
Leveraging Predictions in Smoothed Online Convex Optimization via Gradient-based Algorithms

Yingying Li
SEAS

Harvard University
Cambridge, MA, USA. 02138
yingyingli@g.harvard.edu

Na Li
SEAS

Harvard University
Cambridge, MA, USA. 02138
nali@seas.harvard.edu

Abstract

We consider online convex optimization with time-varying stage costs and additional switching costs. Since the switching costs introduce coupling across all stages, multi-step-ahead (long-term) predictions are incorporated to improve the online performance. However, longer-term predictions tend to suffer from lower quality. Thus, a critical question is: *how to reduce the impact of long-term prediction errors on the online performance?* To address this question, we introduce a gradient-based online algorithm, Receding Horizon Inexact Gradient (RHIG), and analyze its performance by dynamic regrets in terms of the temporal variation of the environment and the prediction errors. RHIG only considers at most W -step-ahead predictions to avoid being misled by worse predictions in the longer term. The optimal choice of W suggested by our regret bounds depends on the tradeoff between the variation of the environment and the prediction accuracy. Additionally, we apply RHIG to a well-established stochastic prediction error model and provide expected regret and concentration bounds under correlated prediction errors. Lastly, we numerically test the performance of RHIG on quadrotor tracking problems.

1 Introduction

In this paper, we consider online convex optimization (OCO) with switching costs, also known as “smoothed” OCO (SOCO) in the literature [1–5]. The stage costs are time-varying but the decision maker (agent) has access to noisy predictions on the future costs. Specifically, we consider stage cost function $f(x_t; \theta_t)$ parameterized by a time-varying parameter $\theta_t \in \Theta$. At each stage $t \in \{1, 2, \dots, T\}$, the agent receives the predictions of the future parameters $\theta_{t|t-1}, \dots, \theta_{T|t-1}$, takes an action $x_t \in \mathbb{X}$, and suffers the stage cost $f(x_t; \theta_t)$ plus a switching cost $d(x_t, x_{t-1})$. The switching cost $d(x_t, x_{t-1})$ penalizes the changes in the actions between consecutive stages. This problem enjoys a wide range of applications. For example, in the data center management problems [6, 7], the switching cost captures the switch on/off costs of the servers [7], and noisy predictions on future electricity prices and network traffic are available for the center manager [8, 9]. Other applications include smart building [10, 11], robotics [12], smart grid [13], connected vehicles [14], optimal control [15], etc.

Unlike OCO [16], the switching costs considered in SOCO introduce coupling among all stages, so multi-step-ahead predictions are usually used for promoting the online performance. However, in most cases, predictions are not accurate, and longer-term predictions tend to suffer lower quality. Therefore, it is crucial to study *how to use the multi-step-ahead predictions effectively*, especially, *how to reduce the impact of long-term prediction errors on the online performance*.

Recent years have witnessed a growing interest in studying SOCO with predictions. However, most literature avoids the complicated analysis on noisy multi-step-ahead predictions by considering a

rather simplified prediction model: the costs in the next W stages are accurately predicted with no errors while the costs beyond the next W stages are adversarial and not predictable at all [3,6,7,15,17]. This first-accurate-then-adversarial model is motivated by the fact that long-term predictions are much worse than the short-term ones, but it fails to capture the gradually increasing prediction errors as one predicts further into the future. Several online algorithms have been proposed for this model, e.g. the optimization-based algorithm AFHC [6], the gradient-based algorithm RHGD [17], etc. Moreover, there have been a few attempts to consider noisy multi-step-ahead predictions in SOCO. In particular, [1] proposes a stochastic prediction error model to describe the correlation among prediction errors. This stochastic model generalizes stochastic filter prediction errors. Later, [18] proposes an optimization-based algorithm CHC, which generalizes AFHC and MPC [19], and analyzes its performance based on the stochastic model in [1].

However, many important questions remain unresolved for SOCO with noisy predictions. For example, though the discussions on the stochastic model in [1,18] are insightful, there still lacks a general understanding on the effects of prediction errors on SOCO without any (stochastic model) assumptions. Moreover, most methods in the literature [1,6,18] require fully solving multi-stage optimization programs at each stage; it is unclear whether any gradient-based algorithm, which is more computationally efficient, would work for SOCO with noisy multi-step-ahead predictions.

Our contributions. In this paper, we introduce a gradient-based online algorithm Receding Horizon Inexact Gradient (RHIG). It is a straightforward extension of RHGD, which was designed for the simple first-accurate-then-adversarial prediction model in [17]. In RHIG, the agent can choose to utilize only $W \geq 0$ steps of future predictions, where W is a tunable parameter for the agent.

We first analyze the dynamic regret of RHIG by considering general prediction errors without any (stochastic model) assumptions. Our regret bound depends on both the errors of the utilized predictions, i.e. k -step-ahead prediction errors for $k \leq W$; and the temporal variation of the environment $V_T = \sum_{t=1}^T \sup_{x \in \mathbb{X}} |f(x; \theta_t) - f(x; \theta_{t-1})|$. Interestingly, the regret bound shows that the optimal choice of W depends on the tradeoff between the variation of environment V_T and the prediction errors, that is, a large W is preferred when V_T is large while a small W is preferred when the prediction errors are large. Further, the k -step prediction errors have an exponentially decaying influence on the regret bound as k increases, indicating that RHIG effectively reduces the negative impact of the noisy multi-step-ahead predictions.

We then consider the stochastic prediction error model in [1,18] to analyze the performance of RHIG under correlated prediction errors. We provide an expected regret bound and a concentration bound on the regret. In both bounds, the long-term correlation among prediction errors has an exponentially decaying effect, indicating RHIG's good performance even with strongly correlated prediction errors.

Finally, we numerically test RHIG on online quadrotor tracking problems. Numerical experiments show that RHIG outperforms AFHC and CHC especially under larger prediction errors. Besides, we show that RHIG is robust to unforeseen shocks in the future.

Additional related work: There is a related line of work on predictable OCO (without switching costs) [20–23]. In this case, stage decisions are fully decoupled and only one-step-ahead predictions are relevant. The proposed algorithms include OMD [21,22], DMD [20], AOMD [23], whose regret bounds depend on one-step prediction errors [20,22,23] and V_T if dynamic regret is concerned [23].

Besides, it is worth mentioning the related online decision making problems with coupling across stages, e.g. OCO with memory [24,25], online optimal control [15,26–28], online Markov decision processes [29–31], etc. Leveraging inaccurate predictions in these problems is also worth exploring.

Notation: $\Pi_{\mathbb{X}}$ denotes the projection onto set \mathbb{X} . $\mathbb{X}^T = \mathbb{X} \times \dots \times \mathbb{X}$ is a Cartesian product. ∇_x denotes the gradient with x . $\sum_{t=0}^k a_t = 0$ if $k < 0$. $\|\cdot\|_F$ and $\|\cdot\|$ are Frobenius norm and L_2 norm.

2 Problem Formulation

Consider stage cost function $f(x_t; \theta_t)$ with a time-varying parameter $\theta_t \in \Theta$ and a switching cost $d(x_t, x_{t-1})$ that penalize the changes in the actions between stages. The total cost in horizon T is: $C(x; \theta) = \sum_{t=1}^T [f(x_t; \theta_t) + d(x_t, x_{t-1})]$, where $x_t \in \mathbb{X} \subseteq \mathbb{R}^n$, $\theta_t \in \Theta \subseteq \mathbb{R}^p$, and we denote $x := (x_1^\top, \dots, x_T^\top)^\top$, $\theta = (\theta_1^\top, \dots, \theta_T^\top)^\top$. The switching cost enjoys many applications as discussed in Section 1. The presence of switching costs $d(x_t, x_{t-1})$ couples decisions among stages. Therefore,

all parameters in horizon T , i.e. $\theta_1, \dots, \theta_T$, are needed to minimize $C(\mathbf{x}; \boldsymbol{\theta})$. However, in practice, only predictions are available ahead of the time and the predictions are often inaccurate, especially the long-term predictions. This may lead to wrong decisions and degrade the online performance. In this paper, we aim at designing an online algorithm to use prediction effectively and unveil the unavoidable influences of the prediction errors on the online performance.

Prediction models. In this paper, we denote the prediction of the future parameter θ_τ obtained at the beginning of stage t as $\theta_{\tau|t-1}$ for $t \leq \tau \leq T$. The initial predictions $\theta_{1|0}, \dots, \theta_{T|0}$ are usually available before the problem starts. We call $\theta_{t|t-k}$ as k -step-ahead predictions of parameter θ_t and let $\delta_t(k)$ denote the k -step-prediction error, i.e.

$$\delta_t(k) := \theta_t - \theta_{t|t-k}, \quad \forall 1 \leq k \leq t. \quad (1)$$

For notation simplicity, we define $\theta_{t|t} := \theta_t$ for $t \leq T$, and thus $\delta_t(k) = \delta_t(t)$ for $k \geq t$. Further, we denote the vector of k -step prediction errors of all stages as follows

$$\boldsymbol{\delta}(k) = (\delta_1(k)^\top, \dots, \delta_T(k)^\top)^\top \in \mathbb{R}^{pT}, \quad \forall 1 \leq k \leq T. \quad (2)$$

It is commonly observed that the number of lookahead steps heavily influences the prediction accuracy and in most cases long-term prediction errors are usually larger than short-term ones.

We will first consider the general prediction errors without additional assumptions on $\delta_t(k)$. Then, we will carry out a more insightful discussion for the case when the prediction error $\|\delta_t(k)\|$ is non-decreasing with the number of look-ahead steps k . Further, it is also commonly observed that the prediction errors are correlated. To study how the correlation among prediction errors affect the algorithm performance, we adopt the stochastic model of prediction errors in [1]. The stochastic model is a more general version of the prediction errors for Wiener filter, Kalman filter, etc. In Section 5, we will review this stochastic model and analyze the performance under this model.

Protocols. We summarize the protocols of our online problem below. We consider that the agent knows the function form $f(\cdot; \cdot)$ and $d(\cdot, \cdot)$ a priori. For each stage $t = 1, 2, \dots, T$, the agent

- receives the predictions $\theta_{t|t-1}, \dots, \theta_{T|t-1}$ at the beginning of stage;¹
- selects x_t based on the predictions and the history, i.e. $\theta_1, \dots, \theta_{t-1}, \theta_{t|t-1}, \dots, \theta_{T|t-1}$;
- suffers $f(x_t; \theta_t) + d(x_t, x_{t-1})$ at the end of stage after true θ_t is revealed.

Performance metrics. This paper considers (expected) dynamic regret [23]. The benchmark is the optimal solution \mathbf{x}^* in hindsight when $\boldsymbol{\theta}$ is known, i.e. $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; \boldsymbol{\theta})$, where $\mathbf{x}^* = ((x_1^*)^\top, \dots, (x_T^*)^\top)^\top$. Notice that \mathbf{x}^* depends on $\boldsymbol{\theta}$ but we omit $\boldsymbol{\theta}$ for brevity. Let $\mathbf{x}^{\mathcal{A}}$ denote the actions selected by the online algorithm \mathcal{A} . The dynamic regret of \mathcal{A} with parameter $\boldsymbol{\theta}$ is defined as

$$\text{Reg}(\mathcal{A}) = C(\mathbf{x}^{\mathcal{A}}; \boldsymbol{\theta}) - C(\mathbf{x}^*; \boldsymbol{\theta}) \quad (3)$$

When considering stochastic prediction errors, we define the expectation of the dynamic regret: $\mathbb{E}[\text{Reg}(\mathcal{A})] = \mathbb{E}[C(\mathbf{x}^{\mathcal{A}}; \boldsymbol{\theta}) - C(\mathbf{x}^*; \boldsymbol{\theta})]$, where the expectation is taken with respect to the randomness of the prediction error as well as the randomness of θ_t if applicable.

Lastly, we consider the following assumptions throughout this paper.

Assumption 1. $f(x; \theta)$ is α strongly convex and l_f smooth with respect to $x \in \mathbb{X}$ for any $\theta \in \Theta$. $d(x, x')$ is convex and l_d smooth with respect to $x, x' \in \mathbb{X}$.

Assumption 2. $\nabla_x f(x; \theta)$ is h -Lipschitz continuous with respect to θ for any x , i.e.

$$\|\nabla_x f(x; \theta_1) - \nabla_x f(x; \theta_2)\| \leq h \|\theta_1 - \theta_2\|, \quad \forall x \in \mathbb{X}, \theta_1, \theta_2 \in \Theta.$$

Assumption 1 is common in convex optimization literature [32]. Assumption 2 ensures a small prediction error on θ only causes a small error in the gradient. Without such an assumption, little can be achieved with noisy predictions. Lastly, we note that these assumptions are for the purpose of theoretical regret analysis. The designed algorithm would apply for general convex smooth functions.

¹If only W -step-ahead predictions are received, we define $\theta_{t+\tau|t-1} := \theta_{t+W-1|t-1}$ for $\tau \geq W$.

3 Receding Horizon Inexact Gradient (RHIG)

This section introduces our online algorithm Receding Horizon Inexact Gradient (RHIG). It is based on a promising online algorithm RHGD [17] designed for an over-simplified prediction model: at stage t , the next W -stage parameters $\{\theta_\tau\}_{\tau=t}^{t+W-1}$ are exactly known but parameters beyond W steps are adversarial and totally unknown. We will first briefly review RHGD and then introduce our RHIG as an extension of RHGD to handle the inaccurate multi-step-ahead predictions.

3.1 Preliminary: RHGD with accurate lookahead window

RHGD is built on the following observation: the k -th iteration of offline gradient descent (GD) on the total cost $C(\mathbf{x}; \boldsymbol{\theta})$ for stage variable $x_\tau(k)$, i.e.,

$$x_\tau(k) = \Pi_{\mathbb{X}}[x_\tau(k-1) - \eta \nabla_{x_\tau} C(\mathbf{x}(k-1); \boldsymbol{\theta})], \quad \forall 1 \leq \tau \leq T, \quad (4)$$

where $\nabla_{x_\tau} C(\mathbf{x}; \boldsymbol{\theta}) = \nabla_{x_\tau} f(x_\tau; \theta_\tau) + \nabla_{x_\tau} d(x_\tau, x_{\tau-1}) + \nabla_{x_\tau} d(x_{\tau+1}, x_\tau) \mathbb{1}_{(\tau \leq T-1)}$,

only requires neighboring stage variables $x_{\tau-1}(k-1)$, $x_\tau(k-1)$, $x_{\tau+1}(k-1)$ and local parameter θ_τ , instead of all variables $\mathbf{x}(k-1)$ and all parameters $\boldsymbol{\theta}$. This observation allows RHGD [17] (Algorithm 1) to implement the offline gradient (4) for W iterations by only using $\{\theta_\tau\}_{\tau=t}^{t+W-1}$. Specifically, at stage $2 - W \leq t \leq T$, RHGD initializes $x_{t+W}(0)$ by an oracle ϕ (Line 4), where ϕ can be OCO algorithms (e.g. OGD, OMD [16]) that compute $x_{t+W}(0)$ with $\{\theta_t\}_{t=1}^{t+W-1}$.² If $t + W > T$, skip this step. Next, RHGD applies the offline GD (4) to compute $x_{t+W-1}(1)$, $x_{t+W-2}(2)$, \dots , $x_t(W)$, which only uses $\theta_{t+W-1}, \dots, \theta_t$ respectively (Line 5-7). RHGD skips x_τ if $\tau \notin \{1, \dots, T\}$. Finally, RHGD outputs $x_t(W)$, the W -th update of offline GD.

Algorithm 1: Receding Horizon Gradient Descent (RHGD) [17]

- 1: **Inputs:** Initial decision x_0 ; stepsize η ; initialization oracle ϕ
 - 2: Let $x_1(0) = x_0$.
 - 3: **for** $t = 2 - W, \dots, T$ **do**
 - 4: Initialize $x_{t+W}(0)$ by oracle ϕ if $t + W \leq T$.
 - 5: **for** $\tau = \min(t + W - 1, T)$ **downto** $\max(t, 1)$ **do**
 - 6: Update $x_\tau(t + W - \tau)$ by the offline GD on x_τ in (4).
 - 7: Output $x_t(W)$ when $1 \leq t \leq T$.
-

3.2 Our algorithm: RHIG for inaccurate predictions

Algorithm 2: Receding Horizon Inexact Gradient (RHIG)

- 1: **Inputs:** the length of the lookahead horizon: $W \geq 0$; initial decision x_0 ; stepsize η ; initialization oracle ϕ
- 2: Let $x_1(0) = x_0$.
- 3: **for** $t = 2 - W$ to T **do**
- 4: **if** $t + W \leq T$ **then**
- 5: Compute $x_{t+W}(0)$ by the initialization oracle ϕ with inexact information.
- 6: **for** $\tau = \min(t + W - 1, T)$ **downto** $\max(t, 1)$ **do**
- 7: Compute $x_\tau(t + W - \tau)$ based on the prediction $\theta_{\tau|t-1}$ and the inexact partial gradient:

$$x_\tau(k) = \Pi_{\mathbb{X}}[x_\tau(k-1) - \eta g_\tau(x_{\tau-1:\tau+1}(k-1); \theta_{\tau|t-1})], \quad \text{where } k = t + W - \tau. \quad (5)$$

- 8: Output the decision $x_t(W)$ when $1 \leq t \leq T$.
-

With noisy predictions, it is natural to use the prediction $\theta_{\tau|t-1}$ to estimate the future partial gradients,

$$g_\tau(x_{\tau-1:\tau+1}; \theta_{\tau|t-1}) = \nabla_{x_\tau} f(x_\tau; \theta_{\tau|t-1}) + \nabla_{x_\tau} d(x_\tau, x_{\tau-1}) + \nabla_{x_\tau} d(x_{\tau+1}, x_\tau) \mathbb{1}_{(\tau \leq T-1)},$$

and then updates x_τ by the estimated gradients. This motivates Receding Horizon Inexact Gradient (RHIG) in Algorithm 2. Compared with RHGD, RHIG has the following major differences.

²For instance, if OGD is used as the initialization oracle ϕ , then $x_{t+W}(0) = x_{t+W-1}(0) - \xi_{t+W} \nabla_{x_\tau} f(x_{t+W-1}(0); \theta_{t+W-1})$, where ξ_{t+W} denotes the stepsize.

$x_1(0) = x_0$	$x_2(0); \phi$	$x_3(0); \phi$	$x_4(0); \phi$	$t = -1$
$x_1(1); \theta_{1 -1}$	$x_2(1); \theta_{2 0}$	$x_3(1); \theta_{3 1}$	$x_4(1); \theta_{4 2}$	$t = 0$
$x_1(2); \theta_{1 0}$	$x_2(2); \theta_{2 1}$	$x_3(2); \theta_{3 2}$	$x_4(2); \theta_{4 3}$	$t = 1$
				$t = 2$
				$t = 3$
				$t = 4$

Figure 1: Example: RHIG for $W = 2, T = 4$. (Orange) at $t = -1$, let $x_1(0) = x_0$. (Yellow) at $t = 0$, initialize $x_2(0)$ by ϕ , then compute $x_1(1)$ by inexact offline GD (5) with prediction $\theta_{1|-1} = \theta_{1|0}$. (Green) At $t = 1$, initialize $x_3(0)$ by ϕ , and update $x_2(1)$ and $x_1(2)$ by (5) with $\theta_{2|0}$ and $\theta_{1|0}$ respectively. At $t = 2$, initialize $x_4(0)$ by ϕ , then update $x_3(1), x_2(2)$ by inexact offline GD (5) with $\theta_{3|1}$ and $\theta_{2|1}$ respectively. $t = 3, 4$ are similar. Notice that $\mathbf{x}(1) = (x_1(1), \dots, x_4(1))$ is computed by inexact offline gradient with 2-step-ahead predictions, and $\mathbf{x}(2)$ by 1-step-ahead predictions.

- (Line 1) Unlike RHGD, the lookahead horizon length $W \geq 0$ is tunable in RHIG. When selecting $W = 0$, RHIG does not use any predictions in Line 5-7. When selecting $1 \leq W \leq T$, RHIG utilizes at most W -step-ahead predictions $\{\theta_{\tau|t-1}\}_{\tau=t}^{t+W-1}$ in Line 5-7. Specifically, when $W = T$, RHIG utilizes all the future predictions $\{\theta_{\tau|t-1}\}_{\tau=t}^T$. Interestingly, one can also select $W > T$. In this case, RHIG not only utilizes all the predictions but also conducts more computation based on the initial predictions $\{\theta_{\tau|0}\}_{\tau=1}^T$ at $t \leq 0$ (recall that $\theta_{\tau|t-1} = \theta_{\tau|0}$ when $t \leq 0$). Notably, when $W \rightarrow +\infty$, RHIG essentially solves $\arg \min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; \{\theta_{\tau|0}\}_{\tau=1}^T)$ at $t \leq 0$ to serve as warm starts at $t = 1$.³ The choice of W will be discussed in Section 4-5.
- (Line 5) Notice that the oracle ϕ no longer receives θ_{t+W-1} exactly in RHIG, so OCO algorithms need to be modified here. For example, OGD initializes $x_{t+W}(0)$ by prediction $\theta_{t+W-1|t-1}$:

$$x_\tau(0) = \Pi_{\mathbb{X}}[x_{\tau-1}(0) - \xi_\tau \nabla_{x_{\tau-1}} f(x_{\tau-1}(0); \theta_{\tau-1|t-1})], \quad \text{where } \tau = t + W. \quad (6)$$

Besides, we note that since $\theta_{\tau|t-1}$ is available, OGD (6) can also use $\theta_{\tau|t-1}$ to update $x_\tau(0)$. Similarly, OCO algorithms with predictions, e.g. (A)OMD [21, 23], DMD [33], can be applied.

- (Line 7) Instead of exact offline GD in RHGD, RHIG can be interpreted as inexact offline GD with prediction errors. Especially, (5) can be written as $x_\tau(k) = x_\tau(k-1) - \eta \nabla_{x_\tau} C(\mathbf{x}(k-1); \theta_\tau - \delta_\tau(W-k+1))$ by the definition (1). More compactly, we can write RHIG updates as

$$\mathbf{x}(k) = \Pi_{\mathbb{X}^T} [\mathbf{x}(k-1) - \eta \nabla_{\mathbf{x}} C(\mathbf{x}(k-1); \boldsymbol{\theta} - \boldsymbol{\delta}(W-k+1))], \quad \forall 1 \leq k \leq W, \quad (7)$$

where $\nabla_{\mathbf{x}} C(\mathbf{x}(k-1); \boldsymbol{\theta} - \boldsymbol{\delta}(W-k+1))$ is an inexact version of the gradient $\nabla_{\mathbf{x}} C(\mathbf{x}(k-1); \boldsymbol{\theta})$.

Though the design of RHIG is rather straightforward, both theoretical analysis and numerical experiments show promising performance of RHIG even under poor long-term predictions (Section 4-6). Some intuitions are discussed below. By formula (7), as the iteration number k increases, RHIG employs inexact gradients with shorter-term prediction errors $\boldsymbol{\delta}(W-k+1)$. Since shorter-term predictions are often more accurate than the longer-term ones, RHIG gradually utilizes more accurate gradient information as iterations go on, reducing the optimality gap caused by inexact gradients. Further, the longer-term prediction errors used at the first several iterations are compressed by later gradient updates, especially for strongly convex costs where GD enjoys certain contraction property.

Lastly, with a gradient-based ϕ and a finite W , RHIG only utilizes gradient updates at each t and is thus more computationally efficient than AFHC [1] and CHC [18] that solve multi-stage optimization.

4 General Regret Analysis

This section considers general prediction errors *without* stochastic model assumptions and provides dynamic regret bounds and discussions,⁴ before which is a helping lemma on the properties of $C(\mathbf{x}; \boldsymbol{\theta})$.

Lemma 1. $C(\mathbf{x}; \boldsymbol{\theta})$ is α strongly convex and $L = l_f + 2l_d$ smooth with $\mathbf{x} \in \mathbb{X}^T$ for any $\boldsymbol{\theta} \in \Theta^T$.

³For more discussion on $W > T$, we refer the reader to our supplementary file.

⁴The results in this section can be extended to more general time-varying cost functions, i.e. $f_t(\cdot)$, where the prediction errors will be measured by the difference in the gradients, i.e. $\sup_{\mathbf{x} \in \mathbb{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-k}(\mathbf{x})\|$.

The following theorem provides a general regret bound for RHIG with any initialization oracle ϕ .

Theorem 1 (General Regret Bound). *Under Assumption 1-2, for $W \geq 0$, oracle ϕ , $\eta = \frac{1}{2L}$, we have*

$$\text{Reg}(\text{RHIG}) \leq \frac{2L}{\alpha} \rho^W \text{Reg}(\phi) + \zeta \sum_{k=1}^{\min(W,T)} \rho^{k-1} \|\delta(k)\|^2 + \mathbb{1}_{(W>T)} \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2, \quad (8)$$

where $\rho = 1 - \frac{\alpha}{4L}$, $\zeta = \frac{h^2}{\alpha} + \frac{h^2}{2L}$, $\text{Reg}(\phi) = C(x(0); \theta) - C(x^*; \theta)$ and $x(0)$ is computed by ϕ .

The regret bound (8) consists of three terms. The first term $\frac{2L}{\alpha} \rho^W \text{Reg}(\phi)$ depends on ϕ . The second term $\zeta \sum_{k=1}^{\min(W,T)} \rho^{k-1} \|\delta(k)\|^2$ and the third term $\mathbb{1}_{(W>T)} \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2$ depend on the errors of the predictions used in Algorithm 2 (Line 5-7). Specifically, when $W \leq T$, at most W -step-ahead predictions are used, so the second term involves at most W -step-ahead prediction errors $\{\delta(k)\}_{k=1}^W$ (the third term is irrelevant). When $W > T$, RHIG uses all predictions, so the second term includes all prediction errors $\{\delta(k)\}_{k=1}^T$; besides, RHIG conducts more computation by the initial predictions $\{\theta_{t|0}\}_{t=1}^T$ at $t \leq 0$ (see Section 3), causing the third term on the initial prediction error $\|\delta(T)\|^2$.

An example of ϕ : restarted OGD [34]. For more concrete discussions on the regret bound, we consider a specific ϕ , restarted OGD [34], as reviewed below. Consider an epoch size Δ and divide T stages into $\lceil T/\Delta \rceil$ epochs with size Δ . In each epoch k , restart OGD (6) and let $\xi_t = \frac{4}{\alpha j}$ at $t = k\Delta + j$ for $1 \leq j \leq \Delta$. Similar to [34], we define the variation of the environment as $V_T = \sum_{t=1}^T \sup_{x \in \mathbb{X}} |f(x; \theta_t) - f(x; \theta_{t-1})|$, and consider V_T is known and $1 \leq V_T \leq T$.⁵ To obtain a meaningful regret bound, we impose Assumption 3, where condition i) is common in OCO literature [23, 34, 35] and condition ii) requires a small switching cost under a small change of actions.

Assumption 3. *i) There exists $G > 0$ such that $\|\nabla_x f(x; \theta)\| \leq G$, $\forall x \in \mathbb{X}, \theta \in \Theta$. ii) There exists β such that $0 \leq d(x, x') \leq \frac{\beta}{2} \|x - x'\|^2$.*⁶

Theorem 2 (Regret bound of restarted OGD). *Under Assumption 1-3, consider $T > 2$ and $\Delta = \lceil \sqrt{2T/V_T} \rceil$, the initialization based on restarted OGD described above satisfies the regret bound:*

$$\text{Reg}(\text{OGD}) \leq C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2, \quad (9)$$

where $C_1 = \frac{4G^2}{\alpha} + \frac{32\beta G^2}{\alpha^2} + 16$.

Notice that restarted OGD's regret bound (9) consists of two terms: the first term $C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T})$ is consistent with the original regret bound in [34] for strongly convex costs, which increases with the environment's variation V_T ; the second term depends on the $\min(W, T)$ -step prediction error, which is intuitive since OGD (6) in our setting only has access to the inexact gradient $\nabla_{x_{s-1}} f(x_{s-1}(0); \theta_{s-1|s-W-1})$ predicted by the $\min(W, T)$ -step-ahead prediction $\theta_{s-1|s-W-1}$.⁷

Corollary 1 (RHIG with restarted OGD initialization). *Under the conditions in Theorem 1 and 2, RHIG with ϕ based on restarted OGD satisfies*

$$\begin{aligned} \text{Reg}(\text{RHIG}) \leq & \underbrace{\rho^W \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T})}_{\text{Part I}} \\ & + \underbrace{\frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^W \|\delta(\min(W, T))\|^2 + \sum_{k=1}^{\min(W,T)} \zeta \rho^{k-1} \|\delta(k)\|^2 + \mathbb{1}_{(W>T)} \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2}_{\text{Part II}}. \end{aligned}$$

where $\rho = 1 - \frac{\alpha}{4L}$, $\zeta = \frac{h^2}{\alpha} + \frac{h^2}{2L}$, and $C_1 = \frac{4G^2}{\alpha} + \frac{32\beta G^2}{\alpha^2} + 16$.

⁵This is without loss of generality. When V_T is unknown, we can use doubling tricks and adaptive stepsizes to generate similar bounds [23]. $1 \leq V_T \leq T$ can be enforced by defining a proper θ_0 and by normalization.

⁶Other norms work too, only leading to different constant factors in the regret bounds.

⁷We have this error term because we do not impose the stochastic structures of the gradient errors in [34].

The order of the regret bound. The regret bound in Corollary 1 consists of two parts: Part I involves the variation of the environment V_T ; while Part II consists of the prediction errors $\{\delta(k)\}_{k=1}^{\min(W,T)}$. The regret bound can be written as $\tilde{O}(\rho^W \sqrt{V_T T} + \sum_{k=1}^{\min(W,T)} \rho^{k-1} \|\delta(k)\|^2)$. The prediction errors $\|\delta(k)\|^2$ can be either larger or smaller than V_T as mentioned in [23]. When $V_T = o(T)$ and $\|\delta(k)\|^2 = o(T)$ for $k \leq W$, the regret bound is $o(T)$. As a simple example of sublinear regrets, consider θ_{t-1} as the prediction of θ_{t+k} ($k \geq 0$) at time t , then $\|\delta(k)\|^2 = O(V_T)$ under proper assumptions, so when $V_T = o(T)$, the regret is $o(T)$.

Impact of V_T . The environment variation V_T only shows up in the Part I of the regret bound in Corollary 1. Fixing V_T , notice that Part I decays exponentially with the lookahead window W . This suggests that the impact of the environment variation V_T on the regret bound decays exponentially when one considers a larger lookahead window W , which is intuitive since long-term thinking/planning allows early preparation for changes in the future and thus mitigates the negative impact of the environment variation.

Impact of $\delta(k)$. Part II in Corollary 1 includes the prediction error terms in (9) and in Theorem 1. Notably, for both $W \leq T$ and $W \geq T$, the factor in front of $\|\delta(k)\|^2$ is dominated by ρ^{k-1} for $1 \leq k \leq \min(W, T)$, which decays exponentially with k since $0 \leq \rho < 1$. This suggests that the impact of the total k -step-ahead prediction error $\|\delta(k)\|^2$ decays exponentially with k , which also indicates that our RHIG (implicitly) focuses more on the short-term predictions than the long-term ones. This property benefits RHIG's performance in practice since short-term predictions are usually more accurate and reliable than the long-term ones.

Choices of W . The optimal choice of W depends on the trade-off between V_T and the prediction errors. For more insightful discussions, we consider non-decreasing k -step-ahead prediction errors, i.e. $\|\delta(k)\| \geq \|\delta(k-1)\|$ for $1 \leq k \leq T$ (in practice, longer-term predictions usually suffer worse quality). It can be shown that Part I increases with V_T and Part II increases with the prediction errors. Further, as W increases, Part I decreases but Part II increases.⁸ Thus, when Part I dominates the regret bound, i.e. V_T is large when compared with the prediction errors, selecting a large W reduces the regret bound. On the contrary, when Part II dominates the regret bound, i.e. the prediction errors are large when compared with V_T , a small W is preferred. The choices of W above are quite intuitive: when the environment is drastically changing while the predictions roughly follow the trends, one should use more predictions to prepare for future changes; however, with poor predictions and slowly changing environments, one can ignore most predictions and rely on the understanding of the current environment. Lastly, though we only consider RHIG with restarted OGD, the discussions provide insights for other ϕ .

An upper and a lower bound in a special case. Next, we consider a special case when V_T is much larger than the prediction errors. It can be shown that the optimal regret is obtained when $W \rightarrow +\infty$.

Corollary 2. Consider non-decreasing k -step-ahead prediction errors, i.e. $\|\delta(k)\|^2 \geq \|\delta(k-1)\|^2$ for $1 \leq k \leq T$. When $\sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \geq \frac{2Lh^2\rho + \alpha^2\zeta}{2LC_1(1-\rho)\alpha} \|\delta(T)\|^2$, the regret bound is minimized by letting $W \rightarrow +\infty$. Further, when $W \rightarrow +\infty$, RHIG's regret can be bounded below.

$$\text{Reg}(\text{RHIG}) \leq \frac{\zeta}{1-\rho} \sum_{k=1}^T \rho^{k-1} \|\delta(k)\|^2.$$

Since $\sqrt{V_T T} \log(1 + \sqrt{T/V_T})$ increases with V_T , the condition in Corollary 2 essentially states that V_T is much larger in comparison to all the prediction errors. Interestingly, the bound in Corollary 2 is not affected by V_T , but all prediction errors $\{\|\delta(k)\|^2\}_{k=1}^T$ are involved, though the factor of $\|\delta(k)\|^2$ exponentially decays with k . Next, we show that such dependence on $\|\delta(k)\|^2$ is unavoidable.

Theorem 3 (Lower bound for a special case). For any online algorithm \mathcal{A} , there exists nontrivial $\sum_t f(x_t; \theta_t) + d(x_t, x_{t-1})$ and predictions $\theta_{t|t-k}$ satisfying the condition in Corollary 2, with parameters $\rho_0 = (\frac{\sqrt{L}-\sqrt{\alpha}}{\sqrt{L}+\sqrt{\alpha}})^2$, $\zeta_0 = (\frac{h(1-\sqrt{\rho_0})}{\alpha+\beta})^2 \frac{\alpha(1-2\rho_0)}{2} > 0$, such that the regret satisfies:

$$\text{Reg}(\mathcal{A}) \geq \frac{\zeta_0}{(1-\rho_0)} \sum_{k=1}^T \rho_0^{k-1} \|\delta(k)\|^2.$$

⁸All the monotonicity claims above are verified in the supplementary file and omitted here for brevity.

In Theorem 3, the influence of $\|\delta(k)\|^2$ also decreases exponentially with k , though with a smaller decay factor ρ_0 . It is left as future work to close the gap between ρ and ρ_0 (and between ζ and ζ_0).

5 Stochastic Prediction Errors

In many applications, prediction errors are usually correlated. For example, the predicted market price of tomorrow usually relies on the predicted price of today, which also depends on the price predicted yesterday. Motivated by this, we adopt an insightful and general stochastic model on prediction errors, which was originally proposed in [1]:

$$\delta_t(k) = \theta_t - \theta_{t|t-k} = \sum_{s=t-k+1}^t P(t-s)e_s, \quad \forall 1 \leq k \leq t \quad (10)$$

where $P(s) \in \mathbb{R}^{p \times q}$, $e_1, \dots, e_T \in \mathbb{R}^q$ are independent with zero mean and covariance R_e . Model (10) captures the correlation patterns described above: the errors $\delta_t(k)$ of different predictions on the same parameter θ_t are correlated by sharing common random vectors from $\{e_t, \dots, e_{t-k+1}\}$; and the prediction errors generated at the same stage, i.e. $\theta_{t+k} - \theta_{t+k|t-1}$ for $k \geq 0$, are correlated by sharing common random vectors from $\{e_t, \dots, e_{t+k}\}$. Notably, the coefficient matrix $P(k)$ represents the degree of correlation between the $\delta_t(1)$ and $\delta_t(k)$ and between $\theta_t - \theta_{t|t-1}$ and $\theta_{t+k} - \theta_{t+k|t-1}$.

As discussed in [1, 18], the stochastic model (10) enjoys many applications, e.g. Wiener filters, Kalman filters [36]. For instance, suppose the parameter follows a stochastic linear system: $\theta_t = \gamma\theta_{t-1} + e_t$ with a given θ_0 and random noise $e_t \sim N(0, 1)$. Then $\theta_t = \gamma^k\theta_{t-k} + \sum_{s=t-k+1}^t \gamma^{t-s}e_s$, the optimal prediction of θ_t based on θ_{t-k} is $\theta_{t|t-k} = \gamma^k\theta_{t-k}$, the prediction error $\delta_t(k)$ satisfies the model (10) with $P(t-s) = \gamma^{t-s}$. A large γ causes strong correlation among prediction errors.

Our next theorem bounds the expected regret of RHIG by the degree of correlation $\|P(k)\|_F$.

Theorem 4 (Expected regret bound). *Under Assumption 1-2, $W \geq 0$, $\eta = 1/L$ and initialization ϕ ,*

$$\mathbb{E}[\text{Reg}(\text{RHIG})] \leq \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \sum_{t=0}^{\min(W,T)-1} \zeta \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t - \rho^W}{1 - \rho}$$

where the expectation is taken with respect to $\{e_t\}_{t=1}^T$, $\rho = 1 - \frac{\alpha}{4L}$, $\zeta = \frac{h^2}{\alpha} + \frac{h^2}{2L}$.

The first term in Theorem 4 represents the influence of ϕ while the second term captures the effects of the correlation. We note that the t -step correlation $\|P(t)\|_F^2$ decays exponentially with t in the regret bound, indicating that RHIG efficiently handles the strong correlation among prediction errors.

Next, we provide a regret bound when RHIG employs the restarted OGD oracle as in Section 4.

Corollary 3 (RHIG with restarted OGD). *Under Assumption 1-3, consider the restarted OGD with $\Delta = \lceil \sqrt{2T/\mathbb{E}[V_T]} \rceil$, we obtain*

$$\mathbb{E}[\text{Reg}(\text{RHIG})] \leq \rho^W C_2 \sqrt{\mathbb{E}[V_T] T} \log(1 + \sqrt{T/\mathbb{E}[V_T]}) + \sum_{t=0}^{\min(W,T)-1} \zeta \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t}{1 - \rho},$$

where we define $C_2 = \frac{2LC_1}{\alpha}$.

Notice that large W is preferred with a large environment variation and weakly correlated prediction errors, and vice versa.

Next, we discuss the concentration property. For simplicity, we consider Gaussian vectors $\{e_t\}_{t=1}^T$.⁹

Theorem 5 (Concentration bound). *Under Assumption 1-3, the conditions in Corollary 3, and consider $\mathbb{E}[V_T] = T$. Denote the expected regret bound in Corollary 3 as $\mathbb{E}[\text{RegBdd}]$. Then,*

$$\mathbb{P}(\text{Reg}(\text{RHIG}) \geq \mathbb{E}[\text{RegBdd}] + b) \leq \exp\left(-c \min\left(\frac{b^2}{K^2}, \frac{b}{K}\right)\right), \quad \forall b > 0,$$

where $K = \zeta \sum_{t=0}^{\min(W,T)-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t}{1 - \rho}$ and c is an absolute constant.

⁹Similar results can be obtained for sub-Gaussian random vectors.




Figure 2: (a) and (b): the regrets of RHIG, AFHC and CHC. (c): RHIG’s tracking trajectories.

Theorem 5 shows that the probability of the regret being larger than the expected regret by $b > 0$ decays exponentially with b , indicating a nice concentration property of RHIG. Further, the concentration effect is stronger (i.e. a larger $1/K$) with a smaller degree of correlation $\|P(t)\|_F^2$.

6 Numerical Experiments

We consider online quadrotor tracking of a vertically moving target [37]. We consider (i) a high-level planning problem which is purely online optimization without modeling the physical dynamics; and (ii) a physical tracking problem where simplified quadrotor dynamics are considered [37].

In (i), we consider SOCO: $\min \sum_{t=1}^T \frac{1}{2}(\alpha(x_t - \theta_t)^2 + \beta(x_t - x_{t-1})^2)$, where x_t is quadrotor’s altitude, θ_t is target’s altitude, and $(x_t - x_{t-1})^2$ penalizes a sudden change in the quadrotor’s altitude. The target θ_t follows: $\theta_t = y_t + d_t$, where $y_t = \gamma y_{t-1} + e_t$ is an autoregressive process with noise e_t [38] and $d_t = a \sin(\omega t)$ is a periodic signal. The predictions are the sum of d_t and the optimal predictions of y_t . Notice that a large γ indicates worse long-term predictions. We consider both a small $\gamma = 0.3$ and a large $\gamma = 0.7$ for different levels of errors. We compare RHIG with AFHC [1, 6] and CHC [18]. (See the supplementary material for more details.) Figure 2(a) shows that with small prediction errors, the three algorithms perform similarly well and RHIG is slightly better. Figure 2(b) shows that with large prediction errors, RHIG significantly outperforms AFHC and CHC. Some intuitive explanations are provided below. Firstly, AFHC and CHC are optimization-based methods, while our RHIG is based on gradient descent, which is known to be more robust to errors. Secondly, RHIG implicitly reduces the impact of the (poorer-quality) long-term predictions and focuses more on the (better) short-term ones by using long-term predictions in the first several updates and then using short-term ones in later updates to refine the decisions; while AFHC and CHC treat predictions more equally by taking averages of the optimal solutions computed by both long-term and short-term predictions (see [1, 18] for more details). These two intuitive reasons may explain the better numerical performance of our RHIG when compared with AFHC and CHC.

In (ii), we consider a simplified second-order model of quadrotor vertical flight: $\ddot{x} = k_1 u - g + k_2$, where x, \dot{x}, \ddot{x} are the altitude, velocity and acceleration respectively, u is the control input (motor thrust command), g is the gravitational acceleration, k_1 and k_2 are physical parameters. We consider time discretization and cost function $\sum_{t=1}^T \frac{1}{2}(\alpha(x_t - \theta_t)^2 + \beta u_t^2)$. The target θ_t follows the process in (i), but with a sudden change in d_t at $t_c = 5.6s$, causing large prediction errors at around t_c , which is unknown until t_c . Figure 2(c) plots the quadrotor’s trajectories generated by RHIG with $W = 1, 10$ and shows RHIG’s nice tracking performance even when considering physical dynamics. $W = 10$ performs better first by using more predictions. However, right after t_c , $W = 1$ performs better since the poor prediction quality there degrades the performance. Lastly, the trajectory with $W = 10$ quickly returns to the desired one after t_c , showing the robustness of RHIG to prediction error shocks.

7 Conclusion

This paper studies how to leverage multi-step-ahead noisy predictions in smoothed online convex optimization. We design a gradient-based algorithm RHIG and analyze its dynamic regret under general prediction errors and a stochastic prediction error model. RHIG effectively reduces the impact of multi-step-ahead prediction errors. Future work includes: 1) closing the gap between the upper and the lower bound in Section 4; 2) lower bounds for general cases; 3) online control problems; 4) the convex case analysis without the strong convexity assumption; etc.

Broader Impact

This paper conducts foundational research and theoretical study on online (real-time) decision making problems. In particular, we propose an online algorithm that leverages noisy predictions in online decision making with coupling among stages. This work can be potentially applied to real-time planning problems in e.g. data center management, robotics, smart grids, smart buildings, transportation systems, as well as other online control applications. Though the algorithm is promising and the analysis is insightful, the results are limited by the theoretical assumptions and should be carefully tested and adjusted before being used in real systems. Further, this paper focuses on the efficiency (i.e. the dynamic regret) and does not consider societal issues such as fairness and privacy. Lastly, we see no ethical concerns on this paper.

Acknowledgments and Disclosure of Funding

The work was supported by NSF CAREER 1553407, AFOSR YIP, ONR YIP.

References

- [1] Niangjun Chen, Anish Agarwal, Adam Wierman, Siddharth Barman, and Lachlan LH Andrew. Online convex optimization using predictions. In *ACM SIGMETRICS Performance Evaluation Review*, volume 43, pages 191–204. ACM, 2015.
- [2] Gautam Goel and Adam Wierman. An online algorithm for smoothed regression and lqr control. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2504–2513, 2019.
- [3] Yiheng Lin, Gautam Goel, and Adam Wierman. Online optimization with predictions and non-convex losses. *arXiv preprint arXiv:1911.03827*, 2019.
- [4] Gautam Goel, Yiheng Lin, Haoyuan Sun, and Adam Wierman. Beyond online balanced descent: An optimal algorithm for smoothed online optimization. In *Advances in Neural Information Processing Systems*, pages 1875–1885, 2019.
- [5] Niangjun Chen, Gautam Goel, and Adam Wierman. Smoothed online convex optimization in high dimensions via online balanced descent. In *Conference On Learning Theory*, pages 1574–1594, 2018.
- [6] Minghong Lin, Zhenhua Liu, Adam Wierman, and Lachlan LH Andrew. Online algorithms for geographical load balancing. In *Green Computing Conference (IGCC), 2012 International*, pages 1–10. IEEE, 2012.
- [7] Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking (TON)*, 21(5):1378–1391, 2013.
- [8] Deyun Wang, Hongyuan Luo, Olivier Grunder, Yanbing Lin, and Haixiang Guo. Multi-step ahead electricity price forecasting using a hybrid model based on two-layer decomposition technique and bp neural network optimized by firefly algorithm. *Applied Energy*, 190:390–407, 2017.
- [9] Paulo Cortez, Miguel Rio, Miguel Rocha, and Pedro Sousa. Multi-scale internet traffic forecasting using neural networks and time series methods. *Expert Systems*, 29(2):143–155, 2012.
- [10] Francesco Zanini, David Aienza, Giovanni De Micheli, and Stephen P Boyd. Online convex optimization-based algorithm for thermal management of mpsoCs. In *Proceedings of the 20th symposium on Great lakes symposium on VLSI*, pages 203–208. ACM, 2010.
- [11] Yingying Li, Yujie Tang, Runyu Zhang, and Na Li. Distributed reinforcement learning for decentralized linear quadratic control: A derivative-free policy optimization approach. *arXiv preprint arXiv:1912.09135*, 2019.
- [12] Tomas Baca, Daniel Hert, Giuseppe Loianno, Martin Saska, and Vijay Kumar. Model predictive trajectory tracking and collision avoidance for reliable outdoor deployment of unmanned aerial vehicles. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6753–6760. IEEE, 2018.

- [13] Lingwen Gan, Adam Wierman, Ufuk Topcu, Niangjun Chen, and Steven H Low. Real-time deferrable load control: handling the uncertainties of renewable generation. *ACM SIGMETRICS Performance Evaluation Review*, 41(3):77–79, 2014.
- [14] Jackeline Rios-Torres and Andreas A Malikopoulos. A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps. *IEEE Transactions on Intelligent Transportation Systems*, 18(5):1066–1077, 2016.
- [15] Yingying Li, Xin Chen, and Na Li. Online optimal control with linear dynamics and predictions: Algorithms and regret analysis. In *Advances in Neural Information Processing Systems*, pages 14887–14899, 2019.
- [16] E. Hazan. *Introduction to Online Convex Optimization*. Foundations and Trends(r) in Optimization Series. Now Publishers, 2016.
- [17] Yingying Li, Guannan Qu, and Na Li. Online optimization with predictions and switching costs: Fast algorithms and the fundamental limit. *arXiv preprint arXiv:1801.07780*, 2018.
- [18] Niangjun Chen, Joshua Comden, Zhenhua Liu, Anshul Gandhi, and Adam Wierman. Using predictions in online optimization: Looking forward with an eye on the past. In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*, pages 193–206. ACM, 2016.
- [19] James B. Rawlings David Q. Mayne and Moritz M. Diehl. *Model Predictive Control: Theory, Computation, and Design, 2nd Edition*. Nob Hill Publishing LLC, 2017.
- [20] Eric C Hall and Rebecca M Willett. Dynamical models and tracking regret in online convex programming. In *Proceedings of the 30th International Conference on International Conference on Machine Learning-Volume 28*, pages I–579, 2013.
- [21] Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pages 993–1019, 2013.
- [22] Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems*, pages 3066–3074, 2013.
- [23] Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *Artificial Intelligence and Statistics*, pages 398–406, 2015.
- [24] Oren Anava, Elad Hazan, and Shie Mannor. Online learning for adversaries with memory: price of past mistakes. In *Advances in Neural Information Processing Systems*, pages 784–792, 2015.
- [25] Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Beyond no-regret: Competitive control via online optimization with memory. *arXiv preprint arXiv:2002.05318*, 2020.
- [26] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pages 111–119, 2019.
- [27] Yingying Li, Subhro Das, and Na Li. Online optimal control with affine constraints. *arXiv preprint arXiv:2010.04891*, 2020.
- [28] Gautam Goel, Niangjun Chen, and Adam Wierman. Thinking fast and slow: Optimization decomposition across timescales. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1291–1298. IEEE, 2017.
- [29] Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.
- [30] Yingying Li and Na Li. Online learning for markov decision processes in nonstationary environments: A dynamic regret analysis. In *2019 American Control Conference (ACC)*, pages 1232–1237. IEEE, 2019.
- [31] Yingying Li, Aoxiao Zhong, Guannan Qu, and Na Li. Online markov decision processes with time-varying transition probabilities and rewards. In *ICML workshop on Real-world Sequential Decision Making*, 2019.
- [32] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.

- [33] Eric C Hall and Rebecca M Willett. Online convex optimization in dynamic environments. *IEEE Journal of Selected Topics in Signal Processing*, 9(4):647–662, 2015.
- [34] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.
- [35] Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 7195–7201. IEEE, 2016.
- [36] Thomas Kailath, Ali H Sayed, and Babak Hassibi. *Linear estimation*. Number BOOK. Prentice Hall, 2000.
- [37] Jaime F Fisac, Anayo K Akametalu, Melanie N Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2018.
- [38] James D Hamilton. *Time series analysis*, volume 2. Princeton New Jersey, 1994.
- [39] Mark Rudelson, Roman Vershynin, et al. Hanson-wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18, 2013.

Appendices

The appendices provide additional discussions and proofs of the theoretical results. In particular, Appendix A provides additional discussions on RHIG when $W > T$; next, Appendix B provides a proof of Lemma 1; Appendix C provides a proof of Theorem 1; Appendix D provides a proof of Theorem 2, a proof of Corollary 1, and also proves the claimed properties of the regret bound in Section 4; Appendix E discusses the special case and proves Corollary 2 and Theorem 3; Appendix F considers the stochastic prediction errors and proves Theorem 4, Corollary 3 and Theorem 5; finally, Appendix G provides additional discussions on the numerical experiments.

A Additional Discussions on RHIG when $W > T$

As mentioned in Section 3, in RHIG, one can select the look-ahead horizon $W > T$. To further illustrate this case, we provide an example of RHIG for $T = 3$ and $W = 5$ below.

Further, we explain the case when $W \rightarrow +\infty$ by using the example in Figure 3 for $T = 3$. Similar to the discussion for Figure 3, for general $W > T$, it can be verified that the computed variables at $t = 0$ are $x_1(W - 1), x_2(W - 2), x_3(W - 3)$, which can be viewed as the iterated variables of offline (exact) gradient descent (4) under parameters $(\theta_{1|0}, \theta_{2|0}, \theta_{3|0})$ after $W - 1, W - 2, W - 3$ iterations respectively. For $W \rightarrow +\infty$, $(x_1(W - 1), x_2(W - 2), x_3(W - 3))$ converges to the optimal solution to $\min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; (\theta_{1|0}, \theta_{2|0}, \theta_{3|0}))$. Then, at $t = 1$, RHIG conducts one inexact gradient update for each $x_3(W - 2), x_2(W - 1), x_1(W)$ based on predictions $(\theta_{1|0}, \theta_{2|0}, \theta_{3|0})$ and outputs $x_1(W)$. (When $W \rightarrow +\infty$, $x_1(W)$ is not updated at $t = 1$ since it has converged to $\min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; (\theta_{1|0}, \theta_{2|0}, \theta_{3|0}))$.) At $t = 2$, RHIG conducts one inexact gradient update for both $x_3(W - 1)$ and $x_2(W)$ based on *new* prediction $\theta_{3|1}$ and $\theta_{2|2}$. At $t = 3$, RHIG conducts one exact gradient update for $x_3(W)$ based on *new* prediction $\theta_{3|2}$. This explains the scenario when $W \rightarrow +\infty$.

B Proof of Lemma 1

Firstly, we prove the strong convexity. Since $f(x_t; \theta_t)$ is α -strongly convex with respect to x_t , $\sum_{t=1}^T f(x_t; \theta_t)$ is α -strongly convex with respect to $\mathbf{x} = (x_1^\top, \dots, x_T^\top)^\top$. Since $d(x_t, x_{t-1})$ is convex with respect to (x_t, x_{t-1}) , $\sum_{t=1}^T d(x_t, x_{t-1})$ is also convex with respect to \mathbf{x} . Consequently, $C(\mathbf{x}; \boldsymbol{\theta}) = \sum_{t=1}^T (f(x_t; \theta_t) + d(x_t, x_{t-1}))$ is α -strongly convex with respect to \mathbf{x} .

Next, we prove the smoothness. For any $x_t, y_t \in \mathbb{X}$, by the l_f -smoothness of $f(x_t; \theta_t)$ for all t , we have

$$f(y_t) \leq f(x_t) + \langle \nabla_{x_t} f(x_t; \theta_t), y_t - x_t \rangle + \frac{l_f}{2} \|x_t - y_t\|^2$$

$x_1(0) = x_0$	$x_2(0); \phi$	$x_3(0); \phi$	$t = -4$
$x_1(1); \theta_{1 -4}$	$x_2(1); \theta_{2 -3}$	$x_3(1); \theta_{3 -2}$	$t = -3$
$x_1(2); \theta_{1 -3}$	$x_2(2); \theta_{2 -2}$	$x_3(2); \theta_{3 -1}$	$t = -2$
$x_1(3); \theta_{1 -2}$	$x_2(3); \theta_{2 -1}$	$x_3(3); \theta_{3 0}$	$t = -1$
$x_1(4); \theta_{1 -1}$	$x_2(4); \theta_{2 0}$	$x_3(4); \theta_{3 1}$	$t = 0$
$x_1(5); \theta_{1 0}$	$x_2(5); \theta_{2 1}$	$x_3(5); \theta_{3 2}$	$t = 1$
			$t = 2$
			$t = 3$

Figure 3: Example of RHIG when $W = 5 > T = 3$. (Pink) At $t = 1 - W = -4$, let $x_1(0) = x_0$. (Orange) At $t = -3$, initialize $x_2(0)$ by ϕ , then compute $x_1(1)$ by inexact offline GD (5) with prediction $\theta_{1|t-1} = \theta_{1|-4} = \theta_{1|0}$. (Yellow) At $t = -2$, initialize $x_3(0)$ by ϕ , and update $x_2(1)$ and $x_1(2)$ by (5) with $\theta_{2|-3} = \theta_{2|0}$ and $\theta_{1|-3} = \theta_{1|0}$ respectively. (Green) At $t = -1$, update $x_3(1)$, $x_2(2)$, $x_1(3)$ by inexact offline GD (5) with $\theta_{3|-2} = \theta_{3|0}$, $\theta_{2|-2} = \theta_{2|0}$, and $\theta_{1|-2} = \theta_{1|0}$ respectively. (Dark green) At $t = 0$, update $x_3(2)$, $x_2(3)$, $x_1(4)$ by inexact offline GD (5) with $\theta_{3|-1} = \theta_{3|0}$, $\theta_{2|-1} = \theta_{2|0}$, and $\theta_{1|12} = \theta_{1|0}$ respectively. (Blue) At $t = 1$, update $x_3(3)$, $x_2(4)$, $x_1(5)$ by inexact offline GD (5) with $\theta_{3|0}$, $\theta_{2|0}$, and $\theta_{1|0}$ respectively. Then output $x_1(5)$. (Purple) At $t = 2$, update $x_3(4)$ and $x_2(5)$ by inexact offline GD (5) with $\theta_{3|1}$, $\theta_{2|1}$ respectively and output $x_2(5)$. (Red) At $t = 3$, update $x_3(5)$ by inexact offline GD (5) with $\theta_{3|2}$ and output $x_3(5)$. By recalling that $\theta_{t|\tau} = \theta_{t|0}$ when $\tau < 0$, we note that all the computation at $t \leq 0$ (above the red lines) is based on initial predictions $\{\theta_{1|0}, \theta_{2|0}, \theta_{3|0}\}$. Therefore, the computed variables $x_1(4)$, $x_2(3)$, $x_3(2)$ at $t = 0$ can be viewed as the iterated variables of offline (exact) gradient descent (4) under parameters $\{\theta_{1|0}, \theta_{2|0}, \theta_{3|0}\}$ after 4,3,2 iterations respectively.

By the l_d -smoothness of $d(x_t, x_{t-1})$, for any $x_t, y_t, x_{t-1}, y_{t-1} \in \mathbb{X}$, we have

$$d(y_t, y_{t-1}) \leq d(x_t, x_{t-1}) + \langle \nabla_{x_t} d(x_t, x_{t-1}), y_t - x_t \rangle + \langle \nabla_{x_{t-1}} d(x_t, x_{t-1}), y_{t-1} - x_{t-1} \rangle + \frac{l_d}{2} (\|y_t - x_t\|^2 + \|y_{t-1} - x_{t-1}\|^2)$$

for $t \geq 2$ and $d(y_1, x_0) \leq d(x_1, x_0) + \langle \nabla_{x_1} d(x_1, x_0), y_1 - x_1 \rangle + \frac{l_d}{2} \|x_1 - y_1\|^2$ for $t = 1$.

Therefore, for any $\mathbf{x}, \mathbf{y} \in \mathbb{X}^T$, by summing the smoothness inequalities above over $t = 1, \dots, T$, we obtain

$$C(\mathbf{y}; \boldsymbol{\theta}) \leq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}), \mathbf{y} - \mathbf{x} \rangle + \frac{l_f + 2l_d}{2} \|\mathbf{x} - \mathbf{y}\|^2$$

where we used the fact that $\nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta})$ is composed of partial gradients $\nabla_{x_t} C(\mathbf{x}; \boldsymbol{\theta}) = \nabla_{x_t} f(x_t; \theta_t) + \nabla_{x_t} d(x_t, x_{t-1}) + \mathbb{1}_{(t < T)} \cdot \nabla_{x_t} d(x_{t+1}, x_t)$.

C Proof of Theorem 1

Consider the offline optimization with parameter $\boldsymbol{\theta}$, i.e. $\min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; \boldsymbol{\theta})$. As mentioned in Section 3, RHIG can be interpreted as projected gradient descent on $C(\mathbf{x}; \boldsymbol{\theta})$ with inexact gradients:

$$\mathbf{x}(k+1) = \Pi_{\mathbb{X}^T} [\mathbf{x}(k) - \eta \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k))] \quad (11)$$

where the exact gradient should be $\nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta})$ but the parameter prediction error $\boldsymbol{\delta}(W - k)$ results in inexact gradient $\nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k))$. Notice that when $W - k > T$, by our definition, we have $\boldsymbol{\delta}(W - k) = \boldsymbol{\delta}(T)$.

Consequently, the regret bound of RHIG can be proved based on the convergence analysis of the projected gradient descent with inexact gradients. We note that unlike the classic inexact gradient where the gradient errors are uniformly bounded, RHIG's inexact gradients (11) have different gradient errors at different iterations, thus calling for slightly different convergence analysis.

In the following, we first provide some supportive lemmas, then provide a rigorous proof of Theorem 1.

C.1 Supportive Lemmas

Firstly, we provide a bound on the gradient errors with respect to the errors on the parameters.

Lemma 2 (Gradient prediction error bound). *For any true parameter $\boldsymbol{\theta} \in \Theta^T$ and the predicted parameter $\boldsymbol{\theta}' \in \Theta^T$, the error of the predicted gradient can be bounded below.*

$$\|\nabla_{\mathbf{x}}C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{\mathbf{x}}C(\mathbf{x}; \boldsymbol{\theta})\|^2 \leq h^2\|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2, \quad \forall \mathbf{x} \in \Theta^T$$

Proof. Firstly, we consider the gradient with respect to each stage variable x_t , which is provided by

$$\nabla_{x_t}C(\mathbf{x}; \boldsymbol{\theta}) = \nabla_{x_t}f(x_t; \theta_t) + \nabla_{x_t}d(x_t, x_{t-1}) + \nabla_{x_t}d(x_{t+1}, x_t)\mathbb{1}_{(t \leq T-1)}$$

Noticing that $d(x_t, x_{t-1})$ does not depend on the parameter $\boldsymbol{\theta}$, we obtain the prediction error bound of gradient with respect to x_t as follows.

$$\begin{aligned} \|\nabla_{x_t}C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{x_t}C(\mathbf{x}; \boldsymbol{\theta})\| &= \|\nabla_{x_t}f(x_t; \theta'_t) - \nabla_{x_t}f(x_t; \theta_t)\| \\ &\leq h\|\theta'_t - \theta_t\| \end{aligned}$$

Therefore, the prediction error of the full gradient can be bounded as follows,

$$\begin{aligned} \|\nabla_{\mathbf{x}}C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{\mathbf{x}}C(\mathbf{x}; \boldsymbol{\theta})\|^2 &= \sum_{t=1}^T \|\nabla_{x_t}C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{x_t}C(\mathbf{x}; \boldsymbol{\theta})\|^2 \\ &\leq h^2 \sum_{t=1}^T \|\theta'_t - \theta_t\|^2 = h^2\|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2 \end{aligned}$$

which completes the proof. \square

Next, we provide an equivalent characterization of the projected gradient update with respect to inexact parameters.

Lemma 3 (A representation of inexact projected gradient updates). *For any predicted parameter $\boldsymbol{\theta}'$ and any stepsize η , the projected gradient descent with predicted parameter $\mathbf{x}(k+1) = \Pi_{\mathbb{X}^T}[\mathbf{x}(k) - \eta\nabla_{\mathbf{x}}C(\mathbf{x}(k); \boldsymbol{\theta}')] is equivalent to the following representation.$*

$$\mathbf{x}(k+1) = \arg \min_{\mathbf{x} \in \mathbb{X}^T} \left\{ \langle \nabla_{\mathbf{x}}C(\mathbf{x}(k); \boldsymbol{\theta}'), \mathbf{x} - \mathbf{x}(k) \rangle + \frac{1}{2\eta} \|\mathbf{x} - \mathbf{x}(k)\|^2 \right\}$$

Proof. By the definition of projection, the projected gradient descent with predicted parameter is equivalent to the following.

$$\begin{aligned} \mathbf{x}(k+1) &= \arg \min_{\mathbf{x} \in \mathbb{X}^T} \{ \|\mathbf{x} - \mathbf{x}(k) + \eta\nabla_{\mathbf{x}}C(\mathbf{x}(k); \boldsymbol{\theta}')\|^2 \} \\ &= \arg \min_{\mathbf{x} \in \mathbb{X}^T} \{ \|\mathbf{x} - \mathbf{x}(k)\|^2 + \eta^2 \|\nabla_{\mathbf{x}}C(\mathbf{x}(k); \boldsymbol{\theta}')\|^2 + 2\eta \langle \nabla_{\mathbf{x}}C(\mathbf{x}(k); \boldsymbol{\theta}'), \mathbf{x} - \mathbf{x}(k) \rangle \} \\ &= \arg \min_{\mathbf{x} \in \mathbb{X}^T} \left\{ \frac{1}{2\eta} \|\mathbf{x} - \mathbf{x}(k)\|^2 + \langle \nabla_{\mathbf{x}}C(\mathbf{x}(k); \boldsymbol{\theta}'), \mathbf{x} - \mathbf{x}(k) \rangle \right\} \end{aligned}$$

where the last equality uses the fact that $\eta^2 \|\nabla_{\mathbf{x}}C(\mathbf{x}(k); \boldsymbol{\theta}')\|^2$ does not depend on \mathbf{x} . \square

Lastly, we provide a strong-convexity-type inequality and a smoothness-type inequality under inexact gradients. Both inequalities suffer from additional error terms caused by the parameter prediction error.

Lemma 4 (Strong convexity inequality with errors). *Consider optimization $\min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; \boldsymbol{\theta})$. For any $\mathbf{x}, \mathbf{y} \in \mathbb{X}^T$, for any inexact parameter $\boldsymbol{\theta}'$ and the resulting inexact gradient $\nabla_{\mathbf{x}}C(\mathbf{x}; \boldsymbol{\theta}')$, we have*

$$C(\mathbf{y}; \boldsymbol{\theta}) \geq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}}C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle + \frac{\alpha}{4} \|\mathbf{x} - \mathbf{y}\|^2 - \frac{h^2}{\alpha} \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2$$

Proof. By the strong convexity of $C(\mathbf{x}; \boldsymbol{\theta})$, for any $\mathbf{x}, \mathbf{y} \in \mathbb{X}^T$ and any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta^T$, we obtain the following.

$$\begin{aligned}
C(\mathbf{y}; \boldsymbol{\theta}) &\geq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}), \mathbf{y} - \mathbf{x} \rangle + \frac{\alpha}{2} \|\mathbf{y} - \mathbf{x}\|^2 \\
&= C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle - \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}), \mathbf{y} - \mathbf{x} \rangle + \frac{\alpha}{2} \|\mathbf{y} - \mathbf{x}\|^2 \\
&\geq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle - \|\nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta})\| \|\mathbf{y} - \mathbf{x}\| + \frac{\alpha}{2} \|\mathbf{y} - \mathbf{x}\|^2 \\
&\geq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle - \frac{1}{\alpha} \|\nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta})\|^2 + \frac{\alpha}{4} \|\mathbf{y} - \mathbf{x}\|^2 \\
&\geq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle - \frac{h^2}{\alpha} \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2 + \frac{\alpha}{4} \|\mathbf{y} - \mathbf{x}\|^2
\end{aligned}$$

□

Lemma 5 (Smoothness inequality with errors). *Consider optimization $\min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; \boldsymbol{\theta})$. For any $\mathbf{x}, \mathbf{y} \in \mathbb{X}^T$, for any inexact parameter $\boldsymbol{\theta}'$ and the resulting inexact gradient $\nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}')$, we have*

$$C(\mathbf{y}; \boldsymbol{\theta}) \leq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle + L \|\mathbf{x} - \mathbf{y}\|^2 + \frac{h^2}{2L} \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2$$

Proof. By the smoothness of $C(\mathbf{x}; \boldsymbol{\theta})$, for any $\mathbf{x}, \mathbf{y} \in \mathbb{X}^T$ and any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta^T$, we obtain the following.

$$\begin{aligned}
C(\mathbf{y}; \boldsymbol{\theta}) &\leq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2 \\
&= C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}) - \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2 \\
&\leq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle + \|\nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta})\| \|\mathbf{y} - \mathbf{x}\| + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2 \\
&\leq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle + \frac{1}{2L} \|\nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}') - \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta})\|^2 + L \|\mathbf{y} - \mathbf{x}\|^2 \\
&\leq C(\mathbf{x}; \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}; \boldsymbol{\theta}'), \mathbf{y} - \mathbf{x} \rangle + \frac{h^2}{2L} \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2 + L \|\mathbf{y} - \mathbf{x}\|^2
\end{aligned}$$

□

C.2 Proof of Theorem 1

According to Algorithm 2 and the definition of the regret, we have $\text{Reg}(RHIG) = C(\mathbf{x}(W); \boldsymbol{\theta}) - C(\mathbf{x}^*; \boldsymbol{\theta})$ and $\text{Reg}(\phi) = C(\mathbf{x}(0); \boldsymbol{\theta}) - C(\mathbf{x}^*; \boldsymbol{\theta})$, where $\mathbf{x}^* = \arg \min_{\mathbb{X}^T} C(\mathbf{x}; \boldsymbol{\theta})$. For notational simplicity, we denote $r_k = \|\mathbf{x}(k) - \mathbf{x}^*\|^2$.

Step 1: bound $\text{Reg}(RHIG)$ with r_{W-1} .

$$\begin{aligned}
C(\mathbf{x}(W); \boldsymbol{\theta}) &\leq C(\mathbf{x}(W-1); \boldsymbol{\theta}) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}(W-1); \boldsymbol{\theta} - \boldsymbol{\delta}(1)), \mathbf{x}(W) - \mathbf{x}(W-1) \rangle \\
&\quad + L \|\mathbf{x}(W) - \mathbf{x}(W-1)\|^2 + \frac{h^2}{2L} \|\boldsymbol{\delta}(1)\|^2 \\
&= \min_{\mathbf{x} \in \mathbb{X}^T} \{ \langle \nabla_{\mathbf{x}} C(\mathbf{x}(W-1); \boldsymbol{\theta} - \boldsymbol{\delta}(1)), \mathbf{x} - \mathbf{x}(W-1) \rangle + L \|\mathbf{x} - \mathbf{x}(W-1)\|^2 \} \\
&\quad + C(\mathbf{x}(W-1); \boldsymbol{\theta}) + \frac{h^2}{2L} \|\boldsymbol{\delta}(1)\|^2 \\
&\leq \langle \nabla_{\mathbf{x}} C(\mathbf{x}(W-1); \boldsymbol{\theta} - \boldsymbol{\delta}(1)), \mathbf{x}^* - \mathbf{x}(W-1) \rangle + L \|\mathbf{x}^* - \mathbf{x}(W-1)\|^2 \\
&\quad + C(\mathbf{x}(W-1); \boldsymbol{\theta}) + \frac{h^2}{2L} \|\boldsymbol{\delta}(1)\|^2 \\
&\leq C(\mathbf{x}^*; \boldsymbol{\theta}) + (L - \frac{\alpha}{4}) r_{W-1} + \left(\frac{h^2}{\alpha} + \frac{h^2}{2L} \right) \|\boldsymbol{\delta}(1)\|^2
\end{aligned}$$

where we used Lemma 5 in the first inequality, Lemma 3 and $\eta = \frac{1}{2L}$ in the first equality, Lemma 4 in the last inequality. By rearranging terms, we obtain

$$\text{Reg}(RHIG) = C(\mathbf{x}(W); \boldsymbol{\theta}) - C(\mathbf{x}^*; \boldsymbol{\theta}) \leq L\rho r_{W-1} + \zeta \|\boldsymbol{\delta}(1)\|^2 \quad (12)$$

where $\rho = 1 - \frac{\alpha}{4L}$, $\zeta = \frac{h^2}{\alpha} + \frac{h^2}{2L}$.

Step 2: a recursive inequality between r_{k+1} and r_k .

In the following, we will show that

$$r_{k+1} \leq \rho r_k + \frac{\zeta}{L} \|\boldsymbol{\delta}(W - k)\|^2, \quad \forall 0 \leq k \leq W - 1 \quad (13)$$

Firstly, by (11), $\eta = \frac{1}{2L}$, Lemma 3 and its first-order optimality condition, we have

$$\langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)) + 2L(\mathbf{x}(k+1) - \mathbf{x}(k)), \mathbf{x} - \mathbf{x}(k+1) \rangle \geq 0, \quad \forall \mathbf{x} \in \mathbb{X}^T$$

By substituting $\mathbf{x} = \mathbf{x}^*$ and rearranging terms, we obtain

$$\frac{1}{2L} \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}^* - \mathbf{x}(k+1) \rangle \geq \langle \mathbf{x}(k+1) - \mathbf{x}(k), \mathbf{x}(k+1) - \mathbf{x}^* \rangle \quad (14)$$

Next, we will derive the recursive inequality (13) by using (14).

$$\begin{aligned} r_{k+1} &= \|\mathbf{x}(k+1) - \mathbf{x}^*\|^2 = \|\mathbf{x}(k+1) - \mathbf{x}(k) + \mathbf{x}(k) - \mathbf{x}^*\|^2 \\ &= r_k - \|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2 + 2\langle \mathbf{x}(k+1) - \mathbf{x}(k), \mathbf{x}(k+1) - \mathbf{x}^* \rangle \\ &\leq r_k - \|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2 + \frac{1}{L} \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}^* - \mathbf{x}(k+1) \rangle \\ &= r_k - \|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2 + \frac{1}{L} \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}^* - \mathbf{x}(k) \rangle \\ &\quad + \frac{1}{L} \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}(k) - \mathbf{x}(k+1) \rangle \\ &= r_k + \frac{1}{L} \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}^* - \mathbf{x}(k) \rangle \\ &\quad - \frac{1}{L} (\langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}(k+1) - \mathbf{x}(k) \rangle + L\|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2) \\ &\leq r_k + \frac{1}{L} \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}^* - \mathbf{x}(k) \rangle \\ &\quad - \frac{1}{L} \left(C(\mathbf{x}(k+1); \boldsymbol{\theta}) - C(\mathbf{x}(k); \boldsymbol{\theta}) - \frac{h^2}{2L} \|\boldsymbol{\delta}(W - k)\|^2 \right) \\ &\leq r_k + \frac{1}{L} \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}^* - \mathbf{x}(k) \rangle \\ &\quad - \frac{1}{L} (C(\mathbf{x}^*; \boldsymbol{\theta}) - C(\mathbf{x}(k); \boldsymbol{\theta})) + \frac{h^2}{2L^2} \|\boldsymbol{\delta}(W - k)\|^2 \\ &= r_k - \frac{1}{L} (C(\mathbf{x}^*; \boldsymbol{\theta}) - C(\mathbf{x}(k); \boldsymbol{\theta})) + \langle \nabla_{\mathbf{x}} C(\mathbf{x}(k); \boldsymbol{\theta} - \boldsymbol{\delta}(W - k)), \mathbf{x}(k) - \mathbf{x}^* \rangle \\ &\quad + \frac{h^2}{2L^2} \|\boldsymbol{\delta}(W - k)\|^2 \\ &\leq r_k - \frac{1}{L} \left(\frac{\alpha}{4} \|\mathbf{x}(k) - \mathbf{x}^*\|^2 - \frac{h^2}{\alpha} \|\boldsymbol{\delta}(W - k)\|^2 \right) + \frac{h^2}{2L^2} \|\boldsymbol{\delta}(W - k)\|^2 \\ &= \rho r_k + \frac{\zeta}{L} \|\boldsymbol{\delta}(W - k)\|^2 \end{aligned}$$

which completes the proof of (13).

Step 3: completing the proof by (13) and (12).

By summing (13) over $k = 0, \dots, W - 2$, we obtain

$$r_{W-1} \leq \rho^{W-1} r_0 + \frac{\zeta}{L} (\|\boldsymbol{\delta}(2)\|^2 + \rho \|\boldsymbol{\delta}(3)\|^2 + \dots + \rho^{W-2} \|\boldsymbol{\delta}(W)\|^2)$$

$$\leq \rho^{W-1} \frac{2}{\alpha} (C(\mathbf{x}(0); \boldsymbol{\theta}) - C(\mathbf{x}^*; \boldsymbol{\theta})) + \frac{\zeta}{L} \sum_{k=2}^W \rho^{k-2} \|\boldsymbol{\delta}(k)\|^2$$

By (12), we obtain the regret bound in Theorem 1:

$$\begin{aligned} \text{Reg}(RHIG) &\leq L\rho \rho^{W-1} \frac{2}{\alpha} (C(\mathbf{x}(0); \boldsymbol{\theta}) - C(\mathbf{x}^*; \boldsymbol{\theta})) + \frac{\zeta}{L} \sum_{k=2}^W \rho^{k-2} \|\boldsymbol{\delta}(k)\|^2 + \zeta \|\boldsymbol{\delta}(1)\|^2 \\ &= \frac{2L}{\alpha} \rho^W \text{Reg}(\phi) + \zeta \sum_{k=1}^W \rho^{k-1} \|\boldsymbol{\delta}(k)\|^2 \\ &= \frac{2L}{\alpha} \rho^W \text{Reg}(\phi) + \zeta \sum_{k=1}^{\min(W,T)} \rho^{k-1} \|\boldsymbol{\delta}(k)\|^2 + \zeta \mathbb{1}_{(W>T)} \sum_{k=T+1}^W \rho^{k-1} \|\boldsymbol{\delta}(T)\|^2 \\ &= \frac{2L}{\alpha} \rho^W \text{Reg}(\phi) + \zeta \sum_{k=1}^{\min(W,T)} \rho^{k-1} \|\boldsymbol{\delta}(k)\|^2 + \zeta \mathbb{1}_{(W>T)} \frac{\rho^T - \rho^W}{1 - \rho} \|\boldsymbol{\delta}(T)\|^2 \end{aligned}$$

where we used the fact that $\|\boldsymbol{\delta}(k)\| = \|\boldsymbol{\delta}(T)\|$ when $k > T$.

D Proofs of Theorem 2, Corollary 1 and the claimed properties of the regret bound in Section 4

In this section, we provide a dynamic regret bound for the restarted OGD initialization rule in Section 4, based on which we prove Corollary 1. To achieve this, we will first establish a static regret bound for OGD initialization (6). The proof is inspired by [34].

For notational simplicity, we slightly abuse the notation and let x_t denote $x_t(0)$ generated by OGD. Further, by the definition of the prediction errors $\delta_{t-1}(W)$ for $W \geq 1$, we can write the initialization rule (6) as the following, which can be interpreted as OGD with inexact gradients:

$$x_t = \Pi_{\mathbb{X}}[x_{t-1} - \xi_t \nabla_x f(x_{t-1}; \theta_{t-1} - \delta_{t-1}(\min(W, T)))], \quad t \geq 2; \quad (15)$$

and $x_1 = x_0$. Here, we used the facts that $\theta_{t-1|t-W-1} = \theta_{t-1} - \delta_{t-1}(W)$ and $\delta_{t-1}(W) = \delta_{t-1}(T)$ for $W > T$.

D.1 Static regret bound for OGD with inexact gradients

In this section, we consider the OGD with inexact gradients (15) with diminishing stepsize $\xi_t = \frac{4}{\alpha t}$ for $t \geq 1$. We will prove its static regret bound below.

Theorem 6 (Static regret of OGD with inexact gradients). *Consider the OGD with inexact gradients (15) with diminishing stepsize $\xi_t = \frac{4}{\alpha t}$ for $t \geq 1$ and any x_0 . Then, for $z^* = \arg \min_{z \in \mathbb{X}} \sum_{t=1}^T f(z; \theta_t)$, we have the following static regret bound:*

$$\sum_{t=1}^T [f(x_t; \theta_t) - f(z^*; \theta_t)] \leq \frac{2G^2}{\alpha} \log(T+1) + \sum_{t=1}^T \frac{h^2}{\alpha} \|\delta_t(\min(W, T))\|^2$$

Further, the total switching cost can be bounded by:

$$\sum_{t=1}^T d(x_t, x_{t-1}) \leq \frac{16G^2\beta}{\alpha^2}$$

Proof. Firstly, we prove the static regret bound. Define $q_t = \|x_t - z^*\|^2$. Then, for $t \geq 1$, we have the following.

$$\begin{aligned} q_{t+1} &= \|x_{t+1} - z^*\|^2 \leq \|x_t - \xi_{t+1} \nabla_x f(x_t; \theta_t - \delta_t(\min(W, T))) - z^*\|^2 \\ &= q_t + \xi_{t+1}^2 \|\nabla_x f(x_t; \theta_t - \delta_t(\min(W, T)))\|^2 - 2\xi_{t+1} \langle x_t - z^*, \nabla_x f(x_t; \theta_t - \delta_t(\min(W, T))) \rangle \end{aligned}$$

$$\begin{aligned} &\leq q_t + \xi_{t+1}^2 G^2 - 2\xi_{t+1} \langle x_t - z^*, \nabla_x f(x_t; \theta_t) \rangle \\ &\quad - 2\xi_{t+1} \langle x_t - z^*, \nabla_x f(x_t; \theta_t - \delta_t(\min(W, T))) - \nabla_x f(x_t; \theta_t) \rangle \end{aligned}$$

where the last inequality uses Assumption 3(i). By rearranging terms, we obtain

$$\langle x_t - z^*, \nabla_x f(x_t; \theta_t) \rangle \leq \frac{q_t - q_{t+1}}{2\xi_{t+1}} + \frac{\xi_{t+1}}{2} G^2 - \langle x_t - z^*, \nabla_x f(x_t; \theta_t - \delta_t(\min(W, T))) - \nabla_x f(x_t; \theta_t) \rangle \quad (16)$$

By the strong convexity of $f(x; \theta_t)$, we have $f(z^*; \theta_t) \geq f(x_t; \theta_t) + \langle z^* - x_t, \nabla_x f(x_t; \theta_t) \rangle + \frac{\alpha}{2} \|z^* - x_t\|^2$. By rearranging terms and by (16), we obtain

$$\begin{aligned} f(x_t; \theta_t) - f(z^*; \theta_t) &\leq \langle x_t - z^*, \nabla_x f(x_t; \theta_t) \rangle - \frac{\alpha}{2} \|z^* - x_t\|^2 \\ &\leq \frac{q_t - q_{t+1}}{2\xi_{t+1}} + \frac{\xi_{t+1}}{2} G^2 - \langle x_t - z^*, \nabla_x f(x_t; \theta_t - \delta_t(\min(W, T))) - \nabla_x f(x_t; \theta_t) \rangle - \frac{\alpha}{2} \|z^* - x_t\|^2 \\ &\leq \frac{q_t - q_{t+1}}{2\xi_{t+1}} + \frac{\xi_{t+1}}{2} G^2 + \|x_t - z^*\| \|\nabla_x f(x_t; \theta_t - \delta_t(\min(W, T))) - \nabla_x f(x_t; \theta_t)\| - \frac{\alpha}{2} \|z^* - x_t\|^2 \\ &\leq \frac{q_t - q_{t+1}}{2\xi_{t+1}} + \frac{\xi_{t+1}}{2} G^2 + \frac{1}{\alpha} \|\nabla_x f(x_t; \theta_t - \delta_t(\min(W, T))) - \nabla_x f(x_t; \theta_t)\|^2 - \frac{\alpha}{4} \|z^* - x_t\|^2 \\ &\leq \frac{q_t - q_{t+1}}{2\xi_{t+1}} + \frac{\xi_{t+1}}{2} G^2 + \frac{h^2}{\alpha} \|\delta_t(\min(W, T))\|^2 - \frac{\alpha}{4} q_t \end{aligned}$$

where we used $ab \leq \frac{\epsilon}{2} a^2 + \frac{1}{2\epsilon} b^2$ for any $a, b \in \mathbb{R}$ and any $\epsilon > 0$ in the second last inequality and Assumption 2 in the last inequality. By summing over $t = 1, \dots, T$, we obtain

$$\begin{aligned} \sum_{t=1}^T [f(x_t; \theta_t) - f(z^*; \theta_t)] &\leq \sum_{t=2}^T \left(\frac{1}{2\xi_{t+1}} - \frac{1}{2\xi_t} - \frac{\alpha}{4} \right) q_t + \left(\frac{1}{2\xi_2} - \frac{\alpha}{4} \right) q_1 - \frac{1}{\xi_{T+1}} q_{T+1} \\ &\quad + \sum_{t=1}^T \frac{\xi_{t+1}}{2} G^2 + \sum_{t=1}^T \frac{h^2}{\alpha} \|\delta_t(\min(W, T))\|^2 \\ &\leq \log(T+1) \frac{2G^2}{\alpha} + \sum_{t=1}^T \frac{h^2}{\alpha} \|\delta_t(\min(W, T))\|^2 \end{aligned}$$

which completes the proof of the static regret bound.

Next, we bound the switching costs. By Assumption 3(ii), we have

$$\begin{aligned} \sum_{t=1}^T d(x_t, x_{t-1}) &\leq \sum_{t=1}^T \frac{\beta}{2} \|x_t - x_{t-1}\|^2 \\ &\leq \sum_{t=1}^T \frac{\beta}{2} \|\xi_t \nabla_x f(x_{t-1}; \theta_{t-1} - \delta_{t-1}(\min(W, T)))\|^2 \\ &\leq \frac{\beta G^2}{2} \sum_{t=1}^T \xi_t^2 \leq \frac{16\beta G^2}{\alpha^2} \end{aligned}$$

□

D.2 Proof of Theorem 2: dynamic regret bound for restarted OGD with inexact gradients

We denote the set of stages in epoch k as $\mathcal{T}_k = \{k\Delta + 1, \dots, \min(k\Delta + \Delta, T)\}$ for $k = 0, \dots, \lceil T/\Delta \rceil - 1$. We introduce $z_k^* = \arg \min_{z \in \mathbb{X}} \sum_{t \in \mathcal{T}_k} [f(z; \theta_t)]$ for all k ; $y_t^* = \arg \min_{x_t \in \mathbb{X}} f(x_t; \theta_t)$ for all t ; and $x^* = \arg \min_{x \in \mathbb{X}^T} \sum_{t=1}^T [f(x_t; \theta_t) + d(x_t, x_{t-1})]$. The dynamic regret of the restarted OGD with inexact gradients can be bounded as follows.

$$\text{Reg}(\text{OGD}) = \sum_{t=1}^T [f(x_t; \theta_t) + d(x_t, x_{t-1})] - \sum_{t=1}^T [f(x_t^*; \theta_t) + d(x_t^*, x_{t-1}^*)]$$

$$\begin{aligned}
&\leq \sum_{t=1}^T [f(x_t; \theta_t) + d(x_t, x_{t-1})] - \sum_{t=1}^T [f(x_t^*; \theta_t)] \\
&\leq \sum_{t=1}^T [f(x_t; \theta_t) + d(x_t, x_{t-1})] - \sum_{t=1}^T [f(y_t^*; \theta_t)] \\
&= \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(x_t; \theta_t) + d(x_t, x_{t-1}) - f(y_t^*; \theta_t)] \\
&= \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(x_t; \theta_t) - f(z_k^*; \theta_t)] + \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} d(x_t, x_{t-1}) \\
&\quad + \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(z_k^*; \theta_t) - f(y_t^*; \theta_t)] \\
&\leq \lceil T/\Delta \rceil \log(\Delta + 1) \frac{2G^2}{\alpha} + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2 + \lceil T/\Delta \rceil \frac{16\beta G^2}{\alpha^2} \\
&\quad + \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(z_k^*; \theta_t) - f(y_t^*; \theta_t)]
\end{aligned}$$

where the first inequality uses Assumption 3, the second inequality uses the optimality of y_t^* , the last inequality uses Theorem 6 and the fact that the OGD considered here restarts at the beginning of each epoch k and repeats the stepsizes defined in Theorem 6, thus satisfying the static regret bound and switching cost bound in Theorem 6 within each epoch.

Now, it suffices to bound $\sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(z_k^*; \theta_t) - f(y_t^*; \theta_t)]$. By the optimality of z_k^* , we have that

$$\sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(z_k^*; \theta_t) - f(y_t^*; \theta_t)] \leq \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(y_{k\Delta+1}^*; \theta_t) - f(y_t^*; \theta_t)].$$

We define $V^k = \sum_{t \in \mathcal{T}_k} \sup_{x \in \mathbb{X}} |f(x; \theta_t) - f(x; \theta_{t-1})|$. Then, for any $t \in \mathcal{T}_k$, we obtain

$$\begin{aligned}
f(y_{k\Delta+1}^*; \theta_t) - f(y_t^*; \theta_t) &= f(y_{k\Delta+1}^*; \theta_t) - f(y_{k\Delta+1}^*; \theta_{k\Delta+1}) + f(y_{k\Delta+1}^*; \theta_{k\Delta+1}) - f(y_t^*; \theta_{k\Delta+1}) \\
&\quad + f(y_t^*; \theta_{k\Delta+1}) - f(y_t^*; \theta_t) \\
&\leq 2V^k.
\end{aligned}$$

By summing over $t \in \mathcal{T}_k$ and $k = 0, \dots, \lceil T/\Delta \rceil - 1$, we obtain

$$\sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} [f(z_k^*; \theta_t) - f(y_t^*; \theta_t)] \leq 2\Delta V_T$$

Combining the bounds above yields the desired bound on the dynamic regret of OGD below by letting $\Delta = \lceil \sqrt{2T/V_T} \rceil$:

$$\begin{aligned}
\text{Reg}(OGD) &\leq \lceil T/\Delta \rceil \log(\Delta + 1) \frac{2G^2}{\alpha} + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2 + \lceil T/\Delta \rceil \frac{16\beta G^2}{\alpha^2} + 2\Delta V_T \\
&\leq \left(\sqrt{\frac{V_T T}{2}} + 1 \right) \log(2 + \sqrt{2T/V_T}) \frac{2G^2}{\alpha} + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2 + \left(\sqrt{\frac{V_T T}{2}} + 1 \right) \frac{16\beta G^2}{\alpha^2} \\
&\quad + 2(\sqrt{2V_T T} + V_T) \\
&\leq (\sqrt{V_T T/2} + 1) \log(2 + \sqrt{2T/V_T}) \left(\frac{2G^2}{\alpha} + \frac{16\beta G^2}{\alpha^2} + 2(2 + \sqrt{2}) \right) + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2 \\
&\leq \sqrt{V_T T} \log(2 + \sqrt{2T/V_T}) \left(\frac{2G^2}{\alpha} + \frac{16\beta G^2}{\alpha^2} + 2(2 + \sqrt{2}) \right) + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2 \\
&\leq \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \left(\frac{4G^2}{\alpha} + \frac{32\beta G^2}{\alpha^2} + 4(2 + \sqrt{2}) \right) + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2
\end{aligned}$$

$$\leq \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \left(\frac{4G^2}{\alpha} + \frac{32\beta G^2}{\alpha^2} + 16 \right) + \frac{h^2}{\alpha} \|\delta(\min(W, T))\|^2$$

where we used the facts that $1 \leq V_T \leq T$, $T > 2$, $\log(2 + \sqrt{2T/V_T}) \leq 2 \log(1 + \sqrt{T/V_T})$ and $4(2 + \sqrt{2}) < 16$.

D.3 Proof of Corollary 1

The proof is straightforward by substituting restarted OGD's regret bound in Theorem 2 into the general regret bound in Theorem 1, that is,

$$\begin{aligned} \text{Reg}(RHIG) &\leq \rho^W \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \\ &\quad + \frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^W \|\delta(\min(W, T))\|^2 + \sum_{k=1}^{\min(W, T)} \zeta \rho^{k-1} \|\delta(k)\|^2 + \mathbb{1}_{(W > T)} \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2. \end{aligned}$$

D.4 Proofs of the monotonicity claims in the discussion of Corollary 1.

In Section 4, when discussing **Choices of W** , we claim that ‘‘Part I increases with V_T and Part II increases with the prediction errors. Further, as W increases, Part I decreases but Part II increases.’’ For completeness, we prove this claim below.

Properties of Part I $\rho^W \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T})$: Since $0 < \rho < 1$, it is straightforward that Part I monotonically decreases with W . Next, consider function $g(x) = x \log(1 + \frac{b}{x})$ for $x, b > 0$. Since $g'(x) = \frac{x}{x+b} - 1 - \log(\frac{x}{x+b}) \geq 0$ by $y - 1 \geq \log(y)$ for any $y > 0$, function $g(x)$ monotonically increases with x . Therefore, for any fixed W , Part I monotonically increases with $\sqrt{V_T}$ and thus V_T .

Properties of Part II $\frac{2Lh^2}{\alpha^2} \rho^W \|\delta(\min(W, T))\|^2 + \sum_{k=1}^{\min(W, T)} \zeta \rho^{k-1} \|\delta(k)\|^2 + \mathbb{1}_{(W > T)} \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2$: It is straightforward that Part II monotonically increases with $\{\|\delta(k)\|^2\}_{k=1}^W$. Next, we discuss the monotonicity with respect to W . We first consider $W \leq T$. In this case, Part II is equal to Part II(W) := $\frac{2Lh^2}{\alpha^2} \rho^W \|\delta(W)\|^2 + \sum_{k=1}^W \zeta \rho^{k-1} \|\delta(k)\|^2$. Notice that

$$\begin{aligned} \text{Part II}(W) - \text{Part II}(W-1) &= \frac{2Lh^2}{\alpha^2} \rho^W \|\delta(W)\|^2 + \zeta \rho^{W-1} \|\delta(W)\|^2 - \frac{2Lh^2}{\alpha^2} \rho^{W-1} \|\delta(W-1)\|^2 \\ &\geq \left(\frac{2Lh^2}{\alpha^2} \rho + \zeta - \frac{2Lh^2}{\alpha^2} \right) (\rho^{W-1} \|\delta(W-1)\|^2) \\ &= \left(\frac{h^2}{2\alpha} + \frac{h^2}{2L} \right) (\rho^{W-1} \|\delta(W-1)\|^2) > 0 \end{aligned}$$

where we used $\|\delta(W)\|^2 \geq \|\delta(W-1)\|^2$, $\rho = 1 - \frac{\alpha}{4L}$, $\zeta = \frac{h^2}{\alpha} + \frac{h^2}{2L}$. Therefore, Part II is monotonically increasing with W for $W \leq T$. Besides, we consider $W > T$. In this case, Part II is equal to Part II(W) := $\frac{2Lh^2}{\alpha^2} \rho^W \|\delta(T)\|^2 + \sum_{k=1}^T \zeta \rho^{k-1} \|\delta(k)\|^2 + \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2$. Notice that, when $W > T$, we have

$$\begin{aligned} \text{Part II}(W) - \text{Part II}(W-1) &= \frac{2Lh^2}{\alpha^2} (\rho^W - \rho^{W-1}) \|\delta(T)\|^2 + \frac{\rho^{W-1} - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2 \\ &= \left(\frac{2Lh^2}{\alpha^2} (\rho - 1) + \zeta \right) (\rho^{W-1} \|\delta(T)\|^2) \\ &= \left(\frac{h^2}{2\alpha} + \frac{h^2}{2L} \right) (\rho^{W-1} \|\delta(W-1)\|^2) > 0 \end{aligned}$$

In conclusion, Part II increases with W for $W \geq 1$.

E Analysis on the special case in Section 4

E.1 Proof of Corollary 2

For notational simplicity, let $R(W)$ denote the regret bound in Corollary 1 given lookahead horizon W , i.e.

$$R(W) = \rho^W \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) + \frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^W \|\delta(\min(W, T))\|^2 + \sum_{k=1}^{\min(W, T)} \zeta \rho^{k-1} \|\delta(k)\|^2 + \mathbb{1}_{(W > T)} \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2.$$

We will show that $R(W) \leq R(W - 1)$ for $W \geq 1$. Firstly, we consider $W \leq T$. In this case, we have $R(W) = \rho^W \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) + \frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^W \|\delta(W)\|^2 + \sum_{k=1}^W \zeta \rho^{k-1} \|\delta(k)\|^2$. Notice that

$$\begin{aligned} R(W) - R(W - 1) &= (\rho^W - \rho^{W-1}) \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) + \frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^W \|\delta(W)\|^2 \\ &\quad + \zeta \rho^{W-1} \|\delta(W)\|^2 - \frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^{W-1} \|\delta(W - 1)\|^2 \\ &\leq \rho^{W-1} \left(\left(\frac{2L}{\alpha} \frac{h^2}{\alpha} \rho + \zeta \right) \|\delta(W)\|^2 - (1 - \rho) \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \right) \\ &\leq 0 \end{aligned}$$

when the following condition holds for any $W \leq T$.

$$\left(\frac{2L}{\alpha} \frac{h^2}{\alpha} \rho + \zeta \right) \|\delta(W)\|^2 \leq (1 - \rho) \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \quad (17)$$

Next, we consider $W > T$. In this case, we have $R(W) = \rho^W \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) + \frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^W \|\delta(T)\|^2 + \sum_{k=1}^T \zeta \rho^{k-1} \|\delta(k)\|^2 + \frac{\rho^T - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2$. Therefore,

$$\begin{aligned} R(W) - R(W - 1) &= (\rho^W - \rho^{W-1}) \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) + \frac{2L}{\alpha} \frac{h^2}{\alpha} (\rho^W - \rho^{W-1}) \|\delta(T)\|^2 \\ &\quad + \frac{\rho^{W-1} - \rho^W}{1 - \rho} \zeta \|\delta(T)\|^2 \\ &\leq \rho^{W-1} \left(\zeta \|\delta(T)\|^2 - (1 - \rho) \frac{2L}{\alpha} C_1 \sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \right) \\ &\leq 0 \end{aligned}$$

given the condition (17).

In conclusion, we have $R(W) \leq R(W - 1)$ for $W \geq 1$ and the $R(W)$ is minimized by letting $W \rightarrow +\infty$. Further, when $W \rightarrow +\infty$, we have the following bound.

$$\lim_{W \rightarrow +\infty} R(W) = \sum_{k=1}^T \zeta \rho^{k-1} \|\delta(k)\|^2 + \frac{\rho^T}{1 - \rho} \zeta \|\delta(T)\|^2 \leq \frac{\zeta}{1 - \rho} \sum_{k=1}^T \rho^{k-1} \|\delta(k)\|^2$$

E.2 Proof of Theorem 3

Without loss of generality, we consider $n = 1$. It is straightforward to generalize the proof to $n > 1$ cases. The proof is based on constructing a special cost function where the lower bound holds.

Consider cost function $f(x_t; \theta_t) = \frac{\alpha}{2} (x_t^2 - 2\theta_t x_t)$ and $d(x_t, x_{t-1}) = \frac{\beta}{2} \|x_t - x_{t-1}\|^2$ on $\mathbb{X} = [-1/2, 1/2]$, where $l_f = \alpha$, $l_d = 2\beta$, $L = \alpha + 4\beta$ and $h = \alpha$. Let $\alpha > 1$ and $\frac{\beta}{\alpha} < 4 + 3\sqrt{2}$ so that $\rho_0 < 1/2$. Let $\theta_t \in \mathbb{X}$ for all t , then we have $G = \sup_{x \in \mathbb{X}} \|\alpha(x - \theta_t)\| = \alpha$. Let $x_0 = 0$.

Consider a random θ_t :

$$\theta_t = \mu_t + e_1^t + \dots + e_t^t, \quad \forall 1 \leq t \leq T,$$

where e_τ^t are independent variables across $1 \leq t \leq T$ and $1 \leq \tau \leq t$. Let the support of e_τ^t be $[-\frac{1}{8t}, \frac{1}{8t}]$ and let $\mu_t = (-1)^t \frac{1}{4}$, so $\theta_t \in \mathbb{X}$ is \mathbb{X} and $\frac{1}{8} \leq \theta_t \leq \frac{3}{8}$ if t is even and $-\frac{3}{8} \leq \theta_t \leq -\frac{1}{8}$ if t is odd. Consider predictions at time τ as $\theta_{t|\tau} = \mu_t + e_1^t + \dots + e_\tau^t$ for any $0 \leq \tau \leq t$. Therefore, $\delta_t(t - \tau) = e_{\tau+1}^t + \dots + e_t^t$ and $\|\delta_t(t - \tau)\| \leq \frac{t-\tau}{8t} \leq 1/8$. According to our construction, we have that $V_T = \sum_{t=1}^T \sup_{x \in \mathbb{X}} \alpha \|(\theta_t - \theta_{t-1})x\| \geq \alpha T/8$; and $\|\delta(k)\|^2 \leq \frac{T}{64}$ for any $k \geq 1$. Then, it is straightforward to verify that the constructed cost functions and predictions satisfy $\sqrt{V_T T} \log(1 + \sqrt{T/V_T}) \geq \frac{2Lh^2\rho + \alpha^2\zeta}{2LC_1(1-\rho)\alpha} \|\delta(k)\|^2$ for any $k \geq 1$.

Notice that knowing $\theta_{t|0}, \dots, \theta_{t|\tau}$ is equivalent with knowing $\mu_t, e_1^t, \dots, e_\tau^t$. Therefore, let filtration \mathcal{F}_t denote all the information at t provide by the predictions and the history, then \mathcal{F}_t is generated by $\theta_1, \dots, \theta_{t-1}$ and $\mu_s, e_1^s, \dots, e_{t-1}^s$ for $s \geq t$. Notice that $\mathbb{E}[\theta_\tau | \mathcal{F}_t] = \theta_{\tau|t-1}$ for $\tau \geq t$. Besides, for any online algorithm \mathcal{A} , we have that $x_t^{\mathcal{A}}$ is measurable in \mathcal{F}_t .

Since $\theta_t \in \mathbb{X}$ for all t , it can be shown that the optimal solution $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{X}^T} C(\mathbf{x}; \boldsymbol{\theta})$ is an interior point of \mathbb{X}^T and thus satisfies the first-order optimality condition $\mathbf{x}^* = A\boldsymbol{\theta}$, where A is the inverse of the Hessian matrix of $C(\mathbf{x}; \boldsymbol{\theta})$. Equivalently, we have $x_t^* = \sum_{\tau=1}^T a_{t,\tau} \theta_\tau$. Further, Lemma 5 in [17] shows that $a_{t,\tau}^2 \geq c_2 \rho_0^{\tau-t}$ for $\tau \geq t$, where $c_2 = (\frac{\alpha}{\alpha+\beta})^2 (1 - \sqrt{\rho_0})^2$.

Since $x_t^{\mathcal{A}}$ is measurable in \mathcal{F}_t , by the projection theory, we have

$$\mathbb{E}[\|x_t^{\mathcal{A}} - x_t^*\|^2] \geq \mathbb{E}[\|\mathbb{E}[x_t^* | \mathcal{F}_t] - x_t^*\|^2].$$

Notice that

$$\begin{aligned} \mathbb{E}[x_t^* | \mathcal{F}_t] &= a_{t,1}\theta_1 + \dots + a_{t,t-1}\theta_{t-1} + a_{t,t}\mathbb{E}[\theta_t | \mathcal{F}_t] + \dots + a_{t,T}\mathbb{E}[\theta_T | \mathcal{F}_t] \\ &= a_{t,1}\theta_1 + \dots + a_{t,t-1}\theta_{t-1} + a_{t,t}\theta_{t|t-1} + a_{t,T}\theta_{T|t-1} \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}[\|\mathbb{E}[x_t^* | \mathcal{F}_t] - x_t^*\|^2] &= \mathbb{E}[\|a_{t,t}\delta_t(1) + \dots + a_{t,T}\delta_T(T-t+1)\|^2] \\ &= a_{t,t}^2 \mathbb{E}[\|\delta_t(1)\|^2] + \dots + a_{t,T}^2 \mathbb{E}[\|\delta_T(T-t+1)\|^2] \\ &\geq c_2 (\mathbb{E}[\|\delta_t(1)\|^2] + \dots + \rho_0^{T-t} \mathbb{E}[\|\delta_T(T-t+1)\|^2]) \end{aligned}$$

where we used the independence among the prediction errors and $a_{t,\tau}^2 \geq c_2 \rho_0^{\tau-t}$ for $\tau \geq t$.

Summing over t leads to the following.

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[\|x_t^{\mathcal{A}} - x_t^*\|^2] &\geq \sum_{t=1}^T \mathbb{E}[\|\mathbb{E}[x_t^* | \mathcal{F}_t] - x_t^*\|^2] \\ &\geq \sum_{t=1}^T (a_{t,t}^2 \mathbb{E}[\|\delta_t(1)\|^2] + \dots + a_{t,T}^2 \mathbb{E}[\|\delta_T(T-t+1)\|^2]) \\ &\geq \sum_{t=1}^T c_2 \sum_{k=1}^{T-t+1} \rho_0^{k-1} \mathbb{E}[\|\delta_{k+t-1}(k)\|^2] \\ &= c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^{T+1-k} \mathbb{E}[\|\delta_{k+t-1}(k)\|^2] \\ &= c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^T \mathbb{E}[\|\delta_t(k)\|^2] - c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^{k-1} \mathbb{E}[\|\delta_t(k)\|^2] \\ &= c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^T \mathbb{E}[\|\delta_t(k)\|^2] - c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^{k-1} \mathbb{E}[\|\delta_t(t)\|^2] \\ &= c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^T \mathbb{E}[\|\delta_t(k)\|^2] - c_2 \sum_{t=1}^{T-1} \mathbb{E}[\|\delta_t(t)\|^2] \sum_{k=t+1}^T \rho_0^{k-1} \\ &\geq c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^T \mathbb{E}[\|\delta_t(k)\|^2] - c_2 \sum_{k=1}^{T-1} \mathbb{E}[\|\delta_k(k)\|^2] \frac{\rho_0^k}{1-\rho_0} \end{aligned}$$

$$\geq c_2 \sum_{k=1}^T \rho_0^{k-1} \sum_{t=1}^T \mathbb{E}[\|\delta_t(k)\|^2] \frac{1-2\rho_0}{1-\rho_0}$$

where we used $\delta_t(k) = \delta_t(t)$ for $k \geq t$ in the third equality and change the counting index from t to k in the second last inequality.

By strong convexity, we have $\mathbb{E}[\text{Reg}(\mathcal{A})] \geq \frac{\alpha}{2} \mathbb{E} \|\mathbf{x}^{\mathcal{A}} - \mathbf{x}^*\|^2 \geq c_2 \frac{\alpha}{2} \frac{1-2\rho_0}{1-\rho_0} \sum_{k=1}^T \rho_0^{k-1} \mathbb{E}[\|\boldsymbol{\delta}(k)\|^2]$. Therefore, there must exist a scenario such that $\text{Reg}(\mathcal{A}) \geq c_2 \frac{\alpha}{2} \frac{1-2\rho_0}{1-\rho_0} \sum_{k=1}^T \rho_0^{k-1} \|\boldsymbol{\delta}(k)\|^2$. Since $h = \alpha$ in our construction, we complete the proof.

F Stochastic Regret Analysis

F.1 Proof of Theorem 4

By taking expectation on both sides of the regret bound in Theorem 1, we have

$$\mathbb{E}[\text{Reg}(RHIG)] \leq \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \zeta \sum_{k=1}^{\min(W,T)} \rho^{k-1} \mathbb{E}[\|\boldsymbol{\delta}(k)\|^2] + \mathbb{1}_{(W>T)} \frac{\rho^T - \rho^W}{1-\rho} \zeta \mathbb{E}[\|\boldsymbol{\delta}(T)\|^2], \quad (18)$$

Therefore, it suffices to bound $\mathbb{E}[\|\boldsymbol{\delta}(k)\|^2]$ for $1 \leq k \leq T$. By $\boldsymbol{\delta}(k) = (\delta_1(k), \dots, \delta_T(k))$, $\delta_t(k) = \theta_t - \theta_{t|t-k} = P(0)e_t + \dots + P(k-1)e_{t-k+1}$ for $k \leq t$ and $\delta_t(k) = \tilde{\delta}_t(t)$ for $k > t$, we have

$$\boldsymbol{\delta}(k) = \mathbf{M}_k \mathbf{e}, \quad 1 \leq k \leq T \quad (19)$$

where we define $\mathbf{e} = (e_1^\top, \dots, e_T^\top)^\top \in \mathbb{R}^{qT}$ and

$$\mathbf{M}_k = \begin{bmatrix} P(0) & 0 & \dots & \dots & \dots & 0 \\ P(1) & P(0) & \dots & \dots & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ P(k-1) & \dots & P(1) & P(0) & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & P(k-1) & \dots & P(1) & P(0) \end{bmatrix}.$$

Let \mathbf{R}_e denote the covariance matrix of \mathbf{e} , i.e.

$$\mathbf{R}_e = \begin{bmatrix} R_e & 0 & \dots & 0 \\ 0 & R_e & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & R_e \end{bmatrix}$$

Then, for $k \leq T$, we have

$$\begin{aligned} \mathbb{E}[\|\boldsymbol{\delta}(k)\|^2] &= \mathbb{E}[\mathbf{e}^\top \mathbf{M}_k^\top \mathbf{M}_k \mathbf{e}] = \mathbb{E}[\text{tr}(\mathbf{e} \mathbf{e}^\top \mathbf{M}_k^\top \mathbf{M}_k)] \\ &= \text{tr}(\mathbf{R}_e \mathbf{M}_k^\top \mathbf{M}_k) \\ &\leq \|\mathbf{R}_e\|_2 \|\mathbf{M}_k\|_F^2 = \|\mathbf{R}_e\|_2 \sum_{t=0}^{k-1} (T-t) \|P(t)\|_F^2 \end{aligned}$$

where the first inequality is by $\text{tr}(AB) \leq \|A\|_2 \text{tr}(B)$ for any symmetric matrices A, B , and $\|\text{diag}(R_e, \dots, R_e)\|_2 = \|\mathbf{R}_e\|_2$ and $\text{tr}(A^\top A) = \|A\|_F^2$ for any matrix A . In addition, for $k \geq T$, we have $\mathbb{E}[\|\boldsymbol{\delta}(k)\|^2] \leq \|\mathbf{R}_e\|_2 \sum_{t=0}^{T-1} (T-t) \|P(t)\|_F^2$. In conclusion, for any $k \geq 1$, we have

$$\mathbb{E}[\|\boldsymbol{\delta}(k)\|^2] \leq \|\mathbf{R}_e\|_2 \sum_{t=0}^{\min(k,T)-1} (T-t) \|P(t)\|_F^2 \quad (20)$$

When $W \leq T$, substituting the bounds on $\mathbb{E}[\|\delta(k)\|^2]$ into (18) yields the bound on the expected regret below.

$$\begin{aligned}\mathbb{E}[\text{Reg}(RHIG)] &\leq \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \zeta \sum_{k=1}^W \rho^{k-1} \|R_e\|_2 \sum_{t=0}^{k-1} (T-t) \|P(t)\|_F^2 \\ &= \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \zeta \sum_{t=0}^{W-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \sum_{k=t+1}^W \rho^{k-1} \\ &= \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \zeta \sum_{t=0}^{W-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t - \rho^W}{1 - \rho}\end{aligned}$$

When $W \geq T$, substituting the bounds on $\mathbb{E}[\|\delta(k)\|^2]$ into (18) yields the bound on the expected regret below.

$$\begin{aligned}\mathbb{E}[\text{Reg}(RHIG)] &\leq \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \zeta \sum_{k=1}^T \rho^{k-1} \|R_e\|_2 \sum_{t=0}^{k-1} (T-t) \|P(t)\|_F^2 \\ &\quad + \zeta \frac{\rho^T - \rho^W}{1 - \rho} \|R_e\|_2 \sum_{t=0}^{T-1} (T-t) \|P(t)\|_F^2 \\ &= \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \zeta \sum_{t=0}^{T-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \left(\sum_{k=t+1}^T \rho^{k-1} + \frac{\rho^T - \rho^W}{1 - \rho} \right) \\ &= \frac{2L}{\alpha} \rho^W \mathbb{E}[\text{Reg}(\phi)] + \zeta \sum_{t=0}^{T-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t - \rho^W}{1 - \rho}\end{aligned}$$

In conclusion, we have the regret bound for general $W \geq 0$ below.

$$\mathbb{E}[\text{Reg}(RHIG)] \leq \frac{2L}{\alpha} \rho^W \text{Reg}(\phi) + \zeta \sum_{t=0}^{\min(W,T)-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t - \rho^W}{1 - \rho}$$

F.2 Proof of Corollary 3

Before the proof, we note that we cannot apply the expected regret bound in [34] directly due to the major differences in the problem formulation as discussed below. Firstly, the expected regret definition considered in this paper is different from that in [34] because the true cost function parameter θ_t in our case is also random and taken expectation on, while the true cost function in [34] is deterministic and the expectation is only taken on the random gradient noises. Besides, [34] considers unbiased gradient estimation while our gradient estimation $\nabla_{x_t} f(x_t; \theta_{t\tau})$ can be biased. Further, [34] considers independent gradient noises at each stage t , while our gradient noises are correlated due to the correlation among prediction errors. Therefore, we have to revise the original proof in [34] for a new regret bound for our setting.

Similar to the proof of Theorem 2, we denote the set of stages in epoch k as $\mathcal{T}_k = \{k\Delta + 1, \dots, \min(k\Delta + \Delta, T)\}$ for $k = 0, \dots, \lceil T/\Delta \rceil - 1$; and introduce $z_k^* = \arg \min_{z \in \mathbb{X}} \sum_{t \in \mathcal{T}_k} [f(z; \theta_t)]$, $y_t^* = \arg \min_{x \in \mathbb{X}} f(x; \theta_t)$, $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{X}^T} \sum_{t=1}^T [f(x_t; \theta_t) + d(x_t, x_{t-1})]$. Notice that z_k^*, y_t^*, x_t^* are all random variables depending on θ . The expected dynamic regret of OGD can be bounded as follows.

$$\begin{aligned}\mathbb{E}[\text{Reg}(OGD)] &= \sum_{t=1}^T \mathbb{E}[f(x_t; \theta_t) + d(x_t, x_{t-1})] - \sum_{t=1}^T \mathbb{E}[f(x_t^*; \theta_t) + d(x_t^*, x_{t-1}^*)] \\ &\leq \sum_{t=1}^T \mathbb{E}[f(x_t; \theta_t) + d(x_t, x_{t-1})] - \sum_{t=1}^T \mathbb{E}[f(x_t^*; \theta_t)] \\ &\leq \sum_{t=1}^T \mathbb{E}[f(x_t; \theta_t) + d(x_t, x_{t-1})] - \sum_{t=1}^T \mathbb{E}[f(y_t^*; \theta_t)]\end{aligned}$$

$$\begin{aligned}
&= \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(x_t; \theta_t) + d(x_t, x_{t-1}) - f(y_t^*; \theta_t)] \\
&= \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(x_t; \theta_t) - f(z_k^*; \theta_t)] + \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[d(x_t, x_{t-1})] \\
&\quad + \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(z_k^*; \theta_t) - f(y_t^*; \theta_t)] \\
&\leq \lceil T/\Delta \rceil \log(\Delta + 1) \frac{2G^2}{\alpha} + \frac{h^2}{\alpha} \mathbb{E} \|\delta(\min(W, T))\|^2 + \lceil T/\Delta \rceil \frac{16\beta G^2}{\alpha^2} \\
&\quad + \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(z_k^*; \theta_t) - f(y_t^*; \theta_t)]
\end{aligned}$$

where the first inequality uses Assumption 3, the second inequality uses the optimality of y_t^* , the last inequality follows from taking expectation on the regret bounds in Theorem 6 and the fact that the OGD considered here restarts at the beginning of each epoch k and repeats the stepsizes defined in Theorem 6, thus satisfying the static regret bound and switching cost bound in Theorem 6 within each epoch.

Now, it suffices to bound $\sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(z_k^*; \theta_t) - f(y_t^*; \theta_t)]$. By the optimality of z_k^* , we have that

$$\sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(z_k^*; \theta_t) - f(y_t^*; \theta_t)] \leq \sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(y_{k\Delta+1}^*; \theta_t) - f(y_t^*; \theta_t)].$$

We define $\mathbb{E}[V^k] = \sum_{t \in \mathcal{T}_k} \mathbb{E}[\sup_{x \in \mathbb{X}} |f(x; \theta_t) - f(x; \theta_{t-1})|]$.¹⁰ Then, for any $t \in \mathcal{T}_k$, we obtain

$$\begin{aligned}
&\mathbb{E}[f(y_{k\Delta+1}^*; \theta_t) - f(y_t^*; \theta_t)] \\
&= \mathbb{E}[f(y_{k\Delta+1}^*; \theta_t) - f(y_{k\Delta+1}^*; \theta_{k\Delta+1})] + \mathbb{E}[f(y_{k\Delta+1}^*; \theta_{k\Delta+1}) - f(y_t^*; \theta_{k\Delta+1})] \\
&\quad + \mathbb{E}[f(y_t^*; \theta_{k\Delta+1}) - f(y_t^*; \theta_t)] \\
&\leq 2\mathbb{E}[V^k].
\end{aligned}$$

By summing over $t \in \mathcal{T}_k$ and $k = 0, \dots, \lceil T/\Delta \rceil - 1$, we obtain

$$\sum_{k=0}^{\lceil T/\Delta \rceil - 1} \sum_{t \in \mathcal{T}_k} \mathbb{E}[f(z_k^*; \theta_t) - f(y_t^*; \theta_t)] \leq 2\Delta \mathbb{E}[V_T]$$

Similar to the proof of Corollary 2, by applying the bounds above and $\Delta = \lceil \sqrt{2T/\mathbb{E}[V_T]} \rceil$, we obtain the desired bound on the expected dynamic regret of OGD for our setting, i.e.

$$\mathbb{E}[\text{Reg}(OGD)] \leq C_1 \sqrt{\mathbb{E}[V_T]T} \log(1 + \sqrt{T/\mathbb{E}[V_T]}) + \frac{h^2}{\alpha} \mathbb{E}[\|\delta(\min(W, T))\|^2]$$

where $C_1 = \frac{4G^2}{\alpha} + \frac{32\beta G^2}{\alpha^2} + 16$.

Consequently, by applying Theorem 4 and the bound on $\mathbb{E}[\|\delta(W)\|^2]$ in (20), we have the following results.

$$\mathbb{E}[\text{Reg}(RHIG)] \leq C_2 \rho^W \sqrt{\mathbb{E}[V_T]T} \log(1 + \sqrt{T/\mathbb{E}[V_T]}) + \rho^W \frac{2h^2 L}{\alpha^2} \sum_{t=0}^{\min(W, T) - 1} \|R_e\|_2 (T - t) \|P(t)\|_F^2$$

¹⁰Notice that y_t^* is a random variable depending on θ_t for all t . Therefore, in the inequalities below, $\mathbb{E}[f(y_{k\Delta+1}^*; \theta_t) - f(y_{k\Delta+1}^*; \theta_{k\Delta+1})]$ can be larger than the term $\sup_{x \in \mathbb{X}} \mathbb{E}[f(x; \theta_t) - f(x; \theta_{k\Delta+1})]$, where x is restricted to only deterministic variables. Thus, the expectation operator \mathbb{E} must be outside the sup operator in our definition of the expected variation of the environment (see $\mathbb{E}[V^k]$ and $\mathbb{E}[V_T]$).

$$\begin{aligned}
& + \zeta \sum_{t=0}^{\min(W,T)-1} \|R_e\|_2(T-t) \|P(t)\|_F^2 \frac{\rho^t - \rho^W}{1-\rho} \\
& \leq C_2 \rho^W \sqrt{\mathbb{E}[V_T] T} \log(1 + \sqrt{T/\mathbb{E}[V_T]}) + \zeta \sum_{t=0}^{\min(W,T)-1} \|R_e\|_2(T-t) \|P(t)\|_F^2 \frac{\rho^t}{1-\rho}
\end{aligned}$$

by $\frac{2h^2L}{\alpha^2} - \frac{\zeta}{1-\rho} < 0$, where $C_2 = \frac{2L}{\alpha} C_1$.

F.3 Proof of Theorem 5

The proof relies on the Hanson-Wright inequality in [39].¹¹

Proposition 1 (Hanson-Wright Inequality [39]). *Consider random Gaussian vector $\mathbf{u} = (u_1, \dots, u_n)^\top$ with u_i i.i.d. following $N(0, 1)$. There exists an absolute constant $c > 0$,¹² such that*

$$\mathbb{P}(\mathbf{u}^\top \mathbf{A} \mathbf{u} \geq \mathbb{E}[\mathbf{u}^\top \mathbf{A} \mathbf{u}] + b) \leq \exp\left(-c \min\left(\frac{b^2}{\|\mathbf{A}\|_F^2}, \frac{b}{\|\mathbf{A}\|_2}\right)\right), \quad \forall b > 0$$

Now, we are ready for the proof. Notice that by $\mathbb{E}[V_T] = T$ and $V_T \leq T$, we have $V_T = T$. Then, for any realization of the random vectors $\{e_t\}_{t=1}^T$, our regret bound in Section 4 still holds, i.e.

$$\begin{aligned}
\text{Reg}(RHIG) & \leq \rho^W \frac{2L}{\alpha} C_1 T \log(2) + \frac{2L}{\alpha} \frac{h^2}{\alpha} \rho^W \|\delta(\min(W, T))\|^2 + \sum_{k=1}^{\min(W,T)} \zeta \rho^{k-1} \|\delta(k)\|^2 \\
& \quad + \mathbb{1}_{(W>T)} \frac{\rho^T - \rho^W}{1-\rho} \zeta \|\delta(T)\|^2
\end{aligned}$$

where both $\text{Reg}(RHIG)$ and $\|\delta(k)\|^2$ are random variables here. Let $R(W)$ denote the regret bound. By the proof of Corollary 3, it is straightforward that $\mathbb{E}[R(W)]$ is no more than $\mathbb{E}[\text{Regbdd}]$ defined in Theorem 5.

From (19) in the proof of Theorem 4, we have that $\delta(k) = \mathbf{M}_k \mathbf{e} = \mathbf{M}_k \mathbf{R}_e^{1/2} \mathbf{u}$, where \mathbf{u} is a standard Gaussian vector for $k \leq T$; and $\delta(k)^\top = \mathbf{M}_T \mathbf{R}_e^{1/2} \mathbf{u}$ for $k \geq T$.

When $W \leq T$, we have the following formula for the regret bound $R(W)$.

$$\begin{aligned}
R(W) & = \rho^W \frac{2L}{\alpha} C_1 T \log(2) + \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|\delta(W)\|^2 + \zeta \sum_{k=1}^W \rho^{k-1} \|\delta(k)\|^2 \\
& = \rho^W \frac{2L}{\alpha} C_1 T \log(2) \\
& \quad + \underbrace{\mathbf{u}^\top \left(\rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \mathbf{R}_e^{1/2} \mathbf{M}_W^\top \mathbf{M}_W \mathbf{R}_e^{1/2} + \zeta \sum_{k=1}^W \rho^{k-1} \mathbf{R}_e^{1/2} \mathbf{M}_k^\top \mathbf{M}_k \mathbf{R}_e^{1/2} \right) \mathbf{u}}_{\mathbf{A}_W}
\end{aligned}$$

We bound $\|\mathbf{A}_W\|_F$ below.

$$\begin{aligned}
\|\mathbf{A}_W\|_F & \leq \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|\mathbf{R}_e^{1/2} \mathbf{M}_W^\top \mathbf{M}_W \mathbf{R}_e^{1/2}\|_F + \zeta \sum_{k=1}^W \rho^{k-1} \|\mathbf{R}_e^{1/2} \mathbf{M}_k^\top \mathbf{M}_k \mathbf{R}_e^{1/2}\|_F \\
& \leq \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|\mathbf{M}_W \mathbf{R}_e^{1/2}\|_F^2 + \zeta \sum_{k=1}^W \rho^{k-1} \|\mathbf{M}_k \mathbf{R}_e^{1/2}\|_F^2 \\
& = \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \text{tr}(\mathbf{R}_e^{1/2} \mathbf{M}_W^\top \mathbf{M}_W \mathbf{R}_e^{1/2}) + \zeta \sum_{k=1}^W \rho^{k-1} \text{tr}(\mathbf{R}_e^{1/2} \mathbf{M}_k^\top \mathbf{M}_k \mathbf{R}_e^{1/2})
\end{aligned}$$

¹¹Here we use the fact that $\|X_i\|_\varphi = 1$ where $\|\cdot\|_\varphi$ is the subGaussian norm defined in [39].

¹²An absolute constant refers to a quantity that does not change with anything.

$$\begin{aligned}
&= \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \text{tr}(\mathbf{R}_e \mathbf{M}_W^\top \mathbf{M}_W) + \zeta \sum_{k=1}^W \rho^{k-1} \text{tr}(\mathbf{R}_e \mathbf{M}_k^\top \mathbf{M}_k) \\
&\leq \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|R_e\|_2 \text{tr}(\mathbf{M}_W^\top \mathbf{M}_W) + \zeta \|R_e\|_2 \sum_{k=1}^W \rho^{k-1} \text{tr}(\mathbf{M}_k^\top \mathbf{M}_k) \\
&= \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|R_e\|_2 \|\mathbf{M}_W\|_F^2 + \zeta \|R_e\|_2 \sum_{k=1}^W \rho^{k-1} \|\mathbf{M}_k\|_F^2 \\
&= \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|R_e\|_2 \sum_{t=0}^{W-1} (T-t) \|P(t)\|_F^2 + \zeta \|R_e\|_2 \sum_{k=1}^W \rho^{k-1} \sum_{t=0}^{k-1} (T-t) \|P(t)\|_F^2 \\
&= \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|R_e\|_2 \sum_{t=0}^{W-1} (T-t) \|P(t)\|_F^2 + \zeta \|R_e\|_2 \sum_{t=0}^{W-1} \sum_{k=t+1}^W \rho^{k-1} (T-t) \|P(t)\|_F^2 \\
&= \rho^W \frac{2L}{\alpha} \frac{h^2}{\alpha} \|R_e\|_2 \sum_{t=0}^{W-1} (T-t) \|P(t)\|_F^2 + \zeta \|R_e\|_2 \sum_{t=0}^{W-1} \frac{\rho^t - \rho^W}{1-\rho} (T-t) \|P(t)\|_F^2 \\
&\leq \zeta \sum_{t=0}^{W-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t}{1-\rho}
\end{aligned}$$

where we used (20) and $\zeta = \frac{h^2}{\alpha} + \frac{h^2}{2L}$ and $\rho = 1 - \frac{\alpha}{4L}$.

When $W > T$, we have the following formula for the regret bound $R(W)$.

$$\begin{aligned}
R(W) &= \rho^W \frac{2L}{\alpha} C_1 T \log(2) + \frac{2h^2 L}{\alpha^2} \rho^W \|\delta(T)\|^2 + \zeta \sum_{k=1}^T \rho^{k-1} \|\delta(k)\|^2 + \zeta \|\delta(T)\|^2 \frac{\rho^T - \rho^W}{1-\rho} \\
&= \rho^W \frac{2L}{\alpha} C_1 T \log(2) \\
&\quad + \underbrace{\mathbf{u}^\top \left(\left(\frac{2h^2 L}{\alpha^2} \rho^W + \zeta \frac{\rho^T - \rho^W}{1-\rho} \right) \mathbf{R}_e^{1/2} \mathbf{M}_T^\top \mathbf{M}_T \mathbf{R}_e^{1/2} + \zeta \sum_{k=1}^T \rho^{k-1} \mathbf{R}_e^{1/2} \mathbf{M}_k^\top \mathbf{M}_k \mathbf{R}_e^{1/2} \right) \mathbf{u}}_{\mathbf{A}_W}
\end{aligned}$$

Similarly, we bound $\|\mathbf{A}_W\|_F$ below.

$$\begin{aligned}
\|\mathbf{A}_W\|_F &\leq \left(\frac{2h^2 L}{\alpha^2} \rho^W + \zeta \frac{\rho^T - \rho^W}{1-\rho} \right) \|\mathbf{R}_e^{1/2} \mathbf{M}_T^\top \mathbf{M}_T \mathbf{R}_e^{1/2}\|_F + \zeta \sum_{k=1}^T \rho^{k-1} \|\mathbf{R}_e^{1/2} \mathbf{M}_k^\top \mathbf{M}_k \mathbf{R}_e^{1/2}\|_F \\
&\leq \left(\frac{2h^2 L}{\alpha^2} \rho^W + \zeta \frac{\rho^T - \rho^W}{1-\rho} \right) \|R_e\|_2 \|\mathbf{M}_T\|_F^2 + \zeta \sum_{k=1}^T \rho^{k-1} \|R_e\|_2 \|\mathbf{M}_k\|_F^2 \\
&\leq \left(\frac{2h^2 L}{\alpha^2} \rho^W + \zeta \frac{\rho^T - \rho^W}{1-\rho} \right) \|R_e\|_2 \sum_{t=0}^{T-1} (T-t) \|P(t)\|_F^2 + \zeta \sum_{k=1}^T \rho^{k-1} \|R_e\|_2 \sum_{t=0}^{k-1} (T-t) \|P(t)\|_F^2 \\
&\leq \left(\frac{2h^2 L}{\alpha^2} \rho^W + \zeta \frac{\rho^T - \rho^W}{1-\rho} \right) \|R_e\|_2 \sum_{t=0}^{T-1} (T-t) \|P(t)\|_F^2 + \zeta \|R_e\|_2 \sum_{t=0}^{T-1} \frac{\rho^t - \rho^T}{1-\rho} (T-t) \|P(t)\|_F^2 \\
&\leq \zeta \sum_{t=0}^{T-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t}{1-\rho}
\end{aligned}$$

In conclusion, for any $W \geq 1$, we have that $R(W) = \rho^W \frac{2L}{\alpha} C_1 T \log(2) + \mathbf{u}^\top \mathbf{A}_W \mathbf{u}$, and $\|\mathbf{A}_W\|_F \leq \zeta \sum_{t=0}^{\min(W,T)-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t}{1-\rho}$. Further, we have $\|\mathbf{A}_W\|_2 \leq \|\mathbf{A}_W\|_F$. Therefore, by Proposition 1, we prove the concentration bound below. For any $b > 0$,

$$\mathbb{P}(\text{Reg}(RHIG) \geq \mathbb{E}[\text{Regbdd}] + b) \leq \mathbb{P}(R(W) \geq \mathbb{E}[\text{Regbdd}] + b)$$

$$\begin{aligned}
&\leq \mathbb{P}(R(W) \geq \mathbb{E}[R(W)] + b) \\
&= \mathbb{P}(\mathbf{u}^\top \mathbf{A}_W \mathbf{u} \geq \mathbb{E}[\mathbf{u}^\top \mathbf{A}_W \mathbf{u}] + b) \\
&\leq \exp\left(-c \min\left(\frac{b^2}{K^2}, \frac{b}{K}\right)\right)
\end{aligned}$$

where $K = \zeta \sum_{t=0}^{\min(T,W)-1} \|R_e\|_2 (T-t) \|P(t)\|_F^2 \frac{\rho^t}{1-\rho}$.

G More details of the numerical experiments

(i) the high-level planning problem. The parameters are: $e_t \sim N(0, 1)$ i.i.d., $T = 20$, $\alpha = 1$, $\beta = 0.5$, $x_0 = 10$, $a = 4$, $\omega = 0.5$, $\eta = 0.5$, $\xi_t = 1$, CHC's commitment level $v = 3$. The regret is averaged over 200 iterations.

(ii) the physical tracking problem. Consider the second-order system

$$\ddot{x} = k_1 u + g + k_2$$

where x is altitude, \dot{x} is velocity, \ddot{x} is acceleration, etc.

The discrete-time version is

$$\frac{x_{t+1} - 2x_t + x_{t-1}}{\Delta^2} = k_1 u_t - g + k_2$$

which is equivalent to

$$u_t = \frac{1}{k_1} \left(\frac{x_{t+1} - 2x_t + x_{t-1}}{\Delta^2} - (-g + k_2) \right).$$

The cost function at stage t is

$$\frac{\alpha}{2} (x_t - \theta_t)^2 + \frac{\beta}{2} u_t^2.$$

We can write it in terms of x_t , that is,

$$\frac{\alpha}{2} (x_t - \theta_t)^2 + \frac{\beta}{2} \frac{1}{k_1^2} \left(\frac{x_{t+1} - 2x_t + x_{t-1}}{\Delta^2} - (-g + k_2) \right)^2.$$

Notice that the switching cost is not $d(x_t, x_{t-1})$ but $d(x_{t+1}, x_t, x_{t-1})$, but we still have the local coupling property of the gradients and we can still apply RHIG.

The experiment parameters are provided below. Consider horizon 10 seconds and time discretization $\Delta = 0.1$ s. Let $k_1 = 1$, $k_2 = 1$, $\alpha = 1$, $\beta = 1 \times 10^{-5}$, $x_0 = 1$ m, $g = 9.8$ m/s². Let $e_t \sim N(0, 0.5^2)$ i.i.d. for all t . Consider $d_t = 0.9 \sin(0.2t) + 1$ before $t \leq 5.6$ s and $d_t = 0.3 \sin(0.2t) + 1$ afterwards. Let $\gamma = 0.6$, $\xi_t = 1$, $\eta = 1/L$ and $L \approx 2.6$.