

Online Contextual Learning with Limited Feedback

Sayak Ray Chowdhury

Boston University, USA

SAYAK@BU.EDU

Aditya Gangrade

Carnegie Mellon University, USA

AGANGRA2@ANDREW.CMU.EDU

Ashok Cutkosky

Boston University, USA

CUTKOSKY@BU.EDU

Venkatesh Saligrama

Boston University, USA

SRV@BU.EDU

Abstract

We consider a smooth online *contextual* convex optimization problem, where the learner, given the context, has flexibility to decide whether to predict, get feedback and suffer a loss, or to abstain and suffer no loss but get no feedback in every round. The learner is compared against a competitor that gets feedback, and suffers a loss, at every round. This paves the way for designing an abstention rate for the learner that obtains *zero* cumulative regret. We design an epoch-based online gradient descent algorithm, where, at every epoch, abstention rates are computed adaptively in order to minimize abstentions while ensuring zero regret. We bound the excess abstention needed by our algorithm compared to a certain benchmark abstention schedule. The ratio of excess abstentions to the benchmark in every epoch m decays at a rate approximately $O\left(\max_{c=1}^K \sqrt{\frac{1}{\tau_{c,m}}}\right)$, where K is the number of contexts and $\tau_{c,m}$ is the total number of observations for context c in epoch m .

Keywords: Online learning, Limited feedback

1. Introduction

Consider a resource constrained (e.g., low-power, battery-limited) edge device, such as a sensor or a smartphone that receives a stream of learning tasks (e.g., classification, regression). Due to resource limitations, such a device cannot locally implement most modern day accurate but resource-intensive models (e.g., deepNNs) that are needed to make accurate decisions. Instead the device has access to such a complex model implemented on a cloud server, to which it can send queries in order to maintain accuracy. This incurs communication costs such as latency and battery drain. By leveraging this ability, we should hope that our device would perform *better* than the best possible on-device model (which could be interpreted as the best compression of the cloud model). Of course, we need to do this while minimizing the number of queries to the cloud, i.e., minimizing the total resource consumption (Xu et al., 2014; Nan and Saligrama, 2017). Furthermore, the complexity of the tasks may vary depending on some external factors. For example, regression tasks assigned to a sensor can be hard during some specific times of the day, and can be easy during other times. These scenarios can be modeled by incorporating contextual information in the learning tasks.

We are interested in devising on-device learning algorithms that address these situations. However, we also recognize an additional wrinkle: in many such on-device prediction tasks, obtaining any sort of feedback on the predictions is extremely difficult (e.g. a smartphone user may not indicate whether a dining recommendation was good or not). However, if we believe that complex cloud model is highly accurate, we can use a query to the cloud not only to obtain better predictions, but also as a way of obtaining *labels*. Thus, the decision to query or not is based not only on problem “difficulty”, but also on a desire to gather data.

Selective online optimization is a paradigm of relevance to such settings (Cesa-Bianchi et al., 2005; Alon et al., 2015). At each round, the learning setup allows an agent to *optionally* observe feedback (typically, a convex and smooth loss function; depending on the generated context at that round) before picking a decision with some probability. This allows us to model a scenario in which feedback is expensive or scarce. In a different twist, we allow the agent to also obtain this expensive information *before* taking any action. Hence, the agent enjoys flexibility to pick the a posteriori best decision corresponding to the observed feedback, and abstain from suffering a loss. Conceptually, this captures querying the cloud model and picking an accurate decision for the learning task of that round. If the agent decides to not query/abstain, then it doesn’t receive any feedback, and it picks the decision suggested by a standard online learning algorithm (e.g. online gradient descent), which is run locally on the edge device.

Given a certain amount of prediction error budget, we would like to minimize the total number of abstentions made by the learning. Thus, there should be some ideal pareto-frontier of possible performances. However, we will focus on a particular intuitively-desirable budget: the total loss of the learning should be *strictly smaller* than the loss of the best-in-hindsight on-device model. That is, given this additional flexibility of abstaining, we would like to ensure a negative regret for the learner while minimizing the total number of abstentions (i.e. number of queries to the cloud).

Our problem is related to that of online selective classification (Cortes et al., 2018; Li et al., 2011; Gangrade et al., 2021). In this setting, data (feature, label) is generated sequentially by an adversary; and a learner uses the features to produce a decision which can either be a label, or an abstention. Typically, feedback is provided if the learner abstains; and the learner is said to incur a mistake if it didn’t abstain and its prediction did not equal the true label. The emphasis is on ensuring very few mistakes, and bound the regret suffered compared to a classifier that makes no mistakes, while abstaining the fewest number of times. However, it is not straightforward to extend these results to the contextual learning setting since it presents an additional challenge of adapting abstention probabilities with contexts.

We construct an epoch-based online gradient descend (OGD) scheme with doubling epoch lengths. At the beginning of every epoch, we adaptively compute abstention probabilities for every context without having to rely on the future loss functions. With these abstention probabilities, we query the cloud model to ensure non-positive regret of OGD in the current epoch. With the emphasis on non-positive regret, we bound the learner’s excess abstentions over an optimal abstention strategy that has the knowledge of cumulative competitor loss for all contexts. We show that at every epoch m , the excess abstentions

decays at a rate approximately $O\left(\max_{c=1}^K \sqrt{\frac{1}{\tau_{c,m}}}\right)$, where K is the number of contexts and $\tau_{c,m}$ is the total number of observations for context c in epoch m .

2. Preliminaries

We consider a modification of the online convex optimization problem with *contexts* and *abstentions*. First, we describe a more standard setting without abstentions before describing our full setting. The learning game proceeds as follows. At each round $s \geq 1$, we (i) observe a context $c \in \mathcal{C}$, where $|\mathcal{C}| = K$, (ii) play a point $w_{c,s} \in W \subset \mathbb{R}^d$ from some closed and non-empty set W of diameter D , (iii) observe convex loss function $f_{c,s} : \mathbb{R}^d \rightarrow \mathbb{R}$, and (iv) pay a loss of value $f_{c,s}(w_{c,s})$. The aim of this game is to minimize the cumulative regret with respect to any competitor $u = (u_1, \dots, u_K) \in W^K$, where $u_c \in W$ is the competitor associated with context c . The cumulative regret after T rounds is defined as

$$R(u) = \sum_{c=1}^K R_c(u_c), \quad R_c(u_c) = \sum_{s=1}^{\tau_{c,T}} (f_{c,s}(w_{c,s}) - f_{c,s}(u_c)),$$

where $\tau_{c,T}$ denotes the number of times context c arrives in T rounds.

In what follows, we will work with losses that are smooth and stochastic. Specifically, we assume that each loss function $f_{c,s}$ is H -smooth in open sets containing W and satisfies $f_{c,s}(w) \in [0, 1]$ for all $w \in \mathbb{R}$. This implies $\|\nabla f_{c,s}(w)\|^2 \leq 2Hf_{c,s}(w)$ for all w, c and s . Furthermore, for each context c , the loss functions $\{f_{c,s}\}_{s \geq 1}$ are drawn i.i.d. from some distribution \mathcal{D}_c supported over $[0, 1]^W$. We also assume that contexts generated independently of all other randomness with (known) probabilities $\pi = (\pi_1, \dots, \pi_K)$ so that $\sum_{c=1}^K \pi_c = 1$.

Adaptive learning rate for OGD. Define, for each context $c \in \mathcal{C}$, and competitor $u_c \in W$, its cumulative loss in T rounds $L_{c,T}^* = \sum_{s=1}^{\tau_{c,T}} f_{c,s}(u_c)$. Then, standard analysis of OGD (Orabona, 2019) for losses corresponding to context c with learning rate $\eta_{c,s} = \frac{D}{\sqrt{2 \sum_{s'=1}^s \|\nabla f_{c,s'}(w_{c,s'})\|^2}}$, $s = 1, \dots, \tau_{c,T}$, yields the regret

$$R_c(u_c) \leq 2D \sqrt{H \sum_{s=1}^{\tau_{c,T}} f_{c,s}(w_{c,s})} = \sqrt{2} \min_{\eta_c > 0} \frac{D^2}{2\eta_c} + H\eta_c \sum_{s=1}^{\tau_{c,T}} f_{c,s}(w_{c,s}), \quad (1)$$

which can be further expressed in terms of $L_{c,T}^*$ as $R_c(u_c) \leq 4HD^2 + 4D\sqrt{HL_{c,T}^*}$.

Abstentions. In this paper we make a modification to the contextual online convex optimization game: after seeing the context c , our learner may *optionally* choose to see the loss $f_{c,s}$ *before* picking $w_{c,s}$ (we call this choice “abstaining”, or “receiving feedback”). Otherwise, we get *no* feedback. Further, we make the assumption that for all s , there is some $w_{c,s}^* \in W$ such that $f_{c,s}(w_{c,s}^*) = 0$. We will be focusing our attention on randomized strategies which abstain with some c -dependent probability p_c , and always play $w_{c,s}^*$ to suffer zero loss when abstaining. We can show that a simple variant of adaptive gradient descent with such randomized feedback obtains:

$$\mathbb{E}[R(u)] \leq \sum_{c=1}^K \left(\frac{4HD^2}{p_c} + 4D\sqrt{\frac{HL_{c,T}^*}{p_c}} - p_c L_{c,T}^* \right), \quad (2)$$

where the expectation is over the randomness introduced due to abstentions. From this, it is clear that appropriate choices of p_c will lead to negative regret.

The key problem is determining what p_c should be: contexts with large $L_{c,T}^*$ are inherently “difficult” to predict on-device, and so we should increase p_c . Note that this “difficulty” arises from two different phenomena. If a context c is very common (has high π_c), then $L_{c,T}^*$ becomes larger. However, even for a relatively rare context, $L_{c,T}^*$ may still be large if it is impossible for the on-device model to accurately predict in this context.

Note that this problem may be interesting even for a *single* context. However, when multiple contexts are involved then it is a significantly richer study: we now have the ability to “hedge” different contexts against each other. For example, by abstaining slightly more than may be naively required on a “hard”, but possibly slightly less common context, we obtain a more negative regret on that context, which can be used to cancel out a positive regret on an easier but more common context, saving a large number of abstentions on this easier context in the process. Thus, the goal is to determine how to balance these abstention rates p_c given that we do not know $L_{c,T}^*$ apriori.

3. Algorithm and Result

We consider an epoch-based strategy, where the length of epoch m is given by τ_m with $\tau_{m+1} = 2\tau_m$, $m \geq 1$. Intuitively, in each epoch, we select conservative p_c using an estimate of L_c^* generated in the *previous* epochs. Then, we will obtain negative regret in this epoch while refining our estimates of L_c^* in order to reduce p_c further in subsequent epochs.

Let $\tau_{c,m}$ denote the number of occurrence of context c in epoch m and $L_{c,m}^* := \sum_{s=1}^{\tau_{c,m}} f_{c,s}(u_c)$ denotes the cumulative loss of a competitor u_c in epoch m . Given probabilities $p = (p_1, \dots, p_K)$, total number of abstentions in epoch m is given by $A_m(p) = \sum_{c=1}^K \pi_c p_c \tau_m$. Furthermore, from (2), the regret in epoch m , denoted by $R_m(u; p)$, satisfies $\mathbb{E}[R_m(u; p)] \leq \sum_{c=1}^K \left(\frac{4HD^2}{p_c} + 4D\sqrt{\frac{HL_{c,m}^*}{p_c}} - p_c L_{c,m}^* \right)$. If the cumulative competitor loss $L_{c,m}^*$, $c = 1, \dots, K$, were known apriori, the abstention probabilities guaranteeing non-positive regret in epoch m could have been computed as follows:

$$p_m^* = \operatorname{argmin}_{0 < p \leq 1} A_m(p) \quad \text{s.t.} \quad \sum_{c=1}^K \left(\frac{4HD^2}{p_c} + 4D\sqrt{\frac{HL_{c,m}^*}{p_c}} - p_c L_{c,m}^* \right) \leq 0. \quad (3)$$

Since the competitor losses are not known in advance, we estimate those using data generated in epoch $m-1$. To this end, let $I_{c,m-1,s}$ denotes the indicator of abstaining for context c at round $s \leq \tau_{c,m-1}$, which occurs with probability $\hat{p}_{c,m-1}$. Define the estimate $\hat{L}_{c,m} = 2 \sum_{s=1}^{\tau_{c,m-1}} \frac{I_{c,m-1,s}}{\hat{p}_{c,m-1}} f_{c,s}(w_{c,s})$. Also, define, for $\delta \in (0, 1]$, $\lambda_c \in (0, 1)$, the deviations

$$\Delta_{c,m}^U = \alpha_{c,m} + 2\beta_{c,m} - 4\sqrt{2\tau_{c,m-1} \log(8K/\delta)}, \quad \Delta_{c,m}^L = \alpha_{c,m} + 2\lambda_c \beta_{c,m} + 2D^2 H / (1 - \lambda_c),$$

where $\alpha_{c,m} = 2\sqrt{2\pi_c \tau_{m-1} \log(8K/\delta)} + \frac{2}{3} \log(8K/\delta)$, $\beta_{c,m} = \frac{1 - \hat{p}_{c,m-1}}{\hat{p}_{c,m-1}} \sqrt{2\tau_{c,m-1} \log(8K/\delta)}$.

Using these estimates, we compute abstention probabilities $\hat{p}_m = (\hat{p}_{1,m}, \dots, \hat{p}_{K,m})$ as

$$\hat{p}_m = \operatorname{argmin}_{0 < p \leq 1} A_m(p) \quad \text{s.t.} \quad \sum_{c=1}^K \left(\frac{4HD^2}{p_c} + 4D\sqrt{\frac{H(\hat{L}_{c,m} + \Delta_{c,m}^U)}{p_c}} - p_c(\lambda_c \hat{L}_{c,m} - \Delta_{c,m}^L) \right) \leq 0. \quad (4)$$

Excess abstentions. Define, for each context c , the competitor $u_c = \operatorname{argmin}_{w \in W} \mathbb{E}[f_{c,s}(w)]$. First, we obtain the following concentration bound for the competitor loss $\tau_{c,m} L_{c,m}^*$.

Algorithm 1 Epoch-Based OGD with abstentions

- 1: **Parameter:** Confidence level $\delta \in (0, 1]$, constants $\lambda_1, \dots, \lambda_K \in (0, 1)$
 - 2: **Initialize:** $\hat{p}_{c,1} = 1$ for all c , initial epoch length τ_1
 - 3: **for** epoch $m = 1, 2, \dots$ **do**
 - 4: Initialize $w_{c,1} \in W$ and $\tau_{c,m} = 0$ for all c
 - 5: **for** time $s = 1, 2, \dots, \tau_m$ **do**
 - 6: Observe context c_s and set $\tau_{c_s,m} = \tau_{c_s,m} + 1$
 - 7: Sample $I_{c_s,m,\tau_{c_s,m}} \sim \mathcal{B}(\hat{p}_{c_s,m})$
 - 8: **if** $I_{c_s,m,\tau_{c_s,m}} = 1$ **then**
 - 9: Output $w_{c_s,\tau_{c_s,m}}^*$ and pay loss 0
 - 10: **else**
 - 11: Output $w_{c_s,\tau_{c_s,m}}$ and pay loss $f_{c_s,\tau_{c_s,m}}(w_{c_s,\tau_{c_s,m}})$
 - 12: **end if**
 - 13: Set $w_{c_s,\tau_{c_s,m}+1} = w_{c_s,\tau_{c_s,m}} - \frac{D}{\sqrt{2 \sum_{s'=1}^{\tau_{c_s,m}} \|\nabla f_{c_s,s'}(w_{c_s,s'})\|^2}} \nabla f_{c_s,\tau_{c_s,m}}(w_{c_s,\tau_{c_s,m}})$
 - 14: **end for**
 - 15: Set size of next epoch $\tau_{m+1} = 2\tau_m$
 - 16: Compute estimates $\hat{L}_{c,m+1}$, $\Delta_{c,m+1}^U$, $\Delta_{c,m+1}^L$ for all c
 - 17: Compute abstention rates $\hat{p}_{c,m+1}$ for all c using (4)
 - 18: **end for**
-

Lemma 1. For any $m \geq 1$, $\delta \in (0, 1]$, $\lambda_1, \dots, \lambda_K \in (0, 1)$, with probability at least $1 - \delta$,

$$\forall c = 1, \dots, K, \quad \lambda_c \hat{L}_{c,m} - \Delta_{c,m}^L \leq L_{c,m}^* \leq \hat{L}_{c,m} + \Delta_{c,m}^U.$$

Armed with this result, we have the following bound on regret and excess abstentions.

Theorem 2. Define, for each $c \in [K]$ and $m \geq 1$, $\varepsilon_{c,m}^U = \frac{\Delta_{c,m}^U}{L_{c,m}^*}$ and $\varepsilon_{c,m}^L = \frac{\Delta_{c,m}^L}{L_{c,m}^*}$. Then, for each epoch $m \geq 1$, with probability at least $1 - \delta$, we have

$$\mathbb{E}[R_m(u; \hat{p}_m)] \leq 0, \quad \frac{A_m(\hat{p}_m) - A_m(p_m^*)}{A_m(p_m^*)} \leq \max_{c=1}^K \left(1 + \frac{1}{\lambda_c} \right) \left(\frac{(1 - \lambda_c) + (\lambda_c \varepsilon_{c,m}^U + \varepsilon_{c,m}^L)}{\lambda_c - (\lambda_c \varepsilon_{c,m}^U + \varepsilon_{c,m}^L)} \right).$$

Observe that $\varepsilon_{c,m}^U \approx O(1/\sqrt{\tau_{c,m}})$ and $\varepsilon_{c,m}^L \approx O(\lambda_c/\sqrt{\tau_{c,m}})$. Hence, we the excess factor in abstentions corresponding to context c behaves approximately as $O(a_{c,m} + b_{c,m}/\sqrt{\tau_{c,m}})$ for appropriate constants $a_{c,m}, b_{c,m}$ (depending on $\lambda_c, D, H, K, \delta, \pi_c, \hat{p}_{c,m-1}$). Therefore, the total excess abstentions in epoch m decays at a rate roughly $O(\max_{c=1}^K 1/\sqrt{\tau_{c,m}})$.

Remark 3. Instead of epoch-based scheme, one can consider abstaining for first half of the horizon T , and use data generated to compute abstention rates for the next half. As a corollary of Theorem 2, this scheme attains $O(\max_{c=1}^K 1/\sqrt{\tau_{c,T}})$ bound on excess abstentions.

Proof. We work under the good event of Lemma 1. We have the expected regret

$$\begin{aligned} \mathbb{E}[R_m(u; \hat{p}_m)] &\leq \sum_{c=1}^K \left(\frac{4HD^2}{\hat{p}_{c,m}} + 4D \sqrt{\frac{HL_{c,m}^*}{\hat{p}_{c,m}}} - \hat{p}_{c,m} L_{c,m}^* \right) \\ &\leq \sum_{c=1}^K \left(\frac{4HD^2}{\hat{p}_{c,m}} + 4D \sqrt{\frac{H(\hat{L}_{c,m} + \Delta_{c,m}^U)}{\hat{p}_{c,m}}} - \hat{p}_{c,m} (\lambda_c \hat{L}_{c,m} - \Delta_{c,m}^L) \right) \leq 0, \end{aligned}$$

where the last inequality follows from definition of \hat{p}_m . Denote, for each context c , an abstention probability (assuming it is well-defined)

$$\tilde{p}_{c,m} := p_{c,m}^* \frac{1/\lambda_c + \varepsilon_{c,m}^U + \varepsilon_{c,m}^L/\lambda_c}{\lambda_c - \lambda_c \varepsilon_{c,m}^U - \varepsilon_{c,m}^L} = p_{c,m}^* \frac{1 + \lambda_c \varepsilon_{c,m}^U + \varepsilon_{c,m}^L}{\lambda_c (\lambda_c - \lambda_c \varepsilon_{c,m}^U - \varepsilon_{c,m}^L)}.$$

It holds that $\tilde{p}_m = (\tilde{p}_{1,m}, \dots, \tilde{p}_{K,m})$ is a feasible point of (4). To see this, note that

$$\begin{aligned} & \sum_{c=1}^K \left(\frac{4HD^2}{\tilde{p}_{c,m}} + 4D \sqrt{\frac{H(\hat{L}_{c,m} + \Delta_{c,m}^U)}{\tilde{p}_{c,m}}} - \tilde{p}_{c,m} (\lambda_c \hat{L}_{c,m} - \Delta_{c,m}^L) \right) \\ & \leq \sum_{c=1}^K \left(\frac{4HD^2}{\tilde{p}_{c,m}} + 4D \sqrt{\frac{H(L_{c,m}^*/\lambda_c + \Delta_{c,m}^U + \Delta_{c,m}^L/\lambda_c)}{\tilde{p}_{c,m}}} - \tilde{p}_{c,m} (\lambda_c L_{c,m}^* - \lambda_c \Delta_{c,m}^U - \Delta_{c,m}^L) \right) \\ & \leq \sum_{c=1}^K \left(\frac{4HD^2}{\tilde{p}_{c,m}} + 4D \sqrt{\frac{HL_{c,m}^* (1/\lambda_c + \varepsilon_{c,m}^U + \varepsilon_{c,m}^L/\lambda_c)}{\tilde{p}_{c,m}}} - \tilde{p}_{c,m} L_{c,m}^* (\lambda_c - \lambda_c \varepsilon_{c,m}^U - \varepsilon_{c,m}^L) \right) \\ & = \sum_{c=1}^K \left(\frac{4HD^2}{p_{c,m}^*} \frac{\lambda_c - \lambda_c \varepsilon_{c,m}^U - \varepsilon_{c,m}^L}{1/\lambda_c + \varepsilon_{c,m}^U + \varepsilon_{c,m}^L/\lambda_c} + 4D \sqrt{\frac{HL_{c,m}^* (\lambda_c - \lambda_c \varepsilon_{c,m}^U - \varepsilon_{c,m}^L)}{p_{c,m}^*}} - p_{c,m}^* L_{c,m}^* \left(\frac{1}{\lambda_c} + \varepsilon_{c,m}^U + \frac{\varepsilon_{c,m}^L}{\lambda_c} \right) \right) \\ & \leq \sum_{c=1}^K \left(\frac{4HD^2}{p_{c,m}^*} + 4D \sqrt{\frac{HL_{c,m}^*}{p_{c,m}^*}} - p_{c,m}^* L_{c,m}^* \right) \leq 0, \end{aligned}$$

where the second last inequality is due to $\lambda_c < 1$ and the last inequality follows from (3). The bound on expected abstentions follows by noting that $A_m(\hat{p}_m) \leq A_m(\tilde{p}_m)$. \square

Proof sketch of Lemma 1. Since we consider doubling epochs, we first control the difference between competitor loss $L_{c,m}^*$ in epoch m and twice the competitor loss $L_{c,m}^*$ in epoch $m-1$. To this end, Bernstein's inequality yields with probability at least $1 - 2\delta$,

$$|L_{c,m}^* - 2L_{c,m-1}^*| = \left| \sum_{s=1}^{\tau_{c,m}} f_{c,s}(u_c) - 2 \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(u_c) \right| \leq 2\sqrt{2\pi_c \tau_{m-1} \log(2/\delta)} + \frac{2}{3} \log(2/\delta).$$

Since the competitor loss $L_{c,m-1}^*$ is unknown, we control its deviation from the loss of OGD. First, by Azuma-Hoeffding bound, we get $\sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(w_{c,s}) - f_{c,s}(u_c) \geq -2\sqrt{2\tau_{c,m-1} \log(2/\delta)}$, with probability at least $1 - \delta$. Next, by regret bound of OGD (see (1)), for any $\eta_c > 0$, we get $\sum_{s=1}^{\tau_{c,m-1}} (f_{c,s}(w_{c,s}) - f_{c,s}(u_c)) \leq \sqrt{2} \left(\frac{D^2}{2\eta_c} + H\eta_c \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(w_{c,s}) \right)$. Choosing an $\eta_c > 0$ such that $\lambda_c := 1 - \sqrt{2}H\eta_c > 0$, we obtain

$$2\lambda_c \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(w_{c,s}) - \frac{2D^2H}{1-\lambda_c} \leq 2 \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(u_c) \leq 2 \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(w_{c,s}) - 4\sqrt{2\tau_{c,m-1} \log(2/\delta)}.$$

Since we get to see feedback in epoch $m-1$ only with probability $\hat{p}_{c,m-1}$, we now need to control the loss of OGD with its importance-weighted version. To this end, by Azuma-Hoeffding inequality, we obtain $\left| \sum_{s=1}^{\tau_{c,m-1}} \left(1 - \frac{I_{c,m-1,s}}{\hat{p}_{c,m-1}} \right) f_{c,s}(w_{c,s}) \right| \leq \frac{1-\hat{p}_{c,m-1}}{\hat{p}_{c,m-1}} \sqrt{2\tau_{c,m-1} \log(2/\delta)}$, with probability at least $1 - \delta$. Putting everything together, taking a union bound over all contexts and replacing δ with $\delta/(4K)$, the result follows.

Conclusion. We conclude by noting some future directions: (a) when losses are generated by an adversary, and (b) when context probabilities need to be learned adaptively.

Acknowledgments

SRC is grateful to a CISE postdoctoral fellowship of Boston University.

References

- Noga Alon, Nicolo Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. In *Conference on Learning Theory*, pages 23–35. PMLR, 2015.
- Nicolo Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Minimizing regret with label efficient prediction. *IEEE Transactions on Information Theory*, 51(6):2152–2162, 2005.
- Corinna Cortes, Giulia DeSalvo, Claudio Gentile, Mehryar Mohri, and Scott Yang. Online learning with abstention. In *international conference on machine learning*, pages 1059–1067. PMLR, 2018.
- Aditya Gangrade, Anil Kag, Ashok Cutkosky, and Venkatesh Saligrama. Online selective classification with limited feedback. *Advances in Neural Information Processing Systems*, 34, 2021.
- Lihong Li, Michael L Littman, Thomas J Walsh, and Alexander L Strehl. Knows what it knows: a framework for self-aware learning. *Machine learning*, 82(3):399–443, 2011.
- Feng Nan and Venkatesh Saligrama. Adaptive classification for prediction under a budget. *Advances in neural information processing systems*, 30, 2017.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Zhixiang Xu, Matt J Kusner, Kilian Q Weinberger, Minmin Chen, and Olivier Chapelle. Classifier cascades and trees for minimizing feature evaluation cost. *The Journal of Machine Learning Research*, 15(1):2113–2144, 2014.

Appendix

Proof of Equation (2)

Fix a context $c \in [K]$. Let $I_{c,s}$ be the indicator of whether we get feedback on round s so that $\mathbb{E}[I_{c,s}] = p_c$. Then, we run OGD using losses $\widehat{f}_{c,s} = \frac{I_{c,s}}{p_c} f_{c,s}$. Note that these losses are $\frac{H}{p_c}$ smooth. Further, on rounds where $I_{c,s} = 1$, we can ignore the output of the online gradient descent and instead play $w_{c,s}^*$. Then, the regret of OGD would be

$$\mathbb{E} \left[\sum_{s=1}^{\tau_{c,T}} \widehat{f}_{c,s}(w_{c,s}) - \widehat{f}_{c,s}(u_c) \right] \leq \frac{4HD^2}{p_c} + 4D \sqrt{\frac{HL_{c,T}^*}{p_c}},$$

where the expectation is over the randomness due to abstentions. However, our regret in expectation is given by

$$\begin{aligned} \mathbb{E}[R_c(u_c)] &\leq \frac{4HD^2}{p_c} + 4D \sqrt{\frac{HL_{c,T}^*}{p_c}} + \mathbb{E} \left[\sum_{s=1}^{\tau_{c,T}} I_{c,s} (f_{c,s}(w_{c,s}^*) - f_{c,s}(w_{c,s})) \right] \\ &\leq \frac{4HD^2}{p_c} + 4D \sqrt{\frac{HL_{c,T}^*}{p_c}} - \mathbb{E} \left[\sum_{s=1}^{\tau_{c,T}} p_c f_{c,s}(w_{c,s}) \right] \\ &\leq \frac{4HD^2}{p_c} + 4D \sqrt{\frac{HL_{c,T}^*}{p_c}} - \mathbb{E} \left[\sum_{s=1}^{\tau_{c,T}} p_c (f_{c,s}(w_{c,s}) - f_{c,s}(u_c)) - p_c \sum_{s=1}^{\tau_{c,T}} f_{c,s}(u) \right] \\ &= \frac{4HD^2}{p_c} + 4D \sqrt{\frac{HL_{c,T}^*}{p_c}} - p_c \mathbb{E}[R_c(u_c)] - \mathbb{E} \left[p_c \sum_{s=1}^{\tau_{c,T}} f_{c,s}(u_c) \right]. \end{aligned}$$

Rearranging $-p_c \mathbb{E}[R_c(u_c)]$ and dividing by $1 + p_c$, we get

$$\begin{aligned} \mathbb{E}[R_c(u_c)] &\leq \frac{1}{1 + p_c} \left(\frac{4HD^2}{p_c} + 4D \sqrt{\frac{HL_{c,T}^*}{p_c}} - p_c \mathbb{E} \left[\sum_{s=1}^{\tau_{c,T}} f_{c,s}(u_c) \right] \right) \\ &\leq \frac{4HD^2}{p_c} + 4D \sqrt{\frac{HL_{c,T}^*}{p_c}} - p_c L_{c,T}^*. \end{aligned}$$

The result now follows by summing over all $c = 1, \dots, K$.

Proof of Lemma 1

Consider the random variable $Z = \sum_{s=1}^{\tau_{c,m}} f_{c,s}(u_c) - 2 \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(u_c)$. Let $J_{c,m,s}$ denote the indicator of observing context c at round s of epoch m . With an abuse of notation, we write

$$Z = \underbrace{\sum_{s=1}^{\tau_m/2} J_{c,m,s} f_{c,s}(u_c) - \sum_{s=1}^{\tau_{m-1}} J_{c,m-1,s} f_{c,s}(u_c)}_{Z_1} + \underbrace{\sum_{s=\tau_m/2+1}^{\tau_m} J_{c,m,s} f_{c,s}(u_c) - \sum_{s=1}^{\tau_{m-1}} J_{c,m-1,s} f_{c,s}(u_c)}_{Z_2}.$$

Using independence of associated random variables and $\tau_m = 2\tau_{m-1}$, we note that Z_1 is sum of τ_{m-1} of random variables, independent of each other, with each having mean zero and absolute value upper bounded by 1. Furthermore, each having variance upper bounded by $2\pi_c$. Then, by Bernstein's inequality, with probability at least $1 - \delta$, $|Z_1| \leq$

$\sqrt{2\pi_c\tau_{m-1}\log(2/\delta)} + \frac{1}{3}\log(2/\delta)$. A similar result holds for Z_2 . Combining both, we get

$$\left| \sum_{s=1}^{\tau_{c,m}} f_{c,s}(u_c) - 2 \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(u_c) \right| \leq |Z_1| + |Z_2| \leq 2\sqrt{2\pi_c\tau_{m-1}\log(2/\delta)} + \frac{2}{3}\log(2/\delta),$$

with probability at least $1 - 2\delta$. Now, by regret upper bound of OGD (1), for any $\eta_c > 0$,

$$2 \sum_{s=1}^{\tau_{c,m-1}} (f_{c,s}(w_{c,s}) - f_{c,s}(u_c)) \leq \sqrt{2} \left(\frac{D^2}{\eta_c} + 2H\eta_c \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(w_{c,s}) \right).$$

Furthermore, for the competitor $u_c = \operatorname{argmin}_{w \in W} \mathbb{E}[f_{c,s}(w)]$, we obtain

$$\sum_{s=1}^{\tau_{c,m-1}} (f_{c,s}(w_{c,s}) - f_{c,s}(u_c)) \geq \underbrace{\sum_{s=1}^{\tau_{c,m-1}} (f_{c,s}(w_{c,s}) - \mathbb{E}[f_{c,s}(w_{c,s})])}_{Z_3} + \underbrace{\sum_{s=1}^{\tau_{c,m-1}} (\mathbb{E}[f_{c,s}(u_c)] - f_{c,s}(u_c))}_{Z_4}.$$

Now, using Azuma-Hoeffding's inequality for martingale difference sequences, we obtain $Z_3 \geq -2\sqrt{2\tau_{c,m-1}\log(1/\delta)}$ with probability at least $1 - \delta$. A similar bound holds for Z_4 . Combining both using a union bound, we have with probability at least $1 - \delta$,

$$\sum_{s=1}^{\tau_{c,m-1}} (f_{c,s}(w_{c,s}) - f_{c,s}(u_c)) \geq -2\sqrt{2\tau_{c,m-1}\log(2/\delta)}.$$

Now, choose an $\eta_c > 0$ such that $\lambda_c := 1 - \sqrt{2}H\eta_c > 0$. Then, we have the following upper and lower bounds on $\sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(u)$:

$$2\lambda_c \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(w_{c,s}) - \frac{2D^2H}{1-\lambda_c} \leq 2 \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(u_c) \leq 2 \sum_{s=1}^{\tau_{c,m-1}} f_{c,s}(w_{c,s}) - 4\sqrt{2\tau_{c,m-1}\log(2/\delta)}.$$

Now, for $s = 1, 2, \dots, \tau_{m-1}$, define the random variable $X_s = \left(1 - \frac{I_{c,m-1,s}}{\hat{p}_{c,m-1}}\right) f_{c,s}(w_{c,s})$. Note that $\{X_s\}_s$ is a martingale difference sequence. Also, $|X_s|^2 \leq \left(\frac{1 - \hat{p}_{c,m-1}}{\hat{p}_{c,m-1}}\right)^2$. Hence, by Azuma-Hoeffding inequality, with probability at least $1 - \delta$,

$$\left| \sum_{s=1}^{\tau_{c,m-1}} \left(1 - \frac{I_{c,m-1,s}}{\hat{p}_{c,m-1}}\right) f_{c,s}(w_{c,s}) \right| \leq \frac{1 - \hat{p}_{c,m-1}}{\hat{p}_{c,m-1}} \sqrt{2\tau_{c,m-1}\log(2/\delta)}.$$

Putting everything together, with probability at least $1 - 4\delta$, we obtain

$$\begin{aligned} L_{c,m}^* &\leq 2 \sum_{s=1}^{\tau_{c,m-1}} \frac{I_{c,m-1,s}}{\hat{p}_{c,m-1}} f_{c,s}(w_s) + \alpha_{c,m} + 2\beta_{c,m} - 4\sqrt{2\tau_{c,m-1}\log(2/\delta)}, \\ L_{c,m}^* &\geq 2\lambda_c \sum_{s=1}^{\tau_{c,m-1}} \frac{I_{c,m-1,s}}{\hat{p}_{c,m-1}} f_{c,s}(w_s) - \alpha_{c,m} - 2\lambda_c\beta_{c,m} - 2D^2H/(1-\lambda_c). \end{aligned}$$

The result follows by taking a union bound over all contexts and replacing δ with $\delta/(4K)$.