FROM SWITCHING TO DYNAMIC REGRET: A SIMPLE REDUCTION VIA UNBIASED RANDOM SEQUENCES

Anonymous authors
Paper under double-blind review

ABSTRACT

In non-stationary online learning, dynamic regret has attracted increasing attention as a measure of how well an online learner performs against a time-varying comparator sequence. Despite considerable advances, existing methods to attain optimal dynamic regret bounds for strongly convex losses and exp-concave losses remain analytically intricate or computationally demanding. In this paper, we present a *simple* framework that reduces dynamic regret minimization to switching regret minimization. The key idea is to construct, for any comparator sequence, an auxiliary random sequence that is unbiased at each round, with controlled variance and a manageable number of switches. Combining this construction with suitable surrogate losses, we decompose dynamic regret into the expected switching regret against the infrequently-changing random sequence and a variance-controlled approximation error. Consequently, *any* algorithm with a switching regret guarantee can be plugged into our framework to obtain a corresponding dynamic regret bound. In particular, by instantiating the framework with existing switching-regret algorithms, we can obtain an $\widetilde{\mathcal{O}}(T^{1/3}P_T^{2/3})$ dynamic regret bound for strongly convex losses and $\widetilde{\mathcal{O}}(d^{2/3}T^{1/3}P_T^{2/3})$ for exp-concave losses, where P_T denotes the path-length and d denotes the dimension. For general convex losses, the same reduction recovers the $\mathcal{O}(\sqrt{T(1+P_T)})$ dynamic regret bound. Notably, all our findings match the known minimax-optimal bounds, up to logarithmic factors, and the modularity of our framework makes it easy to understand and implement.

1 INTRODUCTION

Non-stationary online learning investigates sequential decision-making in evolving environments, with many real-world applications such as online recommendation and traffic scheduling (Hazan, 2016). A standard framework for studying such problems is Online Convex Optimization (OCO) (Shalev-Shwartz, 2012; Orabona, 2019), which can be viewed as a repeated game between a learner against the environment. Formally, at each round t, the learner chooses a decision $\mathbf{x}_t$ from a convex set $\mathcal{X}$, and then suffers a loss $f_t(\mathbf{x}_t)$ with the convex function $f_t(\cdot) : \mathcal{X} \to \mathbb{R}$ revealed by the environment. The goal of the learner is to minimize the dynamic regret (Zinkevich, 2003):

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) = \sum\nolimits_{t=1}^{T} f_t(\mathbf{x}_t) - \sum\nolimits_{t=1}^{T} f_t(\mathbf{u}_t), \tag{1}$$

which is defined as the difference between the cumulative loss incurred by the learner and that of *any* possible comparators $\mathbf{u}_1, \ldots, \mathbf{u}_T \in \mathcal{X}$. By restricting the comparator sequence, (1) recovers different metrics in OCO. For example, choosing the best fixed comparator, with $\mathbf{u}_t = \mathbf{u} \in \arg\min_{\mathbf{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x})$ for all t, yields the static regret (Cesa-Bianchi & Lugosi, 2006), and choosing the best piecewise-constant comparator sequence with a bounded number of switches delivers the switching regret (Pasteris et al., 2024; Yang et al., 2026).

Over the years, there have been a variety of investigations into dynamic regret minimization owing to its generality (Hall & Willett, 2013; Besbes et al., 2015; Jadbabaie et al., 2015; Mokhtari et al., 2016; Yang et al., 2016; Wei et al., 2016; Zhang et al., 2018b; Baby & Wang, 2019; Zhang et al., 2020; Cutkosky, 2020; Wan et al., 2021; Zhang et al., 2021; Wan et al., 2023; Wang et al., 2024). In particular, the seminal work of Zinkevich (2003) introduces the path length P_T, defined in (2), to capture the fluctuations of the comparator sequence, and establishes an $\mathcal{O}(\sqrt{T}(1+P_T))$

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: Comparisons of methods for dynamic regret minimization. Abbreviations: cvx → convex, str-cvx → strongly convex, exp-concave → exponentially concave. The last column indicates whether the proposed method is computationally efficient.

Method	Regret Bounds			Efficient
	cvx	str-cvx	exp-concave	
Zhang et al. (2018a)	$\mathcal{O}(\sqrt{T(1+P_T)})$	–	–	✓
Baby & Wang (2021)	–	$\widetilde{\mathcal{O}}(d^{7/3}T^{1/3}P_T^{2/3})$	$\widetilde{\mathcal{O}}(d^{23/6}T^{1/3}P_T^{2/3})$	✓
Baby & Wang (2022)	–	$\widetilde{\mathcal{O}}(d^{2/3}T^{1/3}P_T^{2/3})$	$\widetilde{\mathcal{O}}(d^{10/3}T^{1/3}P_T^{2/3})$	✓
Zhang et al. (2025)	–	–	$\widetilde{\mathcal{O}}(dT^{1/3}P_T^{2/3})$	✗
This work	$\mathcal{O}(\sqrt{T(1+P_T)})$	$\widetilde{\mathcal{O}}(T^{1/3}P_T^{2/3})$	$\widetilde{\mathcal{O}}(d^{2/3}T^{1/3}P_T^{2/3})$	✓

dynamic regret bound for general convex losses. Subsequently, Zhang et al. (2018a) improve this result to the provably optimal $\mathcal{O}(\sqrt{T(1+P_T)})$ bound. For strongly convex losses and exp-concave losses, Baby & Wang (2021) establish $\widetilde{\mathcal{O}}(d^{7/3}T^{1/3}P_T^{2/3})$ and $\widetilde{\mathcal{O}}(d^{23/6}T^{1/3}P_T^{2/3})$ bounds, respectively, through delicate analyses of the KKT conditions on offline optimal comparator sequences. Using similar analyses, Baby & Wang (2022) reduce the dimension dependence for strongly convex losses and exp-concave losses to $d^{2/3}$ and $d^{10/3}$, respectively. Recently, Zhang et al. (2025) further improve the dimension dependence for exp-concave losses, and obtain an $\widetilde{\mathcal{O}}(dT^{1/3}P_T^{2/3})$ dynamic regret bound. However, their method requires Kullback-Leibler (KL) projections onto constrained Gaussian mixtures, incurring substantial computational costs and complicating implementation.

In this paper, we develop a *simple* framework that reduces dynamic regret minimization to switching regret minimization, avoiding both intricate analyses of offline optimal comparator sequences and computationally intractable KL projections. Specifically, we first show that, for *any* comparator sequence, one can construct an auxiliary random sequence that follows it in expectation with bounded variance and a controlled expected number of switches. Combined with suitable surrogate losses, these properties allow us to decompose dynamic regret to the expected switching regret against the infrequently-changing auxiliary sequence and a variance-controlled approximation error. Consequently, existing algorithms with switching regret guarantees can be used as black boxes to obtain dynamic regret bounds. Importantly, the auxiliary random sequence is introduced solely for the analysis and need not be explicitly generated by the learner. In practice, the learner only runs the chosen switching-regret algorithm on the surrogate losses, without any access to or operations on the comparator sequence. Our main contributions are summarized as follows.

- We show that, for any comparator sequence, one can construct an unbiased auxiliary random sequence with bounded variance and few switches in expectation, enabling a reduction from dynamic to switching regret with a variance-controlled approximation error.
- Algorithmically, we develop a simple and modular framework that reuses existing switching-regret algorithms as black boxes, without modifying their internal updates.
- Theoretically, our framework yields dynamic regret bounds of $\widetilde{\mathcal{O}}(T^{1/3}P_T^{2/3})$ for strongly convex losses, $\widetilde{\mathcal{O}}(d^{2/3}T^{1/3}P_T^{2/3})$ for exp-concave losses, and $\mathcal{O}(\sqrt{T(1+P_T)})$ for general convex losses. In particular, we improve the dimension dependence for strongly convex losses and exp-concave losses. All our results attain the optimal dynamic regret bounds, up to logarithmic factors.

2 Related Work

In this section, we briefly review both dynamic and switching regret minimization in OCO.

Dynamic regret minimization. In the literature, there are two different forms of dynamic regret. One is the general case (1) proposed by Zinkevich (2003), who introduces the path length

$$P_T = \sum_{t=2}^{T} \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2, \tag{2}$$

to establish the first dynamic regret bound of $\mathcal{O}(\sqrt{T}(1+P_T))$ for Online Gradient Descent (OGD). Subsequently, Zhang et al. (2018a) improve this result to the minimax-optimal $\mathcal{O}(\sqrt{T(1+P_T)})$

bound by aggregating multiple OGD instances via the Hedge Algorithm (Freund & Schapire, 1997). To exploit the curvature of losses, Baby & Wang (2021) develop analyses based on the KKT conditions of offline optimal comparator sequence, and establish $\widetilde{\mathcal{O}}(d^{7/3}T^{1/3}P_T^{2/3})$ and $\widetilde{\mathcal{O}}(d^{23/6}T^{1/3}P_T^{2/3})$ bounds for strongly convex losses and exp-concave losses, respectively. Later, following similar analyses, Baby & Wang (2022) achieve improved bounds of $\widetilde{\mathcal{O}}(d^{2/3}T^{1/3}P_T^{2/3})$ for strongly convex losses and $\widetilde{\mathcal{O}}(d^{10/3}T^{1/3}P_T^{2/3})$ for exp-concave losses, with tighter dimension dependences. Recently, Zhang et al. (2025) develop a mixability-based analysis for dynamic regret minimization and extend it to exp-concave losses, obtaining an $\widetilde{\mathcal{O}}(dT^{1/3}P_T^{2/3})$ bound with linear dimension dependence. However, their method requires computationally expensive KL projections onto constrained Gaussian mixtures. We present comparisons of these methods in Table 1.

Another line of research focuses on the *worst-case* dynamic regret, in which the comparators in (1) are chosen as the per-round minimizers $\mathbf{u}_t^\star \in \arg\min_{\mathbf{u}\in\mathcal{X}} f_t(\mathbf{u})$. To quantify the difficulty of tracking these minimizers, existing studies establish regret bounds in terms of the temporal variation of the loss functions or the minimizer sequence (Besbes et al., 2015; Jadbabaie et al., 2015; Yang et al., 2016; Mokhtari et al., 2016; Zhang et al., 2017; Baby & Wang, 2019; Wan et al., 2021; 2023; Xu & Zhang, 2024). In this paper, we focus on the general form of dynamic regret (1), and aim to establish the dynamic regret bounds that adapt to the path-length P_T of *any* comparator sequence.

Switching regret minimization. Switching regret is defined as the difference between the cumulative loss incurred by the learner and that of the best piecewise-constant comparator sequence (Pasteris et al., 2024). Formally, let $\mathcal{S} = \{\mathcal{I}_1, \ldots, \mathcal{I}_{|\mathcal{S}|}\}$ be an arbitrary partition of the time horizon T into contiguous intervals. The switching regret with respect to $\mathcal{S}$ is defined as

$$\text{SW-Reg}_T(\mathcal{S}) = \sum\nolimits_{\mathcal{I}\in\mathcal{S}} \left[\sum\nolimits_{t\in\mathcal{I}} f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{X}} \sum\nolimits_{t\in\mathcal{I}} f_t(\mathbf{x})\right], \tag{3}$$

which is the sum of the static regret over each interval $\mathcal{I} \in \mathcal{S}$. The comparator remains fixed within each interval and switches at most $|\mathcal{S}| - 1$ times. The learner does not know the segmentation in advance, and its switching regret guarantee should hold simultaneously for all possible segmentations.

To minimize (3), Pasteris et al. (2024) propose RESET for general convex losses, which recursively aggregates OGD experts over a segment tree and achieves an $\mathcal{O}(\sum_{\mathcal{I}\in\mathcal{S}} \sqrt{|\mathcal{I}|})$ bound. Later, Yang et al. (2026) propose IRESET, which exploits loss curvature by combining a meta-algorithm with a second-order bound and suitable expert algorithms. With OGD as the expert algorithm for strongly convex losses and Online Newton Step (ONS) (Hazan et al., 2007) for exp-concave losses, IRESET attains switching-regret bounds of $\mathcal{O}(\sum_{\mathcal{I}\in\mathcal{S}} \log^2(|\mathcal{I}|))$ and $\mathcal{O}(d\sum_{\mathcal{I}\in\mathcal{S}} \log^2(|\mathcal{I}|))$, respectively.

3 Reduction via an Unbiased Random Sequence

In this section, we first introduce necessary preliminaries, including assumptions and definitions. Then, we present the reduction from dynamic regret minimization to switching regret minimization.

3.1 Preliminaries

Similar to previous studies in OCO, we list standard assumptions and basic definitions below.

Assumption 1. *The convex set $\mathcal{X}$ contains $\mathbf{0}$ and belongs to a ball with radius R, i.e., $\mathcal{X} \subseteq B(\mathbf{0}, R)$.*

Assumption 2. *At each round t, the loss function $f_t(\cdot)$ is G-Lipschitz over $\mathcal{X}$, i.e.,*

$$\forall \mathbf{x}, \mathbf{y} \in \mathcal{X},\ |f_t(\mathbf{x}) - f_t(\mathbf{y})| \le G\|\mathbf{x} - \mathbf{y}\|_2.$$

Definition 1. *A function $f : \mathcal{X} \to \mathbb{R}$ is convex if*

$$\forall \mathbf{x}, \mathbf{y} \in \mathcal{X},\ f(\mathbf{y}) \ge f(\mathbf{x}) + \nabla f(\mathbf{x})^\top(\mathbf{y} - \mathbf{x}).$$

Definition 2. *A function $f : \mathcal{X} \to \mathbb{R}$ is λ-strongly convex if*

$$\forall \mathbf{x}, \mathbf{y} \in \mathcal{X},\ f(\mathbf{y}) \ge f(\mathbf{x}) + \nabla f(\mathbf{x})^\top(\mathbf{y} - \mathbf{x}) + \frac{\lambda}{2}\|\mathbf{x} - \mathbf{y}\|_2^2.$$

Definition 3. *A function $f : \mathcal{X} \to \mathbb{R}$ is α-exp-concave if $\exp(-\alpha f(\cdot))$ is concave on $\mathcal{X}$.*

3.2 Reducing Dynamic Regret to Switching Regret

To clarify our motivation, we first examine the reduction from dynamic regret to switching regret based on a deterministic sequence, but obtain suboptimal guarantees. Motivated by its limitation, we then present our reduction that ensures optimal bounds for three types of losses.

Deterministic sequence. From (1) and (3), we observe that switching regret is a special case of dynamic regret, where comparators are restricted to a sequence with few changes. Hence, a natural idea is to reduce dynamic regret to switching regret by replacing the original comparator sequence with a deterministic *piecewise-constant* one. On the one hand, the piecewise-constant sequence should have few switches to ensure a small switching regret bound. On the other hand, it should stay close to the original sequence to keep the approximation error manageable. However, balancing these goals yields suboptimal bounds, as illustrated below for strongly convex losses.

Let $\boldsymbol{z}_{1:T}$ be a piecewise-constant approximation of $\mathbf{u}_{1:T}$ with tolerance $\varepsilon > 0$, and let $J(\boldsymbol{z}_{1:T})$ denote its number of switches. We initialize $\boldsymbol{z}_1 = \mathbf{u}_1$. For each $t \geq 2$, we set $\boldsymbol{z}_t = \boldsymbol{z}_{t-1}$ if $\|\boldsymbol{z}_{t-1} - \mathbf{u}_t\|_2 \leq \varepsilon$; otherwise, we set $\boldsymbol{z}_t = \mathbf{u}_t$. This construction ensures $J(\boldsymbol{z}_{1:T}) \leq P_T/\varepsilon$ and $\|\boldsymbol{z}_t - \mathbf{u}_t\|_2 \leq \varepsilon$ at every round. By adding and subtracting $f_t(\boldsymbol{z}_t)$ at each round, we obtain

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) = \underbrace{\sum\nolimits_{t=1}^{T} \big[f_t(\mathbf{x}_t) - f_t(\boldsymbol{z}_t)\big]}_{\text{switching regret}} + \underbrace{\sum\nolimits_{t=1}^{T} \big[f_t(\boldsymbol{z}_t) - f_t(\mathbf{u}_t)\big]}_{\text{approximation error}}, \tag{4}$$

where the first term represents the *switching regret* against the comparator sequence $\boldsymbol{z}_{1:T}$, and the second term is the *approximation error* incurred by approximating $\mathbf{u}_{1:T}$ with $\boldsymbol{z}_{1:T}$. To bound the first term, we run an existing switching-regret minimization algorithm for strongly convex losses on $f_1, \ldots, f_T$, obtaining the following bound (Yang et al., 2026):

$$\sum\nolimits_{t=1}^{T} \big[f_t(\mathbf{x}_t) - f_t(\boldsymbol{z}_t)\big] = \mathcal{O}\big((1 + J(\boldsymbol{z}_{1:T})) \log^2 T\big) = \mathcal{O}\big((1 + P_T/\varepsilon) \log^2 T\big), \tag{5}$$

which depends on the number of switches $J(\boldsymbol{z}_{1:T})$ and becomes tighter as this number decreases. To bound the second term, we use the Lipschitz continuity of f_t in Assumption 2 to obtain

$$\sum\nolimits_{t=1}^{T} \big[f_t(\boldsymbol{z}_t) - f_t(\mathbf{u}_t)\big] \leq \sum\nolimits_{t=1}^{T} G\|\boldsymbol{z}_t - \mathbf{u}_t\|_2 \leq GT\varepsilon. \tag{6}$$

Substituting (5) and (6) into (4), we obtain

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) = \mathcal{O}\left(\left(1 + \frac{P_T}{\varepsilon}\right) \log^2 T + GT\varepsilon\right) = \widetilde{\mathcal{O}}(\sqrt{TP_T}), \tag{7}$$

where we set the tolerance $\varepsilon = \widetilde{\mathcal{O}}(\sqrt{P_T/T})$. This result is worse than the optimal $\widetilde{\mathcal{O}}(T^{1/3}P_T^{2/3})$ bound for strongly convex losses (Baby & Wang, 2021). The same analysis also yields suboptimal bounds of $\mathcal{O}(T^{2/3}P_T^{1/3})$ and $\widetilde{\mathcal{O}}(\sqrt{dTP_T})$ for general convex and exp-concave losses, respectively. The essential limitation is that the deterministic construction does not yield a favorable trade-off between the switching regret and the approximation error.

Technical contributions. To overcome this limitation, we exploit the curvature of f_t to construct surrogate losses whose regret upper-bounds the original regret. The resulting surrogate losses admit a quadratic expansion around the original comparator, allowing us to decompose the approximation error into a *linear* term and a *quadratic* term. We then introduce an auxiliary random sequence that is unbiased at each round, with its variance bounded by ε^2 and the expected number of switches bounded by P_T/ε. The unbiasedness eliminates the linear term in expectation, and the variance controls the quadratic term, yielding an $\mathcal{O}(T\varepsilon^2)$ approximation error. By choosing ε appropriately to balance this error with the switching-regret bound, we obtain the optimal dynamic regret bound up to logarithmic factors. Below, we present these two ingredients and the resulting reduction.

We define the surrogate losses ℓ_t on an enlarged domain $\mathcal{Y} \supseteq \mathcal{X}$ and require them to be H-smooth:

$$\ell_t(\mathbf{v}) \leq \ell_t(\mathbf{u}) + \langle \nabla \ell_t(\mathbf{u}), \mathbf{v} - \mathbf{u} \rangle + \frac{H}{2}\|\mathbf{v} - \mathbf{u}\|_2^2 \tag{8}$$

for all $\mathbf{u}, \mathbf{v} \in \mathcal{Y}$ and every t, with a common constant $H \geq 0$. The following lemma establishes the existence of an auxiliary random sequence with the desired properties.

Lemma 1. *Let $\mathcal{X} \subseteq B(\mathbf{0}, R)$ and $\mathcal{Y} = B(\mathbf{0}, 2R)$. For any comparator sequence $\mathbf{u}_{1:T} \in \mathcal{X}$ and any $0 < \varepsilon \leq R$, we can construct an auxiliary random sequence $\mathbf{V}_{1:T}$ taking values in $\mathcal{Y}$ such that*

$$\mathbb{E}[\mathbf{V}_t] = \mathbf{u}_t, \qquad \mathbb{E}\left[\|\mathbf{V}_t - \mathbf{u}_t\|_2^2\right] \leq \varepsilon^2, \qquad \mathbb{E}[J(\mathbf{V}_{1:T})] \leq \frac{P_T}{\varepsilon},$$

where the first two properties hold for every round and the expectations are taken over $\mathbf{V}_{1:T}$.

Note that the random sequence in Lemma 1 may lie outside the original domain $\mathcal{X}$, so we define the surrogate losses on the enlarged domain $\mathcal{Y}$, which provably controls of the original dynamic regret. Specifically, by projecting predictions of the learner onto $\mathcal{X}$ and applying a suitable gradient correction, the dynamic regret under f_t is upper-bounded by that under ℓ_t against the same comparator sequence (Cutkosky & Orabona, 2018). We provide the detailed construction in Section 4. In the following, we hence focus on analyzing the dynamic regret of surrogate losses over $\mathcal{Y}$.

Let $\mathbf{y}_t \in \mathcal{Y}$ denote the predictions of a switching-regret minimization algorithm run on the surrogate losses $\ell_1, \ldots, \ell_T$. Analogously to (4), we use the auxiliary random sequence $\mathbf{V}_{1:T}$ from Lemma 1 as an intermediate comparator. By adding and subtracting $\ell_t(\mathbf{V}_t)$ at each round and taking expectations over the randomness of $\mathbf{V}_{1:T}$, we decompose the dynamic regret under the surrogate losses into the expected switching regret against $\mathbf{V}_{1:T}$ and the expected approximation error:

$$\sum\nolimits_{t=1}^{T}\big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{u}_t)\big] = \mathbb{E}\left[\sum\nolimits_{t=1}^{T}\big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{V}_t)\big]\right] + \mathbb{E}\left[\sum\nolimits_{t=1}^{T}\big[\ell_t(\mathbf{V}_t) - \ell_t(\mathbf{u}_t)\big]\right]. \quad (9)$$

For λ-strongly convex losses, the corresponding surrogate losses are also λ-strongly convex. We can therefore run the existing switching-regret minimization algorithm (Yang et al., 2026) on the surrogate loss sequence $\ell_1, \ldots, \ell_T$ to obtain a bound analogous to (5):

$$\mathbb{E}\left[\sum\nolimits_{t=1}^{T}\big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{V}_t)\big]\right] \leq \mathbb{E}\left[\mathcal{O}\big((1 + J(\mathbf{V}_{1:T}))\log^2 T\big)\right] = \mathcal{O}\big((1 + P_T/\varepsilon)\log^2 T\big). \quad (10)$$

By the property of ℓ_t in (8), the approximation error is bounded by

$$\mathbb{E}\left[\sum_{t=1}^{T}\big[\ell_t(\mathbf{V}_t) - \ell_t(\mathbf{u}_t)\big]\right] \leq \mathbb{E}\left[\sum_{t=1}^{T}\langle\nabla\ell_t(\mathbf{u}_t), \mathbf{V}_t - \mathbf{u}_t\rangle\right] + \frac{H}{2}\mathbb{E}\left[\sum_{t=1}^{T}\|\mathbf{V}_t - \mathbf{u}_t\|_2^2\right] \leq \frac{HT\varepsilon^2}{2}, \quad (11)$$

where the last inequality follows from Lemma 1: the linear term vanishes by unbiasedness, i.e., $\mathbb{E}[\mathbf{V}_t - \mathbf{u}_t] = \mathbf{0}$, and the quadratic term is controlled by the variance bound, i.e., $\mathbb{E}[\|\mathbf{V}_t - \mathbf{u}_t\|_2^2] \leq \varepsilon^2$.

Combining (9), (10) and (11) delivers

$$\sum_{t=1}^{T}\big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{u}_t)\big] = \mathcal{O}\left(\left(1 + \frac{P_T}{\varepsilon}\right)\log^2 T + HT\varepsilon^2\right) = \widetilde{\mathcal{O}}\left(T^{1/3}P_T^{2/3}\right), \quad (12)$$

where we set the tolerance $\varepsilon = \widetilde{\mathcal{O}}(T^{-1/3}P_T^{1/3})$. Note that this bound attains the optimal guarantee for strongly convex losses, up to logarithmic factors (Baby & Wang, 2021).

The following theorem formalizes this reduction for any learner with a switching-regret guarantee.

Theorem 2. *Let $\mathcal{X} \subseteq B(\mathbf{0}, R)$ and $\mathcal{Y} = B(\mathbf{0}, 2R)$ for some $R > 0$. Suppose the surrogate losses $\ell_t : \mathcal{Y} \to \mathbb{R}$ satisfy* (8) *with some $H \geq 0$. An online learner run on $\ell_1, \ldots, \ell_T$ with predictions $\mathbf{y}_t \in \mathcal{Y}$ guarantees*

$$\sum\nolimits_{t=1}^{T}\big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{v}_t)\big] \leq \Phi\big(T, J(\mathbf{v}_{1:T})\big)$$

for any comparator sequence $\mathbf{v}_{1:T} \in \mathcal{Y}$, where $J(\mathbf{v}_{1:T})$ denotes the number of switches in $\mathbf{v}_{1:T}$, and $\Phi(T, J)$ is the switching regret bound of the learner over T rounds against comparator sequences with J switches. Then, for any comparator sequence $\mathbf{u}_{1:T} \in \mathcal{X}$ and any $0 < \varepsilon \leq R$, we have

$$\sum_{t=1}^{T}\big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{u}_t)\big] \leq \mathbb{E}\big[\Phi(T, J(\mathbf{V}_{1:T}))\big] + \frac{HT\varepsilon^2}{2}, \quad (13)$$

where $\mathbf{V}_{1:T}$ is the auxiliary random sequence constructed in Lemma 1. The expectation is taken only over its auxiliary randomness, with surrogate losses and predictions of the learner held fixed.

Algorithm 1 Switching-to-Dynamic Regret Minimization

Input: Horizon T, domain $\mathcal{X}$, bounds R and G, loss type, and curvature parameter (if applicable).
1: Select the enlarged domain $\mathcal{Y}$ and the switching-regret learner $\mathcal{L}$ according to the loss type.
2: Initialize $\mathcal{L}$ on $\mathcal{Y}$.
3: **for** $t = 1, \dots, T$ **do**
4: Obtain the auxiliary prediction $\mathbf{y}_t \in \mathcal{Y}$ from $\mathcal{L}$.
5: Play $\mathbf{x}_t = \Pi_{\mathcal{X}}(\mathbf{y}_t)$ and observe $\nabla f_t(\mathbf{x}_t)$.
6: Construct the surrogate loss ℓ_t as specified in Sections 4.1–4.3.
7: Update $\mathcal{L}$ with ℓ_t.
8: **end for**

Remark. Theorem 2 provides a unified reduction for all three loss classes. For λ-strongly convex and α-exp-concave losses, the surrogate losses in Sections 4.1 and 4.2 satisfy (8) with $H = \lambda$ and $H = \beta G^2$, respectively, yielding an $\mathcal{O}(T\varepsilon^2)$ approximation error. For general convex losses, the surrogate loss in Section 4.3 satisfies the same condition with $H = 0$, so the approximation error vanishes. In each case, the dynamic regret bound follows by substituting the corresponding switching-regret guarantee into (13) and choosing ε appropriately. The instantiations below yield bounds that are optimal in T and P_T, up to logarithmic factors, through this common reduction.

4 METHODS

We now instantiate the reduction in Theorem 2 for strongly convex, exp-concave, and general convex losses. To implement the reduction, we construct surrogate losses on the enlarged domain $\mathcal{Y}$ using gradient feedback at feasible decisions in $\mathcal{X}$. These losses retain the curvature needed by the switching-regret learner, satisfy (8), and ensure that their regret upper-bounds the original regret. This allows us to translate the guarantees from Section 3.2 into dynamic regret bounds for the actual decisions. The auxiliary random sequence is used only in the analysis and need not be generated.

Algorithm 1 summarizes the procedure. We first select the enlarged domain $\mathcal{Y}$ and the switching-regret learner $\mathcal{L}$ according to the loss type, and initialize $\mathcal{L}$ on $\mathcal{Y}$ (Lines 1–2). At each round t, we obtain an auxiliary prediction $\mathbf{y}_t \in \mathcal{Y}$ from $\mathcal{L}$ (Line 4), play its projection $\mathbf{x}_t = \Pi_{\mathcal{X}}(\mathbf{y}_t)$, and observe $\nabla f_t(\mathbf{x}_t)$ (Line 5). We then construct the surrogate loss ℓ_t for the given loss type and update $\mathcal{L}$ with it, using normalization when required (Lines 6–7). Below, we specify the configurations and present the corresponding dynamic regret guarantees for the three loss classes.

4.1 STRONGLY CONVEX LOSSES

Initially, we consider the case where f_t is λ-strongly convex, and choose the surrogate loss

$$\ell_t^{sc}(\mathbf{y}) = \langle \mathbf{h}_t, \mathbf{y} - \mathbf{y}_t \rangle + \frac{\lambda}{2}\|\mathbf{y} - \mathbf{x}_t\|_2^2, \tag{14}$$

where the corrected gradient is set as (Cutkosky, 2020)

$$\mathbf{h}_t = \nabla f_t(\mathbf{x}_t) - \mathbf{1}\{c_t < 0\}\, c_t \mathbf{n}_t,$$

with $c_t = \langle \nabla f_t(\mathbf{x}_t), \mathbf{n}_t \rangle$ and

$$\mathbf{n}_t = \begin{cases} (\mathbf{y}_t - \mathbf{x}_t)/\|\mathbf{y}_t - \mathbf{x}_t\|_2, & \mathbf{y}_t \neq \mathbf{x}_t, \\ 0, & \mathbf{y}_t = \mathbf{x}_t. \end{cases}$$

It can be readily verified that (14) satisfies (8) with $H = \lambda$ and enjoys the following property:

$$f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \leq \ell_t(\mathbf{y}_t) - \ell_t(\mathbf{u}) \tag{15}$$

for every $\mathbf{u} \in \mathcal{X}$ and $\mathbf{y}_t \in \mathcal{Y}$, with $\mathbf{x}_t = \Pi_{\mathcal{X}}(\mathbf{y}_t)$ (Yang et al., 2024). Therefore, we can conveniently derive dynamic regret bounds for the original losses f_t by analyzing that for the surrogate losses ℓ_t.

Next, we examine the diameter of the domains $\mathcal{X}$ and $\mathcal{Y}$. For any $\mathbf{x}, \mathbf{u} \in \mathcal{X}$, the λ-strong convexity of f_t and Assumption 2 imply that

$$\lambda\|\mathbf{x} - \mathbf{u}\|_2^2 \leq \langle \nabla f_t(\mathbf{x}) - \nabla f_t(\mathbf{u}), \mathbf{x} - \mathbf{u} \rangle \leq 2G\|\mathbf{x} - \mathbf{u}\|_2,$$

which indicates $\mathrm{diam}(\mathcal{X}) \leq 2G/\lambda$. In other words, $\mathcal{X}$ is contained in a ball $B(\mathbf{0}, 2G/\lambda)$ and we choose the enlarged set $\mathcal{Y} = B(\mathbf{0}, 4G/\lambda)$ with the radius $4G/\lambda$. We then bound the gradient norm of the surrogate loss in (14). Since $\|\mathbf{h}_t\|_2 \leq G$ and $\mathbf{x}_t \in \mathcal{X} \subseteq \mathcal{Y}$, for any $\mathbf{y} \in \mathcal{Y}$ we have

$$\|\nabla \ell_t^{sc}(\mathbf{y})\|_2 = \|\mathbf{h}_t + \lambda(\mathbf{y} - \mathbf{x}_t)\|_2 \leq G + \lambda \cdot 8G/\lambda = 9G.$$

With these configurations, we instantiate the switching-regret learner $\mathcal{L}$ with IRESET-OGD (Yang et al., 2026), and present the dynamic regret bound for strongly convex losses below.

Theorem 3. *Under Assumptions 1 and 2, suppose f_t is λ-strongly convex. With the configurations above, Algorithm 1 ensures, for every comparator sequence $\mathbf{u}_{1:T} \in \mathcal{X}$,*

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) = \mathcal{O}\left(\frac{G^2}{\lambda}(1+\log T) + \left(\frac{G^4 T P_T^2}{\lambda}\right)^{1/3}(1+\log T)^{4/3}\right) = \widetilde{\mathcal{O}}\left(T^{1/3}P_T^{2/3}\right).$$

Remark. Theorem 3 matches the minimax optimal bound, up to logarithmic factors (Baby & Wang, 2021) and eliminates the dimension dependence in previous results (Baby & Wang, 2022). Notably, our analysis uses the existing switching-regret guarantee of IRESET-OGD as a black box and combines it with the properties of the auxiliary random sequence in Lemma 1. This yields a simple proof without analyzing the structure of an offline optimal comparator sequence.

4.2 Exp-Concave Losses

In this part, we consider the case where f_t is α-exp-concave, and choose the surrogate loss:

$$\ell_t^{\exp}(\mathbf{y}) = \langle \mathbf{h}_t, \mathbf{y} - \mathbf{y}_t \rangle + \frac{\beta}{2}\langle \mathbf{h}_t, \mathbf{y} - \mathbf{y}_t \rangle^2, \tag{16}$$

in which $\beta = \min\{1/(32GR), \alpha/2\}$. We can also verify that (16) satisfies (8) with $H = \beta G^2$ and (15) (Mhammedi et al., 2019). Moreover, the value of β delivers $R \leq 1/(32\beta G)$ and hence $\mathcal{X} \subseteq B(\mathbf{0}, 1/(32\beta G))$. We therefore set the enlarged domain to $\mathcal{Y} = B(\mathbf{0}, 1/(16\beta G))$, with the radius $1/(16\beta G)$. Since $\|\mathbf{h}_t\|_2 \leq G$, we have

$$\|\nabla \ell_t^{\exp}(\mathbf{y})\|_2 = |1 + \beta\langle \mathbf{h}_t, \mathbf{y} - \mathbf{y}_t \rangle| \, \|\mathbf{h}_t\|_2 \leq \left(1 + \beta G \cdot \frac{1}{8\beta G}\right) G = \frac{9}{8}G \leq \sqrt{2}G.$$

Instantiating the switching-regret learner with IRESET-ONS (Yang et al., 2026) yields the following dynamic regret bound for exp-concave losses.

Theorem 4. *Under Assumptions 1 and 2, suppose f_t is α-exp-concave. With the configurations above, Algorithm 1 ensures, for every comparator sequence $\mathbf{u}_{1:T} \in \mathcal{X}$,*

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) = \mathcal{O}\left(\frac{d}{\beta}(1+\log T) + \left(\frac{d^2 G^2 T P_T^2}{\beta}\right)^{1/3}(1+\log T)^{4/3}\right) = \widetilde{\mathcal{O}}\left(d^{2/3}T^{1/3}P_T^{2/3}\right).$$

Remark. Theorem 4 achieves the minimax-optimal bound for exp-concave losses, up to additional logarithmic factors (Baby & Wang, 2021). Notably, the dependence on d is inherited from the switching regret guarantee of IRESET with ONS experts, and our reduction introduces no additional dependence on d. In other words, any improvement in the dimension dependence of the switching regret guarantee can be carried over to the dynamic regret bound through our reduction.

Remark. Compared with Zhang et al. (2025), Theorem 4 improves the dimension dependence from d to $d^{2/3}$. Moreover, our framework requires only running a single instance of *any* existing switching-regret algorithm over the surrogate loss (16), with its predictions projected onto $\mathcal{X}$. This construction avoids the complicate KL projections and makes it easier to understand and implement.

4.3 General Convex Losses

We further investigate the case where f_t is general convex, and use the following surrogate loss:

$$\ell_t^{cov}(\mathbf{y}) = \langle \mathbf{h}_t, \mathbf{y} - \mathbf{y}_t \rangle, \tag{17}$$

which also satisfies (8) with $H = 0$, as well as (15) (Cutkosky, 2020). Moreover, its gradient satisfies $\|\nabla \ell_t^{\text{cov}}(\mathbf{y})\|_2 = \|\mathbf{h}_t\|_2 \leq G$ by Assumption 2. Therefore, we set the enlarged domain $\mathcal{Y} = B(\mathbf{0}, 2R)$ and normalize the surrogate losses to $[0, 1]$. Then, instantiating the switching-regret learner with RESET (Pasteris et al., 2024) yields the following dynamic regret bound.

Theorem 5. *Under Assumptions 1 and 2, suppose f_t is general convex. With the configuration above, Algorithm 1 ensures, for every comparator sequence $\mathbf{u}_{1:T} \in \mathcal{X}$,*

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) = \mathcal{O}\left(G\sqrt{T(R^2 + RP_T)}\right) = \mathcal{O}\left(\sqrt{T(1+P_T)}\right). \tag{18}$$

Remark. Theorem 5 matches the optimal result for general convex losses (Zhang et al., 2018a). Theorems 3, 4, and 5 together show that the same reduction applies to all three types of losses with optimal dynamic regret bounds, highlighting the simplicity and generality of our framework.

5 Theoretical Analysis

Due to the limitation of space, we only provide the proof of Lemma 1and Theorem 3. The omitted proofs can be found in the supplementary material. Throughout these proofs, all expectations are taken over the auxiliary random sequence $\mathbf{V}_{1:T}$ constructed in Lemma 1.

5.1 Proof of Lemma 1

We first show how to construct an auxiliary random sequence $\mathbf{V}_{1:T}$ for the comparator sequence $\mathbf{u}_{1:T}$. The basic idea is to keep the auxiliary sequence unchanged with high probability when $\mathbf{u}_{1:T}$ moves only slightly. When an update occurs, we move to a point that compensates for this small update probability. Specifically, given $\varepsilon > 0$, we initialize $\mathbf{V}_1 = \mathbf{u}_1$ and define, for each $t \geq 2$,

$$\boldsymbol{\delta}_t = \mathbf{u}_t - \mathbf{u}_{t-1}, \qquad \Delta_t = \|\boldsymbol{\delta}_t\|_2, \qquad p_t = \min\{1, \Delta_t/\varepsilon\}.$$

If $\Delta_t = 0$, we keep $\mathbf{V}_t = \mathbf{V}_{t-1}$. Otherwise, independently of previous random choices, we set

$$\mathbf{V}_t = \begin{cases} \mathbf{u}_{t-1} + \boldsymbol{\delta}_t/p_t, & \text{with probability } p_t, \\ \mathbf{V}_{t-1}, & \text{with probability } 1 - p_t. \end{cases} \tag{19}$$

When $0 < \Delta_t < \varepsilon$, the sequence updates with probability $p_t = \Delta_t/\varepsilon$ to a point at distance ε from $\mathbf{u}_{t-1}$ in the direction of $\boldsymbol{\delta}_t$. The factor $1/p_t$ in (19) makes up for the small update probability, so that the mean still moves by $\boldsymbol{\delta}_t$. When $\Delta_t \geq \varepsilon$, we simply set $\mathbf{V}_t = \mathbf{u}_t$.

We now state the properties of this construction. Let $J(\mathbf{V}_{1:T}) = \sum_{t=2}^{T} \mathbf{1}\{\mathbf{V}_t \neq \mathbf{V}_{t-1}\}$ denote the number of switches, so that the sequence $\mathbf{V}_{1:T}$ has $J(\mathbf{V}_{1:T}) + 1$ constant pieces. Let $B(\mathbf{0}, R) = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \leq R\}$ denote the Euclidean ball of radius R. Since the update point can lie beyond $\mathbf{u}_t$, we allow the random sequence to take values in a larger ball than the original comparator.

Next, we show that every $\mathbf{V}_t$ lies in the enlarged domain $\mathcal{Y} = B(\mathbf{0}, 2R)$. The initial point $\mathbf{V}_1 = \mathbf{u}_1$ lies in $\mathcal{X} \subseteq B(\mathbf{0}, R)$. At an update, the new point is either $\mathbf{u}_t$ when $p_t = 1$, or a point at distance $\varepsilon \leq R$ from $\mathbf{u}_{t-1}$ when $0 < p_t < 1$. In both update cases $\mathbf{V}_t$ lies in $B(\mathbf{0}, 2R)$. If no update occurs at round t, this property still holds, since $\mathbf{V}_t = \mathbf{V}_{t-1} \in B(\mathbf{0}, 2R)$.

Then, we prove the unbiasedness by induction. The claim holds at $t = 1$. If $p_t = 0$, both the comparator and the random sequence remain unchanged. If $p_t > 0$, the independence of the new random choice and the induction hypothesis give

$$\mathbb{E}[\mathbf{V}_t] = (1 - p_t)\mathbf{u}_{t-1} + p_t\left(\mathbf{u}_{t-1} + \frac{\boldsymbol{\delta}_t}{p_t}\right) = \mathbf{u}_t.$$

To bound the variance, let $W_t = \mathbb{E}[\|\mathbf{V}_t - \mathbf{u}_t\|_2^2]$, with $W_1 = 0$. If $p_t = 0$, then $W_t = W_{t-1}$. For $p_t > 0$, expanding the two cases in (19) and using $\mathbb{E}[\mathbf{V}_{t-1} - \mathbf{u}_{t-1}] = \mathbf{0}$ yield

$$W_t = (1 - p_t)(W_{t-1} + \Delta_t^2) + p_t\left(\frac{1 - p_t}{p_t}\right)^2 \Delta_t^2 = (1 - p_t)W_{t-1} + \frac{1 - p_t}{p_t}\Delta_t^2.$$

If $p_t = 1$, then $W_t = 0$. If $0 < p_t < 1$, then $\Delta_t = p_t\varepsilon$. Thus, assuming $W_{t-1} \leq \varepsilon^2$, we obtain

$$W_t \leq (1 - p_t)\varepsilon^2 + p_t(1 - p_t)\varepsilon^2 = (1 - p_t^2)\varepsilon^2 \leq \varepsilon^2.$$

This proves the variance bound for all t. Finally, the sequence can switch only when an update occurs. Therefore, it holds that

$$\mathbb{E}[J(\mathbf{V}_{1:T})] = \sum\nolimits_{t=2}^{T} \Pr(\mathbf{V}_t \neq \mathbf{V}_{t-1}) \leq \sum\nolimits_{t=2}^{T} p_t \leq \frac{1}{\varepsilon}\sum\nolimits_{t=2}^{T} \Delta_t = \frac{P_T}{\varepsilon}.$$

5.2 Proof of Theorem 3

Let $\mathcal{S}(\mathbf{V}_{1:T})$ denote the partition of $\mathbf{V}_{1:T}$ into maximal constant intervals, so that $|\mathcal{S}(\mathbf{V}_{1:T})| - 1 = J(\mathbf{V}_{1:T})$. Define its switching regret under the surrogate losses as

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) = \sum\nolimits_{t=1}^{T} \big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{V}_t)\big].$$

We state the switching-regret bound of IRESET-OGD for strongly convex losses (Yang et al., 2026).

Lemma 6. *Let $\mathcal{Y}$ be a convex domain contained in $B(\mathbf{0}, 4G/\lambda)$ and $\ell_1, \dots, \ell_T$ be the surrogate losses in* (14) *with gradient norms at most $9G$. Then, IRESET-OGD ensures, for every realization of $\mathbf{V}_{1:T}$ in $\mathcal{Y}$*

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le [1 + 4(2 + \log_2 T)J(\mathbf{V}_{1:T})] \left[2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T))\right], \tag{20}$$

where

$$\begin{aligned} \Xi &= 144\Gamma_T \frac{G^2}{\lambda}\left(2 + \frac{1}{\sqrt{\ln 2}}\right) + \frac{81\Gamma_T^2 G^2}{2\lambda \ln 2}, \\ \Gamma_T &= 3\ln 2 + \ln\left(1 + \frac{1}{e}(1 + \ln(2T+1))\right). \end{aligned} \tag{21}$$

Let $\Phi(T, J)$ denote the right-hand side of (20), with $J(\mathbf{V}_{1:T})$ replaced by J. For any $0 < \varepsilon \le 2G/\lambda$, applying Theorem 2 with this choice of Φ and $H = \lambda$, together with $\mathbb{E}[J(\mathbf{V}_{1:T})] \le P_T/\varepsilon$, yields

$$\begin{aligned} \text{D-Reg}_T(\mathbf{u}_{1:T}) &\overset{(15)}{\le} \sum\nolimits_{t=1}^{T} \big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{u}_t)\big] \overset{(13)}{\le} \mathbb{E}\big[\Phi(T, J(\mathbf{V}_{1:T}))\big] + \frac{\lambda T \varepsilon^2}{2} \\ &\le \left(1 + \frac{4(2 + \log_2 T)P_T}{\varepsilon}\right)\left[2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T))\right] + \frac{\lambda T \varepsilon^2}{2}, \end{aligned} \tag{22}$$

of which the optimal minimizer is

$$\varepsilon_* = \left\{\frac{4(2 + \log_2 T)P_T}{\lambda T}\left[2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T))\right]\right\}^{1/3}.$$

If $\varepsilon_* \le 2G/\lambda$, substituting it into (22) gives

$$\begin{aligned} \text{D-Reg}_T(\mathbf{u}_{1:T}) &\le 2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T)) \\ &\quad + \frac{3}{2}(\lambda T)^{1/3}\left\{4(2 + \log_2 T)P_T\left[2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T))\right]\right\}^{2/3}. \end{aligned} \tag{23}$$

If $\varepsilon_* > 2G/\lambda$, according to $\operatorname{diam}(\mathcal{X}) \le 2G/\lambda$, we have

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) \le G\sum\nolimits_{t=1}^{T} \|\mathbf{x}_t - \mathbf{u}_t\|_2 \le \frac{2G^2 T}{\lambda} \le \lambda T \varepsilon_*^2,$$

where the last step is due to $\varepsilon_* \ge 2G/\lambda$, and $\lambda T \varepsilon_*^2$ is bounded by the right-hand side of (23). Thus, (23) holds for every comparator sequence. Combining the two cases finishes the proof.

6 Conclusion

In this paper, we study dynamic regret minimization in non-stationary online learning and propose a novel framework that reduces this problem to switching regret minimization. With an unbiased random sequence and suitable surrogate losses, our framework allows us to use off-the-shelf switching-regret algorithms to establish the dynamic regret bounds of $\widetilde{\mathcal{O}}(T^{1/3}P_T^{2/3})$ for strongly convex losses and $\widetilde{\mathcal{O}}(d^{2/3}T^{1/3}P_T^{2/3})$ for exp-concave losses. The same framework also ensures an $\mathcal{O}(\sqrt{T(1+P_T)})$ dynamic regret bound for general convex losses. All our findings are optimal for three types of losses, by a *simple* and *unified* analysis. Algorithmically, the auxiliary unbiased random sequence is used only for analysis, and need not be generated explicitly in practice. Our framework can thus use *any* algorithm with switching-regret guarantees as a black box, without modifying its internal updates, to achieve corresponding dynamic regret bounds.

AI USE STATEMENT

In this work, we used generative AI tools to polish the writing and assist with checking mathematical proofs. In particular, GPT-6-Astra helps us find a way to construct the auxiliary random sequence. We take full responsibility for the final content of this work, including text and claims developed with the assistance of generative AI.

REFERENCES

Dheeraj Baby and Yu-Xiang Wang. Online forecasting of total-variation-bounded sequences. In *Advances in Neural Information Processing Systems 32*, pp. 11071–11081, 2019.

Dheeraj Baby and Yu-Xiang Wang. Optimal dynamic regret in exp-concave online learning. In *Proceedings of the 34th Conference on Learning Theory*, pp. 359–409, 2021.

Dheeraj Baby and Yu-Xiang Wang. Optimal dynamic regret in proper online learning with strongly convex losses and beyond. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, pp. 1805–1845, 2022.

Omar Besbes, Yonatan Gur, and Assaf J. Zeevi. Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244, 2015.

Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

Ashok Cutkosky. Parameter-free, dynamic, and strongly-adaptive online learning. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 2250–2259, 2020.

Ashok Cutkosky and Francesco Orabona. Black-box reductions for parameter-free online learning in banach spaces. In *Proceedings of the 31st Conference on Learning Theory*, pp. 1493–1529, 2018.

Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.

Eric C. Hall and Rebecca M. Willett. Dynamical models and tracking regret in online convex programming. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 579–587, 2013.

Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3–4):157–325, 2016.

Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Foundations and Trends in Machine Learning*, 69(2):169–192, 2007.

Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, pp. 398–406, 2015.

Zakaria Mhammedi, Wouter M. Koolen, and Tim van Erven. Lipschitz adaptivity with multiple learning rates in online learning. In *Proceedings of the 32nd Conference on Learning Theory*, pp. 2490–2511, 2019.

Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *55th IEEE Conference on Decision and Control*, pp. 7195–7201, 2016.

Francesco Orabona. A modern introduction to online learning. *ArXiv e-prints*, arXiv:1912.13213v6, 2019.

Stephen Pasteris, Chris Hicks, Vasilios Mavroudis, and Mark Herbster. Online convex optimisation: The optimal switching regret for all segmentations simultaneously. In *Advances in Neural Information Processing Systems 37*, pp. 78278–78298, 2024.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.

Yuanyu Wan, Bo Xue, and Lijun Zhang. Projection-free online learning in dynamic environments. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pp. 10067–10075, 2021.

Yuanyu Wan, Lijun Zhang, and Mingli Song. Improved dynamic regret for online frank-wolfe. In *Proceedings of the 36th Conference on Learning Theory*, pp. 3304–3327, 2023.

Yibo Wang, Wenhao Yang, Wei Jiang, Shiyin Lu, Bing Wang, Haihong Tang, Yuanyu Wan, and Lijun Zhang. Non-stationary projection-free online learning with dynamic and adaptive regret guarantees. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, pp. 15671–15679, 2024.

Chen-Yu Wei, Yi-Te Hong, and Chi-Jen Lu. Tracking the best expert in non-stationary stochastic environments. In *Advances in Neural Information Processing Systems 29*, pp. 3972–3980, 2016.

Zhipan Xu and Lijun Zhang. Online non-convex learning in dynamic environments. In *Advances in Neural Information Processing Systems 37*, pp. 51930–51962, 2024.

Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng Yi. Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient. In *Proceedings of the 33rd International Conference on Machine Learning*, pp. 449–457, 2016.

Wenhao Yang, Yibo Wang, Peng Zhao, and Lijun Zhang. Universal online convex optimization with 1 projection per round. In *Advances in Neural Information Processing Systems 37*, 2024.

Wenhao Yang, Yibo Wang, Yuanyu Wan, and Lijun Zhang. Logarithmic switching regret for online convex optimization. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026. To appear.

Lijun Zhang, Tianbao Yang, Jinfeng Yi, Rong Jin, and Zhi-Hua Zhou. Improved dynamic regret for non-degenerate functions. In *Advances in Neural Information Processing Systems 30*, pp. 732–741, 2017.

Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems 31*, pp. 1323–1333, 2018a.

Lijun Zhang, Tianbao Yang, Rong Jin, and Zhi-Hua Zhou. Dynamic regret of strongly adaptive methods. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 5877–5886, 2018b.

Lijun Zhang, Shiyin Lu, and Tianbao Yang. Minimizing dynamic regret and adaptive regret simultaneously. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pp. 309–319, 2020.

Lijun Zhang, Wei Jiang, Shiyin Lu, and Tianbao Yang. Revisiting smoothed online learning. In *Advances in Neural Information Processing Systems 34*, pp. 13599–13612, 2021.

Yu-Jie Zhang, Peng Zhao, and Masashi Sugiyama. Non-stationary online learning for curved losses: Improved dynamic regret via mixability. In *Proceedings of the 42nd International Conference on Machine Learning*, pp. 77044–77068, 2025.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, pp. 928–936, 2003.

A PROOFS

In this part, we present the proofs of Theorems 2, 4 and 5, followed by those of Lemma 6, 7 and 8.

A.1 PROOF OF THEOREM 2

We bound the two terms in (9) separately. For the first term, applying the assumed switching-regret guarantee to each realization of $\mathbf{V}_{1:T}$ and taking expectations gives

$$\mathbb{E}\left[\sum_{t=1}^{T}\big[\ell_t(\mathbf{y}_t)-\ell_t(\mathbf{V}_t)\big]\right]\le\mathbb{E}\big[\Phi(T,J(\mathbf{V}_{1:T}))\big]. \tag{24}$$

For the second term, applying (8) with $\mathbf{u}=\mathbf{u}_t$ and $\mathbf{v}=\mathbf{V}_t$ yields

$$\mathbb{E}\big[\ell_t(\mathbf{V}_t)-\ell_t(\mathbf{u}_t)\big]\le\langle\nabla\ell_t(\mathbf{u}_t),\mathbb{E}[\mathbf{V}_t-\mathbf{u}_t]\rangle+\frac{H}{2}\mathbb{E}\big[\|\mathbf{V}_t-\mathbf{u}_t\|_2^2\big]\le\frac{H\varepsilon^2}{2}, \tag{25}$$

where the last inequality follows from the unbiasedness and variance bound in Lemma 1.

Summing (25) over all rounds and combining it with (24) and (9) gives (13), completing the proof.

A.2 PROOF OF THEOREM 4

We provide the switching regret bound of IRESET-ONS for exp-concave losses (Yang et al., 2026).

Lemma 7. *Let $\mathcal{Y}$ be a convex domain contained in $B(\mathbf{0},1/(16\beta G))$ and $\ell_1,\dots,\ell_T$ be the surrogate losses in* (16) *with gradient norms at most $\sqrt{2}G$. Then, IRESET-ONS ensures, for every realization of $\mathbf{V}_{1:T}$ in $\mathcal{Y}$*

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T})\le[1+4(2+\log_2 T)J(\mathbf{V}_{1:T})]\left[2\Xi_{\exp}+\frac{5d}{\beta}\left(4+\frac{\sqrt{2}}{8}\right)(1+\ln(2T))\right], \tag{26}$$

where

$$\Xi_{\exp}=\frac{1}{\beta}\left[\frac{\sqrt{2}}{4}\Gamma_T\left(2+\frac{1}{\sqrt{\ln 2}}\right)+\frac{4\Gamma_T^2}{\ln 2}\right], \tag{27}$$

with Γ_T defined in (21)*.*

For any $0<\varepsilon\le 1/(32\beta G)$, combining Theorem 2 with Lemma 7 and $\mathbb{E}[J(\mathbf{V}_{1:T})]\le P_T/\varepsilon$ from Lemma 1 yields

$$\begin{aligned}\text{D-Reg}_T(\mathbf{u}_{1:T})&\overset{(15)}{\le}\sum_{t=1}^{T}\big[\ell_t(\mathbf{y}_t)-\ell_t(\mathbf{u}_t)\big]\overset{(13)}{\le}\mathbb{E}\big[\Phi(T,J(\mathbf{V}_{1:T}))\big]+\frac{\beta G^2T\varepsilon^2}{2}\\&\le\left(1+\frac{4(2+\log_2 T)P_T}{\varepsilon}\right)\left[2\Xi_{\exp}+\frac{5d}{\beta}\left(4+\frac{\sqrt{2}}{8}\right)(1+\ln(2T))\right]+\frac{\beta G^2T\varepsilon^2}{2},\end{aligned} \tag{28}$$

of which the optimal minimizer is

$$\varepsilon_*=\left\{\frac{4(2+\log_2 T)P_T}{\beta G^2T}\left[2\Xi_{\exp}+\frac{5d}{\beta}\left(4+\frac{\sqrt{2}}{8}\right)(1+\ln(2T))\right]\right\}^{1/3}. \tag{29}$$

If $\varepsilon_*\le 1/(32\beta G)$, substituting it into (28) and enlarging the coefficient $3/2$ to 64 gives

$$\begin{aligned}\text{D-Reg}_T(\mathbf{u}_{1:T})\le{}&2\Xi_{\exp}+\frac{5d}{\beta}\left(4+\frac{\sqrt{2}}{8}\right)(1+\ln(2T))\\&+64(\beta G^2T)^{1/3}\left[4(2+\log_2 T)P_T\right]^{2/3}\\&\times\left[2\Xi_{\exp}+\frac{5d}{\beta}\left(4+\frac{\sqrt{2}}{8}\right)(1+\ln(2T))\right]^{2/3}.\end{aligned} \tag{30}$$

If $\varepsilon_* > 1/(32\beta G)$, according to $\mathrm{diam}(\mathcal{X}) \le 2R$ and $\beta GR \le 1/32$, we have

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) \le G\sum_{t=1}^{T}\|\mathbf{x}_t - \mathbf{u}_t\|_2 \le 2GRT \le \frac{T}{16\beta} < 64\beta G^2 T\varepsilon_*^2,$$

where the last step is due to $\varepsilon_* > 1/(32\beta G)$, and $64\beta G^2 T\varepsilon_*^2$ is bounded by the right-hand side of (30). Thus, (30) holds for every comparator sequence. Combining the two cases finishes the proof.

A.3 Proof of Theorem 5

We present the switching regret bound of RESET for general convex losses (Pasteris et al., 2024).

Lemma 8. *Let $\mathcal{Y}$ be a convex domain contained in $B(\mathbf{0}, 2R)$ and $\ell_1, \dots, \ell_T$ be the surrogate losses in* (17) *with gradient norms at most G. Then, RESET with OGD as its expert algorithm ensures, for every realization of $\mathbf{V}_{1:T}$ in $\mathcal{Y}$,*

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le 4\sqrt{2}GR\left[\frac{\sqrt{2}}{\sqrt{2}-1} + \frac{\sqrt{8\ln 2}}{3-2\sqrt{2}}\right]\sqrt{T\big(1 + J(\mathbf{V}_{1:T})\big)}. \tag{31}$$

For any $0 < \varepsilon \le R$, combining Theorem 2 with Lemma 8 and $\mathbb{E}[J(\mathbf{V}_{1:T})] \le P_T/\varepsilon$ from Lemma 1 yields

$$\begin{aligned}
\text{D-Reg}_T(\mathbf{u}_{1:T}) &\overset{(15)}{\le} \sum_{t=1}^{T}\big[\ell_t(\mathbf{y}_t) - \ell_t(\mathbf{u}_t)\big] \overset{(13)}{\le} \mathbb{E}\big[\Phi(T, J(\mathbf{V}_{1:T}))\big] \\
&\le 4\sqrt{2}GR\left[\frac{\sqrt{2}}{\sqrt{2}-1} + \frac{\sqrt{8\ln 2}}{3-2\sqrt{2}}\right]\sqrt{T\big(1 + \mathbb{E}[J(\mathbf{V}_{1:T})]\big)} \\
&\le 4\sqrt{2}GR\left[\frac{\sqrt{2}}{\sqrt{2}-1} + \frac{\sqrt{8\ln 2}}{3-2\sqrt{2}}\right]\sqrt{T(1 + P_T/\varepsilon)},
\end{aligned}$$

where $H = 0$ for the surrogate loss in (17), and the third inequality follows from Lemma 8 and Jensen's inequality. Choosing $\varepsilon = R$ gives

$$\text{D-Reg}_T(\mathbf{u}_{1:T}) \le 4\sqrt{2}\left[\frac{\sqrt{2}}{\sqrt{2}-1} + \frac{\sqrt{8\ln 2}}{3-2\sqrt{2}}\right]G\sqrt{T(R^2 + RP_T)}.$$

A.4 Proof of Lemma 6

For a sequence $\mathbf{V}_{1:T}$, let $\mathcal{S} = \mathcal{S}(\mathbf{V}_{1:T}) = \{\mathcal{I}_1, \dots, \mathcal{I}_{|\mathcal{S}|}\}$ be its partition set, so that $|\mathcal{S}| - 1 = J(\mathbf{V}_{1:T})$. We define $\text{SW-Reg}_T^{\ell}(\mathcal{S})$ as in (3), with f_t, $\mathbf{x}_t$, and $\mathcal{X}$ replaced by ℓ_t, $\mathbf{y}_t$, and $\mathcal{Y}$, respectively. The same definition applies to other loss sequences and time horizons. We apply the tree bound of Yang et al. (2026, Lemma 5) to the surrogate losses in (14), and then consider $|\mathcal{S}| = 1$ and $|\mathcal{S}| > 1$ separately.

The surrogate has Hessian λI and gradient norms at most $9G$ on $\mathcal{Y}$, whose diameter is at most $8G/\lambda$. To handle arbitrary T, run IRESET-OGD with internal horizon $T^+ = 2^{\lceil \log_2 T \rceil} < 2T$. For the analysis, choose $\widehat{\mathbf{y}} \in \arg\min_{\mathbf{y}\in\mathcal{Y}} \sum_{t\in\mathcal{I}_{|\mathcal{S}|}} \ell_t(\mathbf{y})$ and extend the loss sequence by

$$\ell_t(\mathbf{y}) = \frac{\lambda}{2}\|\mathbf{y} - \widehat{\mathbf{y}}\|_2^2, \qquad T < t \le T^+.$$

These additional losses are λ-strongly convex with gradient norms at most $8G \le 9G$, so the source bound also applies to the extended sequence. Extend only the final interval to T^+ and denote the resulting partition by $\mathcal{S}^+$. Since $\mathbf{V}_{1:T}$ is constant on each interval of $\mathcal{S}$, its cumulative loss on that interval is no smaller than that of the best fixed comparator. Thus, $\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le \text{SW-Reg}_T^{\ell}(\mathcal{S})$.

We next compare the switching regret before and after extending the horizon. The added losses are nonnegative and vanish at $\widehat{\mathbf{y}}$, a minimizer of the original cumulative loss on the final interval. Hence, extending this interval leaves its minimum cumulative loss unchanged. The other intervals are unaffected, and the learner's cumulative loss can only increase, since its predictions on $[T]$ remain unchanged. Consequently,

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le \text{SW-Reg}_T^{\ell}(\mathcal{S}) \le \text{SW-Reg}_{T^+}^{\ell}(\mathcal{S}^+).$$

On each dyadic interval $\mathcal{I}$, the OGD expert with predictions $\mathbf{w}_t$ and step sizes $1/[\lambda(t - \min \mathcal{I} + 1)]$ satisfies

$$\sum_{t\in\mathcal{I}} \ell_t(\mathbf{w}_t) - \min_{\mathbf{y}\in\mathcal{Y}} \sum_{t\in\mathcal{I}} \ell_t(\mathbf{y}) \le \frac{81G^2}{\lambda}(1 + \ln|\mathcal{I}|).$$

The initial term covers singleton intervals. Let $\mathcal{F}$ be the collection of maximal dyadic intervals contained in the intervals of $\mathcal{S}^+$. Applying the global tree bound of Yang et al. (2026) gives

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le \text{SW-Reg}_{T^+}^{\ell}(\mathcal{S}^+) \le \sum\nolimits_{\mathcal{I}\in\mathcal{F}} \left[2\Xi + \frac{81G^2}{\lambda}(1 + \ln|\mathcal{I}|)\right]. \tag{32}$$

Here the source horizon factor at T^+ is at most Γ_T, since $T^+ < 2T$. Substituting the gradient bound $9G$ and diameter bound $8G/\lambda$ into the source meta-regret constant gives exactly Ξ in (21).

If $|\mathcal{S}| = 1$, the only maximal dyadic interval is the root $[T^+]$. Using $T^+ < 2T$ in (32), we obtain

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le 2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T)).$$

If $|\mathcal{S}| > 1$, each segment $\mathcal{I} \in \mathcal{S}^+$ contains at most two maximal dyadic intervals at each scale, hence at most $2(1 + \lfloor\log_2|\mathcal{I}|\rfloor)$ in total. Since every extended segment has length less than $2T$, $|\mathcal{F}| \le 2|\mathcal{S}|(2 + \log_2 T)$. Applying (32) and $|\mathcal{S}| \le 2(|\mathcal{S}| - 1)$ gives

$$\begin{aligned}\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) &\le 2|\mathcal{S}|(2 + \log_2 T)\left[2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T))\right] \\ &\le 4(|\mathcal{S}| - 1)(2 + \log_2 T)\left[2\Xi + \frac{81G^2}{\lambda}(1 + \ln(2T))\right].\end{aligned}$$

Combining the two cases and using $|\mathcal{S}| - 1 = J(\mathbf{V}_{1:T})$ proves (20), which finishes the proof.

A.5 Proof of Lemma 7

For a sequence $\mathbf{V}_{1:T}$, let $\mathcal{S} = \mathcal{S}(\mathbf{V}_{1:T})$ be its partition set, so $|\mathcal{S}| - 1 = J(\mathbf{V}_{1:T})$. We bound the switching regret using maximal dyadic intervals, treating $|\mathcal{S}| = 1$ and $|\mathcal{S}| > 1$ separately.

For the surrogate loss in (16), $\operatorname{diam}(\mathcal{Y}) \le 1/(8\beta G)$ gives

$$\begin{aligned}|\beta\langle\mathbf{h}_t, \mathbf{y} - \mathbf{y}_t\rangle| &\le \frac{1}{8}, \\ \|\nabla\ell_t(\mathbf{y})\|_2 = |1 + \beta\langle\mathbf{h}_t, \mathbf{y} - \mathbf{y}_t\rangle|\,\|\mathbf{h}_t\|_2 &\le \frac{9}{8}G \le \sqrt{2}G.\end{aligned}$$

Together with $\nabla^2\ell_t(\mathbf{y}) = \beta\mathbf{h}_t\mathbf{h}_t^\top$, this gives $\nabla^2\ell_t(\mathbf{y}) \succeq (\beta/4)\nabla\ell_t(\mathbf{y})\nabla\ell_t(\mathbf{y})^\top$, so ℓ_t is $\beta/4$-exp-concave. The exp-concavity inequality of Yang et al. (2026, Lemma 1) therefore gives

$$\ell_t(\mathbf{v}) \ge \ell_t(\mathbf{u}) + \langle\nabla\ell_t(\mathbf{u}), \mathbf{v} - \mathbf{u}\rangle + \frac{\beta}{16}\langle\nabla\ell_t(\mathbf{u}), \mathbf{v} - \mathbf{u}\rangle^2, \qquad \mathbf{u}, \mathbf{v} \in \mathcal{Y},$$

since $\min\{\beta/8, \beta/\sqrt{2}\} = \beta/8$. Using this coefficient, the gradient bound $\sqrt{2}G$, and the diameter bound $1/(8\beta G)$ in the meta-regret calculation of Yang et al. (2026) gives Ξ_{exp} in (27).

Initialize each ONS expert at a point in $\mathcal{Y}$ with $A_0 = (64G)^2 I$ and step size $8/\beta$. On an interval $\mathcal{I}$, let $\mathbf{w}_t$ denote the expert's prediction at round $t \in \mathcal{I}$. The ONS analysis (Hazan et al., 2007), with the initialization term retained, gives, for every $\mathbf{u} \in \mathcal{Y}$,

$$\sum\nolimits_{t\in\mathcal{I}}\big[\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})\big] \le \frac{4}{\beta}\left[d\ln\left(1 + \frac{|\mathcal{I}|}{2048}\right) + 1\right] \le \frac{5d}{\beta}\left(4 + \frac{\sqrt{2}}{8}\right)(1 + \ln|\mathcal{I}|). \tag{33}$$

The last inequality follows from $d \ge 1$ and $\ln(1 + |\mathcal{I}|/2048) \le \ln|\mathcal{I}| + 1/2048$, and includes singleton intervals.

Run IRESET-ONS with internal horizon $T^+ = 2^{\lceil\log_2 T\rceil} < 2T$. For the analysis, append zero losses, obtained from (16) by setting $\mathbf{h}_t = \mathbf{0}$ for $t > T$, and extend only the final segment of $\mathcal{S}$ to obtain

$\mathcal{S}^+$. This continuation preserves the predictions on $[T]$, the cumulative losses, and the number of segments, so

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le \text{SW-Reg}_T^{\ell}(\mathcal{S}) = \text{SW-Reg}_{T^+}^{\ell}(\mathcal{S}^+).$$

The source horizon factor at T^+ is bounded by Γ_T because $T^+ < 2T$.

Let $\mathcal{F}$ be the partition into maximal dyadic intervals contained in the segments of $\mathcal{S}^+$, using the complete binary tree on $[T^+]$. Applying the recursion of Yang et al. (2026, Lemmas 8–9) to $\mathcal{S}^+$, together with (33), yields

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le \sum\nolimits_{\mathcal{I}\in\mathcal{F}} \left[2\Xi_{\exp} + \frac{5d}{\beta}\left(4 + \frac{\sqrt{2}}{8}\right)(1 + \ln|\mathcal{I}|)\right]. \tag{34}$$

If $|\mathcal{S}| = 1$, the only maximal dyadic interval is the root $[T^+]$. Using $T^+ < 2T$ in (34), we obtain

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le 2\Xi_{\exp} + \frac{5d}{\beta}\left(4 + \frac{\sqrt{2}}{8}\right)(1 + \ln(2T)).$$

If $|\mathcal{S}| > 1$, each segment $\mathcal{I} \in \mathcal{S}^+$ contains at most two maximal dyadic intervals at each scale, hence at most $2(1 + \lfloor \log_2 |\mathcal{I}| \rfloor)$ in total. Since every extended segment has length less than $2T$, $|\mathcal{F}| \le 2|\mathcal{S}|(2 + \log_2 T)$. Applying (34) and $|\mathcal{S}| \le 2(|\mathcal{S}| - 1)$ gives

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le 4(|\mathcal{S}| - 1)(2 + \log_2 T)\left[2\Xi_{\exp} + \frac{5d}{\beta}\left(4 + \frac{\sqrt{2}}{8}\right)(1 + \ln(2T))\right].$$

Combining the two cases and using $|\mathcal{S}| - 1 = J(\mathbf{V}_{1:T})$ proves (26).

A.6 Proof of Lemma 8

For a sequence $\mathbf{V}_{1:T}$, let $\mathcal{S} = \mathcal{S}(\mathbf{V}_{1:T})$ be its partition set, so that $|\mathcal{S}| - 1 = J(\mathbf{V}_{1:T})$. We apply Pasteris et al. (2024, Theorem 2.2) to the normalized surrogate losses in (17). For the surrogate in (17), the normalized losses are

$$\widetilde{\ell}_t(\mathbf{y}) = \frac{\langle \mathbf{h}_t, \mathbf{y}\rangle + 2R\|\mathbf{h}_t\|_2}{4GR} \in [0,1], \qquad \mathbf{y} \in \mathcal{Y}.$$

Their gradient norms are at most $1/(4R)$, and normalization divides regret by $4GR$. For an OGD expert on an interval $\mathcal{I}$, use step size $\eta_{\mathcal{I}} = 16R^2/\sqrt{|\mathcal{I}|}$ and projection onto $\mathcal{Y}$. Let $\mathbf{w}_t$ denote this expert's prediction at round $t \in \mathcal{I}$. Since $\operatorname{diam}(\mathcal{Y}) \le 4R$, starting from any point in $\mathcal{Y}$ gives

$$\sum\nolimits_{t\in\mathcal{I}} \widetilde{\ell}_t(\mathbf{w}_t) - \min_{\mathbf{u}\in\mathcal{Y}} \sum\nolimits_{t\in\mathcal{I}} \widetilde{\ell}_t(\mathbf{u}) \le \frac{16R^2}{2\eta_{\mathcal{I}}} + \frac{\eta_{\mathcal{I}}|\mathcal{I}|}{32R^2} = \sqrt{|\mathcal{I}|}.$$

Thus the source theorem applies with expert regret constant $\gamma = 1$, including $|\mathcal{I}| = 1$.

For an arbitrary horizon T, use the internal horizon $T^+ = 2^{\lceil \log_2 T\rceil} < 2T$ and append zero losses after round T. Extend only the final interval of $\mathcal{S}$ to T^+ and denote the resulting partition by $\mathcal{S}^+$, so that $|\mathcal{S}^+| = |\mathcal{S}|$ and $\sum_{\mathcal{I}\in\mathcal{S}^+} |\mathcal{I}| = T^+$. The continuation leaves all predictions on $[T]$ unchanged, and zero losses do not change the learner's cumulative loss or any interval's minimum cumulative loss. Applying Theorem 2.2 on $[T^+]$ therefore gives

$$\text{SW-Reg}_T^{\ell}(\mathbf{V}_{1:T}) \le 4GR\,\text{SW-Reg}_{T^+}^{\widetilde{\ell}}(\mathcal{S}^+) \le 4GR\left[\frac{\sqrt{2}}{\sqrt{2}-1} + \frac{\sqrt{8\ln 2}}{3-2\sqrt{2}}\right]\sum\nolimits_{\mathcal{I}\in\mathcal{S}^+}\sqrt{|\mathcal{I}|}. \tag{35}$$

Finally, Cauchy–Schwarz and $T^+ < 2T$ yield

$$\sum\nolimits_{\mathcal{I}\in\mathcal{S}^+}\sqrt{|\mathcal{I}|} \le \sqrt{|\mathcal{S}|\sum\nolimits_{\mathcal{I}\in\mathcal{S}^+}|\mathcal{I}|} \le \sqrt{2T\big(1 + J(\mathbf{V}_{1:T})\big)}.$$

Substituting this inequality into (35) proves (31), which finishes the proof.