

Optimal Dynamic Regret in Exp-Concave Online Learning

Dheeraj Baby

Yu-Xiang Wang

Department of Computer Science, UC Santa Barbara

DHEERAJ@UCSB.EDU

YUXIANGW@CS.UCSB.EDU

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

We consider the problem of the [Zinkevich \(2003\)](#)-style dynamic regret minimization in online learning with *exp-concave* losses. We show that whenever improper learning is allowed, a Strongly Adaptive online learner achieves the dynamic regret of $\tilde{O}^*(n^{1/3}C_n^{2/3} \vee 1)$ where C_n is the *total variation* (a.k.a. *path length*) of the an arbitrary sequence of comparators that may not be known to the learner ahead of time. Achieving this rate was highly nontrivial even for square losses in 1D where the best known upper bound was $O(\sqrt{nC_n} \vee \log n)$ ([Yuan and Lamperski, 2019](#)). Our new proof techniques make elegant use of the intricate structures of the primal and dual variables imposed by the KKT conditions and could be of independent interest. Finally, we apply our results to the classical statistical problem of *locally adaptive non-parametric regression* ([Mammen, 1991](#); [Donoho and Johnstone, 1998](#)) and obtain a stronger and more flexible algorithm that do not require any statistical assumptions or any hyperparameter tuning.

Keywords: Non-stationary Online Learning, Dynamic Regret, Strongly Adaptive methods, Online Non-parametric Regression

1. Introduction

We consider a generic online learning framework which is modelled as an interactive n step game between a learner and adversary. At each time step t , the learner predicts a $\mathbf{p}_t \in \mathcal{D} \subseteq \mathbb{R}^d$. Then the adversary reveals a loss function $f_t : \mathbb{R}^d \rightarrow \mathbb{R}$. The objective of the learner is to minimise its regret against a predefined set of strategies \mathcal{W} that is known to the learner before the start of the game. We call a learning algorithm to be *proper* when $\mathcal{D} = \mathcal{W}$. Further when $\mathcal{D} = \mathcal{W}$ are convex sets and the losses f_t are convex in \mathcal{D} , the generic learning framework reduces to the one studied in Online Convex Optimization (OCO) ([Hazan, 2016](#)). On the other hand, we call the learning algorithm to be *improper* when $\mathcal{D} \supset \mathcal{W}$. A commonly used metric to measure the performance of the learner is its *static* regret defined as

$$R_n = \sum_{t=1}^n f_t(\mathbf{p}_t) - \inf_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^n f_t(\mathbf{w}).$$

A sub-linear static regret implies that the average loss incurred by the learner converges to that of the best comparator strategy in hindsight.

A canonical example of an improper algorithm can be found in an online linear regression setting where $f_t(\mathbf{u}) = (y_t - \mathbf{x}_t^T \mathbf{u})^2$ with $|y_t| \leq 1$, $\|\mathbf{x}_t\|_2 \leq 1$ and we are interested in controlling the static regret against a set of linear predictors with bounded norm, $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^d : \|\mathbf{w}\|_2 \leq 1\}$. One popular learning algorithm in this framework is the Vovk-Azoury-Warmuth (VAW) forecaster ([Vovk, 1997](#); [Azoury and Warmuth, 2004](#)) (or see Section 11.8 in ([Cesa-Bianchi and Lugosi, 2006](#))).

The VAW forecaster attains an $O(d \log n)$ static regret against \mathcal{W} . However predictions of VAW at time t denoted by \mathbf{u}_t may not necessarily satisfy $\|\mathbf{u}_t\|_2 \leq 1$ hence making it an improper algorithm.

The notion of static regret is not befitting for non-stationary environments – such as financial markets – where it could be inappropriate to compete against a fixed comparator due to the changes in the dynamics of the environment. The work of (Zinkevich, 2003) introduces the notion of *dynamic* regret defined as

$$R_{\mathbf{w}_1, \dots, \mathbf{w}_n}^n := \sum_{t=1}^n f_t(\mathbf{p}_t) - f_t(\mathbf{w}_t), \quad (1)$$

for *any* sequence of comparators \mathbf{w}_t in \mathcal{W} . The dynamic regret bounds are usually expressed in literature as a function of number of time steps and some path variation metric that captures the degree of non-stationarity in the comparator sequence. In this paper, we study the following path variation:

$$TV(\mathbf{w}_1, \dots, \mathbf{w}_n) := \sum_{t=2}^n \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1.$$

The maximum dynamic regret against all comparator sequences whose path variation is bounded by a number C_n can then be defined as

$$R_n(C_n) := \sup_{\substack{\mathbf{w}_1, \dots, \mathbf{w}_n \\ TV(\mathbf{w}_1, \dots, \mathbf{w}_n) \leq C_n}} R_{\mathbf{w}_1, \dots, \mathbf{w}_n}^n.$$

There is a complementary body of work on Strongly Adaptive (SA) algorithms (Daniely et al., 2015) where the static regret in *any* sub-interval of $[n] := \{1, \dots, n\}$ is controlled (see Section 2 for a review). Hence SA algorithms have the nice property of being globally and locally optimal. The work of (Zhang et al., 2018b) exploits this property of SA algorithms to control their dynamic regret in terms of a variational metric that measures how much the losses f_t change over time. In particular, whenever the losses have extra curvature properties such as strong convexity or exp-concavity, they show that one can get fast dynamic regret rates. However, it was unclear if SA methods can lead to optimal dynamic regret guarantees in terms of the path length of the comparator sequence — an open question raised in (Zhang et al., 2018b).

The works of (Zhang et al., 2018a) and (Yuan and Lamperski, 2019) attains a dynamic regret of $O^*(\sqrt{n(1+C_n)})$ and $O^*(\sqrt{nC_n} \vee \log n)$ respectively, where $O^*(\cdot)$ hides dependence on the dimension and $(a \vee b) = \max\{a, b\}$. However, we show a lower bound of $\Omega^*(n^{1/3}C_n^{2/3} \vee \log n)$ in Proposition 11 applicable to the case when losses are strongly convex / exp-concave. Hence, there is a large gap between this lower bound and existing upper bounds. In this work, we show that whenever improper learning is allowed and when the loss functions are strongly convex / exp-concave, one can leverage SA algorithms to attain the *sharp* rate of $\tilde{O}^*(n^{1/3}C_n^{2/3} \vee \log n)$ for $R_n(C_n)$ where $\tilde{O}^*(\cdot)$ hides dependence in the dimension and factors of $\log n$ (see section 4 for formal statements and complete list of assumptions). Further, the SA algorithms need not require the apriori knowledge of C_n to attain this rate.

As a concrete use case, we show that our results have interesting implications to the problem of *online* Total Variation (TV) denoising. The offline version of TV-denoising problem has seen many influential applications in the signal processing community. For example, algorithms that use TV-regularization has been deployed in every cellphone, digital camera and medical imaging devices (we refer readers to the book (Chambolle et al., 2010) and the references therein) as well as other

tasks beyond the context of images such as change-point detection, semisupervised learning and graph partitioning.

We proceed to formally introduce the non-parametric regression problem behind TV-denoising. Define a non-parametric class of TV bounded sequences as

$$\mathcal{TV}(C_n) := \left\{ (w_1, \dots, w_n) : \sum_{t=2}^n |w_t - w_{t-1}| \leq C_n \right\},$$

where $\sum_{t=2}^n |w_t - w_{t-1}|$ is termed as the TV of the sequence $w_{1:n} := (w_1, \dots, w_n)$. In the offline TV-denoising problem we are given n observations of the form $y_t = w_t + \epsilon_t$ where ϵ_t are iid zero mean subgaussian noise, $t \in [n]$ and $w_{1:n}$ is an unknown sequence in $\mathcal{TV}(C_n)$. We are interested in coming up with estimates \hat{w}_t such that $R^{\mathcal{TV}}(C_n) := \mathbb{E} [\sum_{t=1}^n (\hat{w}_t - w_t)^2]$ is controlled. Several non-parametric regression algorithms such as Trend Filtering (Tibshirani, 2014) are known to achieve a near minimax optimal rate of $\tilde{O}(n^{1/3} C_n^{2/3})$ for $R^{\mathcal{TV}}(C_n)$ where $\tilde{O}(\cdot)$ hides dependence on factors of $\log n$.

We can instantiate an online version of the above non-parametric regression problem behind TV-denoising into our learning framework with slight modifications. We consider a TV class with bounded sequences

$$\mathcal{TV}^B(C_n) := \left\{ w_{1:n} : \sum_{t=2}^n |w_t - w_{t-1}| \leq C_n, |w_t| \leq B \forall t \in [n] \right\}. \quad (2)$$

When viewed through our online learning framework, we take $f_t(x) = (y_t - x)^2$ where $|y_t| \leq B$, $\mathcal{D} = \mathcal{W} = [-B, B]$. Labels $y_{1:n}$ is a fixed sequence in contrast to the stochastic noise setting discussed earlier, and we are hoping to compete with the best approximation from sequences in $\mathcal{TV}^B(C_n)$ for all $C_n \geq 0$ at the same time. We remark that to compete with the entire $\mathcal{TV}(C_n)$ class it is sufficient to compete with $\mathcal{TV}^B(C_n)$ due to the property $|y_t| \leq B$. We show in Section 3 that by using appropriate SA algorithms, one can attain a dynamic regret of $R_n(C_n) = \tilde{O}(n^{1/3} C_n^{2/3})$. This in turn implies the minimax estimation rate in the iid stochastic setting (see Appendix A for details). Further our results have the added advantage of providing an oracle inequality. We conclude this section by summarizing our key contributions below.

- We show that Follow-the-Leading-History (FLH) algorithm (Hazan and Seshadhri, 2007) with Follow The Leader (FTL) as base learners can achieve the optimal minimax regret (modulo $\log n$ factors) of $\tilde{O}(n^{1/3} C_n^{2/3} B^{4/3} \vee B^2 \log n)$ for the problem of online non-parametric regression with TV bounded sequences – $\mathcal{TV}^B(C_n)$ – as the reference class. The policy is *adaptive* to the TV budget C_n . Further, we demonstrate that the same policy is minimax optimal for smoother non-parametric sequence classes such as Sobolev class or Holder class.
- When improper learning is allowed and when the loss functions revealed by the adversary are exp-concave, strongly smooth and Lipschitz on a box that encloses the set of comparators \mathcal{W} , (see Section 4) we show that FLH with ONS as base learners attains a dynamic regret of $\tilde{O}\left(d^{3.5} (n^{1/3} C_n^{2/3} \vee 1)\right)$ when $C_n \geq 1/n$ and $O(d^{1.5} \log n)$ otherwise, without prior knowledge of C_n – the path variation of the *comparator sequence*. This rate is shown to be minimax optimal modulo polynomial factors of $\log n$ and d .

- The proof of the regret bound is facilitated by exploiting a number of distinct structures of primal and dual variables in KKT conditions of the optimization problem solved by the offline oracle. We believe that this style of analysis can be useful in bounding the regret of online algorithms in a broader context.

2. Related Work

We begin by recalling works that are most relevant to our setting. We reserve the term *OCO setting* when $\mathcal{W} = \mathcal{D}$ and loss functions are convex in \mathcal{D} .

For an arbitrary comparator sequence in \mathcal{W} denoted by $\mathbf{w}_{1:n} := (\mathbf{w}_1, \dots, \mathbf{w}_n)$, (Zinkevich, 2003) introduces a path variational defined as

$$P_n(\mathbf{w}_1, \dots, \mathbf{w}_n) = \sum_{t=1}^n \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_2. \quad (3)$$

They show that in the OCO setting, the Online Gradient Descent (OGD) algorithm can attain a dynamic regret (Eq.(1)) of $O(\sqrt{n}(1 + P_n))$, but if P_n is known¹, $O(\sqrt{n}(1 + P_n))$ can be achieved by simply increasing the learning rate appropriately. By hedging over a collection of OGD algorithms defined by exponential grid of step sizes, (Zhang et al., 2018a) proposes an algorithm that achieves a faster rate of $O(\sqrt{n}(1 + P_n))$ which is shown to be minimax optimal when the loss functions are convex. (Yuan and Lamperski, 2019) proposes strategies that can attain regret rates of $O(\sqrt{nP_n} \vee \log n)$ and $O(\sqrt{dnP_n} \vee d \log n)$ for strongly convex and exp-concave losses respectively. However, this regret rate is only optimal when P_n approaches n or $P_n = O(1/n)$.

(Besbes et al., 2015) introduces the functional variation defined as

$$D_n := \sum_{t=2}^n \max_{\mathbf{w} \in \mathcal{W}} |f_t(\mathbf{w}) - f_{t-1}(\mathbf{w})|. \quad (4)$$

They show that by using a restarted variant of OGD, one can attain the dynamic regret rate of $O(n^{2/3}D_n^{1/3})$ and $\tilde{O}(\sqrt{nD_n})$ for convex and strongly convex losses respectively using noisy-gradient feedback. This setting is incompatible to ours as it exploits *smoothness* in f_1, \dots, f_n while we allow f_1, \dots, f_n to be arbitrary. Moreover, they need to know D_n .

There is a parallel line of work (Hazan and Seshadhri, 2007; Daniely et al., 2015; Adamskiy et al., 2016) that focuses on controlling the static regret in any sub-interval of $[n]$. In particular, (Daniely et al., 2015) proposes the notion of Strongly Adaptive algorithms. An algorithm is said to be *Strongly Adaptive* (SA) if for every continuous interval $I \subseteq [n]$, the static regret incurred by the algorithm is $O(\text{poly}(\log n)R^*(|I|))$ where $R^*(|I|)$ is the value of minimax static regret incurred in an interval of length $|I|$. In this viewpoint, the algorithms proposed by (Hazan and Seshadhri, 2007) for strongly convex / exp-concave losses are in fact Strongly Adaptive.

(Zhang et al., 2018b) shows that SA methods enjoys a dynamic regret of $\tilde{O}(n^{2/3}D_n^{1/3})$ for convex functions and $\tilde{O}(\sqrt{nD_n})$ and $\tilde{O}(\sqrt{dnD_n})$ for strongly convex and exp-concave losses respectively without prior knowledge of D_n . We refer the reader to Appendix A for a discussion on various other dynamic regret minimization strategies such as (Jadbabaie et al., 2015; Yang et al., 2016; Mokhtari et al., 2016; Chen et al., 2018; Zhao et al., 2020).

1. In a sense that we are to only compete with sequences with path length $\leq P_n$, rather than *simultaneously* competing with all sequences $P_n > 0$.

The setting of learning with squared error losses we consider in Section 3 can be regarded as an online version of the batch Total Variation denoising problem. The corresponding offline problem has been studied extensively in the non-parametric regression literature. Many algorithms such as Wavelet Smoothing (Donoho and Johnstone, 1998), Locally Adaptive Regression Splines (van de Geer, 1990) and Trend Filtering (Kim et al., 2009; Tibshirani, 2014; Wang et al., 2014, 2016; Guntuboyina et al., 2017) have been shown to achieve the optimal minimax rates of $\tilde{O}(n^{1/3}C_n^{2/3})$ under squared error loss where n is the number of samples and C_n is the TV of the ground truth. All of these estimators have a key property of *local adaptivity* where the estimators are able to detect abrupt local fluctuations in the ground truth signal and adjust the amount of smoothing to be applied which is essential for optimally estimating TV bounded sequences that can exhibit spatially in-homogeneous degree of smoothness.

(Baby and Wang, 2019; Baby et al., 2021) studies the problem of estimating TV bounded sequences in an online stochastic optimization framework. They assume that the labels revealed by the adversary is the noisy realization of a ground truth sequence that belongs to a $TV(C_n)$ ball. However, the absence of such statistical assumptions on revealed labels in our setting makes the problem significantly more challenging. Interestingly, a lower bound from (Baby and Wang, 2019) implies that the meta-hedge algorithm of (Zhang et al., 2018a) requires $\Omega(\sqrt{nP_n})$ dynamic regret even if the loss functions are strongly convex, despite the fact that OGD achieves $O(\log n)$ static regret. Extension to higher order TV classes are considered in (Baby and Wang, 2020).

We refer the reader to Appendix A for an elaborate description on how our TV-denoising framework fits under the umbrella of online non-parametric regression framework developed by (Rakhlin and Sridharan, 2014) and others (Gaillard and Gerchinovitz, 2015; Koolen et al., 2015; Kotłowski et al., 2016).

3. Performance guarantees for squared error losses

In this section, we focus on the online TV-denoising problem which is a special case of our online learning framework with squared error losses as discussed in Section 1. This will help to build the intuitions behind the analysis for general exp-concave losses as well. All unspecified proofs of this section are deferred to Appendix C. We consider the following interaction protocol.

- At time $t \in [n]$ learner predicts $x_t \in \mathcal{D} = [-B, B]$.
- Adversary reveals a label $y_t \in [-B, B]$.
- Learner suffers loss $(y_t - x_t)^2$.

We define the comparator class as the set of TV bounded sequences that takes values in $\mathcal{W} = [-B, B]$ as in Eq.(2). The performance of the learner is measured using dynamic regret against the sequences that belongs to $\mathcal{TV}^B(C_n)$, for all $C_n > 0$ simultaneously.

The main SA method that we will be relying on throughout this paper is the FLH algorithm from (Hazan and Seshadhri, 2007). We provide a description of this algorithm in Appendix B for completeness. We have the following regret guarantee for FLH with Follow-the-Leader (FTL) as base learners (in this case, FTL is equivalent to simple online averaging).

Theorem 1 *Let x_t be the prediction at time t of FLH with learning rate $\zeta = 1/(8B^2)$ and base learners as FTL. Then for any comparator $(w_1, \dots, w_n) \in \mathcal{TV}^B(C_n)$*

$$\sum_{t=1}^n (y_t - x_t)^2 - (y_t - w_t)^2 = \tilde{O}\left(n^{1/3} C_n^{2/3} B^{4/3} \vee B^2\right),$$

where the labels obey $|y_t| \leq B$, $\tilde{O}(\cdot)$ hides dependence on logarithmic factors of horizon n and $a \vee b := \max\{a, b\}$.

Remark 2 (Adaptivity to C_n and (non-stochastic) oracle inequality) *We remark that FLH-FTL does not require C_n as an input thus Theorem 1 implies the following oracle inequality*

$$\sum_{t=1}^n (y_t - x_t)^2 \leq \min_{w_1, \dots, w_n} \sum_{t=1}^n (y_t - w_t)^2 + \tilde{O}\left(n^{1/3} \text{TV}(w_{1:n})^{2/3} B^{4/3} \vee B^2\right).$$

Such result is not known for any algorithm even in the offline case when y_1, \dots, y_n is known. Notice that w_t does not need to be constrained because $-B \leq y_t \leq B$.

The strongest oracle inequality for TV-denoising to our knowledge is that of (Guntuboyina et al., 2017; Ortelli and van de Geer, 2019), which shows that the fused-lasso estimator with tuning parameter λ obeys $\sum_{t=1}^n (y_t - x_t)^2 \leq \min_{w_1, \dots, w_n} \sum_{t=1}^n (y_t - w_t)^2 + O(\lambda \text{TV}(w_{1:n}))$, under additional stochastic assumptions of y_t . Our results eliminate the need to choose hyperparameter λ all together and achieve the same rate achievable by the optimal choice of λ .

For the sake of clarity we next present the strategy we adopt for proving Theorem 1. We also highlight the main technical challenges that are needed to be overcome along the way. This is followed by some useful lemmas and proof of the main theorem in Section 3.2.

3.1. Proof strategy for Theorem 1

Let u_1, \dots, u_n be the *offline optimal* sequence (see Lemma 3) in $\mathcal{TV}^B(C_n)$ which attains the minimum cumulative squared error loss. Note that this offline optimal can depend on the entire sequence of labels y_1, \dots, y_n chosen by the adversary.

Consider a partitioning of $[n]$ into M sub-intervals $\{[i_s, i_t]\}_{i=1}^M$. We will also use the number i to refer to the interval $[i_s, i_t]$. For the interval i , define the quantities: $n_i = i_t - i_s + 1$, $\bar{y}_i = \frac{1}{n_i} \sum_{j=i_s}^{i_t} y_j$, $\bar{u}_i = \frac{1}{n_i} \sum_{j=i_s}^{i_t} u_j$.

We start by the following regret decomposition.

$$R_n = \sum_{i=1}^M \underbrace{\sum_{j=i_s}^{i_t} (x_j - y_j)^2 - (y_j - \bar{y}_i)^2}_{T_{1,i}} + \sum_{i=1}^M \underbrace{\sum_{j=i_s}^{i_t} (y_j - \bar{y}_i)^2 - (y_j - \bar{u}_i)^2}_{T_{2,i}} + \sum_{i=1}^M \underbrace{\sum_{j=i_s}^{i_t} (y_j - \bar{u}_i)^2 - (y_j - u_j)^2}_{T_{3,i}} \quad (5)$$

Now the task of bounding R_n reduces to bounding $T_{1,i}, T_{2,i}, T_{3,i}$ for each bin and adding them up across all M bins. Let C_i be the TV within bin i incurred by the offline optimal. In Lemma 5, we exhibit a partitioning \mathcal{P} of $[n]$ into $M = O(n^{1/3} C_n^{2/3} B^{-2/3})$ bins such that $C_i \leq B/\sqrt{n_i}$ for each bin.

Due to strong adaptivity of FLH, the term $T_{1,i} = O(B^2 \log n)$ since it is the static regret against the fixed comparator \bar{y}_i . Hence adding them across all bins in the partition \mathcal{P} yields $\sum_{i=1}^M T_{1,i} = \tilde{O}(n^{1/3} C_n^{2/3} B^{4/3})$.

By exploiting the KKT conditions satisfied by the offline optimal and using strong smoothness, we show in Lemma 9 that $T_{3,i}$ can be at-most $O(n_i C_i^2 + \lambda C_i)$ in general. Here $\lambda \geq 0$ is the optimal dual variable arising from the KKT conditions (Lemma 3). Since $C_i = O(B/\sqrt{n_i})$ for bins in the partition \mathcal{P} , we have $n_i C_i^2 = O(B^2)$. However, it is not possible to bound $\lambda C_i = O(1)$ since λ can be even $\Theta(n)$ in some cases (See Example 21 in Appendix C).

This is where the term $T_{2,i}$ plays a crucial role. Note that since \bar{y}_i is the minimizer of $g(x) = \sum_{j=i_s}^{i_t} (y_j - x)^2$, we conclude that $T_{2,i} \leq 0$. For simplicity of exposition, let's assume that $T_{2,i} < 0$, deferring formal arguments for the general case to Section 3.2. We show that this negative term diminishes the λC_i arising from the bound on $T_{3,i}$ to a quantity that is $O(1)$. Specifically, $T_{2,i} + T_{3,i} = O(B^2)$ even though individually $|T_{2,i}|, |T_{3,i}|$ can be *very large*. The desired regret bound now follows by summing it across all $M = O(n^{1/3} C_n^{2/3} B^{-2/3})$ bins in \mathcal{P} .

3.2. Regret Analysis

Define the sign function as $\text{sign}(x) = 1$ if $x > 0$; -1 if $x < 0$; and some $u \in [-1, 1]$ if $x = 0$. For a vector $\mathbf{x} \in \mathbb{R}^d$, $\text{sign}(\mathbf{x}) \in \mathbb{R}^d$ is defined by the coordinate-wise application of this rule. We start by presenting a sequence of useful lemmas.

Lemma 3 (characterization of offline optimal) Consider the following convex optimization problem (where $\tilde{z}_1, \dots, \tilde{z}_{n-1}$ are introduced as dummy variables)

$$\begin{aligned} \min_{\tilde{u}_1, \dots, \tilde{u}_n, \tilde{z}_1, \dots, \tilde{z}_{n-1}} \quad & \frac{1}{2} \sum_{t=1}^n (y_t - \tilde{u}_t)^2 \\ \text{s.t.} \quad & \tilde{z}_t = \tilde{u}_{t+1} - \tilde{u}_t \quad \forall t \in [n-1], \\ & \sum_{t=1}^{n-1} |\tilde{z}_t| \leq C_n \end{aligned} \tag{6a}$$

Let $u_1, \dots, u_n, z_1, \dots, z_{n-1}$ be the optimal primal variables and let $\lambda \geq 0$ be the optimal dual variable corresponding to the last constraint (6a). By the KKT conditions, we have

- **stationarity:** $y_t = u_t - \lambda(s_t - s_{t-1})$, where $s_t \in \partial|z_t|$ (a subgradient). Specifically, $s_t = \text{sign}(u_{t+1} - u_t)$ if $|u_{t+1} - u_t| > 0$ and s_t is some value in $[-1, 1]$ otherwise. For convenience of notations later, we also define $s_n = s_0 = 0$.
- **complementary slackness:** $\lambda (\sum_{t=2}^n |u_t - u_{t-1}| - C_n) = 0$.

Remark 4 We enumerate some elementary observations about the optimal primal variables in Lemma 3 that will be used throughout.

- P1 For any time point t , if the optimal solution $u_{t+1} > u_t$, then $s_t = 1$. Similarly $s_t = -1$ whenever $u_{t+1} < u_t$. If $u_t = u_{t+1}$, the s_t can be any number in $[-1, 1]$.
- P2 Consider a sub-interval $[a, b]$ with $2 \leq a \leq n-1$ such that the optimal solution jumps at both the end points. i.e $u_k \neq u_{k-1}$ for $k \in \{b+1, a\}$. Define $\Delta s_{a \rightarrow b} := s_b - s_{a-1}$. Then either $|\Delta s_{a \rightarrow b}| = 0$ or $|\Delta s_{a \rightarrow b}| = 2$ since $s_{a-1} \in \{-1, 1\}$ and $s_b \in \{-1, 1\}$.

P3 Consider a sub-interval $[1, b]$ with $b < n$ such that $u_{b+1} \neq u_b$. Then $|\Delta s_{1 \rightarrow b}| = 1$ since $s_0 = 0$ by convention (Lemma 3). Similarly for a sub-interval $[a, n]$ with $a > 1$, such that $u_{a-1} \neq u_a$, we have $|\Delta s_{a \rightarrow n}| = 1$.

Terminology. We will refer to the optimal primal variables u_1, \dots, u_n in Lemma 3 as the *offline optimal sequence* in this section.

Next, we exhibit a useful partitioning scheme of the interval $[n]$.

Lemma 5 (key partition) Initialize $\mathcal{Q} \leftarrow \Phi$. Starting from time 1, spawn a new bin $[i_s, i_t]$ whenever $\sum_{j=i_s+1}^{i_t+1} |u_j - u_{j-1}| > B/\sqrt{n_i}$, where $n_i = i_t - i_s + 2$. Add the spawned bin $[i_s, i_t]$ to \mathcal{Q} . Consider the following post processing routine.

1. Initialize $\mathcal{P} \leftarrow \Phi$.
2. For $i \in [|\mathcal{Q}|]$:
 - if $u_{i_t} = u_{i_t+1}$:
 - (a) Let p be the largest time point with $u_{p:i_t}$ being constant and let q be the smallest time point with $u_{i_t+1:q}$ being constant.
 - (b) Add bin $[i_s, p-1]$ to \mathcal{P} .
 - (c) If $(i+1)_t > q$ then add $[p, q]$ to \mathcal{P} and set $(i+1)_s \leftarrow q+1$.
 - (d) Goto Step 2.
 - Add $[i_s, i_t]$ to \mathcal{P} . Goto Step 2.

Let $M := |\mathcal{P}|$. We have $M = O\left(1 \vee n^{1/3} C_n^{2/3} B^{-2/3}\right)$. Further for any bin $[i_s, i_t] \in \mathcal{P}$, it holds that $\sum_{j=i_s+1}^{i_t} |u_j - u_{j-1}| \leq B/\sqrt{n_i}$ where $n_i = i_t - i_s + 1$.

Remark 6 Consider a bin $[i_s, i_t] \in \mathcal{P}$. Let $\Delta s_i := s_{i_t} - s_{i_s-1}$. By virtue of the post processing routine of Lemma 5, the bin $[i_s, i_t]$ will conform to either of the cases P2 or P3 in Remark 4. So we have $|\Delta s_i| > 0$ implies $|\Delta s_i| \geq 1$.

We emphasize that the bins $[i_s, i_t]$ we consider in Eq. (5) belong to the partition \mathcal{P} of Lemma 5. We proceed to bound $T_{1,i}$, $T_{2,i}$ and $T_{3,i}$ in the regret decomposition of Eq.(5).

Lemma 7 (bounding $T_{1,i}$) Assume that we run FLH with the settings described in Theorem 1. For any bin i we have $T_{1,i} = O(B^2 \log n)$

Lemma 8 (bounding $T_{2,i}$) Define $C_i := \sum_{j=i_s+1}^{i_t} |u_j - u_{j-1}|$, the TV within bin i incurred by the offline optimal solution. Let $\Delta s_i := s_{i_t} - s_{i_s-1}$ and $n_i := i_t - i_s + 1$. We have $T_{2,i} \leq \frac{-\lambda^2 (\Delta s_i)^2}{n_i}$.

Lemma 9 (bounding $T_{3,i}$) Let C_i and Δs_i be as in Lemma 8.

Case(a) If $|\Delta s_i| > 0$ then $T_{3,i} \leq B^2 + 6\lambda C_i$.

Case(b) If $\Delta s_i = 0$ with $s_{i_{s-1}} = s_{i_t} = 1$ and the offline optimal \mathbf{u} is non-decreasing within bin i , then $T_{3,i} \leq B^2$.

Case(c) If $\Delta s_i = 0$ with $s_{i_{s-1}} = s_{i_t} = -1$ and the offline optimal \mathbf{u} is non-increasing within bin i , then $T_{3,i} \leq B^2$.

Proof of Theorem 1 Tree diagrams that represent the flow of arguments in the proof is displayed in Fig.3 and 4 in Appendix C. We start from the regret decomposition in Eq. (5).

Case (a) in Lemma 9. First we handle case(a) in Lemma 9 where $|\Delta s_i| > 0$. Define $T_i := T_{1,i} + T_{2,i} + T_{3,i}$. From Lemmas 7, 8 and 9 we have

$$\begin{aligned} T_i &\leq O(B^2 \log n) - \frac{\lambda^2(\Delta s_i)^2}{n_i} + B^2 + 6\lambda C_i \\ &\leq O(B^2 \log n) - \frac{\lambda^2(\Delta s_i)^2}{n_i} + 6\lambda C_i \\ &\stackrel{(a)}{\leq} O(B^2 \log n) + \frac{9n_i C_i^2}{(\Delta s_i)^2} - \left(\frac{\lambda \Delta s_i}{\sqrt{n_i}} - \frac{3C_i \sqrt{n_i}}{\Delta s_i} \right)^2 \\ &\stackrel{(b)}{\leq} O(B^2 \log n) + 9B^2 \\ &\leq O(B^2 \log n), \end{aligned}$$

where line (a) is obtained by completing the square. For line (b) we dropped the negative term used Remark 6 to conclude $|\Delta s_i| \geq 1$. Further $n_i C_i^2 \leq B^2$ for bins in the partition \mathcal{P} of Lemma 5.

Case (b) and (c) in Lemma 9. To handle case (b) and case (c) in Lemma 9 where $\Delta s_i = 0$ and monotonic, we have $T_{1,i} = O(B^2 \log n)$ due to Lemma 7, $T_{2,i} \leq 0$ due to Lemma 8 and $T_{3,i} \leq B^2$ due to Lemma 9. So $T_i \leq O(B^2 \log n) + B^2 \leq O(B^2 \log n)$.

Other cases:

(A1) Consider the case when $\Delta s_i = 0$ with $s_{i_{s-1}} = s_{i_t} = -1$ and the offline optimal \mathbf{u} is non-decreasing within bin i . If the sequence is constant within the bin, then trivially we have $T_i = O(B^2 \log n)$ due to Strongly Adaptivity of FLH. Otherwise, we split the original bin into two sub-bins $[i_s, k]$ and $[k+1, i_t]$ such that $s_k = 1$ with $u_{k+1} > u_k$. See config (a) in Fig.1 for an illustration. Then the two sub-bins falls into the category of case (a) in Lemma 9. By bounding the regret within each sub-bin separately by following the previous arguments for case (a) and adding them up, we can get $T_i \leq O(B^2 \log n)$ regret for the original bin. The arguments for the case when $\Delta s_i = 0$ with $s_{i_{s-1}} = s_{i_t} = 1$ and the offline optimal \mathbf{u} is non-increasing within bin i are similar.

(A2) To handle the case when $\Delta s_i = 0$ and the optimal sequence is *not* monotonic, we split the bin into two parts. Consider the case $s_{i_t} = s_{i_{s-1}} = 1$. We can split $\mathbf{u}_{i_s:i_t}$ as $\mathbf{u}_{i_s:k}$ and $\mathbf{u}_{k+1:i_t}$ such that the sequence $\mathbf{u}_{i_s:k}$ is non-decreasing and $s_k = -1$ with $u_k > u_{k+1}$. See config (b) in Fig.1 for an illustration. Notice that both the sub-bins $\mathbf{u}_{i_s:k}$ and $\mathbf{u}_{k+1:i_t}$ now falls into the category of case(a) in Lemma 9. Adding the bounds within these sub-bins by following the treatment for case (a) above yields $T_i \leq O(B^2 \log n)$. The arguments for the scenario $s_{i_t} = s_{i_{s-1}} = -1$ are similar.

Now the theorem follows by summing $\sum_{i=1}^M T_i$ for the $M = O(1 \vee n^{1/3} C_n^{2/3} B^{-2/3})$ bins in the partition \mathcal{P} of Lemma 5. \blacksquare

The previous results generalize to online TV-denoising framework in higher dimensions.




Figure 1: Examples of configurations referred in the proof of Theorem 1. The blue dots corresponds to the offline optimal sequence.

Proposition 10 (Extension to higher dimensions) Consider a protocol where at each time the learner predicts a vector $\mathbf{x}_t \in \mathbb{R}^d$ after which the adversary reveals \mathbf{y}_t such that $\|\mathbf{y}_t\|_\infty \leq B$. Consider a comparator sequence of vectors $\mathbf{w}_1, \dots, \mathbf{w}_n$ such that $TV(\mathbf{w}_{1:n}) := \sum_{t=2}^n \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1 \leq C_n$. By running d instances of FLH with learning rate $\zeta = 1/(8B^2)$ and FTL as base learners, where instance $i, i \in [d]$, predicts $\mathbf{x}_t[i]$ at time t , we have

$$R_n(\mathbf{w}_{1:n}) := \sum_{j=1}^n \|\mathbf{y}_t - \mathbf{x}_t\|_2^2 - \|\mathbf{y}_t - \mathbf{w}_t\|_2^2 = \tilde{O} \left(dB^2 \log n \vee d^{1/3} n^{1/3} C_n^{2/3} B^{4/3} \right).$$

Proposition 11 (Lower bound) Assume the protocol and notations of Proposition 10. For any algorithm, we have

$$\sup_{\mathbf{w}_{1:n}: TV(\mathbf{w}_{1:n}) \leq C_n} R_n(\mathbf{w}_{1:n}) = \Omega \left(dB^2 \log n \vee d^{1/3} n^{1/3} C_n^{2/3} B^{4/3} \right).$$

By comparing the upper and lower bounds, we conclude that the FLH-FTL strategy in Proposition 10 is minimax optimal (modulo log factors) wrt all parameters d, n, B and C_n .

Remark 12 Several other non-parametric sequence classes such as the Holder ball $\mathcal{H}^B(B'_n) = \{w_{1:n} : \|Dw_{1:n}\|_\infty \leq B'_n, \|w_{1:n}\|_\infty \leq B\}$ and Sobolev ball $\mathcal{S}^B(C'_n) = \{w_{1:n} : \|Dw_{1:n}\|_2 \leq C'_n, \|w_{1:n}\|_\infty \leq B\}$ can be shown to be embedded inside a $\mathcal{TV}^B(C_n)$ ball for appropriate choices of C_n, B_n and B'_n (see (Baby and Wang, 2019)) with all classes having the same minimax rates of estimation in the iid setting. So the minimax optimality on TV ball for FLH with FTL as base learners implies minimax optimality on the embedded Holder and Sobolev balls as well.

4. Performance guarantees for exp-concave losses

We begin by listing all the assumptions we make about the loss functions.

EC-1 Without loss of generality, we assume $\mathbf{0} \in \mathcal{W}$. Let $B := \sup_{\mathbf{x} \in \mathcal{W}} \|\mathbf{x}\|_\infty$. Define $\mathcal{D}^- := \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_\infty \leq B\}$. The loss functions $f_t(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ are G Lipschitz in \mathcal{D}^- .

- EC-2 The loss functions are β strongly smooth in $\mathcal{D} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_\infty \leq B + G\}$. i.e $f_t(\mathbf{y}) \leq f_t(\mathbf{x}) + (\mathbf{y} - \mathbf{x})^T \nabla f_t(\mathbf{x}) + \frac{\beta}{2} \|\mathbf{x} - \mathbf{y}\|_2^2$, for all $\mathbf{x}, \mathbf{y} \in \mathcal{D}$. We assume without loss of generality that $\beta \geq 1$.
- EC-3 The loss functions are α exp-concave in \mathcal{D} . i.e $f_t(\mathbf{y}) \geq f_t(\mathbf{x}) + (\mathbf{y} - \mathbf{x})^T \nabla f_t(\mathbf{x}) + \frac{\alpha}{2} ((\mathbf{y} - \mathbf{x})^T \nabla f_t(\mathbf{x}))^2$ for all $\mathbf{x}, \mathbf{y} \in \mathcal{D}$.
- EC-4 The loss functions $f_t(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ are G^\dagger Lipschitz in \mathcal{D} .

Below, we give an example of a family of loss functions that satisfy the above assumptions.

Example 13 (Generalized linear models) Let $f_t(\mathbf{x}) = g(\mathbf{v}_t^T \mathbf{x})$, where $g : \mathbb{R} \rightarrow \mathbb{R}$ is a convex function and \mathbf{v}_t is a feature vector. Let $\|\mathbf{v}_t\|_2 \leq R$. Assume that for all $\mathbf{x} \in \mathcal{D}^-$ we have $|g'_t(\mathbf{v}_t^T \mathbf{x})| \leq a$. Further for all $\mathbf{x} \in \mathcal{D}$, let $|g'_t(\mathbf{v}_t^T \mathbf{x})| \leq a^+$, $g''_t(\mathbf{v}_t^T \mathbf{x}) \leq b$, $g''_t(\mathbf{v}_t^T \mathbf{x}) \geq c > 0$. Then Assumptions EC 1-5 are satisfied by the losses f_t with $G = aR$, $\beta = bR^2$, $\alpha = c/((a^+)^2)$ and $G^\dagger = Ra^+$.

We are interested in characterizing the maximum dynamic regret

$$R_n^+(C_n) := \sup_{\substack{\mathbf{w}_1, \dots, \mathbf{w}_n \in \mathcal{D}^- \\ \sum_{t=2}^n \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1 \leq C_n}} \sum_{t=1}^n f_t(\mathbf{x}_t) - f_t(\mathbf{w}_t),$$

where \mathbf{x}_t are the predictions of the learner. Since $\mathcal{W} \subseteq \mathcal{D}^-$, the dynamic regret against comparators in \mathcal{D}^- trivially upperbounds the dynamic regret against \mathcal{W} . The algorithms that we study throughout this section are improper in the sense that the predictions of the algorithms belong to $\mathcal{D} \supset \mathcal{W}$.

Before diving into the details, we remark that our main focus is to get optimal dependence on n and C_n . The dimension d is considered as a constant problem parameter and we do not try to optimize its polynomial dependence. All unspecified proofs of this section are given in Appendix D.

We have the following regret guarantee for exp-concave losses.

Theorem 14 By using the base learner as ONS with parameter $\zeta = \min \left\{ \frac{1}{4G^\dagger(2B\sqrt{d}+2G/\beta)}, \alpha \right\}$, decision set \mathcal{D} and choosing learning rate $\eta = \alpha$, FLH obeys $R_n^+(C_n) = \tilde{O} \left(d^{3.5} (n^{1/3} C_n^{2/3} \vee 1) \right)$ if $C_n > 1/n$ and $O(d^{1.5} \log n)$ otherwise. Here $a \vee b := \max\{a, b\}$ and $\tilde{O}(\cdot)$ hides dependence on the constants B, G, G^\dagger, α and factors of $\log n$.

Proof [proof sketch] Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be the offline optimal sequence such that $\sum_{t=1}^n f_t(\mathbf{u}_t)$ is minimum across all sequences that obeys: (a) $\sum_{t=2}^n \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_1 \leq C_n$; (b) $\mathbf{u}_t \in \mathcal{D}^-$ for all $t \in [n]$ (see Lemma 29 in Appendix D for more details).

Let \mathcal{P} be a partition of $[n]$ into $M = O^*(n^{1/3} C_n^{2/3})$ bins obtained by a similar scheme in Lemma 5 where within each bin, we have $\sum_{j=i_s+1}^{i_t} \|\mathbf{u}_j - \mathbf{u}_{j-1}\|_1 \leq B/\sqrt{n_i}$. Let $[i_s, i_t]$ denote the i^{th} bin in \mathcal{P} and let n_i be its length. Define $\bar{\mathbf{u}}_i = \frac{1}{n_i} \sum_{j=i_s}^{i_t} \mathbf{u}_j$ and $\dot{\mathbf{u}}_i = \bar{\mathbf{u}}_i - \frac{1}{n_i \beta} \sum_{j=i_s}^{i_t} \nabla f_j(\bar{\mathbf{u}}_i)$ where β is as in Assumption EC-2. Let \mathbf{x}_j be the prediction made by FLH at time j . We start with following regret decomposition.

$$R_n^+(C_n) \leq \underbrace{\sum_{i=1}^M \sum_{j=i_s}^{i_t} f_j(\mathbf{x}_j) - f_j(\dot{\mathbf{u}}_i)}_{T_{1,i}} + \underbrace{\sum_{i=1}^M \sum_{j=i_s}^{i_t} f_j(\dot{\mathbf{u}}_i) - f_j(\bar{\mathbf{u}}_i)}_{T_{2,i}} + \underbrace{\sum_{i=1}^M \sum_{j=i_s}^{i_t} f_j(\bar{\mathbf{u}}_i) - f_j(\mathbf{u}_j)}_{T_{3,i}} \quad (7)$$

Unlike the squared error case, for the term $T_{1,i}$, we *do not* compete with the minimizer of $g(\mathbf{x}) := \sum_{j=i_s}^{i_t} f_j(\mathbf{x})$. Instead we compete with $\hat{\mathbf{u}}_i$ which is obtained by a one-step gradient descent of $g(\mathbf{x})$ from the point $\bar{\mathbf{u}}_i$ where the step size is set as $1/(n_i\beta)$.

Recall that the purpose of \bar{y}_i in Eq. (5) was to make $T_{2,i}$ non-positive thereby facilitating potential cancellation of terms arising from the bound on $T_{3,i}$. Since $g(\mathbf{x})$ is $n_i\beta$ strongly smooth, by the well known *descent lemma* in first order optimization (eg. see Eq. 3.5 in (Bubeck, 2015)), we can bound $T_{2,i}$ in Eq. (7) with a “sufficiently negative” term $-\frac{1}{2n_i\beta}\|\nabla g(\bar{\mathbf{u}}_i)\|^2$ as well. Also, observe that

$$\|\hat{\mathbf{u}}\|_\infty \leq \|\bar{\mathbf{u}}_i\|_\infty + \frac{\sum_{j=i_s}^{i_t} \|\nabla f_j(\bar{\mathbf{u}})\|_\infty}{n_i\beta} \leq B + G,$$

where in the last line we used the fact $\bar{\mathbf{u}}_i \in \mathcal{D}^-$ and the Lipschitzness assumption in EC-1 along with $\beta > 1$ by assumption EC-2. So in $T_{1,i}$ the comparator term $\hat{\mathbf{u}}_i \in \mathcal{D}$. The base learners of the FLH produce predictions in \mathcal{D} to compete with such a comparator hence making the overall algorithm improper. We do not project $\hat{\mathbf{u}}$ to the set \mathcal{W} , because doing so appears to make $T_{2,i}$ not negative enough to adequately diminish the terms arising from $T_{3,i}$.

Rest of the proof proceeds by introducing lemmas analogous to the squared error case, carefully bounding $T_{1,i} + T_{2,i} + T_{3,i}$ for each bin in \mathcal{P} and summing them up across all bins. However, we remark that the analysis is significantly more involved in comparison to that of squared error case due to dual variables introduced by the additional constraint that $\mathbf{u}_t \in \mathcal{D}^-$.

We first present the proof for the 1D-exp-concave case in Appendix D.1, which illustrates how boundedness constraints are handled by the structures in the KKT-conditions (Lemma 23) and by discussing various combinations (see Fig. 6-8). Then we present the full proof for the higher-dimensional exp-concave losses in Appendix D.2, where the structure becomes too complex for us to enumerate all combinations. We address this by constructing an iterative algorithm that generates bins and prove that the algorithm is guaranteed to find a partition with cardinality $O^*(n^{1/3}C_n^{2/3})$ that satisfies a number of additional properties that give rise to the regret bound we claim. ■

Proposition 15 *For strongly convex losses, the regret bound can be improved to $\tilde{O}\left(d^2(n^{1/3}C_n^{2/3} \vee 1)\right)$ if $C_n > 1/n$ and $O(\log n)$ otherwise by using OGD as base learners in the FLH procedure. See Appendix D.2 for a proof.*

By comparing with the lower bound in Proposition 11 we conclude that the dynamic regret bound of Theorem 14 is minimax optimal (up to $\log n$ factors) in n and C_n .

Remark 16 (Implications in statistical methodology.) *Example 13 and Theorem 14 extends the locally-adaptive nonparametric regression theory that are typically studied for square loss to an arbitrary strongly convex / exp-concave loss while allowing covariates (exogenous variables) to be modeled. Moreover, the method enjoys strong oracle inequalities (e.g. Remark 2) that certifies the predictive performance in a fully agnostic / model-misspecified setting with no stochastic assumptions. In addition, the method does not introduce additional tuning parameters at all.*

5. Conclusion and further discussions

In this paper, we considered the problem of dynamic regret minimization with exp-concave losses and showed that SA methods are minimax optimal (modulo factors of $\log n$ and d) in a setting where

improper learning is allowed. To the best of our knowledge this is the first work that attains optimal dynamic regret rates under this setting. The resulting algorithms are adaptive to the path variation of the comparator sequence. Further, our results have far reaching consequences in locally adaptive non-parametric regression as mentioned in Remark 16.

An open problem to investigate is if SA methods can still perform optimally in a proper learning setting. If we consider a very restrictive setup where the loss functions are exp-concave and for each function, at-least one of the global optimal points lie in the comparator set \mathcal{W} , it is indeed the case. An example of this scenario is the squared error loss $f_t(x) = (y_t - x)^2$ with $|y_t| \leq B$ and $\mathcal{W} = [-B, B]$ as in the TV-denoising setup. On the other hand, if there exists an SA learner that can guarantee $O(\log n)$ static regret against any point in \mathbb{R}^d in any time interval, then our results provides optimal proper learning when $\mathcal{D} = \mathcal{W} = \mathbb{R}^d$.

Acknowledgments

The research was partially supported by NSF Award #2029626, #2007117 and a start-up grant from UCSB CS department.

References

- Dmitry Adamskiy, Wouter M. Koolen, Alexey Chernov, and Vladimir Vovk. A closer look at adaptive regret. *Journal of Machine Learning Research*, 2016.
- Katy S. Azoury and Manfred K. Warmuth. Relative loss bounds for on-line density estimation with the exponential family of distributions. *Machine Learning*, 43:211–246, 2004.
- Dheeraj Baby and Yu-Xiang Wang. Online forecasting of total-variation-bounded sequences. In *Neural Information Processing Systems (NeurIPS)*, 2019.
- Dheeraj Baby and Yu-Xiang Wang. Adaptive online estimation of piecewise polynomial trends. *Neural Information Processing Systems (NeurIPS)*, 2020.
- Dheeraj Baby, Xuandong Zhao, and Yu-Xiang Wang. An optimal reduction of tv-denoising to adaptive online learning. *AISTATS*, 2021.
- Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, 8:231–357, 2015.
- Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006. ISBN 0521841089.
- Antonin Chambolle, Vicent Caselles, Daniel Cremers, Matteo Novaga, and Thomas Pock. An introduction to total variation for image analysis. *Theoretical foundations and numerical methods for sparse recovery*, 9(263-340):227, 2010.
- Xi Chen, Yining Wang, and Yu-Xiang Wang. Non-stationary stochastic optimization under lp, q-variation measures. 2018.

- Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. In *International Conference on Machine Learning*, pages 1405–1411, 2015.
- David L Donoho and Iain M Johnstone. Minimax estimation via wavelet shrinkage. *The annals of Statistics*, 26(3):879–921, 1998.
- Pierre Gaillard and Sébastien Gerchinovitz. A chaining algorithm for online nonparametric regression. In *Conference on Learning Theory*, pages 764–796, 2015.
- Adityanand Guntuboyina, Donovan Lieu, Sabyasachi Chatterjee, and Bodhisattva Sen. Adaptive risk bounds in univariate total variation denoising and trend filtering. 2017.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Elad Hazan and Comandur Seshadhri. Adaptive algorithms for online decision problems. In *Electronic colloquium on computational complexity (ECCC)*, volume 14, 2007.
- Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. 2007.
- Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *Artificial Intelligence and Statistics*, pages 398–406, 2015.
- Seung-Jean Kim, Kwangmoo Koh, Stephen Boyd, and Dmitry Gorinevsky. ℓ_1 trend filtering. *SIAM Review*, 51(2):339–360, 2009.
- Wouter M Koolen, Alan Malek, Peter L Bartlett, and Yasin Abbasi. Minimax time series prediction. In *Advances in Neural Information Processing Systems (NIPS’15)*, pages 2557–2565. 2015.
- Wojciech Kotłowski, Wouter M. Koolen, and Alan Malek. Online isotonic regression. In *Annual Conference on Learning Theory (COLT-16)*, volume 49, pages 1165–1189. PMLR, 2016.
- Enno Mammen. Nonparametric regression under qualitative smoothness assumptions. *Annals of Statistics*, 19(2):741–759, 1991.
- Aryan Mokhtari, Shahin Shahrampour, A. Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 7195–7201, 2016.
- Francesco Orтели and Sara van de Geer. Prediction bounds for (higher order) total variation regularized least squares. *arXiv preprint arXiv:1904.10871*, 2019.
- Alexander Rakhlin and Karthik Sridharan. Online non-parametric regression. In *Conference on Learning Theory*, pages 1232–1264, 2014.
- Ryan J Tibshirani. Adaptive piecewise polynomial estimation via trend filtering. *Annals of Statistics*, 42(1):285–323, 2014.
- Sara van de Geer. Estimating a regression function. *Annals of Statistics*, 18(2):907–924, 1990.

- V. Vovk. Competitive on-line linear regression. In *NIPS*, 1997.
- Volodya Vovk. Competitive on-line statistics. *International Statistical Review*, 69(2):213–248, 2001.
- Yu-Xiang Wang, Alex Smola, and Ryan Tibshirani. The falling factorial basis and its statistical applications. In *International Conference on Machine Learning (ICML-14)*, pages 730–738, 2014.
- Yu-Xiang Wang, James Sharpnack, Alex Smola, and Ryan J Tibshirani. Trend filtering on graphs. *Journal of Machine Learning Research*, 17(105):1–41, 2016.
- 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 *International Conference on Machine Learning (ICML-16)*, pages 449–457, 2016.
- Jianjun Yuan and Andrew Lamperski. Trading-off static and dynamic regret in online least-squares and beyond. 2019.
- Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems (NeurIPS-18)*, pages 1323–1333, 2018a.
- Lijun Zhang, Tianbao Yang, Zhi-Hua Zhou, et al. Dynamic regret of strongly adaptive methods. In *International Conference on Machine Learning (ICML-18)*, pages 5877–5886, 2018b.
- Peng Zhao, Y. Zhang, L. Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. *NeurIPS*, 2020.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning (ICML-03)*, pages 928–936, 2003.

Appendix A. More on Related Work

Throughout this section, we refer to the variationals P_n in Eq.(3) and D_n in Eq.(4) where the arguments are dropped for brevity. In the OCO setting, when the environment is benign, (Zhao et al., 2020) replaces the \sqrt{n} dependence in the regret of $O(\sqrt{n(1 + P_n)})$ attained by (Zhang et al., 2018b) with problem dependent quantities that could be much smaller than \sqrt{n} . Although the linear smoother lower bound in Proposition 2 of (Baby and Wang, 2019) would imply that an OEGD (Online Extra Gradient Descent) (Zhao et al., 2020) expert with any learning rate sequences require $\Omega(\sqrt{nP_n})$ dynamic regret for the 1D-TV-denoising problem.

Interestingly in (Yuan and Lamperski, 2019) the authors mention that even in the one-dimensional setting, a lower bound on dynamic regret for strongly convex / exp-concave losses that holds uniformly for the entire range $0 \leq P_n \leq n$ is unknown. However, we find that one can combine the existing lower bounds on univariate TV-denoising in a stochastic setting (Donoho and Johnstone, 1998) with the lower bound construction of (Vovk, 2001) (or see Theorem 11.9 in (Cesa-Bianchi and Lugosi, 2006)) for online learning with squared error losses to obtain an $\Omega(\log n \vee n^{1/3}C_n^{2/3})$ in one dimensions (see Appendix C for details). In this work, we show that SA methods can achieve a regret that matches this lower bound (modulo polynomial factors of dimension and $\log n$) when losses are strongly convex / exp-concave.

When the loss functions are strongly convex, (Mokhtari et al., 2016) studies the dynamic regret against the comparator points that are the unique minimizers of the revealed losses in the set \mathcal{W} ($= \mathcal{D}$). Specifically when $\mathbf{w}_t^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{W}} f_t(\mathbf{x})$, and $C_n^* := \sum_{t=2}^n \|\mathbf{w}_t^* - \mathbf{w}_{t-1}^*\|_2$, they show that OGD can be used to get the rate of $O(1 + C_n^*)$ for the dynamic regret against the sequence $\mathbf{w}_{1:n}^*$. However, as noted in (Zhang et al., 2018a), that even though this implies an $O(1 + C_n^*)$ bound on the dynamic regret against arbitrary comparator sequences in Eq.(1), the resulting bound can be overly pessimistic. As an example, in TV-denoising, $f_t(x) = (y_t - x)^2$ where $y_t = w_t + \text{Noise}$. Even if w_t obeys that $\text{TV}(w_{1:n}) = O(1)$, we would still have $E[C_n^*] = \sum_{t=2}^n E[|y_t - y_{t-1}|] \geq \Omega(n)$, thus the $O(1 + C_n^*)$ regret bound does not imply any non-trivial bounds in our setting, e.g., if we take the comparator sequence to be w_1, \dots, w_n .

Different variational measures capture different aspects of the online learning problem and are not comparable in general. (Jadbabaie et al., 2015) introduces a policy that attains dynamic regret in terms of D_n and P_n simultaneously. Various other interesting variational measures and strategies to control dynamic regret can be found in the works of (Yang et al., 2016; Chen et al., 2018).

The seminal work of (Hazan and Seshadhri, 2007) introduces the notion of weakly adaptive regret which is defined as the maximum static regret incurred by the learning algorithm in any continuous interval. They propose algorithms that obtain static regret guarantees of $\tilde{O}(\sqrt{n})$ for convex losses and $\tilde{O}(1)$ and $\tilde{O}(d)$ for strongly convex and exp-concave losses respectively. This has been further developed in (Adamskiy et al., 2016). However, one drawback of weakly adaptive regret stems from its trivial regret guarantees on short intervals. For example, an $\tilde{O}(\sqrt{n})$ static regret guarantee on an interval of length \sqrt{n} is meaningless. This drawback is overcome by the notion of Strongly Adaptive regret as discussed in Section 2 by taking into account the length of the interval where the static regret is computed.

(Zhang et al., 2018b) shows that SA methods enjoys a dynamic regret of $\tilde{O}(n^{2/3}D_n^{1/3})$ for convex functions and $\tilde{O}(\sqrt{nD_n})$ and $\tilde{O}(\sqrt{dnD_n})$ for strongly convex and exp-concave losses respectively in an OCO setting. Thus when combined with our results, we can conclude that SA methods are

simultaneously optimal wrt to the dynamic regret based on D_n and P_n in an application that allows improper learning.

Our online TV-denoising setting studied in Section 3 can be cast into the framework of (Rakhlin and Sridharan, 2014). They study the regret against non-parametric function classes under squared error loss as follows

$$R'_n = \sum_{t=1}^n (\hat{y}_t - y_t)^2 - \inf_{f \in \mathcal{F}} \sum_{t=1}^n (f(x_t) - y_t)^2.$$

Our TV-denoising setting becomes identical to (Rakhlin and Sridharan, 2014) if one takes the comparator class \mathcal{F} to be space of TV bounded functions and when the features x_t are revealed in an isotonic order: $x_1 \leq \dots \leq x_n$. All the results can be trivially extended to the case of arbitrary covariates that may be non-isotonically revealed by maintaining online averages across all intervals in a Geometric Cover on $[n]$ and using the specialist aggregation scheme in (Adamskiy et al., 2016) (see (Baby et al., 2021) for an illustration of this idea in a stochastic setting). We do not follow this path for the sake of simplicity of exposition.

The results of (Rakhlin and Sridharan, 2014) establish the minimax rates for the quantity R'_n when \mathcal{F} is taken to be a Besov ball. It is known that a TV ball is sandwiched between two Besov spaces (see for eg. (Donoho and Johnstone, 1998)) that have the same minimax rate for R'_n . Hence results of (Rakhlin and Sridharan, 2014) establishes that minimax regret of our problem is $\tilde{O}(n^{1/3})$. However their bounds don't capture the correct dependence on C_n and are obtained by non-constructive arguments. In contrary we obtain upper bounds with optimal dependence on both n and C_n by an efficient algorithm.

(Kotłowski et al., 2016) proposes a policy that achieves a rate of $\tilde{O}(n^{1/3})$ for R'_n when \mathcal{F} is the family of isotonic functions that take values in $[0, 1]$. This class is indeed a subset of $\mathcal{TV}^B(1)$. They exploit the property that the optimal isotonic function is piecewise constant and within a constant section, it takes the value equal to mean of labels y_t within that section. However for our case the offline problem solved by the oracle is an instance of a constrained fused LASSO which doesn't yield such nice closed form expression for value of optimal function within a constant section.

(Gaillard and Gerchinovitz, 2015) proposes a novel chaining algorithm that achieves optimal rate for R'_n when \mathcal{F} is the family of Holder smooth functions. The functions residing in this class are spatially homogeneous and more regular than the TV class. We show that our policy is also optimal for regret against Holder ball embedded within a \mathcal{TV}^B space (see Remark 12). Interestingly the generic forecaster they proposed can be shown to yield the optimal $\tilde{O}(n^{1/3})$ rate for our problem. However, the run-time of that policy is exponential.

Our setting is closely related to the setup studied in (Koolen et al., 2015). Their setting can be viewed as competing against Sobolev sequences which are more regular than TV bounded sequences. We show that our policy is also optimal for regret against Sobolev ball embedded within a TV bounded space (see Remark 12).

We now proceed to explain how the model agnostic regret guarantees presented in this paper imply minimax statistical estimation rate in a stochastic setting. When applied to squared error losses, the FLH-FTL procedure can yield

$$\sum_{t=1}^n (y_t - x_t)^2 - (y_t - w_t)^2 = \tilde{O}(n^{1/3} C_n^{2/3}), \quad (8)$$

where x_t are the predictions of FLH-FTL procedure and $w_{1:n} \in \mathcal{TV}^B(C_n)$ (see Eq.(2)). We demonstrate how this implies minimax estimation in an iid setting which is the usual subject of study in non-parametric regression among the statistics community. In the stochastic setting we have the following observation model:

$$y_t = w_t + \epsilon_t, t \in [n]$$

for some fixed $w_{1:n} \in \mathcal{TV}^B(C_n)$ and ϵ_t are iid zero mean subgaussian noise with magnitude at-most B (We could relax the boundedness to be obeyed with high probability). We have,

$$\begin{aligned} \sum_{t=1}^n E[(y_t - x_t)^2] - E[(y_t - w_t)^2] &= \sum_{t=1}^n E[(x_t - w_t)^2] - 2E[\epsilon_t(x_t - w_t)] + E[\epsilon_t^2] - E[\epsilon_t^2] \\ &=_{(a)} \sum_{t=1}^n E[(x_t - w_t)^2] - 2E[\epsilon_t]E[(x_t - w_t)] \\ &= \sum_{t=1}^n E[(x_t - w_t)^2] \\ &=_{(b)} \tilde{O}(n^{1/3}C_n^{2/3}), \end{aligned}$$

where line (a) is due to the fact that x_t and ϵ_t are mutually independent and line (b) is due to Eq.(8). From (Baby and Wang, 2019), this is indeed the minimax rate of estimating $w_{1:n}$ under the stochastic setting.

Thus we conclude that the model agnostic regret guarantees presented in this paper implies minimax estimation rate in a stochastic setting and hence the former is strictly stronger.

Appendix B. Preliminaries

In this section, we recall the Follow-the-Leading-History (FLH) algorithm from (Hazan and Seshadhri, 2007) along with some basic definitions.

Definition 17 (Strong convexity) Loss functions f_t are said to be H strongly convex in the domain \mathcal{D} if it satisfies

$$f_t(\mathbf{y}) \geq f_t(\mathbf{x}) + (\mathbf{y} - \mathbf{x})^T \nabla f_t(\mathbf{x}) + \frac{H}{2} \|\mathbf{x} - \mathbf{y}\|^2,$$

for all $\mathbf{x}, \mathbf{y} \in \mathcal{D}$.

FLH enjoys the following guarantee against any base learner.

Proposition 18 (Hazan and Seshadhri, 2007) Suppose the loss functions are exp-concave with parameter α . For any interval $I = [r, s]$ in time, the algorithm FLH Fig.2 with learning rate $\zeta = \alpha$ gives $O(\alpha^{-1}(\log r + \log |I|))$ regret against the base learner in hindsight.

Definition 19 ((Daniely et al., 2015)) An algorithm is said to be Strongly Adaptive (SA) if for every contiguous interval $I \subseteq [n]$, the static regret incurred by the algorithm is $O(\text{poly}(\log n)\Gamma^*(|I|))$ where $\Gamma^*(|I|)$ is the value of minimax static regret incurred in an interval of length $|I|$.

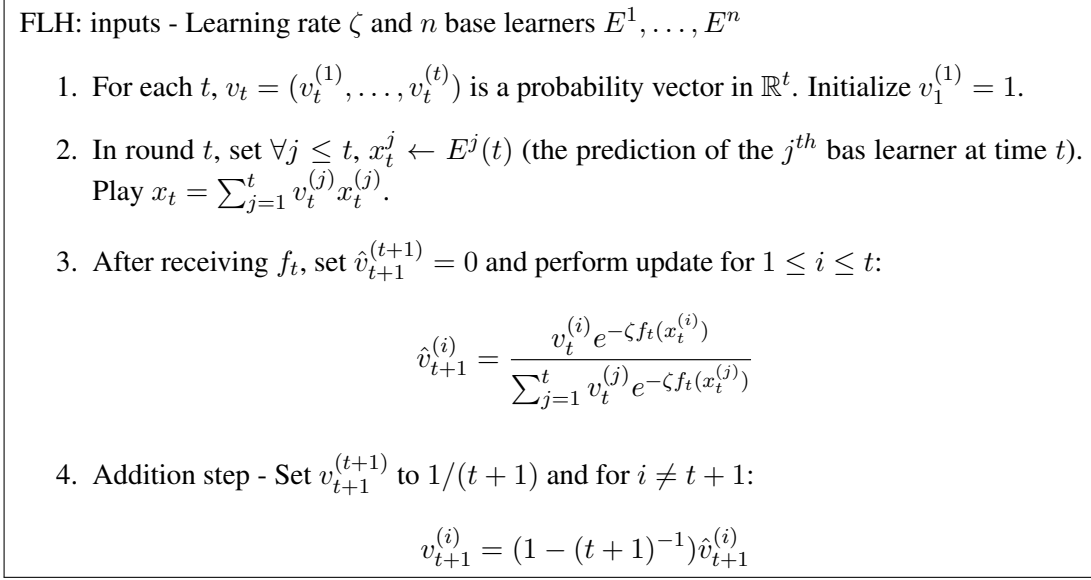


Figure 2: FLH algorithm

It is known from (Hazan et al., 2007) that OGD and ONS achieves static regret of $O(\log n)$ and $O(d \log n)$ for strongly convex and exp-concave losses respectively. Hence in view of Proposition 18 and Definition 19, we can conclude that:

- FLH with OGD as base learners is an SA algorithm for strongly convex losses.
- FLH with ONS as base learners is an SA algorithm for exp-concave losses. (We treat dimension d as a constant problem parameter and consider minimaxity only wrt n .)

We have the following guarantee on runtime.

Proposition 20 (Hazan and Seshadhri, 2007) *Let ρ be the per round run time of base learners and r_n be the static regret suffered by the base learners over n rounds. Then FLH procedure has a runtime of $O(\rho n)$ per round. To improve the runtime one can use AFLH procedure from (Hazan and Seshadhri, 2007) that incurs $O(\rho \log n)$ runtime overhead per round and suffers $O(r_n \log n)$ static regret in any interval.*

Appendix C. Proofs for Section 3

We start by providing an example of a scenario where λ in Lemma 3 can scale linearly with n .

Example 21 *Consider the $\mathcal{TV}(C_n)$ class with $C_n = 1$ and $n \geq 6$. Let the offline optimal be given by the step sequence $u_1 = \dots = u_{(n/2)-1} = 0$ and $u_{n/2} = \dots = u_n = 1$. Our aim is to generate a sequence of labels y_t such that this sequence \mathbf{u} is indeed the offline optimal in the class $\mathcal{TV}(1)$ along with the property that the optimal dual variable λ scales linearly with the horizon n .*

Clearly we must have $s_{(n/2)-1} = 1$. For some appropriate parameter ϵ , consider the following sign assignment:

- $s_{(n/2)-2} = 1 - \epsilon, s_{(n/2)-3} = 1 - 2\epsilon, \dots, s_1 = 1 - ((n/2) - 2)\epsilon,$
- $s_{n-1} = \epsilon, s_{n-2} = 2\epsilon, \dots, s_{n/2} = (n/2)\epsilon.$

By setting $\epsilon = 2/n$ for $n \geq 6$, we get a consistent sign assignment because $s_t \in [-1, 1]$ for all $1 \leq t \leq (n/2) - 2$ which corresponds to the portion where $u_t = 0$; $s_{(n/2)-1} = 1$; and $s_t \in [-1, 1]$ for all $n/2 \leq t \leq n - 1$ which corresponds to the portion where $u_t = 1$.

By taking $\lambda = n/2$ the adversary can generate labels y_t according to the stationarity condition in Lemma 3 as follows:

- $y_1 = -2,$
- $y_t = -1, \text{ for } 2 \leq t \leq (n/2) - 1,$
- $y_{n/2} = 1,$
- $y_t = 2, \text{ for } (n/2) + 1 \leq t \leq n.$

Since the TV of the sequence \mathbf{u} is 1, the complementary slackness is also satisfied. Thus we conclude that if the labels $y_t \in [-2, 2]$ are generated as above, the offline optimal sequence in $\mathcal{TV}(1)$ class is given by the step sequence \mathbf{u} . Furthermore, the optimal dual variable $\lambda = n/2$ scales linearly with the horizon.

Lemma 3 (characterization of offline optimal) Consider the following convex optimization problem (where $\tilde{z}_1, \dots, \tilde{z}_{n-1}$ are introduced as dummy variables)

$$\begin{aligned} \min_{\tilde{u}_1, \dots, \tilde{u}_n, \tilde{z}_1, \dots, \tilde{z}_{n-1}} \quad & \frac{1}{2} \sum_{t=1}^n (y_t - \tilde{u}_t)^2 \\ \text{s.t.} \quad & \tilde{z}_t = \tilde{u}_{t+1} - \tilde{u}_t \quad \forall t \in [n-1], \\ & \sum_{t=1}^{n-1} |\tilde{z}_t| \leq C_n \end{aligned} \tag{6a}$$

Let $u_1, \dots, u_n, z_1, \dots, z_{n-1}$ be the optimal primal variables and let $\lambda \geq 0$ be the optimal dual variable corresponding to the last constraint (6a). By the KKT conditions, we have

- **stationarity:** $y_t = u_t - \lambda(s_t - s_{t-1})$, where $s_t \in \partial|z_t|$ (a subgradient). Specifically, $s_t = \text{sign}(u_{t+1} - u_t)$ if $|u_{t+1} - u_t| > 0$ and s_t is some value in $[-1, 1]$ otherwise. For convenience of notations later, we also define $s_n = s_0 = 0$.
- **complementary slackness:** $\lambda(\sum_{t=2}^n |u_t - u_{t-1}| - C_n) = 0$.

Proof We can form the Lagrangian of the optimization problem as:

$$\mathcal{L}(\tilde{\mathbf{u}}, \tilde{\mathbf{z}}, \tilde{\mathbf{v}}, \tilde{\lambda}) = \frac{1}{2} \sum_{t=1}^n (y_t - \tilde{u}_t)^2 + \tilde{\lambda} \left(\sum_{t=1}^{n-1} |\tilde{z}_t| - C_n \right) + \sum_{t=1}^{n-1} \tilde{v}_t (\tilde{u}_{t+1} - \tilde{u}_t - \tilde{z}_t),$$

for dual variables $\tilde{\lambda} > 0$ and $\tilde{v} \in \mathbb{R}^{n-1}$ unconstrained. Let the $(\mathbf{u}, \mathbf{z}, \mathbf{v}, \lambda)$ be the optimal primal and dual variables. By stationarity conditions, we have

$$u_t - y_t = v_t - v_{t-1},$$

where we take $v_0 = v_n = 0$ and

$$v_t = \lambda s_t$$

Combining the above two equations and the complementary slackness rule yields the lemma. \blacksquare

Lemma 5 (key partition) Initialize $\mathcal{Q} \leftarrow \Phi$. Starting from time 1, spawn a new bin $[i_s, i_t]$ whenever $\sum_{j=i_s+1}^{i_t+1} |u_j - u_{j-1}| > B/\sqrt{n_i}$, where $n_i = i_t - i_s + 2$. Add the spawned bin $[i_s, i_t]$ to \mathcal{Q} . Consider the following post processing routine.

1. Initialize $\mathcal{P} \leftarrow \Phi$.
2. For $i \in [|\mathcal{Q}|]$:
 - if $u_{i_t} = u_{i_t+1}$:
 - (a) Let p be the largest time point with $u_{p:i_t}$ being constant and let q be the smallest time point with $u_{i_t+1:q}$ being constant.
 - (b) Add bin $[i_s, p-1]$ to \mathcal{P} .
 - (c) If $(i+1)_t > q$ then add $[p, q]$ to \mathcal{P} and set $(i+1)_s \leftarrow q+1$.
 - (d) Goto Step 2.
 - Add $[i_s, i_t]$ to \mathcal{P} . Goto Step 2.

Let $M := |\mathcal{P}|$. We have $M = O\left(1 \vee n^{1/3} C_n^{2/3} B^{-2/3}\right)$. Further for any bin $[i_s, i_t] \in \mathcal{P}$, it holds that $\sum_{j=i_s+1}^{i_t} |u_j - u_{j-1}| \leq B/\sqrt{n_i}$ where $n_i = i_t - i_s + 1$.

Proof Let's use the notation $TV[a, b]$ to denote the TV incurred by the optimal solution sequence in the interval $[a, b]$. Let $\mathcal{Q} = \{[t_1, \bar{t}_1], \dots, [t_N, \bar{t}_N]\}$ with $t_1 := 1$ and $\bar{t}_N := n$. Let $n_j := \bar{t}_j - t_j + 1$. We have,

$$\sum_{j=1}^{N-1} TV[t_j, \bar{t}_j + 1] \leq C_n.$$

By construction we have $TV[t_j, \bar{t}_j + 1] > \nu/\sqrt{n_j}$. So,

$$\begin{aligned} C_n &\geq \sum_{j=1}^{N-1} \nu/\sqrt{n_j} \\ &\geq (N-1)^{3/2} \nu/\sqrt{n}, \end{aligned}$$

where the last line follows by Jensen's inequality. Rearranging gives the bound on $N = O\left(1 \vee n^{1/3} C_n^{2/3} B^{-2/3}\right)$. Now the post processing step only increases the number of bins by $O(N)$. Thus we get $M = O\left(1 \vee n^{1/3} C_n^{2/3} B^{-2/3}\right)$. \blacksquare

Lemma 7 (bounding $T_{1,i}$) Assume that we run FLH with the settings described in Theorem 1. For any bin i we have $T_{1,i} = O(B^2 \log n)$

Proof Note that FTL with squared error losses outputs predictions which are online averages of the past labels that the algorithm has seen so far. Hence the predictions of all base learners as well as FLH belong to the interval $[-B, B]$. It is known that (see for eg. (Cesa-Bianchi and Lugosi, 2006), Chapter 3) squared error losses are $1/(8B^2)$ exp-concave in the interval $[-B, B]$. Further FTL with squared error losses suffers only logarithmic regret of $O(B^2 \log n)$ ((Cesa-Bianchi and Lugosi, 2006), Chapter 3).

Hence due to the adaptive regret bound of FLH (Theorem 3.2 in (Hazan and Seshadhri, 2007)) by setting the learning rate $\zeta = 1/(8B^2)$, we have that the static regret of FLH in any interval $[i_s, i_t]$ is also $O(\log n)$. This proves the lemma. \blacksquare

Lemma 8 (bounding $T_{2,i}$) Define $C_i := \sum_{j=i_s+1}^{i_t} |u_j - u_{j-1}|$, the TV within bin i incurred by the offline optimal solution. Let $\Delta s_i := s_{i_t} - s_{i_s-1}$ and $n_i := i_t - i_s + 1$. We have $T_{2,i} \leq \frac{-\lambda^2 (\Delta s_i)^2}{n_i}$.

Proof From the stationarity conditions in Lemma 3, we can write

$$\bar{u}_i - \bar{y}_i = \frac{\lambda \Delta s_i}{n_i}. \quad (9)$$

Further,

$$\begin{aligned} \sum_{j=i_s}^{i_t} (y_j - \bar{y}_i)^2 - (y_j - \bar{u}_i)^2 &= n_i (\bar{u}_i - \bar{y}_i)^2 + 2 \sum_{j=i_s}^{i_t} (y_j - \bar{u}_i) (\bar{u}_i - \bar{y}_i) \\ &= -n_i (\bar{u}_i - \bar{y}_i)^2 \end{aligned}$$

Now plugging in Eq. (9) yields the lemma. \blacksquare

Lemma 9 (bounding $T_{3,i}$) Let C_i and Δs_i be as in Lemma 8.

Case(a) If $|\Delta s_i| > 0$ then $T_{3,i} \leq B^2 + 6\lambda C_i$.

Case(b) If $\Delta s_i = 0$ with $s_{i_s-1} = s_{i_t} = 1$ and the offline optimal \mathbf{u} is non-decreasing within bin i , then $T_{3,i} \leq B^2$.

Case(c) If $\Delta s_i = 0$ with $s_{i_s-1} = s_{i_t} = -1$ and the offline optimal \mathbf{u} is non-increasing within bin i , then $T_{3,i} \leq B^2$.

Proof Applying stationarity conditions, we have

$$\begin{aligned}
T_{3,i} &= \sum_{j=i_s}^{i_t} (y_j - \bar{u}_i)^2 - (y_j - u_j)^2 \\
&= \sum_{j=i_s}^{i_t} (u_j - \bar{u}_i)(2y_j - \bar{u}_i - u_j) \\
&= \sum_{j=i_s}^{i_t} (u_j - \bar{u}_i)(2y_j - 2u_j + u_j - \bar{u}_i) \\
&= \sum_{j=i_s}^{i_t} (u_j - \bar{u}_i)^2 + 2\lambda(u_j - \bar{u}_i)(s_{j-1} - s_j) \\
&\leq n_i C_i^2 + \sum_{j=i_s}^{i_t} 2\lambda(u_j - \bar{u}_i)(s_{j-1} - s_j), \tag{10}
\end{aligned}$$

where in the last line we used $|u_j - \bar{u}_i| \leq C_i$. Also observe that $n_i C_i^2 \leq B^2$ for bins in the partition \mathcal{P} by Lemma 5. Now by expanding the second term followed by a regrouping of the terms in the summation, we can write

$$\begin{aligned}
\sum_{j=i_s}^{i_t} 2\lambda(u_j - \bar{u}_i)(s_{j-1} - s_j) &= 2\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) + 2\lambda \sum_{j=i_s+1}^{i_t} |u_j - u_{j-1}| \\
&= 2\lambda C_i + 2\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) \tag{11}
\end{aligned}$$

Now we discuss the three cases.

Case (a) When $|\Delta s_i| > 0$, then by triangle inequality we have

$$2\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) \leq 4\lambda C_i.$$

Case (b) In this case we have $2\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) = \lambda(u_{i_s} - u_{i_t}) = -2\lambda C_i$ since the sequence is non-decreasing within the bin. Hence this term cancels with the corresponding additive term of $2\lambda C_i$ in Eq. (11).

Case (c) By similar logic as in case (b) we can once again write

$$2\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) = -2\lambda C_i.$$

Substituting the bound of each case into (10), we obtain the expression as stated. \blacksquare

Proposition 10 (Extension to higher dimensions) Consider a protocol where at each time the learner predicts a vector $\mathbf{x}_t \in \mathbb{R}^d$ after which the adversary reveals \mathbf{y}_t such that $\|\mathbf{y}_t\|_\infty \leq B$. Consider a comparator sequence of vectors $\mathbf{w}_1, \dots, \mathbf{w}_n$ such that $TV(\mathbf{w}_{1:n}) := \sum_{t=2}^n \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1 \leq C_n$. By running d instances of FLH with learning rate $\zeta = 1/(8B^2)$ and FTL as base learners, where instance $i, i \in [d]$, predicts $\mathbf{x}_t[i]$ at time t , we have

$$R_n(\mathbf{w}_{1:n}) := \sum_{j=1}^n \|\mathbf{y}_t - \mathbf{x}_t\|_2^2 - \|\mathbf{y}_t - \mathbf{w}_t\|_2^2 = \tilde{O}\left(dB^2 \log n \vee d^{1/3} n^{1/3} C_n^{2/3} B^{4/3}\right).$$

Proof Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be the offline optimal sequence. Let $C_n[k] = \sum_{t=2}^n |\mathbf{u}_t[k] - \mathbf{u}_{t-1}[k]|$ be its TV allocated to coordinate k . WLOG, let's assume the FLH for coordinates $k \in [k']$ for $k' \leq d$ incurs $\tilde{O}(n^{1/3}(C_n[k])^{2/3}B^{4/3})$ regret and the regret incurred by FLH for coordinates $k > k'$ is $O(\log n)$. Since squared error losses decomposes coordinate-wise, we have

$$\begin{aligned} R_n(\mathbf{w}_{1:n}) &\leq \sup_{\mathbf{w}_{1:n}:TV(\mathbf{w}_{1:n})\leq C_n} R_n(\mathbf{w}_{1:n}) \\ &= R_n(\mathbf{u}_{1:n}) \\ &= (d - k')B^2 \log n + \sum_{k=1}^{k'} \tilde{O}\left(n^{1/3}(C_n[k])^{2/3}B^{4/3}\right) \\ &\leq (d - k')B^2 \log n + \tilde{O}\left(n^{1/3}(k')^{1/3}B^{4/3}\left(\sum_{k=1}^{k'} C_n[k]\right)^{2/3}\right), \end{aligned}$$

where the last line follows by Holder's inequality $\mathbf{x}^T \mathbf{y} \leq \|\mathbf{x}\|_3 \|\mathbf{y}\|_{3/2}$, where we treat \mathbf{x} as just a vector of ones in $\mathbb{R}^{k'}$. The above expression can be further upper bounded by $\tilde{O}\left(2dB^2 \log n \vee 2d^{1/3}n^{1/3}C_n^{2/3}B^{4/3}\right)$. ■

Proposition 11 (Lower bound) *Assume the protocol and notations of Proposition 10. For any algorithm, we have*

$$\sup_{\mathbf{w}_{1:n}:TV(\mathbf{w}_{1:n})\leq C_n} R_n(\mathbf{w}_{1:n}) = \Omega\left(dB^2 \log n \vee d^{1/3}n^{1/3}C_n^{2/3}B^{4/3}\right).$$

Proof Consider a fixed (but unknown) sequence $\mathbf{u}_1, \dots, \mathbf{u}_n$ such that $TV(\mathbf{u}_{1:n}) \leq C_n$ with $\|\mathbf{u}_t\|_\infty \leq B/2$ and TV along the coordinate $k \in [d]$, $TV(\mathbf{u}_{1:n}[k]) \leq C_n/d$ for all k . Let the labels be $\mathbf{y}_t = \mathbf{u}_t + \boldsymbol{\epsilon}_t$ where each coordinate of $\boldsymbol{\epsilon}_t$ is generated by iid $U[-B/2, B/2]$. Further $\boldsymbol{\epsilon}_1, \dots, \boldsymbol{\epsilon}_n$ are also iid. Then by the results of (Donoho and Johnstone, 1998), for any prediction strategy that produces outputs \mathbf{x}_t , we have

$$\begin{aligned} \sup_{\mathbf{w}_{1:n}:TV(\mathbf{w}_{1:n})\leq C_n} R_n(\mathbf{w}_{1:n}) &\geq \sum_{k=1}^d \sum_{t=1}^n E[(\mathbf{y}_t[k] - \mathbf{x}_t[k])^2 - (\mathbf{y}_t[k] - \mathbf{u}_t[k])^2] \\ &=_{(a)} \sum_{k=1}^d \sum_{t=1}^n E[(\mathbf{u}_t[k] - \mathbf{x}_t[k])^2] \\ &= \sum_{k=1}^d \Omega(n^{1/3}(C_n/d)^{2/3}B^{4/3}) \\ &= \Omega(d^{1/3}n^{1/3}C_n^{2/3}B^{4/3}), \end{aligned}$$

where in line (a) we used the fact that $\mathbf{x}_t[k]$ is independent of $\mathbf{y}_t[k]$ and $\mathbf{y}_t[k] - \mathbf{u}_t[k] \sim U[-B/2, B/2]$.

The $dB^2 \log n$ part of the lower bound is implied by the lower bound construction of Vovk (Vovk, 2001) (or cf. proof of Theorem 11.9 in (Cesa-Bianchi and Lugosi, 2006)).

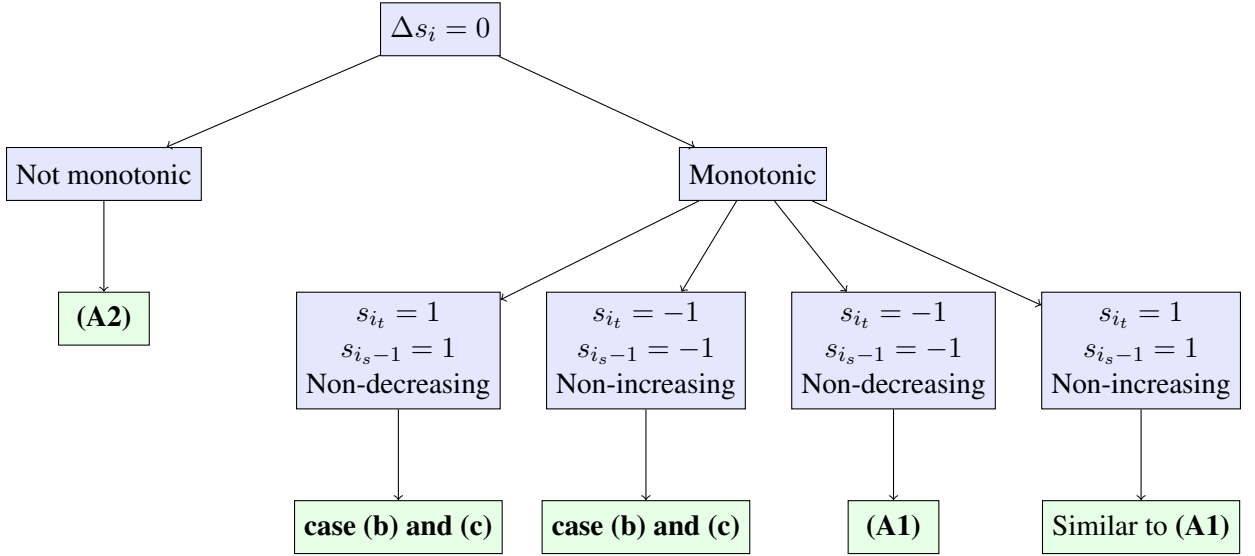


Figure 3: Various configurations of the optimal sequence within a bin $[i_s, i_t]$ with $\Delta s_i = 0$. The leaf nodes indicate the labels of the paragraphs in the Proof of Theorem 1 to handle each scenario. ■




Figure 4: A configuration of optimal sequence within a bin $[i_s, i_t]$ with $|\Delta s_i| \neq 0$. The leaf node indicate the label of the paragraph in the Proof of Theorem 1 to handle this scenario.

Close comparison to lower bound in (Baby and Wang, 2019). For the case of 1D forecasting of TV bounded sequences, (Baby and Wang, 2019) consider a stochastic setting where the labels obey $y_t = w_t + \epsilon_t$ for some iid σ subgaussian noise ϵ_t and $w_t \in \mathcal{TV}^B(C_n)$. They provide a lower bound of $\tilde{\Omega} \left((nB^2 \wedge n\sigma^2 \wedge n^{1/3}C_n^{2/3}\sigma^{4/3}) + (nB^2 \wedge BC_n) + B^2 \right)$ where $(a \wedge b) = \min\{a, b\}$. In accordance with the proof of Proposition 11, we can take $\sigma = B/2$ and $w_{1:n} \in \mathcal{TV}^{B/2}(C_n)$ to translate this lower bound into our setting for 1D case to get a lower bound of:

$$R_n(C_n) = \tilde{\Omega} \left((nB^2 \wedge n^{1/3}C_n^{2/3}B^{4/3}) + (nB^2 \wedge BC_n) + B^2 \right). \quad (12)$$

Any learner must have to incur $O(B^2)$ loss in the first round. Combining this with the upper bound in Theorem 1 along with the trivial regret bound of $O(nB^2)$ we can get a refined regret upper bound of:

$$R_n(C_n) = \tilde{O} \left((nB^2 \wedge n^{1/3}C_n^{2/3}B^{4/3}) + B^2 \right). \quad (13)$$

Comparing Eq.(12) and (13) seems to falsely suggest that during the regime where $n^{1/3}C_n^{2/3}B^{4/3} < BC_n < nB^2$ upper bound in Eq.(13) is smaller than the lower bound in Eq.(12). But $n^{1/3}C_n^{2/3}B^{4/3} < BC_n$ happens when $C_n > nB$, in which case $BC_n < nB^2$ is not satisfied. Hence we conclude that this regime is not realisable implying no contradictions.

Close comparison to lower bound in (Yuan and Lamperski, 2019). Proposition 1 of (Yuan and Lamperski, 2019) considers squared error losses in 1D and show that when $C_n = n^{\frac{2+\gamma}{4-\gamma}}$ for all $\gamma \in (0, 1)$, the dynamic regret obeys

$$R_n(C_n) = \Omega\left(\log n \vee (nC_n)^{\gamma/2}\right).$$

We proceed to show that our lower bound of $\Omega(\log n \vee n^{1/3}C_n^{2/3})$ is tighter than this. Whenever $C_n = n^{\frac{2+\gamma}{4-\gamma}}$, we have

$$(nC_n)^{\gamma/2} = n^{\frac{3\gamma}{4-\gamma}},$$

and,

$$n^{1/3}C_n^{2/3} = n^{\frac{8+\gamma}{12-3\gamma}}.$$

It can be verified that for all $\gamma \in (0, 1)$, $n^{\frac{3\gamma}{4-\gamma}} \leq n^{\frac{8+\gamma}{12-3\gamma}}$ making our lower bound tighter.

Appendix D. Proofs for Section 4

D.1. One dimensional setting

In the section, we adopt all the notations used in Section 3. For the sake of simplicity of exposition, we first present the results in one dimensional setting and extend it later to higher dimensions. We have the following guarantee in one dimension.

Theorem 22 ($d = 1$) *By using the base learner as ONS with parameter $\zeta = \min\left\{\frac{1}{4G^+(2B+2G/\beta)}, \alpha\right\}$ and decision set \mathcal{D} and choosing learning rate $\eta = \alpha$, FLH guarantees a dynamic regret $R_n(C_n) = \tilde{O}\left(n^{1/3}C_n^{2/3} \vee \log n\right)$.*

We start the analysis by inspecting the KKT conditions.

Lemma 23 (*characterization of offline optimal*) *Consider the following convex optimization problem.*

$$\begin{aligned} & \min_{\tilde{u}_1, \dots, \tilde{u}_n, \tilde{z}_1, \dots, \tilde{z}_{n-1}} \sum_{t=1}^n f_t(\tilde{u}_t) \\ & \text{s.t.} \quad \tilde{z}_t = \tilde{u}_{t+1} - \tilde{u}_t \quad \forall t \in [n-1], \\ & \quad \quad \sum_{t=1}^{n-1} |\tilde{z}_t| \leq C_n, \\ & \quad \quad -B \leq \tilde{u}_t \quad \forall t \in [n], \\ & \quad \quad \tilde{u}_t \leq B \quad \forall t \in [n], \end{aligned} \tag{14a}$$

$$-B \leq \tilde{u}_t \quad \forall t \in [n], \tag{14b}$$

$$\tilde{u}_t \leq B \quad \forall t \in [n], \tag{14c}$$

Let $u_1, \dots, u_n, z_1, \dots, z_{n-1}$ be the optimal primal variables and let $\lambda \geq 0$ be the optimal dual variable corresponding to the constraint (14a). Further, let $\gamma_t^- \geq 0, \gamma_t^+ \geq 0$ be the optimal dual variables that correspond to constraints (14b) and (14c) respectively for all $t \in [n]$. By the KKT conditions, we have

- **stationarity:** $\nabla f_t(u_t) = \lambda (s_t - s_{t-1}) + \gamma_t^- - \gamma_t^+$, where $s_t \in \partial|z_t|$ (a subgradient). Specifically, $s_t = \text{sign}(u_{t+1} - u_t)$ if $|u_{t+1} - u_t| > 0$ and s_t is some value in $[-1, 1]$ otherwise. For convenience of notations later, we also define $s_n = s_0 = 0$.
- **complementary slackness:** (a) $\lambda (\sum_{t=2}^n |u_t - u_{t-1}| - C_n) = 0$; (b) $\gamma_t^-(u_t + B) = 0$ and $\gamma_t^+(u_t - B) = 0$ for all $t \in [n]$

Terminology. We will refer to the optimal primal variables u_1, \dots, u_n in Lemma 23 as the *offline optimal sequence* in this section.

Next, we record an easy corollary of Lemma 5.

Corollary 24 (key partition) Assume the notations of Lemma 5. Create a partition of \mathcal{P} of $[n]$ with the procedure mentioned in Lemma 5. Then for any $[i_s, i_t] \in \mathcal{P}$, we have

- (TV constraint) $\sum_{j=i_s+1}^{i_t} |u_j - u_{j-1}| \leq B/\sqrt{n_i}$,
- (Bins bound) $M := |\mathcal{P}| = O(n^{1/3}C_n^{2/3})$.
- (Structural property) If $i_s > 1$ then $u_{i_s} \neq u_{i_s-1}$. Similarly if $i_t < n$ then $u_{i_t} \neq u_{i_t+1}$.

Now we make an important observation regarding the dual variables γ_j^- and γ_j^+ . The following property will be used several times in the proofs to follow.

Lemma 25 Define $\Gamma_i^+ := \sum_{j=i_s}^{i_t} \gamma_j^+$ and $\Gamma_i^- := \sum_{j=i_s}^{i_t} \gamma_j^-$. Consider a bin $[i_s, i_t] \in \mathcal{P}$, where \mathcal{P} is the partition of $[n]$ constructed in Corollary 24. Then at-least one of the following is always satisfied.

- $\gamma_j^- = 0$ for all $j \in [i_s, i_t]$.
- $\gamma_j^+ = 0$ for all $j \in [i_s, i_t]$.

Consequently we have $\sum_{j=i_s}^{i_t} |\gamma_j^-| + |\gamma_j^+| = |\Gamma_i^- - \Gamma_i^+|$, for any bin $[i_s, i_t] \in \mathcal{P}$.

Proof From the properties of the partition \mathcal{P} in Corollary 24, we have that the TV of the offline optimal incurred within each bin is at-most $B/\sqrt{n_i} \leq B$. Hence within bin $[i_s, i_t] \in \mathcal{P}$, if the optimal sequence attains the value $-B$ at some time point, it can never attain the value B and vice-versa. So due to complementary slackness rule in Lemma 23, either $\gamma_j^+ = 0$ or $\gamma_j^- = 0$ uniformly for all $j \in [i_s, i_t]$. The last line in the statement of lemma follows by recalling that $\gamma_j^- \geq 0$ and $\gamma_j^+ \geq 0$ from Lemma 23. ■

For convenience, we recall here the regret decomposition of Eq.(7) specified to one dimensional setting. Let \mathcal{P} be a partition of $[n]$ into M bins as specified in Corollary 24. Let $[i_s, i_t]$ denote the i^{th} bin in \mathcal{P} and let n_i be its length. Define $\bar{u}_i = \frac{1}{n_i} \sum_{j=i_s}^{i_t} u_j$ and $\dot{u}_i = \bar{u}_i - \frac{1}{n_i \beta} \sum_{j=i_s}^{i_t} \nabla f_j(\bar{u}_i)$

where β is as in Assumption EC-2. Let x_j be the prediction made by FLH at time j . We start with following regret decomposition.

$$R_n(C_n) \leq \sum_{i=1}^M \underbrace{\sum_{j=i_s}^{i_t} f_j(x_j) - f_j(\dot{u}_i)}_{T_{1,i}} + \sum_{i=1}^M \underbrace{\sum_{j=i_s}^{i_t} f_j(\dot{u}_i) - f_j(\bar{u}_i)}_{T_{2,i}} + \sum_{i=1}^M \underbrace{\sum_{j=i_s}^{i_t} f_j(\bar{u}_i) - f_j(u_j)}_{T_{3,i}}.$$

We proceed to bound the terms $T_{1,i}, T_{2,i}, T_{3,i}$ for the bins that belong to the partition \mathcal{P} .

Lemma 26 (bounding $T_{1,i}$) *Let the experts in FLH be the ONS algorithms with parameter $\zeta = \min \left\{ \frac{1}{4G^\dagger(2B+2G)}, \alpha \right\}$ and decision set \mathcal{D} . Also choose learning rate $\eta = \alpha$, for FLH. Then for any bin $[i_s, i_t]$ we have,*

$$\begin{aligned} \sum_{j=i_s}^{i_t} f_j(x_j) - f_j(\dot{u}_i) &= O \left(BG^\dagger \log n + GG^\dagger \log n + \frac{\log n}{\alpha} \right) \\ &= O(\log n). \end{aligned}$$

Proof First we proceed to bound $|\dot{u}_i|$. Since $|\nabla f_j(u_j)| \leq G$ by Assumption EC-1, we have

$$\begin{aligned} |\dot{u}_i| &\leq |\bar{u}_i| + \frac{G}{\beta} \\ &\leq B + G, \end{aligned}$$

since $\beta \geq 1$ by Assumption EC-2. For any $x \in \mathcal{D}$, we have $|x - \dot{u}_i| \leq 2B + 2G$ by triangle inequality.

By Assumption EC-4 we have $|\nabla f_j(x)| \leq G^\dagger$ for any $x \in \mathcal{D}$. Also, recall that by Assumption EC-3, the loss functions f_j are α exp-concave in the domain \mathcal{D} . Let p_j be the predictions of ONS in the interval $[i_s, i_t]$. If we choose $\zeta = \min \left\{ \frac{1}{4G^\dagger(2B+2G)}, \alpha \right\}$ as the parameter of the ONS, Theorem 2 of (Hazan et al., 2007) implies that

$$\begin{aligned} \sum_{j=i_s}^{i_t} f_j(p_j) - f_j(\dot{u}_i) &= O \left(BG^\dagger \log n + GG^\dagger \log n \right) \\ &= O(\log n). \end{aligned}$$

Now the Lemma is implied by the SA regret bound of FLH (Theorem 3.2 of (Hazan and Seshadhri, 2007)). ■

Lemma 27 (bounding $T_{2,i}$) *For a bin $[i_s, i_t] \in \mathcal{P}$, let C_i, n_i and Δs_i be as in Lemma 8 and Γ_i^+, Γ_i^- be as in Lemma 25. We have*

$$\sum_{j=i_s}^{i_t} f_j(\dot{u}_i) - f_j(\bar{u}_i) \leq \frac{-(\lambda \Delta s_i + \Gamma_i^- - \Gamma_i^+)^2}{2n_i \beta} + \lambda |\Delta s_i| C_i + |\Gamma_i^- - \Gamma_i^+| C_i.$$

Proof We start with the short proof the descent lemma. Let $g(x)$ be a L strongly smooth function. Let $x^+ = x - \mu \nabla f(x)$ for some $\mu > 0$. Then we have

$$\begin{aligned} g(x^+) - g(x) &\leq (\nabla g(x))^2 \left(\frac{L}{2} \mu^2 - \mu \right) \\ &= \frac{-(\nabla g(x))^2}{2L}, \end{aligned}$$

by choosing $\mu = 1/L$. By taking $g(x) = \sum_{j=i_s}^{i_t} f_j(x)$ and noting that g is $n_i \beta$ gradient Lipschitz due to Assumption EC-2, we get

$$\begin{aligned} T_{2,i} &:= \sum_{j=i_s}^{i_t} f_j(\bar{u}_i) - f_j(\bar{u}_i) \\ &\leq \frac{-\left(\sum_{j=i_s}^{i_t} \nabla f_j(\bar{u}_i)\right)^2}{2n_i \beta} \\ &= \frac{-1}{2n_i \beta} \left(\sum_{j=i_s}^{i_t} \nabla f_j(u_j) + \nabla f_j(\bar{u}_i) - \nabla f_j(u_j) \right)^2 \\ &\leq \frac{-1}{2n_i \beta} \left(\sum_{j=i_s}^{i_t} \nabla f_j(u_j) \right)^2 + \frac{1}{n_i \beta} \left| \sum_{j=i_s}^{i_t} \nabla f_j(u_j) \right| \left| \sum_{j=i_s}^{i_t} \nabla f_j(\bar{u}_i) - \nabla f_j(u_j) \right|. \end{aligned}$$

From the KKT conditions in Lemma 23 we have $\sum_{j=i_s}^{i_t} \nabla f_j(u_j) = \lambda \Delta s_i + \Gamma_i^- - \Gamma_i^+$. Since f_j are β -gradient Lipschitz and $|\bar{u}_i - u_j| \leq C_i$, we also have

$$\left| \sum_{j=i_s}^{i_t} \nabla f_j(\bar{u}_i) - \nabla f_j(u_j) \right| \leq n_i \beta C_i.$$

Substituting these we get,

$$T_{2,i} \leq \frac{-(\lambda \Delta s_i + \Gamma_i^- - \Gamma_i^+)^2}{2n_i \beta} + \lambda |\Delta s_i| C_i + |\Gamma_i^- - \Gamma_i^+| C_i. \quad \blacksquare$$

Lemma 28 (bounding $T_{3,i}$) For a bin $[i_s, i_t] \in \mathcal{P}$, let C_i, n_i and Δs_i be as in Lemma 8 and Γ_i^+, Γ_i^- be as in Lemma 25.

case(a) If $|\Delta s_i| > 0$ then we have,

$$\sum_{j=i_s}^{i_t} f_j(\bar{u}_i) - f_j(u_j) \leq \frac{\beta n_i C_i^2}{2} + 3\lambda C_i + |\Gamma_i^- - \Gamma_i^+| C_i.$$

case(b) If $\Delta s_i = 0$ with $s_{i_s-1} = s_{i_t} = 1$ and the offline optimal \mathbf{u} is non-decreasing within bin i with $-B < u_i < B$ for all $i \in [i_s, i_t]$, then

$$\sum_{j=i_s}^{i_t} f_j(\bar{u}_i) - f_j(u_j) \leq \frac{\beta n_i C_i^2}{2}.$$

case(c) If $\Delta s_i = 0$ with $s_{i_s-1} = s_{i_t} = -1$ and the offline optimal \mathbf{u} is non-increasing within bin i with $-B < u_i < B$ for all $i \in [i_s, i_t]$, then

$$\sum_{j=i_s}^{i_t} f_j(\bar{u}_i) - f_j(u_j) \leq \frac{\beta n_i C_i^2}{2}.$$

Proof Due to strong smoothness, we have

$$\begin{aligned} T_{3,i} &:= \sum_{j=i_s}^{i_t} f_j(\bar{u}_i) - f_j(u_j) \\ &\leq \sum_{j=i_s}^{i_t} \nabla f_j(u_j)(\bar{u}_i - u_j) + \frac{\beta}{2}(\bar{u}_i - u_j)^2 \\ &\leq \frac{\beta n_i C_i^2}{2} + \sum_{j=i_s}^{i_t} \nabla f_j(u_j)(\bar{u}_i - u_j). \end{aligned}$$

Now by expanding the second term and using the structure of gradients as in Lemma 23 followed by a regrouping of the terms in the summation we can write,

$$\begin{aligned} \sum_{j=i_s}^{i_t} \nabla f_j(u_j)(\bar{u}_i - u_j) &= \lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) + \lambda \sum_{j=i_s+1}^{i_t} |u_j - u_{j-1}| \\ &\quad + \sum_{j=i_s}^{i_t} (\gamma_j^- - \gamma_j^+)(\bar{u}_i - u_j) \\ &\leq \lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) + \lambda C_i + |\Gamma_i^- - \Gamma_i^+| C_i, \quad (15) \end{aligned}$$

where the last line follows due to Lemma 25 and $|\bar{u}_i - u_j| \leq C_i$ for all $j \in [i_s, i_t]$.

Now we consider three cases in the statement of the lemma.

case (a) When $|\Delta s_i| > 0$, then by triangle inequality we have $\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) \leq 2\lambda C_i$.

case (b) In this case we have

$\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) = \lambda(u_{i_s} - u_{i_t}) = -\lambda C_i$ since the sequence is non-decreasing within the bin. Hence this term cancels with the corresponding additive term of λC_i in Eq. (15). Further $\gamma_j^- = \gamma_j^+ = 0$ since $-B < u_j < B$ for all $j \in [i_s, i_t]$.

case (c) By similar logic as in case (b) we can once again write

$$\lambda(s_{i_s-1}(u_{i_s} - \bar{u}_i) - s_{i_t}(u_{i_t} - \bar{u}_i)) = -\lambda C_i.$$

Putting everything together now yields the lemma. ■

Proof of Theorem 22. The strategy of the proof is to bound the regret incurred within each time interval $[i_s, i_t] \in \mathcal{P}$ where \mathcal{P} is as in Corollary 24 and add them up towards the end. We annotate several key paragraphs for the purposes of referring the arguments contained in them at later points.

If the the partition \mathcal{P} contains only one bin, then we split it into at-most two bins $[1, a]$ and $[a + 1, n]$ such that the optimal sequence is constant within $[1, a]$ and hence regret incurred within this bin is $\tilde{O}(1)$ by Strong Adaptivity of FLH. The regret incurred in the bin $[a + 1, i_t]$ can be bounded by using the arguments below. So in what follows we assume for a bin $[i_s, i_t]$ either $i_s > 1$ or $i_t < n$.

By virtue of Lemma 25, any bin $[i_s, i_t] \in \mathcal{P}$ will have either $\gamma_j^- = 0$ for all $j \in [i_s, i_t]$ or $\gamma_j^+ = 0$ for all $j \in [i_s, i_t]$. Below we bound the regret for bins with $\gamma_j^+ = 0$ uniformly for all $j \in [i_s, i_t]$. The arguments for the alternate case where $\gamma_j^- = 0$ follows similarly. Figures 6, 7 and 8 sketch the floor plan of the proof pictorially. Throughout the proof, we will use the properties in Corollary 24 in conjunction with the observations in Remark 6.

(S1): Consider a bin with $\Delta s_i = 0$ with $s_{i_t} = s_{i_s-1} = 1$ and the optimal sequence is non-decreasing within the bin. By the structural property of Corollary 24, this happens when $u_{i_s} > u_{i_s-1}$, $u_{i_t+1} > u_{i_t}$ where $1 < i_s < i_t < n$. Since the sequence is non-decreasing, it never attains $-B$ within this bin. Hence this is the same situation as in case (b) of Lemma 28. We have $T_{1,i} = \tilde{O}(1)$ due to Lemma 26. $T_{2,i} = 0$ due to Lemma 27 as $\Gamma_i^+ = \Gamma_i^- = 0$ since the sequence never attains $\pm B$ within the current bin combined with the fact that $\Delta s_i = 0$. $T_{3,i} = O(1)$ due to Lemma 28 combined with the fact that $C_i \leq B/\sqrt{n_i}$ due to Corollary 24. So the total regret within the current bin is bounded by $T_{1,i} + T_{2,i} + T_{3,i} = \tilde{O}(1)$.

The total regret for a bin satisfying case (c) of Lemma 28 can be bound using similar arguments as above.

The three cases where (i) $\Delta s_i = 0$ with $s_{i_s-1} = s_{i_t} = -1$ and the offline optimal \mathbf{u} is non-decreasing within bin i ; (ii) $\Delta s_i = 0$ with $s_{i_s-1} = s_{i_t} = 1$ and the offline optimal \mathbf{u} is non-increasing within bin i and (iii) $\Delta s_i = 0$ and \mathbf{u} is not monotonic will be covered shortly in the arguments to follow.

Consider a bin with $|\Delta s_i| > 0$ and $\gamma_j^+ = 0$ uniformly. From Lemmas 27 and 28 and using the fact that $|\Delta s_i| \leq 2$ we have,

$$T_{2,i} + T_{3,i} \leq \frac{\beta n_i C_i^2}{2} + \underbrace{\frac{-\lambda^2 (\Delta s_i)^2}{2n_i \beta}}_{(1)} + 7\lambda C_i + \underbrace{\frac{(-\Gamma_i^-)^2}{2n_i \beta} + 2\Gamma_i^- C_i}_{(2)} - \frac{\lambda \Delta s_i \Gamma_i^-}{n_i \beta}.$$

By completing the squares with the terms (1) and (2) in the above display and dropping the negative terms, we get

$$\begin{aligned} T_{2,i} + T_{3,i} &\leq \frac{\beta n_i C_i^2}{2} + \frac{49 C_i^2 n_i \beta}{2 (\Delta s_i)^2} + 2\beta n_i C_i^2 - \frac{\lambda \Delta s_i \Gamma_i^-}{n_i \beta} \\ &\leq 27B^2 \beta - \frac{\lambda \Delta s_i \Gamma_i^-}{n_i \beta}, \end{aligned}$$

where in last line we used the facts that $C_i \leq B/\sqrt{n_i}$ by Corollary 24 and $|\Delta s_i| > 1$ whenever $|\Delta s_i| \neq 0$ by Remark 6.

Define $T_i := \sum_{j=i_s}^{i_t} f_j(x_j) - f_j(u_j)$. Notice that:

- **(A1):** When $\Gamma_i^- = 0$ and $|\Delta s_i| > 0$, combining Lemma 26 we have $T_i = \tilde{O}(1)$;

- **(A2):** Similarly when $\Delta s_i > 0$, we get $T_i = \tilde{O}(1)$ as $\Gamma_i^- \geq 0$ by Lemma 23.

In what follows, we try to split an original bin $[i_s, i_t]$ with $\Delta s_i < 0$ into sub-bins that satisfy the above conditions (A1) or (A2).

If optimal sequence is uniformly constant, we can appeal to the static regret guarantee of FLH to get logarithmic regret over n rounds. So we assume that the optimal sequence is not constant uniformly in the analysis below.

Next, we consider the case when $\Delta s_i < 0$. We start with the following observation.

(B1): Consider a bin $[i_s, i_t]$ that satisfies the structural property in Corollary 24. When either $i_s > 1$ or $i_t < n$ and $\Delta s_i < 0$, then $s_{i_t} \in \{-1, 0\}$ and $s_{i_s-1} \in \{0, 1\}$ with at-least one of them being non-zero.

Since by our assumption $|\mathcal{P}| > 1$, i_s and i_t can't be 1 and n simultaneously. So for any bin $[i_s, i_t]$ with $\Delta s_i < 0$, observation (B1) has to be satisfied.

When $\Delta s_i < 0$, we can have three cases as follows.

Case (1): If the optimal solution is constant (i.e $C_i = 0$) within the bin i . Then we trivially get $T_i = \tilde{O}(1)$.

Case (2): If the optimal solution is monotonic within bin i (see config (a) in Fig.5 for an example of this configuration). Then we split the original bin $[i_s, i_t]$ into at-most 2 bins. Let j_1, j_2 be such that $u_m = -B \forall m \in [i_s, j_1 - 1] \cup [j_2 + 1, i_t]$ and $u_{j_1} > -B, u_{j_2} > -B$. If $u_{i_s} > -B$, then $j_1 = i_s$ and $[i_s, j_1 - 1]$ is viewed as an empty interval. Similar logic applies for the right interval $[j_2 + 1, i_t]$. Since the optimal sequence is monotonic within $[i_s, i_t]$, either $j_1 = i_s$ or $j_2 = i_t$. Without loss of generality let's assume that $j_2 = i_t$. We proceed to bound the regret incurred within each of the two sub-bins separately.

Let's annotate bin $[i_s, j_1 - 1]$ by $i^{(1)}$ and bin $[j_1, i_t]$ by $i^{(2)}$. For the bin $i^{(1)}$, the optimal solution is constant and hence the regret $T_{i^{(1)}} = \tilde{O}(1)$. For the bin $i^{(2)}$, notice that $\gamma_j^- = 0 \forall j \in [j_1, i_t]$ since the sequence is monotonic with $u_{j_1} > -B$ and since our assumption $j_2 = i_t$ implies $u_{j_2} > -B$. Hence we have $\Gamma_{i^{(2)}}^- = 0$. Since $s_{j_1-1} \in \{0, 1\}$ and by observation (B1), $s_{i_t} \in \{-1, 0\}$ with at-least one of them being non-zero, we have $|\Delta s_i| \neq 0$. Hence the bin $i^{(2)}$ falls into the category (A1). So $T_{i^{(2)}} = \tilde{O}(1)$. Adding the regret incurred in each sub-bin separately yields $T_i = \tilde{O}(1)$.

Case (3): Consider the alternate case where we have $\Delta s_i < 0$ and the sequence is not monotonic (see config (b) in Fig. 5 for an example of this configuration). We split the original bin $[i_s, i_t]$ into at-most three sub-bins $[i_s, j_1 - 1], [j_1, j_2], [j_2 + 1, i_t]$ such that (i) If $u_{i_s} = -B$, then $u_m = -B \forall m \in [i_s, j_1 - 1]$ and $u_{j_1} > -B$. If $u_{i_s} > -B$, then we take $j_1 = i_s$ and view $[i_s, j_1 - 1]$ as empty interval. (ii) j_2 is the smallest point in $[j_1, i_t]$ such that $s_{j_2} = -1$ and $u_{j_2} > u_{j_2+1}$.

Let's annotate bins $[i_s, j_1 - 1], [j_1, j_2], [j_2 + 1, i_t]$ by $i^{(1)}, i^{(2)}, i^{(3)}$ respectively. If bin $i^{(1)}$ is not empty, then we have $T_{i^{(1)}} = \tilde{O}(1)$ since \mathbf{u} is constant within that bin.

Since $\Delta s_i < 0$, we must have $s_{j_1-1} \in \{0, 1\}$ even if $j_1 = i_s$. By construction the sequence \mathbf{u} never attains the value $-B$ in the bin $i^{(2)}$ since $u_{j_1} > -B$ and j_2 is the first time point since j_1 after which the optimal sequence jumps downwards. So we have $\Gamma_{i^{(2)}} = 0$. Further we also have $|\Delta s_{i^{(2)}}| > 0$ within bin $i^{(2)}$. So we get $T_{i^{(2)}} = \tilde{O}(1)$ since $i^{(2)}$ falls into category (A1)

For simplicity let's assume that $u_{i_t} > -B$, otherwise we can create another bin that ends at time i_t where optimal solution assumes a constant value of $-B$ and proceed with similar arguments as before to bound the regret in the constant interval.

(S2): If the sequence \mathbf{u} is not monotonic in $i^{(3)}$, we split the bin $i^{(3)}$ into two parts $[j_2 + 1, j_3], [j_3 + 1, i_t]$ such that j_3 is the largest point in $[j_2 + 1, i_t]$ with $s_{j_3} = 1$ and $u_{j_3} < u_{j_3+1}$.

Let's annotate the bins $[j_2 + 1, j_3]$, $[j_3 + 1, i_t]$ by $q^{(1)}$, $q^{(2)}$ respectively. We have $\Delta s_{q^{(1)}} > 0$ since $s_{j_3} = 1$ and $s_{j_2} = -1$. Hence the bins $q^{(1)}$ falls into the category (A2) mentioned before and we get $T_{q^{(1)}} = \tilde{O}(1)$. Notice that $s_{i_t} \in \{-1, 0\}$ as $\Delta s_i < 0$. Since j_3 is the largest point in $[j_2 + 1, i_t]$ with $s_{j_3} = 1$ and it is assumed before that $u_{i_t} > -B$, we conclude that the sequence in the interval $q^{(2)}$ is a non-increasing sequence that never attains the value $-B$. So $\Gamma_{q^{(2)}}^- = 0$. Further we have $|\Delta s_{q^{(2)}}| > 1$. So $T_{q^{(2)}} = \tilde{O}(1)$ since $q^{(2)}$ falls into the category (A1). We pause to remark that the arguments we used to bound the regret in the bin $i^{(3)}$ can be used to bound the regret of any bin $[r_s, r_t] \in \mathcal{P}$ with $\Delta s_r = 0$ and the sequence \mathbf{u} being not monotonic within bin r .

Note that since $u_{j_2+1} < u_{j_2}$, bin $i^{(3)}$ satisfies the structural property of Corollary 24. So if the sequence \mathbf{u} is non-increasing in bin $i^{(3)}$ and $s_{i_t} = -1$, it fits into case (c) of Lemma 28. So we can bound $T_{i^{(3)}} = \tilde{O}(1)$ using arguments presented in (S1).

If the sequence \mathbf{u} is monotonic in bin $i^{(3)}$ and $s_{i_t} = 0$ (which happens when $i_t = n$), then we have $\Delta s_{i^{(3)}} = 0 - (-1) = 1 > 0$. So bin $i^{(3)}$ falls into the category (A2) mentioned before. Hence the regret $T_{i^{(3)}} = \tilde{O}(1)$.

(S3): If the sequence \mathbf{u} is non-decreasing in bin $i^{(3)}$, we split the bin into two intervals $[j_2 + 1, k]$, $[k + 1, i_t]$ such that k is any point in $[j_2 + 1, i_t]$ with $s_k = 1$ and $u_{k+1} > u_k$. (This configuration is similar to that of config (a) in Fig. 1). Annotate $[j_2 + 1, k]$, $[k + 1, i_t]$ by $q^{(1)}$, $q^{(2)}$ respectively. In bin $q^{(1)}$ we have $\Delta s_{q^{(1)}} = 2$ and hence $T_{q^{(1)}} = \tilde{O}(1)$ since $q^{(1)}$ falls into the category (A2). Within bin $q^{(2)}$ due to the assumption that $u_{i_t} > -B$, we have $\Gamma_{q^{(2)}}^- = 0$. We also have $|\Delta s_{q^{(2)}}| > 0$ and consequently $q^{(2)}$ falls into category (A1). So we have $T_{q^{(2)}} = \tilde{O}(1)$. We pause to remark that the arguments we used to bound the regret in bin $i^{(3)}$ for the case where \mathbf{u} is non-decreasing, can also be used to bound the regret of any bin $[r_s, r_t]$ with $\Delta s_r = 0$ and $s_{r_t} = s_{r_s-1} = -1$ and the sequence \mathbf{u} is non-decreasing. The regret for the alternate case where $\Delta s_r = 0$ and $s_{r_t} = s_{r_s-1} = 1$ and the sequence \mathbf{u} is non-increasing can be bounded similarly using a mirrored argument.

So summarizing, in case (3) we get $T_i = \tilde{O}(1)$. Since the intermediate splitting operations can only increase the number of bins to at-most $6M$, adding the regret across all $O(M)$ bins in Corollary 24 yields the Theorem. ■




Figure 5: Examples of configurations referred in the proof of Theorem 22. The blue dots corresponds to the offline optimal sequence.

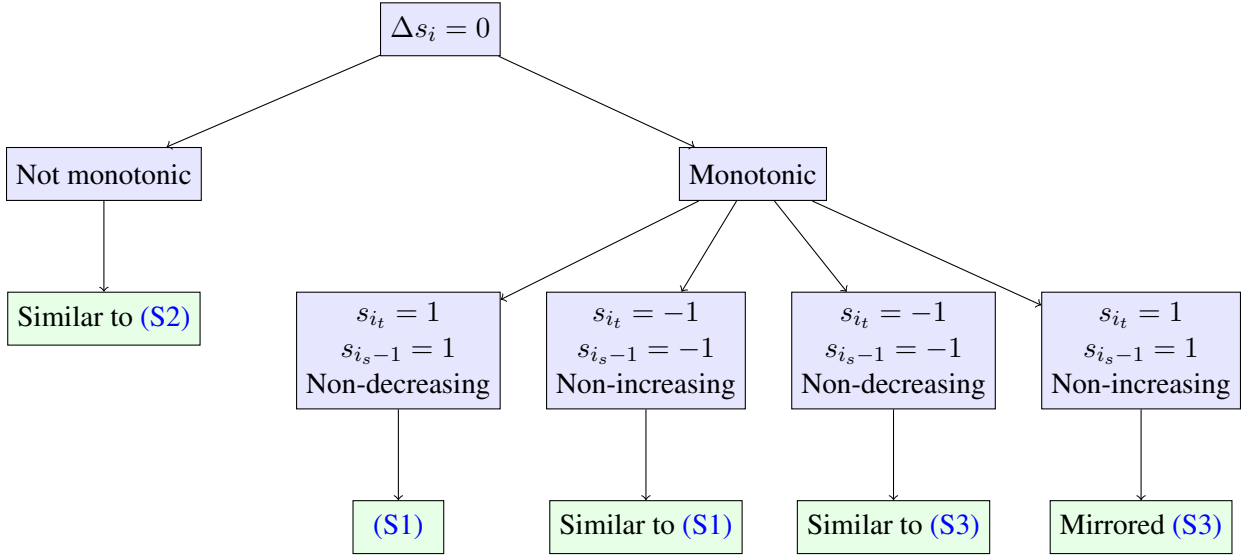


Figure 6: Various configurations of the optimal sequence within a bin $[i_s, i_t]$ with $\Delta s_i = 0$. The leaf nodes indicate the arguments used in the proof of Theorem 22 to handle each scenario.

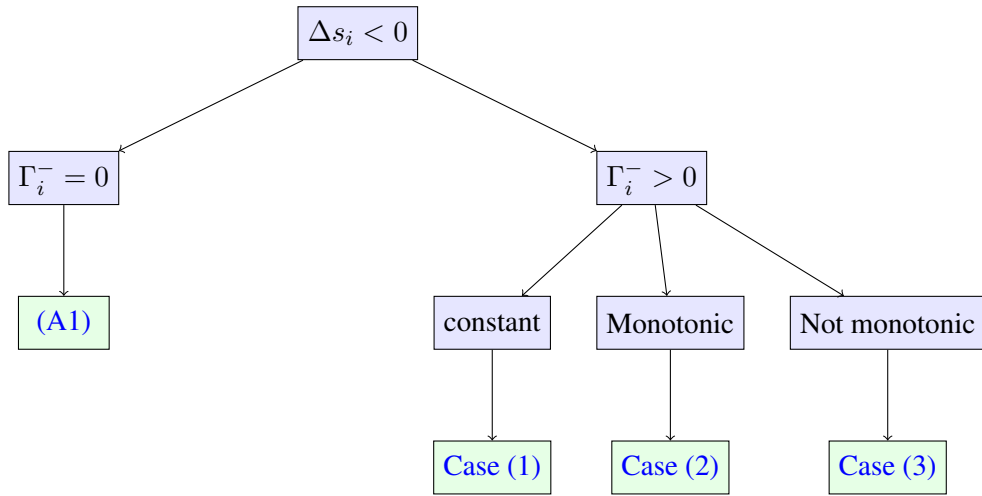


Figure 7: Various configurations of optimal sequence within a bin $[i_s, i_t]$ with $\Delta s_i < 0$. The leaf nodes indicate the arguments used in the proof of Theorem 22 to handle each scenario.




Figure 8: A configuration of optimal sequence within a bin $[i_s, i_t]$ with $\Delta s_i > 0$. The leaf node indicate the arguments used in the proof of Theorem 22 to handle each scenario.

D.2. Multi dimensional setting

We start by inspecting the KKT conditions.

Lemma 29 (*characterization of offline optimal*) Consider the following convex optimization problem.

$$\begin{aligned}
 \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_n, \tilde{\mathbf{z}}_1, \dots, \tilde{\mathbf{z}}_{n-1} \quad & \min \sum_{t=1}^n f_t(\tilde{\mathbf{u}}_t) \\
 \text{s.t.} \quad & \tilde{\mathbf{z}}_t = \tilde{\mathbf{u}}_{t+1} - \tilde{\mathbf{u}}_t \quad \forall t \in [n-1], \\
 & \sum_{t=1}^{n-1} \|\tilde{\mathbf{z}}_t\|_1 \leq C_n, \tag{16a}
 \end{aligned}$$

$$\|\tilde{\mathbf{u}}_t\|_\infty \leq B \quad \forall t \in [n], \tag{16b}$$

Let $\mathbf{u}_1, \dots, \mathbf{u}_n, \mathbf{z}_1, \dots, \mathbf{z}_{n-1} \in \mathbb{R}^d$ be the optimal primal variables and let $\lambda \geq 0$ be the optimal dual variable corresponding to the constraint (16a). Further, let $\gamma_t^+, \gamma_t^- \in \mathbb{R}^d$ with $\gamma_t^+ \geq 0$ and $\gamma_t^- \geq 0$ be the optimal dual variables that correspond to constraint (16b). Specifically for $k \in [d]$, $\gamma_t^+[k]$ corresponds to the dual variable for the constraint $\mathbf{u}_t[k] \leq B$ induced by the relation (16b). Similarly $\gamma_t^-[k]$ corresponds to the constraint $-B \leq \mathbf{u}_t[k]$. By the KKT conditions, we have

- **stationarity:** $\nabla f_t(\mathbf{u}_t) = \lambda(\mathbf{s}_t - \mathbf{s}_{t-1}) + \gamma_t^- - \gamma_t^+$, where $\mathbf{s}_t \in \partial|\mathbf{z}_t|$ (a subgradient). Specifically, $\mathbf{s}_t[k] = \text{sign}(\mathbf{u}_{t+1}[k] - \mathbf{u}_t[k])$ if $|\mathbf{u}_{t+1}[k] - \mathbf{u}_t[k]| > 0$ and $\mathbf{s}_t[k]$ is some value in $[-1, 1]$ otherwise. For convenience of notations later, we also define $\mathbf{s}_n = \mathbf{s}_0 = \mathbf{0}$.
- **complementary slackness:** (a) $\lambda(\sum_{t=2}^n \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_1 - C_n) = 0$; (b) $\gamma_t^-[k](\mathbf{u}_t[k] + B) = 0$ and $\gamma_t^+[k](\mathbf{u}_t[k] - B) = 0$ for all $t \in [n]$ and all $k \in [d]$.

The proof of the above lemma is similar to the 1D case and hence omitted.

Terminology. We will refer to the optimal primal variables $\mathbf{u}_1, \dots, \mathbf{u}_n$ in Lemma 29 as the *offline optimal sequence* in this section.

Next, we claim the existence of a partitioning of $[n]$ with some useful properties.

Lemma 30 (*key partition*) There exist a partitioning \mathcal{P} of $[n]$ into $M = O(dn^{1/3}C_n^{2/3})$ intervals viz $\{[i_s, i_t]\}_{i=1}^M$ such that for any interval $[i_s, i_t] \in \mathcal{P}$, $C_i \leq B/\sqrt{n_i}$ where $C_i := \sum_{j=i_s+1}^{i_t} \|\mathbf{u}_j - \mathbf{u}_{j-1}\|_1$ and n_i is the length of the interval.

Define $\Gamma_i^+ := \sum_{j=i_s}^{i_t} \gamma_j^+$ and $\Gamma_i^- := \sum_{j=i_s}^{i_t} \gamma_j^-$. Let $\Delta \mathbf{s}_i = \mathbf{s}_{i_t} - \mathbf{s}_{i_s-1}$, where \mathbf{s} is as defined in Section 3.2. We also have that each bin $[i_s, i_t] \in \mathcal{P}$ satisfies at-least one of the following properties.

Property 1 Across each coordinate $k \in [d]$, the sequence $\mathbf{u}_j[k]$, $j \in [i_s, i_t]$ is either non-decreasing or non-increasing.

Property 2 $\|\lambda \Delta \mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 \geq \lambda/4$.

The proof of the above lemma is deferred to Section E.

We recall Eq.(7) here for convenience. Let \mathcal{P} be a partition of $[n]$ into M bins obtained in Lemma 30 Let $[i_s, i_t]$ denote the i^{th} bin in \mathcal{P} and let n_i be its length. Define $\bar{\mathbf{u}}_i = \frac{1}{n_i} \sum_{j=i_s}^{i_t} \mathbf{u}_j$ and $\dot{\mathbf{u}}_i = \bar{\mathbf{u}}_i - \frac{1}{n_i \beta} \sum_{j=i_s}^{i_t} \nabla f_j(\bar{\mathbf{u}}_i)$ where β is as in Assumption EC-2. Let \mathbf{x}_j be the prediction made by FLH at time j . We start with following regret decomposition.

$$R_n(C_n) \leq \underbrace{\sum_{i=1}^M \sum_{j=i_s}^{i_t} f_j(\mathbf{x}_j) - f_j(\dot{\mathbf{u}}_i)}_{T_{1,i}} + \underbrace{\sum_{i=1}^M \sum_{j=i_s}^{i_t} f_j(\dot{\mathbf{u}}_i) - f_j(\bar{\mathbf{u}}_i)}_{T_{2,i}} + \underbrace{\sum_{i=1}^M \sum_{j=i_s}^{i_t} f_j(\bar{\mathbf{u}}_i) - f_j(\mathbf{u}_j)}_{T_{3,i}}.$$

Lemma 31 (*bounding $T_{1,i}$*) Let the experts in FLH be the ONS algorithms with parameter $\zeta = \min \left\{ \frac{1}{4G^\dagger(2B\sqrt{d}+2G\sqrt{d})}, \alpha \right\}$ and decision set \mathcal{D} . Also choose learning rate $\eta = \alpha$, for FLH. Then for any bin $[i_s, i_t]$ we have,

$$\begin{aligned} \sum_{j=i_s}^{i_t} f_j(\mathbf{x}_j) - f_j(\dot{\mathbf{u}}_i) &= O \left(d^{3/2} B G^\dagger \log n + d^{3/2} G G^\dagger \log n + \frac{\log n}{\alpha} \right) \\ &= O(d^{3/2} \log n), \end{aligned}$$

where $\mathbf{x}_j \in \mathbb{R}^d$ are the outputs of FLH.

Proof First we proceed to bound $\|\dot{\mathbf{u}}_i\|_\infty$. Since $\|\nabla f_j(\mathbf{u}_j)\|_2 \leq G$ by Assumption EC-1, we have

$$\begin{aligned} \|\dot{\mathbf{u}}_i\|_\infty &\leq \|\bar{\mathbf{u}}_i\|_\infty + \frac{G}{\beta} \\ &\leq B + \frac{G}{\beta} \\ &\leq B + G, \end{aligned}$$

where we used $\beta > 1$ from Assumption EC-2.

For any $\mathbf{x} \in \mathcal{D}$, we have $\|\mathbf{x} - \dot{\mathbf{u}}_i\|_2 \leq 2B\sqrt{d} + 2G\sqrt{d}$ by triangle inequality and the fact $\|\mathbf{y}\|_2 \leq \sqrt{d} \|\mathbf{y}\|_\infty$.

By Assumption EC-4 we have $\|\nabla f_j(\mathbf{x})\|_2 \leq G^\dagger$ for any $\mathbf{x} \in \mathcal{D}$. Also, recall that by Assumption EC-3, the loss functions f_j are α exp-concave in the domain \mathcal{D} . Let \mathbf{p}_j , $j \in [i_s, i_t]$ be the predictions of an ONS algorithm when run in the interval $[i_s, i_t]$. If we choose $\zeta = \min \left\{ \frac{1}{4G^\dagger(2B\sqrt{d}+2G\sqrt{d})}, \alpha \right\}$ as the parameter of the ONS, Theorem 2 of (Hazan et al., 2007) implies that

$$\begin{aligned} \sum_{j=i_s}^{i_t} f_j(\mathbf{p}_j) - f_j(\dot{\mathbf{u}}_i) &= O \left(d^{3/2} B G^\dagger \log n + d^{3/2} G G^\dagger \log n + \frac{\log n}{\alpha} \right) \\ &= O \left(d^{3/2} \log n \right). \end{aligned}$$

Now the Lemma is implied by the SA regret bound of FLH (Theorem 3.2 of (Hazan and Seshadhri, 2007)). \blacksquare

For strongly convex, losses the term $T_{1,i}$ can enjoy a better bound.

Lemma 32 (bounding $T_{1,i}$ for strongly convex losses) *Suppose that the losses are H strongly convex. Take experts in FLH as OGD with step size $1/(Hn)$ and decision set \mathcal{D} . Also choose learning rate $\eta = H/(G^\dagger)^2$, for FLH. Then for any bin $[i_s, i_t]$ we have,*

$$\sum_{j=i_s}^{i_t} f_j(\mathbf{x}_j) - f_j(\dot{\mathbf{u}}_i) = O\left(\frac{(G^\dagger)^2 \log n}{H}\right),$$

where $\mathbf{x}_j \in \mathbb{R}^d$ are the outputs of FLH.

Proof [Proof Sketch] The lemma follows by using the regret bound of OGD with strongly convex losses from (Hazan et al., 2007) and following similar lines of arguments as in Lemma 31. \blacksquare

We state the next lemma to be generically valid for any bin which is not necessarily a member of \mathcal{P} .

Some notations. For a bin $[a, b]$, introduce the notations $\Delta \mathbf{s}_{a \rightarrow b} := s(\mathbf{u}_{b+1} - \mathbf{u}_b) - s(\mathbf{u}_a - \mathbf{u}_{a-1})$, $\Gamma_{a \rightarrow b}^+ = \sum_{j=a}^b \gamma_j^+$ and $\Gamma_{a \rightarrow b}^- := \sum_{j=a}^b \gamma_j^-$. $n_{a \rightarrow b} := b - a + 1$. $\bar{\mathbf{u}}_{a \rightarrow b} = \frac{1}{n_{a \rightarrow b}} \sum_{j=a}^b \mathbf{u}_j$ and $\dot{\mathbf{u}}_{a \rightarrow b} = \bar{\mathbf{u}}_{a \rightarrow b} - \frac{1}{\beta n_{a \rightarrow b}} \sum_{j=a}^b \nabla f_j(\bar{\mathbf{u}}_{a \rightarrow b})$.

Lemma 33 *For any bin $[a, b]$, we have*

$$\begin{aligned} T_{2,[a,b]} &:= \sum_{j=a}^b f_j(\dot{\mathbf{u}}_{a \rightarrow b}) - f_j(\bar{\mathbf{u}}_i) \\ &\leq \frac{-\|\lambda \Delta \mathbf{s}_{a \rightarrow b} + \Gamma_{a \rightarrow b}^- - \Gamma_{a \rightarrow b}^+\|_2^2}{2n_{a \rightarrow b} \beta} + \|\lambda \Delta \mathbf{s}_{a \rightarrow b} + \Gamma_{a \rightarrow b}^- - \Gamma_{a \rightarrow b}^+\|_1 C_{a \rightarrow b}. \end{aligned}$$

Proof Let $g(\mathbf{x})$ be a α -strongly smooth function. Let $\mathbf{x}^+ = \mathbf{x} - \mu \nabla f(\mathbf{x})$ for some $\mu > 0$. Then we have

$$\begin{aligned} g(\mathbf{x}^+) - g(\mathbf{x}) &\leq \|\nabla g(\mathbf{x})\|_2^2 \left(\frac{\alpha}{2} \mu^2 - \mu \right) \\ &= \frac{-\|\nabla g(\mathbf{x})\|_2^2}{2\alpha}, \end{aligned}$$

by choosing $\mu = 1/\alpha$. By taking $g(x) = \sum_{j=a}^b f_j(\mathbf{x})$ and noting that g is $n_i\beta$ gradient Lipschitz due to Assumption SC-2, we get

$$\begin{aligned}
 T_{2,[a,b]} &:= \sum_{j=a}^b f_j(\hat{\mathbf{u}}_{a \rightarrow b}) - f_j(\bar{\mathbf{u}}_{a \rightarrow b}) \\
 &\leq \frac{-\left\| \sum_{j=a}^b \nabla f_j(\bar{\mathbf{u}}_{a \rightarrow b}) \right\|_2^2}{2n_{a \rightarrow b}\beta} \\
 &= \frac{-1}{2n_{a \rightarrow b}\beta} \left\| \sum_{j=a}^b \nabla f_j(\mathbf{u}_j) + \nabla f_j(\bar{\mathbf{u}}_{a \rightarrow b}) - \nabla f_j(\mathbf{u}_j) \right\|_2^2 \\
 &\leq \frac{-1}{2n_{a \rightarrow b}\beta} \left\| \sum_{j=a}^b \nabla f_j(\mathbf{u}_j) \right\|_2^2 + \frac{1}{n_{a \rightarrow b}\beta} \left\| \sum_{j=a}^b \nabla f_j(\mathbf{u}_j) \right\|_1 \left\| \sum_{j=a}^b \nabla f_j(\bar{\mathbf{u}}_{a \rightarrow b}) - \nabla f_j(\mathbf{u}_j) \right\|_2,
 \end{aligned}$$

where we used $\langle \mathbf{x}, \mathbf{y} \rangle \leq \|\mathbf{x}\|_2 \|\mathbf{y}\|_2 \leq \|\mathbf{x}\|_1 \|\mathbf{y}\|_2$ and dropped a negative term from expanding the squared norm. From the KKT conditions in Lemma 29 we have $\sum_{j=a}^b \nabla f_j(\mathbf{u}_j) = \lambda \Delta \mathbf{s}_{a \rightarrow b} + \Gamma_{a \rightarrow b}^- - \Gamma_{a \rightarrow b}^+$. Since f_j are β -gradient Lipschitz and $\|\bar{\mathbf{u}}_{a \rightarrow b} - \mathbf{u}_j\|_2 \leq \|\bar{\mathbf{u}}_{a \rightarrow b} - \mathbf{u}_j\|_1 \leq C_{a \rightarrow b}$, we also have

$$\left\| \sum_{j=a}^b \nabla f_j(\bar{\mathbf{u}}_{a \rightarrow b}) - \nabla f_j(\mathbf{u}_j) \right\|_2 \leq n_{a \rightarrow b} \beta C_{a \rightarrow b}.$$

Substituting these we get the statement of the lemma. \blacksquare

Lemma 34 For any bin $[i_s, i_t] \in \mathcal{P}$, we have

$$\sum_{j=i_s}^{i_t} f_j(\bar{\mathbf{u}}_i) - f_j(\mathbf{u}_j) \leq \frac{\beta n_i C_i^2}{2} + 5\lambda C_i + \|\lambda \Delta \mathbf{s}_i + \Gamma_i^- - \Gamma_i^+\|_2 C_i.$$

Proof Due to strong smoothness, we have

$$\begin{aligned}
 T_{3,i} &:= \sum_{j=i_s}^{i_t} f_j(\bar{\mathbf{u}}_i) - f_j(\mathbf{u}_j) \\
 &\leq_{(a)} \sum_{j=i_s}^{i_t} \langle \nabla f_j(\mathbf{u}_j), \bar{\mathbf{u}}_i - \mathbf{u}_j \rangle + \frac{\beta}{2} \|\bar{\mathbf{u}}_i - \mathbf{u}_j\|_1^2 \\
 &\leq \frac{\beta n_i C_i^2}{2} + \sum_{j=i_s}^{i_t} \langle \nabla f_j(\mathbf{u}_j), \bar{\mathbf{u}}_i - \mathbf{u}_j \rangle,
 \end{aligned} \tag{17}$$

where in line (a) we used $\|\bar{\mathbf{u}}_i - \mathbf{u}_j\|_2 \leq \|\bar{\mathbf{u}}_i - \mathbf{u}_j\|_1$.

Further,

$$\begin{aligned} \sum_{j=i_s}^{i_t} \langle \nabla f_j(\mathbf{u}_j), \bar{\mathbf{u}}_i - \mathbf{u}_j \rangle &= \lambda (\langle \mathbf{s}_{i_s-1}, \mathbf{u}_{i_s} - \bar{\mathbf{u}}_i \rangle - \langle \mathbf{s}_{i_t}, \mathbf{u}_{i_t} - \bar{\mathbf{u}}_i \rangle) \\ &\quad + \lambda \sum_{j=i_s+1}^{i_t} \|\mathbf{u}_j - \mathbf{u}_{j-1}\|_1 + \sum_{j=i_s}^{i_t} \langle \gamma_j^- - \gamma_j^+, \bar{\mathbf{u}}_i - \mathbf{u}_j \rangle \end{aligned}$$

By triangle and Holder's inequalities, the first two terms can be bounded by $3\lambda C_i$ (recall that $\|\mathbf{u}_t - \bar{\mathbf{u}}_i\|_1 \leq C_i$ for all $t \in [i_s, i_t]$). Let's proceed to bound the last term in the above display. From Lemma 30, we have $C_i \leq B/\sqrt{n_i}$. So the TV incurred across each coordinate of the optimal solution is at-most B . Using similar arguments as in Lemma 25, the complementary slackness in Lemma 29 implies that for each $k \in [d]$, if $\gamma_j^- [k] > 0$ for at-least one $j \in [i_s, i_t]$ then $\gamma_j^+ [k] = 0$ for all $j \in [i_s, i_t]$. Similarly for each $k \in [d]$, if $\gamma_j^+ [k] > 0$ for at-least one $j \in [i_s, i_t]$ then $\gamma_j^- [k] = 0$ for all $j \in [i_s, i_t]$. This observation allows us to write,

$$\sum_{j=i_s}^{i_t} |\gamma_j^- [k] - \gamma_j^+ [k]| = |\Gamma_i^- [k] - \Gamma_i^+ [k]| \quad (18)$$

Define $C_i^k := \sum_{j=i_s+1}^{i_t} \|\mathbf{u}_j[k] - \mathbf{u}_{j-1}[k]\|$. We have,

$$\begin{aligned} \sum_{j=i_s}^{i_t} \langle \gamma_j^- - \gamma_j^+, \bar{\mathbf{u}}_i - \mathbf{u}_j \rangle &= \sum_{k=1}^d \sum_{j=i_s}^{i_t} (\gamma_j^- [k] - \gamma_j^+ [k]) (\bar{\mathbf{u}}_i[k] - \mathbf{u}_j[k]) \\ &\leq_{(a)} \sum_{k=1}^d |\Gamma_i^- [k] - \Gamma_i^+ [k]| C_i^k \\ &= \sum_{k=1}^d (\lambda \Delta \mathbf{s}_i[k] \text{sign}(\Gamma_i^- [k] - \Gamma_i^+ [k]) + \text{sign}(\Gamma_i^- [k] - \Gamma_i^+ [k]) (\Gamma_i^- [k] - \Gamma_i^+ [k])) C_i^k \\ &\quad - \sum_{k=1}^d \lambda \Delta \mathbf{s}_i[k] \text{sign}(\Gamma_i^- [k] - \Gamma_i^+ [k]) C_i^k \\ &\leq_{(b)} \|\lambda \Delta \mathbf{s}_i + \Gamma_i^- - \Gamma_i^+\|_2 C_i + 2\lambda C_i, \end{aligned}$$

where in line (a) we applied $\langle \mathbf{x}, \mathbf{y} \rangle \leq \|\mathbf{x}\|_1 \|\mathbf{y}\|_\infty$ along with the Eq. (18). In line (b) we applied $\langle \mathbf{x}, \mathbf{y} \rangle \leq \|\mathbf{x}\|_2 \|\mathbf{y}\|_1$ for the first term and $\langle \mathbf{x}, \mathbf{y} \rangle \leq \|\mathbf{x}\|_\infty \|\mathbf{y}\|_1$ for the second term. Putting everything together yields the Lemma. \blacksquare

Lemma 35 *Let `splitMonotonic` in Fig.9 be run with an input $[i_s, i_t]$. Then the partition \mathcal{S} it return obeys $|\mathcal{S}| = O(d)$.*

Proof From the psuedo-code in Fig. 9 it is obvious that each coordinate can contribute to increasing the bin count by $O(1)$. Hence the overall bin count in \mathcal{S} is $O(d)$. \blacksquare

An illustrative example of the input and output of `splitMonotonic` is given in Fig. 10.

`splitMonotonic`: Inputs - (1) an interval $[i_s, i_t]$ such that the offline optimal is monotonic across each coordinate $k \in [d]$; (2) offline optimal sequences $\mathbf{u}_{1:n}$ and the sequence of subgradients (dual variables) $\mathbf{s}_{1:n-1}$ (recall that $s_0 = s_n = 0$ by convention.).

1. Initialize $\mathcal{T} \leftarrow \Phi, \mathcal{S} \leftarrow \Phi$.
2. Add i_s, i_t to \mathcal{T} .
3. For each coordinate $k \in [d]$:
 - (a) If $\mathbf{u}[k]$ is constant in $[i_s, i_t]$, then skip the current coordinate.
 - (b) Initialize $z_1 \leftarrow i_s, z_2 \leftarrow i_t$.
 - (c) If $\mathbf{u}_{i_s}[k] = \pm B$, let z_1 be the first time point in $[i_s, i_t]$ where $\mathbf{u}_{z_1}[k] \neq \pm B$. Add $z_1 - 1, z_1$ to \mathcal{T} .
 - (d) If $\mathbf{u}_{i_t}[k] = \pm B$, let z_2 be the last time point in $[i_s, i_t]$ where $\mathbf{u}_{z_2}[k] \neq \pm B$. Add $z_2, z_2 + 1$ to \mathcal{T} .
 - (e) If $\mathbf{u}[k]$ is non-decreasing in $[i_s, i_t]$ then let $p \geq z_1$ be the first point with $\mathbf{s}_{p-1}[k] = 1$. If $p > z_1$, add $p - 1, p$ to \mathcal{T} .
 - (f) If $\mathbf{u}[k]$ is non-decreasing in $[i_s, i_t]$ then let $q \leq z_2$ be the last point with $\mathbf{s}_q[k] = 1$. If $q < z_2$, add $q, q + 1$ to \mathcal{T} .
 - (g) If $\mathbf{u}[k]$ is non-increasing in $[i_s, i_t]$ then let $p \geq z_1$ be the first point with $\mathbf{s}_{p-1}[k] = -1$. If $p > z_1$, add $p - 1, p$ to \mathcal{T} .
 - (h) If $\mathbf{u}[k]$ is non-increasing in $[i_s, i_t]$ then let $q \leq z_2$ be the last point with $\mathbf{s}_q[k] = -1$. If $q < z_2$, add $q, q + 1$ to \mathcal{T} .
4. For each entry t in \mathcal{T} :
 - (a) If t appears more than 2 times, delete some occurrences of t such that t only appears 2 times in \mathcal{T} .
5. Sort \mathcal{T} in non-decreasing order. For each consecutive points $s, t \in \mathcal{T}$, add $[s, t]$ to \mathcal{S} .
6. Return the partition \mathcal{S} .

Figure 9: `splitMonotonic` procedure




Figure 10: An example of the partitioning created by `splitMonotonic` (See Fig. 9). The partition \mathcal{S} returned by `splitMonotonic` is $\{[i_s, a_1 - 1], [a_1, a_2 - 1], [a_2, a_3 - 1], [a_3, i_t]\}$. Blue dots indicate the offline optimal sequence.

Theorem 14 By using the base learner as ONS with parameter $\zeta = \min \left\{ \frac{1}{4G^\dagger(2B\sqrt{d}+2G/\beta)}, \alpha \right\}$, decision set \mathcal{D} and choosing learning rate $\eta = \alpha$, FLH obeys $R_n^+(C_n) = \tilde{O} \left(d^{3.5} (n^{1/3} C_n^{2/3} \vee 1) \right)$ if $C_n > 1/n$ and $O(d^{1.5} \log n)$ otherwise. Here $a \vee b := \max\{a, b\}$ and $\tilde{O}(\cdot)$ hides dependence on the constants B, G, G^\dagger, α and factors of $\log n$.

Proof Consider a bin $[i_s, i_t] \in \mathcal{P}$. By Lemma 30, the bin has to satisfy one of the two Properties. Let's first focus on the scenario where $[i_s, i_t]$ satisfies Property 2.

Combining the results of Lemmas 33, 34 we can write,

$$\begin{aligned}
 T_{2,i} + T_{3,i} &\leq \frac{-\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2^2}{2n_i\beta} + \|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 C_i \\
 &\quad + \frac{\beta n_i C_i^2}{2} + 5\lambda C_i + \|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 C_i \\
 &\stackrel{(a)}{\leq} \frac{\beta B^2}{2} - \frac{\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2^2}{2n_i\beta} + 7(\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 \vee \lambda) C_i \\
 &= \frac{\beta B^2}{2} - \left(\frac{\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2}{\sqrt{2n_i\beta}} - \frac{7C_i\sqrt{n_i\beta}(\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 \vee \lambda)}{\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2\sqrt{2}} \right)^2 \\
 &\quad + \frac{49n_i\beta C_i^2}{2} \left(\frac{\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 \vee \lambda}{\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2} \right)^2 \\
 &\stackrel{(b)}{\leq} \frac{\beta B^2}{2} + 392\beta n_i C_i^2 \\
 &\leq 393\beta B^2,
 \end{aligned}$$

where in line (a) we used $C_i \leq B/\sqrt{n_i}$ for partitions in \mathcal{P} (Lemma 30). In line (b) we used $\frac{\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 \vee \lambda}{\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2} \geq 4$ since $\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 \geq \lambda/4$ by Property 2 of Lemma 30.

Now using Lemma 31, for the bins $[i_s, i_t]$ that satisfy property 2, we can write

$$T_{1,i} + T_{2,i} + T_{3,i} = \tilde{O}(d^{1.5}). \quad (19)$$

Now suppose that the bin $[t, \bar{t}]$ satisfies Property 1 in Lemma 30. In this case, via a call to `splitMonotonic` function with the input interval as $[t, \bar{t}]$, we split the original bin into $O(d)$ sub-bins (see Lemma 35). Further for a fixed k , if $\mathbf{u}_j[k]$, $j \in [t, \bar{t}]$ is non-decreasing, then we can group those consecutive sub-bins into at-most three categories: (a) a section of time where $\mathbf{u}_j[k]$ is constant; (b) a section of time where $\mathbf{u}_j[k]$ is non-decreasing; (c) a section of time where $\mathbf{u}_j[k]$ is constant.

We proceed to define these sections formally (where p, m, q are indices defined for convenience)

- For section (a) let $\mathcal{A} = \{[t, t-p-1], [t-p, \bar{t}-p], \dots, [t_0, \bar{t}_0]\}$
- For section (b) let $\mathcal{B} = \{[t_1, \bar{t}_1], \dots, [t_m, \bar{t}_m]\}$
- For section (c) let $\mathcal{C} = \{[t_{m+1}, \bar{t}_{m+1}], \dots, [\bar{t}_q + 1, \bar{t}]\}$

As mentioned before, these sections are constructed so that the offline optimal satisfy the following properties.

- (i) $\mathbf{u}_j[k]$ $j \in [t, \bar{t}_0]$ is constant.
- (ii) $\mathbf{u}_{t_1}[k] > \mathbf{u}_{t_1-1}[k]$ and $\mathbf{u}_{t_{m+1}}[k] > \mathbf{u}_{t_{m+1}-1}[k]$.
- (iii) $\mathbf{u}_j[k]$ $j \in [t_1, \bar{t}_m]$ is non-decreasing.

(iv) $\mathbf{u}_j[k] j \in [\underline{t}_{m+1}, \bar{t}]$ is constant.

We remark that the grouping may be different for different coordinates k . Further some of \mathcal{A}, \mathcal{B} or \mathcal{C} can be empty. In the example we gave in Fig. 10:

- For coordinate 1 $\mathcal{A} = \phi, \mathcal{B} = \phi, \mathcal{C} = \{[i_s, a_1 - 1], [a_1, a_2 - 1], [a_2, a_3 - 1], [a_3, i_t]\}$.
- For coordinate 2 $\mathcal{A} = [i_s, a_1 - 1], \mathcal{B} = \phi, \mathcal{C} = \{[a_1, a_2 - 1], [a_2, a_3 - 1], [a_3, i_t]\}$.
- For coordinate 3 $\mathcal{A} = \{[i_s, a_1 - 1], [a_1, a_2 - 1]\}, \mathcal{B} = \{[a_2, a_3 - 1]\}, \mathcal{C} = \{[a_3, i_t]\}$

We will focus on the aforementioned scenario where $\mathbf{u}_j[k], j \in [\underline{t}, \bar{t}]$ is non-decreasing. The arguments for the case where $\mathbf{u}_j[k], j \in [\underline{t}, \bar{t}]$ is non-increasing are similar. Further similar to the proof of Theorem 22, we give arguments for the case where $\gamma_j^+[k] = 0$ for all j in the interval $[\underline{t}, \bar{t}]$ stating that arguments for the case $\gamma_j^-[k] = 0$ uniformly in $[\underline{t}, \bar{t}]$ are similar.

From Lemma 33, we have

$$\sum_{j=a}^b f_j(\hat{\mathbf{u}}_{a \rightarrow b}) - f_j(\bar{\mathbf{u}}_{a \rightarrow b}) \leq \frac{-\|\lambda \Delta \mathbf{s}_{a \rightarrow b} + \mathbf{\Gamma}_{a \rightarrow b}^- - \mathbf{\Gamma}_{a \rightarrow b}^+\|_2^2}{2n_{a \rightarrow b} \beta} + \|\lambda \Delta \mathbf{s}_{a \rightarrow b} + \mathbf{\Gamma}_{a \rightarrow b}^- - \mathbf{\Gamma}_{a \rightarrow b}^+\|_1 C_{a \rightarrow b}. \quad (20)$$

Observe that the relation in Eq. (17) holds for any generic bin $[a, b]$ that may not be a member of \mathcal{P} (replacing $C_i, n_i, \bar{\mathbf{u}}_i$ with $C_{a \rightarrow b}, n_{a \rightarrow b}, \bar{\mathbf{u}}_{a \rightarrow b}$). So

$$T_{3,[a,b]} := \sum_{j=a}^b f_j(\bar{\mathbf{u}}_{a \rightarrow b}) - f_j(\mathbf{u}_j) \leq \sum_{k=1}^d \frac{\beta n_{a \rightarrow b} C_{a \rightarrow b}^2}{2d} + \sum_{j=a}^b \langle \nabla f_j(\mathbf{u}_j), \bar{\mathbf{u}}_{a \rightarrow b} - \mathbf{u}_j \rangle. \quad (21)$$

Note that Eq. (20) and (21) decompose coordinate-wise. So for the bin $[\underline{t}, \bar{t}] \in \mathcal{P}$ where the optimal sequence is monotonic across each coordinate, our strategy is to bound

$$\begin{aligned} \mathbf{S}_{a \rightarrow b}[k] := & \frac{-(\lambda \Delta \mathbf{s}_{a \rightarrow b}[k] + \mathbf{\Gamma}_{a \rightarrow b}^-[k] - \mathbf{\Gamma}_{a \rightarrow b}^+[k])^2}{2n_{a \rightarrow b} \beta} + |\lambda \Delta \mathbf{s}_{a \rightarrow b}[k] + \mathbf{\Gamma}_{a \rightarrow b}^-[k] - \mathbf{\Gamma}_{a \rightarrow b}^+[k]| C_{a \rightarrow b} \\ & + \frac{\beta n_{a \rightarrow b} C_{a \rightarrow b}^2}{2d} + \sum_{j=a}^b \nabla f_j(\mathbf{u}_j)[k] (\bar{\mathbf{u}}_{a \rightarrow b}[k] - \mathbf{u}_j[k]), \end{aligned} \quad (22)$$

for each $k \in [d]$ and $[a, b] \in \mathcal{A} \cup \mathcal{B} \cup \mathcal{C}$ and finally adding them across all coordinates to bound $\sum_{k=1}^d \mathbf{S}_{a \rightarrow b}[k]$. Doing so will result in a bound on $T_{2,[a,b]} + T_{3,[a,b]}$. Further, $T_{1,[a,b]}$ can be bound by strongly adaptive regret. This enables us to bound $\sum_{[a,b] \in \mathcal{A} \cup \mathcal{B} \cup \mathcal{C}} T_{1,[a,b]} + T_{2,[a,b]} + T_{3,[a,b]}$ thereby leading to a regret bound in the parent bin $[\underline{t}, \bar{t}] \in \mathcal{P}$ which was the input interval for the call to `splitMonotonic` that we started with.

Let $C_{[a,b]}[k]$ be the TV of offline optimal incurred in the interval any interval $[a, b]$ along coordinate k . First we focus on the bins in \mathcal{B} . If \mathcal{B} is not empty, then $\gamma_j[k] = 0 \forall j \in [\underline{t}_1 \rightarrow \bar{t}_m]$ due to property (ii) and (iii) above. By using the stationarity conditions in Lemma 29, we can write

$$\begin{aligned} \sum_{[a,b] \in \mathcal{B}} \sum_{j=a}^b \nabla f_j(\mathbf{u}_j)[k] (\bar{\mathbf{u}}_{a \rightarrow b}[k] - \mathbf{u}_j[k]) &= \lambda C_{\underline{t}_1 \rightarrow \bar{t}_m}[k] + \lambda (s_{\underline{t}_1-1}[k] \mathbf{u}_{\underline{t}_1}[k] - s_{\bar{t}_m}[k] \mathbf{u}_{\bar{t}_m}[k]) \\ &+ \sum_{[a,b] \in \mathcal{B}} \lambda \bar{\mathbf{u}}_{a \rightarrow b}[k] \Delta \mathbf{s}_{a \rightarrow b}[k]. \end{aligned}$$

So we have,

$$\begin{aligned}
 \sum_{[a,b] \in \mathcal{B}} \mathbf{S}_{a \rightarrow b}[k] &\leq \frac{\beta n_{\underline{t}_1 \rightarrow \bar{t}_m} C_{\underline{t}_1 \rightarrow \bar{t}_m}^2}{2d} + \lambda C_{\underline{t}_1 \rightarrow \bar{t}_m}[k] + \lambda (s_{\underline{t}_1-1}[k] \mathbf{u}_{\underline{t}_1}[k] - s_{\bar{t}_m}[k] \mathbf{u}_{\bar{t}_m}[k]) \\
 &\quad + \sum_{[a,b] \in \mathcal{B}} \frac{-\lambda^2 (\Delta \mathbf{s}_{a \rightarrow b}[k])^2}{2n_{a \rightarrow b} \beta} + \lambda |\Delta \mathbf{s}_{a \rightarrow b}[k]| C_{\underline{t}_1 \rightarrow \bar{t}_m} + \lambda \bar{\mathbf{u}}_{a \rightarrow b}[k] \Delta \mathbf{s}_{a \rightarrow b}[k] \\
 &\leq_{(a)} \frac{\beta n_{\underline{t} \rightarrow \bar{t}} C_{\underline{t} \rightarrow \bar{t}}^2}{2d} + \lambda C_{\underline{t} \rightarrow \bar{t}}[k] + \lambda (s_{\underline{t}_1-1}[k] \mathbf{u}_{\underline{t}_1}[k] - s_{\bar{t}_m}[k] \mathbf{u}_{\bar{t}_m}[k]) \\
 &\quad + \sum_{[a,b] \in \mathcal{B}} \frac{-\lambda^2 (\Delta \mathbf{s}_{a \rightarrow b}[k])^2}{2n_{a \rightarrow b} \beta} + \lambda |\Delta \mathbf{s}_{a \rightarrow b}[k]| C_{\underline{t} \rightarrow \bar{t}} + \lambda \bar{\mathbf{u}}_{a \rightarrow b}[k] \Delta \mathbf{s}_{a \rightarrow b}[k] \\
 &\leq_{(b)} \frac{\beta B^2}{2d} + \sum_{[a,b] \in \mathcal{B}} \frac{-\lambda^2 (\Delta \mathbf{s}_{a \rightarrow b}[k])^2}{2n_{a \rightarrow b} \beta} + \lambda |\Delta \mathbf{s}_{a \rightarrow b}[k]| C_{\underline{t} \rightarrow \bar{t}} + \lambda \bar{\mathbf{u}}_{a \rightarrow b}[k] \Delta \mathbf{s}_{a \rightarrow b}[k],
 \end{aligned}$$

where in line (a) we used the fact that $C_{\underline{t}_1 \rightarrow \bar{t}_m} \leq C_{\underline{t} \rightarrow \bar{t}}$ and $n_{\underline{t}_1 \rightarrow \bar{t}_m} \leq n_{\underline{t} \rightarrow \bar{t}}$ since $[\underline{t}_1, \bar{t}_m]$ is contained within $[\underline{t}, \bar{t}]$. In line (b) we used $C_{\underline{t} \rightarrow \bar{t}} \leq B/\sqrt{n_{\underline{t} \rightarrow \bar{t}}}$ (since $[\underline{t}, \bar{t}] \in \mathcal{P}$) along with the fact that

$$\begin{aligned}
 \lambda (s_{\underline{t}_1-1}[k] \mathbf{u}_{\underline{t}_1}[k] - s_{\bar{t}_m}[k] \mathbf{u}_{\bar{t}_m}[k]) &= -\lambda C_{\underline{t}_1 \rightarrow \bar{t}_m} \text{ since} \\
 s_{\underline{t}_1-1}[k] = s_{\bar{t}_m} &= 1 \text{ due to property (ii) and (iii) above.}
 \end{aligned}$$

Define $\check{\mathbf{u}}_{\mathcal{B}} := \frac{1}{|\mathcal{B}|} \sum_{[a,b] \in \mathcal{B}} \bar{\mathbf{u}}_{a \rightarrow b}$. Observe that since $s_{\underline{t}_1-1}[k] = s_{\bar{t}_m}[k] = 1$, we can write $\sum_{[a,b] \in \mathcal{B}} \Delta \mathbf{s}_{a \rightarrow b}[k] = 0$ by the telescoping structure.

By noting that we can subtract $0 = \lambda \check{\mathbf{u}}_{\mathcal{B}}[k] \sum_{(a,b) \in \mathcal{B}} \Delta \mathbf{s}_{a \rightarrow b}[k]$ and that $|\bar{\mathbf{u}}_{a \rightarrow b}[k] - \check{\mathbf{u}}_{\mathcal{B}}[k]| \leq C_{\underline{t} \rightarrow \bar{t}}$, we have

$$\begin{aligned}
 \sum_{[a,b] \in \mathcal{B}} \mathbf{S}_{a \rightarrow b}[k] &\leq \frac{\beta B^2}{2d} + \sum_{[a,b] \in \mathcal{B}} \frac{-\lambda^2 (\Delta \mathbf{s}_{a \rightarrow b}[k])^2}{2n_{a \rightarrow b} \beta} + \lambda |\Delta \mathbf{s}_{a \rightarrow b}[k]| C_{\underline{t} \rightarrow \bar{t}} + \lambda (\bar{\mathbf{u}}_{a \rightarrow b}[k] - \check{\mathbf{u}}_{\mathcal{B}}[k]) \Delta \mathbf{s}_{a \rightarrow b}[k] \\
 &\leq \frac{\beta B^2}{2d} + \sum_{[a,b] \in \mathcal{B}} \frac{-\lambda^2 (\Delta \mathbf{s}_{a \rightarrow b}[k])^2}{2n_{a \rightarrow b} \beta} + 2\lambda |\Delta \mathbf{s}_{a \rightarrow b}[k]| C_{\underline{t} \rightarrow \bar{t}} \\
 &= \frac{\beta B^2}{2d} + \sum_{[a,b] \in \mathcal{B}} - \left(\frac{\lambda \Delta \mathbf{s}_{a \rightarrow b}[k]}{\sqrt{2n_{a \rightarrow b} \beta}} - C_{\underline{t} \rightarrow \bar{t}} \sqrt{2n_{a \rightarrow b} \beta} \right)^2 + 2\beta n_{a \rightarrow b} C_{\underline{t} \rightarrow \bar{t}}^2 \\
 &\leq \frac{\beta B^2}{2d} + 2\beta n_{\underline{t} \rightarrow \bar{t}} C_{\underline{t} \rightarrow \bar{t}}^2 \\
 &\leq \frac{\beta B^2}{2d} + 2\beta B^2 \\
 &\leq 3\beta B^2.
 \end{aligned}$$

Next, we address bins present in \mathcal{A} and \mathcal{C} . We provide the arguments for bounding $\sum_{[a,b] \in \mathcal{A}} \mathbf{S}_{a \rightarrow b}[k]$. Bounding the sum for bins in \mathcal{C} can be done using similar arguments.

Observe that by property (i) above, the sequence $\mathbf{u}_j[k]$ for $j \in [i_s, \bar{t}_0]$ is a constant. So the last term in Eq. (22) is zero for any $\mathbf{S}_{a \rightarrow b}[k]$ where $[a, b] \in \mathcal{A}$. Now proceeding similar to above by completing the squares and dropping the negative terms, we get

$$\begin{aligned}
 \sum_{[a,b] \in \mathcal{A}} \mathbf{S}_{a \rightarrow b}[k] &\leq \sum_{[a,b] \in \mathcal{A}} \left(\frac{-\left(\lambda \Delta \mathbf{s}_{a \rightarrow b}[k] + \mathbf{\Gamma}_{a \rightarrow b}^-[k] - \mathbf{\Gamma}_{a \rightarrow b}^+[k]\right)^2}{2n_{a \rightarrow b} \beta} \right. \\
 &\quad \left. + \left| \lambda \Delta \mathbf{s}_{a \rightarrow b}[k] + \mathbf{\Gamma}_{a \rightarrow b}^-[k] - \mathbf{\Gamma}_{a \rightarrow b}^+[k] \right| C_{a \rightarrow b} + \frac{\beta n_{a \rightarrow b} C_{a \rightarrow b}^2}{2d} \right) \\
 &= \sum_{[a,b] \in \mathcal{A}} \left(- \left(\frac{\lambda \Delta \mathbf{s}_{a \rightarrow b}[k] + \mathbf{\Gamma}_{a \rightarrow b}^-[k] - \mathbf{\Gamma}_{a \rightarrow b}^+[k]}{\sqrt{2n_{a \rightarrow b} \beta}} - C_{a \rightarrow b} \sqrt{\frac{n_{a \rightarrow b} \beta}{2}} \right)^2 \right. \\
 &\quad \left. + \frac{n_{a \rightarrow b} \beta C_{a \rightarrow b}^2}{2} + \frac{\beta n_{a \rightarrow b} C_{a \rightarrow b}^2}{2d} \right) \\
 &\leq \sum_{[a,b] \in \mathcal{A}} n_{a \rightarrow b} \beta C_{\underline{t} \rightarrow \bar{t}}^2 \\
 &\leq n_{\underline{t} \rightarrow \bar{t}} \beta C_{\underline{t} \rightarrow \bar{t}}^2 \\
 &\leq \beta B^2.
 \end{aligned}$$

Similarly it can be shown that $\sum_{[a,b] \in \mathcal{C}} \mathbf{S}_{a \rightarrow b}[k] = O(1)$. Recalling that $|\mathcal{A}| + |\mathcal{B}| + |\mathcal{C}| = O(d)$ we have

$$T_{2,[\underline{t}, \bar{t}]} + T_{3,[\underline{t}, \bar{t}]} \leq \sum_{k=1}^d \sum_{[a,b] \in \mathcal{A} \cup \mathcal{B} \cup \mathcal{C}} \mathbf{S}_{a \rightarrow b}[k] = O(d).$$

From Lemma 31 we have

$$T_{1,[\underline{t}, \bar{t}]} = \tilde{O}(d^{2.5}), \quad (23)$$

for bins $[\underline{t}, \bar{t}] \in \mathcal{P}$ that satisfy property 2 in Lemma 30.

Comparing Eq. (19) and (23) we conclude that

$$T_{1,i} + T_{2,i} + T_{3,i} = \tilde{O}(d^{2.5}), \quad (24)$$

for all bins $[i_s, i_t]$ in the partition \mathcal{P} of Lemma 30. Since $|\mathcal{P}| = O(dn^{1/3}C_n^{2/3})$, adding the above bound across all bins leads to the theorem.

If $C_n \leq 1/n$, then we have

$$\begin{aligned}
 \sum_{t=1}^n f_t(\mathbf{x}_j) - f_t(\mathbf{u}_t) &\leq \sum_{t=1}^n f_t(\mathbf{x}_j) - f_t(\mathbf{u}_1) + \sum_{t=1}^n f_t(\mathbf{u}_1) - f_t(\mathbf{u}_t) \\
 &\leq_{(a)} \tilde{O}(d^{1.5}) + G^\dagger n C_n \\
 &= \tilde{O}(d^{1.5})
 \end{aligned}$$

where line (a) follows from the fact that f_t is G^\dagger Lipschitz in \mathcal{D} . ■

Proposition 15 *For strongly convex losses, the regret bound can be improved to $\tilde{O}\left(d^2(n^{1/3}C_n^{2/3} \vee 1)\right)$ if $C_n > 1/n$ and $O(\log n)$ otherwise by using OGD as base learners in the FLH procedure. See Appendix D.2 for a proof.*

Proof [Proof Sketch] First we consider the case where the offline optimal is monotonic in each coordinate of a bin in \mathcal{P} . The static regret in any bin for strongly convex losses is $O(\log n)$ by Lemma 32 (as opposed to $\tilde{O}(d^{1.5})$ for exp-concave losses). Hence Eq.(23) can be re-written as $T_{1,[\underline{t}, \bar{t}]} = \tilde{O}(d)$. By following similar arguments as in proof of Theorem 14, we can re-write Eq.(24) as

$$T_{1,i} + T_{2,i} + T_{3,i} = \tilde{O}(d).$$

If the offline optimal is not monotonic, in each coordinate, we can write

$$T_{1,i} + T_{2,i} + T_{3,i} = \tilde{O}(1),$$

by following similar arguments for the corresponding case in the proof of Theorem 14.

Finally we sum across all $|\mathcal{P}| = O(dn^{1/3}C_n^{2/3})$. The case $C_n \leq 1/n$ can be handled similar to that of the exp-concave case. ■

Appendix E. Technical Lemmas

We start by describing a partitioning procedure namely `generateBins`.

<p><code>generateBins</code>: Inputs - the offline optimal sequence.</p> <p>Step 1 Initialize $\mathcal{Q} \leftarrow \Phi$. Starting from time 1, spawn a new bin $[i_s, i_t]$ whenever $\sum_{j=i_s+1}^{i_t+1} \ \mathbf{u}_j - \mathbf{u}_{j-1}\ _1 > B/\sqrt{n_i}$, where $n_i = i_t - i_s + 1$. Add the spawned bin $[i_s, i_t]$ to \mathcal{Q}.</p> <p>Step 2 Initialize $\mathcal{P} \leftarrow \Phi, \mathcal{R} \leftarrow \Phi$.</p> <p>Step 3 For each bin $[i_s, i_t] \in \mathcal{Q}$:</p> <ul style="list-style-type: none"> (a) Let $\Delta \mathbf{s}_i = \mathbf{s}_{i_t} - \mathbf{s}_{i_s-1}$. $\mathbf{\Gamma}_i^+ = \sum_{j=i_s}^{i_t} \gamma_j^+$. $\mathbf{\Gamma}_i^- = \sum_{j=i_s}^{i_t} \gamma_j^-$. (b) If for each $k \in [d]$, the sequence \mathbf{u}_k is monotonic in $[i_s, i_t]$, then remove $[i_s, i_t]$ from \mathcal{Q} and add it to \mathcal{P}. (c) If there exists one coordinate $k \in [d]$ such that $\mathbf{s}_{i_s-1}[k] \in [-1, -1/4]$ and $\mathbf{s}_{i_t}[k] \in [0, 1]$ and $\gamma_j^+[k] = 0 \forall j \in [i_s, i_t]$, then remove $[i_s, i_t]$ from \mathcal{Q} and add it to \mathcal{P}. Goto Step 3. (d) If there exists one coordinate $k \in [d]$ such that $\mathbf{s}_{i_s-1}[k] \in [-1/4, 0]$ and $\mathbf{s}_{i_t}[k] \in [1/4, 1]$ and $\gamma_j^+[k] = 0 \forall j \in [i_s, i_t]$, then remove $[i_s, i_t]$ from \mathcal{Q} and add it to \mathcal{P}. Goto Step 3. (e) If there exists one coordinate $k \in [d]$ such that \mathbf{u}_k is non-monotonic in $[i_s, i_t]$ and $\mathbf{s}_{i_s-1}[k] \in [-1/4, 1]$ and $\mathbf{s}_{i_t}[k] \in [-1/4, 1]$ and $\gamma_j^+[k] = 0 \forall j \in [i_s, i_t]$ then: <ul style="list-style-type: none"> i. Initialize $z \leftarrow i_s$. Remove $[i_s, i_t]$ from \mathcal{Q}.
--

- ii. if $\mathbf{u}_{i_s}[k] = -B$, then split $[i_s, i_t]$ into $[i_s, a]$ and $[a + 1, i_t]$ where a is the first time point within $[i_s, i_t]$ such that $\mathbf{u}_a[k] > -B$. Add $[i_s, a - 1]$ to \mathcal{Q} . Set $z \leftarrow a$.
 - iii. Let j be the first time in $[z, i_t]$ such that $\mathbf{s}_{j-1}[k] = -1$ with $\mathbf{u}_j[k] < \mathbf{u}_{j-1}[k]$. Add $[z, j - 1]$ and $[j, i_t]$ to \mathcal{P} . Goto Step 3.
- (f) If there exists one coordinate $k \in [d]$ such that \mathbf{u}_k is non-monotonic in $[i_s, i_t]$ and $\mathbf{s}_{i_s-1}[k], \mathbf{s}_{i_t}[k] \in [-1/4, 1/4]$ and $\gamma_j^+[k] = 0 \forall j \in [i_s, i_t]$ then:
- i. Initialize $z \leftarrow i_s$. Remove $[i_s, i_t]$ from \mathcal{Q} .
 - ii. if $\mathbf{u}_{i_s}[k] = -B$, then split $[i_s, i_t]$ into $[i_s, a]$ and $[a + 1, i_t]$ where a is the first time point within $[i_s, i_t]$ such that $\mathbf{u}_a[k] > -B$. Add $[i_s, a]$ to \mathcal{Q} . Set $z \leftarrow a + 1$.
 - iii. Let j be the first time in $[z, i_t]$ such that $\mathbf{s}_{j-1}[k] = -1$ with $\mathbf{u}_j[k] < \mathbf{u}_{j-1}[k]$. Add $[z, j - 1]$ and $[j, i_t]$ to \mathcal{P} . Goto Step 3.
- (g) If there exists one coordinate $k \in [d]$ \mathbf{u}_k is non-monotonic in $[i_s, i_t]$ and such that $\mathbf{s}_{i_s-1}[k] \in [-1, -1/4]$ and $\mathbf{s}_{i_t}[k] \in [-1, 1/4]$ and $\gamma_j^+[k] = 0 \forall j \in [i_s, i_t]$ then:
- i. Initialize $z \leftarrow i_t$. Remove $[i_s, i_t]$ from \mathcal{Q} .
 - ii. If $\mathbf{u}_{i_t}[k] = -B$, then split $[i_s, i_t]$ into $[i_s, a]$ and $[a + 1, i_t]$ where a is the last time point within $[i_s, i_t]$ such that $\mathbf{u}_a[k] > -B$. Add $[a + 1, i_t]$ to \mathcal{Q} . Set $z \leftarrow a$.
 - iii. Let j be the last time in $[i_s, z]$ such that $\mathbf{s}_{j-1}[k] = 1$ with $\mathbf{u}_j[k] > \mathbf{u}_{j-1}[k]$. Add $[i_s, j - 1]$ and $[j, z]$ to \mathcal{P} . Goto Step 3.
- (h) If there exists a coordinate $k \in [d]$ such that \mathbf{u}_k is non-monotonic in $[i_s, i_t]$ and $\mathbf{s}_{i_s-1}[k] \in [-1/4, 1]$ and $\mathbf{s}_{i_t}[k] \in [-1, 1/4]$ and $\gamma_j^+[k] = 0, \forall j \in [i_s, i_t]$ then:
- i. Initialize $p \leftarrow i_s - 1$. Remove $[i_s, i_t]$ from \mathcal{Q} .
 - ii. If $\mathbf{u}_{i_s}[k] = -B$, then let p be the largest point in $[i_s, i_t]$ such that $\mathbf{u}_t[k] = -B \forall t \in [i_s, p]$. Add $[i_s, p]$ to \mathcal{Q} .
 - iii. Let j be the first point in $[p + 1, i_t]$ with $\mathbf{s}_{j-1}[k] = -1$ with $\mathbf{u}_{j-1}[k] > -B$ and $\mathbf{u}_j[k] < \mathbf{u}_{j-1}[k]$. Add $[p + 1, j - 1]$ to \mathcal{P} .
 - iv. If $\mathbf{u}_r[k]$ is monotonic in $[j, i_t]$, add $[j, i_t]$ to \mathcal{Q} . Goto Step 3.
 - v. Initialize $q \leftarrow i_t + 1$.
 - vi. If $\mathbf{u}_{i_t}[k] = -B$, let q be smallest point in $[j, i_t]$ such that $\mathbf{u}_r[k] = -B \forall r \in [q, i_t]$. Add $[q, i_t]$ to \mathcal{Q} .
 - vii. Let h be the last time point in $[j, q - 1]$ such that $\mathbf{s}_{h-1}[k] = 1$ with $\mathbf{u}_h[k] > \mathbf{u}_{h-1}[k]$. Add $[j, h - 1]$ to \mathcal{P} .
 - viii. If $h < q - 1$, add $[h, q - 1]$ to \mathcal{P} .
 - ix. Goto Step 3.
- (i) If there exists one coordinate $k \in [d]$ such that $\mathbf{s}_{i_s-1}[k] \in [0, 1]$ and $\mathbf{s}_{i_t}[k] \in [-1, -1/4]$ and $\gamma_j^- = 0 \forall j \in [i_s, i_t]$, then remove $[i_s, i_t]$ from \mathcal{Q} and add it to \mathcal{P} . Goto Step 3.

- (j) If there exists one coordinate $k \in [d]$ such that $\mathbf{s}_{i_s-1}[k] \in [1/4, 1]$ and $\mathbf{s}_{i_t}[k] \in [-1/4, 0]$ and $\gamma_j^- = 0 \forall j \in [i_s, i_t]$, then remove $[i_s, i_t]$ from \mathcal{Q} and add it to \mathcal{P} . Goto Step 3.
- (k) If there exists one coordinate $k \in [d]$ such that \mathbf{u}_k is non-monotonic in $[i_s, i_t]$ and $\mathbf{s}_{i_s-1}[k], \mathbf{s}_{i_t}[k] \in [-1, 1/4]$ and $\gamma_j^- = 0 \forall j \in [i_s, i_t]$ and there exists a coordinate $j \in [i_s, i_t]$ such that $\mathbf{s}_{j-1}[k] = 1$ and $\mathbf{u}_{j-1}[k] < B$ then:
- i. Initialize $p \leftarrow i_s$. Remove $[i_s, i_t]$ from \mathcal{Q} .
 - ii. If $\mathbf{u}_{i_s}[k] = B$, then let p be the largest point in $[i_s, i_t]$ such that $\mathbf{u}_t[k] = B \forall t \in [i_s, p]$. Add $[i_s, p]$ to \mathcal{Q} .
 - iii. Let j be the first point in $[p+1, i_t]$ such that $\mathbf{s}_{j-1}[k] = 1$ with $\mathbf{u}_{j-1}[k] < \mathbf{u}_j[k]$. Add $[p+1, j-1]$ to \mathcal{P} . Add $[j, i_q]$ to \mathcal{Q} . Goto Step 3.
- (l) If there exists one coordinate $k \in [d]$ such that \mathbf{u}_k is non-monotonic in $[i_s, i_t]$ and $\mathbf{s}_{i_s-1}[k], \mathbf{s}_{i_t}[k] \in [-1/4, 1]$ and $\gamma_j^- = 0 \forall j \in [i_s, i_t]$ and there exists a $j \in [i_s, i_t]$ such that $\mathbf{u}_j[k] - \mathbf{u}_{j-1}[k] = -1$ and $\mathbf{u}_j[k] < B$ then:
- i. Initialize $z_1 \leftarrow i_s, z_2 \leftarrow i_t$. Remove $[i_s, i_t]$ from \mathcal{Q} .
 - ii. If $\mathbf{u}_{i_s}[k] = B$, then let p_1 be the last point in $[i_s, i_t]$ such that $\mathbf{u}_t[k] = B \forall t \in [i_s, p_1]$. Set $z_1 \leftarrow p_1 + 1$. Add $[i_s, p_1]$ to \mathcal{Q} .
 - iii. If $\mathbf{u}_{i_t}[k] = B$, then let p_2 be the smallest point in $[i_s, i_t]$ such that $\mathbf{u}_t[k] = B \forall t \in [p_2, i_t]$. Set $z_2 \leftarrow p_2 - 1$. Add $[p_2, i_t]$ to \mathcal{Q} .
 - iv. Let j be the last point in $[z_1, z_2]$ such that $\mathbf{s}_{j-1}[k] = -1$ and $\mathbf{u}_j[k] < B$ with $\mathbf{u}_{j-1}[k] > \mathbf{u}_j[k]$. Add $[z_1, j-1]$ and $[j, z_2]$ to \mathcal{P} . Goto Step 3.
- (m) If there exists one coordinate $k \in [d]$ such that \mathbf{u}_k is non-monotonic in $[i_s, i_t]$ and $\mathbf{s}_{i_s-1}[k] \in [-1, 1/4]$ and $\mathbf{s}_{i_t}[k] \in [-1/4, 1]$ and $\gamma_j^- = 0 \forall j \in [i_s, i_t]$ then:
- i. Initialize $p \leftarrow i_s - 1$. Remove $[i_s, i_t]$ from \mathcal{Q} .
 - ii. If $\mathbf{u}_{i_s}[k] = B$, then let p be the last time point such that $\mathbf{u}_t[k] = B \forall t \in [i_s, p]$. Add $[i_s, p]$ to \mathcal{Q} .
 - iii. Let j be the first point in $[p+1, i_t]$ such that $\mathbf{s}_{j-1}[k] = 1$ with $\mathbf{u}_{j-1}[k] < B$ with $\mathbf{u}_{j-1}[k] < \mathbf{u}_j[k]$. Add $[p+1, j-1]$ to \mathcal{P} .
 - iv. If $\mathbf{u}_r[k]$ is monotonic in $[j, i_t]$, add $[j, i_t]$ to \mathcal{Q} . Goto Step 3.
 - v. Initialize $q \leftarrow i_t + 1$.
 - vi. If $\mathbf{u}_{i_t}[k] = B$, let q be smallest point in $[j, i_t]$ such that $\mathbf{u}_r[k] = B \forall r \in [q, i_t]$. Add $[q, i_t]$ to \mathcal{Q} .
 - vii. Let h be the last time point in $[j, q-1]$ such that $\mathbf{s}_{h-1}[k] = -1$ with $\mathbf{u}_{h-1}[k] > \mathbf{u}_h[k]$. Add $[j, h-1]$ to \mathcal{P} .
 - viii. If $h < q-1$, add $[h, q-1]$ to \mathcal{P} .
 - ix. Goto Step 3.

Step 4 Return \mathcal{P} .

Lemma 36 *The partitioning routine generateBins halts. Further we have $|\mathcal{P}| = O(dn^{1/3}C_n^{2/3})$.*

Proof We need to argue that the loop in Step 3 halts.

FACT1: Notice that in the loop of Step 3, we add a bin to \mathcal{Q} only if $\mathbf{u}[\tilde{k}]$ is monotonic in that bin for a coordinate \tilde{k} . Once such a bin is added, in the later steps we do not create new bins across the previous coordinate \tilde{k} .

FACT2: Step 1 ensures that within each bin we consider, a TV of at-most B will only be incurred. Due to Lemma 25, this TV constraint implies that both $\gamma^+[k]$ and $\gamma^-[k]$ cannot be simultaneously non-zero within any bin $[i_s, i_t]$. Consequently we consider all possible configurations of such TV constrained bins in Steps 3(b-m). See Table 1 for a comprehensive summary.

Combining the previous two facts, we conclude that any time point in $[n]$ will be into some bin in \mathcal{P} or \mathcal{R} in at-most d (maybe non-consecutive) iterations of the loop in Step 3.

By using similar arguments as in proof of Lemma 5, we have $|\mathcal{Q}| = O(n^{1/3}C_n^{2/3})$ after Step 1 gets finished. Due to FACT1, the loop in Step 3 can split a bin that was originally present in \mathcal{Q} at the end of Step 1 into at-most $O(d)$ sub-bins. Hence $|\mathcal{P}|$ can be $O(dn^{1/3}C_n^{2/3})$ after Step 3. \blacksquare

Lemma 37 *Let \mathcal{P} be the partition produced by `generateBins`. Consider a bin $[i_s, i_t] \in \mathcal{P}$. Using the notations of Lemma 33, the bin $[i_s, i_t]$ satisfy one of the following properties.*

- *Property 1: Across each coordinate $k \in [d]$, the sequence $\mathbf{u}_j[k], j \in [i_s, i_t]$ is non-decreasing or non-increasing. Or,*
- *Property 2: $\|\lambda\Delta\mathbf{s}_i + \mathbf{\Gamma}_i^- - \mathbf{\Gamma}_i^+\|_2 \geq \lambda/4$.*

Proof To prove the properties satisfied by each bin in \mathcal{P} we inspect the steps in `generateBins` and verify the stated properties. Below when we refer the coordinate k , we mean the same coordinate that is used by the corresponding steps in `generateBins`. For a bin $[a, b]$, we also recall the notations $\Delta\mathbf{s}_{a \rightarrow b}$, $\mathbf{\Gamma}_{a \rightarrow b}^+$ and $\mathbf{\Gamma}_{a \rightarrow b}^-$. We use the short hands $\Delta\mathbf{s}_i, \mathbf{\Gamma}_i^+, \mathbf{\Gamma}_i^-$ as in Lemma 33 for a bin referred by $[i_s, i_t]$.

1. For the bin added in Step 3(c) we have $\Delta\mathbf{s}_i[k] > 1/4$. Also $\mathbf{\Gamma}_i^-[k] - \mathbf{\Gamma}_i^+[k] = \mathbf{\Gamma}_i^-[k] \geq 0$. Hence Property 2 is verified.
2. Step 3(d) can be verified as above.
3. In Step 3(e), we add $[z, j-1]$ and $[j, i_t]$ to \mathcal{P} . By construction, the $\mathbf{u}[k]$ solution do not attain the value $-B$ in $[z, j-1]$. Hence $\mathbf{\Gamma}_{z \rightarrow j-1}^- - \mathbf{\Gamma}_{z \rightarrow j-1}^+ = 0$. Since $\mathbf{s}_{z-1}[k] \in [-1/4, 1]$ and $\mathbf{s}_{j-1}[k] = -1$, we have $\left| \lambda\Delta\mathbf{s}_{z \rightarrow j-1}[k] + \mathbf{\Gamma}_{z \rightarrow j-1}^- - \mathbf{\Gamma}_{z \rightarrow j-1}^+ \right| > \lambda/4$. Hence Property 2 is verified for $[z, j-1]$. For the bin $[j, i_t]$ we have $\gamma_r^+[k] = 0, \forall r \in [j, i_t]$. Since $\mathbf{s}_{j-1}[k] = -1$ and $\mathbf{s}_2[k] \in [-1/4, 1]$, we have $\Delta\mathbf{s}_{j \rightarrow i_t}[k] \geq 1/4$. So $\lambda\Delta\mathbf{s}_{j \rightarrow i_t}[k] + \mathbf{\Gamma}_{j \rightarrow i_t}^- - \mathbf{\Gamma}_{j \rightarrow i_t}^+ \geq \lambda\Delta\mathbf{s}_{j \rightarrow i_t}[k] \geq \lambda/4$. Thus Property 2 is verified for $[j, i_t]$. See Fig. 11 for an example of this configuration.
4. Step 3(f) can be verified using similar arguments as above.
5. In Step 3(g) we add $[i_s, j-1]$ and $[j, z]$ to \mathcal{P} . Since $\mathbf{s}_1[k] \in [-1, -1/4]$ and $\mathbf{s}_{j-1}[k] = 1$. So we have $\Delta\mathbf{s}_{i_s \rightarrow j-1} \geq 0$ and hence Property 2 is satisfied for $[i_s, j-1]$. By construction there $\mathbf{u}[k]$ do not attain the value $-B$ in $[j, z]$. So $\mathbf{\Gamma}_{j \rightarrow z}^- - \mathbf{\Gamma}_{j \rightarrow z}^+ = 0$. Since $\mathbf{s}_z[k] \in [-1, 1/4]$ and $\mathbf{s}_{j-1}[k] = 1$, we conclude that Property 2 is satisfied for $[j, z]$. See Fig. 12 for an example of this configuration.

6. By construction of Step 3(h) $\gamma_r^-[k] = 0 \forall r \in [p+1, j-1]$. Thus $\Gamma_{p+1 \rightarrow j-1}^-[k] - \Gamma_{p+1 \rightarrow j-1}^+[k] = 0$. Since $s_{j-1}[k] = -1$ and $s_p[k] \in [-1/4, 1]$, we can conclude that Property 2 is verified for $[p+1, j-1]$. Since $s_{h-1}[k] = 1$ we have $\lambda \Delta s_{j \rightarrow h-1}[k] + \Gamma_{j \rightarrow h-1}^-[k] - \Gamma_{j \rightarrow h-1}^+[k] \geq \lambda \Delta s_{j \rightarrow h-1}[k] \geq \lambda/4$. Thus Property 2 is verified for $[j, h-1]$. By construction, $\mathbf{u}[k]$ do not attain the value $-B$ in $[h, q-1]$. Hence $\Gamma_{h \rightarrow q-1}^- - \Gamma_{h \rightarrow q-1}^+ = 0$. Since $s_{q-1}[k] \in [-1, 1/4]$ and $s_{h-1}[k] = 1$, Property 2 is satisfied for $[h, q-1]$. See Fig. 13 for an example of this configuration.

The Properties stated in the Lemma can be verified for all bins that get added to \mathcal{P} in steps 3(i-m) using similar arguments as above. ■

Proof of Lemma 30. The proof is completed by the partitioning produced by `generateBins` and results of Lemmas 36 and 37. ■

$s_{i_s-1}[k]$	$s_{i_t}[k]$	$\Gamma_i^+[k] = 0$	$\Gamma_i^-[k] = 0$
[-1,-1/4]	[-1,-1/4]	(g)	(k)
	[-1/4,0]	(g)	(k)
	[0,1/4]	(g)	(k)
	[1/4,1]	(c)	(m)
[-1/4,0]	[-1,-1/4]	(h)	(k)
	[-1/4,0]	(f)	(l)
	[0,1/4]	(f)	(l)
	[1/4,1]	(d)	(l)
[0,1/4]	[-1,-1/4]	(h)	(k)
	[-1/4,0]	(f)	(k)
	[0,1/4]	(f)	(k)
	[1/4,1]	(e)	(l)
[1/4,1]	[-1,-1/4]	(h)	(i)
	[-1/4,0]	(h)	(j)
	[0,1/4]	(h)	(l)
	[1/4,1]	(e)	(l)

Table 1: Various configurations of a non-monotonic coordinate within a bin $[i_s, i_t]$ and their assignments to the corresponding steps of `generateBins` routine for the cases $\gamma_j^+[k] = 0$ for all $j \in [i_s, i_t]$ and $\gamma_j^-[k] = 0$ for all $j \in [i_s, i_t]$.




Figure 11: An example of a configuration corresponding Step 3(e) of `generateBins`. Here $z = i_s$.




Figure 12: An example of a configuration corresponding Step 3(g) of `generateBins`. Here $z = i_t$.




Figure 13: An example of a configuration corresponding Step 3(h) of `generateBins`. Here $p + 1 = i_s$, $q = i_t + 1$.