

Gradient Methods with Online Scaling

Part II. Practical Aspects

Ya-Chi Chu^{*1}, Wenzhi Gao^{†2}, Yinyu Ye^{‡2,3}, and Madeleine Udell^{§2,3}

¹Department of Mathematics, Stanford University

²ICME, Stanford University

³Department of Management Science and Engineering, Stanford University

September 13, 2025

Abstract

Part I of this work [14] establishes online scaled gradient methods (**OSGM**)¹, a framework that utilizes online convex optimization to adapt stepsizes in gradient methods. This paper focuses on the practical aspects of **OSGM**. We leverage the **OSGM** framework to design new adaptive first-order methods and provide insights into their empirical behavior. The resulting method, **OSGM-Best**, matches the performance of quasi-Newton variants while requiring less memory and cheaper iterations. We also extend **OSGM** to nonconvex optimization and outline directions that connect **OSGM** to existing branches of optimization theory and practice.

1 Introduction

Consider gradient descent applied to a smooth convex problem $f^* := \min_{x \in \mathbb{R}^n} f(x)$:

$$x^{k+1} = x^k - P_k \nabla f(x^k), \quad (1)$$

where $P_k \in \mathbb{R}^{n \times n}$ is a matrix stepsize. The first part of this work [14] introduces **OSGM**, a framework that uses online learning to adjust stepsize $\{P_k\}$ on the fly and achieves problem-dependent acceleration.

We briefly summarize the idea of **OSGM** introduced in Part I [14]. The analyses of algorithms for smooth optimization often show that each iteration yields some improvement. For instance, in linearly convergent methods, a common improvement is a contraction of the form $\frac{f(x^{k+1}) - f^*}{f(x^k) - f^*} \leq 1 - \frac{1}{\kappa}$. Here $\kappa \geq 1$ often denotes some condition number. If this contraction ratio bound holds uniformly for all k , chaining the per iteration improvements over K steps yields the rate

$$\frac{f(x^{K+1}) - f^*}{f(x^1) - f^*} = \prod_{k=1}^K \frac{f(x^{k+1}) - f^*}{f(x^k) - f^*} \leq (1 - \frac{1}{\kappa})^K.$$

Many proofs first establish improvement at every iteration and then chain these improvements to prove convergence. However, this approach requires monotonic improvement. **OSGM** instead chains the contraction first:

$$\frac{f(x^{K+1}) - f^*}{f(x^1) - f^*} = \prod_{k=1}^K \frac{f(x^{k+1}) - f^*}{f(x^k) - f^*} \leq (\frac{1}{K} \sum_{k=1}^K \frac{f(x^{k+1}) - f^*}{f(x^k) - f^*})^K. \quad (2)$$

^{*}ycchu97@stanford.edu, equal contribution

[†]gwz@stanford.edu, equal contribution

[‡]yeye@stanford.edu

[§]udell@stanford.edu

¹This paper extends two previous works [13] and [5].

To guarantee convergence, it suffices to have the *average* contraction ratio $\frac{1}{K} \sum_{k=1}^K \frac{f(x^{k+1}) - f^*}{f(x^k) - f^*}$ be sufficiently small. In other words, convergence is possible even if some iterations worsen the objective. When the iterates $\{x^k\}$ are generated via the update (1), the contraction ratio attributed to one iteration, starting from x , can be written as $r_x(P) = \frac{f(x - P \nabla f(x)) - f^*}{f(x) - f^*}$, which depends on the stepsize matrix. In view of (2), it suffices to select $\{P_k\}$ sequentially to minimize $\frac{1}{K} \sum_{k=1}^K r_{x^k}(P_k)$, the average contraction ratio. This perspective formulates stepsize selection as a sequential decision-making problem, a central topic in online learning [27, 18]. If $\{P_k\}$ are chosen by a no-regret online learning algorithm, we obtain the guarantee: $\frac{1}{K} \sum_{k=1}^K r_{x^k}(P_k) \leq \frac{1}{K} \sum_{k=1}^K r_{x^k}(\hat{P}) + o(1)$ for any benchmark stepsize \hat{P} , including the optimal preconditioner [15, 31]. The relation further yields the convergence guarantee

$$\frac{f(x^{K+1}) - f^*}{f(x^1) - f^*} \leq (\frac{1}{K} \sum_{k=1}^K r_{x^k}(\hat{P}) + o(1))^K \quad \text{for any } \hat{P}.$$

More generally, **OSGM** models stepsize selection in a gradient-based method as a sequential decision-making problem between a stepsize scheduling agent and a landscape agent. At each iteration:

1. Scheduler makes decision P_k and proposes an update $x^{k+1/2} = x^k - P_k \nabla f(x^k)$.
2. Landscape chooses the next iterate $x^{k+1} = \mathcal{M}(x^k, x^{k+1/2})$ and provides feedback $\ell_{x^k}(P_k)$ to the scheduler.
3. Scheduler updates the stepsize using an online learning algorithm $\mathcal{A}(P_k, \{\ell_{x^j}\}_{j \leq k})$.

Different instantiations of $(\ell_x, \mathcal{M}, \mathcal{A})$ yield different algorithm variants, each with its own convergence properties. Part I of this work develops the theoretical foundations of **OSGM**, providing specific instantiations that achieve both global and local convergence guarantees.

1.1 Contributions

Building on the theoretical foundations established in Part I, the second part focuses on the practical aspects of **OSGM** and outlines several future directions.

Insights into algorithm behavior. While Part I focuses on complexity analysis, this work investigates the empirical behavior of **OSGM**, with a particular focus on **OSGM-H** [14]. We develop a more detailed analysis that explains the empirical behavior of **OSGM-H** and other hypergradient-type methods [2]. These insights lead to concrete algorithmic recommendations.

A practical variant of OSGM. Part I developed **OSGM** variants based on gradient descent. This paper shows how to use the same ideas to choose stepsizes in other gradient-based methods, such as heavy-ball momentum. The strategies above, when combined with heavy-ball momentum, give rise to **OSGM-Best**: a robust variant of **OSGM** that performs competitively with **L-BFGS** [25] while requiring less memory and cheaper iterations.

Nonconvex optimization. Although **OSGM** is grounded in online convex optimization, we demonstrate that the analysis can be extended to smooth nonconvex problems by introducing regularization in the stepsize space. Under the Polyak-Łojasiewicz (PL) condition, the resulting **OSGM** variants provably adapt to locally convex regions of the landscape. Experiments on standard nonconvex benchmarks support these theoretical claims.

Connection to other optimization techniques. We highlight connections between **OSGM** and a range of ideas in optimization, including the Barzilai-Borwein step [1], accelerated gradient descent [8], proximal gradient methods [3], and the performance estimation problem (PEP) [35]. These connections highlight the potential of **OSGM** as a valuable addition to the modern optimization toolbox.

1.2 Notations

We use $\|\cdot\|$ to denote the Euclidean norm. Letters A, a denote matrices and scalars. $\|A\|_F := \sqrt{\sum_{ij} a_{ij}^2}$ denotes the matrix Frobenius norm. We use $\langle \cdot, \cdot \rangle$ for Euclidean or Frobenius inner product. $\Pi_{\mathcal{P}}[\cdot]$ denotes the orthogonal

projection onto a closed convex set \mathcal{P} . We use $\mathcal{X}^* = \{x : f(x) = f^*\}$ to denote the optimal set of f and \mathcal{S}^* to denote the set of stationary points of f ; $\text{dist}(P, \mathcal{P}) := \|P - \Pi_{\mathcal{P}}[P]\|_F$ denotes the distance between a point P and a closed convex set \mathcal{P} ; $\text{diam}(\mathcal{P}) = \max_{X, Y \in \mathcal{P}} \|X - Y\|_F$ denotes the diameter of set \mathcal{P} in Frobenius norm. We use superscript x^k to index algorithm iterates and subscript P_k to index the stepsize sequence. A function f satisfies μ -PL condition if $f(x) - f^* \leq \frac{1}{2\mu} \|\nabla f(x)\|^2$. To emphasize that the feedback functions have stepsize P as their argument, we will use $h_x(P)$ instead of h_x in most contexts.

Structure of the paper. This paper is organized as follows. **Section 2** provides insights into the convergence behavior of OSGM-H observed in the literature. These insights, combined with a heavy-ball momentum enhancement, constitute OSGM-Best, a variant that yields competitive practical performance. **Section 3** generalizes OSGM to smooth nonconvex optimization. **Section 4** presents numerical experiments on both convex and nonconvex optimization problems. **Section 5** outlines several future directions and shows the connection between OSGM and existing optimization theory.

2 Algorithmic behavior and ingredients of practical OSGM

In this section, we motivate the key ingredients of OSGM-Best that improve the practical convergence of OSGM. We provide insights into each ingredient through either rigorous analysis or illustrative examples. Our analysis focuses on OSGM-H, OSGM with hypergradient feedback, although most of the insights extend naturally to other variants of OSGM. We briefly recap OSGM-H (**Algorithm 1**), OSGM using the hypergradient feedback [14]

$$h_x(P) = \frac{f(x - P\nabla f(x)) - f(x)}{\|\nabla f(x)\|^2}.$$

Algorithm 1: Monotone OSGM-H	Algorithm 2: Classic-HDM
1 input: $x^1, P_1 \in \mathbb{R}^{n \times n}$, online gradient stepsize $\eta > 0$ 2 for $k = 1, 2, \dots$ do 3 $x^{k+1/2} = x^k - P_k \nabla f(x^k)$ 4 Let x^{k+1} satisfy $f(x^{k+1}) \leq \min\{f(x^{k+1/2}), f(x^k)\}$ 5 $P_{k+1} = P_k - \eta \nabla h_{x^k}(P_k)$ 6 end	1 input: $x^1, P_1 \in \mathcal{P} = \mathbb{R}^{n \times n}$, $\eta > 0$ 2 for $k = 1, 2, \dots$ do 3 $P_{k+1} = P_k - \eta \nabla h_{x^k}(P_k)$ 4 $x^{k+1} = x^k - P_{k+1} \nabla f(x^k)$ 5 end

The feedback $h_x(P)$ is motivated by the classical descent lemma $f(x - \frac{1}{L} \nabla f(x)) \leq -\frac{1}{2L} \|\nabla f(x)\|^2$ and is closely related to the classic hypergradient descent heuristic (**Algorithm 2**) studied in the literature [2, 32]. However, there are two key differences: the monotonicity condition imposed on Line 4 and the order in which P_k and x^k are updated. These seemingly subtle distinctions turn out to be critical for an effective variant of OSGM-H in both theory and in practice. In this section, we detail our design and address the following aspects:

- *Monotone landscape action for OSGM-H.* Is it necessary to enforce monotonicity in OSGM-H?
- *Limitations of hypergradient descent in the literature.* How does the order of the updates in OSGM-H matter?
- *Heavy-ball momentum and modified feedback.* Can momentum accelerate the convergence of OSGM-H?

After discussing each ingredient in turn, we combine them into one recipe: OSGM-Best.

2.1 Monotone landscape action for hypergradient feedback

The convergence analysis for OSGM-H in Part I [14] always requires a monotone landscape action to ensure at least $f(x^{k+1}) \leq f(x^k)$. In practice, OSGM-H is still observed to converge without enforcing monotonicity, but it exhibits periodic “spikes” in the objective (**Figure 1** (left)). Similar spiky patterns have also been observed in other methods, such as heavy-ball momentum [7] and the Borwein-Barzilai method [11].

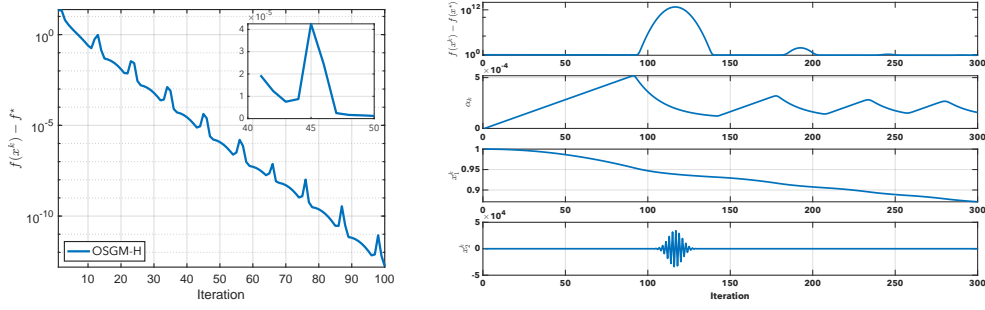


Figure 1: Spiky behavior of Vanilla OSGM-H on a quadratic function.

In this section, we show that this requirement is not an artifact of analysis. Actually, one can construct an example for which the spike of OSGM-H can be arbitrarily high and can happen arbitrarily late. This precludes the possibility of establishing a standard convergence rate in terms of K without the monotone landscape action.

Example. Consider a simple 2D quadratic function parametrized by $\kappa \gg 1$:

$$\min_{x=(x_1, x_2) \in \mathbb{R}^2} f(x) = \frac{1}{2}x_1^2 + \frac{\kappa}{2}x_2^2.$$

The objective f is κ -smooth and 1-strongly convex with condition number κ and $x^* = (0, 0)$. The minimax optimal stepsize for this problem is $\alpha^* = \frac{2}{1+\kappa} \approx \frac{2}{\kappa}$, and stepsizes larger than α^* can lead to divergence. Scalar version of OSGM-H (without enforcing monotonicity) reads

$$\begin{aligned} \begin{pmatrix} x_1^{k+1} \\ x_2^{k+1} \end{pmatrix} &= \begin{pmatrix} (1-\alpha_k)x_1^k \\ (1-\kappa\alpha_k)x_2^k \end{pmatrix}, \\ \alpha_{k+1} &= \alpha_k + \eta \underbrace{\left[\frac{(x_1^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2} (1-\alpha_k) + \frac{\kappa^2(x_2^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2} (1-\kappa\alpha_k) \right]}_{\Delta_k}. \end{aligned} \quad (3)$$

The sign of stepsize change Δ_k is determined by both $1 - \alpha_k$ and $1 - \kappa\alpha_k$. When $\alpha_k > \frac{1}{\kappa}$, $1 - \kappa\alpha_k < 0$ from the second coordinate will try reducing the stepsize, while $1 - \alpha_k$ from the first coordinate tends to increase the stepsize until $\alpha_k \geq 1 \gg \alpha^*$. However, these two terms are weighted by quantities that depend on the location of x^k on the landscape. In particular, if $x_2^k \ll \frac{1}{\kappa}x_1^k$, then $\frac{\kappa^2(x_2^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2} \rightarrow 0$, and the information from the second coordinate $1 - \kappa\alpha_k$ will be disregarded. It can therefore lead to overly large stepsize and divergence. Suppose OSGM-H is initialized such that $x_1^k \gg \kappa x_2^k$. Then each spike in **Figure 1** can be decomposed into three phases.

Myopic stepsize increase. When $x_1^k \gg \kappa x_2^k$, $\Delta_k \approx 1 - \alpha_k$, and OSGM-H drives the stepsize to 1 to accelerate the convergence in the direction of x_1 . However, the magnitude of x_2 explodes since $|1 - \kappa\alpha_k| \gg 1$.

Divergence and recovery. Though the divergence of x_2 starts from a small initial value, it finally becomes large enough that the second term $\frac{\kappa^2(x_2^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2}$ in (3) dominates, decreasing α_k . However, OSGM-H requires many iterations to decrease α_k to $\frac{2}{\kappa}$, and x_2^k continues to increase during this phase, resulting in a spike in the objective.

Periodic spiky behavior. After OSGM-H decreases the stepsize below $\frac{2}{\kappa}$, the second coordinate x_2^k decreases sharply. However, x_2^k decreases much faster than x_1^k does since $1 - \kappa\alpha_k \ll 1 - \alpha_k$ if $\alpha_k = \mathcal{O}(\frac{1}{\kappa})$. Then OSGM-H returns to the state $x_1^k \gg \kappa x_2^k$, and repeats the same divergence-recovery loop. Notice that the magnitude $|x_1|$ consistently decreases (**Figure 1** (right)), since α_k remains less than 1 and thus $|1 - \alpha_k| < 1$. Hence, the algorithm ultimately converges after several loops. But this analysis also precludes a standard convergence proof, as we can pick $\delta \rightarrow 0$ to construct a bad starting point $(1, \delta)$ for any (sufficiently large) iteration K so that a spike happens around K and $f(x^K) > f(x^1)$.

To quantify the spiky behavior, suppose **OSGM-H** starts from $(x_1^1, x_2^1) = (1, \delta)$ for some $\delta > 0$ to be determined later, and the stepsize is initialized by $\alpha_1 = 0$. **Lemma 2.1** shows that **OSGM-H** increases stepsize by at least a constant factor of $\eta \leq \frac{1}{\kappa}$ when α_k is small enough and $x_1^k \gg \kappa x_2^k$.

Lemma 2.1. *Let $\kappa \geq 2$. If $\alpha_k \leq \frac{1}{2}$ and $|x_1^k| \geq \sqrt{2}\kappa^{3/2}|x_2^k|$, then $\alpha_{k+1} \geq \alpha_k + \frac{\eta}{4}$.*

Assume κ is large enough to have a constant a such that $\frac{2}{\kappa} \ll \frac{a}{2} < \frac{1}{2}$, say $a = \frac{1}{3}$. Note that if the condition in **Lemma 2.1** is always satisfied, **OSGM-H** takes at most $\lceil \frac{4a}{\eta} \rceil$ iterations to increase the stepsize from $\alpha_1 = 0$ to at least a .

During the stepsize increasing period, the magnitude $|x_2^k|$ is at most expanded by a factor $|1 - \kappa\alpha_k| < \kappa - 1$ since $\alpha_k < 1$. Therefore, we can choose $\delta > 0$ small enough such that the iterates $\{(x_1^k, x_2^k)\}_{k \leq K}$ always satisfy the condition in **Lemma 2.1** but (x_1^{K+1}, x_2^{K+1}) violates the same condition. Then **OSGM-H** starts to decrease the stepsize after iteration $K + 1$, but the stepsize $\alpha_{K+1} > a \gg \frac{2}{\kappa}$ and thus x_2 continues to expand, causing a spike in the objective. **OSGM-H** takes $\Theta(\frac{1}{\eta\kappa})$ iterations to decrease the stepsize from a to $\frac{a}{2}$ since for $\alpha_k \leq a$, every iteration can decrease the stepsize by at most $\eta(\kappa a - 1)$:

$$\alpha_{k+1} - \alpha_k = \eta \frac{(1-\alpha_k)(x_1^k)^2 + \kappa^2(1-\kappa\alpha_k)(x_2^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2} \geq -\eta(\kappa a - 1).$$

Every iteration during this phase expands x_2 by at least a factor $|1 - \alpha_k\kappa| \geq \frac{a\kappa}{2} - 1 = \frac{\kappa}{6} - 1$ and the spike grows as $\Omega(\kappa^{\frac{1}{\eta\kappa}})$. Hence, given any K , we can pick δ, η , and κ so that a spike appears around K (**Figure 1** (right)).

2.2 Limitations of Classic-HDM

The classic version of hypergradient descent (**Classic-HDM**, **Algorithm 2**) [2, 32] uses a different order for the primal update and stepsize update compared to **OSGM-H**:

Classic-HDM	OSGM-H
$P_{k+1} = P_k - \eta \nabla h_{x^k}(P_k)$	$x^{k+1} = x^k - P_k \nabla f(x^k)$
$x^{k+1} = x^k - P_{k+1} \nabla f(x^k)$	$P_{k+1} = P_k - \eta \nabla h_{x^k}(P_k)$

To facilitate our analysis, for this section only, we drop the monotone landscape action for **OSGM-H**. This simplification does not compromise the intuitions we want to convey. **Classic-HDM** updates the stepsize using the feedback function at x^k first and then applies the newly obtained stepsize P_{k+1} to the gradient step on x^k ; while **OSGM-H** directly applies the stepsize P_k . From an online learning perspective, the decisions in **Classic-HDM** utilize information from the future. Intuitively, it should yield better performance. However, the performance of **Classic-HDM** is inferior to **OSGM-H**. **Figure 2** (left) illustrates the behavior of **Classic-HDM** applied to a quadratic function. There are two observations:

1. **OSGM-H** converges faster than **Classic-HDM** under the same parameter settings.
2. The stepsize $\{\alpha_k\}$ in **Classic-HDM** converges, while $\{\alpha_k\}$ keeps oscillating in **OSGM-H**.

The convergence of $\{\alpha_k\}$ was observed by [2] and was partially explained in [5] in the case where P is a full-dimensional matrix. This section further characterizes this phenomenon on a toy quadratic problem:

$$\min_{x \in \mathbb{R}^n} \frac{1}{2} \langle x, Ax \rangle,$$

where $A \succ 0$ has distinct eigenvalues. We will demonstrate that $\{\alpha_k\}$ in **Classic-HDM** stabilizes around the minimax optimal stepsize $\alpha^* = \frac{2}{L+\mu}$. Besides, the iterates $\{x^k\}$ generated by **Classic-HDM** alternates between two extremal eigen directions of A and lead to slow convergence. On the other hand, in **OSGM-H**, $\{\alpha_k\}$ does not stabilize and keeps adapting.




Figure 2: Left: stepsizes $\{\alpha_k\}$ converge to $\alpha^* = \frac{2}{L+\mu}$ in **Classic-HDM** and oscillate around $\alpha^* = \frac{2}{L+\mu}$ in **OSGM**. Right: as α_k converges, **Classic-HDM** will converge to x^* along two directions.

Dynamical system formulation and orbits. We can further simplify the analysis of **Classic-HDM** and **OSGM-H** by working on the basis that consists of the eigenvectors of A . Let $A = Q\Lambda Q^\top$ be the eigendecomposition. Transform the variables to $z := Qx$. Then **Classic-HDM** and **OSGM-H** with scalar stepsize $P_k = \alpha_k I$ read as

$$\begin{array}{ll}
 \text{Classic-HDM} & \text{OSGM-H} \\
 \alpha_{k+1} = \alpha_k + \eta - \alpha_k \eta \frac{\langle z^k, \Lambda^3 z^k \rangle}{\langle z^k, \Lambda^2 z^k \rangle} & z^{k+1} = (I - \alpha_k \Lambda) z^k \\
 z^{k+1} = (I - \alpha_{k+1} \Lambda) z^k & \alpha_{k+1} = \alpha_k + \eta - \alpha_k \eta \frac{\langle z^k, \Lambda^3 z^k \rangle}{\langle z^k, \Lambda^2 z^k \rangle}
 \end{array}$$

Since the problem is homogeneous, we consider the dynamical system whose state is given by the stepsize α and the direction $\hat{z} := \frac{z}{\|z\|}$:

$$\mathcal{F}_{\text{HDM}}(\alpha, \hat{z}) := (\alpha^+(\hat{z}), \frac{(I - \alpha^+(\hat{z})\Lambda)\hat{z}}{\|(I - \alpha^+(\hat{z})\Lambda)\hat{z}\|}) \quad \text{and} \quad \mathcal{F}_{\text{OSGM}}(\alpha, \hat{z}) := (\alpha^+(\hat{z}), \frac{(I - \alpha\Lambda)\hat{z}}{\|(I - \alpha\Lambda)\hat{z}\|}) \quad (4)$$

where the updated stepsize $\alpha^+(\hat{z}) := \alpha + \eta - \alpha\eta \frac{\langle \hat{z}, \Lambda^3 \hat{z} \rangle}{\langle \hat{z}, \Lambda^2 \hat{z} \rangle}$ starting at state (α, \hat{z}) is the same in both algorithms. Assume A has distinct eigenvalues for simplicity. We show that both \mathcal{F}_{HDM} and $\mathcal{F}_{\text{OSGM}}$ have a period-two orbit.

Lemma 2.2 (Orbit of the dynamical system). *Suppose A has distinct eigenvalues between $[\mu, L]$: $L = \lambda_1 > \lambda_2 > \dots > \lambda_{n-1} > \lambda_n = \mu$. Let q_i denote the unit eigenvectors of A associated with eigenvalues λ_i . Then the dynamical systems \mathcal{F}_{HDM} and $\mathcal{F}_{\text{OSGM}}$ have a common period-2 orbit:*

$$\{(\alpha^*, \hat{z}_1^*), (\alpha^*, \hat{z}_2^*)\} := \{(\frac{2}{L+\mu}, \gamma e_1 + \delta e_n), (\frac{2}{L+\mu}, -\gamma e_1 + \delta e_n)\},$$

where $e_1 = (1, 0, \dots, 0)$, $e_n = (0, \dots, 0, 1)$, $\gamma = \frac{\mu}{\sqrt{L^2 + \mu^2}}$, and $\delta = \pm \frac{L}{\sqrt{L^2 + \mu^2}}$.

Remark 1. Both **Classic-HDM** and **OSGM-H** have additional period-two orbits when the iterates produced fall in a subspace orthogonal to certain eigenvectors of A . For example, when $\langle q_1, z^1 \rangle = 0$, then the orbit can be derived based on $L = \lambda_2$. Without loss of generality, we focus on the orbit generated by the iterations that satisfy $|\langle q_1, z^k \rangle| > 0$, $|\langle q_n, z^k \rangle| > 0$ for all k . In this case, the algorithm iterations always have non-vanishing components along both top and bottom eigen directions.

Although both **Classic-HDM** and **OSGM-H** share the same orbits, they exhibit different stability properties around the orbit.

Theorem 2.1. *Let \mathcal{J}_{HDM} and $\mathcal{J}_{\text{OSGM-H}}$ denote the Jacobians of \mathcal{F}_{HDM} and $\mathcal{F}_{\text{OSGM}}$ respectively. For any $\eta \in (0, \frac{1}{L}]$, their spectral radius satisfy*

$$\rho(\mathcal{J}_{\text{HDM}}(\alpha^*, \hat{z}_1^*) \mathcal{J}_{\text{HDM}}(\alpha^*, \hat{z}_2^*)) < 1 \quad \text{and} \quad \rho(\mathcal{J}_{\text{OSGM}}(\alpha^*, \hat{z}_1^*) \mathcal{J}_{\text{OSGM}}(\alpha^*, \hat{z}_2^*)) > 1$$

In other words, the orbit $\{(\alpha^, \hat{z}_1^*), (\alpha^*, \hat{z}_2^*)\}$ is stable for \mathcal{F}_{HDM} but unstable for $\mathcal{F}_{\text{OSGM}}$.*

Theorem 2.1 indicates that $\{\alpha_k\}$ in **Classic-HDM** will get trapped around α^* and the iterates $\{x^k\}$ approach the optimal solution $x^* = 0$ from two specific directions $d_1 := \gamma q_1 + \delta q_n$ and $d_2 := -\gamma q_1 + \delta q_n$ (**Figure 2** (right)). This spectral property has been extensively discussed in the literature [20, 6], and a common belief is that sticking to the orbit results in slow convergence [39]. In particular, we have

$$\frac{d}{d\alpha} h_{x^* + \lambda d_1}(\alpha^*) = \frac{d}{d\alpha} h_{x^* + \lambda d_2}(\alpha^*) = 0$$

for all $\lambda \neq 0$, which means α^* is exactly the optimal stepsize for this trajectory. As a result, **Classic-HDM** will stop adapting $\{\alpha_k\}$: **Classic-HDM** drives the iterates to a trajectory with no regret. However, despite **Classic-HDM** learns the optimal stepsize on this trajectory, this trajectory is so bad that even the optimal stepsize still offers only slow convergence. In this respect, instability of **OSGM-H** keeps $\{\alpha_k\}$ oscillating between large and small stepsizes and balances different spectral components of x^k . The idea of alternating between large and small stepsizes to leave the orbit has been explicitly used in the literature [9, 38] by designing specific stepsize schedulers. In contrast, **OSGM-H** automatically achieves this goal.

2.3 Heavy-ball momentum and potential function-based feedback

OSGM framework begins with the gradient-based update $x^{k+1} = x^k - P_k \nabla f(x^k)$. In contrast, modern first-order methods with good empirical performance [23] typically add a momentum component

$$x^{k+1} = x^k - P_k \nabla f(x^k) + \beta_k (x^k - x^{k-1}), \quad (5)$$

where $x^k - x^{k-1}$ is known as heavy-ball momentum [7, 29, 30] and $\beta_k \in \mathcal{B} \subseteq \mathbb{R}$ is the extrapolation coefficient. To adapt (P, β) jointly, a straightforward approach defines feedback based on the objective value [2] as in **OSGM-H**:

$$\frac{f(x - P \nabla f(x) + \beta(x - x^-)) - f(x)}{\|\nabla f(x)\|^2}.$$

However, this direct extension does not yield global convergence for two reasons: 1) due to the non-monotone behavior of standard heavy-ball momentum [7], there is no universal (P^*, β^*) that ensures descent of the objective value; 2) a momentum method can take a large step even when the gradient norm is small, so it may happen that $\|x - x^-\| \gg \|\nabla f(x)\|^2$ and hence $\frac{f(x - P \nabla f(x) + \beta(x - x^-)) - f(x)}{\|\nabla f(x)\|^2}$ can have an arbitrarily large Lipschitz constant (P, β) depending on x, x^- . Feedback based on the function value alone does not faithfully reflect the quality of (P, β) and cannot guide the scheduler constructively. To identify a suitable feedback, recall that heavy-ball momentum admits a potential function [7]

$$\varphi_\omega(x, x^-) := f(x) - f^* + \frac{\omega}{2} \|x - x^-\|^2. \quad (6)$$

This potential function monotonically decreases if (P, β) are appropriately chosen.

Lemma 2.3 (Heavy-ball potential [7, Theorem 2]). *Suppose f is L -smooth and convex. Given any $x, x^- \notin \mathcal{X}^*$, denote $x^+ := x - P \nabla f(x) + \beta(x - x^-)$ with $P = \alpha I$, $\alpha \leq \frac{1}{L}$, and $\beta \in [0, \sqrt{1 - \alpha L}]$. Then the potential φ_ω with $\omega = \frac{1 - \alpha L}{2\alpha}$ satisfies*

$$\varphi_\omega(x^+, x) \leq \varphi_\omega(x, x^-) - \frac{\alpha}{2} [\|\nabla f(x)\|^2 + \frac{1 - \alpha L - \beta^2}{\alpha^2} \|x - x^-\|^2].$$

Lemma 2.3 resembles the descent lemma: the progress of each heavy-ball step in reducing φ_ω is proportional to $\|\nabla f(x)\|^2 + \frac{\tau}{2} \|x - x^-\|^2$ for $\tau = \frac{1 - \alpha L - \beta^2}{\alpha^2} > 0$. Therefore, we define the heavy-ball feedback as

$$h_{x, x^-}(P, \beta) := \frac{\varphi_\omega(x^+(P, \beta), x) - \varphi_\omega(x, x^-)}{\|\nabla f(x)\|^2 + \frac{\tau}{2} \|x - x^-\|^2}, \quad (7)$$

where $x^+(P, \beta) = x - P \nabla f(x) + \beta(x - x^-)$ and $\tau > 0$ is a parameter.

The heavy-ball feedback (7) depends on two points x and x^- , which we denote by a state $z = (z_1, z_2) = (x, x^-)$. With slight abuse of notation, we denote $f(z) := f(z_1)$ and $\nabla f(z) := \nabla f(z_1)$. We will also slightly abuse

the notation of the gradient by directly concatenating the gradients with respect to P and β : $\nabla h_z(P, \beta) = (\nabla_P h_z(P, \beta), \nabla_\beta h_z(P, \beta))$. The pushforward operation and the heavy-ball feedback are written with respect to the state z as

$$z^+(P, \beta) := (z_1 - P\nabla f(z) + \beta(z_1 - z_2), z_1), \quad h_z(P, \beta) := \frac{\varphi_\omega(z^+(P, \beta)) - \varphi_\omega(z)}{\|\nabla f(z)\|^2 + \frac{\tau}{2}\|z_1 - z_2\|^2}.$$

Using the superscript k for the iteration count, the first component of the state $z^k := (z_1^k, z_2^k)$ satisfies $z_1^k = x^k$; however, the second component z_2^k need not be equal to x^{k-1} . We will utilize this flexibility to implement the monotone landscape action outlined below. In the product space $\mathcal{P} \times \mathcal{B}$ of stepsize and momentum, define the norm and its associated dual norm as

$$\|(P, \beta)\| := \sqrt{\|P\|_F^2 + \frac{1}{L^2}\beta^2}, \quad \|(P, \beta)\|_* := \sqrt{\|P\|_F^2 + L^2\beta^2}. \quad (8)$$

The weight $\frac{1}{L^2}$ ensures both $\|P\|_F^2$ and $\frac{1}{L^2}\beta^2$ have the same physical dimension since P lies in the space of inverse Hessian and β is dimensionless. We prove properties of feedback $h_z(P, \beta)$ with respect to this norm.

Lemma 2.4 (Properties of heavy-ball feedback). *Suppose f is convex and L -smooth and let $\tau \geq 2L^2$. Then for any state $z = (z_1, z_2)$ with $z_1, z_2 \notin \mathcal{X}^*$, the heavy-ball feedback $h_z(P, \beta)$ is jointly convex in (P, β) and $(L + \omega)$ -smooth in (P, β) . Moreover, the gradient of $h_z(P, \beta)$ takes the form*

$$\nabla h_z(P, \beta) = \left(-\frac{[\nabla f(z^+(P, \beta)) + \omega(-P\nabla f(z) + \beta(z_1 - z_2))]\nabla f(z)^\top}{\|\nabla f(z)\|^2 + \frac{\tau}{2}\|z_1 - z_2\|^2}, \frac{\langle [\nabla f(z^+(P, \beta)) + \omega(-P\nabla f(z) + \beta(z_1 - z_2))], z_1 - z_2 \rangle}{\|\nabla f(z)\|^2 + \frac{\tau}{2}\|z_1 - z_2\|^2} \right).$$

At each iteration, define the progress with respect to the heavy-ball feedback:

$$b_k := \frac{\varphi_\omega(z^{k+1}) - \varphi_\omega(z^k)}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2}\|z_1^k - z_2^k\|^2}.$$

Theorem 2.2 shows a reduction from the function value gap to the cumulative progress $\sum_{k=1}^K b_k$.

Theorem 2.2 (Heavy-ball reduction). *Choose the parameters $\omega > 0$ and $\tau > 0$ such that $\tau \geq 2L\omega$. Let $\{z^k\}$ be any sequence of states such that $\varphi_\omega(z^{k+1}) \leq \varphi_\omega(z^k)$ and let the initial state $z^1 = (z_1^1, z_2^1)$ satisfy $z_1^1 = z_2^1$.*

- If f is convex, then

$$f(z^{K+1}) - f^* \leq \frac{f(z^1) - f^*}{1 + V \sum_{k=1}^K -b_k}, \quad (9)$$

where $V := \min \left\{ \frac{f(z^1) - f^*}{4\Delta^2}, \frac{L}{2} \right\}$ and $\Delta := \max_{x \in \{x: f(x) \leq f(x^1)\}} \min_{x^* \in \mathcal{X}^*} \|x - x^*\|$.

- If f is μ -strongly convex, then

$$f(z^{K+1}) - f^* \leq [f(z^1) - f^*](1 - \frac{\mu}{K} \sum_{k=1}^K -b_k)^K. \quad (10)$$

The heavy-ball reduction (**Theorem 2.2**) requires monotonicity of the potential $\varphi_\omega(z^k)$, instead of function value $f(z^k)$. Hence, given the proposal state from the scheduler

$$z^{k+1/2} := (z^k)^+ = (z_1^k - P_k \nabla f(z^k) + \beta_k(z_1^k - z_2^k), z_1^k), \quad (11)$$

the monotone landscape action for OSGM-HB must ensure the potential is nonincreasing:

$$\varphi_\omega(z^{k+1}) \leq \min\{\varphi_\omega(z^{k+1/2}), \varphi_\omega(z^k)\},$$

where z^{k+1} is the state returned by the monotone landscape action. We may implement this monotone landscape action by the null step $z^{k+1} = \arg \min \{\varphi_\omega(z^{k+1/2}), \varphi_\omega(z^k)\}$, which simply returns whichever state has the smaller potential value. In practice, the lookahead action described in **Section 2.4** is often preferred.

2.4 Algorithm design of OSGM-Best

We put together the insights above and propose OSGM-HB, a variant of OSGM that demonstrates competitive performance in practice. OSGM-Best utilizes the lookahead landscape action, tailored to heavy-ball feedback, to suppress regret and thereby accelerate convergence.

Part I [14] introduced the lookahead landscape action, which takes an additional gradient step on the proposed iterate $x^{k+1/2}$ to suppress the regret of online gradient descent on feedback. Since the heavy-ball feedback is defined with respect to the potential function, the additional gradient step should be taken with respect to the potential function $\varphi_\omega(\cdot, z_2^{k+1/2})$. As a result, the lookahead landscape action for OSGM-HB is

$$\begin{aligned}\mathcal{M}_{\text{look}}(z^{k+1/2}) &:= (z_1^{k+1/2} - \frac{1}{L+\omega} \nabla_1 \varphi_\omega(z^{k+1/2}), z_2^{k+1/2}) \\ &= (z_1^{k+1/2} - \frac{1}{L+\omega} [\nabla f(z^{k+1/2}) + \omega(z_1^{k+1/2} - z_2^{k+1/2})], z_2^{k+1/2}) \\ &= ((1 - \frac{\omega}{L+\omega})z_1^{k+1/2} + \frac{\omega}{L+\omega}z_2^{k+1/2} - \frac{1}{L+\omega} \nabla f(z^{k+1/2}), z_2^{k+1/2}),\end{aligned}$$

where $\nabla_1 \varphi_\omega(z)$ denotes the gradient with respect to z_1 and the stepsize $\frac{1}{L+\omega}$ is the inverse of the smoothness constant $L + \omega$ of the potential function $\varphi_\omega(z)$ with respect to z_1 . **Monotone Lookahead OSGM-HB** determines the next state z^{k+1} by composing the lookahead and monotone landscape actions:

$$\varphi_\omega(z^{k+1}) \leq \min\{\varphi_\omega(\mathcal{M}_{\text{look}}(z^{k+1/2})), \varphi_\omega(z^k)\}.$$

Algorithm 3 summarizes **Monotone Lookahead OSGM-HB**, with monotone landscape action implemented by the null step. The global convergence of **Monotone Lookahead OSGM-HB** is established in **Theorem 2.3**.

Algorithm 3: OSGM-Best

```

1 input Initial state  $z^1 = (z_1^1, z_2^1)$  with  $z_1^1 = z_2^1$ , initial stepsize  $P_1 \in \mathbb{R}^{n \times n}$  and momentum parameter
    $\beta_1 \in \mathbb{R}$ , online gradient descent stepsizes  $\eta_P, \eta_\beta > 0$ 
2 for  $k = 1, 2, \dots$  do
3    $z^{k+1/2} = (z_1^k - P_k \nabla f(z^k) + \beta_k(z_1^k - z_2^k), z_1^k)$  ▷ State proposal from scheduler
4    $z_{\text{look}}^{k+1} = (z_1^{k+1/2} - \frac{1}{L+\omega} \nabla_1 \varphi_\omega(z^{k+1/2}), z_2^{k+1/2})$  ▷ Lookahead action
5   Choose  $z^{k+1}$  satisfying  $\varphi_\omega(z^{k+1}) \leq \min\{\varphi_\omega(z_{\text{look}}^{k+1}), \varphi_\omega(z^k)\}$  ▷ Monotone action by null step
6    $P_{k+1} = P_k - \eta_P \nabla_P h_{z^k}(P_k, \beta_k)$  ▷ Online gradient update on  $P$ 
7    $\beta_{k+1} = \beta_k - \eta_\beta \nabla_\beta h_{z^k}(P_k, \beta_k)$  ▷ Online gradient update on  $\beta$ 
8 end

```

Theorem 2.3 (Global convergence). *Suppose f is L -smooth. Choose $\omega = 3L$, $\tau = 16L^2$ and online gradient stepsizes $\eta_P = \frac{\eta_\beta}{L^2} = \frac{1}{2L}$. Then for any benchmark stepsize $\hat{P} \in \mathbb{R}^{n \times n}$ and momentum parameter $\hat{\beta} \in \mathbb{R}$,*

$$\begin{aligned}f(x^{K+1}) - f^* &\leq \frac{f(x^1) - f^*}{KV \max\{\frac{1}{K} \sum_{k=1}^K -h_{z^k}(\hat{P}, \hat{\beta}) - \frac{L}{4K} \|(P_1 - \hat{P}, \beta_1 - \hat{\beta})\|^2, 0\} + 1}, \quad (\text{convex}) \\ f(x^{K+1}) - f^* &\leq [f(x^1) - f^*](1 - \mu \max\{\frac{1}{K} \sum_{k=1}^K -h_{z^k}(\hat{P}, \hat{\beta}) - \frac{L}{4K} \|(P_1 - \hat{P}, \beta_1 - \hat{\beta})\|^2, 0\})^K. \\ &\hspace{15em} (\mu\text{-strongly convex})\end{aligned}$$

where V and Δ are constants defined in **Theorem 2.2**. In particular, with $(P_1, \beta_1) = (\frac{1}{4L}I, \frac{1}{2})$, we have

$$\begin{aligned}f(x^{K+1}) - f^* &\leq \frac{8L[f(x^1) - f^*]}{KV}, \quad (\text{convex}) \\ f(x^{K+1}) - f^* &\leq [f(x^1) - f^*](1 - \frac{1}{8\kappa})^K. \quad (\mu\text{-strongly convex})\end{aligned}$$

3 OSGM for smooth nonconvex optimization

OSGM can be extended to smooth nonconvex optimization problems, but requires modifying the feedback. The analysis of OSGM in Part I [14] and **Section 2** assumes a convex objective f , which guarantees convexity of feedback and enables hypergradient [14, Theorem 4.2] and heavy-ball (**Theorem 2.2**) progress reduction.

One possible approach to handle nonconvex f is to apply the proximal point method [4] with a sufficiently large regularizer $\gamma > L$:

$$x^{k+1} = \arg \min_x \{f(x) + \frac{\gamma}{2}\|x - x^k\|^2\}. \quad (12)$$

When f is nonconvex but L -smooth, the proximal subproblem (12) is $(L - \gamma)$ -weakly convex and becomes strongly convex if $\gamma > L$. Then OSGM can serve as a subroutine to solve the subproblem (12). However, this proximal point approach requires a double-loop algorithm: an outer loop updates the proximal center x^k and an inner loop solves the proximal subproblem with OSGM. In contrast, this section introduces a different approach to smooth nonconvex optimization by stepsize regularization, avoiding the need for a double-loop algorithm. Observe that $\frac{1}{L}I$ is a safe stepsize in smooth nonconvex optimization. Hence, we can regularize the feedback function $h_x(P)$ by $\|P - \frac{1}{L}I\|_F^2$. This regularizer vanishes when the stepsize $\hat{P} = \frac{1}{L}I$ is used as a benchmark, and an online learning method using this regularized feedback at least competes with the performance of $\frac{1}{L}I$. Regularization in stepsize space requires only a fixed center $\frac{1}{L}I$ and avoids the second loop.

Structure of the section. The roadmap of this section resembles Part I: **Section 3.1** introduces the regularized feedback and establishes its analytical properties; **Section 3.2** introduces two progress reductions for general nonconvex and PL functions and several landscape actions. Finally, in **Section 3.3** we design and analyze OSGM-H variants for nonconvex optimization. For simplicity, we will again focus on the hypergradient feedback, while the results for ratio feedback hold similarly for nonconvex PL functions.

3.1 Regularized feedback design and analysis

Given a regularization parameter $\lambda \geq 0$, define the regularized hypergradient feedback

$$h_x^\lambda(P) := h_x(P) + \frac{\lambda}{2}\|P - \frac{1}{L}I\|_F^2 = \frac{f(x - P\nabla f(x)) - f(x)}{\|\nabla f(x)\|^2} + \frac{\lambda}{2}\|P - \frac{1}{L}I\|_F^2. \quad (13)$$

The regularization parameter λ is chosen to neutralize the *local nonconvexity* of h_x for $P \in \mathcal{P}$. Since $h_x(P)$ is L -smooth, choosing $\lambda = L$ suffices since it guarantees h_x^λ is globally convex. However, this strong regularizer forces P to stay close to the conservative stepsize $\frac{1}{L}I$. In practice, we want to choose λ as small as possible to retain adaptivity, especially when x enters a locally convex region. In other words, we need an estimate for the local weak convexity constant of $h_x(P)$. With additional assumptions on f , it is possible.

Lemma 3.1 (Nonconvexity of h_x). *For any $x \notin \mathcal{S}^* := \{x : \|\nabla f(x)\| = 0\}$, $h_x(P)$ has the following properties:*

1. *If f is L -smooth, then $h_x(P)$ has weak convexity constant $\rho \leq \rho_1 := L$.*
2. *If f is L -smooth, has H -Lipschitz Hessian and $\text{diam}(\mathcal{P}) \leq D$, then $h_x(P)$ has weak convexity constant*

$$\rho \leq \rho_2 := \max\{-\lambda_{\min}(\nabla^2 h_x(0 \cdot I)) + HD\|\nabla f(x)\|, 0\}$$

when $P \in \mathcal{P}$.

3. *Moreover, if f is also μ -PL, then $\rho \leq \rho_3 := (\frac{H}{\mu} + HD)\|\nabla f(x)\|$.*

The intuition behind **Lemma 3.1** is simple: when x is close to stationarity and $\text{diam}(\mathcal{P}) < \infty$, then $\nabla^2 h_x(P)$ does not change much within $P \in \mathcal{P}$. Hence, we can bound the worst-case weak convexity constant of $h_x(P)$ for any $P \in \mathcal{P}$ by evaluating $\lambda_{\min}(\nabla^2 h_x(\bar{P}))$ for some fixed $\bar{P} \in \mathcal{P}$. In particular, if f is μ -PL, $\lambda_{\min}(\nabla^2 h_x(\bar{P}))$ can also be bounded using $\|\nabla f(x)\|$. **Lemma 3.1** takes $\bar{P} = 0 \cdot I$.

Remark 2. Computing $\lambda_{\min}(\nabla^2 h_x(\bar{P}))$ can be computationally expensive in practice. One exception is when $P = \alpha I$ is a scalar multiple of the identity. In this case the second-order derivative of $h_x(\alpha I)$ with respect to α is $\langle \frac{\nabla f(x)}{\|\nabla f(x)\|}, \nabla^2 f(x - \alpha \nabla f(x)) \frac{\nabla f(x)}{\|\nabla f(x)\|} \rangle$ and can be computed by evaluating the Hessian-vector product, without forming the Hessian matrix [28].

Since $h_x^\lambda(P)$ adds a strongly convex regularizer to $h_x(P)$, its analytic properties follow from those of $h_x(P)$.

Lemma 3.2 (Properties of h_x^λ). *For any $x \notin \mathcal{S}^*$, the regularized feedback $h_x^\lambda(P)$ has gradient*

$$\nabla h_x^\lambda(P) = -\frac{\nabla f(x - P \nabla f(x)) \nabla f(x)^\top}{\|\nabla f(x)\|^2} + \lambda(P - \frac{1}{L}I).$$

and $h_x^\lambda(\frac{1}{L}I) \leq -\frac{1}{2L}$ for all λ . Moreover, if f is L -smooth, has H -Lipschitz Hessian and $\text{diam}(\mathcal{P}) \leq D$, then $h_x^\lambda(P)$ is convex and $(L + \lambda)$ -smooth for all $\lambda \geq \min\{\rho_1, \rho_2, \rho_3\}$.

Finally, the definition of minimax stepsize does not immediately carry over to the nonconvex case; therefore, we use $\frac{1}{L}I$ as the benchmark stepsize in the theoretical analyses.

3.2 Nonconvex progress reduction and landscape action

Recall the per iteration progress associated with the hypergradient feedback $h_k := \frac{f(x^{k+1}) - f(x^k)}{\|\nabla f(x^k)\|^2}$. **Theorem 3.1** reduces the stationarity/optimality guarantees to the cumulative progress. Notably, these two reductions do not assume the convexity of f .

Theorem 3.1 (Nonconvex reductions). *Let $\{x^k\}$ be any sequence such that $x^k \notin \mathcal{S}^*$ and $f(x^{k+1}) \leq f(x^k)$. Then*

$$\min_{1 \leq k \leq K} \|\nabla f(x^k)\|^2 \leq \frac{f(x^1) - f^*}{K} \frac{1}{\frac{1}{K} \sum_{k=1}^K -h_k}. \quad (14)$$

Moreover, if f is μ -PL, then

$$f(x^{K+1}) - f^* \leq [f(x^1) - f^*](1 - \frac{2\mu}{K} \sum_{k=1}^K -h_k)^K. \quad (15)$$

Landscape actions \mathcal{M} and progress. As in Part I, the landscape action $\mathcal{M}(x^{k+1/2}, x^k)$ determines the choice of the next iterate in OSGM. Recall the two landscape actions in OSGM-H:

- *Monotone.* x^{k+1} satisfies $f(x^{k+1}) \leq \min\{f(x^{k+1/2}), f(x^k)\}$.
- *Monotone lookahead.* x^{k+1} satisfies $f(x^{k+1}) \leq \min\{f(x^{k+1/2} - \frac{1}{L}\nabla f(x^{k+1/2})), f(x^k)\}$.

For each action, we associate the progress and feedback in **Lemma 3.3**.

Lemma 3.3 (Feedback and progress). *Let f be L -smooth and suppose $\lambda_k \leq L$. Then each of the above two landscape actions guarantees the following relation between the feedback and per iteration progress:*

- *Monotone.* $h_k \leq \min\{h_{x^k}^{\lambda_k}(P_k) - \frac{\lambda_k}{2}\|P_k - \frac{1}{L}I\|_F^2, 0\}$.
- *Monotone lookahead.* $h_k \leq \min\{h_{x^k}^{\lambda_k}(P_k) - \frac{1}{4L}\|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2, 0\}$.

3.3 Algorithm design and analysis

Putting the results together, we instantiate **Monotone Lookahead OSGM-H** for nonconvex optimization with

$\ell_x(P) := h_x^\lambda(P)$ Monotone Lookahead landscape: $f(x^{k+1}) \leq \min\{f(x^{k+1/2} - \frac{1}{L}\nabla f(x^{k+1/2})), f(x^k)\}$ $\mathcal{A} := \text{OGD}$.

The algorithm is **Monotone Lookahead OSGM-H** with regularization (**Algorithm 4**).

Algorithm 4: Monotone Lookahead OSGM-H for nonconvex optimization

```

1 input initial point  $x^1$ , initial stepsize  $P_1 \in \mathcal{P}$ , online gradient stepsize  $\eta > 0$ 
2 for  $k = 1, 2, \dots$  do
3    $x^{k+1/2} = x^k - P_k \nabla f(x^k)$ 
4   Choose  $x^{k+1}$  that satisfies  $f(x^{k+1}) \leq \min\{f(x^{k+1/2} - \frac{1}{L} \nabla f(x^{k+1/2})), f(x^k)\}$ 
5   Choose  $0 \leq \lambda_k \leq L$  based on Lemma 3.1 such that  $h_{x^k}^{\lambda_k}(P)$  is convex within  $\mathcal{P}$ 
6    $P_{k+1} = \Pi_{\mathcal{P}}[P_k - \eta \nabla h_{x^k}^{\lambda_k}(P_k)]$ 
7 end

```

Recall that the gradient of the regularized hypergradient feedback in Line 6 of **Algorithm 4** takes the form

$$\nabla h_x^\lambda(P) = -\frac{\nabla f(x - P \nabla f(x)) \nabla f(x)^\top}{\|\nabla f(x)\|^2} + \lambda(P - \frac{1}{L} I)$$

and \mathcal{P} can be chosen as a simple bounded set (e.g., an element-wise box) to allow inexpensive orthogonal projection. Due to the difficulty of the nonconvex setting, we restrict our analysis to global convergence only.

Theorem 3.2 (Convergence). *Let f be L -smooth. For any benchmark stepsize $\hat{P} \in \mathcal{P}$, Monotone Lookahead OSGM-H (**Algorithm 4**) with $\eta = 1/(4L)$ satisfies*

$$\begin{aligned} \min_{1 \leq k \leq K} \|\nabla f(x^k)\|^2 &\leq \frac{f(x^1) - f^*}{K} \frac{1}{\frac{1}{K} \sum_{k=1}^K [-h_{x^k}(\hat{P}) - \frac{\lambda_k}{2} \|\hat{P} - \frac{1}{L} I\|_F^2] - \frac{L}{K} \|P_1 - \hat{P}\|_F^2} \quad (\text{nonconvex}) \\ f(x^{K+1}) - f^* &\leq [f(x^1) - f^*] (1 - 2\mu \max\{\frac{1}{K} \sum_{k=1}^K [-h_{x^k}(\hat{P}) - \frac{\lambda_k}{2} \|\hat{P} - \frac{1}{L} I\|_F^2] - \frac{L}{K} \|P_1 - \hat{P}\|_F^2, 0\})^K. \quad (\mu\text{-PL}) \end{aligned}$$

In particular, if $P_1 = \frac{1}{L} I$, then

$$\begin{aligned} \min_{1 \leq k \leq K} \|\nabla f(x^k)\|^2 &\leq \frac{2L[f(x^1) - f^*]}{K} \quad (\text{nonconvex}) \\ f(x^{K+1}) - f^* &\leq [f(x^1) - f^*] (1 - \frac{1}{\kappa})^K. \quad (\mu\text{-PL}) \end{aligned}$$

Theorem 3.2 does not show an improved convergence rate compared to gradient descent using constant stepsize $\frac{1}{L} I$. However, when the iterates enter a locally convex region (e.g., approaching a local minimum), λ_k will automatically vanish, and we recover adaptivity in the convex case.

4 Numerical experiments

This section conducts numerical experiments to validate the empirical performance of hypergradient descent. We compare OSGM-Best with different adaptive optimization algorithms.

4.1 Experiment setup

Dataset. We use 38 datasets from LIBSVM for convex optimization tasks and 47 unconstrained problems from CUTEst dataset [16].

Algorithm benchmark. We benchmark the following algorithms.

- GD. Vanilla gradient descent.
- GD-HB. Gradient descent with heavy-ball momentum. [29]
- AGD-CVX. The smooth convex version of accelerated gradient descent (Nesterov momentum). [8]
- AGD-SCVX. The smooth strongly convex version of accelerated gradient descent. [8]

- **Adam**. Adaptive momentum estimation. [23]
- **AdaGrad**. Adaptive (sub)gradient method. [10]
- **BFGS**. BFGS from `scipy` [26, 36].
- **L-BFGS-Mk**. L-BFGS with memory size k in `scipy`.
- Practical variant **OSGM-Best** uses as memory 7 vectors of size n , comparable to memory for L-BFGS-M1.

Algorithm configuration.

- Stepsize in GD, GD-HB, AGD-CVX, and AGD-SCVX are all set to $1/L$.
- The momentum parameter in GD-HB is chosen within the set $\{0.1, 0.5, 0.9, 0.99\}$.
- The Adam stepsize is chosen within the set $\{1/L, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$. $\beta_1 = 0.9, \beta_2 = 0.999$.
- The AdaGrad stepsize is chosen within the set $\{1/L, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$.
- BFGS, L-BFGS-Mk use default parameters in `scipy`.
- OSGM-Best uses the default parameters without tuning.

Testing configurations.

- 1) *Maximum oracle access*. For convex problems, we allow a maximum of 1000 gradient oracles for each algorithm. For nonconvex problems, we allow a maximum of 2000 gradient oracles.
- 2) *Initial point*. For convex problems, All the algorithms are initialized from the same starting point generated from the normal distribution $\mathcal{N}(0, I_n)$ and normalized to have unit length. For nonconvex problems, we use the starting point provided by CUTEst.
- 3) *Stopping criterion*. Algorithms stop if $\|\nabla f\|_\infty \leq 10^{-3}$.

Table 1: Number of solved problems for each algorithm.

Algorithm/Problem	SVM (38) \uparrow	Logistic Regression (38) \uparrow	CUTEst (47) \uparrow
GD	5	2	13
GD-HB	13	10	20
AGD-CVX	13	6	20
AGD-SCVX	10	9	15
Adam	31	16	24
AdaGrad	11	11	18
L-BFGS-M1	27	18	35
L-BFGS-M3	33	19	37
L-BFGS-M5	35	26	37
L-BFGS-M10	35	31	36
BFGS	37	36	38
OSGM-Best	35	35	34

4.2 Convex optimization

For each algorithm, we record the number of successfully solved instances ($\|\nabla f\|_\infty \leq 10^{-3}$ within 1000 gradient oracles). **Table 1** summarizes the detailed statistics. The number of instances solved by **OSGM-Best** is comparable to that of L-BFGS-M10.

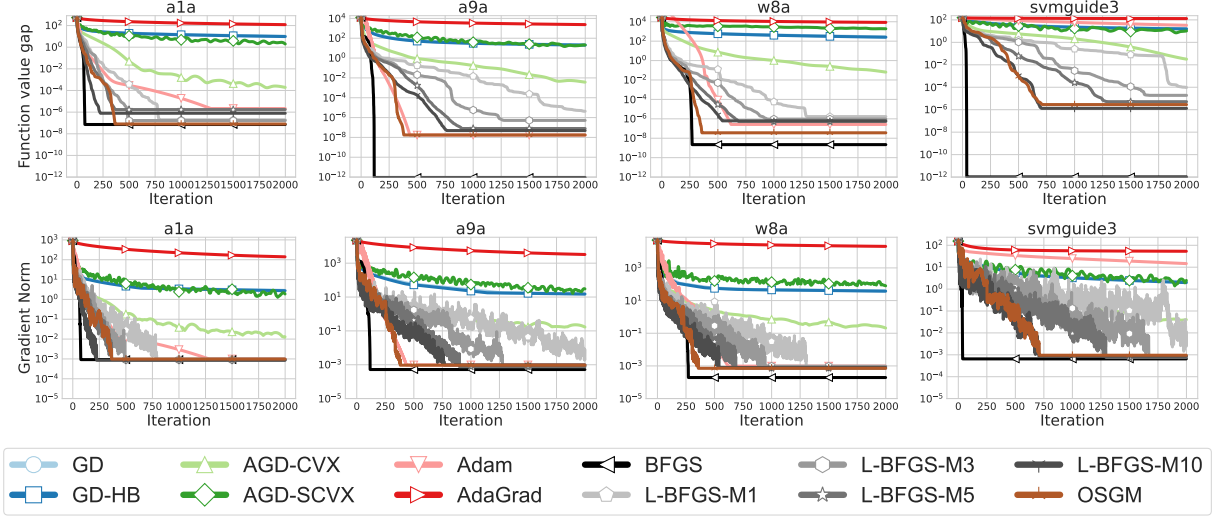


Figure 3: Support vector-machine problems. First row: function value gap. Second row: gradient norm.

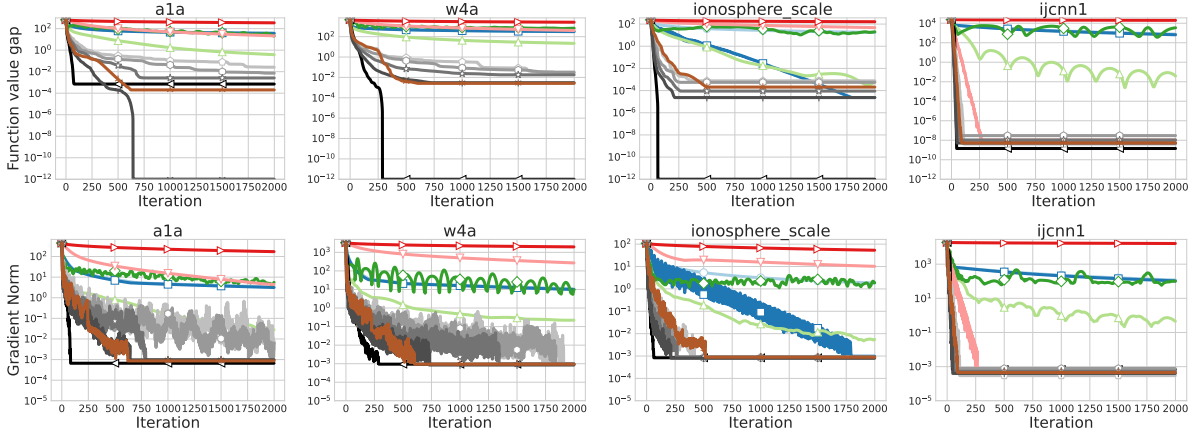


Figure 4: Logistic regression problems. First row: function value gap. Second row: gradient norm.

Support Vector Machine. Figure 5 shows the function value gap and gradient norm plots on sample test instances on support vector machine problems. The optimal value for each instance is obtained by running BFGS until $\|\nabla f\|_\infty \leq 10^{-3}$. We see that the practical variant of OSGM-Best achieves a significant speedup over other adaptive first-order methods. In particular, OSGM-Best often matches L-BFGS-M5 and L-BFGS-M10, while its memory usage is closer to L-BFGS-M1. Notably, Adam also achieves competitive performance in several instances.

Logistic Regression. In logistic regression (Figure 4), OSGM-Best still compares well with L-BFGS-M5 and is significantly faster than other adaptive first-order methods.

4.3 Nonconvex optimization

On nonconvex problems, OSGM-Best still significantly outperforms standard adaptive first-order methods, and its performance is comparable to L-BFGS.

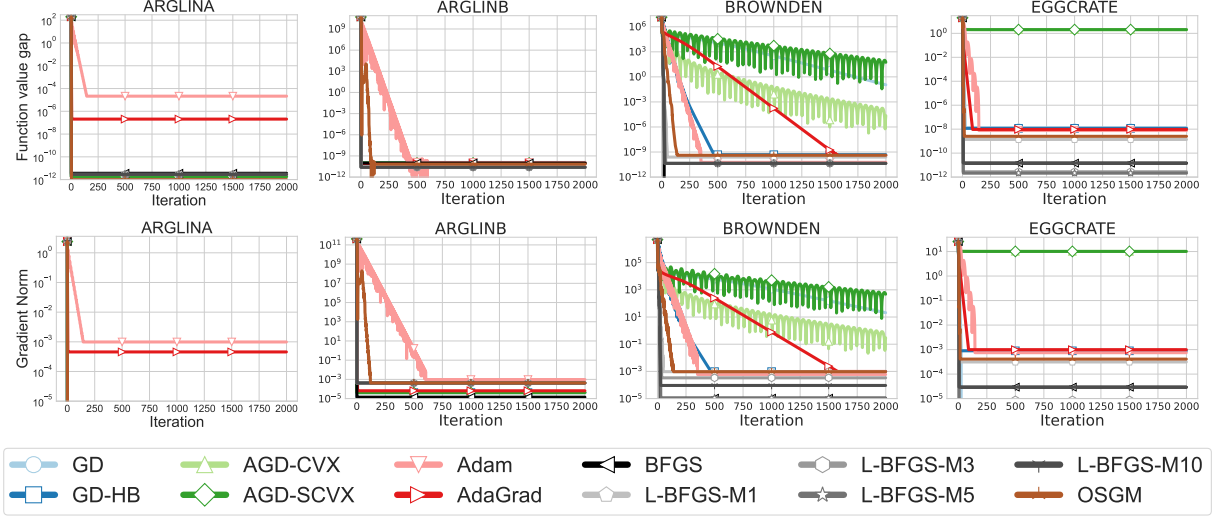


Figure 5: CUTEst problems. First row: function value gap. Second row: gradient norm.

5 Future directions and open problems

In this section, we sketch extensions of OSGM and provide several possible future directions.

5.1 Barzilai-Borwein stepsize and online proximal point

The Barzilai-Borwein stepsize (also known as the BB step) is a widely used practical stepsize selection strategy [1]. There are two types of BB steps:

$$\alpha_k := \frac{\|x^k - x^{k-1}\|^2}{\langle \nabla f(x^k) - \nabla f(x^{k-1}), x^k - x^{k-1} \rangle} \quad \text{and} \quad \alpha'_k := \frac{\langle \nabla f(x^k) - \nabla f(x^{k-1}), x^k - x^{k-1} \rangle}{\|\nabla f(x^k) - \nabla f(x^{k-1})\|^2}.$$

In the context of convex quadratic minimization, it is known that α_k is the steepest descent stepsize *at the previous iteration* in terms of function value [39, Section 2.1],

$$\alpha_k = \arg \min_{\alpha} f(x^{k-1} - \alpha \nabla f(x^{k-1})),$$

and similarly, α'_k is the steepest descent stepsize at the previous iteration in terms of the gradient norm

$$\alpha'_k = \arg \min_{\alpha} \|\nabla f(x^{k-1} - \alpha \nabla f(x^{k-1}))\|.$$

Both BB steps fit into OSGM framework by rewriting them as the minimizers of feedback functions:

$$\alpha_{k+1} = \arg \min_{\alpha} h_{x^k}(\alpha) := \frac{f(x^k - \alpha \nabla f(x^k)) - f(x^k)}{\|\nabla f(x^k)\|^2} \quad \text{and} \quad \alpha'_{k+1} = \arg \min_{\alpha} g_{x^k}(\alpha) := \frac{\|\nabla f(x^k - \alpha \nabla f(x^k))\|}{\|\nabla f(x^k)\|},$$

where $h_x(\alpha)$ is the hypergradient feedback and $g_x(\alpha)$ is the gradient norm feedback discussed in [13]. The BB steps can be derived from the minimization of one of the feedback functions, and we may further develop this perspective to create variants of the BB step powered by online learning. For example, we can write the BB step as the online proximal point update (also known as implicit online learning [24]) with $\eta \rightarrow \infty$

$$\alpha_{k+1} = \arg \min_{\alpha} \{h_{x^k}(\alpha) + \frac{1}{2\eta}(\alpha - \alpha_k)^2\}. \quad (16)$$

To our knowledge, the connection between the BB step and online learning has not been previously explored. Related ideas have primarily appeared in the literature under the guise of “retarded” or “delayed” gradient updates [12, 39]. We believe this new perspective can motivate extensions or improvements on the BB step. For example, using finite values of η above would yield stabilized algorithms similar to the BB step, with a pathway towards a provable convergence guarantee through this connection to the online proximal point method and OSGM. Alternatively, this perspective clarifies how to develop momentum or diagonal variants of the BB step. In Appendix C.1, we illustrate a convergence analysis of OSGM-H based on (16).

Question 1. *The BB step is guaranteed to converge on quadratics, but no theoretical guarantees are available for general functions without complex linesearch [37]. Is it possible to use this online learning perspective to design a competitive variant of the BB step with a theoretically guaranteed convergence rate for general functions?*

Finally, for quadratic f , the BB step uses the steepest descent stepsize from the previous iteration. In contrast, the steepest descent method knows the future ahead of time and applies the “optimal” stepsize. Although the BB step can be unstable when f has a large condition number, it typically outperforms steepest descent. This phenomenon, where knowing the future hinders convergence, aligns with our observation in Section 2.2. To formally state the analogy, steepest descent is to BB step what Classic-HDM is to OSGM-H.

5.2 OSGM-R without knowing f^*

Although OSGM-H has more appealing practical performance, only OSGM-R theoretically improves the complexity of gradient descent. However, OSGM-R has a drawback of requiring knowledge of f^* .

Question 2. *Is it possible to obtain $\mathcal{O}(\kappa^* \log(1/\varepsilon))$ complexity without knowledge of f^* ?*

This issue also appears in well-known methods such as the Polyak stepsize [19]. Two strategies from the Polyak stepsize literature can be borrowed: a double-loop algorithm that searches for f^* can maintain a provable convergence rate, but introduces an additional $\log(1/\varepsilon)$ factor in the bound [19]. On the other hand, methods that dynamically estimate f^* , such as the variable target or moving target methods [17, 22], typically cannot recover a nonasymptotic convergence rate in the smooth strongly convex case. Here, we demonstrate an additional (partial) solution to this question in the special case when a dual bound on the primal objective f can be obtained using the Fenchel trick. Suppose the original optimization problem takes the form

$$\min_{x \in \mathbb{R}^n} f(x) := f_1(x) + f_2(Ax)$$

where $A \in \mathbb{R}^{m \times n}$. Then, using the Fenchel dual trick, we can write the dual problem as

$$\max_{y \in \mathbb{R}^m} g(y) := -f_1^*(-A^\top y) - f_2^*(y).$$

Suppose $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $g : \mathbb{R}^m \rightarrow \mathbb{R}$ are both smooth strongly convex, and that first-order oracles $f(x), \nabla f(x), g(y), \nabla g(y)$ can be efficiently evaluated. Examples satisfying these two conditions include ridge regression, where $f(x) = \frac{1}{2}\|Ax - b\|^2 + \frac{\lambda}{2}\|x\|^2, g(y) = -\frac{1}{2\lambda}\|A^\top y\|^2 - \langle b, y \rangle - \frac{1}{2}\|y\|^2$ and certain regularized generalized linear models [33]. Then strong duality holds $\min_x f(x) = f^* = \max_y g(y)$, and we define the duality gap function $F(z) = F(x, y) := f(x) - g(y)$. Since $F(z^*) = 0$, we can apply OSGM to the duality gap function F to achieve $\mathcal{O}(\kappa_F^* \log(1/\varepsilon))$ complexity, where κ_F^* is the optimal condition number of the gap function. This idea can also be applied to remove f^* in the Polyak stepsize.

5.3 OSGM for accelerated gradient descent

OSGM works well with heavy-ball momentum. Can it similarly improve the performance of Nesterov momentum, for example, by using AGD as our base algorithm instead of gradient descent?

Question 3. *Is there a method that (with knowledge of f^*) achieves (asymptotic) $\mathcal{O}(\sqrt{\kappa^*} \log(1/\varepsilon))$ complexity with only access to a first-order oracle for f ?*

Algorithm 5: Accelerated gradient descent

```

1 input initial point  $x^1 = z^1$ 
2 for  $k = 1, 2, \dots$  do
3    $y^k = x^k + \frac{1}{\sqrt{\kappa}+1}(z^k - x^k)$ 
4    $x^{k+1} = y^k - \frac{1}{L}\nabla f(y^k)$ 
5    $z^{k+1} = (1 - \frac{1}{\sqrt{\kappa}})z^k + \frac{1}{\sqrt{\kappa}}(y^k - \frac{1}{\mu}\nabla f(y^k))$ 
6 end

```

Algorithm 6: Preconditioned AGD

```

1 input initial point  $x^1 = z^1$ 
2 for  $k = 1, 2, \dots$  do
3    $y^k = x^k + \frac{1}{\sqrt{\kappa}+1}(z^k - x^k)$ 
4    $x^{k+1} = y^k - P\nabla f(y^k)$ 
5    $z^{k+1} = (1 - \frac{1}{\sqrt{\kappa}})z^k + \frac{1}{\sqrt{\kappa}}(y^k - \kappa P\nabla f(y^k))$ 
6 end

```

To be concrete, consider the AGD and preconditioned AGD iterations (**Algorithm 5**, **Algorithm 6**), where the standard x update in Line 4 is replaced by $x^{k+1} = y^k - P_k \nabla f(y^k)$. We believe that the biggest challenge in combining OSGM with AGD lies in feedback design: as in the heavy-ball method, AGD does not guarantee monotonicity in $f(x)$. To faithfully reflect the quality of a stepsize P , we likely require another potential function to measure the quality of a stepsize P . Unfortunately, potential functions from the literature are all impossible to evaluate efficiently. In fact, they all require knowing the optimum x^* in advance [8]! For example, $\phi(x, z) := f(x) - f^* + \frac{\mu}{2}\|z - x^*\|^2$ is a valid potential function [8]:

$$[f(x^{k+1}) - f^* + \frac{\mu}{2}\|z^{k+1} - x^*\|^2] \leq (1 - \frac{1}{\sqrt{\kappa}})[f(x^k) - f^* + \frac{\mu}{2}\|z^k - x^*\|^2].$$

In view of this problem, the challenge is to design a potential function for AGD that does not depend on x^* . For quadratic functions, such a potential function exists (**Theorem 5.1**).

Theorem 5.1 (Potential function for AGD on quadratic). *Suppose f is a strongly-convex quadratic, and that the iterates $\{(x^k, z^k)\}$ are generated by AGD with a fixed value of the momentum β . Define the potential*

$$\varphi_\mu(x, z) := f(z) - f^* + \frac{1}{2\mu}\|\nabla f(x)\|^2.$$

Then

$$\varphi_\mu(x^{k+1}, z^{k+1}) \leq \left[\frac{(\kappa-1)^2}{\kappa(\sqrt{\kappa}+1)^2} \left(1 + \frac{1}{2\kappa} + \frac{\sqrt{4\kappa+1}}{2\kappa}\right) \right] \varphi_\mu(x^k, z^k) \leq (1 - \frac{1}{\sqrt{\kappa}})\varphi_\mu(x^k, z^k).$$

This potential still requires knowledge of f^* , but as we have seen in **Section 5.2**, it is often possible to relax this requirement as well. Then following the same techniques as OSGM-R [14], we can define ratio feedback $r_{x,z}(P) := \frac{\varphi_\mu(x^+(P), z^+(P))}{\varphi_\mu(x, z)}$, where $x^+(P)$ and $z^+(P)$ are the iterates after one iteration of **Algorithm 6**. Define $\hat{\kappa}$ such that

$$1 - \frac{1}{\sqrt{\hat{\kappa}}} = \min_P \max_{x,z} r_{x,z}(P),$$

where $\hat{\kappa} \leq \kappa$ since $r_{x,z}(\frac{1}{L}) \leq 1 - \frac{1}{\sqrt{\kappa}}$ for any x, z . Then an $\mathcal{O}(\sqrt{\hat{\kappa}} \log(1/\varepsilon))$ complexity can be obtained. However, whether it is possible to design an algorithm that achieves $\mathcal{O}(\sqrt{\kappa^*} \log(1/\varepsilon))$ with only blackbox access to f remains unclear.

Feedback design through performance estimation. OSGM provides a general mechanism to accelerate iterative methods when a computable potential (Lyapunov) function is known. Although such potential functions can often be obtained from literature (for example, **Section 2** demonstrates this approach in the case of heavy-ball momentum), it can be challenging to find a computable potential function that does not rely on x^* .

Question 4. *Is there a systematic way to design computable potential functions for iterative methods (that do not rely on the exact solution x^*)?*

Recent advances in first-order methods have shown that finding a potential function can be reduced to a small-scale semidefinite optimization problem through performance estimation (PEP) [34, 35]. The advantage of the PEP framework is that it can numerically find a tight worst-case potential function for some problem class (e.g., for some given smoothness constant L and strongly convexity constant μ). If PEP can find a potential

function that does not depend on x^* , then **OSGM** can be directly applied since it only requires a numerically computable potential function. It is appealing to accelerate first-order methods in practice by applying **OSGM** to the potential functions generated by performance estimation.

5.4 Convex composite optimization and proximal gradient

OSGM works primarily in the smooth unconstrained setting. **Section 3** shows one way to extend **OSGM** to smooth nonconvex optimization. Here, we discuss challenges and opportunities for **OSGM** in nonsmooth optimization.

Question 5. *Can **OSGM** be used to accelerate the solution of convex composite optimization or general nonsmooth convex optimization problems?*

For example, consider the convex composite optimization problem $\min_x \varphi(x) := f(x) + w(x)$, where f is smooth convex and w is possibly nonsmooth but prox-friendly: that is, for any $\gamma > 0$, the proximal mapping

$$\text{prox}_{w/\gamma}(x) := \arg \min_z \{w(z) + \frac{\gamma}{2}\|z - x\|^2\}$$

can be efficiently evaluated. The operator $\mathcal{G}_\gamma(x) := \gamma(x - \text{prox}_{w/\gamma}(x - \frac{1}{\gamma}\nabla f(x)))$, known as the gradient map, generalizes gradient to this composite setting, and the parameter $\gamma > 0$ generalizes the stepsize. The proximal gradient method updates the iterate as

$$x^{k+1} = x^k - \frac{1}{\gamma}\mathcal{G}_\gamma(x^k).$$

One natural idea to accelerate convergence is to adapt the parameter γ using **OSGM**. However, a basic implementation of this idea does not work, as the ratio or hypergradient feedback is not a convex function of γ (or $1/\gamma$) due to the complex dependence from $\mathcal{G}_\gamma(x)$. A partial solution in this context separates the stepsize multiplying the gradient map from the parameter inside the map: fix $\gamma = L$ in the gradient map and consider the update

$$x^{k+1} = x^k - P_k \mathcal{G}_L(x^k)$$

with stepsize P_k . This iteration directly corresponds to gradient descent. We can show an improved convergence rate in the following two settings:

Case 1. If $f(x) = 0$, then $\mathcal{G}_L(x)$ is the gradient of the Moreau envelope $\varphi^{1/L}(x) = w^{1/L}(x) := \min_z \{w(z) + \frac{L}{2}\|z - x\|^2\}$, which is smooth and convex. The update is essentially gradient descent on the envelope $w^{1/L}$, and all the **OSGM** results apply. If w is not prox-friendly, we can solve the proximal subproblem inexactly with an additional subroutine and accelerate gradient descent on the envelope.

Case 2. When f is a general smooth convex function, we can define the proximal hypergradient feedback

$$h_x(P) := \frac{\varphi(x - P\mathcal{G}_L(x)) - \varphi(x)}{\|\mathcal{G}_L(x)\|^2} = \frac{f(x - P\mathcal{G}_L(x)) + w(x - P\mathcal{G}_L(x)) - [f(x) + w(x)]}{\|\mathcal{G}_L(x)\|^2}.$$

Since w is nonsmooth, $h_x(P)$ will not be Lipschitz continuous when $\|\mathcal{G}_L(x)\| \rightarrow 0$ and online subgradient method cannot guarantee sublinear regret. To address this issue, we consider online proximal gradient update

$$P_{k+1} = \arg \min_P \left\{ \frac{f(x^k - P_k \mathcal{G}_L(x^k)) + \langle \nabla f(x^k - P_k \mathcal{G}_L(x^k)) \mathcal{G}_L(x^k)^\top, P - P_k \rangle}{\|\mathcal{G}_L(x^k)\|^2} + \frac{w(x^k - P \mathcal{G}_L(x^k))}{\|\mathcal{G}_L(x^k)\|^2} + \frac{1}{2\eta} \|P - P_k\|_F^2 \right\},$$

where $h_x(P)$ has been partially linearized. This subproblem is a proximal update in P with respect to the (scaled, translated) function $\tilde{w}(P) = w(x - P\mathcal{G}_L(x))$. This function often, but not always, inherits from w a proximal mapping that can be efficiently evaluated. If P is a diagonal matrix and w is separable in the coordinates, the proximal operator of \tilde{w} can be evaluated with the same complexity as the proximal operator of w . Combined with the monotone landscape action, we can obtain a convergence guarantee in this case in terms of $\|\mathcal{G}_L(x)\|$: $\min_{1 \leq k \leq K} \|\mathcal{G}_L(x^k)\|^2 = \mathcal{O}(\frac{1}{K})$.

Theorem 5.2 (Informal). *Suppose f is L -smooth convex and w is convex. Then given any benchmark stepsize $\hat{P} \in \mathcal{P}$, a variant of **OSGM** achieves*

$$\min_{1 \leq k \leq K} \|\mathcal{G}_L(x^k)\|^2 \leq \frac{\varphi(x^1) - \varphi(x^*)}{K} \frac{1}{\frac{1}{K} \max\{\sum_{k=1}^K h_{x^k}(\hat{P}) - \frac{L}{2} \|P_1 - \hat{P}\|_F^2, 0\}}.$$

In particular, taking $P_1 = \hat{P} = \frac{1}{L}I$ yields $\min_{1 \leq k \leq K} \|\mathcal{G}_L(x^k)\|^2 \leq \frac{2L[\varphi(x^1) - \varphi^]}{K}$.*

A detailed analysis is available in the appendix.

6 Conclusions

The second part of our paper explores the practical aspects of **OSGM** and outlines several promising directions for future research. We design and analyze **OSGM-Best**, a practical variant of **OSGM** that integrates insights into algorithmic behavior with heavy-ball momentum. We also extend the **OSGM** framework to smooth nonconvex optimization problems. Numerical experiments demonstrate that **OSGM** achieves competitive performance across a range of real-world optimization tasks.

References

- [1] Jonathan Barzilai and Jonathan M Borwein. Two-point step size gradient methods. *IMA journal of numerical analysis*, 8(1):141–148, 1988. (cited on 2, 15)
- [2] Atilim Gunes Baydin, Robert Cornish, David Martinez Rubio, Mark Schmidt, and Frank Wood. Online learning rate adaptation with hypergradient descent. In *International Conference on Learning Representations*, 2018. (cited on 2, 3, 5, 7)
- [3] Amir Beck. *First-order methods in optimization*. SIAM, 2017. (cited on 2, 40)
- [4] Xinyi Chen and Elad Hazan. Open problem: Black-box reductions & adaptive gradient methods. *Proceedings of Machine Learning Research vol*, 196:1–8, 2024. (cited on 10)
- [5] Ya-Chi Chu, Wenzhi Gao, Yinyu Ye, and Madeleine Udell. Provable and practical online learning rate adaptation with hypergradient descent. *arXiv preprint arXiv:2502.11229*, 2025. (cited on 1, 5)
- [6] Yu-Hong Dai and XQ Yang. A new gradient method with an optimal stepsize property. *Computational optimization and applications*, 33(1):73–88, 2006. (cited on 7)
- [7] Marina Danilova, Anastasiia Kulakova, and Boris Polyak. Non-monotone behavior of the heavy ball method. In *Difference Equations and Discrete Dynamical Systems with Applications: 24th ICDEA, Dresden, Germany, May 21–25, 2018*, pages 213–230. Springer, 2020. (cited on 3, 7, 27)
- [8] Alexandre d’Aspremont, Damien Scieur, Adrien Taylor, et al. Acceleration methods. *Foundations and Trends® in Optimization*, 5(1-2):1–245, 2021. (cited on 2, 12, 17, 39)
- [9] Roberta De Asmundis, Daniela di Serafino, Filippo Riccio, and Gerardo Toraldo. On spectral properties of steepest descent methods. *IMA Journal of Numerical Analysis*, 33(4):1416–1435, 2013. (cited on 7)
- [10] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011. (cited on 13)
- [11] Roger Fletcher. On the barzilai-borwein method. In *Optimization and control with applications*, pages 235–256. Springer, 2005. (cited on 3)
- [12] Ana Friedlander, José Mario Martínez, Brigida Molina, and Marcus Raydan. Gradient method with retards and generalizations. *SIAM Journal on Numerical Analysis*, 36(1):275–289, 1998. (cited on 16)

- [13] Wenzhi Gao, Ya-Chi Chu, Yinyu Ye, and Madeleine Udell. Gradient methods with online scaling. *arXiv preprint arXiv:2411.01803*, 2024. (cited on 1, 15, 33)
- [14] Wenzhi Gao, Ya-Chi Chu, Yinyu Ye, and Madeleine Udell. Gradient methods with online scaling part i. theoretical foundations. *arXiv preprint arXiv:2505.23081*, 2025. (cited on 1, 2, 3, 9, 10, 17, 35, 36)
- [15] Wenzhi Gao, Zhaonan Qu, Madeleine Udell, and Yinyu Ye. Scalable approximate optimal diagonal preconditioning. *arXiv preprint arXiv:2312.15594*, 2023. (cited on 2)
- [16] Nicholas IM Gould, Dominique Orban, and Philippe L Toint. Cutest: a constrained and unconstrained testing environment with safe threads for mathematical optimization. *Computational optimization and applications*, 60:545–557, 2015. (cited on 12)
- [17] Robert M Gower, Aaron Defazio, and Michael Rabbat. Stochastic polyak stepsize with a moving target. *arXiv preprint arXiv:2106.11851*, 2021. (cited on 16)
- [18] Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016. (cited on 2)
- [19] Elad Hazan and Sham Kakade. Revisiting the polyak step size. *arXiv preprint arXiv:1905.00313*, 2019. (cited on 16)
- [20] Yakui Huang, Yu-Hong Dai, Xin-Wei Liu, and Hongchao Zhang. Gradient methods exploiting spectral properties. *Optimization Methods and Software*, 35(4):681–705, 2020. (cited on 7)
- [21] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*, pages 795–811. Springer, 2016. (cited on 32)
- [22] Sehun Kim, Hyunsil Ahn, and Seong-Cheol Cho. Variable target value subgradient method. *Mathematical Programming*, 49:359–369, 1990. (cited on 16)
- [23] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. (cited on 7, 13)
- [24] Brian Kulis and Peter L Bartlett. Implicit online learning. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 575–582, 2010. (cited on 15)
- [25] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989. (cited on 2)
- [26] Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 1999. (cited on 13)
- [27] Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019. (cited on 2)
- [28] Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994. (cited on 11)
- [29] Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *Ussr computational mathematics and mathematical physics*, 4(5):1–17, 1964. (cited on 7, 12)
- [30] Boris T Polyak. Introduction to optimization. 1987. (cited on 7)
- [31] Zhaonan Qu, Wenzhi Gao, Oliver Hinder, Yinyu Ye, and Zhengyuan Zhou. Optimal diagonal preconditioning. *Operations Research*, 2024. (cited on 2)

- [32] David Martinez Rubio. Convergence analysis of an adaptive method of gradient descent. *University of Oxford, Oxford, M. Sc. thesis*, 2017. (cited on 3, 5)
- [33] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(1), 2013. (cited on 16)
- [34] Adrien Taylor, Bryan Van Scoy, and Laurent Lessard. Lyapunov functions for first-order methods: Tight automated convergence guarantees. In *International Conference on Machine Learning*, pages 4897–4906. PMLR, 2018. (cited on 17)
- [35] Manu Upadhyaya, Sebastian Banert, Adrien B Taylor, and Pontus Giselsson. Automated tight Lyapunov analysis for first-order methods. *Mathematical Programming*, 209(1):133–170, 2025. (cited on 2, 17)
- [36] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020. (cited on 13)
- [37] Hongchao Zhang and William W Hager. A nonmonotone line search technique and its application to unconstrained optimization. *SIAM journal on Optimization*, 14(4):1043–1056, 2004. (cited on 16)
- [38] Qinnmeng Zou and Frédéric Magoulès. Fast gradient methods with alignment for symmetric linear systems without using cauchy step. *Journal of Computational and Applied Mathematics*, 381:113033, 2021. (cited on 7)
- [39] Qinnmeng Zou and Frédéric Magoulès. Delayed gradient methods for symmetric and positive definite linear systems. *SIAM Review*, 64(3):517–553, 2022. (cited on 7, 15, 16)

Appendix

Table of Contents

A	Proof of results in Section 2	23
A.1	Proof of Lemma 2.1	23
A.2	Proof of Theorem 2.1	23
A.3	Proof of Lemma 2.3	27
A.4	Proof of Lemma 2.4	28
A.5	Proof of Theorem 2.2	29
A.6	Proof of Theorem 2.3	30
B	Proof of results in Section 3	32
B.1	Proof of Lemma 3.1	32
B.2	Proof of Lemma 3.2	33
B.3	Proof of Theorem 3.1	33
B.4	Proof of Lemma 3.3	34
B.5	Proof of Theorem 3.2	34
C	Proof of results in Section 5	35
C.1	Details of Subsection 5.1	35
C.2	Proof of Theorem 5.1	37
C.3	Details of Subsection 5.4	39

A Proof of results in Section 2

A.1 Proof of Lemma 2.1

The assumption $|x_1^k| \geq \sqrt{2}\kappa^{3/2}|x_2^k|$ and that $\kappa \geq 2$ implies

$$(x_1^k)^2 \geq 2\kappa^3(x_2^k)^2 \geq 2\kappa^2(\kappa - 2)(x_2^k)^2. \quad (17)$$

Then we deduce that

$$\begin{aligned} \frac{\alpha_{k+1} - \alpha_k}{\eta} &= \frac{(x_1^k)^2(1 - \alpha_k) + \kappa^2(1 - \kappa\alpha_k)(x_2^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2} && \text{(by (3))} \\ &\geq \frac{1}{4} \frac{2(x_1^k)^2 - 2\kappa^2(\kappa - 2)(x_2^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2} && \text{(by } \alpha_k \leq \frac{1}{2} \text{)} \\ &\geq \frac{1}{4} \frac{(x_1^k)^2}{(x_1^k)^2 + \kappa^2(x_2^k)^2} \geq \frac{1}{4}. && \text{(by (17))} \end{aligned}$$

This completes the proof.

A.2 Proof of Theorem 2.1

The proof of **Theorem 2.1** requires several lemmas. The proof of each lemma is deferred to the the end of the section.

Lemma A.1. *Let $J_{HDM}(\alpha, \hat{z})$ and $J_{OSGM}(\alpha, \hat{z})$ denote the Jacobians of the dynamical systems \mathcal{F}_{HDM} and \mathcal{F}_{OSGM} respectively with $\eta = \frac{1}{L}$. Then each of $J_{HDM}(\alpha^*, \hat{z}_i^*)$ and $J_{OSGM}(\alpha^*, \hat{z}_i^*)$, $i = 1, 2$, is similar to the block diagonal matrix by a common permutation matrix:*

$$J_{HDM}(\alpha^*, \hat{z}_i^*) \sim \text{Diag}(1 - \alpha^* \lambda_2, \dots, 1 - \alpha^* \lambda_{n-1}, J_{HDM}^i) \quad \text{and} \quad J_{OSGM}(\alpha^*, \hat{z}_i^*) \sim \text{Diag}(1 - \alpha^* \lambda_2, \dots, 1 - \alpha^* \lambda_{n-1}, J_{OSGM}^i),$$

where J_{HDM}^i and J_{OSGM}^i , $i = 1, 2$, are 3×3 matrices defined by

$$\begin{aligned} J_{HDM}^1 &:= D \begin{pmatrix} \frac{\kappa-1}{2\kappa} & -\sqrt{\kappa^2+1} \pm \frac{\sqrt{\kappa^2+1}}{\kappa} \\ \frac{-\kappa(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{\kappa}{\kappa^2+1} & \frac{\pm 1}{\kappa^2+1} \\ \frac{\mp(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{\mp 1}{\kappa^2+1} & \frac{-1}{\kappa(\kappa^2+1)} \end{pmatrix} D^{-1}, \quad J_{HDM}^2 := D \begin{pmatrix} \frac{\kappa-1}{2\kappa} & \sqrt{\kappa^2+1} \pm \frac{\sqrt{\kappa^2+1}}{\kappa} \\ \frac{\kappa(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{\kappa}{\kappa^2+1} & \frac{\mp 1}{\kappa^2+1} \\ \frac{\mp(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{\pm 1}{\kappa^2+1} & \frac{-1}{\kappa(\kappa^2+1)} \end{pmatrix} D^{-1}, \\ J_{OSGM}^1 &:= D \begin{pmatrix} \frac{\kappa-1}{2\kappa} & -\sqrt{\kappa^2+1} \pm \frac{\sqrt{\kappa^2+1}}{\kappa} \\ \frac{-\kappa(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{-\kappa^2}{\kappa^2+1} & \frac{\pm \kappa}{\kappa^2+1} \\ \frac{\mp(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{\mp \kappa}{\kappa^2+1} & \frac{1}{\kappa^2+1} \end{pmatrix} D^{-1}, \quad J_{OSGM}^2 := D \begin{pmatrix} \frac{\kappa-1}{2\kappa} & \sqrt{\kappa^2+1} \pm \frac{\sqrt{\kappa^2+1}}{\kappa} \\ \frac{\kappa(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{-\kappa^2}{\kappa^2+1} & \frac{\mp \kappa}{\kappa^2+1} \\ \frac{\mp(\kappa+1)}{(\kappa^2+1)^{3/2}} & \frac{\pm \kappa}{\kappa^2+1} & \frac{1}{\kappa^2+1} \end{pmatrix} D^{-1}, \end{aligned}$$

and $D = \text{Diag}(\frac{\kappa-1}{L(\kappa+1)}, 1, 1)$ is a 3×3 invertible diagonal matrix.

The stability of the period-2 orbit is determined by the spectral radius of the products $J_{HDM}(\alpha^*, \hat{z}_1^*) J_{HDM}(\alpha^*, \hat{z}_2^*)$ and $J_{OSGM}(\alpha^*, \hat{z}_1^*) J_{OSGM}(\alpha^*, \hat{z}_2^*)$, which equals the maximum absolute value of the eigenvalues. By **Lemma A.1**, the spectrum of these two matrices are

$$\begin{aligned} \sigma(J_{HDM}(\alpha^*, \hat{z}_1^*) J_{HDM}(\alpha^*, \hat{z}_2^*)) &= \{(1 - \alpha^* \lambda_2)^2, \dots, (1 - \alpha^* \lambda_{n-1})^2\} \cup \sigma(J_{HDM}^1 J_{HDM}^2); \\ \sigma(J_{OSGM}(\alpha^*, \hat{z}_1^*) J_{OSGM}(\alpha^*, \hat{z}_2^*)) &= \{(1 - \alpha^* \lambda_2)^2, \dots, (1 - \alpha^* \lambda_{n-1})^2\} \cup \sigma(J_{OSGM}^1 J_{OSGM}^2). \end{aligned}$$

Since the eigenvalues of A are distinct, i.e., $L = \lambda_1 > \lambda_2 > \dots > \lambda_{n-1} > \lambda_n = \mu$, and $\alpha^* = \frac{2}{L+\mu}$, the first $n-2$ eigenvalues satisfy

$$-\frac{\kappa-1}{\kappa+1} = 1 - \alpha^* L < 1 - \alpha^* \lambda_i < 1 - \alpha^* \mu = \frac{\kappa-1}{\kappa+1} \quad \text{for } i = 2, \dots, n-1.$$

Then $(1 - \alpha^* \lambda_i)^2 < 1$ for all $i = 2, \dots, n-1$ and the stability of the orbit $\{(\alpha^*, \hat{z}_1^*), (\alpha^*, \hat{z}_2^*)\}$ is determined by the spectral radius of $J_{HDM}^1 J_{HDM}^2$ and $J_{OSGM}^1 J_{OSGM}^2$ for the corresponding system.

Lemma A.2. The 3×3 matrices $J_{\text{HDM}}^1 J_{\text{HDM}}^2$ and $J_{\text{OSGM}}^1 J_{\text{OSGM}}^2$ have spectral radius

$$\rho(J_{\text{HDM}}^1 J_{\text{HDM}}^2) = \max \left\{ -\frac{3\kappa^3 - 7\kappa - 9}{8(\kappa+1)\kappa^2} \pm \left[\frac{\sqrt{-(\kappa-3)^2 \kappa^2 (\kappa+1)^3 (7\kappa-9)(\kappa^2+1)^5}}{8(\kappa^2+1)^{5/2} (\kappa+1)\kappa^3} - \frac{5}{8(\kappa+1)} \right] \right\} \leq 1$$

$$\rho(J_{\text{OSGM}}^1 J_{\text{OSGM}}^2) = \frac{3\kappa+1}{2\kappa} > 1$$

Observe that the spectral radius $\rho(J_{\text{HDM}}^1 J_{\text{HDM}}^2)$ asymptotically increases to $\frac{1}{2}$ when $\kappa \rightarrow \infty$, implying that the eigenvalues of the 3×3 submatrix $J_{\text{HDM}}^1 J_{\text{HDM}}^2$ are not active for the spectral radius of the whole Jacobian when κ is large. In this case, the spectral radius of $J_{\text{HDM}}(\alpha^*, \hat{z}_1^*) J_{\text{HDM}}(\alpha^*, \hat{z}_2^*)$ is the maximum of $(1 - \alpha^* \lambda_2)^2$ and $(1 - \alpha^* \lambda_{n-1})^2$.

Remark 3. While our result is established for $\eta = \frac{1}{L}$, the same quantitative result holds for $\eta = \frac{\theta}{L}, \theta \in (0, 1]$.

A.2.1 Proof of Lemma 2.2

To show the orbit, we simply plug in (α^*, \hat{z}_1^*) and (α^*, \hat{z}_2^*) into the updates of both **HDM** and **OSGM**. Observe that $L^2 \gamma^2 = \mu^2 \delta^2$ and the eigenvectors q_1 and q_n are orthogonal to each other. Then

$$\langle \hat{z}_1^*, \Lambda^3 \hat{z}_1^* \rangle = \langle \gamma e_1 + \delta e_n, L^3 \gamma e_1 + \mu^3 \delta e_n \rangle = L^3 \gamma^2 + \mu^3 \delta^2 = (L + \mu) L^2 \gamma^2, \quad (18)$$

$$\langle \hat{z}_2^*, \Lambda^3 \hat{z}_2^* \rangle = \langle -\gamma e_1 + \delta e_n, -L^3 \gamma e_1 + \mu^3 \delta e_n \rangle = L^3 \gamma^2 + \mu^3 \delta^2 = (L + \mu) L^2 \gamma^2, \quad (19)$$

$$\langle \hat{z}_1^*, \Lambda^2 \hat{z}_1^* \rangle = \langle \gamma e_1 + \delta e_n, L^2 \gamma e_1 + \mu^2 \delta e_n \rangle = L^2 \gamma^2 + \mu^2 \delta^2 = 2L^2 \gamma^2, \quad (20)$$

$$\langle \hat{z}_2^*, \Lambda^2 \hat{z}_2^* \rangle = \langle -\gamma e_1 + \delta e_n, -L^2 \gamma e_1 + \mu^2 \delta e_n \rangle = L^2 \gamma^2 + \mu^2 \delta^2 = 2L^2 \gamma^2, \quad (21)$$

and hence

$$\frac{\langle \hat{z}_1^*, \Lambda^3 \hat{z}_1^* \rangle}{\langle \hat{z}_1^*, \Lambda^2 \hat{z}_1^* \rangle} = \frac{\langle \hat{z}_2^*, \Lambda^3 \hat{z}_2^* \rangle}{\langle \hat{z}_2^*, \Lambda^2 \hat{z}_2^* \rangle} = \frac{L + \mu}{2}. \quad (22)$$

The α -update of both systems \mathcal{F}_{HDM} and $\mathcal{F}_{\text{OSGM}}$ are the same. Starting from either state (α^*, \hat{z}_1^*) and (α^*, \hat{z}_2^*) , the new state of stepsize is

$$[\alpha^*]^+(\hat{z}_i^*) = \alpha^* + \eta - \alpha^* \eta \frac{\langle \hat{z}_i^*, \Lambda^3 \hat{z}_i^* \rangle}{\langle \hat{z}_i^*, \Lambda^2 \hat{z}_i^* \rangle} = \alpha^* + \eta - \alpha^* \eta \frac{L + \mu}{2} = \alpha^*. \quad (23)$$

Since the new state of stepsize remains unchanged to be $\alpha^* = \frac{L + \mu}{2}$, both systems end up with the same new state:

$$\mathcal{F}_{\text{HDM}}(\alpha^*, \hat{z}_i^*) = ([\alpha^*]^+(\hat{z}_i^*), \frac{(I - [\alpha^*]^+(\hat{z}_i^*)\Lambda)\hat{z}_i^*}{\|(I - [\alpha^*]^+(\hat{z}_i^*)\Lambda)\hat{z}_i^*\|}) = (\alpha^*, \frac{(I - \alpha^*\Lambda)\hat{z}_i^*}{\|(I - \alpha^*\Lambda)\hat{z}_i^*\|});$$

$$\mathcal{F}_{\text{OSGM}}(\alpha^*, \hat{z}_i^*) = ([\alpha^*]^+(\hat{z}_i^*), \frac{(I - \alpha^*\Lambda)\hat{z}_i^*}{\|(I - \alpha^*\Lambda)\hat{z}_i^*\|}) = (\alpha^*, \frac{(I - \alpha^*\Lambda)\hat{z}_i^*}{\|(I - \alpha^*\Lambda)\hat{z}_i^*\|}).$$

The updates of direction before the normalization are

$$(I - \alpha^*\Lambda)\hat{z}_1^* = (1 - \frac{2L}{L + \mu})\gamma e_1 + (1 - \frac{2\mu}{L + \mu})\delta e_n = \frac{L - \mu}{L + \mu}(-\gamma e_1 + \delta e_n) = \frac{\kappa - 1}{\kappa + 1}\hat{z}_2^*; \quad (24)$$

$$(I - \alpha^*\Lambda)\hat{z}_2^* = -(1 - \frac{2L}{L + \mu})\gamma e_1 + (1 - \frac{2\mu}{L + \mu})\delta e_n = \frac{L - \mu}{L + \mu}(\gamma e_1 + \delta e_n) = \frac{\kappa - 1}{\kappa + 1}\hat{z}_1^*. \quad (25)$$

Since \hat{x}_1^* and \hat{x}_2^* both have unit norm, after the normalization, we have

$$\mathcal{F}_{\text{HDM}}(\alpha^*, \hat{z}_1^*) = \mathcal{F}_{\text{OSGM}}(\alpha^*, \hat{z}_1^*) = (\alpha^*, \hat{z}_2^*) \quad \text{and} \quad \mathcal{F}_{\text{HDM}}(\alpha^*, \hat{z}_2^*) = \mathcal{F}_{\text{OSGM}}(\alpha^*, \hat{z}_2^*) = (\alpha^*, \hat{z}_1^*).$$

This proves that $\{(\alpha^*, \hat{z}_1^*), (\alpha^*, \hat{z}_2^*)\}$ is an orbit with period 2.

A.2.2 Proof of Lemma A.1

The Jacobians of \mathcal{F}_{HDM} and $\mathcal{F}_{\text{OSGM}}$ are

$$J_{\text{HDM}}(\alpha, \hat{z}) = \begin{pmatrix} 1 - \eta \frac{\langle \hat{z}, \Lambda^3 \hat{z} \rangle}{\langle \hat{z}, \Lambda^2 \hat{z} \rangle} & v(\alpha, \hat{z})^\top \\ \frac{-1}{\|u(\alpha, \hat{z})\|} \Pi(\alpha, \hat{z}) \left(1 - \eta \frac{\langle \hat{z}, \Lambda^3 \hat{z} \rangle}{\langle \hat{z}, \Lambda^2 \hat{z} \rangle} \right) \Lambda \hat{z} & \frac{1}{\|u(\alpha, \hat{z})\|} \Pi(\alpha, \hat{z}) [(I - \alpha^+(\hat{z})\Lambda) - \Lambda \hat{z} v(\alpha, \hat{z})^\top] \end{pmatrix},$$

$$J_{\text{OSGM}}(\alpha, \hat{z}) = \begin{pmatrix} 1 - \eta \frac{\langle \hat{z}, \Lambda^3 \hat{z} \rangle}{\langle \hat{z}, \Lambda^2 \hat{z} \rangle} & v(\alpha, \hat{z})^\top \\ \frac{-1}{\|u(\alpha, \hat{z})\|} \Pi(\alpha, \hat{z}) \Lambda \hat{z} & \frac{1}{\|u(\alpha, \hat{z})\|} \Pi(\alpha, \hat{z}) (I - \alpha \Lambda) \end{pmatrix},$$

where $\Pi(\alpha, \hat{z}) := I - \frac{u(\alpha, \hat{z})u(\alpha, \hat{z})^\top}{\|u(\alpha, \hat{z})\|^2}$, $u(\alpha, \hat{z}) := (I - \alpha \Lambda)\hat{z}$, and $v(\alpha, \hat{z}) := \frac{2\alpha\eta}{\langle \hat{z}, \Lambda^2 \hat{z} \rangle} \Lambda^2 (\frac{\langle \hat{z}, \Lambda^3 \hat{z} \rangle}{\langle \hat{z}, \Lambda^2 \hat{z} \rangle} I - \Lambda)\hat{z}$. To evaluate the Jacobians at (α^*, \hat{z}_1^*) and (α^*, \hat{z}_2^*) , we calculate each block one by one. Note that J_{HDM} and J_{OSGM} only differ in the last row.

(1,1)-block of $J_{\text{HDM}}(\alpha^*, \hat{z}_i^*)$ and $J_{\text{OSGM}}(\alpha^*, \hat{z}_i^*)$. Equation (22) and $\eta = \frac{1}{L}$ imply

$$1 - \eta \frac{\langle \hat{z}_i^*, \Lambda^3 \hat{z}_i^* \rangle}{\langle \hat{z}_i^*, \Lambda^2 \hat{z}_i^* \rangle} = 1 - \frac{L+\mu}{2L} = \frac{\kappa-1}{2\kappa} \quad \text{for } i = 1, 2. \quad (26)$$

(1,2)-block of $J_{\text{HDM}}(\alpha^*, \hat{z}_i^*)$ and $J_{\text{OSGM}}(\alpha^*, \hat{z}_i^*)$. Using (20)–(22) and $\eta = \frac{1}{L}$, we have

$$v(\alpha^*, \hat{z}_i^*) = \frac{2\alpha^*\eta}{\langle \hat{z}_i^*, \Lambda^2 \hat{z}_i^* \rangle} \Lambda^2 (\frac{\langle \hat{z}_i^*, \Lambda^3 \hat{z}_i^* \rangle}{\langle \hat{z}_i^*, \Lambda^2 \hat{z}_i^* \rangle} I - \Lambda) \hat{z}_i^* = \frac{1}{L^3 \gamma^2} \Lambda^2 (I - \alpha^* \Lambda) \hat{z}_i^*,$$

which, together with (24) and (25), implies

$$v(\alpha^*, \hat{z}_1^*) = \frac{\kappa-1}{L^3 \gamma^2 (\kappa+1)} \Lambda^2 \hat{z}_2^* = \frac{\kappa-1}{L^3 \gamma^2 (\kappa+1)} [-L^2 \gamma e_1 + \mu^2 \delta e_n] = -\frac{(\kappa-1)\sqrt{\kappa^2+1}}{L(\kappa+1)} e_1 + \frac{\pm(\kappa-1)\sqrt{\kappa^2+1}}{L\kappa(\kappa+1)} e_n; \quad (27)$$

$$v(\alpha^*, \hat{z}_2^*) = \frac{\kappa-1}{L^3 \gamma^2 (\kappa+1)} \Lambda^2 \hat{z}_1^* = \frac{\kappa-1}{L^3 \gamma^2 (\kappa+1)} [L^2 \gamma e_1 + \mu^2 \delta e_n] = \frac{(\kappa-1)\sqrt{\kappa^2+1}}{L(\kappa+1)} e_1 + \frac{\pm(\kappa-1)\sqrt{\kappa^2+1}}{L\kappa(\kappa+1)} e_n. \quad (28)$$

(2,1)-block of $J_{\text{OSGM}}(\alpha^*, \hat{z}_i^*)$. Equations (24)–(25) imply

$$u(\alpha^*, \hat{z}_1^*) = (I - \alpha^* \Lambda) \hat{z}_1^* = \frac{\kappa-1}{\kappa+1} \hat{z}_2^*, \quad u(\alpha^*, \hat{z}_2^*) = (I - \alpha^* \Lambda) \hat{z}_2^* = \frac{\kappa-1}{\kappa+1} \hat{z}_1^*, \quad \|u(\alpha^*, \hat{z}_1^*)\| = \|u(\alpha^*, \hat{z}_2^*)\| = \frac{\kappa-1}{\kappa+1}, \quad (29)$$

and hence

$$I - \frac{u(\alpha^*, \hat{z}_1^*)u(\alpha^*, \hat{z}_1^*)^\top}{\|u(\alpha^*, \hat{z}_1^*)\|^2} = I - \hat{z}_2^*(\hat{z}_2^*)^\top \quad \text{and} \quad I - \frac{u(\alpha^*, \hat{z}_2^*)u(\alpha^*, \hat{z}_2^*)^\top}{\|u(\alpha^*, \hat{z}_2^*)\|^2} = I - \hat{z}_1^*(\hat{z}_1^*)^\top. \quad (30)$$

Observe that $(\hat{z}_2^*)^\top A \hat{z}_1^* = (\hat{z}_1^*)^\top A \hat{z}_2^* = -L\gamma^2 + \mu\delta^2$, together with (29), we have

$$\begin{aligned} \frac{-1}{\|u(\alpha^*, \hat{z}_1^*)\|} \Pi(\alpha^*, \hat{z}_1^*) \Lambda \hat{z}_1^* &= \frac{-1}{\|u(\alpha^*, \hat{z}_1^*)\|} (I - \frac{u(\alpha^*, \hat{z}_1^*)u(\alpha^*, \hat{z}_1^*)^\top}{\|u(\alpha^*, \hat{z}_1^*)\|^2}) \Lambda \hat{z}_1^* \\ &= -\frac{\kappa+1}{\kappa-1} (\Lambda \hat{z}_1^* - \hat{z}_2^*(\hat{z}_2^*)^\top \Lambda \hat{z}_1^*) \\ &= -\frac{\kappa+1}{\kappa-1} [\Lambda \hat{z}_1^* - (-L\gamma^2 + \mu\delta^2) \hat{z}_2^*] \\ &= -\frac{\kappa+1}{\kappa-1} [(L\gamma e_1 + \mu\delta e_n) - (-L\gamma^2 + \mu\delta^2)(-\gamma e_1 + \delta e_n)] \\ &= -\frac{\kappa+1}{\kappa-1} [(L\gamma - L\gamma^3 + \mu\gamma\delta^2)e_1 + (\mu\delta + L\gamma^2\delta - \mu\delta^3)e_n] \\ &= \frac{-L\kappa(\kappa+1)^2}{(\kappa-1)(\kappa^2+1)^{3/2}} e_1 \mp \frac{L(\kappa+1)^2}{(\kappa-1)(\kappa^2+1)^{3/2}} e_n \end{aligned} \quad (31)$$

$$\begin{aligned}
\frac{-1}{\|u(\alpha^*, \hat{z}_1^*)\|} \Pi(\alpha^*, \hat{z}_2^*) \Lambda \hat{z}_2^* &= \frac{-1}{\|u(\alpha^*, \hat{z}_2^*)\|} \left(I - \frac{u(\alpha^*, \hat{z}_2^*) u(\alpha^*, \hat{z}_2^*)^\top}{\|u(\alpha^*, \hat{z}_2^*)\|^2} \right) \Lambda \hat{z}_2^* \\
&= -\frac{\kappa+1}{\kappa-1} (\Lambda \hat{z}_2^* - \hat{z}_1^* (\hat{z}_1^*)^\top \Lambda \hat{z}_2^*) \\
&= -\frac{\kappa+1}{\kappa-1} [\Lambda \hat{z}_2^* - (-L\gamma^2 + \mu\delta^2) \hat{z}_1^*] \\
&= -\frac{\kappa+1}{\kappa-1} [(-L\gamma e_1 + \mu\delta e_n) - (-L\gamma^2 + \mu\delta^2)(\gamma e_1 + \delta e_n)] \\
&= -\frac{\kappa+1}{\kappa-1} [(-L\gamma + L\gamma^3 - \mu\gamma\delta^2)e_1 + (\mu\delta + L\gamma^2\delta - \mu\delta^3)e_n] \\
&= \frac{L\kappa(\kappa+1)^2}{(\kappa-1)(\kappa^2+1)^{3/2}} e_1 \mp \frac{L(\kappa+1)^2}{(\kappa-1)(\kappa^2+1)^{3/2}} e_n. \tag{32}
\end{aligned}$$

(2, 1)-block of $J_{\text{HDM}}(\alpha^*, \hat{z}_i^*)$. Note that the (2, 1)-block of $J_{\text{HDM}}(\alpha^*, \hat{z}_i^*)$ is a scalar multiple of the (2, 1)-block of $J_{\text{OSGM}}(\alpha^*, \hat{z}_i^*)$ with the scalar $1 - \eta_{\langle \hat{z}_i^*, \Lambda^3 \hat{z}_i^* \rangle}^{\langle \hat{z}_i^*, \Lambda^2 \hat{z}_i^* \rangle}$. Then using (26) and (31)–(32), we have

$$\begin{aligned}
\frac{-1}{\|u(\alpha^*, \hat{z}_1^*)\|} \Pi(\alpha^*, \hat{z}_1^*) (1 - \eta_{\langle \hat{z}_1^*, \Lambda^3 \hat{z}_1^* \rangle}^{\langle \hat{z}_1^*, \Lambda^2 \hat{z}_1^* \rangle}) \Lambda \hat{z}_1^* &= \frac{-L\kappa(\kappa+1)^2}{2\kappa(\kappa^2+1)^{3/2}} e_1 \mp \frac{L(\kappa+1)^2}{2\kappa(\kappa^2+1)^{3/2}} e_n; \\
\frac{-1}{\|u(\alpha^*, \hat{z}_2^*)\|} \Pi(\alpha^*, \hat{z}_2^*) (1 - \eta_{\langle \hat{z}_2^*, \Lambda^3 \hat{z}_2^* \rangle}^{\langle \hat{z}_2^*, \Lambda^2 \hat{z}_2^* \rangle}) \Lambda \hat{z}_2^* &= \frac{L\kappa(\kappa+1)^2}{2\kappa(\kappa^2+1)^{3/2}} e_1 \mp \frac{L(\kappa+1)^2}{2\kappa(\kappa^2+1)^{3/2}} e_n.
\end{aligned}$$

(2, 2)-block of $J_{\text{OSGM}}(\alpha^*, \hat{z}_i^*)$. Let $I_{1,n} := [e_1, \dots, e_n] \in \mathbb{R}^{n \times 2}$ denote the first and last column of an $n \times n$ identity matrix.

Using (29) and (30), we have

$$\begin{aligned}
\frac{1}{\|u(\alpha^*, \hat{z}_1^*)\|} \Pi(\alpha^*, \hat{z}_1^*) (I - \alpha^* \Lambda) &= \frac{\kappa+1}{\kappa-1} (I - \hat{z}_2^* (\hat{z}_2^*)^\top) (I - \alpha^* \Lambda) = \frac{\kappa+1}{\kappa-1} (I - \alpha^* \Lambda) - \frac{\kappa+1}{\kappa-1} \hat{z}_2^* [(I - \alpha^* \Lambda) \hat{z}_2^*]^\top \\
&= \frac{\kappa+1}{\kappa-1} (I - \alpha^* \Lambda) - \hat{z}_2^* (\hat{z}_1^*)^\top \\
&= \frac{\kappa+1}{\kappa-1} (I - \alpha^* \Lambda) - I_{1,n} \begin{pmatrix} -\gamma^2 & -\gamma\delta \\ \gamma\delta & \delta^2 \end{pmatrix} I_{1,n}^\top, \tag{33}
\end{aligned}$$

$$\begin{aligned}
\frac{1}{\|u(\alpha^*, \hat{z}_2^*)\|} \Pi(\alpha^*, \hat{z}_2^*) (I - \alpha^* \Lambda) &= \frac{\kappa+1}{\kappa-1} (I - \hat{z}_1^* (\hat{z}_1^*)^\top) (I - \alpha^* \Lambda) = \frac{\kappa+1}{\kappa-1} (I - \alpha^* \Lambda) - \frac{\kappa+1}{\kappa-1} \hat{z}_1^* [(I - \alpha^* \Lambda) \hat{z}_1^*]^\top \\
&= \frac{\kappa+1}{\kappa-1} (I - \alpha^* \Lambda) - \hat{z}_1^* (\hat{z}_2^*)^\top \\
&= \frac{\kappa+1}{\kappa-1} (I - \alpha^* \Lambda) - I_{1,n} \begin{pmatrix} -\gamma^2 & \gamma\delta \\ -\gamma\delta & \delta^2 \end{pmatrix} I_{1,n}^\top. \tag{34}
\end{aligned}$$

Both (33) and (34) have $1 - \alpha^* \lambda_i$ at the (i, i) -entry for $i = 2, \dots, n-1$ and the only other non-zero entries lie on the four corners indexed by $(1, 1)$, $(1, n)$, $(n, 1)$, and (n, n) .

(2, 2)-block of $J_{\text{HDM}}(\alpha^*, \hat{z}_i^*)$. Equation (23) reads as $[\alpha^*]^+(\hat{z}_i^*) = \alpha^*$ for $i = 1, 2$, and thus the (2, 2)-block of $J_{\text{HDM}}(\alpha^*, \hat{z}_i^*)$ becomes

$$\frac{1}{\|u(\alpha^*, \hat{z}_i^*)\|} \Pi(\alpha^*, \hat{z}_i^*) [(I - [\alpha^*]^+(\hat{z}_i^*) \Lambda) - \Lambda \hat{z}_i^* v(\alpha^*, \hat{z}_i^*)^\top] = \frac{1}{\|u(\alpha^*, \hat{z}_i^*)\|} \Pi(\alpha^*, \hat{z}_i^*) [(I - \alpha^* \Lambda) - \Lambda \hat{z}_i^* v(\alpha^*, \hat{z}_i^*)^\top],$$

in which is first term has been evaluated in (33) and (34). Using (27)–(28) and (31)–(32), the second term is evaluated as

$$-\frac{1}{\|u(\alpha^*, \hat{z}_1^*)\|} \Pi(\alpha^*, \hat{z}_1^*) \Lambda \hat{z}_1^* v(\alpha^*, \hat{z}_1^*)^\top = -\frac{1}{L^3 \gamma^2} I_{1,n} \begin{pmatrix} L\gamma - L\gamma^3 + \mu\gamma\delta^2 \\ \mu\delta + L\gamma^2\delta - \mu\delta^3 \end{pmatrix} \begin{pmatrix} -L^2\gamma \\ \mu^2\delta \end{pmatrix}^\top I_{1,n}^\top; \tag{35}$$

$$-\frac{1}{\|u(\alpha^*, \hat{z}_2^*)\|} \Pi(\alpha^*, \hat{z}_2^*) \Lambda \hat{z}_2^* v(\alpha^*, \hat{z}_2^*)^\top = -\frac{1}{L^3 \gamma^2} I_{1,n} \begin{pmatrix} -L\gamma + L\gamma^3 - \mu\gamma\delta^2 \\ \mu\delta + L\gamma^2\delta - \mu\delta^3 \end{pmatrix} \begin{pmatrix} L^2\gamma \\ \mu^2\delta \end{pmatrix}^\top I_{1,n}^\top, \tag{36}$$

The non-zero entries of (35) and (36) again only lie on the four corners indexed by $(1, 1)$, $(1, n)$, $(n, 1)$, and (n, n) . Hence, the Jacobians $J_{\text{HDM}}(\alpha^*, \hat{z}_i^*)$ and $J_{\text{OSGM}}(\alpha^*, \hat{z}_i^*)$ for $i = 1, 2$ all have the structure

$$\begin{pmatrix} \times & \times & 0 & \cdots & 0 & \times \\ \times & \times & 0 & \cdots & & \times \\ 0 & 0 & 1 - \alpha^* \lambda_2 & & & 0 \\ & & & \ddots & & \\ \vdots & \vdots & & & \ddots & \vdots \\ 0 & 0 & & & 1 - \alpha^* \lambda_{n-1} & 0 \\ \times & \times & 0 & \cdots & 0 & \times \end{pmatrix}.$$

We can simultaneously permute the rows and columns so that all the four Jacobians are similar to a block diagonal matrix consists of $n - 2$ individual diagonal entries $1 - \alpha^* \lambda_2, \dots, 1 - \alpha^* \lambda_{n-1}$ and one 3×3 dense matrix, denoted by J_{HDM}^i and J_{OSGM}^i for $i = 1, 2$. After careful but tedious calculations, the four 3×3 matrices match their definition in the lemma.

A.2.3 Proof of Lemma A.2

The result is obtained by direct computation using Mathematica.

A.3 Proof of Lemma 2.3

The proof is adapted from [7, Theorem 2] for completeness. The relation $x^+ = x - \alpha \nabla f(x) + \beta(x - x^-)$ implies

$$x^+ - x = -\alpha \nabla f(x) + \beta(x - x^-), \quad (37)$$

$$\|x^+ - x\|^2 = \alpha^2 \|\nabla f(x)\|^2 + \beta^2 \|x - x^-\|^2 - 2\alpha\beta \langle \nabla f(x), x - x^- \rangle. \quad (38)$$

The L -smoothness of f implies

$$f(x^+) \leq f(x) + \langle \nabla f(x), x^+ - x \rangle + \frac{L}{2} \|x^+ - x\|^2. \quad (39)$$

Plugging (37) and (38) into (39) gives

$$\begin{aligned} f(x^+) &\leq f(x) - \alpha \|\nabla f(x)\|^2 + \beta \langle \nabla f(x), x - x^- \rangle \\ &\quad + \frac{\alpha^2 L}{2} \|\nabla f(x)\|^2 + \frac{\beta^2 L}{2} \|x - x^-\|^2 - L\alpha\beta \langle \nabla f(x), x - x^- \rangle \\ &= f(x) + (\frac{\alpha^2 L}{2} - \alpha) \|\nabla f(x)\|^2 + \beta(1 - L\alpha) \langle \nabla f(x), x - x^- \rangle + \frac{\beta^2 L}{2} \|x - x^-\|^2. \end{aligned} \quad (40)$$

Multiplying (38) by $\frac{1-\alpha L}{2\alpha}$ and adding to (40), we have

$$\begin{aligned} f(x^+) + \frac{1-\alpha L}{2\alpha} \|x^+ - x\|^2 &\leq f(x) + (\frac{\alpha^2 L}{2} - \alpha) \|\nabla f(x)\|^2 + \beta(1 - L\alpha) \langle \nabla f(x), x - x^- \rangle + \frac{\beta^2 L}{2} \|x - x^-\|^2 \\ &\quad + \frac{\alpha(1-\alpha L)}{2} \|\nabla f(x)\|^2 + \frac{\beta^2(1-\alpha L)}{2\alpha} \|x - x^-\|^2 - \beta(1 - L\alpha) \langle \nabla f(x), x - x^- \rangle \\ &= f(x) + \frac{\beta^2}{2\alpha} \|x - x^-\|^2 - \frac{\alpha}{2} \|\nabla f(x)\|^2 \\ &= f(x) + \frac{1-\alpha L}{2\alpha} \|x - x^-\|^2 - \frac{\alpha}{2} [\|\nabla f(x)\|^2 + \frac{1-\alpha L - \beta^2}{\alpha^2} \|x - x^-\|^2]. \end{aligned} \quad (41)$$

Let $\omega = \frac{1-\alpha L}{2\alpha}$. Then (41) becomes

$$\varphi_\omega(x^+, x) \leq \varphi_\omega(x, x^-) - \frac{\alpha}{2} [\|\nabla f(x)\|^2 + \frac{1-\alpha L - \beta^2}{\alpha^2} \|x - x^-\|^2]$$

and this completes the proof.

A.4 Proof of Lemma 2.4

Notation. Recall that $z = (x, x^-)$. We explicitly use (x, x^-) when establishing the analytic properties of φ_ω .

Convexity. To prove the convexity of $h_{x,x^-}(P, \beta)$ in (P, β) , it suffices to show the joint convexity of

$$\varphi_\omega(x^+(P, \beta), x) := f(x^+(P, \beta)) + \frac{\omega}{2} \|x^+(P, \beta) - x\|^2.$$

Since $x^+(P, \beta) := x - P\nabla f(x) + \beta(x - x^-)$ is affine in (P, β) and f and $\|\cdot\|^2$ are both convex, the composition rule of convex functions implies the joint convexity of $\varphi_\omega(x^+(P, \beta), x)$.

Gradient. The chain rule implies

$$\nabla_P h_{x,x^-}(P, \beta) = \frac{-[\nabla f(x^+(P, \beta)) + \omega(x^+(P, \beta) - x)] \nabla f(x)^\top}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2}, \quad (42)$$

$$\nabla_\beta h_{x,x^-}(P, \beta) = \frac{\langle \nabla f(x^+(P, \beta)) + \omega(x^+(P, \beta) - x), x - x^- \rangle}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2}. \quad (43)$$

Smoothness. The difference of the gradients in P and β are respectively bounded by

$$\begin{aligned} \|\nabla_P h_{x,x^-}(P_1, \beta_1) - \nabla_P h_{x,x^-}(P_2, \beta_2)\|_F &= \frac{\|[\nabla f(x^+(P_1, \beta_1)) - \nabla f(x^+(P_2, \beta_2)) + \omega(x^+(P_1, \beta_1) - x^+(P_2, \beta_2))] \nabla f(x)^\top\|_F}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2} \\ &= \frac{\|\nabla f(x^+(P_1, \beta_1)) - \nabla f(x^+(P_2, \beta_2)) + \omega(x^+(P_1, \beta_1) - x^+(P_2, \beta_2))\| \|\nabla f(x)\|}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2} \end{aligned} \quad (44)$$

$$\begin{aligned} |\nabla_\beta h_{x,x^-}(P_1, \beta_1) - \nabla_\beta h_{x,x^-}(P_2, \beta_2)| &= \frac{|\langle \nabla f(x^+(P_1, \beta_1)) - \nabla f(x^+(P_2, \beta_2)) + \omega(x^+(P_1, \beta_1) - x^+(P_2, \beta_2)), x - x^- \rangle|}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2} \\ &\leq \frac{\|\nabla f(x^+(P_1, \beta_1)) - \nabla f(x^+(P_2, \beta_2)) + \omega(x^+(P_1, \beta_1) - x^+(P_2, \beta_2))\| \|x - x^-\|}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2}, \end{aligned} \quad (45)$$

where (45) uses Cauchy's inequality. Observe that the common norm in the numerator can be bounded by

$$\begin{aligned} &\|\nabla f(x^+(P_1, \beta_1)) - \nabla f(x^+(P_2, \beta_2)) + \omega(x^+(P_1, \beta_1) - x^+(P_2, \beta_2))\| \\ &\leq \|\nabla f(x^+(P_1, \beta_1)) - \nabla f(x^+(P_2, \beta_2))\| + \omega \|x^+(P_1, \beta_1) - x^+(P_2, \beta_2)\| \quad (\text{by } \|a + b\| \leq \|a\| + \|b\|) \\ &\leq (L + \omega) \|x^+(P_1, \beta_1) - x^+(P_2, \beta_2)\| \quad (\text{by } L\text{-smooth of } f) \\ &= (L + \omega) \|(P_2 - P_1) \nabla f(x) + (\beta_1 - \beta_2)(x - x^-)\| \\ &\leq (L + \omega) \|P_1 - P_2\| \|\nabla f(x)\| + |\beta_1 - \beta_2| \|x - x^-\| \quad (\text{by submultiplicativity}) \\ &\leq (L + \omega) \sqrt{\|P_1 - P_2\|_F^2 + \frac{1}{L^2} (\beta_1 - \beta_2)^2} \sqrt{\|\nabla f(x)\|^2 + L^2 \|x - x^-\|^2} \quad (\text{by Cauchy's inequality}) \\ &= (L + \omega) \sqrt{\|\nabla f(x)\|^2 + L^2 \|x - x^-\|^2} \|(P_1, \beta_1) - (P_2, \beta_2)\|. \quad (\text{by definition of } \|(P, \beta)\| \text{ norm}) \end{aligned}$$

Then (44) and (45) can be further bounded by

$$\begin{aligned} \|\nabla_P h_{x,x^-}(P_1, \beta_1) - \nabla_P h_{x,x^-}(P_2, \beta_2)\|_F &\leq (L + \omega) \frac{\sqrt{\|\nabla f(x)\|^2 + L^2 \|x - x^-\|^2} \|\nabla f(x)\|}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2} \|(P_1, \beta_1) - (P_2, \beta_2)\|, \\ |\nabla_\beta h_{x,x^-}(P_1, \beta_1) - \nabla_\beta h_{x,x^-}(P_2, \beta_2)| &\leq (L + \omega) \frac{\sqrt{\|\nabla f(x)\|^2 + L^2 \|x - x^-\|^2} \|x - x^-\|}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2} \|(P_1, \beta_1) - (P_2, \beta_2)\|. \end{aligned}$$

Hence the dual norm $\|\cdot\|_*$ of the gradient difference is bounded by

$$\begin{aligned} &\|(\nabla_P h_{x,x^-}(P_1, \beta_1) - \nabla_P h_{x,x^-}(P_2, \beta_2), \nabla_\beta h_{x,x^-}(P_1, \beta_1) - \nabla_\beta h_{x,x^-}(P_2, \beta_2))\|_*^2 \\ &= \|\nabla_P h_{x,x^-}(P_1, \beta_1) - \nabla_P h_{x,x^-}(P_2, \beta_2)\|_F^2 + L^2 \|\nabla_\beta h_{x,x^-}(P_1, \beta_1) - \nabla_\beta h_{x,x^-}(P_2, \beta_2)\|^2 \\ &\leq (L + \omega)^2 \frac{\|\nabla f(x)\|^2 + L^2 \|x - x^-\|^2}{\|\nabla f(x)\|^2 + \frac{\omega}{2} \|x - x^-\|^2} \|(P_1, \beta_1) - (P_2, \beta_2)\|^2 \\ &\leq (L + \omega)^2 \max\{\frac{4L^4}{\tau^2}, 1\} \|(P_1, \beta_1) - (P_2, \beta_2)\|^2 \quad (\text{by } \frac{a^2 + \gamma b^2}{a^2 + \delta b^2} \leq \max\{\frac{\gamma}{\delta}, 1\}) \\ &= (L + \omega)^2 \|(P_1, \beta_1) - (P_2, \beta_2)\|^2, \quad (\text{by } \tau \geq 2L^2) \end{aligned}$$

which implies $(L + \omega)$ -smoothness of $h_{x,x^-}(P, \beta)$. Finally, as a by-product relation that will be later used, we derive the dual norm $\|\cdot\|_*$ of the gradient is bounded by

$$\begin{aligned} \|(\nabla_P h_{x,x^-}(P, \beta), \nabla_\beta h_{x,x^-}(P, \beta))\|_*^2 &= \|\nabla_P h_{x,x^-}(P, \beta)\|^2 + L^2 \|\nabla_\beta h_{x,x^-}(P, \beta)\|^2 \\ &\leq \frac{\|\nabla f(x^+(P, \beta)) + \omega(x^+(P, \beta) - x)\|^2 [\|\nabla f(x)\|^2 + L^2 \|x - x^-\|^2]}{[\|\nabla f(x)\|^2 + \frac{\tau}{2} \|x - x^-\|^2]^2} \quad (\text{by (42) and (43)}) \\ &\leq \max\left\{\frac{2L^2}{\tau}, 1\right\} \frac{\|\nabla f(x^+(P, \beta)) + \omega(x^+(P, \beta) - x)\|^2}{\|\nabla f(x)\|^2 + \frac{\tau}{2} \|x - x^-\|^2}. \end{aligned} \quad (46)$$

A.5 Proof of Theorem 2.2

Note that the initial state z^1 with $z_2^1 = z_1^1$ implies that the initial potential is the initial suboptimality:

$$\varphi_\omega(z^1) = f(z^1) - f^* + \frac{\omega}{2} \|z_1^1 - z_2^1\|^2 = f(z^1) - f^*.$$

Convex f . Observe that

$$\varphi_\omega(z^{K+1}) = \frac{1}{\sum_{k=1}^K \left[\frac{\varphi_\omega(z^k) - \varphi_\omega(z^{k+1})}{\varphi_\omega(z^{k+1})\varphi_\omega(z^k)} \right] + \frac{1}{\varphi_\omega(z^1)}} = \frac{1}{\sum_{k=1}^K -b_k \left[\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{\varphi_\omega(z^{k+1})\varphi_\omega(z^k)} \right] + \frac{1}{f(z^1) - f^*}}. \quad (47)$$

Note that $b_k \leq 0$ since the algorithm guarantees that the potential is monotone $\varphi_\omega(z^{k+1}) \leq \varphi_\omega(z^k)$. To further upper bound $\varphi_\omega(z^{K+1})$, we need to lower bound the ratio

$$\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{\varphi_\omega(z^{k+1})\varphi_\omega(z^k)} \geq \frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{\varphi_\omega(z^k)^2} = \frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{[f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2]^2}.$$

Case 1. $f(z^k) - f^* \geq \frac{\omega}{2} \|z_1^k - z_2^k\|^2$. Observe that $z_1^k \in \{x : f(x) \leq f(z^1)\}$ for all $k \geq 1$ since the monotonicity of potential guarantees

$$f(z^k) - f^* \leq f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2 = \varphi_\omega(z^k) \leq \varphi_\omega(z^1) = f(z^1) - f^*.$$

Then $\frac{f(z^k) - f^*}{\|\nabla f(z^k)\|} \leq \|z_1^k - x^*\| \leq \Delta := \max_{x \in \{x : f(x) \leq f(z^1)\}} \min_{x^* \in \mathcal{X}^*} \|x - x^*\|$ and hence

$$\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{[f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2]^2} \geq \frac{\|\nabla f(z^k)\|^2}{4[f(z^k) - f^*]^2} \geq \frac{1}{4\Delta^2}. \quad (48)$$

Case 2. $f(z^k) - f^* \leq \frac{\omega}{2} \|z_1^k - z_2^k\|^2$. Observe that

$$\frac{\omega}{2} \|z_1^k - z_2^k\|^2 \leq \varphi_\omega(z^k) \leq \varphi_\omega(z^1) = f(z^1) - f^*. \quad (49)$$

Using (49) and the choice of $\tau \geq 2L\omega$, we have

$$\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{[f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2]^2} \geq \frac{\frac{\tau}{2} \|z_1^k - z_2^k\|^2}{\omega^2 \|z_1^k - z_2^k\|^4} = \frac{\tau}{2\omega^2} \frac{1}{\|z_1^k - z_2^k\|^2} \geq \frac{L}{2[f(z^1) - f^*]}. \quad (50)$$

Using (47) and the two lower bounds (48) and (50), the suboptimality can be bounded by

$$f(z^{K+1}) - f^* \leq \varphi_\omega(z^{K+1}) \leq \frac{f(z^1) - f^*}{1 + V \sum_{k=1}^K -b_k},$$

where $V := \min\left\{\frac{f(z^1) - f^*}{4\Delta^2}, \frac{L}{2}\right\}$.

μ -strongly convex f . Observe that

$$\begin{aligned} \frac{\varphi_\omega(z^{K+1})}{\varphi_\omega(z^1)} &= \prod_{k=1}^K \frac{\varphi_\omega(z^{k+1})}{\varphi_\omega(z^k)} \leq \left(\frac{1}{K} \sum_{k=1}^K \frac{\varphi_\omega(z^{k+1})}{\varphi_\omega(z^k)} \right)^K \\ &= \left(1 + \frac{1}{K} \sum_{k=1}^K \frac{\varphi_\omega(z^{k+1}) - \varphi_\omega(z^k)}{\varphi_\omega(z^k)} \right)^K \\ &= \left(1 - \frac{1}{K} \sum_{k=1}^K -b_k \frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2} \right)^K. \end{aligned} \quad (51)$$

To further upper bound (51), we lower bound the ratio $\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2}$ by case analysis.

Case 1. $f(z^k) - f^* \leq \frac{\omega}{2} \|z_1^k - z_2^k\|^2$. Then

$$\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2} \geq \frac{\frac{\tau}{2} \|z_1^k - z_2^k\|^2}{\omega \|z_1^k - z_2^k\|^2} = \frac{\tau}{2\omega}. \quad (52)$$

Case 2. $f(z^k) - f^* \geq \frac{\omega}{2} \|z_1^k - z_2^k\|^2$. Then using $\frac{\|\nabla f(z^k)\|^2}{f(z^k) - f^*} \geq 2\mu$ to obtain

$$\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2} \geq \frac{\|\nabla f(z^k)\|^2}{2[f(z^k) - f^*]} \geq \mu. \quad (53)$$

Since ω and τ are chosen such that $\frac{\tau}{2\omega} \geq L \geq \mu$, two lower bounds (52)–(53) together imply

$$\frac{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}{f(z^k) - f^* + \frac{\omega}{2} \|z_1^k - z_2^k\|^2} \geq \min\left\{\frac{\tau}{2\omega}, \mu\right\} = \mu,$$

and hence we conclude from (51) that

$$f(z^{K+1}) - f^* \leq \varphi_\omega(z^{K+1}) \leq \varphi_\omega(z^1) \left(1 - \frac{\mu}{K} \sum_{k=1}^K -b_k\right)^K = [f(z^1) - f^*] \left(1 - \frac{\mu}{K} \sum_{k=1}^K -b_k\right)^K.$$

A.6 Proof of Theorem 2.3

Recall that the per iteration progress of OSGM-HB is

$$b_k := \frac{\varphi_\omega(z^{k+1}) - \varphi_\omega(z^k)}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}.$$

The scheduler proposes the state

$$z^{k+1/2} := (z^k)^+ = (z_1^k - P_k \nabla f(z^k) + \beta_k (z_1^k - z_2^k), z_1^k).$$

Monotone Lookahead OSGM-HB determines the next state z^{k+1} such that

$$\varphi_\omega(z^{k+1}) \leq \min\{\varphi_\omega(\mathcal{M}_{\text{look}}(z^{k+1/2})), \varphi_\omega(z^k)\}, \quad (54)$$

where $\mathcal{M}_{\text{look}}(z) := (z_1 - \frac{1}{L+\omega} \nabla_1 \varphi_\omega(z), z_2)$. The next lemma quantifies the improvement on iteration progress by this landscape action.

Lemma A.3. *Suppose $\tau \geq 2L^2$. Then the monotone lookahead landscape action (54) guarantees*

$$b_k \leq \min\{h_{z^k}(P_k, \beta_k) - \frac{1}{2(L+\omega)} \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2, 0\}. \quad (55)$$

Proof. The monotone landscape action guarantees $\varphi_\omega(z^{k+1}) \leq \min\{\varphi_\omega(\mathcal{M}_{\text{look}}(z^{k+1/2})), \varphi_\omega(z^k)\}$, and hence

$$b_k := \frac{\varphi_\omega(z^{k+1}) - \varphi_\omega(z^k)}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2} \leq \min\left\{\frac{\varphi_\omega(\mathcal{M}_{\text{look}}(z^{k+1/2})) - \varphi_\omega(z^k)}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}, 0\right\}. \quad (56)$$

Descent lemma on $(L + \omega)$ -smooth potential $\varphi_\omega(\cdot, z_2^{k+1/2})$ implies

$$\begin{aligned}\varphi_\omega(\mathcal{M}_{\text{look}}(z^{k+1/2})) &\leq \varphi_\omega(z^{k+1/2}) - \frac{1}{2(L+\omega)} \|\nabla_1 \varphi_\omega(z^{k+1/2})\|^2 \\ &= \varphi(z^{k+1/2}) - \frac{1}{2(L+\omega)} \|\nabla f(z^{k+1/2}) + \omega(z_1^{k+1/2} - z_2^{k+1/2})\|^2.\end{aligned}$$

Subtracting both sides by $\varphi_\omega(z^k)$ and dividing by $\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2$, we obtain

$$\begin{aligned}\frac{\varphi_\omega(\mathcal{M}_{\text{look}}(z^{k+1/2})) - \varphi_\omega(z^k)}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2} &\leq \frac{\varphi_\omega(z^{k+1/2}) - \varphi_\omega(z^k)}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2} - \frac{1}{2(L+\omega)} \frac{\|\nabla f(z^{k+1/2}) + \omega(z_1^{k+1/2} - z_2^{k+1/2})\|^2}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2} \\ &= h_{z^k}(P_k, \beta_k) - \frac{1}{2(L+\omega)} \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2,\end{aligned}\tag{57}$$

where (57) uses the inequality of gradient norm from (46) and that $\tau \geq 2L^2$:

$$\|\nabla h_{z^k}(P_k, \beta_k)\|_*^2 \leq \frac{\|\nabla f(z^{k+1/2}) + \omega(z_1^{k+1/2} - z_2^{k+1/2})\|^2}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2}.$$

Combining the two inequalities (56) and (57) completes the proof. \square

To bound the cumulative feedback $\sum_{k=1}^K h_{z^k}(P_k, \beta_k)$, we use the analysis of online gradient descent.

Lemma A.4. *Consider online gradient descent on heavy-ball feedback*

$$\begin{aligned}P_{k+1} &= \Pi_{\mathcal{P}}[P_k - \eta_P \nabla h_{z^k}(P_k, \beta_k)], \\ \beta_{k+1} &= \Pi_{\mathcal{B}}[\beta_k - \eta_\beta \nabla h_{z^k}(P_k, \beta_k)],\end{aligned}$$

with $\eta_P = \eta$ and $\eta_\beta = L^2 \eta$ for some $\eta > 0$. Then for any $(\hat{P}, \hat{\beta}) \in \mathcal{P} \times \mathcal{B}$, the sequence $\{(P_k, \beta_k)\}$ satisfies

$$\sum_{k=1}^K h_{z^k}(P_k, \beta_k) \leq \sum_{k=1}^K h_{z^k}(\hat{P}, \hat{\beta}) + \frac{1}{2\eta} \|(P_1 - \hat{P}, \beta_1 - \hat{\beta})\|^2 + \frac{\eta}{2} \sum_{k=1}^K \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2.$$

Proof. The proof is similar to the analysis of online gradient descent under Frobenius norm. Note that our algorithm finally takes $\mathcal{P} = \mathbb{R}^{n \times n}$, and the analysis here is more general. Observe that

$$\begin{aligned}\|(P_{k+1} - \hat{P}, \beta_{k+1} - \hat{\beta})\|^2 &= \|\Pi_{\mathcal{P}}[P_k - \eta_P \nabla_P h_{z^k}(P_k, \beta_k)] - \hat{P}\|_F^2 + \frac{1}{L^2} \|\Pi_{\mathcal{B}}[\beta_k - \eta_\beta \nabla_\beta h_{z^k}(P_k, \beta_k)] - \hat{\beta}\|^2 \\ &\leq \|P_k - \hat{P} - \eta_P \nabla_P h_{z^k}(P_k, \beta_k)\|_F^2 + \frac{1}{L^2} \|\beta_k - \hat{\beta} - \eta_\beta \nabla_\beta h_{z^k}(P_k, \beta_k)\|^2 \\ &= \|P_k - \hat{P}\|_F^2 - 2\eta_P \langle P_k - \hat{P}, \nabla_P h_{z^k}(P_k, \beta_k) \rangle + \eta_P^2 \|\nabla_P h_{z^k}(P_k, \beta_k)\|^2 \\ &\quad + \frac{1}{L^2} \|\beta_k - \hat{\beta}\|^2 - 2\frac{\eta_\beta}{L^2} \langle \beta_k - \hat{\beta}, \nabla_\beta h_{z^k}(P_k, \beta_k) \rangle + \frac{\eta_\beta^2}{L^2} \|\nabla_\beta h_{z^k}(P_k, \beta_k)\|^2 \\ &= \|(P_k - \hat{P}, \beta_k - \hat{\beta})\|^2 + \eta^2 [\|\nabla_P h_{z^k}(P_k, \beta_k)\|^2 + L^2 \|\nabla_\beta h_{z^k}(P_k, \beta_k)\|^2] \\ &\quad - 2\eta [\langle P_k - \hat{P}, \nabla_P h_{z^k}(P_k, \beta_k) \rangle + \langle \beta_k - \hat{\beta}, \nabla_\beta h_{z^k}(P_k, \beta_k) \rangle] \\ &\leq \|(P_k - \hat{P}, \beta_k - \hat{\beta})\|^2 + \eta^2 \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2 - 2\eta [h_{z^k}(P_k, \beta_k) - h_{z^k}(\hat{P}, \hat{\beta})],\end{aligned}$$

which rearranges to

$$h_{z^k}(P_k, \beta_k) \leq h_{z^k}(\hat{P}, \hat{\beta}) + \frac{1}{2\eta} [\|(P_k - \hat{P}, \beta_k - \hat{\beta})\|^2 - \|(P_{k+1} - \hat{P}, \beta_{k+1} - \hat{\beta})\|^2] + \frac{\eta}{2} \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2.$$

Telescope from $k = 1, \dots, K$ to obtain

$$\sum_{k=1}^K h_{z^k}(P_k, \beta_k) \leq \sum_{k=1}^K h_{z^k}(\hat{P}, \hat{\beta}) + \frac{1}{2\eta} \|(P_1 - \hat{P}, \beta_1 - \hat{\beta})\|^2 + \frac{\eta}{2} \sum_{k=1}^K \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2.$$

\square

Proof of Theorem 2.3. Note that our choice of $\tau = 16L^2$ satisfies the condition $\tau \geq 2L^2$ in **Lemma A.3**. By **Lemma A.3** and **Lemma A.4**, the cumulative progress of **Monotone Lookahead OSGM-HB** with $\eta \leq \frac{1}{L+\omega}$

can be bounded by

$$\begin{aligned}
\sum_{k=1}^K b_k &\leq \sum_{k=1}^K \min\{h_{z^k}(P_k, \beta_k) - \frac{1}{2(L+\omega)} \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2, 0\} && \text{(by Lemma A.3)} \\
&\leq \min\{\sum_{k=1}^K [h_{z^k}(P_k, \beta_k) - \frac{\eta}{2} \|\nabla h_{z^k}(P_k, \beta_k)\|_*^2], 0\} && \text{(by } \eta \leq \frac{\tau}{2L^2(L+\omega)} \text{)} \\
&\leq \min\{\sum_{k=1}^K h_{z^k}(\hat{P}, \hat{\beta}) + \frac{1}{2\eta} \|(P_1 - \hat{P}, \beta_1 - \hat{\beta})\|^2, 0\}. && \text{(by Lemma A.4)}
\end{aligned}$$

Choose the online gradient stepsize $\eta = \frac{1}{L+\omega}$ and denote

$$\rho_K(\hat{P}, \hat{\beta}) := \frac{1}{2\eta} \|(P_1 - \hat{P}, \beta_1 - \hat{\beta})\|^2$$

Then the bound on average progress of **Monotone Lookahead OSGM-HB** becomes

$$\frac{1}{K} \sum_{k=1}^K b_k \leq \min\{\sum_{k=1}^K h_{z^k}(\hat{P}, \hat{\beta}) + \frac{\rho_K(\hat{P}, \hat{\beta})}{K}, 0\}. \quad (58)$$

Our choice of parameters $\omega = 3L$ and $\tau = 16L^2$ also satisfies the condition $\tau \geq 2L\omega$ in **Theorem 2.2**. Plug in the bound (58) into the heavy-ball reductions (9)–(10) (**Theorem 2.2**) to obtain

$$(\text{convex}) \quad f(z^{K+1}) - f^* \leq \frac{f(z^1) - f^*}{KV \max\{\frac{1}{K} \sum_{k=1}^K -h_{z^k}(\hat{P}, \hat{\beta}) - \frac{\rho_K(\hat{P}, \hat{\beta})}{K}, 0\} + 1}, \quad (59)$$

$$(\mu\text{-strongly convex}) \quad f(z^{K+1}) - f^* \leq [f(z^1) - f^*](1 - \mu \max\{\frac{1}{K} \sum_{k=1}^K -h_{z^k}(\hat{P}, \hat{\beta}) - \frac{\rho_K(\hat{P}, \hat{\beta})}{K}, 0\})^K. \quad (60)$$

Lemma 2.3 with $\alpha = \frac{1}{4L}$ and $\beta = \frac{1}{2}$ implies $\omega = 3L$ and guarantees

$$\begin{aligned}
\varphi_w((z^k)^+ (\frac{1}{4L}I, \frac{1}{2})) &= \varphi_w(z_1^k - \frac{1}{4L} \nabla f(z^k) + \frac{1}{2}(z_1^k - z_2^k), z_1^k) \\
&\leq \varphi_w(z^k) - \frac{1}{8L} [\|\nabla f(z^k)\|^2 + 8L^2 \|z_1^k - z_2^k\|^2] \\
&= \varphi_w(z^k) - \frac{1}{8L} [\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2], && \text{(by } \tau = 16L^2 \text{)}
\end{aligned}$$

which rearranges to

$$h_{z^k}(\frac{1}{4L}I, \frac{1}{2}) = \frac{\varphi_w((z^k)^+ (\frac{1}{4L}I, \frac{1}{2})) - \varphi_w(z^k)}{\|\nabla f(z^k)\|^2 + \frac{\tau}{2} \|z_1^k - z_2^k\|^2} \leq -\frac{1}{8L}, \text{ for all } k. \quad (61)$$

Plug in $(\hat{P}, \hat{\beta}) = (\frac{1}{4L}I, \frac{1}{2})$ into (59) and (60) and further bound by (61) completes the proof.

B Proof of results in Section 3

B.1 Proof of Lemma 3.1

Since $h_x(P)$ is L -smooth, it is at most L -weakly convex. The second result needs several auxiliary lemmas.

Lemma B.1 (Hessian of PL-functions). *Suppose f is L -smooth, has H -Lipschitz Hessian, and is μ -PL, then $\lambda_{\min}(\nabla^2 f(x)) \geq -\mu^{-1}H\|\nabla f(x)\|$.*

Proof. Let $x^* = \Pi_{\mathcal{X}^*}[x]$. Using H -Lipschitz Hessian, we have

$$\begin{aligned}
\nabla^2 f(x) &= \nabla^2 f(x^*) + \nabla^2 f(x) - \nabla^2 f(x^*) \\
&\succeq \nabla^2 f(x^*) - \|\nabla^2 f(x) - \nabla^2 f(x^*)\| \cdot I \\
&\succeq \nabla^2 f(x^*) - H\|x - x^*\| \cdot I. && (H\text{-Lipschitz Hessian})
\end{aligned}$$

Since PL condition implies quadratic growth [21], $\frac{1}{2\mu} \|\nabla f(x)\|^2 \geq f(x) - f^* \geq \frac{\mu}{2} \|x - x^*\|^2$ and $\|x - x^*\| \leq \frac{1}{\mu} \|\nabla f(x)\|$. Hence $\nabla^2 f(x) \succeq \nabla^2 f(x^*) - \frac{H}{\mu} \|\nabla f(x)\| \cdot I \succeq -\mu^{-1}H\|\nabla f(x)\| \cdot I$ and this completes the proof. \square

Lemma B.2 (Hessian of $h_x(P)$). *Suppose f has H -Lipschitz Hessian and $x \notin \mathcal{S}^*$, then*

$$\|\nabla^2 h_x(P_1) - \nabla^2 h_x(P_2)\| \leq H \|\nabla f(x)\| \|P_1 - P_2\|.$$

Proof. Recall that $\nabla h_x(P) = \frac{\nabla f(x - P\nabla f(x))\nabla f(x)^\top}{\|\nabla f(x)\|^2}$ and $\nabla^2 h_x(P) = (gg^\top) \otimes \nabla^2 f(x - P\nabla f(x))$ where $g = \frac{\nabla f(x)}{\|\nabla f(x)\|}$. Since $\|g\| = 1$, H -Lipschitz Hessian implies

$$\|\nabla^2 h_x(P_1) - \nabla^2 h_x(P_2)\| \leq \|\nabla^2 f(x - P_1\nabla f(x)) - \nabla^2 f(x - P_2\nabla f(x))\| \leq H \|\nabla f(x)\| \|P_1 - P_2\|$$

and this completes the proof. \square

Proof of Lemma 3.1. Take any fixed $\bar{P} \in \mathcal{P}$. Using **Lemma B.2**, we deduce that

$$\begin{aligned} \nabla^2 h_x(P) &\succeq \nabla^2 h_x(\bar{P}) - H \|\nabla f(x)\| \|P - \bar{P}\| \cdot I \\ &\succeq [\lambda_{\min}(\nabla^2 f(x - \bar{P}\nabla f(x))) - H \|\nabla f(x)\| \|P - \bar{P}\|] \cdot I \\ &\succeq [\lambda_{\min}(\nabla^2 f(x - \bar{P}\nabla f(x))) - H \|\nabla f(x)\| D] \cdot I \end{aligned}$$

When f is additionally μ -PL and $\bar{P} = 0$, we can further apply $\lambda_{\min}(\nabla^2 f(x - \bar{P}\nabla f(x))) \geq -\mu^{-1}H \|\nabla f(x)\|$ from **Lemma B.1** and this completes the proof.

B.2 Proof of Lemma 3.2

The expression of gradient follows by definition $\nabla h_x^\lambda(P) = \nabla h_x(P) + \lambda(P - \frac{1}{L}I)$. $h_x^\lambda(\frac{1}{L}I) = h_x(\frac{1}{L}I) \leq -\frac{1}{2L}$ follows from the descent lemma. Using **Lemma 3.1**, local convexity of $h_x^\lambda(P)$ follows immediately and the smoothness constant follows by adding L , the smoothness of $h_x(P)$, and λ .

B.3 Proof of Theorem 3.1

These reductions have appeared in [13], and we repeat the proofs for completeness. The first relation (14) follows by definition:

$$\begin{aligned} (\sum_{k=1}^K -h_k) \cdot \min_{1 \leq k \leq K} \|\nabla f(x^k)\|^2 &\leq \sum_{k=1}^K -h_k \|\nabla f(x^k)\|^2 && \text{(by } -h_k \geq 0\text{)} \\ &= -[\sum_{k=1}^K f(x^{k+1}) - f(x^k)] && \text{(by definition of } h_k\text{)} \\ &= f(x^1) - f(x^{K+1}) \\ &\leq f(x^1) - f^*. && \text{(by } f(x^{K+1}) \leq f^*\text{)} \end{aligned}$$

and the second relation (15) follows from:

$$\begin{aligned} \frac{f(x^{K+1}) - f^*}{f(x^1) - f^*} &\leq \left(\frac{1}{K} \sum_{k=1}^K \frac{f(x^{k+1}) - f^*}{f(x^k) - f^*} \right)^K && \text{(by AM-GM inequality)} \\ &= \left(1 + \frac{1}{K} \sum_{k=1}^K \frac{f(x^{k+1}) - f(x^k)}{f(x^k) - f^*} \right)^K \\ &= \left(1 + \frac{1}{K} \sum_{k=1}^K \frac{h_k \|\nabla f(x^k)\|^2}{f(x^k) - f^*} \right)^K && \text{(by definition of } h_k\text{)} \\ &\leq \left(1 - \frac{2\mu}{K} \sum_{k=1}^K -h_k \right)^K && \text{(by } h_k \leq 0 \text{ and } \frac{\|\nabla f(x^k)\|^2}{f(x^k) - f^*} \geq 2\mu\text{)} \end{aligned}$$

and this completes the proof.

B.4 Proof of Lemma 3.3

Recall that $x^{k+1} = \mathcal{M}(x^{k+1/2}, x^k)$ and $x^{k+1/2} = x^k - P_k \nabla f(x^k)$. The per iteration progress and (regularized) feedback are

$$h_k = \frac{f(x^{k+1}) - f(x^k)}{\|\nabla f(x^k)\|^2}, \quad h_{x^k}(P_k) = \frac{f(x^{k+1/2}) - f(x^k)}{\|\nabla f(x^k)\|^2}, \quad h_{x^k}^\lambda(P_k) = h_{x^k}(P_k) + \frac{\lambda}{2} \|P_k - \frac{1}{L} I\|_F^2.$$

- *Monotone* landscape satisfies $f(x^{k+1}) \leq \min\{f(x^{k+1/2}), f(x^k)\}$ and thus

$$h_k = \frac{f(x^{k+1}) - f(x^k)}{\|\nabla f(x^k)\|^2} \leq \min\{h_{x^k}(P_k), 0\} = \min\{h_{x^k}^\lambda(P_k) - \frac{\lambda}{2} \|P_k - \frac{1}{L} I\|_F^2, 0\}.$$

- *Monotone lookahead* landscape satisfies $f(x^{k+1}) \leq \min\{f(x^{k+1/2} - \frac{1}{L} \nabla f(x^{k+1/2})), 0\}$ and by descent lemma, $f(x^{k+1/2} - \frac{1}{L} \nabla f(x^{k+1/2})) \leq f(x^{k+1/2}) - \frac{1}{2L} \|\nabla f(x^{k+1/2})\|^2$ and

$$\begin{aligned} h_k &= \frac{f(x^{k+1}) - f(x^k)}{\|\nabla f(x^k)\|^2} \\ &\leq \frac{f(x^{k+1/2}) - f(x^k)}{\|\nabla f(x^k)\|^2} - \frac{1}{2L} \frac{\|\nabla f(x^{k+1/2})\|^2}{\|\nabla f(x^k)\|^2} && \text{(by descent lemma)} \\ &= h_{x^k}(P_k) - \frac{1}{2L} \|\nabla h_{x^k}(P_k)\|_F^2 && \text{(by } \|\nabla h_{x^k}(P_k)\|_F^2 = \frac{\|\nabla f(x^{k+1/2})\|^2}{\|\nabla f(x^k)\|^2} \text{)} \\ &= h_{x^k}^\lambda(P_k) - \frac{\lambda}{2} \|P_k - \frac{1}{L} I\|_F^2 - \frac{1}{2L} \|\nabla h_{x^k}(P_k)\|_F^2 \\ &= h_{x^k}^\lambda(P_k) - \frac{1}{4L} \|\nabla h_{x^k}^\lambda(P_k)\|_F^2 + \frac{1}{4L} [\|\nabla h_{x^k}^\lambda(P_k)\|_F^2 - 2\lambda L \|P_k - \frac{1}{L} I\|_F^2 - 2\|\nabla h_{x^k}(P_k)\|_F^2]. \end{aligned}$$

Notice that when $\lambda \leq L$,

$$\begin{aligned} \|\nabla h_{x^k}^\lambda(P_k)\|_F^2 &= \|\nabla h_{x^k}(P_k) + \lambda(P_k - \frac{1}{L} I)\|_F^2 \\ &\leq 2\|\nabla h_{x^k}(P_k)\|_F^2 + 2\lambda^2 \|P_k - \frac{1}{L} I\|_F^2 && \text{(by } \|P + Q\|_F^2 \leq 2\|P\|_F^2 + 2\|Q\|_F^2 \text{)} \\ &\leq 2\|\nabla h_{x^k}(P_k)\|_F^2 + 2\lambda L \|P_k - \frac{1}{L} I\|_F^2 && \text{(by } \lambda \leq L \text{)} \end{aligned}$$

which implies $h_k \leq h_{x^k}^\lambda(P_k) - \frac{1}{4L} \|\nabla h_{x^k}^\lambda(P_k)\|_F^2$ and completes the proof.

B.5 Proof of Theorem 3.2

Recall the regret guarantee of online gradient descent: for any $\hat{P} \in \mathcal{P}$, we successively deduce that

$$\begin{aligned} \|P_{k+1} - \hat{P}\|_F^2 &= \|\Pi_{\mathcal{P}}[P_k - \eta \nabla h_{x^k}^{\lambda_k}(P_k)] - \hat{P}\|_F^2 \\ &\leq \|P_k - \eta \nabla h_{x^k}^{\lambda_k}(P_k) - \hat{P}\|_F^2 && \text{(by non-expansiveness)} \\ &= \|P_k - \hat{P}\|_F^2 - 2\eta \langle \nabla h_{x^k}^{\lambda_k}(P_k), P_k - \hat{P} \rangle + \eta^2 \|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2 \\ &\leq \|P_k - \hat{P}\|_F^2 - 2\eta [h_{x^k}^{\lambda_k}(P_k) - h_{x^k}^{\lambda_k}(\hat{P})] + \eta^2 \|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2, \end{aligned}$$

where the last inequality holds since $P_k, \hat{P} \in \mathcal{P}$ and our choice of $\{\lambda_k\}$ ensures $h_{x^k}^{\lambda_k}(P)$ is locally convex. Next, we rearrange the terms and obtain

$$h_{x^k}^{\lambda_k}(P_k) - h_{x^k}^{\lambda_k}(\hat{P}) \leq \frac{1}{2\eta} [\|P_k - \hat{P}\|_F^2 - \|P_{k+1} - \hat{P}\|_F^2] + \frac{\eta}{2} \|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2. \quad (62)$$

Back to the algorithm analysis. First, according to **Theorem 3.1**, the convergence of **Monotone Lookahead OSGM-H** follows from bounding the cumulative progress

$$\min_{1 \leq k \leq K} \|\nabla f(x^k)\|^2 \leq \frac{f(x^1) - f^*}{K} \frac{1}{\frac{1}{K} \sum_{k=1}^K -h_k}, \quad (63)$$

$$f(x^{K+1}) - f^* \leq [f(x^1) - f^*] (1 - \frac{2\mu}{K} \sum_{k=1}^K -h_k)^K. \quad (64)$$

Second, the cumulative progress $\sum_{k=1}^K h_k$ under monotone lookahead action is bounded by **Lemma 3.3**

$$\frac{1}{K} \sum_{k=1}^K h_k \leq \frac{1}{K} \sum_{k=1}^K \min\{h_{x^k}^{\lambda_k}(P_k) - \frac{1}{4L} \|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2, 0\} \leq \min\{\frac{1}{K} \sum_{k=1}^K h_{x^k}^{\lambda_k}(P_k) - \frac{1}{4L} \|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2, 0\}. \quad (65)$$

Third, the regret guarantee of online gradient descent (62) bounds the cumulative feedback $\sum_{k=1}^K h_{x^k}^{\lambda_k}(P_k)$ by

$$\sum_{k=1}^K h_{x^k}^{\lambda_k}(P_k) \leq \sum_{k=1}^K h_{x^k}^{\lambda_k}(\hat{P}) + \frac{1}{2\eta} \|P_1 - \hat{P}\|_F^2 + \frac{\eta}{2} \sum_{k=1}^K \|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2. \quad (66)$$

Putting (65) and (66) together, the cumulative progress is bounded by

$$\frac{1}{K} \sum_{k=1}^K h_k \leq \min\{\sum_{k=1}^K h_{x^k}^{\lambda_k}(\hat{P}) + \frac{1}{2\eta} \|P_1 - \hat{P}\|_F^2 + (\frac{\eta}{2} - \frac{1}{4L}) \sum_{k=1}^K \|\nabla h_{x^k}^{\lambda_k}(P_k)\|_F^2, 0\}.$$

When $\eta \leq \frac{1}{2L}$, the gradient norm term vanishes. With $\eta = \frac{1}{2L}$ the relation simplifies to

$$\frac{1}{K} \sum_{k=1}^K h_k \leq \min\{\sum_{k=1}^K h_{x^k}^{\lambda_k}(\hat{P}) + L \|P_1 - \hat{P}\|_F^2, 0\} = \min\{\sum_{k=1}^K h_{x^k}^{\lambda_k}(\hat{P}) + \frac{\lambda}{2} \|P - \frac{1}{L} I\|_F^2 + L \|P_1 - \hat{P}\|_F^2, 0\}.$$

Plugging the relation back into (63) and (64) completes the proof. Also notice that when $P_1 = \hat{P} = \frac{1}{L} I$, $h_x^{\lambda}(\hat{P}) \leq -\frac{1}{2L}$ and $\|P_1 - \hat{P}\|_F^2 = 0$.

C Proof of results in Section 5

C.1 Details of Subsection 5.1

In this section, we replace the online gradient descent step in **Monotone Lookahead OSGM-H** with online proximal point. For the generality of the result, we still use a matrix stepsize P .

Algorithm 7: Monotone Lookahead OSGM-H with online proximal point update

```

1 input initial point  $x^1$ , initial stepsize  $P_1 \in \mathbb{R}^{n \times n}$ , online proximal regularizer  $\eta > 0$ 
2 for  $k = 1, 2, \dots$  do
3    $x^{k+1/2} = x^k - P_k \nabla f(x^k)$ 
4   Choose  $x^{k+1}$  that satisfies  $f(x^{k+1}) \leq \min\{f(x^{k+1/2} - \frac{1}{L} \nabla f(x^{k+1/2})), f(x^k)\}$ 
5    $P_{k+1} = \arg \min_P \{h_{x^k}(P) + \frac{1}{2\eta} \|P - P_k\|_F^2\}$ 
6 end
```

To show the convergence of **Algorithm 7**, we again proceed in three steps: 1) bounding suboptimality with progress 2) bounding progress with feedback 3) bounding feedback with online learning. Since we essentially use the same hypergradient feedback, step 1) and 2) are the same as in Part I. It suffices to prove the regret guarantee of online proximal point and combine it with the lookahead action. Recall the definition of progress $h_k = \frac{f(x^{k+1}) - f(x^k)}{\|\nabla f(x^k)\|^2}$ and feedback $h_{x^k}(P_k) = \frac{f(x^{k+1/2}) - f(x^k)}{\|\nabla f(x^k)\|^2}$. We have, from [14, Lemma 4.1], that

Lemma C.1 (Feedback and progress [14]). *Let f be L -smooth and convex. Then $h_k \leq \min\{h_{x^k}(P_k) - \frac{1}{2L} \|\nabla h_{x^k}(P_k)\|_F^2, 0\}$ with monotone lookahead action.*

The proof requires an additional technical lemma that shows a relation between the value of the feedback function $h_x(P)$ and its Moreau envelope function $h_x^{1/\rho}(P) := \min_Q \{h_x(Q) + \frac{\rho}{2} \|Q - P\|_F^2\}$.

Lemma C.2. *Let $h^{1/\rho}(x) := \min_z \{h(z) + \frac{\rho}{2} \|z - x\|^2\}$ be the Moreau envelope of a differentiable convex function f . Then $h(x) - h^{1/\rho}(x) - \frac{1}{2\rho} \|\nabla h(x)\|^2 \leq 0$.*

Proof. By convexity, we have, for all z , that

$$\begin{aligned} h(z) + \frac{\rho}{2}\|z - x\|^2 &\geq h(x) + \langle \nabla h(x), z - x \rangle + \frac{\rho}{2}\|z - x\|^2 \\ &\geq \min_d \{h(x) + \langle \nabla h(x), d \rangle + \frac{\rho}{2}\|d\|^2\} \\ &= h(x) - \frac{1}{2\rho}\|\nabla h(x)\|^2. \end{aligned}$$

Minimizing the LHS, $h^{1/\rho}(x) = \min_z \{h(z) + \frac{\rho}{2}\|z - x\|^2\} \geq h(x) - \frac{1}{2\rho}\|\nabla h(x)\|^2$. This completes the proof. \square

Using the optimality condition on the proximal subproblem on Line 5 of **Algorithm 7**, for any \hat{P} , we deduce that

$$h_{x^k}(P_{k+1}) + \frac{1}{2\eta}\|P_{k+1} - P_k\|_F^2 \leq h_{x^k}(\hat{P}) + \frac{1}{2\eta}\|\hat{P} - P_k\|_F^2 - \frac{1}{2\eta}\|P_{k+1} - \hat{P}\|_F^2.$$

Re-arranging the terms, $h_{x^k}(P_{k+1}) - h_{x^k}(\hat{P}) \leq \frac{1}{2\eta}\|\hat{P} - P_k\|_F^2 - \frac{1}{2\eta}\|P_{k+1} - P_k\|_F^2 - \frac{1}{2\eta}\|P_{k+1} - \hat{P}\|_F^2$ and

$$\begin{aligned} h_{x^k}(P_k) - h_{x^k}(\hat{P}) &= h_{x^k}(P_k) - h_{x^k}(P_{k+1}) + h_{x^k}(P_{k+1}) - h_{x^k}(\hat{P}) \\ &\leq h_{x^k}(P_k) - h_{x^k}(P_{k+1}) + \frac{1}{2\eta}\|\hat{P} - P_k\|_F^2 - \frac{1}{2\eta}\|P_{k+1} - P_k\|_F^2 - \frac{1}{2\eta}\|P_{k+1} - \hat{P}\|_F^2 \\ &= h_{x^k}(P_k) - [h_{x^k}(P_{k+1}) + \frac{1}{2\eta}\|P_{k+1} - P_k\|_F^2] + \frac{1}{2\eta}[\|\hat{P} - P_k\|_F^2 - \|P_{k+1} - \hat{P}\|_F^2] \\ &= h_{x^k}(P_k) - h_{x^k}^\eta(P_k) + \frac{1}{2\eta}[\|\hat{P} - P_k\|_F^2 - \|P_{k+1} - \hat{P}\|_F^2] \\ &\leq \frac{\eta}{2}\|\nabla h_{x^k}(P_k)\|^2 + \frac{1}{2\eta}[\|\hat{P} - P_k\|_F^2 - \|P_{k+1} - \hat{P}\|_F^2] \end{aligned}$$

Finally, we bound the progress using the monotone lookahead action using **Lemma C.1**:

$$\begin{aligned} \sum_{k=1}^K h_k &\leq \min\{\sum_{k=1}^K h_{x^k}(P_k) - \frac{1}{2L}\|\nabla h_{x^k}(P_k)\|_F^2, 0\} \\ &\leq \min\{\sum_{k=1}^K h_{x^k}(\hat{P}) + (\frac{\eta}{2} - \frac{1}{2L})\|\nabla h_{x^k}(P_k)\|^2 + \frac{1}{2\eta}\|P_1 - \hat{P}\|_F^2, 0\} \end{aligned}$$

and for $\eta = \frac{1}{L}$, $\sum_{k=1}^K h_k \leq \sum_{k=1}^K h_{x^k}(\hat{P}) + \frac{L}{2}\|P_1 - \hat{P}\|_F^2$. Plugging the progress relation back into the reduction in Part I [14, Theorem 4.2] establishes the convergence.

C.2 Proof of Theorem 5.1

By definition of the AGD update, we deduce that

$$\begin{aligned}
y^k &= x^k + \frac{1}{\sqrt{\kappa+1}}(z^k - x^k) \\
&= \frac{\sqrt{\kappa}}{\sqrt{\kappa+1}}x^k + \frac{1}{\sqrt{\kappa+1}}z^k \\
x^{k+1} - x^* &= y^k - \frac{1}{L}\nabla f(y^k) - x^* \\
&= y^k - \frac{1}{L}(Ay^k - Ax^*) - x^* \\
&= (I - \frac{1}{L}A)y^k + \frac{1}{L}Ax^* - x^* \\
&= (I - \frac{1}{L}A)(\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}}x^k + \frac{1}{\sqrt{\kappa+1}}z^k) + \frac{1}{L}Ax^* - x^* \\
&= (I - \frac{1}{L}A)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}}(x^k - x^*) - (I - \frac{1}{L}A)\frac{1}{\sqrt{\kappa+1}}(z^k - x^*) \\
z^{k+1} - x^* &= (1 - \frac{1}{\sqrt{\kappa}})(z^k - x^*) + \frac{1}{\sqrt{\kappa}}(y^k - \frac{1}{\mu}(Ay^k - Ax^*) - x^*) \\
&= (1 - \frac{1}{\sqrt{\kappa}})(z^k - x^*) + \frac{1}{\sqrt{\kappa}}(I - \frac{1}{\mu}A)(y^k - x^*) \\
&= (1 - \frac{1}{\sqrt{\kappa}})(z^k - x^*) + \frac{1}{\sqrt{\kappa}}(I - \frac{1}{\mu}A)(\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}}x^k + \frac{1}{\sqrt{\kappa+1}}z^k - x^*) \\
&= \frac{1}{\sqrt{\kappa+1}}(I - \frac{1}{\mu}A)(x^k - x^*) + \frac{\sqrt{\kappa}}{\sqrt{\kappa+1}}(I - \frac{1}{L}A)(z^k - x^*)
\end{aligned}$$

On the other hand, we can express the potential function as follows:

$$\begin{aligned}
\varphi_\mu(x, z) &= \frac{1}{2}\langle (z - x^*), A(z - x^*) \rangle + \frac{1}{2\mu}\|A(x - x^*)\|^2 \\
&= \frac{1}{2}\|A^{1/2}(z - x^*)\|^2 + \frac{1}{2}\|\frac{1}{\sqrt{\mu}}A(x - x^*)\|^2 \\
&= \frac{1}{2}\left\|\begin{pmatrix} \frac{1}{\sqrt{\mu}}A(x - x^*) \\ A^{1/2}(z - x^*) \end{pmatrix}\right\|^2.
\end{aligned}$$

Using the AGD update, we can further write

$$\begin{pmatrix} \frac{1}{\sqrt{\mu}}A(x^{k+1} - x^*) \\ A^{1/2}(z^{k+1} - x^*) \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{\mu}}A(I - \frac{1}{L}A)(\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}}(x^k - x^*) + \frac{1}{\sqrt{\kappa+1}}(z^k - x^*)) \\ A^{1/2}[\frac{1}{\sqrt{\kappa+1}}(I - \frac{1}{\mu}A)(x^k - x^*) + \frac{\sqrt{\kappa}}{\sqrt{\kappa+1}}(I - \frac{1}{L}A)(z^k - x^*)] \end{pmatrix} = M \begin{pmatrix} \frac{1}{\sqrt{\mu}}A(x^k - x^*) \\ A^{1/2}(z^k - x^*) \end{pmatrix}$$

where $M = \begin{pmatrix} (I - \frac{1}{L}A)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} & \frac{1}{\sqrt{\mu}}(A^{1/2} - \frac{1}{L}A^{3/2})\frac{1}{\sqrt{\kappa+1}} \\ \sqrt{\mu}(A^{-1/2} - \frac{1}{\mu}A^{1/2})\frac{1}{\sqrt{\kappa+1}} & (I - \frac{1}{L}A)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} \end{pmatrix}$. Therefore,

$$\varphi(x^{k+1}, z^{k+1}) = \frac{1}{2}\left\|M \begin{pmatrix} \frac{1}{\sqrt{\mu}}Ax^k \\ A^{1/2}z^k \end{pmatrix}\right\|^2 \leq \frac{1}{2}\|M^\top M\|\left\|\begin{pmatrix} \frac{1}{\sqrt{\mu}}Ax^k \\ A^{1/2}z^k \end{pmatrix}\right\|^2 = \|M^\top M\| \cdot \varphi(x^k, z^k)$$

and it suffices to analyze $\|M^\top M\|$. Since $A \in \mathbb{S}_{++}^n$, without loss of generality, let $A = Q\Lambda^2 Q^\top$ and

$$\begin{aligned} M &= \begin{pmatrix} (I - \frac{1}{L}A)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} & \sqrt{\mu}(A^{-1/2} - \frac{1}{\mu}A^{1/2})\frac{1}{\sqrt{\kappa+1}} \\ \frac{1}{\sqrt{\mu}}(A^{1/2} - \frac{1}{L}A^{3/2})\frac{1}{\sqrt{\kappa+1}} & (I - \frac{1}{L}A)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} \end{pmatrix} \\ &= \begin{pmatrix} Q(I - \frac{1}{L}\Lambda^2)Q^\top\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} & Q(\sqrt{\mu}\Lambda^{-1} - \frac{1}{\sqrt{\mu}}\Lambda)Q^\top\frac{1}{\sqrt{\kappa+1}} \\ \frac{1}{\sqrt{\mu}}Q(\Lambda - \frac{1}{L}\Lambda^3)Q^\top\frac{1}{\sqrt{\kappa+1}} & Q(I - \frac{1}{L}\Lambda^2)Q^\top\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} \end{pmatrix} \\ &= \begin{pmatrix} Q & \\ & Q \end{pmatrix} \begin{pmatrix} (I - \frac{1}{L}\Lambda^2)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} & (\sqrt{\mu}\Lambda^{-1} - \frac{1}{\sqrt{\mu}}\Lambda)\frac{1}{\sqrt{\kappa+1}} \\ \frac{1}{\sqrt{\mu}}(\Lambda - \frac{1}{L}\Lambda^3)\frac{1}{\sqrt{\kappa+1}} & (I - \frac{1}{L}\Lambda^2)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} \end{pmatrix} \begin{pmatrix} Q & \\ & Q \end{pmatrix}^\top \\ &= \begin{pmatrix} Q & \\ & Q \end{pmatrix} N \begin{pmatrix} Q & \\ & Q \end{pmatrix}^\top, \end{aligned}$$

where $N = \begin{pmatrix} (I - \frac{1}{L}\Lambda^2)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} & (\sqrt{\mu}\Lambda^{-1} - \frac{1}{\sqrt{\mu}}\Lambda)\frac{1}{\sqrt{\kappa+1}} \\ \frac{1}{\sqrt{\mu}}(\Lambda - \frac{1}{L}\Lambda^3)\frac{1}{\sqrt{\kappa+1}} & (I - \frac{1}{L}\Lambda^2)\frac{\sqrt{\kappa}}{\sqrt{\kappa+1}} \end{pmatrix}$ and it suffices to analyze $\|N^\top N\| = \|M^\top M\|$.

By direct computation, $N^\top N$ is a 2 by 2 block matrix

$$N^\top N = \frac{1}{(\sqrt{\kappa} + 1)^2} \begin{pmatrix} \frac{L}{\mu} - \frac{1}{\mu}\Lambda^2 - \frac{1}{L\mu}\Lambda^4 + \frac{1}{\mu}\frac{1}{L^2}\Lambda^6 & \sqrt{L}(\Lambda^{-1} - \frac{1}{L}\Lambda - \frac{1}{L\mu}\Lambda^3 + \frac{1}{\mu L^2}\Lambda^5) \\ \sqrt{L}(\Lambda^{-1} - \frac{1}{L}\Lambda - \frac{1}{L\mu}\Lambda^3 + \frac{1}{\mu L^2}\Lambda^5) & \mu\Lambda^{-2} + (\frac{L}{\mu} - 2)I - \frac{1}{\mu}\Lambda^2 + \frac{1}{L\mu}\Lambda^4 \end{pmatrix},$$

where each block is a diagonal matrix. Since permutation does not affect the spectrum, we can permute $N^\top N$ to obtain multiple 2 by 2 matrices, each taking form of

$$\frac{1}{(\sqrt{\kappa} + 1)^2} \begin{pmatrix} \frac{L}{\mu} - \frac{1}{\mu}\lambda^2 - \frac{1}{L\mu}\lambda^4 + \frac{1}{\mu L^2}\lambda^6 & \sqrt{L}\lambda^{-1} - \frac{1}{\sqrt{L}}\lambda - \frac{1}{\sqrt{L}\mu}\lambda^3 + \frac{1}{\mu L^{3/2}}\lambda^5 \\ \sqrt{L}\lambda^{-1} - \frac{1}{\sqrt{L}}\lambda - \frac{1}{\sqrt{L}\mu}\lambda^3 + \frac{1}{\mu L^{3/2}}\lambda^5 & \mu\lambda^{-2} + \frac{L}{\mu} - 2 - \frac{1}{\mu}\lambda^2 + \frac{1}{L\mu}\lambda^4 \end{pmatrix},$$

where λ is one of the diagonal elements of Λ . Hence it suffices to compute

$$\begin{aligned} &\max_{\sqrt{\mu} \leq \lambda \leq \sqrt{L}} \left\| \begin{pmatrix} \frac{L}{\mu} - \frac{1}{\mu}\lambda^2 - \frac{1}{L\mu}\lambda^4 + \frac{1}{\mu L^2}\lambda^6 & \sqrt{L}\lambda^{-1} - \frac{1}{\sqrt{L}}\lambda - \frac{1}{\sqrt{L}\mu}\lambda^3 + \frac{1}{\mu L^{3/2}}\lambda^5 \\ \sqrt{L}\lambda^{-1} - \frac{1}{\sqrt{L}}\lambda - \frac{1}{\sqrt{L}\mu}\lambda^3 + \frac{1}{\mu L^{3/2}}\lambda^5 & \mu\lambda^{-2} + \frac{L}{\mu} - 2 - \frac{1}{\mu}\lambda^2 + \frac{1}{L\mu}\lambda^4 \end{pmatrix} \right\| \\ &= \max_{1 \leq \lambda \leq \sqrt{\kappa}} \left\| \begin{pmatrix} \frac{L}{\mu} - \lambda^2 - \frac{\mu}{L}\lambda^4 + \frac{\mu^2}{L^2}\lambda^6 & \sqrt{\frac{L}{\mu}}\lambda^{-1} - \sqrt{\frac{\mu}{L}}\lambda - \sqrt{\frac{\mu}{L}}\lambda^3 + (\frac{\mu}{L})^{3/2}\lambda^5 \\ \sqrt{\frac{L}{\mu}}\lambda^{-1} - \sqrt{\frac{\mu}{L}}\lambda - \sqrt{\frac{\mu}{L}}\lambda^3 + (\frac{\mu}{L})^{3/2}\lambda^5 & \lambda^{-2} + \frac{L}{\mu} - 2 - \lambda^2 + \frac{\mu}{L}\lambda^4 \end{pmatrix} \right\| \\ &= \max_{1 \leq \lambda \leq \sqrt{\kappa}} \left\| \begin{pmatrix} \kappa - \lambda^2 - \frac{1}{\kappa}\lambda^4 + \frac{1}{\kappa^2}\lambda^6 & \sqrt{\kappa}\lambda^{-1} - \kappa^{-1/2}\lambda - \sqrt{\kappa^{-1/2}}\lambda^3 + \kappa^{-3/2}\lambda^5 \\ \sqrt{\kappa}\lambda^{-1} - \kappa^{-1/2}\lambda - \sqrt{\kappa^{-1/2}}\lambda^3 + \kappa^{-3/2}\lambda^5 & \lambda^{-2} + \kappa - 2 - \lambda^2 + \kappa^{-1}\lambda^4 \end{pmatrix} \right\| \\ &= \max_{1 \leq \lambda \leq \sqrt{\kappa}} \left\| \begin{pmatrix} \kappa - \lambda^2 - \frac{1}{\kappa}\lambda^4 + \frac{1}{\kappa^2}\lambda^6 & \sqrt{\kappa}\lambda^{-1} - \kappa^{-1/2}\lambda - \sqrt{\kappa^{-1/2}}\lambda^3 + \kappa^{-3/2}\lambda^5 \\ \sqrt{\kappa}\lambda^{-1} - \kappa^{-1/2}\lambda - \sqrt{\kappa^{-1/2}}\lambda^3 + \kappa^{-3/2}\lambda^5 & \lambda^{-2} + \kappa - 2 - \lambda^2 + \kappa^{-1}\lambda^4 \end{pmatrix} \right\| \\ &= \max_{1 \leq \lambda \leq \sqrt{\kappa}} \left\| \begin{pmatrix} \frac{(\kappa - \lambda^2)^2(\kappa + \lambda^2)}{\kappa^2} & \frac{(\kappa - \lambda^2)(\kappa - \lambda^4)}{\kappa^{3/2}\lambda} \\ \frac{(\kappa - \lambda^2)(\kappa - \lambda^4)}{\kappa^{3/2}\lambda} & \frac{\lambda^4}{\kappa} + \kappa - \lambda^2 + \frac{1}{\lambda^2} - 2 \end{pmatrix} \right\| \end{aligned}$$

Two eigenvalues of the matrix are respectively

$$\frac{1}{2}(\frac{\lambda^6}{\kappa^2} \pm (\lambda^4 - \kappa) \frac{\sqrt{4(\kappa-1)\kappa^2\lambda^2 + \kappa^2 + 2\kappa(1-2\kappa)\lambda^4 + \lambda^8}}{\kappa^2\lambda^2} + 2\kappa - 2\lambda^2 + \frac{1}{\lambda^2} - 2),$$

whose maximizers are respectively $\sqrt{\kappa}$ and 1, giving

$$\max\{-2 + \frac{1}{\kappa} + \kappa, \frac{(\kappa-1)^2}{\kappa}(1 + \frac{1}{2\kappa} + \frac{\sqrt{4\kappa+1}}{2\kappa})\} \leq \frac{(\kappa-1)^2}{\kappa}(1 + \frac{1}{2\kappa} + \frac{\sqrt{4\kappa+1}}{2\kappa})$$

and multiplying the upper bound by $\frac{1}{(\sqrt{\kappa}+1)^2}$ completes the proof.

Remark 4. As a final remark, we compare the linear convergence rate for the new potential function with two standard convergence rates in the literature: $1 - \frac{1}{\sqrt{\kappa}}$ for AGD and $(\frac{1-\sqrt{\kappa}}{1+\sqrt{\kappa}})^2$ for Chebyshev iteration [8]. As **Figure 6** (left) shows, our new convergence rate is sandwiched between these two rates. **Figure 6** (right) verifies the validity of this new potential function on a toy quadratic minimization problem.




Figure 6: Left: comparison between the convergence rate of the new potential function and classic ones . Right: plots of the new potential function, the classic potential function, and the function value gap

C.3 Details of Subsection 5.4

In this section, we illustrate the convergence analysis of a variant of OSGM-H. We make the following assumptions throughout the section.

A1: f is L -smooth and convex with $\text{dom } f = \mathbb{R}^n$.

A2: The proximal term w is convex.

Similarly, we denote $\varphi^* := \varphi(x^*)$. The challenging part of the analysis is that the proximal term w will introduce non-Lipschitzness of the feedback function if $\|\mathcal{G}_L(x)\| \rightarrow 0$. To this end, we consider **Algorithm 8**.

Algorithm 8: OSGM-H with proximal gradient update

```

1 input initial point  $x^1$ , initial stepsize  $P_1 \in \mathbb{R}^{n \times n}$ , online proximal regularizer  $\eta > 0$ 
2 for  $k = 1, 2, \dots$  do
3    $P_{k+1} = \arg \min_P \left\{ \frac{f(x^k - P_k \mathcal{G}_L(x^k)) + \langle \nabla f(x^k - P_k \mathcal{G}_L(x^k)), \mathcal{G}_L(x^k)^\top, P - P_k \rangle + \omega(x^k - P \mathcal{G}_L(x^k))}{\|\mathcal{G}_L(x^k)\|^2} + \frac{1}{2\eta} \|P - P_k\|_F^2 \right\}$ 
4   Choose  $x^{k+1}$  such that  $\varphi(x^{k+1}) \leq \min\{\varphi(x^k - P_{k+1} \mathcal{G}_L(x^k)), \varphi(x^k)\}$ 
5 end
```

This update corresponds to **Classic-HDM** we discussed in **Section 2.2**. As a first step, we define progress $h_k := \frac{\varphi(x^{k+1}) - \varphi(x^k)}{\|\mathcal{G}_L(x^k)\|^2}$. Note that $h_k \leq 0$ by definition, which enables the following progress reduction.

Lemma C.3 (Nonsmooth reduction). *Under **A1** and **A2**, $\min_{1 \leq k \leq K} \|\mathcal{G}_L(x^k)\|^2 \leq \frac{\varphi(x^1) - \varphi^*}{K} \frac{1}{\frac{1}{K} \sum_{k=1}^K -h_k}$.*

Proof. The proof resembles that of the hypergradient reduction.

$$\begin{aligned}
(\sum_{k=1}^K -h_k) \cdot \min_{1 \leq k \leq K} \|\mathcal{G}_L(x^k)\|^2 &\leq \sum_{k=1}^K -h_k \|\mathcal{G}_L(x^k)\|^2 && \text{(by } -h_k \geq 0) \\
&= -[\sum_{k=1}^K \varphi(x^{k+1}) - \varphi(x^k)] && \text{(by definition of } h_k) \\
&= \varphi(x^1) - \varphi(x^{K+1}) \\
&\leq \varphi(x^1) - \varphi^*. && \text{(by } f(x^{K+1}) \leq f^*)
\end{aligned}$$

□

An important step is to establish online learning guarantees with respect to $\sum_{k=1}^K h_{x^k}(P_{k+1})$. Below, we prove the result for a more general setting.

Lemma C.4 (Prescient online proximal gradient). *Suppose $\{(\ell_k, \pi_k)\}$ are a sequence of convex functions and $\{\ell_k\}$ are L -smooth. Consider the online proximal gradient update*

$$P_{k+1} = \arg \min_{P \in \mathcal{P}} \{ \ell_k(P_k) + \langle \nabla \ell_k(P_k), P - P_k \rangle + \pi_k(P) + \frac{1}{2\eta} \|P - P_k\|_F^2 \}.$$

Then with $\eta \leq \frac{1}{L}$, $\sum_{k=1}^K \ell_k(P_{k+1}) + \pi_k(P_{k+1}) - [\ell_k(\hat{P}) + \pi_k(\hat{P})] \leq \frac{1}{2\eta} \|P_1 - \hat{P}\|_F^2$ for any $\hat{P} \in \mathcal{P}$.

Proof. By the optimality condition of the online update, for any $\hat{P} \in \mathcal{P}$,

$$\begin{aligned} & \langle \nabla \ell_k(P_k), P_{k+1} - P_k \rangle + \pi_k(P_{k+1}) + \frac{1}{2\eta} \|P_{k+1} - P_k\|_F^2 \\ & \leq \langle \nabla \ell_k(P_k), \hat{P} - P_k \rangle + \pi_k(\hat{P}) + \frac{1}{2\eta} \|\hat{P} - P_k\|_F^2 - \frac{1}{2\eta} \|P_{k+1} - \hat{P}\|_F^2. \end{aligned}$$

Re-arranging the terms,

$$\begin{aligned} \|P_{k+1} - \hat{P}\|_F^2 & \leq \|P_k - \hat{P}\|_F^2 - \|P_{k+1} - P_k\|_F^2 + 2\eta \langle \nabla \ell_k(P_k), \hat{P} - P_{k+1} \rangle + 2\eta [\pi_k(\hat{P}) - \pi_k(P_{k+1})] \\ & = \|P_k - \hat{P}\|_F^2 - \|P_{k+1} - P_k\|_F^2 + 2\eta \langle \nabla \ell_k(P_k), \hat{P} - P_k \rangle \\ & \quad + 2\eta \langle \nabla \ell_k(P_k), P_k - P_{k+1} \rangle + 2\eta [\pi_k(\hat{P}) - \pi_k(P_{k+1})] \\ & \leq \|P_k - \hat{P}\|_F^2 - \|P_{k+1} - P_k\|_F^2 - 2\eta [\ell_k(P_k) - \ell_k(\hat{P})] \\ & \quad + 2\eta \langle \nabla \ell_k(P_k), P_k - P_{k+1} \rangle + 2\eta [\pi_k(\hat{P}) - \pi_k(P_{k+1})] \quad (\text{Convexity of } \ell_k) \\ & = \|P_k - \hat{P}\|_F^2 - 2\eta [\ell_k(P_{k+1}) + \pi_k(P_{k+1}) - \ell_k(\hat{P}) - \pi_k(\hat{P})] \\ & \quad + 2\eta [\ell_k(P_{k+1}) - \ell_k(P_k) + \langle \nabla \ell_k(P_k), P_k - P_{k+1} \rangle - \frac{1}{2\eta} \|P_{k+1} - P_k\|_F^2]. \end{aligned}$$

When $\eta \leq \frac{1}{L}$, we have, by smoothness of ℓ_k , that

$$\ell_k(P_{k+1}) - [\ell_k(P_k) + \langle \nabla \ell_k(P_k), P_{k+1} - P_k \rangle] \leq \frac{L}{2} \|P_{k+1} - P_k\|_F^2 \leq \frac{1}{2\eta} \|P_{k+1} - P_k\|_F^2$$

and $\|P_{k+1} - \hat{P}\|_F^2 \leq \|P_k - \hat{P}\|_F^2 - 2\eta [\ell_k(P_{k+1}) + \pi_k(P_{k+1}) - \ell_k(\hat{P}) - \pi_k(\hat{P})]$. Telescoping completes the proof. \square

Chaining **Lemma C.3**, **Lemma C.4** and the fact $h_x(\frac{1}{L}I) \leq -\frac{1}{2L}$ [3, Corollary 10.18] gives the desired result.

Theorem C.1 (Theorem 5.2). *Suppose f is L -smooth convex and w is convex. Then for any benchmark stepsize $\hat{P} \in \mathcal{P}$, **Algorithm 8** achieves*

$$\min_{1 \leq k \leq K} \|\mathcal{G}_L(x^k)\|^2 \leq \frac{\varphi(x^1) - \varphi^*}{K} \frac{1}{\frac{1}{K} \max\{\sum_{k=1}^K h_{x^k}(\hat{P}) - \frac{L}{2} \|P_1 - \hat{P}\|_F^2, 0\}}.$$

In particular, taking $P_1 = \hat{P} = \frac{1}{L}I$ yields $\min_{1 \leq k \leq K} \|\mathcal{G}_L(x^k)\|^2 \leq \frac{2L[\varphi(x^1) - \varphi(x^*)]}{K}$.

Proof. Using **Lemma C.4**, we can bound

$$\sum_{k=1}^K h_k \leq \sum_{k=1}^K h_{x^k}(P_{k+1}) \leq \sum_{k=1}^K h_{x^k}(\hat{P}) + \frac{1}{2\eta} \|P_1 - \hat{P}\|_F^2.$$

Plugging the relation back into **Lemma C.3** and using $h_{x^k}(\frac{1}{L}I) \leq -\frac{1}{2L}$ completes the proof. \square