_{\sigma_i}$ -subGaussian for each $i \in [n]$. The final result then follows by picking $c = \frac{1}{2}$ and $s = \log(1/\delta)$. \square

Assuming Condition 1, the following concentration result also holds for a sequence of nSG random vectors.

Lemma 7 (Lemma 6, Corollary 7 and Corollary 8 in Jin et al. (2019b) combined). *Suppose $(\mathbf{x}_1, \dots, \mathbf{x}_n) \in \mathbb{R}^d$ satisfy Condition 1. Then, there exists an absolute constant c such that for any fixed $\delta \in (0, 1)$, $\theta > 0$, with probability at least $1 - \delta$,*

$$\left\| \sum_{i=1}^n \mathbf{x}_i \right\| \leq c\theta \sum_{i=1}^n \sigma_i^2 + \frac{1}{\theta} \log(2d/\delta).$$

Moreover, there are two corollaries.

1. ((Jin et al., 2019b, Corollary 7)) *When $\{\sigma_i\}$ is deterministic, there exists an absolute constant c such that for any fixed $\delta \in (0, 1)$, with probability at least $1 - \delta$.*

$$\left\| \sum_{i=1}^n \mathbf{x}_i \right\| \leq c \sqrt{\log(2d/\delta) \sum_{i=1}^n \sigma_i^2}$$

2. ((Jin et al., 2019b, Corollary 8)) *Suppose that the $\{\sigma_i\}$ sequence is random. Then, there exists an absolute constant c such that for any fixed $\delta \in (0, 1)$ and $B > b > 0$, with probability at least $1 - \delta$:*

$$\text{either } \sum_{i=1}^n \sigma_i^2 \geq B \quad \text{or} \quad \left\| \sum_{i=1}^n \mathbf{x}_i \right\| \leq c \sqrt{\max \left\{ \sum_{i=1}^n \sigma_i^2, b \right\} \cdot (\log(2d/\delta) + \log(\log(B/b)))}$$

We state here a Bernstein-type concentration inequality for sub-exponential random variables, which we also need.

Lemma 8 (Bernstein concentration inequality). *Consider a sequence of independently distributed σ -subexponential variables $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}$, with mean $\mathbb{E}[\mathbf{x}_i] \leq c'\sigma$ for some $c' > 0$ and each $i \in [n]$. Then, there exists an absolute constant $C > 0$, such that for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$\sum_{i=1}^n \mathbf{x}_i \leq C\sigma(n + \log(1/\delta)). \quad (6)$$

Proof. The result of Eq. (6) follows by applying Bernstein’s inequality to $\sum_{i=1}^n \mathbf{x}_i - \mathbb{E}[\mathbf{x}_i]$ (so each summand is mean-zero). Per Bernstein’s inequality, (cf. Theorem 2.8.1 in Vershynin (2018)), there exists an absolute constant $c > 0$ such that for any $s \geq 0$,

$$\mathbb{P}\left(\sum_{i=1}^n (\mathbf{x}_i - \mathbb{E}[\mathbf{x}_i]) \geq s\right) \leq \exp\left(-c \min\left\{\frac{s^2}{n\sigma^2}, \frac{s}{\sigma}\right\}\right).$$

Pick $s = \sigma\left(n + \frac{\log(1/\delta)}{c}\right)$. Then,

$$\min\left\{\frac{s^2}{n\sigma^2}, \frac{s}{\sigma}\right\} = \min\left\{n + 2\frac{\log(1/\delta)}{c} + \frac{(\log(1/\delta))^2}{c^2n}, n + \frac{\log(1/\delta)}{c}\right\} = n + \frac{\log(1/\delta)}{c}.$$

Continuing, we have that

$$\mathbb{P}\left(\sum_{i=1}^n (\mathbf{x}_i - \mathbb{E}[\mathbf{x}_i]) \geq s\right) \leq \exp\left(-c \min\left\{\frac{s^2}{n\sigma^2}, \frac{s}{\sigma}\right\}\right) \leq \exp\left(-c\left(n + \frac{\log(1/\delta)}{c}\right)\right) \leq \delta.$$

Thus, it follows that with probability at least $1 - \delta$,

$$\sum_{i=1}^n (\mathbf{x}_i - \mathbb{E}[\mathbf{x}_i]) \leq \sigma\left(n + \frac{\log(1/\delta)}{c}\right) \implies \sum_{i=1}^n \mathbf{x}_i \leq \sigma\left(n + \frac{\log(1/\delta)}{c}\right) + nc'\sigma,$$

where implication holds since by assumption, $\mathbb{E}[\mathbf{x}_i] \leq c'\sigma$ for some $c' > 0$. Then, by setting $C = \max\{1 + c', 1/c\}$, the desired result follows. \square

C.3 A NOVEL CONCENTRATION INEQUALITY FOR THE ZERO-ORDER SETTING

In the zeroth-order setting, we will frequently have to bound the norms of terms of the form

$$W_\tau = \sum_{t=0}^{\tau-1} M_t (Z_t Z_t^\top - I) v_t, \quad (7)$$

where M_t is a known and fixed quantity, while Z_t is random, and v_t depends on x_0 and the history of previous $\{Z_j\}_{j=0}^{t-1}$ ’s, and is hence \mathcal{F}_{t-1} -measurable. For our purposes, it suffices to consider $Z_t \sim N(0, I)$.

To see why such a bound will be useful, as mentioned in the main text and as we will see again later in the full proofs, in the analysis of escaping saddle points, we need to bound a term of the form

$$W_{g_0}(\tau) = \eta \sum_{t=0}^{\tau-1} (I - \eta H)^{\tau-1-t} (Z_t Z_t^\top - I) (\nabla f(x_t) - \nabla f(x'_t)),$$

where $H = \nabla^2 f(x_0)$ (assuming that x_0 is an ϵ -saddle point), and x_t and x'_t are two coupled sequences. Comparing with Eq. (7), we see that for the equation above, we can define $M_t = \eta(I - \eta H)^{\tau-1-t}$ (a fixed and known quantity) and $v_t = \nabla f(x_t) - \nabla f(x'_t)$ (clearly, $\nabla f(x_t) - \nabla f(x'_t)$ is \mathcal{F}_{t-1} -measurable). This motivates why we wish to bound terms of the form Eq. (7).

Observe that each $(Z_t Z_t^\top - I)v_t \mid \mathcal{F}_{t-1}$ term is subExponential rather than subGaussian. While it is possible to define norm-subExponential vectors in analogous way to norm-subGaussian vectors, the corresponding moment generating function (MGF) for subExponential random variables is not defined on the entirety of \mathbb{R} . When bounding a sum in the form of $\sum_{t=0}^{\tau-1} (Z_t Z_t^\top - I)v_t$, this creates a subtle but challenging technical issue.

Following the intuition outlined in the main text, we bypass this difficulty by proving the following result. For notational simplicity, we introduce the function

$$\text{lr}(x) := \log(x \log(x)). \quad (8)$$

We now recall Proposition 2 which we first introduced in the main text.

Proposition 2. *Let \mathcal{F}_t , $t \geq -1$ be a filtration. Let $(Z_t)_{t \geq 0}$ be a sequence of random vectors following the distribution $N(0, I)$ such that $Z_t \in \mathcal{F}_t$ and is independent of \mathcal{F}_{t-1} , and let $(v_t)_{t \geq 0}$ be a sequence of random vectors such that $v_t \in \mathcal{F}_{t-1}$. For each $\tau \geq 0$, let*

$$W_\tau = \sum_{t=0}^{\tau-1} M_t (Z_t Z_t^\top - I) v_t,$$

where each M_t is a deterministic matrix of appropriate dimension. Then, there exist some absolute constants $c', C > 0$ such that for any $\tau \in \mathbb{Z}^+$ and $\delta \in (0, 1/e]$, the following statements hold:

1. For any $\theta > 0$, with probability at least $1 - \delta$, we have

$$\|W_\tau\| \leq c' \theta \sum_{t=0}^{\tau-1} \|M_t\|_2^2 d(\text{lr}(C\tau/\delta))^2 \|v_t\|^2 + \frac{1}{\theta} \log(Cd\tau/\delta).$$

2. For any $B > b > 0$, with probability at least $1 - \delta$,

$$\text{either } \sum_{t=0}^{\tau-1} \|M_t\|_2^2 d(\text{lr}(C\tau/\delta))^2 \|v_t\|^2 \geq B$$

$$\text{or } \|W_\tau\| \leq c' \sqrt{\max \left\{ \sum_{t=0}^{\tau-1} \|M_t\|_2^2 d(\text{lr}(C\tau/\delta))^2 \|v_t\|^2, b \right\}} (\log(C\tau d/\delta) + \log(\log(B/b) + 1))$$

Moreover, as is clear from the bounds above, we may pick $C \geq 1$ such that $\log(C/\delta) \geq 1, \forall \delta \in (0, 1/e]$.

Proof. We will focus on proving the first point, since the second follows as a natural corollary of our proof of the first part and the proof of (Jin et al., 2019b, Corollary 8). For simplicity, we shall assume $v_t \neq 0$ in the intermediate steps; extension to the general case is straightforward.

First of all, for $0 \leq \alpha < 1$, let

$$g(\alpha; \delta) = \sqrt{\frac{2}{\pi}} \int_{\alpha}^{\sqrt{2\text{lr}(1/\delta)}} (x^2 - 1) e^{-x^2/2} dx = \sqrt{\frac{2}{\pi}} \left(\alpha e^{-\alpha^2/2} - \frac{\delta \sqrt{2\text{lr}(1/\delta)}}{\log(1/\delta)} \right).$$

It's not hard to see that for a fixed $\delta \in (0, 1/e]$, $g(\alpha; \delta)$ is continuous and strictly increasing over $\alpha \in [0, 1)$. Then, since $\frac{\log x}{x} + 1 \leq x$ for $x \geq 1$, by plugging in $x = \log(1/\delta)$, we get

$$\frac{\text{lr}(1/\delta)}{(\log(1/\delta))^2} = \frac{\log \log(1/\delta) + \log(1/\delta)}{(\log(1/\delta))^2} = \frac{1}{\log(1/\delta)} \left(\frac{\log \log(1/\delta)}{\log(1/\delta)} + 1 \right) \leq 1,$$

which leads to

$$g(2\delta; \delta) = \sqrt{\frac{2}{\pi}} \left(2\delta e^{-2\delta^2} - \frac{\delta \sqrt{2\text{lr}(1/\delta)}}{\log(1/\delta)} \right) \geq \sqrt{\frac{2}{\pi}} \left(2e^{-2/e^2} \delta - \sqrt{2\delta} \right) > 0$$

for $\delta \in (0, 1/e]$. Furthermore, we obviously have $g(0; \delta) < 0$. Therefore $g(\alpha; \delta) = 0$ has a unique solution in $(0, 2\delta)$, which we denote by $\alpha(\delta)$.³ These results imply that, for a random variable Z following the standard normal distribution, we have

$$\mathbb{E} \left[(Z^2 - 1) \mathbb{1}_{\alpha(\delta) \leq |Z| \leq \sqrt{2\text{lr}(1/\delta)}} \right] = \sqrt{\frac{2}{\pi}} \int_{\alpha(\delta)}^{\sqrt{2\text{lr}(1/\delta)}} (x^2 - 1) e^{-x^2/2} dx = g(h(\delta); \delta) = 0$$

³By letting $W_0(x)$ denote the the principal branch of the Lambert W function, it can be shown that

$$\alpha(\delta) = \sqrt{-W_0 \left(-\frac{2\delta^2 \text{lr}(1/\delta)}{(\log(1/\delta))^2} \right)}.$$

and

$$\begin{aligned} \mathbb{P}(\alpha(\delta) \leq |Z| \leq \sqrt{2 \operatorname{lr}(1/\delta)}) &\geq 1 - 2 \left(\frac{1}{\sqrt{2\pi}} \int_{\sqrt{2 \operatorname{lr}(1/\delta)}}^{\infty} e^{-x^2/2} dx + \frac{1}{\sqrt{2\pi}} \int_0^{\alpha(\delta)} e^{-x^2/2} dx \right) \\ &\geq 1 - 2 \left(\frac{1}{2} \exp\left(-\frac{2 \operatorname{lr}(1/\delta)}{2}\right) + \frac{\alpha(\delta)}{\sqrt{2\pi}} \right) = 1 - 2 \left(\frac{\delta}{2 \log(1/\delta)} + \frac{\alpha(\delta)}{\sqrt{2\pi}} \right) \\ &\geq 1 - 2 \left(\frac{\delta}{2} + \frac{2}{\sqrt{2\pi}} \delta \right) \geq 1 - C\delta \end{aligned}$$

for any $\delta \in (0, 1/e]$, where we define the absolute constant $C := 2(1/2 + 2/\sqrt{2\pi})$.

Now we let A_t denote the event

$$A_t = \left\{ \alpha(\delta) \leq \frac{|Z_t^\top v_t|}{\|v_t\|} \leq \sqrt{2 \operatorname{lr}(1/\delta)} \right\}.$$

Since $Z_t^\top v_t / \|v_t\|$ conditioned on \mathcal{F}_{t-1} follows the standard normal distribution, we have

$$\mathbb{P}_{\mathcal{F}_{t-1}}(A_t) \geq 1 - C\delta, \quad (9)$$

and

$$\mathbb{E}_{\mathcal{F}_{t-1}}[v_t^\top (Z_t Z_t^\top - I) v_t \mathbb{1}_{A_t}] = 0.$$

Moreover, for any random vector $u \in \mathcal{F}_{t-1}$ that is orthogonal to v_t , we have

$$\mathbb{E}_{\mathcal{F}_{t-1}}[u^\top (Z_t Z_t^\top - I) v_t \mathbb{1}_{A_t}] = \mathbb{E}_{\mathcal{F}_{t-1}}[u^\top Z_t] \cdot \mathbb{E}_{\mathcal{F}_{t-1}}[Z_t^\top v_t \mathbb{1}_{A_t}] = 0,$$

where we used the fact that $Z_t^\top u$ is independent of $Z_t^\top v_t$ conditioned on \mathcal{F}_{t-1} . Therefore

$$\mathbb{E}_{\mathcal{F}_{t-1}}[(Z_t Z_t^\top - I) v_t \mathbb{1}_{A_t}] = 0.$$

Consider defining then the random variable Q_t by

$$Q_t := (Z_t Z_t^\top - I) v_t \cdot \mathbb{1}_{A_t}.$$

We now show that $Q_t \mid \mathcal{F}_{t-1}$ is norm-subGaussian. Let $u \in \mathbb{R}^d$ with $\|u\| = 1$ be arbitrary. We have

$$\begin{aligned} u^\top Q_t &= u^\top (Z_t Z_t^\top - I) v_t \cdot \mathbb{1}_{A_t} \\ &= u^\top \left(\frac{v_t v_t^\top}{\|v_t\|^2} + I - \frac{v_t v_t^\top}{\|v_t\|^2} \right) (Z_t Z_t^\top - I) v_t \cdot \mathbb{1}_{A_t} \\ &= u^\top v_t \left(\frac{|Z_t^\top v_t|^2}{\|v_t\|^2} - 1 \right) \cdot \mathbb{1}_{A_t} + u^\top \left(I - \frac{v_t v_t^\top}{\|v_t\|^2} \right) (Z_t Z_t^\top - I) v_t \cdot \mathbb{1}_{A_t} \\ &= u^\top v_t \left(\frac{|Z_t^\top v_t|^2}{\|v_t\|^2} - 1 \right) \cdot \mathbb{1}_{A_t} + u_\perp^\top Z_t Z_t^\top v_t \cdot \mathbb{1}_{A_t}, \end{aligned}$$

where we denote $u_\perp = \left(I - \frac{v_t v_t^\top}{\|v_t\|^2} \right) u$. Since

$$\left| u^\top v_t \left(\frac{|Z_t^\top v_t|^2}{\|v_t\|^2} - 1 \right) \cdot \mathbb{1}_{A_t} \right| \leq |u^\top v_t| (2 \operatorname{lr}(1/\delta) - 1),$$

we see that $u^\top v_t \left(\frac{|Z_t^\top v_t|^2}{\|v_t\|^2} - 1 \right) \cdot \mathbb{1}_{A_t}$ conditioned on \mathcal{F}_{t-1} is $|u^\top v_t| (2 \operatorname{lr}(1/\delta) - 1)$ -subGaussian.

Furthermore, since $|u_\perp^\top Z_t Z_t^\top v_t \cdot \mathbb{1}_{A_t}| \leq |Z_t^\top u_\perp| \sqrt{2 \operatorname{lr}(1/\delta)} \|v_t\|$, we have

$$\mathbb{P}_{\mathcal{F}_{t-1}}(|u_\perp^\top Z_t Z_t^\top v_t \cdot \mathbb{1}_{A_t}| \geq s) \leq \mathbb{P}_{\mathcal{F}_{t-1}}(|Z_t^\top u_\perp| \sqrt{2 \operatorname{lr}(1/\delta)} \|v_t\| \geq s),$$

and since $Z_t u_\perp / \|u_\perp\| \mid \mathcal{F}_{t-1}$ follows the standard normal distribution, we see that $u_\perp^\top Z_t Z_t^\top v_t \cdot \mathbb{1}_{A_t}$ is a $\sqrt{2 \operatorname{lr}(1/\delta)} \|u_\perp\| \|v_t\|$ -subGaussian variable. Note that $u^\top Q_t$ is just the sum of $u^\top v_t \left(\frac{|Z_t^\top v_t|^2}{\|v_t\|^2} - 1 \right) \cdot \mathbb{1}_{A_t}$ and $u_\perp^\top Z_t Z_t^\top v_t \cdot \mathbb{1}_{A_t}$, we can conclude that $u^\top Q_t$ is subGaussian with parameter

$$(2 \operatorname{lr}(1/\delta) - 1) |u^\top v_t| + \sqrt{2 \operatorname{lr}(1/\delta)} \|u_\perp\| \|v_t\|$$

$$\begin{aligned} &\leq 2\text{lr}(1/\delta)(|u^\top v_t| + \|u_\perp\| \|v_t\|) \leq 2\sqrt{2}\text{lr}(1/\delta)\sqrt{|u^\top v_t|^2 + \|u_\perp\|^2 \|v_t\|^2} \\ &= 2\sqrt{2}\text{lr}(1/\delta)\|v_t\|, \end{aligned}$$

whenever $\delta \in (0, 1/e]$. Consequently, by (Jin et al., 2019b, Lemma 1), we see that $Q_t \mid \mathcal{F}_{t-1}$ is $8\text{lr}(1/\delta)\sqrt{d}\|v_t\|$ -norm-subGaussian.

It follows easily that $M_t Q_t \mid \mathcal{F}_{t-1}$ is mean-zero and $8\text{lr}(1/\delta)\|M_t\|_2\|v_t\|\sqrt{d}$ -norm-subGaussian. Hence, by (Jin et al., 2019a, Lemma 6), we know that there exists an absolute constant $c > 0$ such that for any $\theta > 0$ and $\delta > 0$, we have that with probability at least $1 - \delta$,

$$\left\| \sum_{t=0}^{\tau-1} M_t Q_t \right\| \leq c\theta \sum_{t=0}^{\tau-1} d(\text{lr}(1/\delta))^2 \|M_t\|_2^2 \|v_t\|^2 + \frac{1}{\theta} \log(2d/\delta).$$

Now, consider denoting the event

$$A := \bigcup_{t=0}^{\tau-1} A_t = \left\{ |Z_t^\top v_t| \in \left(\alpha(\delta)\|v_t\|, \sqrt{2\text{lr}(1/\delta)}\|v_t\| \right), \forall t \in \{0, \dots, \tau-1\} \right\}$$

By the union bound and Eq. (9), we note that

$$\mathbb{P}(A) \geq 1 - \tau C \delta.$$

Moreover, note that on the event A , $\sum_{t=0}^{\tau-1} M_t Q_t = \sum_{t=0}^{\tau-1} M_t (Z_t Z_t^\top - I) v_t$. Hence,

$$\begin{aligned} &\mathbb{P} \left(\left\| \sum_{t=0}^{\tau-1} M_t (Z_t Z_t^\top - I) v_t \right\| \leq c\theta \sum_{t=0}^{\tau-1} d(\text{lr}(1/\delta))^2 \|M_t\|_2^2 \|v_t\|^2 + \frac{1}{\theta} \log(2d/\delta) \right) \\ &\geq \mathbb{P} \left(\left\| \sum_{t=0}^{\tau-1} M_t Y_t \right\| \leq c\theta \sum_{t=0}^{\tau-1} d(\text{lr}(1/\delta))^2 \|M_t\|_2^2 \|v_t\|^2 + \frac{1}{\theta} \log(2d/\delta), \text{ and } A \text{ happens} \right) \\ &\geq 1 - \left(\mathbb{P} \left(\left\| \sum_{t=0}^{\tau-1} M_t Y_t \right\| \geq c\theta \sum_{t=0}^{\tau-1} d(\text{lr}(1/\delta))^2 \|M_t\|_2^2 \|v_t\|^2 + \frac{1}{\theta} \log(2d/\delta) \right) + \mathbb{P}(A^c) \right) \\ &\geq 1 - (\delta + \tau C \delta). \end{aligned}$$

Now, by rescaling δ to $\delta/(C\tau + 1)$, we get the desired result. Note this C is different from the C in the statement of the lemma by an absolute multiplicative factor. \square

C.4 SUB-WEIBULL RANDOM VARIABLES

In our work, we occasionally require bounding sums of heavy-tailed distribution, e.g. higher powers of $\|Z\|$ where $Z \sim N(0, I)$. To this end, we consider the following definition of sub-Weibull random variables.

Definition 7. We say that a random variable $X \in \mathbb{R}$ is *sub-Weibull* (K, α) for some $K, \alpha > 0$,

$$\mathbb{P}(|X| \geq s) \leq 2 \exp(-(s/K)^{1/\alpha}) \quad \forall s \geq 0.$$

For instance, the standard normal distribution is sub-Weibull $(1, \frac{1}{2})$. From the way we define the tail parameter α , the larger the α , the heavier the tail of the distribution.

In our work, we need to show that the sum of sub-Weibull random variables is again sub-Weibull, which is ensured by the following result

Lemma 9. *Suppose X and Y are sub-Weibull (K_X, α) and sub-Weibull (K_Y, α) respectively. Then, XY is sub-Weibull $(C(K_X \cdot K_Y), 2\alpha)$ and $X + Y$ is sub-Weibull $(C(K_X + K_Y), \alpha)$ for some absolute constant $C > 0$.*

A helpful result is the following, which bounds the sum of identically distributed sub-Weibull random variables.

Lemma 10 (Corollary 3.1 in Vladimirova et al. (2020)). *Suppose X_1, \dots, X_n are identically distributed (K', α) sub-Weibull random variables. Then, for some absolute constant $c > 0$, for all $s \geq ncK'$, we have*

$$\mathbb{P}\left(\left|\sum_{i=1}^n X_i\right| \geq s\right) \leq \exp\left(-\left(\frac{s}{ncK'}\right)^{1/\alpha}\right)$$

In our work, we frequently need to bound sums of the k -th power of the norm of a standard d -dimensional Gaussian. We do so using Lemma 10.

Lemma 11. *Suppose $X_i \stackrel{i.i.d.}{\sim} N(0, I_d)$ for $i \in [n]$. Then, for any $k \in \mathbb{Z}^+$, there exists absolute constants $c, C > 0$ such that for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$\left|\sum_{i=1}^n \|X_i\|^{2k}\right| \leq nC c^k d^k (1 + (\log(1/\delta))^k).$$

In particular, for any $\delta \in (0, 1/e)$ such that $\log(1/\delta) \geq 1$, it follows that

$$\left|\sum_{i=1}^n \|X_i\|^{2k}\right| \leq 2nC c^k d^k (\log(1/\delta))^k.$$

Proof. First, observe that for any $j \in [d]$, $(X_i)_j^2$, being subExponential, is $(1, 1)$ -subWeibull. Then, by Lemma 9, $\|X_i\|^2 = \sum_{j=1}^d (X_i)_j^2$ is $(cd, 1)$ for some absolute constant c . Now, it follows from definition of sub-Weibullness in Definition 7 that $\|X_i\|^{2k}$ is $(c^k d^k, k)$ -subWeibull. Hence, applying Lemma 10, we have that there exists absolute constant $C > 0$ such that for any $s \geq nC c^k d^k$,

$$\mathbb{P}\left(\left|\sum_{i=1}^n \|X_i\|^{2k}\right| \geq s\right) \leq \exp\left(-\left(\frac{s}{nC c^k d^k}\right)^{1/k}\right)$$

Choosing $s = (1 + (\log(1/\delta))^k) nC c^k d^k$, we arrive then at the desired result. \square

C.5 SUPERMARTINGALE CONCENTRATION INEQUALITIES

We first state and prove a supermartingale-type concentration inequality of the form we later require.

Lemma 12. *Consider a filtration of sigma-algebras $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \dots \subset \mathcal{F}_{n-1} \subset \mathcal{F}_n$ and a sequence of random variables X_1, \dots, X_n such that $X_i \in \mathcal{F}_i$. Suppose that*

$$\mathbb{P}_{\mathcal{F}_{i-1}}(X_i \leq a) = 1 \quad \text{and} \quad \mathbb{P}_{\mathcal{F}_{i-1}}(X_i \leq -b) \geq p \quad (10)$$

for some $a, b > 0$ and $0 < p \leq \frac{1}{2}$. Then, for any $0 < \mu \leq b$ such that $|-b + \mu| \geq \frac{1-p}{p}(a + \mu)$, we have

$$\mathbb{P}\left(\sum_{i=1}^n X_i \geq -n\mu + s\right) \leq \exp\left(-\frac{s^2}{4n(b - \mu)^2}\right), \quad \forall s > 0.$$

Proof. Observe that by Markov's inequality, for any $\lambda > 0$,

$$\mathbb{P}\left(\sum_{i=1}^n X_i \geq -n\mu + s\right) = \mathbb{P}\left(\exp\left(\lambda \sum_{i=1}^n (X_i + \mu)\right) \geq \exp(\lambda s)\right) \leq \frac{\mathbb{E}\left[\exp\left(\lambda \sum_{i=1}^n (X_i + \mu)\right)\right]}{\exp(\lambda s)}.$$

Now, observe that

$$\begin{aligned} \mathbb{E}\left[\exp\left(\lambda \sum_{i=1}^n (X_i + \mu)\right)\right] &= \mathbb{E}\left[\mathbb{E}_{\mathcal{F}_{n-1}}\left[\exp\left(\lambda \sum_{i=1}^n (X_i + \mu)\right)\right]\right] \\ &= \mathbb{E}\left[\exp\left(\lambda \sum_{i=1}^{n-1} (X_i + \mu)\right) \mathbb{E}_{\mathcal{F}_{n-1}}[\exp(\lambda(X_n + \mu))]\right]. \end{aligned} \quad (11)$$

Let us now compute $\mathbb{E}_{\mathcal{F}_{n-1}}[\exp(\lambda(X_n + \mu))]$:

$$\begin{aligned} & \mathbb{E}_{\mathcal{F}_{n-1}}[\exp(\lambda(X_n + \mu))] \\ &= \int_{(-\infty, -b]} \exp(\lambda(x + \mu)) \mathbb{P}_{\mathcal{F}_{n-1}}(X_n \in dx) + \int_{(-b, a]} \exp(\lambda(x + \mu)) \mathbb{P}_{\mathcal{F}_{n-1}}(X_n \in dx) \\ &\leq \mathbb{P}_{\mathcal{F}_{n-1}}(X_n \leq -b) \exp(\lambda(-b + \mu)) + \mathbb{P}_{\mathcal{F}_{n-1}}(-b < X_n \leq a) \exp(\lambda(a + \mu)) \\ &\leq p \exp(\lambda(-b + \mu)) + (1 - p) \exp(\lambda(a + \mu)). \end{aligned}$$

Then observe that by our choice of μ , $-b + \mu < 0$, and that $|-b + \mu| \geq (a + \mu) \frac{1-p}{p}$. Since we assumed $p \leq \frac{1}{2}$, this means that $\frac{1-p}{p} \geq 1$ and so for any $k \geq 1$,

$$|-b + \mu| \geq (a + \mu) \frac{1-p}{p} \implies |-b + \mu| \geq (a + \mu) \left(\frac{1-p}{p}\right)^{1/k} \implies p|-b + \mu|^k \geq (1-p)(a + \mu)^k.$$

Consequently, by Taylor expansion,

$$\begin{aligned} & p \exp(\lambda(-b + \mu)) + (1 - p) \exp(\lambda(a + \mu)) \\ &= 1 + \sum_{k=1}^{\infty} \frac{\lambda^k (p(-b + \mu)^k + (1-p)(a + \mu)^k)}{k!} \leq 1 + \sum_{k=1}^{\infty} \frac{\lambda^k (p(-b + \mu)^k + p|-b + \mu|^k)}{k!} \\ &= 1 + \sum_{k=1}^{\infty} \frac{\lambda^{2k} \cdot 2p|-b + \mu|^{2k}}{(2k)!} \leq 1 + \sum_{k=1}^{\infty} \frac{\lambda^{2k}|-b + \mu|^{2k}}{(k)!} \\ &= \exp(\lambda^2(-b + \mu)^2), \end{aligned}$$

which leads to

$$\mathbb{E}_{\mathcal{F}_{n-1}}[\exp(\lambda(X_n + \mu))] \leq \exp(\lambda^2(-b + \mu)^2).$$

Now, continuing from Eq. (11), we have that

$$\begin{aligned} \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^n (X_i + \mu) \right) \right] &\leq \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{n-1} (X_i + \mu) \right) \mathbb{E}_{\mathcal{F}_{n-1}} \mathbf{1} [\exp(\lambda(X_n + \mu))] \right] \\ &\leq \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{n-1} (X_i + \mu) \right) \exp(\lambda^2(b - \mu)^2) \right] \\ &\leq \dots \\ &\leq \exp(n\lambda^2(b - \mu)^2). \end{aligned}$$

Thus, for any $\lambda > 0$ and $s \geq 0$,

$$\begin{aligned} \mathbb{P} \left(\sum_{i=1}^n X_i \geq -n\mu + s \right) &\leq \frac{\mathbb{E} [\exp(\lambda(\sum_{i=1}^n (X_i + \mu)))]}{\exp(\lambda s)} \\ &\leq \exp(n\lambda^2(b - \mu)^2 - \lambda s) \end{aligned}$$

By finding the minimizing λ , we find that

$$\mathbb{P} \left(\sum_{i=1}^n X_i \geq -n\mu + s \right) \leq \exp \left(-\frac{s^2}{4n(b - \mu)^2} \right),$$

which completes the proof. \square

We will later require a weakened form of a supermartingale concentration inequality, as stated and proven below.

Proposition 3 (Weakened supermartingale concentration inequality). *Consider a filtration of sigma-algebras $\mathcal{F}_0 \subset \mathcal{F}_1 \cdots \subset \mathcal{F}_n$ and a sequence of random variables X_1, \dots, X_n such that $X_i \in \mathcal{F}_i$. Consider for each $i \in \{1, \dots, n\}$ a bad set B_i where $\mathbb{1}_{B_i} \in \mathcal{F}_{i-1}$, and suppose*

$$\mathbb{P}_{\mathcal{F}_{i-1}}(X_i \mathbb{1}_{B_i^c} \leq a) = 1 \quad \text{and} \quad \mathbb{P}_{\mathcal{F}_{i-1}}(X_i \mathbb{1}_{B_i} \leq -b) \geq p$$

for some $a, b > 0$ and $0 \leq p \leq 1/2$. Then, for any $0 < \mu \leq b$ such that $|-b + \mu| \geq \frac{1-p}{p}(a + \mu)$, we have

$$\mathbb{P}\left(\sum_{i=1}^n X_i \geq -n\mu + s\right) \leq \exp\left(-\frac{s^2}{4n(b-\mu)^2}\right) + \sum_{i=1}^n \mathbb{P}(X_i \in B_i), \quad \forall s > 0.$$

Proof. We define $Q_i := X_i \mathbb{1}_{B_i^c}$. We can then apply Lemma 12 and get

$$\mathbb{P}\left(\sum_{i=1}^n Q_i \geq -n\mu + s\right) \leq \exp\left(-\frac{s^2}{4n(b-\mu)^2}\right).$$

Since $\mathbb{P}(X_i \neq Q_i \text{ for some } i \in [n]) \leq \sum_i \mathbb{P}(X_i \in B_i)$, it follows that

$$\mathbb{P}\left(\sum_{i=1}^n X_i \geq -n\mu + s\right) \leq \exp\left(-\frac{s^2}{4n(b-\mu)^2}\right) + \sum_{i=1}^n \mathbb{P}(X_i \in B_i),$$

which completes the proof. \square

D FUNCTION DECREASE IN LARGE GRADIENT REGIME

In this section, we show that sufficient function decrease can be made across the iterations with large gradients. We first restate and prove the function decrease lemma (Lemma 1), first introduced in the main text. We then provide a detailed roadmap of our proof in the subsequent discussion following the proof of Lemma 1.

Lemma 1 (Function decrease for batch zeroth-order optimization). *Suppose at each time t , the algorithm performs the update step (with batch-size parameter $1 \leq m \leq d$)*

$$x_{t+1} = x_t - \eta \left(g_u^{(m)}(x_t) + Y_t \right),$$

where

$$g_u^{(m)}(x_t) = \frac{1}{m} \sum_{i=1}^m \frac{f(x_t + uZ_{t,i}) - f(x_t - uZ_{t,i})}{2u} Z_{t,i},$$

where each $Z_{t,i}$ is drawn i.i.d from $N(0, I)$, $u > 0$ is the smoothing radius, and $Y_t \sim N(0, \frac{r^2}{d} I)$ with $r > 0$ denoting the perturbation radius.

Then, there exist absolute constants $c_1 > 0, C_1 \geq 1$ such that, for any $T \in \mathbb{Z}^+$ and $T \geq \tau > 0$, $\alpha > 0$ and $\delta \in (0, 1/e]$, upon defining $\mathcal{H}_{0,\tau}(\delta)$ to be the event on which the inequality

$$\begin{aligned} f(x_\tau) - f(x_0) &\leq -\frac{3\eta}{4} \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{\alpha} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 \\ &\quad + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\ &\quad + \eta c_1 r^2 (\alpha + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2 \end{aligned} \quad (3)$$

is satisfied (where $\chi := \log(C_1 d m T / \delta)$), we have

$$\mathbb{P}(\mathcal{H}_{0,\tau}(\delta)) \geq 1 - \frac{(\tau + 4)\delta}{T}, \quad \mathbb{P}(\cap_{\tau'=1}^{\tau} \mathcal{H}_{0,\tau}(\delta)) \geq 1 - \frac{5\tau'\delta}{T}$$

for any $0 \leq \tau' \leq T$.

Proof. First, for each $t \in \{-1, \dots, \tau\}$, we define \mathcal{F}_t to be the sigma-algebra generated by

$$x_0, \quad (\{Z_{0,i}\}_{i=1}^m, \dots, \{Z_{t,i}\}_{i=1}^m), \quad (Y_0, \dots, Y_t).$$

Note that \mathcal{F}_{-1} is the sigma-algebra generated only by x_0 .

By Taylor expansion, for any $x, y \in \mathbb{R}^d$, there exists $\alpha \in [0, 1]$ such that $f(x + y) = f(x) + \langle \nabla f(x), y \rangle + \frac{1}{2} y^\top \nabla^2 f(x + \alpha y) y$. Therefore

$$\frac{f(x_t + uZ_{t,i}) - f(x_t - uZ_{t,i})}{2u} = \langle \nabla f(x_t), Z_{t,i} \rangle + \frac{u}{2} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i}$$

with

$$\tilde{H}_{t,i} = \frac{\nabla^2 f(x_t + \alpha_{i,+} u Z_{t,i}) - \nabla^2 f(x_t - \alpha_{i,-} u Z_{t,i})}{2}$$

for some $\alpha_{i,\pm} \in [0, 1]$, and

$$x_{t+1} = x_t - \eta \left(\frac{1}{m} \sum_{i=1}^m \left(Z_{t,i} Z_{t,i}^\top \nabla f(x_t) + \frac{u}{2} Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right) + Y_t \right) \quad (12)$$

By the ρ -Hessian Lipschitz property of f , it follows that $\|\tilde{H}_{t,i}\| \leq \rho u \|Z_{t,i}\|$

Observe that

$$\begin{aligned} f(x_{t+1}) &\stackrel{(i)}{\leq} f(x_t) + \langle x_{t+1} - x_t, \nabla f(x_t) \rangle + \frac{L}{2} \|x_{t+1} - x_t\|^2 \\ &\stackrel{(ii)}{=} f(x_t) - \eta \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 - \eta \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i}^\top \nabla f(x_t) \cdot Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} - \eta \langle \nabla f(x_t), Y_t \rangle \\ &\quad + \frac{L\eta^2}{2} \left\| \frac{1}{m} \sum_{i=1}^m \left(Z_{t,i} Z_{t,i}^\top \nabla f(x_t) + \frac{u}{2} Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right) + Y_t \right\|^2 \\ &\stackrel{(iii)}{\leq} f(x_t) - \frac{\eta}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \frac{\eta}{m} \sum_{i=1}^m \left(\frac{|Z_{t,i}^\top \nabla f(x_t)|^2}{4} + \frac{u^2 |Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i}|^2}{4} \right) - \eta \langle \nabla f(x_t), Y_t \rangle \\ &\quad + \frac{L\eta^2}{2} \left(2 \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) \right\|^2 + u^2 \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2 + 4 \|Y_t\|^2 \right) \\ &\stackrel{(iv)}{\leq} f(x_t) - \frac{3\eta}{4m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \frac{\eta u^2}{m} \sum_{i=1}^m \frac{u^2 \rho^2 \|Z_{t,i}\|^6}{4} - \eta \langle \nabla f(x_t), Y_t \rangle \\ &\quad + \frac{L\eta^2}{2} \left(2 \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) \right\|^2 + \frac{u^2}{m} \sum_{i=1}^m u^2 \rho^2 \|Z_{t,i}\|^8 + 4 \|Y_t\|^2 \right) \\ &\leq f(x_t) - \frac{3\eta}{4m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \frac{\eta u^4 \rho^2}{4m} \sum_{i=1}^m \|Z_{t,i}\|^6 + \frac{L\eta^2 u^4 \rho^2}{2m} \sum_{i=1}^m \|Z_{t,i}\|^8 - \eta \langle \nabla f(x_t), Y_t \rangle \\ &\quad + \frac{L\eta^2}{2} \left(2 \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) \right\|^2 + 4 \|Y_t\|^2 \right) \end{aligned} \quad (13)$$

Above, to derive (i), we used the L -smoothness of f . To derive (ii), we used the expression for $(x_{t+1} - x_t)$ shown in Eq. (12). To derive (iii), we used the fact that $ab \leq (a^2 + b^2)/2$ for any $a, b \in \mathbb{R}_{\geq 0}$, as well as two applications of the fact that $\|a + b\|^2 \leq 2(\|a\|^2 + \|b\|^2)$ for any two vectors $a, b \in \mathbb{R}^d$. To derive (iv), we used the fact that $\|\tilde{H}_{t,i}\| \leq \rho u \|Z_{t,i}\|$.

To continue from Eq. (13), we first observe that we can rewrite

$$Z_{t,i} Z_{t,i}^\top \nabla f(x_t) = (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) + \nabla f(x_t),$$

so that

$$\left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) \right\|^2 \leq 2 \left\| \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\|^2 + 2 \|\nabla f(x_t)\|^2.$$

Observe that we can apply the bound in Proposition 2 to $\left\| \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\|$, and since $Z_{t,i}$ is independent of \mathcal{F}_{t-1} for all i , we know there exist absolute constants $\mathbf{c}_1 > 0, C_1 \geq 1$ such that for any $\delta \in (0, 1/e]$ and $\theta > 0$, with probability at least $1 - \delta$ conditioned on \mathcal{F}_{t-1} ,

$$\begin{aligned} \left\| \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\| &\leq \mathbf{c}_1 \theta \sum_{i=1}^m d (\text{lr}(C_1 m / \delta))^2 \|\nabla f(x_t)\|^2 + \frac{1}{\theta} \log(C_1 d m / \delta) \\ &= \mathbf{c}_1 \theta m d (\text{lr}(C_1 m / \delta))^2 \|\nabla f(x_t)\|^2 + \frac{1}{\theta} \log(C_1 d m / \delta). \end{aligned} \quad (14)$$

Moreover, since $C_1 \geq 1$, $\log(C_1 d m / \delta)$ and $\text{lr}(C_1 m / \delta)$ both are at least 1 as long as $\delta \in (0, 1/e]$. Observe that conditioned on \mathcal{F}_{t-1} , $\nabla f(x_t)$ is fixed. Hence, we can pick

$$\theta = \frac{1}{\sqrt{\mathbf{c}_1 m d \text{lr}(C_1 d m / \delta)} \|\nabla f(x_t)\|}$$

which is \mathcal{F}_{t-1} -measurable, and plug it into Eq. (14) to find that the probability conditioned on \mathcal{F}_{t-1} of the following event

$$\left\| \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\| \leq 2\sqrt{\mathbf{c}_1} (\text{lr}(C_1 d m / \delta))^{3/2} \sqrt{m d} \|\nabla f(x_t)\| \quad (15)$$

is at least $1 - \delta$. By taking the total expectation, it follows that the event has a total probability at least $1 - \delta$. Thus, with probability at least $1 - \delta$,

$$\begin{aligned} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) \right\|^2 &\leq 2 \left\| \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\|^2 + 2 \|\nabla f(x_t)\|^2 \\ &\leq 4\mathbf{c}_1 (\text{lr}(C_1 d m / \delta))^3 \frac{d}{m} \|\nabla f(x_t)\|^2 + 2 \|\nabla f(x_t)\|^2 \\ &\leq \mathbf{c}_2 (\text{lr}(C_1 d m / \delta))^3 \frac{d}{m} \|\nabla f(x_t)\|^2, \end{aligned} \quad (16)$$

where the last inequality comes from the fact that $\text{lr}(C_1 d m / \delta) \geq 1$, our assumption at the outset of the appendix that $d \geq m$, and denoting $\mathbf{c}_2 := 4\mathbf{c}_1 + 2$.

Denote the event $\tilde{H}_{0,\tau}(\delta)$ as the event that

$$\begin{aligned} f(x_\tau) - f(x_0) &\leq - \sum_{t=0}^{\tau-1} \frac{3\eta}{4m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + L\eta^2 \frac{\mathbf{c}_2 d (\text{lr}(C_1 d m / \delta))^3}{m} \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 \\ &\quad + \frac{\eta u^4 \rho^2}{4m} \sum_{t=0}^{\tau-1} \sum_{i=1}^m \|Z_{t,i}\|^6 + \frac{L\eta^2 u^4 \rho^2}{2m} \sum_{t=0}^{\tau-1} \sum_{i=1}^m \|Z_{t,i}\|^8 \\ &\quad - \eta \sum_{t=0}^{\tau-1} \langle \nabla f(x_t), Y_t \rangle + 2L\eta^2 \sum_{t=0}^{\tau-1} \|Y_t\|^2 \end{aligned} \quad (17)$$

holds.

Now, continuing from Eq. (13), and using the bound in Eq. (16), summing over the iterations from $t = 0$ to $\tau - 1$, we find using the union bound that $\mathbb{P}(\cap_{\tau=1}^{\tau'} \tilde{H}_{0,\tau}(\delta)) \geq 1 - \tau' \delta$, $\mathbb{P}(\tilde{H}_{0,\tau}(\delta)) \geq 1 - \tau \delta$.

Now, by Lemma 6, for any $\delta \in (0, 1)$, $\alpha > 0$, with probability at least $1 - \delta$, there exists an absolute constant $\mathbf{c}_3 > 0$ such that

$$-\eta \sum_{t=0}^{\tau-1} \langle \nabla f(x_t), Y_t \rangle \leq \eta \left(\frac{1}{\alpha} \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 + \mathbf{c}_3 \alpha r^2 \log(1/\delta) \right). \quad (18)$$

Meanwhile, since $Y_t \sim N(0, (r^2/d)I)$, $\|Y_t\|^2$ is sub-exponential with sub-exponential norm cr^2 for some absolute constant $c > 0$, and by Bernstein's inequality (Lemma 8), there exists some absolute constant $\mathbf{c}_4 > 0$ such that

$$\sum_{t=0}^{\tau-1} \|Y_t\|^2 \leq \mathbf{c}_4 r^2 (\tau + \log(1/\delta)) \quad (19)$$

with probability at least $1 - \delta$.

To bound $\sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}\|^6$ and $\sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}\|^8$, both sums of heavy tailed Gaussian moments, we use Lemma 11, which states that for any $k \in \mathbb{Z}^+$ and $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\frac{1}{m} \sum_{t=0}^{\tau-1} \sum_{i=1}^m \|Z_{t,i}\|^{2k} \leq c_5 \tau (c_6)^k d^k (1 + (\log(1/\delta))^k) \quad (20)$$

for some absolute constants $c_5, c_6 > 0$. As in the statement of the proof, using $\chi := \text{lr}(C_1 dm/\delta)$ to ease the notation, denote the event that

$$\begin{aligned} f(x_\tau) - f(x_0) &\leq -\frac{3\eta}{4} \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{\alpha} + \frac{c_2 L \eta^2 \chi^3 d}{m} \right) \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 \\ &\quad + \frac{\tau \eta u^4 \rho^2}{2} \cdot c_5 c_6^3 d^3 \left(\log \frac{1}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_5 c_6^4 d^4 \left(\log \frac{1}{\delta} \right)^4 \\ &\quad + \eta (c_3 \alpha r^2 + 2c_4 \eta L r^2) \log \frac{1}{\delta} + 2c_4 L \eta^2 \tau r^2 \end{aligned}$$

holds as $\mathcal{H}_{0,\tau}(\delta)$.

Plugging Eq. (18), Eq. (19), and Eq. (20) into Eq. (17), by union bound, we see that

$$\mathbb{P}(\cap_{\tau=1}^{\tau'} \mathcal{H}_{0,\tau}(\delta)) \geq 1 - (\tau' + 4\tau')\delta = 1 - 5\tau'\delta, \quad \mathbb{P}(\mathcal{H}_{0,\tau}) \geq 1 - (\tau + 4)\delta.$$

The final result then follows by rescaling δ to $\frac{\delta}{T}$ and denoting $c_1 := \max\{c_2, c_3, 2c_4, c_5 c_6^3/2, c_5 c_6^4\}$. \square

Outline of proof approach. Similar to the first-order setting, our goal is to show that we can arrive at a contradiction $f(x_T) < \min_x f(x)$ when there is a large number of steps at which $\|\nabla f(x_t)\| \geq \epsilon$. Roughly speaking, as Eq. (3) shows, we need to prove a lower bound of the form

$$\sum_{t=0}^{T-1} \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2 \geq \Omega \left(\frac{1}{\alpha} + \frac{c_1 L \eta \chi^3 d}{m} \right) \sum_{t=0}^{T-1} \|\nabla f(x_t)\|^2 \quad (21)$$

for some α which is not too large (an example would be picking α such that it only scales logarithmically in the problem parameters). However, it is tricky to prove such a lower-bound in the zeroth-order setting. In particular, for small batch-sizes m , $\frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2$ could be small even as $\|\nabla f(x_t)\|^2$ is large; this is because for each $i \in [m]$, $Z_{t,i}$ could have a negligible component in the $\nabla f(x_t)$ direction. This necessitates a more careful analysis to prove a bound similar to Eq. (21). We do so using the following approach.

1. Intuitively, whilst for each individual iteration t , $\frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2$ could be small even as $\|\nabla f(x_t)\|^2$ is large, in a small number of (consecutive) iterations $\{t_0, \dots, t_0 + t_f\}$, with high probability, there will be at least one iteration t within $\{t_0, \dots, t_0 + t_f - 1\}$, such that $\frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2 = \Omega(\|\nabla f(x_t)\|^2)$. We formalize this intuition in Lemma 14. Thus, we consider breaking the time-steps into chunks where each chunk has t_f consecutive iterations.
2. Consider any such interval $\{t_0, \dots, t_0 + t_f - 1\}$. There are two cases to consider.
 - (a) The first case is when the gradient throughout all t_f iterations is large enough to dominate the perturbation terms. Intuitively, in this case, it is not hard to see that given appropriate parameter choices, the gradient will change little throughout the t_f iterations. In fact, as we formalize in Lemma 16, for an appropriate choice of t_f and η , we can show that

$$\frac{1}{2} \|\nabla f(x_{t_0})\| \leq \|\nabla f(x_t)\| \leq 2 \|\nabla f(x_{t_0})\| \quad \forall t \in \{t_0, \dots, t_0 + t_f - 1\}.$$

As a result, combined with point 1, we see that

$$\sum_{t=t_0}^{t_0+t_f-1} \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2 \geq \Omega(\|\nabla f(x_{t_0})\|^2).$$

Thus, by choosing α and η judiciously, for such intervals, it is possible to show that

$$\begin{aligned} \sum_{t=t_0}^{t_0+t_f-1} \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}^\top \nabla f(x_t)\|^2 &\geq \Omega(\|\nabla f(x_{t_0})\|^2) \geq \Omega\left(\frac{1}{\alpha} + \frac{c_1 L \eta \chi^3 d}{m}\right) \sum_{t=t_0}^{t_0+t_f-1} \|\nabla f(x_t)\|^2 \\ &= \Omega\left(\frac{1}{\alpha} + \frac{c_1 L \eta \chi^3 d}{m}\right) \Omega\left(t_f \|\nabla f(x_{t_0})\|^2\right) \end{aligned}$$

Thus, in these intervals, it is possible to obtain function improvement on the order of $\eta \Omega(\|\nabla f(x_{t_0})\|^2)$.

- (b) The remaining case is when the gradient is small and dominated by the perturbation terms in any one of the t_f iterations. In this case, as we show in Lemma 17, for each of the t_f iterations, the gradient will be small and on the same scale as the perturbation terms. In turn, by choosing r, u and η appropriately, we can make the perturbation terms small. Thus, whilst these intervals may not contribute to function decrease, they also contribute little in the way of function increase.
3. When there are at least $T/4$ iterations with large gradient (i.e. $\|\nabla f(x_t)\| \geq \epsilon$), assuming t_f divides T , it follows that there are at least $T/(4t_f)$ intervals of length t_f where one iteration in the interval contains a large gradient. By choosing u, r and η appropriately such they are dominated by ϵ , it is possible to show that with high probability, such an interval cannot belong to the second case above, and must instead be from the first case. Since $\|\nabla f(x_t)\| \approx \|\nabla f(x_{t_0})\|$ for each $t \in \{t_0, \dots, t_0 + t_f - 1\}$ in this case, and we know that one of the iterations has a gradient with size at least ϵ , it follows that we make function decrease progress of at least $\eta \Omega(\epsilon^2)$ for such intervals. By appropriately choosing η, u and r to limit the effects of the intervals of the second form, we can then show a contradiction of the form $f(x_T) < f^*$. We demonstrate this formally in Proposition 4.

We formalize our approach in the following series of results. First, for analytical convenience, we prove the following result showing that for any t , the perturbation terms $\|Y_t\|$ and $\frac{1}{m} \sum_{i=1}^m \|Z_{t,i}\|^4$ are bounded with high probability.

Lemma 13. *There exists an absolute constant $c_3 > 0$ such that, for any $t \in \mathbb{N}$, the event*

$$\mathcal{G}_t(\delta) := \left\{ \|Y_t\|^2 \leq c_3^2 r^2 \left(1 + \frac{\log(T/\delta)}{d}\right) \text{ and } \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}\|^4 \leq 2c_3 d^2 \left(\log \frac{T}{\delta}\right)^2 \right\}$$

has probability at least $1 - 2\delta/T$ for any $\delta \in (0, 1/e]$.

Proof. Noting that $Y_t \sim N(0, (r^2/d)I)$, by applying Bernstein's inequality (Lemma 8), it can be shown that with probability at least δ/T ,

$$\|Y_t\|^2 \leq c_3^2 r^2 \left(1 + \frac{\log(T/\delta)}{d}\right),$$

where $c_3 > 0$ is some absolute constant. Then by using Lemma 11, applying the union bound, and redefining the constant c_3 , we complete the proof. \square

Next, in Lemma 14, we show that in a small number of iterations, with high probability, there exists some iteration t such that $\frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 \geq \frac{1}{2} \|\nabla f(x_t)\|^2$.

Lemma 14. *There exists an absolute constant $c_2 \geq 1$ such that, upon defining*

$$t_f(\delta) = \left\lceil \frac{c_2}{m} \log \frac{T}{\delta} \right\rceil, \quad \delta > 0,$$

and defining the event

$$\mathcal{B}_{t_0}(\delta; k) := \bigcup_{t=t_0}^{t_0+k-1} \left\{ \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 \geq \frac{1}{2} \|\nabla f(x_t)\|^2 \right\},$$

we have

$$\mathbb{P}(\mathcal{B}_{t_0}(\delta; k)) \geq 1 - \frac{\delta}{T}.$$

for any $\delta \in (0, 1)$, $t_0 \in \mathbb{N}$ and $k \geq t_f(\delta)$.

Proof. Denote the event

$$E_t = \left\{ \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 < \frac{1}{2} \|\nabla f(x_t)\|^2 \right\}.$$

Observe that, conditioned on \mathcal{F}_{t-1} , the set of random variables $\{\|\nabla f(x_t)\|^2 - |Z_{t,i}^\top \nabla f(x_t)|^2\}_{i=1}^m$ are independent, mean-zero, and subexponential with subexponential norm $\leq c\|\nabla f(x_t)\|^2$ for some absolute constant $c > 0$. Hence

$$\begin{aligned} \mathbb{P}_{\mathcal{F}_{t-1}}(E_t) &= \mathbb{P}_{\mathcal{F}_{t-1}} \left(\frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 < \frac{1}{2} \|\nabla f(x_t)\|^2 \right) \\ &= \mathbb{P}_{\mathcal{F}_{t-1}} \left(\sum_{i=1}^m \left(\|\nabla f(x_t)\|^2 - |Z_{t,i}^\top \nabla f(x_t)|^2 \right) > \frac{m}{2} \|\nabla f(x_t)\|^2 \right) \\ &\leq \exp(-c'm), \end{aligned}$$

where c' is some positive absolute constant. Then, for any $t_0, k \in \mathbb{N}$,

$$\begin{aligned} &\mathbb{P} \left(\frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 < \frac{1}{2} \|\nabla f(x_t)\|^2 \text{ for every } t \in [t_0, t_0 + k) \right) \\ &= \mathbb{E} \left[\prod_{t=t_0}^{t_0+k-1} \mathbb{1}_{E_t} \right] = \mathbb{E} \left[\prod_{t=t_0}^{t_0+k-2} \mathbb{1}_{E_t} \cdot \mathbb{E}_{\mathcal{F}_{t_0+k-2}} [\mathbb{1}_{E_{t_0+k-1}}] \right] \\ &\leq \exp(-c'm) \cdot \mathbb{E} \left[\prod_{t=t_0}^{t_0+k-2} \mathbb{1}_{E_t} \right] \leq \dots \leq \exp(-c'mk). \end{aligned}$$

Therefore, by letting $c_2 = \max\{1, 1/c'\}$ and

$$k \geq t_f(\delta) = \left\lceil \frac{c_2}{m} \log \frac{T}{\delta} \right\rceil,$$

we get

$$\mathbb{P} \left(\frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 < \frac{1}{2} \|\nabla f(x_t)\|^2 \text{ for every } t \in [t_0, t_0 + k) \right) \leq \frac{\delta}{T},$$

which completes the proof. \square

The term $t_f(\delta)$ will frequently appear in the proofs to come; in the sequel we denote

$$t_f(\delta) := \left\lceil \frac{c_2}{m} \log \frac{T}{\delta} \right\rceil, \quad \delta \in (0, 1/e], \quad (22)$$

where $c_2 \geq 1$ is the absolute constant defined in Lemma 14.

We next show that with high probability, the norm difference term $\|\nabla f(x_{t+1}) - \nabla f(x_t)\|$ can be bounded in terms of $\|\nabla f(x_t)\|$ and the perturbation terms $\left\| \frac{u}{2m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|$ as well as $\|Y_t\|$.

Lemma 15. *Define*

$$\mathcal{A}_t(\delta) := \left\{ \left\| \nabla f(x_{t+1}) - \nabla f(x_t) \right\| \leq \frac{\|\nabla f(x_t)\|}{8t_f(\delta)} + \eta L \left(\left\| \frac{u}{2m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right) \right\} \quad (23)$$

where $t_f(\delta)$ is defined in Eq. (22), and let $C_1 \geq 1$ be the corresponding absolute constants defined in Lemma 1. Then there exists an absolute constant $c_4 > 0$ such that, whenever η satisfies

$$\eta L \frac{c_4 (\ln(C_1 dmT/\delta))^{3/2} \sqrt{d}}{\sqrt{m}} \leq \frac{1}{8t_f(\delta)}, \quad (24)$$

we have

$$\mathbb{P}(\mathcal{A}_t(\delta)) \geq 1 - \frac{\delta}{T}$$

for any $\delta \in (0, 1/e]$ and $t \in \mathbb{Z}^+$.

Proof. Since ∇f is L -Lipschitz, following the zeroth-order update step, we see that

$$\|\nabla f(x_{t+1}) - \nabla f(x_t)\| \leq L \|x_{t+1} - x_t\| \quad (25)$$

$$= \eta L \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) + \frac{u}{2m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} + Y_t \right\|. \quad (26)$$

Now, it follows from Eq. (16) (with a slight modification in the absolute constant terms since here the norm is not squared) that there exists some absolute constant $c_4 > 0$ such that for any $\delta \in (0, 1/e]$, we have that with probability at least $1 - \delta/T$, the event

$$\left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) \right\| \leq c_4 (\ln(C_1 dmT/\delta))^{3/2} \sqrt{\frac{d}{m}} \|\nabla f(x_t)\|,$$

Hence, continuing from Eq. (26), it follows that with probability at least $1 - \delta/T$,

$$\begin{aligned} & \|\nabla f(x_{t+1}) - \nabla f(x_t)\| \\ & \leq \eta L \left(c_4 (\ln(C_1 dmT/\delta))^{3/2} \sqrt{\frac{d}{m}} \|\nabla f(x_t)\| + \left\| \frac{u}{2m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right), \end{aligned}$$

and by plugging in the condition Eq. (24), we see that the event

$$\mathcal{A}_t(\delta) = \left\{ \left\| \nabla f(x_{t+1}) - \nabla f(x_t) \right\| \leq \frac{\|\nabla f(x_t)\|}{8t_f(\delta)} + \eta L \left(\left\| \frac{u}{2m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right) \right\}$$

has probability at least $1 - \delta/T$. \square

We show now that if the norm of the gradient dominates the norm of the perturbation terms, and we choose the step-size η sufficiently small, then in a small number of iterations, the norm of the gradient does not change very much. For notational simplicity, we denote the event

$$\mathcal{E}(t_1, t_2, \delta) := \bigcap_{t=t_1}^{t_1+t_2-1} \left\{ \|\nabla f(x_t)\| > 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right) \right\}.$$

Lemma 16. *Let $\delta \in (0, 1/e]$ and $T \in \mathbb{Z}^+$ be such that $T > 2t_f(\delta) + 1$. Consider any positive integer $t'_f \leq 2t_f(\delta)$, and any $t_0 \in \{0, \dots, T - 1 - t'_f\}$. Suppose η satisfies the condition Eq. (24). Then, on the event*

$$\mathcal{E}(t_0, t'_f, \delta) \cap \left(\bigcap_{t=t_0}^{t_0+t'_f-1} \mathcal{A}_t(\delta) \right),$$

we have

$$\frac{1}{2} \|\nabla f(x_0)\| \leq \|\nabla f(x_t)\| \leq 2 \|\nabla f(x_0)\|$$

for all $t \in \{t_0, \dots, t_0 + t'_f - 1\}$.

Proof. By plugging

$$\|\nabla f(x_t)\| > 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right)$$

into the definition of $\mathcal{A}_t(\delta)$, we see that, on the event $\mathcal{E}(t_0, t'_f, \delta) \cap \left(\bigcap_{t=t_0}^{t_0+t'_f-1} \mathcal{A}_t(\delta) \right)$, we have

$$\|\nabla f(x_{t+1}) - \nabla f(x_t)\| \leq \frac{\|\nabla f(x_t)\|}{4t_f(\delta)},$$

and consequently,

$$\left(1 - \frac{1}{4t_f(\delta)} \right) \|\nabla f(x_t)\| \leq \|\nabla f(x_{t+1})\| \leq \left(1 + \frac{1}{4t_f(\delta)} \right) \|\nabla f(x_t)\|,$$

which leads to

$$\left(1 - \frac{1}{4t_f(\delta)} \right)^{t-t_0} \|\nabla f(x_0)\| \leq \|\nabla f(x_t)\| \leq \left(1 + \frac{1}{4t_f(\delta)} \right)^{t-t_0} \|\nabla f(x_0)\|$$

for all $t \in \{t_0, \dots, t_0 + t'_f\}$. Then, since $(1 + 1/(4x))^{2x} \leq 2$ and $(1 - 1/(4x))^{2x} \geq 1/2$ for any $x \geq 1$, noting that $t'_f \leq 2t_f(\delta)$, we get the desired result. \square

Conversely, in the following result, we show that in a small number of consecutive iterations, if the gradient is smaller than the perturbation terms in any one of the iterations, then for each of the iterations in this range, the gradient will be small and be on the same scale as the size of the perturbation terms.

Lemma 17. *Let $\delta \in (0, 1/e]$ and $T \in \mathbb{Z}^+$ be such that $T > 2t_f(\delta) + 1$. Consider any positive integer $t'_f \leq 2t_f(\delta)$, and any $t_0 \in \{0, \dots, T - 1 - t'_f\}$. Suppose η satisfies the condition Eq. (24). Then, on the event*

$$\mathcal{E}^c(t_0, t'_f, \delta) \cap \left(\bigcap_{t=t_0}^{t_0+t'_f-1} \mathcal{A}_t(\delta) \right) \cap \left(\bigcap_{t=t_0}^{t_0+t'_f-1} \mathcal{G}_t(\delta) \right),$$

we have

$$\|\nabla f(x_t)\| \leq c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \quad \forall t \in \{t_0, t_0+1, \dots, t_0+t'_f-1\},$$

where c_5 is some absolute constant.

Proof. Let t' be the first iteration in $\{t_0, t_0 + 1, \dots, t_0 + t'_f - 1\}$ such that

$$\|\nabla f(x_{t'})\| \leq 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t',i} Z_{t',i}^\top \tilde{H}_{t',i} Z_{t',i} \right\| + \|Y_{t'}\| \right). \quad (27)$$

Since we are working on an event which is a subset of $\mathcal{E}^c(t_0, t'_f, \delta)$, t' is well-defined. By $\|\tilde{H}_{t',i}\| \leq \rho u \|Z_{t',i}\|$, we see that

$$\begin{aligned} \|\nabla f(x_{t'})\| &\leq 8t_f(\delta)\eta L \left(\frac{u^2 \rho}{2m} \sum_{i=1}^m \|Z_{t',i}\|^4 + \|Y_{t'}\| \right) \\ &\leq 8t_f(\delta)\eta L \left(c_3 u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + c_3 \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right), \end{aligned}$$

where we used the definition of $\mathcal{G}_t(\delta)$.

Recall that t' is the first time step such that Eq. (27) holds. By deriving similarly as in the proof of Lemma 16, we can show that for any $j \in \{t_0, t_0 + 1, \dots, t' - 1\}$,

$$\|\nabla f(x_j)\| \leq 2\|\nabla f(x_{t'})\| \leq 16t_f(\delta)\eta Lc_3 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right).$$

Meanwhile, for iterations $t \in [t', t_0 + t'_f]$, by using the definitions of $\mathcal{A}_t(\delta)$ and $\mathcal{G}_t(\delta)$, we have

$$\begin{aligned} \|\nabla f(x_{t+1})\| &\leq \left(1 + \frac{1}{8t_f(\delta)}\right) \|\nabla f(x_t)\| + \eta Lc_3 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \\ &= \left(1 + \frac{1}{8t_f(\delta)}\right)^{t+1-t'} \|\nabla f(x_{t'})\| \\ &\quad + \sum_{i=0}^{t-t'} \left(1 + \frac{1}{8t_f(\delta)}\right)^{t-t'-i} \eta Lc_3 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \\ &\leq \left(1 + \frac{1}{8t_f(\delta)}\right)^{t'_f} \|\nabla f(x_{t'})\| \\ &\quad + 8t_f(\delta) \left(\left(1 + \frac{1}{8t_f(\delta)}\right)^{t'_f} - 1 \right) \eta Lc_3 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \\ &\leq e^{1/4} \cdot 8t_f(\delta) \eta Lc_3 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \\ &\quad + 8t_f(\delta)(e^{1/4} - 1) \cdot \eta Lc_3 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \\ &\leq 16t_f(\delta) \eta Lc_3 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right), \end{aligned}$$

where we used $t'_f \leq 2t_f(\delta)$ and the fact that $(1 - 1/(8x))^{2x} \leq e^{1/4}$ for all $x > 0$. By defining $c_5 := 16c_3$, we complete the proof. \square

We next derive a useful result showing that the function change $f(x_\tau) - f(x_0)$ can be decomposed into one component arising from intervals when the gradient dominates noise (which improves function value) and another component arising from intervals with small gradient which may add to function value but whose contributions are bounded in terms of η , u and r . For now, we focus on the case $\tau \geq t_f(\delta)$, since it will be useful to us in proving that there cannot be more than $T/4$ iterations with large gradient.

Lemma 18 (Function change for large τ). *Let $c_1 > 0, c_4 > 0, c_5 > 0, C_1 \geq 1$ be the absolute constants defined in the statements of the previous lemmas. Let $\delta \in (0, 1/e]$, and let $\tau \geq t_f(\delta)$ be arbitrary. Consider splitting $\{0, 1, \dots, \tau - 1\}$ into $K := \lfloor \tau/t_f(\delta) \rfloor$ intervals:*

$$\begin{aligned} J_k &= \{kt_f(\delta), \dots, (k+1)t_f(\delta) - 1\}, \quad 0 \leq k < K - 1, \\ J_{K-1} &= \{(K-1)t_f(\delta), \dots, \tau - 1\}. \end{aligned}$$

Let I_1 denote the set of indices k such that for every time-step t in the interval J_k , the gradient dominates the noise terms as

$$\|\nabla f(x_t)\| > 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right). \quad (28)$$

Suppose we choose η such that

$$\eta \leq \frac{1}{Lt_f(\delta)} \cdot \min \left\{ \frac{\sqrt{m}}{8c_4(\ln(C_1 dmT/\delta))^{3/2} \sqrt{d}}, \frac{m}{128c_1(\ln(C_1 dmT/\delta))^3 d} \right\}. \quad (29)$$

Then, on the event

$$\mathcal{E}_\tau(\delta) := \mathcal{H}_\tau(\delta) \cap \left(\bigcap_{t=0}^{\tau-1} \mathcal{A}_t(\delta) \right) \cap \left(\bigcap_{t=0}^{\tau-1} \mathcal{G}_t(\delta) \right) \cap \left(\bigcap_{k=0}^{K-2} \mathcal{B}_{kt_f(\delta)}(\delta; t_f(\delta)) \right) \cap \mathcal{B}_{(K-1)t_f(\delta)}(\delta; \tau - (K-1)t_f(\delta)),$$

we have the following upper bound on function value change:

$$\begin{aligned} f(x_\tau) - f(x_0) &\leq - \sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 + \tau \frac{c_5^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\ &\quad + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\ &\quad + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2. \end{aligned} \quad (30)$$

Moreover, $\mathbb{P}(\mathcal{E}_\tau(\delta)) \geq 1 - \frac{(5\tau+4)\delta}{T}$.

Proof. Without loss of generality, we may assume that τ is a multiple of $t_f(\delta)$.⁴ Then, any interval $J_k = \{t_0, \dots, t_0 + t_f(\delta) - 1\}$ belongs to one of the following two cases:

Case 1 (Gradient dominates noise): Recall that this means that for every $t \in J_k$, we have

$$\|\nabla f(x_t)\| > 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right).$$

By our choice of η in Eq. (29), we can apply Lemma 16 to get

$$\min_{t \in J_k} \|\nabla f(x_t)\| \geq \frac{1}{4} \max_{t \in J_k} \|\nabla f(x_t)\|.$$

We now consider the two cases when J has fewer than $t_f(\delta)$ iterations and when $J = J_k$ f

Note also that on the event $\mathcal{B}_{kt_f(\delta)}(\delta; t_f(\delta))$, there exists some $t \in J_k$ such that

$$\frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 \geq \frac{1}{2} \|\nabla f(x_t)\|^2.$$

This implies then that

$$\begin{aligned} \frac{1}{4} \sum_{t \in J_k} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 &\geq \frac{1}{4} \min_{t \in J_k} \|\nabla f(x_t)\|^2 \geq \frac{1}{64} \max_{t \in J_k} \|\nabla f(x_t)\|^2 \\ &\geq \frac{1}{64 t_f(\delta)} \sum_{t \in J_k} \|\nabla f(x_t)\|^2. \end{aligned} \quad (31)$$

Thus by setting $\alpha = 128 t_f(\delta)$ in Eq. (3) and by choosing η such that

$$\frac{c_1 L \eta^2 \chi^3 d}{m} \leq \frac{\eta}{\alpha} = \frac{\eta}{128 t_f(\delta)} \iff \eta \leq \frac{m}{128 c_1 L t_f(\delta) d \chi^3},$$

it follows that

$$\begin{aligned} &- \frac{3\eta}{4} \sum_{t \in J_k} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{128 t_f(\delta)} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t \in J_k} \|\nabla f(x_t)\|^2 \\ &= - \frac{3\eta}{4} \sum_{t \in J_k} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \frac{\eta}{64 t_f(\delta)} \sum_{t \in J_k} \|\nabla f(x_t)\|^2 \end{aligned}$$

⁴To accommodate the last interval which has length at most $2t_f(\delta) - 1$, we note that the results we require for the proof, namely Lemma 14, Lemma 16 and Lemma 17, all hold for any interval length $t'_f \leq 2t_f(\delta)$.

$$\begin{aligned}
&\leq -\frac{\eta}{2} \sum_{t \in J_k} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 \\
&\leq -\frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2
\end{aligned} \tag{32}$$

Case 2 (Gradient does not dominate noise): there exists some $t \in J_k$ such that

$$\|\nabla f(x_t)\| \leq 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right).$$

By our choice of η in Eq. (29), we can apply Lemma 17 to get

$$\|\nabla f(x_t)\| \leq c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \quad \forall t \in J_k.$$

Hence, by setting $\alpha = 128t_f(\delta)$ in Eq. (3) and choosing η such that

$$\frac{c_1 L \eta^2 \chi^3 d}{m} \leq \frac{\eta}{\alpha} = \frac{\eta}{128t_f(\delta)},$$

it follows that

$$\begin{aligned}
&\left(\frac{\eta}{128t_f(\delta)} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t \in J_k} \|\nabla f(x_t)\|^2 \\
&\leq \frac{\eta}{64t_f(\delta)} \sum_{t \in J_k} \left(c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \right)^2 \\
&\leq \frac{c_5^2}{64} t_f(\delta)^2 \eta^3 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right)^2
\end{aligned} \tag{33}$$

Without loss of generality, we may assume that τ is a multiple of $t_f(\delta)$.⁵ Then, any interval $J_k = \{t_0, \dots, t_0 + t_f(\delta) - 1\}$ belongs to one of the following two cases:

Having studied the two cases, we may now proceed to use them to complete the proof. Let I_1^c denote the complement of I_1 in $\{0, 1, \dots, K-1\}$. Then,

$$\begin{aligned}
&-\frac{3\eta}{4} \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{\alpha} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 \\
&= \sum_{k \in I_1} \left(-\frac{3\eta}{4} \sum_{t \in J_k} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{128t_f(\delta)} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t \in J_k} \|\nabla f(x_t)\|^2 \right) \\
&\quad + \sum_{k \in I_1^c} \left(-\frac{3\eta}{4} \sum_{t \in J_k} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{128t_f(\delta)} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t \in J_k} \|\nabla f(x_t)\|^2 \right) \\
&\leq -\sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 + \sum_{k \in I_1^c} t_f(\delta) \left(\frac{c_5^2}{64} t_f(\delta)^2 \eta^3 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right)^2 \right) \\
&\leq -\sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 + \tau \frac{c_5^2}{64} t_f(\delta)^2 \eta^3 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right)^2.
\end{aligned} \tag{34}$$

⁵To accommodate the last interval which has length at most $2t_f(\delta) - 1$, we note that the results we require for the proof, namely Lemma 14, Lemma 16 and Lemma 17, all hold for any interval length $t'_f \leq 2t_f(\delta)$.

and so by Eq. (3),

$$\begin{aligned} f(x_\tau) - f(x_0) &\leq - \sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 + \tau \frac{c_5^2}{64} t_f(\delta)^2 \eta^3 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right)^2 \\ &\quad + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\ &\quad + \eta c_1 r^2 (\alpha + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2. \end{aligned}$$

Note that we choose $\alpha = 128t_f(\delta)$. In addition, observe that by our choice of δ (such that $\delta \leq \frac{1}{e}$), it follows that $\sqrt{1 + \frac{\log(T/\delta)}{d}} \leq \sqrt{2 \log(T/\delta)}$.

We can now complete our proof by using the union bound (suppressing the dependence of some of the events on δ for notational simplicity) to derive

$$\begin{aligned} \mathbb{P}(\mathcal{E}_\tau^c) &\leq \mathbb{P}(\mathcal{H}_\tau^c) + \sum_{t=0}^{\tau-1} \mathbb{P}(\mathcal{A}_t^c) + \sum_{t=0}^{\tau-1} \mathbb{P}(\mathcal{G}_t^c) + \sum_{k=0}^{K-1} \mathbb{P}(\mathcal{B}_{kt_f(\delta)}^c(\delta; t_f(\delta))) \\ &\leq \frac{(\tau+4)\delta}{T} + \frac{\tau}{T}\delta + 2\frac{\tau}{T}\delta + \frac{K\delta}{T} \leq \frac{(5\tau+4)}{T}\delta. \quad \square \end{aligned}$$

We are now ready to show that if sufficiently many iterations have a large gradient, then with high probability, the function value of the last iterate $f(x_T)$, will be less than $\min_x f(x)$, a contradiction. Hence this limits the number of iterations that can have a large gradient.

Proposition 4. *Let $c_1 > 0, c_2 \geq 1, c_4 > 0, c_5 > 0, C_1 \geq 1$ be the absolute constants defined in the statements of the previous lemmas, and let $\delta \in (0, 1/e]$ be arbitrary. Suppose we choose u, r, η and T such that*

$$\begin{aligned} u &\leq \frac{\sqrt{\epsilon}}{d\sqrt{\rho} \log(T/\delta)} \cdot \min \left\{ \frac{1}{64c_5^2 c_2}, \frac{1}{2048c_1 c_2} \right\}^{1/4}, \quad r \leq \epsilon \cdot \min \left\{ \frac{1}{8c_5 \sqrt{2} c_2}, \frac{1}{32\sqrt{c_1}} \right\}, \\ \eta &\leq \frac{1}{L t_f(\delta)} \min \left\{ \frac{1}{\log(T/\delta)}, \frac{\sqrt{m}}{8c_4 (\ln(C_1 d m T/\delta))^{3/2} \sqrt{d}}, \frac{m}{128c_1 (\ln(C_1 d m T/\delta))^3 d} \right\}, \\ T &\geq \max \left\{ \frac{256 t_f(\delta) ((f(x_0) - f^*) + \epsilon^2/L)}{\eta \epsilon^2}, 4 \right\}. \end{aligned}$$

Then, with probability at least $1 - 6\delta$, there are at most $T/4$ iterations for which $\|\nabla f(x_t)\| \geq \epsilon$.

Proof. Without loss of generality, we assume that T is a multiple of $t_f(\delta)$, and we similarly split $\{0, 1, \dots, T\}$ into $K = \lfloor T/t_f(\delta) \rfloor$ intervals J_0, \dots, J_{K-1} . Let I_1 denote the set of indices k such that for every $t \in J_k$,

$$\|\nabla f(x_t)\| > 8t_f(\delta)\eta L \left[\left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| \right) + \|Y_t\| \right]. \quad (35)$$

We let I_1^c denote the complement of I_1 in $\{0, 1, \dots, K-1\}$. We denote

$$\mathcal{E}_T(\delta) := \mathcal{H}_T(\delta) \cap \left(\bigcap_{t=0}^{T-1} \mathcal{A}_t(\delta) \right) \cap \left(\bigcap_{t=0}^{T-1} \mathcal{G}_t(\delta) \right) \cap \left(\bigcap_{k=0}^{K-1} \mathcal{B}_{kt_f(\delta)}(\delta; t_f(\delta)) \right).$$

In the remaining part of the proof, unless otherwise stated, we shall always assume that we are working on the event $\mathcal{E}_T(\delta)$.

By Lemma 18 with $\tau = T$ and our choices of η and δ in the statement of the lemma, we have

$$f(x_T) - f(x_0) \leq - \sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 + T \frac{c_5^2}{64} t_f(\delta)^2 \eta^3 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2$$

$$\begin{aligned}
& + T\eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + TL\eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
& + \eta c_1 r^2 (128t_f(\delta) + \eta L) \log \frac{T}{\delta} + Tc_1 L\eta^2 r^2.
\end{aligned} \tag{36}$$

Suppose that there are at least $T/4$ iterations where $\|\nabla f(x_t)\| \geq \epsilon$. Let I_ϵ denote the set of indices k for which there exists some $t \in J_k$ with $\|\nabla f(x_t)\| \geq \epsilon$. Then, by the pigeonhole principle, the set I_ϵ has at least $\lceil T/(4t_f(\delta)) \rceil$ members. Note that, by our choices of the parameters η, u, r , it can be shown that

$$c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) < \epsilon, \tag{37}$$

while by Lemma 17, if k is in I_1^c , we have

$$\|\nabla f(x_t)\| \leq c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \log(T/\delta) + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right), \quad \forall t \in J_k.$$

This implies that $I_\epsilon \subseteq I_1$.

Observe that by Lemma 16, for any $k \in I_1$, we have

$$\frac{1}{2} \|\nabla f(x_{kt_f(\delta)})\| \leq \|\nabla f(x_t)\| \leq 2 \|\nabla f(x_{kt_f(\delta)})\|, \quad \forall t \in J_k.$$

This implies in particular that for any $k \in I_\epsilon$, we have $\min_{t \in J_k} \|\nabla f(x_t)\| \geq \frac{1}{16} \epsilon^2$, and consequently

$$- \sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 \leq - \sum_{k \in I_\epsilon} \frac{\eta}{2} \cdot \frac{\epsilon^2}{16} \leq - \frac{T}{4t_f(\delta)} \cdot \frac{\eta}{2} \cdot \frac{\epsilon^2}{16} = - \frac{T\eta\epsilon^2}{128t_f(\delta)}.$$

Hence, by Eq. (36),

$$\begin{aligned}
f(x_T) - f(x_0) & \leq - \frac{T\eta\epsilon^2}{128t_f(\delta)} + T \frac{c_5^2}{64} t_f(\delta)^2 \eta^3 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\
& \quad + T\eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + T\eta \cdot (\eta L) u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
& \quad + \eta c_1 r^2 (128t_f(\delta) + \eta L) \log \frac{T}{\delta} + T\eta \cdot c_1 \eta L r^2.
\end{aligned} \tag{38}$$

Now, by our choices of u, r and η , we have

$$\begin{aligned}
& T \frac{c_5^2}{64} t_f(\delta)^2 \eta^3 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\
& \leq T\eta \cdot \frac{c_5^2}{32} t_f(\delta)^2 (\eta L)^2 \left(u^4 d^4 \rho^2 \left(\log \frac{T}{\delta} \right)^4 + 2 \log(T/\delta) r^2 \right) \\
& \leq T\eta \cdot \left(\frac{\epsilon^2}{2048c_2 \left(\log \frac{T}{\delta} \right)^2} + \frac{\epsilon^2}{2048c_2 \log(T/\delta)} \right) \leq \frac{T\eta\epsilon^2}{512t_f(\delta)},
\end{aligned}$$

where we used $\log(T/\delta) \geq 1$ and $2c_2 \log(T/\delta) \geq t_f(\delta)$. We also have

$$\begin{aligned}
& T\eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + T\eta \cdot (\eta L) u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 + Tc_1 L\eta^2 r^2 \\
& \leq T\eta \cdot \frac{\epsilon^2}{2048c_2 d \log(T/\delta)} + T\eta \cdot \frac{\epsilon^2}{2048c_2 t_f(\delta) \log(T/\delta)} + T\eta \cdot \frac{\epsilon^2}{1024t_f(\delta) \log(T/\delta)} \\
& \leq \frac{T\eta\epsilon^2}{512t_f(\delta)},
\end{aligned}$$

where we used $c_2 d \log(T/\delta) \geq t_f(\delta)$, $c_2 \geq 1$ and $\log(T/\delta) \geq 1$. Finally,

$$\eta c_1 r^2 (128t_f(\delta) + \eta L) \log \frac{T}{\delta} \leq \frac{(128t_f(\delta) + 1)\epsilon^2}{1024Lt_f(\delta)} < \frac{\epsilon^2}{L}.$$

By plugging these bounds into Eq. (38), we get

$$f(x_T) - f(x_0) < -\frac{T\eta\epsilon^2}{128t_f(\delta)} + \frac{T\eta\epsilon^2}{512t_f(\delta)} + \frac{T\eta\epsilon^2}{512t_f(\delta)} + \frac{\epsilon^2}{L} \leq -\frac{T\eta\epsilon^2}{256t_f(\delta)} + \frac{\epsilon^2}{L}.$$

Therefore, as long as

$$T \geq \frac{256t_f(\delta) ((f(x_0) - f^*) + \epsilon^2/L)}{\eta\epsilon^2},$$

we will get $f(x_T) < f^*$, which is a contradiction. Thus, we can conclude that on the event $\mathcal{E}_T(\delta)$, there are at most $T/4$ iterations for which $\|\nabla f(x_t)\| \geq \epsilon$.

We can now complete our proof by using the union bound (suppressing the dependence of some of the events on δ for notational simplicity) to derive

$$\begin{aligned} \mathbb{P}(\mathcal{E}_T^c) &\leq \mathbb{P}(\mathcal{H}_T^c) + \sum_{t=0}^{T-1} \mathbb{P}(\mathcal{A}_t^c) + \sum_{t=0}^{T-1} \mathbb{P}(\mathcal{G}_t^c) + \sum_{k=0}^{K-1} \mathbb{P}(\mathcal{B}_{kt_f(\delta)}^c(\delta; t_f(\delta))) \\ &\leq \frac{(T+4)\delta}{T} + \delta + 2\delta + \frac{K\delta}{T} \leq 6\delta. \quad \square \end{aligned}$$

E ESCAPING SADDLE POINT

In this section, we first show that the travelling distance of the iterates can be bounded in terms of the function value improvement (Appendix E.2). Utilizing this result, as well as Proposition 2 in Appendix C.3 which provides a concentration bound on the the zeroth-order noise, we then prove that sufficient function value decrease can be made near a saddle point in Appendix E.3.

E.1 KEY QUANTITIES AND NOTATION

We will use γ to denote $-\lambda_{\min}(\nabla^2 f(x_0))$, where we know that $\gamma \geq \sqrt{\rho}\epsilon$.

E.2 IMPROVE OR LOCALIZE

In this subsection, we aim to bound the movement of the iterates across a number of steps in terms of the function value improvement made during these number of steps.

We first state a simple result separating the norm of the difference between $x_{t_0+\tau}$ and x_{t_0} into a few different terms.

Lemma 19. *Consider the perturbed zeroth-order update Algorithm 1. Then, for any $t_0 \in \mathbb{N}$ and $\tau \in \mathbb{N}$,*

$$\|x_{t_0+\tau} - x_{t_0}\|^2 \leq V_1(t_0, \tau) + V_2(t_0, \tau) + V_3(t_0, \tau) + V_4(t_0, \tau), \quad (39)$$

where

$$\begin{aligned} V_1(t_0, \tau) &:= 8\eta^2 \tau \sum_{t=t_0}^{t_0+\tau-1} \|\nabla f(x_t)\|^2, & V_2(t_0, \tau) &:= 8\eta^2 \left\| \sum_{t=t_0}^{t_0+\tau-1} \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\|^2 \\ V_3(t_0, \tau) &:= 4\eta^2 \left\| \sum_{t=t_0}^{t_0+\tau-1} Y_t \right\|^2, & V_4(t_0, \tau) &:= 4\eta^2 \left\| \sum_{t=t_0}^{t_0+\tau-1} \frac{1}{m} \sum_{i=1}^m u Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2. \end{aligned} \quad (40)$$

Proof. For notational convenience, let $t_0 := 0$. Then, applying the form of the perturbed zeroth-order update in Algorithm 1, we get

$$\|x_\tau - x_0\|^2$$

$$\begin{aligned}
&= \left\| \sum_{t=0}^{\tau-1} x_{t+1} - x_t \right\|^2 \\
&= \eta^2 \left\| \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) + \frac{1}{m} \sum_{i=1}^m u Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} + Y_t \right\|^2 \\
&\leq 4\eta^2 \left\| \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \nabla f(x_t) \right\|^2 + 4\eta^2 \left\| \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m u Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2 + 4\eta^2 \left\| \sum_{t=0}^{\tau-1} Y_t \right\|^2 \\
&\leq 4\eta^2 \left\| \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) + \sum_{t=0}^{\tau-1} \nabla f(x_t) \right\|^2 + 4\eta^2 \left\| \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m u Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2 + 4\eta^2 \left\| \sum_{t=0}^{\tau-1} Y_t \right\|^2 \\
&\leq 8\eta^2 \tau \underbrace{\sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2}_{V_1(0,\tau)} + 8\eta^2 \underbrace{\left\| \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\|^2}_{V_2(0,\tau)} + 4\eta^2 \underbrace{\left\| \sum_{t=0}^{\tau-1} Y_t \right\|^2}_{V_3(0,\tau)} + 4\eta^2 \underbrace{\left\| \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m u Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2}_{V_4(0,\tau)}.
\end{aligned}$$

□

We now proceed to bound the terms $V_1(t_0, \tau)$, $V_2(t_0, \tau)$, $V_3(t_0, \tau)$ and $V_4(t_0, \tau)$.

First, we have the following result bounding $V_1(t_0, \tau)$.

Lemma 20. *Let $c_1 > 0, c_2 \geq 1, c_4 > 0, c_5 > 0, C_1 \geq 1$ be the absolute constants defined in the statements of the previous lemmas, and let $\delta \in (0, 1/e]$ be arbitrary.*

Suppose we choose η such that

$$\eta \leq \frac{1}{Lt_f(\delta)} \cdot \min \left\{ \frac{\sqrt{m}}{8c_4(\ln(C_1 dmT/\delta))^{3/2} \sqrt{d}}, \frac{m}{128c_1(\ln(C_1 dmT/\delta))^3 d} \right\}.$$

There are two cases to consider:

1. The first is when $\tau \geq t_f(\delta)$. In this case, split $\{t_0, t_0+1, \dots, t_0+\tau-1\}$ into $K := \lfloor \tau/t_f(\delta) \rfloor$ intervals:

$$\begin{aligned}
J_k &= \{t_0 + kt_f(\delta), \dots, t_0 + (k+1)t_f(\delta) - 1\}, \quad 0 \leq k < K-1, \\
J_{K-1} &= \{t_0 + (K-1)t_f(\delta), \dots, t_0 + \tau - 1\}.
\end{aligned}$$

Then, on the event

$$\mathcal{E}_{t_0, \tau}(\delta) := \mathcal{H}_{t_0, \tau}(\delta) \cap \left(\bigcap_{t=t_0}^{t_0+\tau-1} \mathcal{A}_t(\delta) \right) \cap \left(\bigcap_{t=t_0}^{t_0+\tau-1} \mathcal{G}_t(\delta) \right) \cap \left(\bigcap_{k=0}^{K-2} \mathcal{B}_{t_0+kt_f(\delta)}(\delta; t_f(\delta)) \right) \cap \mathcal{B}_{t_0+(K-1)t_f(\delta)}(\delta; \tau - (K-1)t_f(\delta)),$$

we have that

$$\begin{aligned}
V_1(t_0, \tau) &= 8\eta^2 \tau \sum_{t=t_0}^{t_0+\tau-1} \|\nabla f(x_t)\|^2 \\
&\leq 64\eta\tau t_f(\delta) ((f(x_0) - f(x_\tau)) + N_{u,r}(\tau; \delta)),
\end{aligned}$$

where

$$\begin{aligned}
N_{u,r}(\tau; \delta) &:= \tau \frac{c_5^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\
&\quad + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
&\quad + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2 \\
&\quad + c_5^2 t_f^3(\delta) \eta^3 L^2 \left(u^2 d^2 \rho \log(T/\delta) + \sqrt{2 \log(T/\delta)} r \right)^2. \tag{41}
\end{aligned}$$

2. The second is when $\tau < t_f(\delta)$. Suppose we choose u and r such that

$$u \leq \frac{\sqrt{\epsilon}}{d\sqrt{\rho}\log(T/\delta)} \cdot \min \left\{ \frac{1}{64c_5^2c_2}, \frac{1}{2048c_1c_2} \right\}^{1/4}, \quad r \leq \epsilon \cdot \min \left\{ \frac{1}{8c_5\sqrt{2c_2}}, \frac{1}{32\sqrt{c_1}} \right\}.$$

Suppose the event $\cap_{t=t_0}^{t_0+\tau-1} (\mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta))$ holds. Suppose also that $\|\nabla f(x_{t_0})\| \leq \epsilon$. Then,

$$V_1(t_0, \tau) \leq 32\eta^2\tau^2\epsilon^2 \leq 32\eta^2(t_f(\delta))^2\epsilon^2$$

Proof. 1. We first consider the case where $\tau \geq t_f(\delta)$. Let I_1 denote the set of indices k such that for every time-step t in the interval J_k , the gradient dominates the noise terms as

$$\|\nabla f(x_t)\| > 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right). \quad (42)$$

WLOG, we may assume that $t_0 := 0$, and denote $V_1(\tau) := V_1(0, \tau)$. WLOG, we also assume that τ is a multiple of $t_f(\delta)$. From Lemma 18, on the event that $\mathcal{E}_\tau(\delta)$ holds and by our choice of η , we have

$$\begin{aligned} f(x_\tau) - f(x_0) &\leq - \sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 + \tau \frac{c_5^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\ &\quad + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\ &\quad + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2. \end{aligned}$$

By Lemma 16 (and our choice of η), it follows that for any $k \in I_1$, on the event $\cap_{t \in J_k} \mathcal{A}_t(\delta)$, we have

$$\sum_{t \in J_k} \|\nabla f(x_t)\|^2 \leq 4t_f \min_{t \in J_k} \|\nabla f(x_t)\|^2.$$

Thus, on the event that $\mathcal{E}_\tau(\delta)$ holds, for our choice of η , we have

$$\begin{aligned} \eta \sum_{k \in I_1} \sum_{t \in J_k} \|\nabla f(x_t)\|^2 &\leq 4t_f(\delta)\eta \sum_{k \in I_1} \min_{t \in J_k} \|\nabla f(x_t)\|^2 \\ &\leq 8t_f(\delta) \sum_{k \in I_1} \frac{\eta}{2} \min_{t \in J_k} \|\nabla f(x_t)\|^2 \\ &\leq 8t_f(\delta) \left((f(x_0) - f(x_\tau)) + \tau \frac{c_5^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \right) \\ &\quad + 8t_f(\delta) \left(\tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \right) \\ &\quad + 8t_f(\delta) \left(\eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2 \right). \end{aligned}$$

Similarly, for any $k \in I_1^c$ (where I_1^c denotes the complement of I_1 in $\{0, 1, \dots, K-1\}$, i.e. intervals where the gradient is smaller than than the perturbation terms in some iteration), on the event $(\cap_{t \in J_k} \mathcal{A}_t(\delta)) \cap (\cap_{t \in J_k} \mathcal{G}_t(\delta))$, by Lemma 17 (and our choice of η), we have

$$\|\nabla f(x_t)\| \leq c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \log(T/\delta) + \sqrt{2 \log(T/\delta)} r \right), \quad \forall t \in J_k.$$

On the event that $\mathcal{E}_\tau(\delta)$ holds, this gives us then

$$\eta \sum_{k \in I_1^c} \sum_{t \in J_k} \|\nabla f(x_t)\|^2 \leq \eta \tau \left(c_5^2 t_f^2(\delta) \eta^2 L^2 \left(u^2 d^2 \rho \log(T/\delta) + \sqrt{2 \log(T/\delta)} r \right)^2 \right).$$

Hence, on the event that $\mathcal{E}_\tau(\delta)$ holds, we have that

$$\begin{aligned} \eta \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 &= \eta \sum_{k \in I_1} \sum_{t \in J_k} \|\nabla f(x_t)\|^2 + \eta \sum_{k \in I_1^c} \sum_{t \in J_k} \|\nabla f(x_t)\|^2 \\ &\leq 8t_f(\delta) \left((f(x_0) - f(x_\tau)) + \tau \frac{c_5^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \right) \\ &\quad + 8t_f(\delta) \left(\tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \right) \\ &\quad + 8t_f(\delta) \left(\eta c_1 r^2 (128t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2 \right) \\ &\quad + 8t_f(\delta) \eta \tau \left(c_5^2 t_f^2(\delta) \eta^2 L^2 \left(u^2 d^2 \rho \log(T/\delta) + \sqrt{2 \log(T/\delta)} r \right)^2 \right). \end{aligned}$$

This yields the final result for the case $\tau \geq t_f(\delta)$.

2. We next consider the case where $1 \leq \tau < t_f(\delta)$. Recall the notation that

$$\mathcal{E}(t_0, t_0 + \tau, \delta) := \cap_{t=t_0}^{t_0+\tau-1} \left\{ \|\nabla f(x_t)\| > 8t_f(\delta) \eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right) \right\}$$

There are two cases to consider.

(a) On the event $\mathcal{E}(t_0, t_0 + \tau, \delta) \cap (\cap_{t=t_0}^{t_0+\tau-1} \mathcal{A}_t(\delta))$, we have by Lemma 16 that $\|\nabla f(x_t)\| \leq 2\|\nabla f(x_0)\|$ for each $t \in \{0, 1, \dots, \tau-1\}$. Then,

$$V_1(t_0, \tau) = 8\eta^2 \tau \sum_{t=t_0}^{t_0+\tau-1} \|\nabla f(x_t)\|^2 \leq 8\eta^2 \tau^2 \left(4\|\nabla f(x_0)\|^2 \right) \leq 32\eta^2 \tau^2 \epsilon^2,$$

where the final inequality uses the assumption that $\|\nabla f(x_0)\| \leq \epsilon$.

(b) Suppose the event $\mathcal{E}^c(t_0, t_0 + \tau, \delta) \cap (\cap_{t=t_0}^{t_0+\tau-1} \mathcal{A}_t(\delta)) \cap (\cap_{t=t_0}^{t_0+\tau-1} \mathcal{G}_t(\delta))$ holds. In this case, by Lemma 17, we have that for each $t \in \{t_0, t_0 + 1, \dots, t_0 + \tau - 1\}$

$$\begin{aligned} \|\nabla f(x_t)\| &\leq c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \\ &\leq \epsilon, \end{aligned}$$

where the final inequality follows by our choice of η, u and r (cf. Eq. (37)). Hence,

$$\begin{aligned} V_1(t_0, \tau) &= 8\eta^2 \tau \sum_{t=t_0}^{t_0+\tau-1} \|\nabla f(x_t)\|^2 \\ &\leq 8\eta^2 \tau^2 \left(c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \right)^2 \\ &\leq 8\eta^2 \tau^2 \epsilon^2 < 32\eta^2 \tau^2 \epsilon^2. \end{aligned}$$

The final result for the case $\tau < t_f(\delta)$ then follows. \square

We proceed to bound $V_2(t_0, \tau)$.

Lemma 21. *Let $c_1 > 0, c_2 \geq 1, c_4 > 0, c_5 > 0, C_1 \geq 1$ be the absolute constants defined in the statements of the previous lemmas, and let $\delta \in (0, 1/e]$ be arbitrary and $\tau > 0$ be arbitrary. Suppose we choose η such that*

$$\eta \leq \frac{1}{L t_f(\delta)} \cdot \min \left\{ \frac{\sqrt{m}}{8c_4 (\ln(C_1 d m T / \delta))^{3/2} \sqrt{d}}, \frac{m}{128c_1 (\ln(C_1 d m T / \delta))^3 d} \right\}.$$

Let T_s denote an integer such that $T_s \geq \max\{\tau, t_f(\delta)\}$, and for any $F > 0$, define

$$B(\delta; F) := \frac{8t_f(\delta)(F + N_{u,r}(T_s, \delta))}{\eta} \left(T_s + \frac{d}{m}\right) (\ln(CT^2/\delta))^2, \quad b_\tau(\delta; F) := \frac{t_f(\delta)\tau F}{\eta}.$$

Let $c', C > 0$ denote the same constants as in the statement of Proposition 2. Denote the event that

$$\text{either } \sum_{t=t_0}^{t_0+\tau-1} \frac{d}{m} (\ln(CT^2/\delta))^2 \|\nabla f(x_t)\|^2 \geq B(\delta; F)$$

$$\text{or } \sqrt{\frac{V_2(t_0, \tau)}{8\eta^2}} \leq c' \sqrt{\max\left\{\sum_{t=t_0}^{t_0+\tau-1} \frac{d}{m} (\ln(CT^2/\delta))^2 \|\nabla f(x_t)\|^2, b_\tau(\delta; F)\right\}} \left(\log\left(\frac{CT^2}{\delta}\right) + \log\left(\log\left(\frac{B(\delta; F)}{b_\tau(\delta; F)}\right) + 1\right)\right)$$

holds as $\mathcal{L}_{t_0, \tau}(\delta; F)$ ⁶. We show that $\mathbb{P}(\mathcal{L}_{t_0, \tau}(\delta; F)) \geq 1 - \frac{\delta}{T}$. Finally, denote the event $\mathcal{M}_{t_0, T_s}(F)$ as the event that $f(x_{t_0}) - f(x_{t_0+T_s}) < F$.

Then, on the event $\mathcal{L}_{t_0, \tau}(\delta) \cap \mathcal{E}_{t_0, T_s}(\delta) \cap \mathcal{M}_{t_0, T_s}(F)$ (where $\mathcal{E}_{0, T_s}(\delta)$ is as defined in Lemma 20),

$$V_2(t_0, \tau) \leq 8c'^2 \beta_1(\delta; F) \eta t_f(\delta) \max\left\{\frac{8d}{m} (\ln(CT^2/\delta))^2 (F + N_{u,r}(T_s, \delta)), \tau F\right\}, \quad (43)$$

where

$$\beta_1(\delta; F) := \log\left(\frac{CT^2}{\delta}\right) + \log\left(\log\left(\frac{B(\delta; F)}{b_1(\delta; F)}\right) + 1\right).$$

Proof. We note that $\mathbb{P}(\mathcal{L}_{t_0, \tau}(\delta; F)) \geq 1 - \frac{\delta}{T}$ is a direct consequence of Proposition 2. In the rest of the proof, without loss of generality, we assume that $t_0 = 0$ for notational simplicity. On the event $\mathcal{L}_{0, \tau}(\delta; F) \cap \mathcal{E}_{0, T_s}(\delta) \cap \mathcal{M}_{0, T_s}(F)$, suppose that

$$\begin{aligned} & \sum_{t=0}^{\tau-1} \frac{d}{m} (\ln(CT^2/\delta))^2 \|\nabla f(x_t)\|^2 \geq B(\delta; F) = \frac{8t_f(\delta)(F + N_{u,r}(T_s, \delta))}{\eta} \left(T_s + \frac{d}{m}\right) (\ln(CT^2/\delta))^2 \\ \implies & \eta \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 \geq 8t_f(\delta)(F + N_{u,r}(T_s, \delta)) \\ \implies & \eta \sum_{t=0}^{T_s-1} \|\nabla f(x_t)\|^2 \geq 8t_f(\delta)(F + N_{u,r}(T_s, \delta)) \\ \implies & 8\eta^2 T_s \sum_{t=0}^{T_s-1} \|\nabla f(x_t)\|^2 \geq 64\eta T_s t_f(\delta)(F + N_{u,r}(T_s, \delta)) \\ \implies & 8\eta^2 T_s \sum_{t=0}^{T_s-1} \|\nabla f(x_t)\|^2 \geq 64\eta T_s t_f(\delta)(f(x_0) - f(x_{T_s}) + N_{u,r}(T_s, \delta)), \quad \text{since } f(x_0) - f(x_{T_s}) \leq F \\ \iff & V_1(0, T_s) \geq 64\eta T_s t_f(\delta)(f(x_0) - f(x_{T_s}) + N_{u,r}(T_s, \delta)), \end{aligned}$$

where we note the last equation contradicts Lemma 20. For notational simplicity, denote

$$\beta_\tau(\delta; F) := \log\left(\frac{CT^2}{\delta}\right) + \log\left(\log\left(\frac{B(\delta; F)}{b_\tau(\delta; F)}\right) + 1\right).$$

Observe that β_1 is larger than β_τ for every $\tau \geq 1$. Since $\mathcal{L}_{t_0, \tau}(\delta; F)$ holds, we must have then that

$$\sqrt{\frac{V_2(0, \tau)}{8\eta^2}} \leq c' \sqrt{\max\left\{\sum_{t=0}^{\tau-1} \frac{d}{m} (\ln(CT^2/\delta))^2 \|\nabla f(x_t)\|^2, b_\tau(\delta; F)\right\}} \beta_1(\delta; F).$$

⁶We note that by construction, $B(\delta; F) \geq b_\tau(\delta; F)$

$$\begin{aligned}
& \text{Now, continuing, recalling the definition of } V_1(0, T_s) = 8\eta^2 T_s \sum_{t=0}^{T_s-1} \|\nabla f(x_t)\|^2 \\
V_2(0, \tau) & \leq c'^2 \beta_1(\delta; F) \max \left\{ 8\eta^2 \sum_{t=0}^{\tau-1} \frac{d}{m} (\ln(CT^2/\delta))^2 \|\nabla f(x_t)\|^2, 8\eta^2 b_\tau(\delta; F) \right\} \\
& \leq c'^2 \beta_1(\delta; F) \max \left\{ 8\eta^2 \sum_{t=0}^{T_s-1} \frac{d}{m} (\ln(CT^2/\delta))^2 \|\nabla f(x_t)\|^2, 8\eta^2 b_\tau(\delta; F) \right\} \\
& \leq c'^2 \beta_1(\delta; F) \max \left\{ \frac{d}{m} (\ln(CT^2/\delta))^2 \frac{V_1(0, T_s)}{T_s}, 8\eta t_f(\delta) \tau F \right\} \\
& \stackrel{(i)}{\leq} c'^2 \beta_1(\delta; F) \max \left\{ \frac{d}{m} (\ln(CT^2/\delta))^2 (64\eta t_f(\delta) (f(x_0) - f(x_{T_s}) + N_{u,r}(T_s, \delta))), 8\eta t_f(\delta) \tau F \right\} \\
& \stackrel{(ii)}{\leq} c'^2 \beta_1(\delta; F) \max \left\{ \frac{d}{m} (\ln(CT^2/\delta))^2 (64\eta t_f(\delta) (F + N_{u,r}(T_s, \delta))), 8\eta t_f(\delta) \tau F \right\} \\
& = c'^2 \beta_1(\delta; F) (8\eta t_f(\delta)) \max \left\{ \frac{d}{m} (\ln(CT^2/\delta))^2 (8(F + N_{u,r}(T_s, \delta))), \tau F \right\}.
\end{aligned}$$

We note that (i) is a consequence of Lemma 20, while (ii) comes from our assumption that the event $\mathcal{M}_{t_0, T_s}(F)$ holds, i.e. $f(x_{t_0}) - f(x_{t_0+T_s}) \leq F$.

□

We next bound $V_3(t_0, \tau)$ and $V_4(t_0, \tau)$.

Lemma 22. *Let $c > 0$ denote the same constant in Lemma 7. Consider any arbitrary $0 < \delta \leq 1/e$, and let $\tau \geq t_f(\delta)$ be arbitrary. Let $\mathcal{N}_{t_0, \tau}(\delta)$ denote the event that*

$$V_3(t_0, \tau) := 4\eta^2 \left\| \sum_{t=t_0}^{t_0+\tau-1} Y_t \right\|^2 \leq 4c_6 \eta^2 \tau \log(2dT/\delta) r^2,$$

where $c_6 > 0$ is an absolute constant. Then, by Lemma 7, $\mathbb{P}(\mathcal{N}_{t_0, \tau}(\delta)) \geq 1 - \frac{\delta}{T}$. Denote the event

$$\mathcal{O}_t(\delta) := \left\{ \frac{1}{m} \sum_{i=1}^m \|Z_{t,i}\|^8 \leq c_7 d^4 \left(\log \left(\frac{T}{\delta} \right) \right)^4 \right\},$$

where $c_7 > 0$ is an absolute constant. Then, on the event $\cap_{t=t_0}^{t_0+\tau-1} \mathcal{O}_t(\delta)$, we have

$$V_4(t_0, \tau) \leq 4c_7 \eta^2 \tau^2 \rho^2 u^4 d^4 \left(\log \left(\frac{T}{\delta} \right) \right)^4.$$

Moreover, for each t , $\mathbb{P}(\mathcal{O}_t(\delta)) \geq 1 - \frac{\delta}{T}$.

Proof. The proof for $V_3(t_0, \tau)$ follows directly from Lemma 7, by picking c_6 to be the c that appears in the statement of Lemma 7. Meanwhile, observe that

$$\begin{aligned}
V_4(t_0, \tau) & = 4\eta^2 \left\| \sum_{t=t_0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m u_{Z_{t,i}} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2 \\
& \leq 4\eta^2 \tau \left(\sum_{t=t_0}^{t_0+\tau-1} \left\| \frac{1}{m} \sum_{i=1}^m u_{Z_{t,i}} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2 \right) \\
& \stackrel{(iii)}{\leq} 4\eta^2 \tau \sum_{t=t_0}^{t_0+\tau-1} \frac{1}{m} \sum_{i=1}^m \rho^2 u^4 \|Z_{t,i}\|^8 \\
& \leq 4c_7 \eta^2 \tau^2 \rho^2 u^4 d^4 \left(\log \left(\frac{T}{\delta} \right) \right)^4.
\end{aligned}$$

Above, to derive (iii), we used the bound that $\|\tilde{H}_{t,i}\| \leq \rho u \|Z_{t,i}\|$. The final inequality is a consequence of our assumption that $\cap_{t=t_0}^{t_0+\tau-1} \mathcal{O}_t(\delta)$ holds. Finally, the result that $\mathbb{P}(\mathcal{O}_t(\delta)) \geq 1 - \frac{\delta}{T}$ holds due to Lemma 11, where we note that we may pick the absolute constant c_7 to be equal to $2C c^4$, where $c, C > 0$ are the absolute constants that appear in the statement of Lemma 11. \square

Finally, combining the earlier results, we have the following technical result, which bounds the travelling distance of the iterates in terms of the decrease in function value decrease.

Lemma 23 (Improve or Localize). *Consider the perturbed zeroth-order update Algorithm 1. Let $c' > 0, c_1 > 0, c_2 \geq 1, c_4 > 0, c_5 > 0, c_6 > 0, c_7 > 0, C_1 \geq 1$ be the absolute constants defined in the statements of the previous lemmas, and let $\delta \in (0, 1/e]$ be arbitrary. Consider any $T_s \geq t_f(\delta)$. For any $F > 0$, suppose $f(x_{T_s}) - f(x_0) > -F$, i.e. $f(x_0) - f(x_{T_s}) < F$. Suppose that the event*

$$\mathcal{P}_{t_0, T_s}(\delta, F) := \cap_{\tau=1}^{T_s} (\mathcal{L}_{t_0, \tau}(\delta; F) \cap \mathcal{N}_{t_0, \tau}(\delta)) \cap \left(\cap_{t=t_0}^{t_0+T_s-1} \mathcal{O}_t(\delta) \cap \mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta) \right) \cap \left(\cap_{\tau=t_f(\delta)}^{T_s-1} \mathcal{E}_{t_0, \tau}(\delta) \right)$$

holds, where the events $\mathcal{E}_{t_0, \tau}(\delta), \mathcal{L}_{t_0, \tau}(\delta), \mathcal{N}_{t_0, \tau}(\delta), \mathcal{O}_t(\delta)$ are as defined in Lemma 20, Lemma 21 and Lemma 22, and $\mathcal{G}_t(\delta)$ and $\mathcal{A}_t(\delta)$ are as defined in Lemma 13 and Lemma 15. Suppose we choose u, r and η such that

$$u \leq \frac{\sqrt{\epsilon}}{d\sqrt{\rho} \log(T/\delta)} \cdot \min \left\{ \frac{1}{64c_5^2 c_2}, \frac{1}{2048c_1 c_2} \right\}^{1/4}, \quad r \leq \epsilon \cdot \min \left\{ \frac{1}{8c_5 \sqrt{2} c_2}, \frac{1}{32\sqrt{c_1}} \right\},$$

$$\eta \leq \frac{1}{L t_f(\delta)} \min \left\{ \frac{1}{\log(T/\delta)}, \frac{\sqrt{m}}{8c_4 (\ln(C_1 d m T/\delta))^{3/2} \sqrt{d}}, \frac{m}{128c_1 (\ln(C_1 d m T/\delta))^3 d} \right\}.$$

Suppose $\eta \leq \min \left\{ 1, \frac{1}{t_f(\delta)}, \frac{1}{t_f \delta L} \right\}$. Suppose also we pick u and r small enough such that

$$u \leq \frac{r^{1/2}}{d \log(T/\delta) \rho^{1/2}}, \quad r^2 \leq \min \left\{ \frac{F}{\eta T_s \log(T/\delta) \left(\frac{65c_5^2}{8} + 132c_1 + 1 \right)}, \frac{F}{4c_6 \log(2dT/\delta) + 4c_7 \eta T_s} \right\}.$$

Then, for each $\tau \in \{0, 1, \dots, T_s\}$, we have that

$$\|x_{t_0+\tau} - x_{t_0}\|^2 \leq \phi_{T_s}(\delta, F),$$

where

$$\phi_{T_s}(\delta, F) \leq \max \{ 128\eta T_s t_f(\delta) F, 32\eta^2 (t_f(\delta))^2 \epsilon^2 \} + 8c'^2 \beta_1(\delta; F) \eta t_f(\delta) \max \left\{ \frac{16d}{m} (\ln(CT^2/\delta))^2 F, T_s F \right\} + T_s \eta t_f(\delta) F,$$

where $\beta_1(\delta; F)$ is defined as in Lemma 21. Moreover, $\mathbb{P}(\mathcal{P}_{t_0, T_s}(\delta, F)) \geq 1 - \frac{12T_s \delta}{T}$.

Proof. We recall that

$$\begin{aligned} \|x_{t_0+\tau} - x_{t_0}\|^2 &\leq \underbrace{8\eta^2 \tau \sum_{t=t_0}^{t_0+\tau-1} \|\nabla f(x_t)\|^2}_{V_1(t_0, \tau)} + \underbrace{8\eta^2 \left\| \sum_{t=t_0}^{t_0+\tau-1} \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t) \right\|^2}_{V_2(t_0, \tau)} \\ &\quad + \underbrace{4\eta^2 \left\| \sum_{t=t_0}^{t_0+\tau-1} Y_t \right\|^2}_{V_3(t_0, \tau)} + \underbrace{4\eta^2 \left\| \sum_{t=t_0}^{t_0+\tau-1} \frac{1}{m} \sum_{i=1}^m u Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\|^2}_{V_4(t_0, \tau)} \end{aligned}$$

By Lemma 20, Lemma 21, and Lemma 22, which bound $V_1(t_0, \tau), V_2(t_0, \tau)$, and $V_3(t_0, \tau), V_4(t_0, \tau)$ respectively, on the event $\mathcal{P}_{t_0, T_s}(\delta, F)$, we have, for any $0 \leq \tau \leq T_s$,

$$\begin{aligned} \|x_\tau - x_0\|^2 &\leq V_1(0, \tau) + V_2(0, \tau) + V_3(0, \tau) + V_4(0, \tau) \\ &\leq \max \{ 64\eta \tau t_f(\delta) (F + N_{u,r}(\tau; \delta)), 32\eta^2 (t_f(\delta))^2 \epsilon^2 \} \end{aligned}$$

$$\begin{aligned}
& + 8c'^2\beta_1(\delta; F)\eta t_f(\delta) \max \left\{ \frac{8d}{m} (\ln(CT^2/\delta))^2 (F + N_{u,r}(T_s, \delta)), \tau F \right\} \\
& + 4c_6\eta^2\tau \log(2dT/\delta)r^2 + 4c_7\eta^2\tau^2\rho^2u^4d^4 (\log(T/\delta))^4,
\end{aligned}$$

where $N_{u,r}(\tau; \delta)$ is defined as in Lemma 20.

For the simplified bound (which does not contain $N_{u,r}(\tau; \delta)$), it remains for us to show that our choice of u and r ensures that $N_{u,r}(T_s, \delta) \leq F$ and

$$4c_6\eta^2T_s \log(2dT/\delta)r^2 + 4c_7\eta^2T_s^2\rho^2u^4d^4 (\log(T/\delta))^4 \leq \eta T_s t_f(\delta) F.$$

First, our choice of u ensures that

$$u^4d^4\rho^2(\log(T/\delta))^4 \leq r^2.$$

Next, recall that

$$\begin{aligned}
N_{u,r}(\tau; \delta) & := \tau \frac{c_5^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\
& + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
& + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2 \\
& + c_5^2 t_f^3(\delta) \eta^3 L^2 \left(u^2 d^2 \rho \log(T/\delta) + \sqrt{2 \log(T/\delta)} r \right)^2.
\end{aligned}$$

Recalling our choice of η such that

$$\eta \leq \min \left\{ 1, \frac{1}{t_f(\delta)}, \frac{1}{t_f(\delta)L} \right\},$$

it follows that

$$\begin{aligned}
N_{u,r}(T_s; \delta) & \leq \eta T_s r^2 \left(\frac{8c_5^2}{64} \log(T/\delta) + 2c_1 + 2c_1 + (128c_1 + 1) \log(T/\delta) + c_1 + 8c_5^2 \log(T/\delta) \right) \\
& \leq \eta T_s r^2 \log(T/\delta) \left(\frac{65c_5^2}{8} + 132c_1 + 1 \right) \leq F,
\end{aligned}$$

where the last inequality follows choosing r such that $r^2 \leq \frac{F}{\eta T_s \log(T/\delta) \left(\frac{65c_5^2}{8} + 132c_1 + 1 \right)}$. Similarly,

we have

$$\begin{aligned}
& 4c_6\eta^2T_s \log(2dT/\delta)r^2 + 4c_7\eta^2T_s^2\rho^2u^4d^4 (\log(T/\delta))^4 \\
& \leq \eta T_s t_f(\delta) (4c_6\eta \log(2dT/\delta)r^2 + 4c_7\eta T_s \rho^2u^4d^4 (\log(T/\delta))^4) \\
& \leq \eta T_s t_f(\delta) (4c_6\eta \log(2dT/\delta)r^2 + 4c_7\eta T_s r^2)
\end{aligned}$$

By choosing r such that

$$r^2 \leq \frac{F}{4c_6 \log(2dT/\delta) + 4c_7\eta T_s},$$

it follows that

$$4c_6\eta^2T_s \log(2dT/\delta)r^2 + 4c_7\eta^2T_s^2\rho^2u^4d^4 (\log(T/\delta))^4 \leq \eta T_s t_f(\delta) F,$$

as desired.

We next lower bound the probability of

$$\mathcal{P}_{t_0, T_s}(\delta, F) := \cap_{\tau=1}^{T_s} (\mathcal{L}_{t_0, \tau}(\delta; F) \cap \mathcal{N}_{t_0, \tau}(\delta)) \cap \left(\cap_{t=t_0}^{t_0+T_s-1} \mathcal{O}_t(\delta) \cap \mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta) \right) \cap \left(\cap_{\tau=t_f(\delta)}^{T_s} \mathcal{E}_{t_0, \tau}(\delta) \right).$$

Observe that

$$\cap_{\tau=t_f(\delta)}^{T_s} \mathcal{E}_{t_0, \tau}(\delta)$$

$$\begin{aligned}
&= \cap_{\tau=t_f(\delta)}^{T_s} \left(\mathcal{H}_{t_0, \tau}(\delta) \cap \left(\bigcap_{t=t_0}^{t_0+\tau-1} \mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta) \right) \cap \left(\bigcap_{k=0}^{K-2} \mathcal{B}_{t_0+kt_f(\delta)}(\delta; t_f(\delta)) \right) \cap \mathcal{B}_{t_0+(K-1)t_f(\delta)}(\delta; \tau - (K-1)t_f(\delta)) \right) \\
&= \cap_{\tau=t_f(\delta)}^{T_s} \left(\mathcal{H}_{t_0, \tau}(\delta) \cap \left(\bigcap_{k=0}^{K-2} \mathcal{B}_{t_0+kt_f(\delta)}(\delta; t_f(\delta)) \right) \cap \mathcal{B}_{t_0+(K-1)t_f(\delta)}(\delta; \tau - (K-1)t_f(\delta)) \right) \cap \left(\bigcap_{t=t_0}^{T_s-1} \mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta) \right).
\end{aligned}$$

Note this implies that $\cap_{\tau=t_f(\delta)}^{T_s} \mathcal{E}_{t_0, \tau}(\delta) \cap \left(\bigcap_{t=t_0}^{T_s-1} \mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta) \right) = \cap_{\tau=t_f(\delta)}^{T_s} \mathcal{E}_{t_0, \tau}(\delta)$. We note that by Lemma 1,

$$\mathbb{P} \left(\left(\cap_{\tau=t_f(\delta)}^{T_s} \mathcal{H}_{t_0, \tau}(\delta) \right)^c \right) \leq \frac{5T_s\delta}{T}.$$

Meanwhile, we note that

$$\cap_{t=t_0}^{T_s-1} \mathcal{B}_t(\delta; t_f(\delta)) \subseteq \cap_{\tau=t_f(\delta)}^{T_s} \left(\left(\bigcap_{k=0}^{K-2} \mathcal{B}_{t_0+kt_f(\delta)}(\delta; t_f(\delta)) \right) \cap \mathcal{B}_{t_0+(K-1)t_f(\delta)}(\delta; \tau - (K-1)t_f(\delta)) \right).$$

Hence, by Lemma 14, we have that

$$\begin{aligned}
&\mathbb{P} \left(\left(\cap_{\tau=t_f(\delta)}^{T_s} \left(\left(\bigcap_{k=0}^{K-2} \mathcal{B}_{t_0+kt_f(\delta)}(\delta; t_f(\delta)) \right) \cap \mathcal{B}_{t_0+(K-1)t_f(\delta)}(\delta; \tau - (K-1)t_f(\delta)) \right) \right)^c \right) \\
&\leq \mathbb{P} \left(\left(\cap_{t=t_0}^{T_s-1} \mathcal{B}_t(\delta; t_f(\delta)) \right)^c \right) \leq \frac{T_s\delta}{T}.
\end{aligned}$$

Meanwhile, by Lemma 13 and Lemma 15, we may bound

$$\mathbb{P} \left(\left(\bigcap_{t=t_0}^{T_s-1} \mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta) \right)^c \right) \leq \frac{T_s\delta}{T} + \frac{2T_s\delta}{T} = \frac{3T_s\delta}{T}.$$

Hence, it follows that

$$\mathbb{P} \left(\left(\cap_{\tau=t_f(\delta)}^{T_s} \mathcal{E}_{t_0, \tau}(\delta) \cap \left(\bigcap_{t=t_0}^{T_s-1} \mathcal{A}_t(\delta) \cap \mathcal{G}_t(\delta) \right) \right)^c \right) \leq \frac{5T_s\delta}{T} + \frac{T_s\delta}{T} + \frac{3T_s\delta}{T} = \frac{9T_s\delta}{T}.$$

Meanwhile, it follows from our results in the preceding lemmas that

$$\mathbb{P} \left(\left(\cap_{\tau=1}^{T_s} (\mathcal{L}_{t_0, \tau}(\delta; F) \cap \mathcal{N}_{t_0, \tau}(\delta)) \cap \left(\bigcap_{t=t_0}^{T_s-1} \mathcal{O}_t(\delta) \right) \right)^c \right) \leq \frac{3T_s\delta}{T}.$$

Hence, it follows that $\mathbb{P}(\mathcal{P}_{t_0, T_s}(\delta, F)) \geq 1 - \frac{12T_s\delta}{T}$.

□

E.3 PROVING FUNCTION VALUE DECREASE NEAR SADDLE POINT

We next build on the technical result earlier to prove that each time we are near the saddle point, there is a constant probability of making significant function value decrease. We briefly provide a high-level proof outline below. In our proof, we introduce a coupling argument connecting two closely-related sequences both starting from the saddle, differing only in the sign of their perturbative term along the minimum eigendirection of the Hessian at the saddle. Specifically, when function decrease from a saddle is not sufficiently large, due to the earlier technical result, we know that the coupled sequences will remain within a radius ϕ of the original saddle for a large number (which we will denote as T_s) of iterations. We then utilize this fact to show that the difference of the coupled sequence will (with some constant probability) grow exponentially large, eventually moving out of their specified radius ϕ within T_s iterations, leading to a contradiction.

Our first result formally introduces the coupling, setting the stage for the rest of our arguments. For notational convenience, in this section, unless otherwise specified, we will often assume that the initial iterate x_0 is an ϵ -saddle point.

Lemma 3. *Suppose x_0 is an ϵ -approximate saddle point. Without loss of generality, suppose that the minimum eigendirection of $H := \nabla^2 f(x_0)$ is the e_1 direction, and let γ to denote $-\lambda_{\min}(\nabla^2 f(x_0))$ (note $\gamma \geq \sqrt{\rho\epsilon}$). Consider the following coupling mechanism, where we run the zeroth-order gradient dynamics, starting with x_0 , with two isotropic noise sequences, Y_t and Y'_t respectively, where $(Y_t)_1 = -(Y'_t)_1$, and $(Y_t)_j = (Y'_t)_j$ for all other $j \neq 1$. Suppose that the sequence $\{Z_{t,i}\}_{t \in T, i \in [m]}$ is the same for both sequences. Let $\{x_t\}$ denote the sequence with the $\{Y_t\}$ noise sequence, and let the $\{x'_t\}$ denote the sequence with the $\{Y'_t\}$ noise sequence, where*

$$x'_{t+1} = x'_t - \eta \left(\frac{1}{m} \sum_{i=1}^m \left(Z_{t,i} Z_{t,i}^\top \nabla f(x'_t) + \frac{u}{2} Z_{t,i} Z_{t,i}^\top \tilde{H}'_{t,i} Z_{t,i} \right) + Y'_t \right), \quad x'_0 = x_0,$$

and $\tilde{H}'_{t,i} := \frac{H'_{t,i,+} - H'_{t,i,-}}{2}$, with $H'_{t,i,+} = \nabla^2 f(x'_t + \alpha'_{t,i,+} u Z'_i)$ for some $\alpha'_{t,i,+} \in [0, 1]$, and $H'_{t,i,-} = \nabla^2 f(x'_t - \alpha'_{t,i,-} u Z'_i)$ for some $\alpha'_{t,i,-} \in [0, 1]$. Then, for any $t \geq 0$,

$$\begin{aligned} \hat{x}_{t+1} := x_{t+1} - x'_{t+1} &= - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}_{g_0}(\tau)}_{W_{g_0}(t+1)} - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} (\bar{H}_\tau - H) \hat{x}_\tau}_{W_H(t+1)} \\ &\quad - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}_u(\tau)}_{W_u(t+1)} - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{Y}_\tau}_{W_p(t+1)} \end{aligned}$$

where

$$\begin{aligned} \xi_{g_0}(t) &= \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t), \quad \xi'_{g_0}(t) = \frac{1}{m} \sum_{i=1}^m (Z_{t,i} (Z_{t,i})^\top - I) \nabla f(x'_t), \quad \hat{\xi}_{g_0}(t) = \xi_{g_0}(t) - \xi'_{g_0}(t), \\ \xi_u(t) &= \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}_{t,i} Z_{t,i}, \quad \xi'_u(t) = \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}'_{t,i} Z_{t,i}, \quad \hat{\xi}_u(t) = \xi_u(t) - \xi'_u(t), \\ \hat{Y}_t &= Y_t - Y'_t, \quad \bar{H}_t = \int_0^1 \nabla^2 f(ax_t + (1-a)x'_t) da. \end{aligned}$$

Proof. Observe that

$$\begin{aligned} \hat{x}_{t+1} &:= x_{t+1} - x'_{t+1} \\ &= x_t - \eta (\nabla f(x_t) + \xi_{g_0}(t) + \xi_u(t) Y_t) - [x'_t - \eta (\nabla f(x'_t) + \xi'_{g_0}(t) + \xi'_u(t) + Y'_t)] \\ &= \hat{x}_t - \eta [(\nabla f(x_t) - \nabla f(x'_t)) + (\xi_{g_0}(t) - \xi'_{g_0}(t)) + (\xi_u(t) - \xi'_u(t)) + (Y_t - Y'_t)] \\ &= \hat{x}_t - \eta H \hat{x}_t - \eta (\bar{H}_t - H) \hat{x}_t - \eta \hat{\xi}_{g_0}(t) - \eta \hat{\xi}_u(t) - \eta \hat{Y}_t \\ &= - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}_{g_0}(\tau)}_{W_{g_0}(t+1)} - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} (\bar{H}_\tau - H) \hat{x}_\tau}_{W_H(t+1)} - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}_u(\tau)}_{W_u(t+1)} - \underbrace{\eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{Y}_\tau}_{W_p(t+1)} \end{aligned}$$

where

$$\begin{aligned} \xi_{g_0}(t) &= \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t), \quad \xi'_{g_0}(t) = \frac{1}{m} \sum_{i=1}^m (Z_{t,i} (Z_{t,i})^\top - I) \nabla f(x'_t), \quad \hat{\xi}_{g_0}(t) = \xi_{g_0}(t) - \xi'_{g_0}(t), \\ \xi_u(t) &= \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}_{t,i} Z_{t,i}, \quad \xi'_u(t) = \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}'_{t,i} Z_{t,i}, \quad \hat{\xi}_u(t) = \xi_u(t) - \xi'_u(t), \\ \hat{Y}_t &= Y_t - Y'_t, \quad \bar{H}_t = \int_0^1 \nabla^2 f(ax_t + (1-a)x'_t) da. \end{aligned}$$

To derive the final equality, we utilized the fact that $x'_0 = x_0$. This completes our proof. \square

Suppose x_0 is an ϵ -saddle point. Recall that $\gamma > 0$ denotes $-\lambda_{\min}(\nabla^2 f(x_0))$, where we know that $\gamma \geq \sqrt{\rho\epsilon}$.

$$\gamma \geq \bar{\psi} := \begin{cases} \min\{\psi, 1, L\} & \text{if } f(\cdot) \text{ is } (\epsilon, \psi)\text{-strict saddle for any } \psi > \sqrt{\rho\epsilon} \\ \sqrt{\rho\epsilon} & \text{otherwise.} \end{cases}$$

In the sequel, for any $t \geq 0$, it is helpful to define the quantities

$$\beta(t)^2 := \frac{(1 + \eta\gamma)^{2t}}{(\eta\gamma)^2 + 2\eta\gamma}, \quad \alpha(t)^2 := \frac{(1 + \eta\gamma)^{2t} - 1}{(\eta\gamma)^2 + 2\eta\gamma}. \quad (44)$$

We next introduce some probabilistic events (and their implications) which, if true, can be used to bound the sizes of $\|W_{g_0}(t+1)\|$, $\|W_u(t+1)\|$, $\|W_v(t+1)\|$ (and as we will see in the next result, indirectly bound $\|W_H(t+1)\|$). These bounds will be useful in the final proof of making function value progress near a saddle point.

Lemma 24. *We assume $\delta \in (0, 1/e]$ throughout the lemma. Suppose that we pick u, r and η as specified in Lemma 23. Suppose $T_s \geq t_f(\delta)$. Suppose also that*

$$f(x_{T_s}) - f(x_0) > -F, \quad f(x'_{T_s}) - f(x_0) > -F.$$

Then, we have the following results.

1. Let $\mathcal{S}_\phi(\delta)$ denote the event

$$\mathcal{S}_\phi(\delta) := \left\{ \max\{\|x_t - x_0\|^2, \|x'_t - x_0\|^2\} \leq \phi_{T_s}(\delta, F), \quad \forall 0 \leq t \leq T_s \right\}.$$

In addition, let $\mathcal{S}_u(\delta)$ denote the event

$$\mathcal{S}_u(\delta) := \left\{ \|W_u(t+1)\| \leq \eta\beta(t+1) \frac{\sqrt{3}}{\sqrt{\eta\psi}} (2c_3\rho d^2 (\log(T/\delta))^2) u^2, \quad \forall 0 \leq t \leq T_s - 1 \right\},$$

where c_3 is the same absolute constant as the c_3 in the preceding lemmas. Then,

$$\mathbb{P}(\mathcal{S}_\phi(\delta) \cap \mathcal{S}_u(\delta)) \geq 1 - \frac{24T_s\delta}{T}.$$

2. Consider defining the event $\mathcal{R}_t(\delta)$, which is the event where

$$\text{either } \sum_{\tau=0}^t (1 + \eta\gamma)^{2(t-\tau)} \frac{dL^2}{m} \|x_\tau - x'_\tau\|^2 (\ln(CT^2/\delta))^2 \geq G_{T_s}(\delta, F), \text{ or}$$

$$\|W_{g_0}(t+1)\|$$

$$\leq c' \eta \sqrt{\max\left\{ \left(\ln\left(\frac{CT^2}{\delta}\right) \right)^2 \sum_{\tau=0}^t \frac{dL^2}{m} (1 + \eta\gamma)^{2(t-\tau)} \|x_\tau - x'_\tau\|^2, g(t+1) \right\} \left(\log\left(\frac{CdT^2}{\delta}\right) + \log\left(\log\left(\frac{G_{T_s}(\delta, F)}{g(t+1)}\right)\right) + 1 \right)}$$

normalize holds. Above, c', C refer to the same constants as in Proposition 2, and

$$G_{T_s}(\delta, F) := 8 \sum_{\tau=0}^{T_s-1} (1 + \eta\gamma)^{2\tau} \frac{dL^2}{m} (\ln(CT^2/\delta))^2 \phi_{T_s}(\delta, F) + \left(\frac{\beta(T_s)\eta r}{60\sqrt{d}} \right)^2, \quad g(t+1) := \left(\frac{\beta(t+1)\eta r}{60\sqrt{d}} \right)^2.$$

Then, $\mathbb{P}(\mathcal{R}_t(\delta)) \geq 1 - \frac{\delta}{T}$. Suppose the event

$$\left(\bigcap_{t=0}^{T_s-1} \mathcal{R}_t(\delta) \right) \cap \mathcal{S}_\phi(\delta)$$

holds. Then, the event $\mathcal{S}_{g_0}(\delta)$ holds, where

$$\mathcal{S}_{g_0}(\delta) := \bigcap_{t=0}^{T_s-1} \mathcal{S}_{g_0, t}(\delta),$$

and $\mathcal{S}_{g_0, t}(\delta)$ is defined as

$$\mathcal{S}_{g_0, t}(\delta) := \left\{ \|W_{g_0}(t+1)\| \leq \zeta_1(\delta, F) c' \eta \sqrt{\max\left\{ \left(\ln\left(\frac{CT^2}{\delta}\right) \right)^2 \sum_{\tau=0}^t \frac{dL^2}{m} (1 + \eta\gamma)^{2(t-\tau)} \|x_\tau - x'_\tau\|^2, g(t+1) \right\}} \right\}$$

where

$$\zeta_1(\delta, F) := \left(\log\left(\frac{CdT^2}{\delta}\right) + \log\left(\log\left(\frac{G_{T_s}(\delta, F)}{g(1)}\right)\right) + 1 \right).$$

3. In addition, let $\mathcal{S}_p(\delta)$ denote the event

$$\mathcal{S}_p(\delta) := \left\{ \|W_p(t+1)\| \leq \frac{2\sqrt{2\log(T/\delta)}\beta(t+1)\eta r}{\sqrt{d}} \quad \forall 0 \leq t \leq T_s - 1 \right\}.$$

Then, $\mathbb{P}(\mathcal{S}_p(\delta)) \geq 1 - \frac{T_s\delta}{T}$.

Proof. We consider the three claims separately.

1. Note that our assumptions satisfy the conditions required in Lemma 23. Hence, by Lemma 23, on the event $\mathcal{P}_{0,T_s}(\delta, F)$, we have that $\|x_\tau - x_0\|^2 \leq \phi_{T_s}(\delta, F)$. Simultaneously, on the event $\mathcal{P}_{0,T_s}(\delta, F)$, we know that $\cap_{t=0}^{T_s-1} \mathcal{G}_t(\delta)$ holds, i.e.

$$\frac{1}{m} \sum_{i=1}^m \|Z_{t,i}\|^4 \leq 2c_3 d^2 (\log(T/\delta))^2, \quad \forall 0 \leq t \leq T_s - 1. \quad (45)$$

Thus, for $W_u(t+1)$, we have that

$$\begin{aligned} \|W_u(t+1)\| &= \left\| \eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}_u(\tau) \right\| \\ &\leq \left\| \eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \xi_u(\tau) \right\| + \left\| \eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{\xi}'_u(\tau) \right\| \\ &\leq \eta \sum_{\tau=0}^t (1 + \eta\gamma)^{t-\tau} \left(\left\| \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}_{t,i} Z_{t,i} \right\| + \left\| \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}'_{t,i} Z_{t,i} \right\| \right) \\ &\leq \eta \sum_{\tau=0}^t (1 + \eta\gamma)^{t-\tau} \frac{\rho}{m} \sum_{i=1}^m \|Z_{t,i}\|^4 u^2 \\ &\stackrel{(iv)}{\leq} \eta \sum_{\tau=0}^t (1 + \eta\gamma)^{t-\tau} \rho (2c_3)^2 d^2 (\log(T/2/\delta))^2 u^2 \\ &\leq \eta \frac{(1 + \eta\gamma)^{t+1}}{\eta\gamma} (2c_3 \rho C d^2 (\log(T/\delta))^2) u^2 \\ &\stackrel{(v)}{=} \eta \beta(t+1) \frac{\sqrt{(\eta\gamma)^2 + 2\eta\gamma}}{\eta\gamma} (2c_3 \rho d^2 (\log(T/\delta))^2) u^2 \\ &\leq \eta \beta(t+1) \frac{\sqrt{3}}{\sqrt{\eta\gamma}} (2c_3 \rho d^2 (\log(T/\delta))^2) u^2 \\ &\stackrel{(vi)}{\leq} \eta \beta(t+1) \frac{\sqrt{3}}{\sqrt{\eta\psi}} (2c_3 \rho d^2 (\log(T/\delta))^2) u^2 \end{aligned}$$

where the inequality in (iv) holds due to Eq. (45), the equality in (v) holds due to the definition of $\beta(t+1)$, and the inequality in (vi) used the fact that $\gamma \geq \psi$.

Hence the event

$$\cap_{t=0}^{T_s} \left\{ \|x_t - x_0\|^2 \leq \phi_{T_s}(\delta, F) \text{ and } \right\} \cap \mathcal{S}_u(\delta)$$

holds with probability at least $1 - \frac{12T_s\delta}{T}$.

Note that by the coupling, the distribution of x'_τ is the same as that of x_τ . Thus, by the assumption $f(x'_{T_s}) - f(x_0) > -F$, it follows by a similar argument that the bound $\|x'_\tau - x_0\|^2 \leq \phi_{T_s}(\delta, F)$ also holds with probability at least $1 - \frac{12T_s\delta}{T}$. The claim then follows by an application of the union bound.

2. For the second claim, observe first that the claim $\mathbb{P}(\mathcal{R}_t(\delta)) \geq 1 - \frac{\delta}{T}$ is a consequence of Proposition 2. Suppose next that $f(x_{T_s}) - f(x_0) > -F$. Then, by definition of the event $\mathcal{S}_\phi(\delta)$, we know that

$$\|x_\tau - x_0\|^2 \leq \phi_{T_s}(\delta, F), \quad \|x'_\tau - x_0\|^2 \leq \phi_{T_s}(\delta, F)$$

where $\phi_{T_s}(\delta, F)$ is as defined in Lemma 23.

Suppose now that $\mathcal{R}_t(\delta)$ holds true, and suppose for contradiction that

$$\begin{aligned} & \sum_{\tau=0}^t (1 + \eta\gamma)^{2(t-\tau)} \frac{dL^2}{m} \|x_\tau - x'_\tau\|^2 (\ln(CT^2/\delta))^2 \\ & \geq G_{T_s}(\delta, F) \\ & = 8 \sum_{\tau=0}^{T_s-1} (1 + \eta\gamma)^{2\tau} \frac{dL^2}{m} (\ln(CT^2/\delta))^2 \phi_{T_s}(\delta, F) + \left(\frac{\beta(T_s)\eta r}{60\sqrt{d}} \right)^2. \end{aligned}$$

This implies that there exists some $0 \leq \tau \leq t \leq T_s$ such that $\|x_\tau - x'_\tau\|^2 \geq 8\phi_{T_s}(\delta, F)$. However, we also know that on the event $\mathcal{S}_\phi(\delta)$,

$$\|x_\tau - x'_\tau\|^2 \leq 2\|x_\tau - x_0\|^2 + 2\|x'_\tau - x_0\|^2 \leq 4\phi_{T_s}(\delta, F).$$

This leads to a contradiction. We must then have that

$$\|W_{g_0}(t+1)\| \leq \zeta_1(\delta, F) c' \eta \sqrt{\max \left\{ \left(\ln \left(\frac{CT^2}{\delta} \right) \right)^2 \sum_{\tau=0}^t \frac{dL^2}{m} (1 + \eta\gamma)^{2(t-\tau)} \|x_\tau - x'_\tau\|^2, g(t+1) \right\}}$$

where

$$\zeta_1(\delta, F) := \sqrt{\log \left(\frac{CdT^2}{\delta} \right) + \log \left(\log \left(\frac{G(\delta, F)}{g(1)} \right) + 1 \right)}$$

3. Observe that

$$W_p(t+1) = \eta \sum_{\tau=0}^t (I - \eta H)^{t-\tau} \hat{Y}_\tau = \eta \sum_{\tau=0}^t (1 + \eta\gamma)^{t-\tau} (2(Y_\tau)_1),$$

which means that $W_p(t+1)$ is a 1-dimensional Gaussian with variance

$$\eta^2 \sum_{\tau=0}^t (1 + \eta\gamma)^{2(t-\tau)} \frac{4r^2}{d} = \frac{4\eta^2 r^2}{d} \frac{(1 + \eta\gamma)^{2(t+1)} - 1}{2\eta\gamma + (\eta\gamma)^2} = \frac{4\eta^2 r^2 \alpha(t+1)^2}{d}. \quad (46)$$

Since $\alpha(t+1) \leq \beta(t+1)$, using the subGaussianity of a Gaussian distribution, it follows that for any t , with probability at least $1 - \delta/T$,

$$\|W_p(t+1)\| \leq \frac{2\sqrt{2 \log(T/\delta)} \beta(t+1) \eta r}{\sqrt{d}}.$$

□

For any $F > 0$, we are now ready to show that the algorithm makes a function decrease of F with $\Omega(1)$ probability near an ϵ -saddle point.

Proposition 5. *Suppose that x_{t_0} is an ϵ -approximate saddle point. Let $c' > 0, c_1 > 0, c_2 \geq 1, c_4 > 0, c_5 > 0, c_6 > 0, c_7 > 0, C_1 \geq 1$ be the absolute constants defined in the statements of the previous lemmas, and let $\delta \in (0, 1/e]$ be arbitrary. Consider any $F > 0$. As in the statement of Lemma 23, suppose we choose u, r and η such that*

$$u \leq \frac{\sqrt{\epsilon}}{d\sqrt{\rho} \log(T/\delta)} \cdot \min \left\{ \frac{1}{64c_5^2 c_2}, \frac{1}{2048c_1 c_2} \right\}^{1/4}, \quad r \leq \epsilon \cdot \min \left\{ \frac{1}{8c_5 \sqrt{2c_2}}, \frac{1}{32\sqrt{c_1}} \right\},$$

$$\eta \leq \frac{1}{Lt_f(\delta)} \min \left\{ \frac{1}{\log(T/\delta)}, \frac{\sqrt{m}}{8c_4(\ln(C_1 dmT/\delta))^{3/2} \sqrt{d}}, \frac{m}{128c_1(\ln(C_1 dmT/\delta))^3 d} \right\}.$$

Suppose we pick

$$T_s = \max \left\{ \left\lceil \frac{L}{\eta\bar{\psi}} \right\rceil, t_f(\delta), 4 \right\}, \quad (47)$$

where

$$\iota = \max \left\{ \log \left(2\sqrt{\phi_{T_s}(\delta, F)} \frac{20\sqrt{d}\sqrt{\eta^2\gamma^2 + 2\eta\gamma}}{\eta r} \right), 1 \right\},$$

$$\bar{\psi} := \begin{cases} \min\{\psi, 1, L\} & \text{if } f(\cdot) \text{ is } (\epsilon, \psi)\text{-strict saddle for any } \psi > \sqrt{\rho\epsilon} \\ \sqrt{\rho\epsilon} & \text{otherwise.} \end{cases}$$

Suppose in addition that u, η also satisfy the conditions

$$u \leq \sqrt{\frac{r\sqrt{\eta\bar{\psi}}}{120\sqrt{3}c_3\sqrt{d}\rho d^2(\log(T/\delta))^2}},$$

$$\eta \leq \max \left\{ \frac{1}{c'c_9\zeta_1(\delta, F)}, \frac{m\bar{\psi}}{360\iota(c')^2c_9^2dL^2(\ln(\frac{CT^2}{\delta}))^2\zeta_1(\delta, F)^2}, \frac{1}{2\bar{\psi}} \right\},$$

where $\zeta_1(\delta, F)$ is as defined in Lemma 23, $c', c_3, C > 0$ are the same constants as in the previous results, and $c_9 = 2\sqrt{2} + \frac{1}{20}$. Suppose also that $\phi_{T_s}(\delta, F)$ satisfies the bound

$$\phi_{T_s}(\delta, F) \leq \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2. \quad (48)$$

Then, with probability at least $\frac{1}{3} - \frac{13T_s\delta}{T}$, $f(x_{t_0+T_s}) - f(x_{t_0}) \leq -F$.

Proof of Proposition 5. Without loss of generality, we assume that $t_0 = 0$. By Lemma 3, we have

$$\begin{aligned} & \hat{x}_{t+1} \\ := & x_{t+1} - x'_{t+1} \\ = & -\eta \underbrace{\sum_{\tau=t_0}^t (I - \eta H)^{t-\tau} \hat{\xi}_{g_0}(\tau)}_{W_{g_0}(t+1)} - \eta \underbrace{\sum_{\tau=t_0}^t (I - \eta H)^{t-\tau} (\bar{H}_\tau - H) \hat{x}_\tau}_{W_H(t+1)} - \eta \underbrace{\sum_{\tau=t_0}^t (I - \eta H)^{t-\tau} \hat{\xi}_u(\tau)}_{W_u(t+1)} - \eta \underbrace{\sum_{\tau=t_0}^t (I - \eta H)^{t-\tau} \hat{Y}_\tau}_{W_p(t+1)} \end{aligned}$$

where

$$\xi_{g_0}(t) = \frac{1}{m} \sum_{i=1}^m (Z_{t,i} Z_{t,i}^\top - I) \nabla f(x_t), \quad \xi'_u(t) = \frac{1}{m} \sum_{i=1}^m (Z_{t,i} (Z_{t,i})^\top - I) \nabla f(x'_t), \quad \hat{\xi}_{g_0}(t) = \xi_{g_0}(t) - \xi'_{g_0}(t),$$

$$\xi_u(t) = \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}'_{t,i} Z_{t,i}, \quad \xi'_u(t) = \frac{1}{m} \sum_{i=1}^m \frac{u}{2} Z_{t,i} Z_{t,i} \tilde{H}'_{t,i} Z_{t,i}, \quad \hat{\xi}_u(t) = \xi_u(t) - \xi'_u(t),$$

$$\hat{Y}_t = Y_t - Y'_t, \quad \bar{H}_t = \int_0^1 \nabla^2 f(ax_t + (1-a)x'_t) da.$$

Recall that we define for $t \geq 0$,

$$\beta(t)^2 := \frac{(1 + \eta\gamma)^{2t}}{(\eta\gamma)^2 + 2\eta\gamma}, \quad \alpha(t)^2 := \frac{(1 + \eta\gamma)^{2t} - 1}{(\eta\gamma)^2 + 2\eta\gamma}.$$

Throughout the proof, we suppose for contradiction that

$$f(x_{T_s}) - f(x_0) > -F, \quad f(x'_{T_s}) - f(x_0) > -F,$$

and assume the event

$$\left(\bigcap_{t=0}^{T_s-1} \mathcal{R}_t(\delta)\right) \cap \mathcal{S}_\phi(\delta) \cap \mathcal{S}_u(\delta) \cap \mathcal{S}_p(\delta)$$

holds, where the events intersected are defined in Lemma 24. Then, by Lemma 24, the event $\mathcal{S}_{g_0}(\delta)$ (also defined in Lemma 24) holds⁷.

Consider the following induction argument, where we seek to show that there exists an absolute constant $c_9 > 0$ such that for every $t \in \{0, 1, \dots, T_s\}$,

$$\|x_t - x'_t\| \leq c_9 \log(T/\delta) \frac{\beta(t)\eta r}{\sqrt{d}}, \text{ and } \max\{\|W_{g_0}(t)\|, \|W_H(t)\|, \|W_u(t)\|\} \leq \frac{\beta(t+1)\eta r}{\sqrt{d}} \quad (49)$$

Combined with a lower bound on $\|W_p(t+1)\|$ (which makes use of the property that $W_p(t+1)$ is a 1-dimensional Gaussian), we will then use the inductive claim in Eq. (49) to show that

$$\|W_p(T_s)\| \geq 2 \left(\|W_{g_0}(T_s)\| + \|W_H(T_s)\| + \|W_u(T_s)\| \right).$$

Since $W_p(t+1)$ is a 1-dimensional Gaussian random variable with a standard deviation that grows exponentially with t , by our choice of T_s , we will see that $\|x_{T_s} - x'_{T_s}\|$ is larger than what expect (since our assumptions imply that $\max\{\|x_{T_s} - x_0\|^2, \|x'_{T_s} - x_0\|^2\} \leq \phi_{T_s}(\delta, F)$, i.e. x_{T_s} and x'_{T_s} both remain close to x_0 and hence close to each other). This yields a contradiction, implying that on the event we assumed to hold, i.e.

$$\left(\bigcap_{t=0}^{T_s-1} \mathcal{R}_t(\delta)\right) \cap \mathcal{S}_\phi(\delta) \cap \mathcal{S}_p(\delta)$$

the assumption

$$f(x_{T_s}) - f(x_0) > -F, \quad \text{and } f(x'_{T_s}) - f(x_0) > -F$$

is not true, i.e. one of the sequences must have made function value progress of at least F .

We proceed to prove Eq. (49). Observe that the claim holds for the base case $t = 0$; this is true since $x_0 = x'_0$. Now suppose that this holds for all $\tau \leq t$. We will seek to show that Eq. (49) holds for $t+1$ as well. We do so by bounding the norms of $W_{g_0}(t+1)$, $W_H(t+1)$, $W_u(t+1)$ and $W_p(t+1)$ respectively.

1. (Bounding $\|W_{g_0}(t+1)\|$) Since the event $\mathcal{S}_{g_0}(\delta)$ holds, it follows that for each $0 \leq t \leq T_s - 1$, we have that

$$\|W_{g_0}(t+1)\| \leq \zeta_1(\delta, F) c' \eta \sqrt{\max\left\{\left(\text{lr}\left(\frac{CT^2}{\delta}\right)\right)^2 \sum_{\tau=0}^t \frac{dL^2}{m} (1+\eta\gamma)^{2(t-\tau)} \|x_\tau - x'_\tau\|^2, g(t+1)\right\}}$$

where

$$\zeta_1(\delta, F) := \left(\log\left(\frac{CdT^2}{\delta}\right) + \log\left(\log\left(\frac{G_{T_s}(\delta, F)}{g(1)}\right)\right) + 1 \right),$$

and the terms $G_{T_s}(\delta, F)$ and $g(1)$ are defined as in Lemma 24. Recall by the inductive claim in Eq. (49) that there exists $c_9 > 0$ such that

$$\|x_\tau - x'_\tau\| \leq c_9 \log(T/\delta) \frac{\beta(t)\eta r}{\sqrt{d}} \quad \forall 0 \leq \tau \leq t.$$

Hence, it follows that

$$\|W_{g_0}(t+1)\| \leq c' \zeta_1(\delta, F) \eta \max\left\{\sqrt{t+1} \left(\text{lr}\left(\frac{CT^2}{\delta}\right)\right) \frac{c_9 \sqrt{d} L}{\sqrt{m}} \frac{\beta(t)\eta r}{\sqrt{d}}, \frac{\beta(t+1)\eta r}{60\sqrt{d}}\right\}.$$

⁷We may also directly assume that $\mathcal{S}_{g_0}(\delta)$ also holds, but our way of reasoning prevents double counting of probabilities.

Hence, noting the choice of T_s in Eq. (47), by choosing η such that

$$c'c_9\zeta_1(\delta, F)\eta\sqrt{T_s} \left(\ln \left(\frac{CT^2}{\delta} \right) \right) \frac{\sqrt{d}L}{\sqrt{m}} \leq \frac{1}{60} \iff \eta \leq \frac{m\bar{\psi}}{360\iota(c')^2c_9^2dL^2 \left(\ln \left(\frac{CT^2}{\delta} \right) \right)^2 \zeta_1(\delta, F)^2}, \text{ and} \quad (50)$$

$$c'c_9\zeta_1(\delta, F)\eta \leq 1.$$

it follows that

$$\|W_{g_0}(t+1)\| \leq \frac{\beta(t+1)\eta r}{60\sqrt{d}}.$$

2. Meanwhile, the term $W_H(t+1)$ can be bounded as follows. By the inductive assumption in Eq. (49), we have that

$$\|\hat{x}_\tau\| = \|x_\tau - x'_\tau\| \leq c_9 \log(T/\delta) \frac{\beta(\tau)\eta r}{\sqrt{d}} \quad \forall 0 \leq \tau \leq t.$$

Moreover, on the event our proof assumes, we know that

$$\max \left\{ \|x_\tau - x_0\|^2, \|x'_\tau - x_0\|^2 \right\} \leq \phi_{T_s}(\delta, F).$$

Thus, using the ρ -Hessian Lipschitz property, we have

$$\begin{aligned} \|W_H(t+1)\| &= \eta \left\| \sum_{\tau=0}^t (I - \eta H)^{t-\tau} (\bar{H}_\tau - H) \hat{x}_\tau \right\| \\ &\leq \eta \sum_{\tau=0}^t (1 + \eta\gamma)^{t-\tau} \rho \sqrt{\phi_{T_s}(\delta, F)} \frac{c_9 \log(T/\delta) \beta(\tau)\eta r}{\sqrt{d}} \\ &\leq c_9(t+1) \log(T/\delta) \eta \rho \sqrt{\phi_{T_s}(\delta, F)} \frac{\beta(t)\eta r}{\sqrt{d}} \\ &\leq c_9 T_s \log(T/\delta) \eta \rho \sqrt{\phi_{T_s}(\delta, F)} \frac{\beta(t)\eta r}{\sqrt{d}}. \end{aligned}$$

Given our choice of T_s in Eq. (47), if

$$c_9 T_s \log(T/\delta) \eta \rho \sqrt{\phi_{T_s}(\delta, F)} \leq \frac{1}{60} \iff \phi_{T_s}(\delta, F) \leq \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2$$

it follows that

$$\|W_H(t+1)\| \leq \frac{\beta(t+1)\eta r}{60\sqrt{d}}.$$

3. Meanwhile, for $W_u(t+1)$, since the event $\mathcal{S}_u(\delta)$ holds, we have that

$$\|W_u(t+1)\| \leq \eta\beta(t+1) \frac{\sqrt{3}}{\sqrt{\eta\bar{\psi}}} (2c_3\rho d^2 (\log(T/\delta))^2) u^2.$$

Now, by picking

$$\eta\beta(t+1) \frac{\sqrt{3}}{\sqrt{\eta\bar{\psi}}} (2c_3\rho d^2 (\log(T/\delta))^2) u^2 \leq \frac{\beta(t+1)\eta r}{60\sqrt{d}} \iff u \leq \sqrt{\frac{r\sqrt{\eta\bar{\psi}}}{120\sqrt{3}c_3\sqrt{d}\rho d^2 (\log(T/\delta))^2}},$$

it follows that with probability $1 - \delta/T$, $\|W_u(t+1)\| \leq \frac{\beta(t+1)\eta r}{60\sqrt{d}}$.

4. Meanwhile, observe that since $\mathcal{S}_p(\delta)$ holds, it follows that

$$W_p(t+1) \leq \frac{2\sqrt{2\log(T/\delta)}\beta(t+1)\eta r}{\sqrt{d}}.$$

Combining the bounds for W_{g_0} , W_p , W_H and W_u , it follows that

$$\begin{aligned} \|\hat{x}_{t+1}\| &\leq \|W_{g_0}(t+1)\| + \|W_p(t+1)\| + \|W_H(t+1)\| + \|W_u(t+1)\| \\ &\leq \frac{\beta(t+1)\eta r}{\sqrt{d}} \left(\frac{1}{60} + \frac{1}{60} + \frac{1}{60} + 2\sqrt{2\log(T/\delta)} \right) \\ &\leq \frac{\beta(t+1)\eta r}{\sqrt{d}} \left(\frac{1}{20} + 2\sqrt{2} \right) \log(T/\delta), \end{aligned}$$

where the final inequality uses the fact that $0 < \delta \leq 1/e$ (which implies $\log(T/\delta) \geq 1$). Hence, we see that the first part of the inductive claim of Eq. (49) holds with the constant $c_9 := \frac{1}{20} + 2\sqrt{2}$, and the second part follows naturally as a consequence of our argument above.

Meanwhile, observe that for any η such that $\eta\bar{\psi} \leq \frac{1}{2}$, we have that $(1 + \eta\gamma)^{\frac{1}{\eta\bar{\psi}}} \geq 2$. Thus, by choosing η such that $\eta\bar{\psi} \leq \frac{1}{2}$, we have that for any $t \geq \frac{1}{\eta\bar{\psi}}$,

$$\alpha(t+1)^2 \geq \frac{1}{2}\beta(t+1)^2.$$

Hence, following Eq. (46), by choosing $T_s \geq \frac{1}{\eta\bar{\psi}}$, $W_p(T_s)$ is a 1-dimensional Gaussian with variance at least $\frac{2\eta^2 r^2 \beta(T_s)}{d}$, such that with probability at least $2/3$,

$$\|W_p(T_s)\| \geq \frac{\beta(T_s)\eta r}{10\sqrt{d}}.$$

Simultaneously, we know that on the event

$$\left(\bigcap_{t=0}^{T_s-1} \mathcal{R}_t(\delta) \right) \cap \mathcal{S}_\phi(\delta) \cap \mathcal{S}_u(\delta) \cap \mathcal{S}_p(\delta),$$

we have

$$\|W_{g_0}(T_s)\| + \|W_H(T_s)\| + \|W_u(T_s)\| \leq \frac{3\beta(T_s)\eta r}{60\sqrt{d}} = \frac{\beta(T_s)\eta r}{20\sqrt{d}}.$$

We note that by Lemma 24, we have

$$\mathbb{P} \left(\left(\bigcap_{t=0}^{T_s-1} \mathcal{R}_t(\delta) \right) \cap \mathcal{S}_\phi(\delta) \cap \mathcal{S}_u(\delta) \cap \mathcal{S}_p(\delta) \right) \geq 1 - \left(\frac{24T_s\delta}{T} + \frac{T_s\delta}{T} + \frac{T_s\delta}{T} \right) = 1 - \frac{26T_s\delta}{T}.$$

Thus, with probability at least $2/3 - \frac{26T_s\delta}{T}$, we have

$$\|\hat{x}_{T_s}\| \geq \frac{1}{2}\|W_p(T_s)\| \geq \frac{\beta(T_s)\eta r}{20\sqrt{d}}$$

Thus, choosing $T_s \geq \frac{\iota}{\eta\bar{\psi}}$, where

$$\iota = \max \left\{ \log \left(2\sqrt{\phi_{T_s}(\delta, F)} \frac{20\sqrt{d}\sqrt{\eta^2\gamma^2 + 2\eta\gamma}}{\eta r} \right), 1 \right\},$$

noting that if $\eta\bar{\psi} \leq 1/2$, then $(1 + \eta\gamma)^{\frac{1}{\eta\bar{\psi}}} \geq (1 + \eta\bar{\psi})^{\frac{1}{\eta\bar{\psi}}} \geq 2$, we have that with probability at least $2/3 - \frac{26T_s\delta}{T}$,

$$\begin{aligned} \|\hat{x}_{T_s}\| &\geq \frac{\beta(T_s)\eta r}{20\sqrt{d}} = \frac{\eta r}{20\sqrt{d}} \frac{(1 + \eta\gamma)^{T_s}}{\sqrt{2\eta\gamma + (\eta\gamma)^2}} \\ &\geq \frac{\eta r}{20\sqrt{d}} \frac{(1 + \eta\gamma)^{\frac{\log \left(2\sqrt{\phi_{T_s}(\delta, F)} \frac{20\sqrt{d}\sqrt{\eta^2\gamma^2 + 2\eta\gamma}}{\eta r} \right)}}{\eta\bar{\psi}}}{\sqrt{2\eta\gamma + (\eta\gamma)^2}} \\ &\geq \frac{\eta r}{20\sqrt{d}\sqrt{2\eta\gamma + (\eta\gamma)^2}} 2^{\log \left(2\sqrt{\phi_{T_s}(\delta, F)} \frac{20\sqrt{d}\sqrt{\eta^2\gamma^2 + 2\eta\gamma}}{\eta r} \right)} > 2\sqrt{\phi_{T_s}(\delta, F)} > 2\sqrt{\phi(T_s, \delta)}. \end{aligned}$$

Thus, at least one of $\|x_{T_s} - x_0\|$ and $\|x'_{T_s} - x_0\|$ is larger than $\sqrt{\phi(T_s, \delta)}$, a contradiction. Since the two sequences have the same distribution, it follows that with probability at least $1/3 - \frac{13T_s\delta}{T}$, $f(x_{T_s}) - f(x_0) \leq -F$. \square

In the result above, we require an upper bound on the norm of $\phi_{T_s}(\delta, F)$ to hold (i.e. equation 48), which in turn necessitates an upper bound on F , the function value improvement we can expect to make. Below, we show how to choose F to be as large as possible (up to constants and logarithmic factors) whilst still satisfying equation 48, assuming that u, r and η are chosen appropriately small such that the dominant term of $\|\phi_{T_s}(\delta, F)\|$ scales with F .

Lemma 25. *Consider choosing F such that*

$$F = \frac{1}{2} \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2 \frac{1}{\eta T_s t_f(\delta) (129 + 8c'^2\beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))}.$$

Suppose $\eta \leq \min \left\{ 1, \frac{1}{t_f(\delta)}, \frac{1}{t_f\delta L} \right\}$. Suppose we pick u and r small enough such that

$$u \leq \frac{r^{1/2}}{d \log(T/\delta)\rho^{1/2}}, \quad r^2 \leq \min \left\{ \frac{F\bar{\psi}}{2\iota \log(T/\delta) \left(\frac{65c_2^2}{8} + 6c_1 + 1 \right)}, \frac{F}{4c_6 \log(2dT/\delta) + \frac{8c_7\iota}{\bar{\psi}}} \right\}.$$

Then, $N_{u,r}(T_s, \delta) \leq F$, and that

$$4c_6\eta^2 T_s \log(2dT/\delta)r^2 + 4c_7\eta^2 T_s^2 \rho^2 u^4 d^4 (\log(T/\delta))^4 \leq \eta T_s t_f(\delta)F.$$

Suppose in addition η is small enough so that

$$32\eta^2 (t_f(\delta))^2 \epsilon^2 \leq \frac{1}{2} \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2.$$

Suppose also that $\bar{\psi} \leq 1^8$ and $\eta \leq \frac{m}{d}$, so that $T_s \geq \frac{\iota}{\eta\bar{\psi}} \geq \frac{d}{m}$. Then, the condition in Eq. (48) will be satisfied.

Proof. We note that since $\frac{\iota}{\bar{\psi}} \leq T_s \leq \frac{2\iota}{\bar{\psi}}$, it follows by our choice of r that r also satisfies the condition

$$r^2 \leq \min \left\{ \frac{F}{\eta T_s \log(T/\delta) \left(\frac{65c_2^2}{8} + 6c_1 + 1 \right)}, \frac{F}{4c_6 \log(2dT/\delta) + 4c_7\eta T_s} \right\}.$$

Hence, our choice of η, u and r satisfies the conditions in Lemma 23, and it follows then that

$$\phi_{T_s}(\delta, F) \leq \max \left\{ 128\eta T_s t_f(\delta)F, 32\eta^2 (t_f(\delta))^2 \epsilon^2 \right\} + 8c'^2\beta_1(\delta; F)\eta t_f(\delta) \max \left\{ \frac{16d}{m} (\ln(CT^2/\delta))^2 F, T_s F \right\} + T_s \eta t_f(\delta)F,$$

where $\beta_1(\delta; F)$ is as defined in Lemma 21.

The condition in Eq. (48) requires that

$$\phi_{T_s}(\delta, F) \leq \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2.$$

By our choice of η such that

$$32\eta^2 (t_f(\delta))^2 \epsilon^2 \leq \frac{1}{2} \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2,$$

it suffices for us to show that

$$\begin{aligned} \frac{1}{2} \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2 &\geq 128\eta T_s t_f(\delta)F + 8c'^2\beta_1(\delta; F)\eta t_f(\delta) \max \left\{ \frac{16d}{m} (\ln(CT^2/\delta))^2 F, T_s F \right\} + \eta T_s t_f(\delta)F \\ &= 129\eta T_s t_f(\delta)F + 8c'^2\beta_1(\delta; F)\eta t_f(\delta) \max \left\{ \frac{16d}{m} (\ln(CT^2/\delta))^2 F, T_s F \right\}. \end{aligned}$$

⁸Without loss of generality, we may set $\bar{\psi} = 1$ if $f(\cdot)$ is (ϵ, ψ) -strict saddle for any $\psi > 1$.

By our assumption, we know that $T_s \geq \frac{d}{m}$. Thus, further simplifying indicates that it suffices for us to show

$$\frac{1}{2} \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2 \geq 129\eta T_s t_f(\delta) F + 8c'^2\beta_1(\delta; F)\eta t_f(\delta) \max \{16T_s(\ln(CT^2/\delta))^2 F, T_s F\}. \quad (51)$$

By choosing F such that

$$F \leq \frac{1}{2} \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2 \frac{1}{\eta T_s t_f(\delta) (129 + 8c'^2\beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))},$$

we see that Eq. (51) is satisfied. \square

Remark 2. Suppose without loss of generality that $T_s = \frac{t}{\eta\bar{\psi}}$. Then, as a consequence of Lemma 25, we note that the amortized function value progress of decreasing function value by F over T_s iterations is

$$\begin{aligned} \frac{F}{T_s} &= \frac{1}{2} \left(\frac{\bar{\psi}}{60c_9\iota\rho \log(T/\delta)} \right)^2 \frac{1}{\eta T_s^2 t_f(\delta) (129 + 8c'^2\beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))} \\ &= \eta \frac{\bar{\psi}^4}{\rho^2} \left(\frac{1}{2t^2 (60c_9\iota \log(T/\delta))^2 (t_f(\delta)) (129 + 8c'^2\beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))} \right) \end{aligned}$$

F PROVING THE MAIN RESULT (INFORMAL STATEMENT IN THEOREM 1, FULL STATEMENT IN THEOREM 2)

In this section, we prove our main result. First, we need an additional result (Lemma 26) showing that with high probability, we can bound the function value increase if a saddle appears within $t_f(\delta)$ iterations immediately after we have had T_s iterations after the previous saddle. We note that such a bound is necessary because our earlier result upper bounding function increase in τ iterations (see Lemma 18) focused on the case where $\tau \geq t_f(\delta)$. Next, we state and prove Theorem 2, which is the precise version of Theorem 1 in the main text.

Lemma 26 (Function change for small τ). *Let $c_1 > 0, c_4 > 0, c_5 > 0, C_1 \geq 1$ be the absolute constants defined in the statements of the previous lemmas. Let $\delta \in (0, 1/e]$, and suppose $\tau < t_f(\delta)$.*

Let J denote the interval $\{0, 1, \dots, \tau - 1\}$ where $\tau < t_f(\delta)$.

Suppose we choose η such that

$$\eta \leq \frac{1}{Lt_f(\delta)} \cdot \min \left\{ \frac{\sqrt{m}}{8c_4(\ln(C_1 dmT/\delta))^{3/2} \sqrt{d}}, \frac{m}{128c_1(\ln(C_1 dmT/\delta))^3 d} \right\}. \quad (52)$$

Suppose also we pick u, r and η as prescribed in the statement of Proposition 4.

Suppose that $\min_{t \in J} \|\nabla f(x_t)\| \leq \epsilon$. Then, on the event

$$\mathcal{D}_\tau(\delta) := \mathcal{H}_{0,\tau}(\delta) \cap \left(\bigcap_{t=0}^{\tau-1} \mathcal{A}_t(\delta) \right) \cap \left(\bigcap_{t=0}^{\tau-1} \mathcal{G}_t(\delta) \right),$$

we have the following upper bound on function value change:

$$\begin{aligned} f(x_\tau) - f(x_0) &\leq \frac{\eta}{4} \epsilon^2 + t_f(\delta) \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + t_f(\delta) L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\ &\quad + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + t_f(\delta) c_1 L \eta^2 r^2. \end{aligned}$$

Moreover, $\mathbb{P}(\mathcal{D}_\tau(\delta)) \geq 1 - \frac{(4t_f(\delta)+4)\delta}{T}$.

Proof. Throughout the proof, we assume that the event $\mathcal{D}_\tau(\delta)$ holds.

Let J denote $\{0, 1, \dots, \tau - 1\}$ where $\tau < t_f(\delta)$. Then, J belongs to one of the two following cases.

Case 1 (Gradient dominates noise): Recall that this means that for every $t \in J$, we have

$$\|\nabla f(x_t)\| > 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right).$$

By our choice of η in Eq. (29), we can apply Lemma 16 to get

$$\min_{t \in J} \|\nabla f(x_t)\| \geq \frac{1}{4} \max_{t \in J} \|\nabla f(x_t)\|.$$

Thus by setting $\alpha = 128t_f(\delta)$ in Eq. (3) and by choosing η such that

$$\frac{c_1 L \eta^2 \chi^3 d}{m} \leq \frac{\eta}{\alpha} = \frac{\eta}{128t_f(\delta)} \iff \eta \leq \frac{m}{128c_1 L t_f(\delta) d \chi^3},$$

it follows that

$$\begin{aligned} & -\frac{3\eta}{4} \sum_{t \in J} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{128t_f(\delta)} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t \in J} \|\nabla f(x_t)\|^2 \\ &= -\frac{3\eta}{4} \sum_{t \in J} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \frac{\eta}{64t_f(\delta)} \sum_{t \in J} \|\nabla f(x_t)\|^2 \\ &\leq \frac{\eta}{64t_f(\delta)} \sum_{t \in J} \|\nabla f(x_t)\|^2 \\ &\leq \frac{\eta}{64t_f(\delta)} \sum_{t \in J} \max_{t \in J} \|\nabla f(x_t)\|^2 \\ &\leq \frac{16\eta}{64t_f(\delta)} \sum_{t \in J} \min_{t \in J} \|\nabla f(x_t)\|^2 \leq \frac{\eta}{4} \min_{t \in J} \|\nabla f(x_t)\|^2 \leq \frac{\eta}{4} \epsilon^2, \end{aligned} \tag{53}$$

where the final bound holds since we assumed $\min_{t \in J} \|\nabla f(x_t)\| \leq \epsilon$.

Case 2 (Gradient does not dominate noise): there exists some $t \in J$ such that

$$\|\nabla f(x_t)\| \leq 8t_f(\delta)\eta L \left(\frac{u}{2} \left\| \frac{1}{m} \sum_{i=1}^m Z_{t,i} Z_{t,i}^\top \tilde{H}_{t,i} Z_{t,i} \right\| + \|Y_t\| \right).$$

By our choice of η in Eq. (29), we can apply Lemma 17 to get

$$\|\nabla f(x_t)\| \leq c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \quad \forall t \in J.$$

Note that, by our choices of the parameters η, u, r , it can be shown that

$$c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) < \epsilon,$$

Hence, by setting $\alpha = 128t_f(\delta)$ in Eq. (3) and choosing η such that

$$\frac{c_1 L \eta^2 \chi^3 d}{m} \leq \frac{\eta}{\alpha} = \frac{\eta}{128t_f(\delta)},$$

it follows that

$$\left(\frac{\eta}{128t_f(\delta)} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t \in J} \|\nabla f(x_t)\|^2$$

$$\begin{aligned}
&\leq \frac{\eta}{64t_f(\delta)} \sum_{t \in J} \left(c_5 t_f(\delta) \eta L \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{1 + \frac{\log(T/\delta)}{d}} r \right) \right)^2 \\
&\leq \frac{\eta}{64t_f(\delta)} \sum_{t \in J} \epsilon^2 \\
&\leq \frac{\eta}{64} \epsilon^2
\end{aligned} \tag{54}$$

Combining both cases above (Eq. (53) and Eq. (54)), we see that for the choice $\alpha = 128t_f(\delta)$, the bound

$$-\frac{3\eta}{4} \sum_{t \in J} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{128t_f(\delta)} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t \in J} \|\nabla f(x_t)\|^2 \leq \frac{\eta}{4} \epsilon^2 \tag{55}$$

always holds.

Recall by Eq. (3) that we have

$$\begin{aligned}
f(x_\tau) - f(x_0) &\leq -\frac{3\eta}{4} \sum_{t=0}^{\tau-1} \frac{1}{m} \sum_{i=1}^m |Z_{t,i}^\top \nabla f(x_t)|^2 + \left(\frac{\eta}{\alpha} + \frac{c_1 L \eta^2 \chi^3 d}{m} \right) \sum_{t=0}^{\tau-1} \|\nabla f(x_t)\|^2 \\
&\quad + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
&\quad + \eta c_1 r^2 (\alpha + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2.
\end{aligned}$$

By plugging in Eq. (55) above, as well as the choice $\alpha = 128t_f(\delta)$, we see that

$$\begin{aligned}
f(x_\tau) - f(x_0) &\leq \frac{\eta}{4} \epsilon^2 + t_f(\delta) \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + t_f(\delta) L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
&\quad + \eta c_1 r^2 (128t_f(\delta) + \eta L) \log \frac{T}{\delta} + t_f(\delta) c_1 L \eta^2 r^2.
\end{aligned}$$

We can now complete our proof by using the union bound (suppressing the dependence of some of the events on δ for notational simplicity) to derive

$$\begin{aligned}
\mathbb{P}(\mathcal{D}_\tau^c) &\leq \mathbb{P}(\mathcal{H}_\tau^c) + \sum_{t=0}^{\tau-1} \mathbb{P}(\mathcal{A}_t^c) + \sum_{t=0}^{\tau-1} \mathbb{P}(\mathcal{G}_t^c) \\
&\leq \frac{(\tau+4)\delta}{T} + \frac{\tau}{T} \delta + 2 \frac{\tau}{T} \delta \leq \frac{(4t_f(\delta)+4)}{T} \delta \quad \square
\end{aligned}$$

Armed with Proposition 5 and Lemma 25, we are now ready to show for T sufficiently large, with high probability, there can be no more than $T/4$ ϵ -saddle points. Combined with Proposition 4, this yields the following result.

Theorem 2. *Suppose we pick u, r, η such that they satisfy the conditions in Proposition 5 and Lemma 25. Suppose F is chosen as prescribed in Lemma 25. Suppose that $\bar{\psi} \leq 1$, so that $T_s \geq \frac{L}{\eta \bar{\psi}} \geq \frac{d}{mL}$ ⁹. Suppose we pick T_s as prescribed in Proposition 5. Suppose in addition we pick r such that*

$$r^2 \leq \min \left\{ \frac{\epsilon^2}{4(130c_1 t_f(\delta) + c_1 \log(T/\delta) + c_1)}, \frac{F \bar{\psi}}{80\iota \log(T/\delta) \left(\frac{65c_1^2}{8} + 132c_1 + 1 \right)} \right\}.$$

Suppose also that we choose η such that

⁹Recall we focus on the case $\bar{\psi} \leq L$, since otherwise, by the L -Lipschitz assumption, $\lambda_{\min}(\nabla^2 f(x)) \geq -L$ for all $x \in \mathbb{R}^d$, i.e. ϵ -first order stationary points are also ϵ -second order stationary points.

$$\eta \leq \frac{0.1 \bar{\psi}}{2\epsilon^2} \frac{1}{2l} \left(\frac{\bar{\psi}}{60c_9 t \rho \log(T/\delta)} \right)^2 \frac{1}{t_f(\delta) (129 + 8c'^2 \beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))}$$

Suppose

$$T \geq \left\{ \frac{256t_f(\delta) ((f(x_0) - f^*) + \epsilon^2/L)}{\eta\epsilon^2}, \frac{\varphi\rho^2 (f(x_0) - f^*)}{\eta\bar{\psi}^4}, 256 \lceil \frac{l}{\eta\bar{\psi}} \rceil, 256t_f(\delta), 1024 \right\}, \quad (56)$$

where

$$\varphi := 20 (2l^2 (60c_9 t \log(T/\delta))^2 (t_f(\delta)) (129 + 8c'^2 \beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))).$$

Then, with probability at least $1 - 22\delta$, there are at least $T/2$ ϵ -approximate second order stationary points.

Proof. Consider defining the following sequence of stopping times:

$$\tau_1 = \inf_t \{t \leq T : \|\nabla f(x_t)\| < \epsilon, \lambda_{\min}(\nabla^2 f(x_t)) \leq -\sqrt{\rho\epsilon}\},$$

$$\tau_{i+1} = \inf_t \{t \leq T : t > \tau_i + T_s, \|\nabla f(x_t)\| < \epsilon, \lambda_{\min}(\nabla^2 f(x_t)) \leq -\sqrt{\rho\epsilon}\}, \quad \forall 1 \leq i \leq \lfloor T/T_s \rfloor.$$

□

We note that if $\tau_i = T$, then $\tau_j = T$ for any $j > i$. Let N_s denote the (random) number of saddle points encountered in T iterations.

We observe that we can decompose the function change as

$$\begin{aligned} & f(x_T) - f(x_0) \\ &= (f(x_{\tau_{N_s}}) - f(x_0)) + (f(x_T) - f(x_{\tau_{N_s}})) \\ &= (f(x_{\tau_1}) - f(x_0)) + \sum_{i=1}^{N_s} (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) + \sum_{i=1}^{N_s-1} (f(x_{\tau_{i+1}}) - f(x_{\tau_i+T_s})) + (f(x_T) - f(x_{\tau_{N_s}})) \\ &= \underbrace{\sum_{i=1}^{N_s} (f(x_{\tau_i+T_s}) - f(x_{\tau_i}))}_{U_1} + \underbrace{(f(x_{\tau_1}) - f(x_0)) + \sum_{i=1}^{N_s-1} (f(x_{\tau_{i+1}}) - f(x_{\tau_i+T_s})) + (f(x_T) - f(x_{\tau_{N_s}}))}_{U_2}. \end{aligned}$$

We first consider U_1 . Letting $x_j := x_T$ for any $j \geq T$, we have that

$$\sum_{i=1}^{N_s} f(x_{\tau_i+T_s}) - f(x_{\tau_i}) = \sum_{i=1}^{\lfloor T/T_s \rfloor} (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T}$$

Now, by Eq. (30), observe that with probability at least $1 - \frac{(5T_s+4)\delta}{T} \geq 1 - \frac{6T_s\delta}{T}$ (note $T_s \geq 4$), for any $1 \leq i \leq T/T_s$, we have that

$$\begin{aligned} (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T} &\leq \tau \frac{c_5^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\ &\quad + \tau \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\ &\quad + \eta c_1 r^2 (128t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau c_1 L \eta^2 r^2 \\ &:= M_{u,r,T_s}. \end{aligned}$$

Suppose we pick u, r such that $M_{u,r,T_s} \leq 0.1F$. Recall from Proposition 5 that with probability at least $1/3 - \frac{13T_s\delta}{T}$, $(f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T} \leq -F$. Choosing δ such that $1/3 - \frac{13T_s\delta}{T} \geq 0.3$, and letting $\mu = 0.1F$, we note that $|-F + \mu| = 0.9F \geq \frac{0.7}{0.3}0.2F \geq \frac{0.7}{0.3}(M_{u,r,T_s} + \mu)$.

Now, let \mathcal{E}_{τ_i} denote the bad event on which

$$\text{neither } (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T} \leq -F, \quad \text{nor } (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T} \leq M_{u,r,T_s} \leq 0.1F.$$

We know that \mathcal{E}_{τ_i} has probability at most $\frac{6T_s\delta}{T}$. Let $\mathcal{E}_\tau := \cup_{i=1}^{\lfloor T/T_s \rfloor} \mathcal{E}_{\tau_i}$, such that $\mathbb{P}(\mathcal{E}_\tau) \leq 6\delta$. Then, by applying the weakened supermartingale inequality in Proposition 3, we have

$$\mathbb{P}\left(\sum_{i=1}^{\lfloor T/T_s \rfloor} (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T} \geq -N_s 0.9F + s\right) \leq \mathbb{E}\left[\exp\left(-\frac{s^2}{4N_s F^2}\right)\right] + \mathbb{P}(\mathcal{E}_\tau) \leq \exp\left(-\frac{s^2}{4(T/T_s)F^2}\right) + 6\delta.$$

Now, pick $s = 2F\sqrt{\sqrt{\log(1/\delta)}T/T_s}$, then

$$\mathbb{P}\left(\sum_{i=1}^{\lfloor T/T_s \rfloor} (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T} \geq -N_s 0.9F + 2F\sqrt{\sqrt{\log(1/\delta)}T/T_s}\right) \leq 7\delta.$$

Note that supposing for contradiction that there are at least $T/4$ saddles, we must then have that $N_s \geq T/(4T_s)$, such that

$$-N_s 0.9F + 2F\sqrt{\sqrt{\log(1/\delta)}T/T_s} \leq F(-0.9T/(4T_s) + (2\sqrt{\sqrt{\log(1/\delta)}T/T_s})) \leq F(-0.1T/T_s),$$

where we may ensure the last inequality by picking T/T_s such that

$$T/T_s \geq \left(\frac{2}{0.125}\right)^2 \sqrt{\log(1/\delta)} = 256\sqrt{\log(1/\delta)}.$$

Note that our choice of T ensures this.

Thus, with probability at least $1 - 7\delta$,

$$U_1 = \sum_{i=1}^{\lfloor T/T_s \rfloor} (f(x_{\tau_i+T_s}) - f(x_{\tau_i})) 1_{\tau_i < T} \leq -(0.1T/T_s)F.$$

Next, we bound the summand U_2 . Recall that

$$U_2 = (f(x_{\tau_1}) - f(x_0)) + \sum_{i=1}^{N_s-1} (f(x_{\tau_{i+1}}) - f(x_{\tau_i+T_s})).$$

Without loss of generality, we may analyze each of the summands $f(x_{\tau_{i+1}}) - f(x_{\tau_i+T_s})$ in the same way as we treat $(f(x_{\tau_1}) - f(x_0))$. Let us then consider the summand $f(x_{\tau_1}) - f(x_0)$. There are two cases to consider.

1. The first is when $\tau_1 < t_f(\delta)$. In this case, since we know that $\|\nabla f(x_{\tau_1})\| \leq \epsilon$ (as x_{τ_1} is an ϵ -saddle point), it follows by Lemma 26 that

$$\begin{aligned} f(x_{\tau_1}) - f(x_0) &\leq \frac{\eta}{4}\epsilon^2 + t_f(\delta)\eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta}\right)^3 + t_f(\delta)L\eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta}\right)^4 \\ &\quad + \eta c_1 r^2 (128t_f(\delta) + \eta L) \log \frac{T}{\delta} + t_f(\delta)c_1 L\eta^2 r^2 \end{aligned}$$

with probability at least $1 - \frac{(4t_f(\delta)+4)\delta}{T}$.

2. The second case is when $\tau_1 \geq t_f(\delta)$. In this case, by Lemma 18, we have that

$$f(x_{\tau_1}) - f(x_0) \leq \tau_1 \frac{c_2^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta}\right)^2 + \sqrt{2\log(T/\delta)}r\right)^2$$

$$\begin{aligned}
& + \tau_1 \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + \tau_1 L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
& + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + \tau_1 c_1 L \eta^2 r^2.
\end{aligned}$$

with probability at least $1 - \frac{(5\tau_1+4)\delta}{T}$.

By our choice of u , we know that

$$\begin{aligned}
& t_f(\delta) \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + t_f(\delta) L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + t_f(\delta) c_1 L \eta^2 r^2 \\
& \leq t_f(\delta) r^2 c_1 + t_f(\delta) r^2 c_1 + c_1 r^2 (128 t_f(\delta) + 1) \log(T/\delta) + c_1 r^2 \\
& = r^2 (130 c_1 t_f(\delta) + c_1 \log(T/\delta) + c_1).
\end{aligned}$$

Hence, by picking r such that

$$r \leq \frac{\epsilon^2}{4(130 c_1 t_f(\delta) + c_1 \log(T/\delta) + c_1)},$$

it follows that

$$\begin{aligned}
\frac{\eta \epsilon^2}{4} & \geq t_f(\delta) \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + t_f(\delta) L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
& + \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + t_f(\delta) c_1 L \eta^2 r^2.
\end{aligned}$$

Then, if $\tau_1 < t_f(\delta)$, with probability at least $1 - \frac{(5t_f(\delta)+4)\delta}{T}$,

$$f(x_{\tau_1}) - f(x_0) \leq \frac{\eta \epsilon^2}{2}.$$

Suppose also that we pick r such that

$$r^2 \leq \frac{F \sqrt{\rho \epsilon}}{80 \iota \log(T/\delta) \left(\frac{65 c_2^2}{8} + 132 c_1 + 1 \right)} \leq \frac{F}{40 \eta T_s \log(T/\delta) \left(\frac{65 c_2^2}{8} + 132 c_1 + 1 \right)}.$$

Then, it can be verified that

$$\begin{aligned}
\frac{F}{40} \frac{T}{T_s} & \geq T \frac{c_2^2}{64} \eta^3 t_f(\delta)^2 L^2 \left(u^2 d^2 \rho \left(\log \frac{T}{\delta} \right)^2 + \sqrt{2 \log(T/\delta)} r \right)^2 \\
& + T \eta u^4 \rho^2 \cdot c_1 d^3 \left(\log \frac{T}{\delta} \right)^3 + T L \eta^2 u^4 \rho^2 \cdot c_1 d^4 \left(\log \frac{T}{\delta} \right)^4 \\
& + \frac{T}{T_s} \eta c_1 r^2 (128 t_f(\delta) + \eta L) \log \frac{T}{\delta} + T c_1 L \eta^2 r^2.
\end{aligned}$$

Then, by a union bound, it follows that with probability at least $1 - 9\delta$,

$$\begin{aligned}
U_2 & = (f(x_{\tau_1}) - f(x_0)) + \sum_{i=1}^{N_s-1} (f(x_{\tau_{i+1}}) - f(x_{\tau_i+T_s})) \\
& \leq \frac{T}{T_s} \frac{\eta \epsilon^2}{2} + \frac{F}{40} \frac{T}{T_s}
\end{aligned}$$

Therefore, by the union bound, with probability at least $1 - 16\delta$,

$$f(x_{\tau_{N_s}}) - f(x_0) = U_1 + U_2 \leq \frac{T}{T_s} \left(-0.1F + \eta \epsilon^2/2 + \frac{F}{40} \right)$$

By recalling our choice of F in Lemma 25, by choosing η such that

$$\eta \leq \frac{0.1 \bar{\psi}}{2\epsilon^2} \frac{1}{2\iota} \frac{1}{2} \left(\frac{\bar{\psi}}{60 c_9 \iota \rho \log(T/\delta)} \right)^2 \frac{1}{t_f(\delta) (129 + 8c'^2 \beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))}$$

$$\leq \frac{0.1}{2\epsilon^2} \frac{1}{2} \left(\frac{\sqrt{\rho\epsilon}}{60c_9\iota\rho \log(T/\delta)} \right)^2 \frac{1}{\eta T_s t_f(\delta) (129 + 8c'^2\beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1))} = \frac{0.1F}{2\epsilon^2},$$

it follows that with probability at least $1 - 16\delta$,

$$\begin{aligned} f(x_{\tau_{N_s}}) - f(x_0) &= U_1 + U_2 \\ &\leq \frac{T}{T_s} \left(-0.1F + \eta\epsilon^2/2 + \frac{F}{40} \right) \\ &\leq \frac{T}{T_s} (-0.1F + 0.1F/4 + 0.1F/4) = \frac{T}{T_s} (-0.05F). \end{aligned}$$

Choose T such that

$$-(0.05T/T_s)F \leq -(f(x_0) - f^*) \iff T \geq \frac{20T_s(f(x_0) - f^*)}{F} \geq \frac{\varphi\rho^2(f(x_0) - f^*)}{\eta\psi^4}$$

yields a contradiction, where

$$\varphi := 20 \left(2\iota^2(60c_9\iota \log(T/\delta))^2 (t_f(\delta)) (129 + 8c'^2\beta_1(\delta; F) (16(\ln(CT^2/\delta))^2 + 1)) \right)$$

Hence, with probability at least $1 - 16\delta$, there cannot be more than $T/4$ saddle points. In addition, with probability at least $1 - 6\delta$, by Proposition 4, there cannot be more than $T/4$ iterates with $\|\nabla f(x_t)\| \geq \epsilon$. Hence, with probability at least $1 - 22\delta$, there are at least $T/2$ ϵ -approximate second order stationary points.

G SIMULATIONS

We test the performance of our proposed algorithm with two-point estimators (ZOPGD-2pt) against existing zeroth-order benchmarks using the *octopus function* (proposed in Du et al. (2017)) of varying dimensions. It is known that the octopus function defined on \mathbb{R}^d , which chains d saddle points sequentially, takes exponential (in d) time for exact gradient descent to escape; it has thus emerged as a popular benchmark to evaluate and compare the performance of algorithms that seek to escape saddle points. In our experiments, we compare the performance of our two-point estimator algorithm (ZOPGD-2pt) with PAGD (Algorithm 1 in Vlatakis-Gkaragkounis et al. (2019)) and ZO-GD-NCF (see Zhang et al. (2022)), which are the only two existing zeroth-order algorithms that have (a) a $\tilde{O}(d/\epsilon^2)$ sample complexity for escaping saddle points (with the latter algorithm yielding the tightest bounds), and (b) performed the best empirically on escaping saddle points (see the simulation results in Zhang et al. (2022)). We note that both PAGD and ZO-GD-NCF have to use $2d$ function evaluations per iteration to estimate the gradient while our algorithm only needs to use 2 function evaluations. In our plots, we plot the function value against the number of function evaluations. For completeness, we also plot the performance of exact gradient descent (normalized such that its x -axis is also the number of function queries).

We tested the algorithms for $d = 10$ and $d = 30$. To account for the stochasticity in the algorithms, for each algorithm, we computed the average and standard deviation over 30 trials, and plotted the mean trajectory with an additional band that represents 1.5 times the standard deviation. For our algorithms' hyperparameters, we picked

$$\eta = \frac{1}{4dL}, u = 10^{-2}, r = 0.05, m = 1 \text{ (i.e. two-point estimator)} \quad (57)$$

For PAGD, we used the hyperparameters listed in their paper, and for ZO-GD-NCF, we used the code from their Neurips submission. We note in particular that both methods used the step-size $\frac{1}{4L}$. For initialization, we chose a random x_0 near the saddle point at the origin, drawn from $N(0, 10^{-3}I_{d \times d})$ ¹⁰ (fixed for all trials and all algorithms).

As we can see in Fig. 1, in both cases, our algorithm reaches the global minimum of the octopus function in significantly fewer function evaluations than PAGD and ZO-GD-NCF (approximately 2.5

¹⁰Using the random seed in our code, we note that $\|\nabla f(x_0)\| = 0.011$ for $d = 10$ and $\|\nabla f(x_0)\| = 0.030$ for $d = 30$.

times faster than ZO-GD-NCF, and approximately 3 times faster than PAGD), despite our algorithm only using 2 function evaluations per iteration compared to $2d$ function evaluations per iteration for both PAGD and ZO-GD-NCF. As a sanity check, we note that the number of function evaluations required for PAGD and ZO-GD-NCF to reach the global minimum approximately matches that in Figure 1 of Zhang et al. (2022); here the correspondence is only approximate since Zhang et al. (2022) only plots one trial while we compute the mean and standard deviation of 30 trials.

This result suggests that in addition to the theoretical convergence guarantees, there might also be empirical benefits to using two-point estimators versus existing $2d$ -point estimators in the zeroth-order escaping saddle point literature.

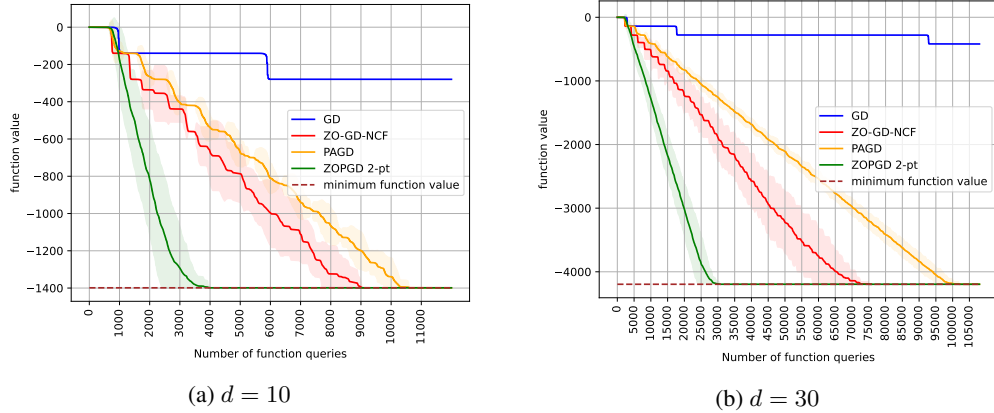


Figure 1: Performance on toy octopus function, with $\tau = e, L = e, \gamma = 1$ (Here, τ, L, γ are parameters determining the properties of f . Our parameter choice is consistent with that in Zhang et al. (2022). See Du et al. (2017) for details about the definitions of τ, L and γ .)