
Online Selective Classification with Limited Feedback

Aditya Gangrade
Boston University
gangrade@bu.edu

Anil Kag
Boston University
anilkag@bu.edu

Ashok Cutkosky
Boston University
ashok@cutkosky.com

Venkatesh Saligrama
Boston University
srv@bu.edu

Abstract

Motivated by applications to resource-limited and safety-critical domains, we study selective classification in the online learning model, wherein a predictor may abstain from classifying an instance. For example, this may model an adaptive decision to invoke more resources on this instance. Two salient aspects of the setting we consider are that the data may be non-realizable, due to which abstention may be a valid long-term action, and that feedback is only received when the learner abstains, which models the fact that reliable labels are only available when the resource intensive processing is invoked.

Within this framework, we explore strategies that make few mistakes, while not abstaining too many times more than the best-in-hindsight error-free classifier from a given class. That is, the one that makes no mistakes, while abstaining the fewest number of times. We construct simple versioning-based schemes for any $\mu \in (0, 1]$, that make most T^μ mistakes while incurring $\tilde{O}(T^{1-\mu})$ excess abstention against adaptive adversaries. We further show that this dependence on T is tight, and provide illustrative experiments on realistic datasets.

1 Introduction

Consider a low-power or battery-limited edge device, such as a sensor or a smart-speaker that receives a stream of classification requests. Due to the resource limitations, such a device cannot implement modern models that are needed for accurate decisions. Instead the device has access (e.g. via an internet connection) to an accurate but resource-intensive model implemented on a cloud server, and may send queries to the cloud server in order to retain accuracy. Of course, this incurs costs such as latency and battery drain due to communication. The ideal operation of such a device should thus be to learn a rule that classifies ‘easy’ instances locally, while sending harder ones to the cloud, thus maintaining accuracy whilst minimising the net resource consumption [Xu+14; NS17].

Selective classification [Cho57; Cho70] is a classical paradigm of relevance to such settings. The setup allows a predictor to abstain from classifying some instances (without incurring a mistake). This abstention models adaptive decisions to invoke more resource-intensive methods on subtle cases, like in the above example. The solution concept is relevant widely - for instance, it is relevant to adaptively recommending further (and costly) tests rather than offering a diagnosis in a medical scenario, or to recommending a human review instead of an alarm-or-not decision in security contexts. Two aspects of such settings are of particular interest to us. Firstly, the cheaper methods are typically not sufficient to realise the true labels, due to which abstention may be a long-term necessity. Secondly, a-priori reliable labels can only be obtained by invoking the resource intensive option, and thus feedback on whether a non-abstaining decision was correct is unavailable.

We propose online selective classification, with an emphasis on ensuring very few mistakes, to account for the need for very accurate decisions. Concretely, an adversary sequentially produces contexts and labels (X_t, Y_t) , and the learner uses the X_t s to produce a decision \hat{Y}_t that may either be one of K classes, or an abstention, which we represent as \perp . Feedback in the form of Y_t is provided if and only if $\hat{Y}_t = \perp$, and the learner incurs a mistake if \hat{Y}_t was non-abstaining and did not equal Y_t .

With the emphasis on controlling the total number of mistakes, we study regrets achievable when compared to the behaviour of the best-in-hindsight error-free selective classifier from a given class - that is, one that makes no mistakes, while abstaining the fewest number of times. Notice that our situation is non-realizable, and therefore this competitor may abstain in the long-run. The two metrics of importance here are the number of mistakes the learner makes, and its excess abstention over this competitor. An effective learner must control both abstention and mistakes, and it is not enough to make one small, e.g. a learner that makes a lot of mistakes but incurs a very negative excess abstention is no good. This *simultaneous* control of two regrets raises particular challenges.

We construct a simple scheme that, when competing against finite classes, simultaneously guarantees $O(T^\mu)$ mistakes and $O(T^{1-\mu})$ excess abstentions against adaptive adversaries (for any $\mu \in [0, 1]$), and show that these rates are Pareto-tight [OR94]. We further show that against stochastic adversaries, the same rates can be attained with improved dependence of the regret bounds on the size of the class, and we also describe schemes that enjoy similar improvements against adaptive adversaries, but at the cost of the T -dependence of the regret bounds. The main schemes randomly abstain at a given rate in order to gain information, and otherwise play \hat{Y}_t consistent with the ‘version space’ of classifiers that have not been observed to make mistakes. For the adversarial case, the analysis of the scheme relies on a new ‘adversarial uniform law of large numbers’ (ALLN) to argue that such methods cannot incur too many mistakes. This ALLN uses a self-normalised martingale concentration bound, and further yields an adaptive continuous approximation guarantee for the Bernoulli-sampling sketch in the sense of Ben-Eliezer & Yogev [BY20; Al+21]. The theoretical exploration is complemented by illustrative experiments that implement our scheme on two benchmark datasets.

1.1 Related Work

Selective classification has been well studied in the batch setting, and many theoretical and methodological results have appeared [e.g. HW06; BW08; EW10; WE11; KKM12; CDM16; Lei14; GE19; GKS21]. These batch results do not have strong implications for the online setting.

Cortes et al. have studied selective classification in the online setting [Cor+18], but with two differences from our setting. Firstly, rather than individually controlling mistakes and abstentions, the regret is defined according to the Chow loss, which adds up the number of mistakes and c times the number of abstentions, where c is a fixed cost parameter. Secondly (and more importantly) it is assumed that feedback is provided only when the learner *does not* abstain, rather than only when it does. This difference arises from the underlying situations being modelled - Cortes et al. view the abstention as a decision given to a user in which case no feedback is forthcoming, while we view it as a decision to invoke further processing. Both of the scenarios are reasonable, and so both of these explorations are valid, however it is unclear what implications one set of results have for the other.

A similar decision and feedback model as ours was proposed by Li et al. in the ‘knows what it knows’ (KWIK) framework [LLWS11]. The KWIK model, however, fundamentally views abstentions as a short term action, typically arguing that only a finite number of these are made. This is viable since Li et al. study this model in an essentially realizable setting, wherein the optimal labels are known to be essentially realised by a given class - notice that in such a case, a single abstention at an instance x determines what value should be played there in the long run. Our interest however lies in the situation where this data cannot be represented in such a way, and such strategies are not viable since the labels may be noisy. Our work thus generalises the KWIK setting to non-realizable data, and to situations wherein abstention is a valid long-term action, as motivated in the introduction, by studying behaviour against competitors that may abstain.¹

While Szita and Szepesvári have extended the KWIK formulation to the agnostic case in a regression setting [SS11], this work also focuses on limiting the number of abstentions to be finite rather than

¹The KWIK model also bears other significant differences. It posits an input parameter ε , and requires that the learner either abstains, or produces an ε -accurate response. A notion of competitor is not invoked, and rather than studying regret, the number of abstentions needed to achieve this ε -accuracy is studied.

long-run abstentions. Concretely it is assumed that $Y_t = g(X_t) + \text{noise}$, for some function g , and the learner knows a class \mathcal{H} , and a bound Δ such some $h \in \mathcal{H}$ is Δ -close to g (in an appropriate norm). Using the knowledge of Δ , they describe schemes that have limited abstention, but at the cost of mistakes, by producing responses \hat{Y}_t that are up to $(2 + o(1))\Delta$ separated from Y_t . In contrast, in our formulation, contexts X_t for which no function in \mathcal{H} can represent the ground truth g well would always be abstained upon. In addition to this work, trade-offs between mistakes and abstentions in a relaxed version of the KWIK framework have been considered [ZC16; SZB10; DZ13], and in particular the agnostic case has been explored by Zhang and Chaudhuri [ZC16], but unlike our situation this relaxed KWIK model requires full-feedback to be available whether or not the learner abstains. Neu and Zhivotovskiy [NZ20] also work in this relaxed model, and show that when comparing the standard loss of a *non-abstaining* classifier against the Chow loss of an abstaining learner, regrets independent of time can be obtained.

Due to the limited feedback, our setting is related to partial-monitoring [LS20, Ch. 37]. Viewing actions as choices over functions, our setting has feedback graphs [MS11] that connect abstaining actions to every other action and themselves. The novelty with respect to partial-monitoring arises from the fact that we individually control two notions of losses, rather than a single one. It's unclear how to apply the generic partial-monitoring setup to this situation - indeed, naïvely, our game is only weakly observable in the sense of Alon et al. [ACDK15], and one would expect $\Omega(T^{2/3})$ regrets, while we can control both mistakes and excess abstention to $\tilde{O}(\sqrt{T})$. A limited feedback setting where two ‘losses’ are individually controlled is label-efficient prediction [CLS05], where a learner must query in order to get feedback. However, in our setting, abstentions are both a way to gather feedback, and also necessary to prevent mistakes. That is, our competitor may abstain regularly, but makes few mistakes, while in this prior work the competitor does not abstain, but may make many mistakes. The resulting scenario is both qualitatively and quantitatively distinct, e.g. in label-efficient prediction, the smallest symmetric rate of number of queries and excess mistakes is again $\Theta(T^{2/3})$.

2 Setting, and Problem Formulation

Setup Let \mathcal{X} be a feature space, \mathcal{Y} a finite set of labels, and \mathcal{F} a finite class of selective classifiers, which are $\mathcal{Y} \cup \{\perp\}$ valued. For simplicity, we assume that \mathcal{F} contains the all abstaining classifier (i.e. the function f_\perp such that $\forall x, f_\perp(x) = \perp$). We will denote $|\mathcal{F}| = N$. The setting may be described as a game between a learner and an adversary (or more prosaically, a data generating mechanism) proceeding in T rounds. Also for simplicity, we will assume that T is known to both the learner and the adversary in advance. The objects in this game are the context process, $X_t \in \mathcal{X}$, the label process $Y_t \in \mathcal{Y}$, the action process $\hat{Y}_t \in \mathcal{Y} \cup \{\perp\}$ and the feedback process $Z_t \in \mathcal{Y} \cup \{*\}$, where $* \notin \mathcal{Y}$ is a trivial symbol. The information sets of the adversary and learner up to the t th round are respectively $\mathcal{H}_{t-1}^{\mathcal{X}} := \{(X_s, Y_s, \hat{Y}_s) : s < t\}$, and $\mathcal{H}_{t-1}^{\mathcal{Z}} := \{(X_s, \hat{Y}_s, Z_s) : s < t\}$.

The Game For each round $t \in [1 : T]$, the adversary produces a context and a label (X_t, Y_t) on the basis its history $\mathcal{H}_{t-1}^{\mathcal{X}}$. The learner observes only the context, X_t , and on the basis of this and its history $\mathcal{H}_{t-1}^{\mathcal{Z}}$, produces an action \hat{Y}_t . We will say that this action is an abstention if $\hat{Y}_t = \perp$, and that it is a prediction otherwise. If the action was an abstention, set $Z_t = Y_t$, and otherwise to $*$. The learner then observes Z_t , and the round concludes. Notice that since Z_t is a deterministic function of Y_t and \hat{Y}_t , and since the adversary observes both, $\mathcal{H}_{t-1}^{\mathcal{Z}}$ can be deterministically generated from $\mathcal{H}_{t-1}^{\mathcal{X}}$. Due to the same reason, \hat{Y}_t and Y_t are conditionally independent given $(X_t, \mathcal{H}_{t-1}^{\mathcal{X}})$.

Adversaries are characterised by a sequence of conditional laws on (X_t, Y_t) given $\mathcal{H}_{t-1}^{\mathcal{X}}$ (and T, \mathcal{F}). In the following we will explicitly consider two classes of such laws:

- **Stochastic Adversary:** (X_t, Y_t) are drawn according to a fixed law, P , unknown to the learner, independently of $\mathcal{H}_{t-1}^{\mathcal{X}}$.
- **Adaptive Adversary:** (X_t, Y_t) are arbitrary random variables with $\mathcal{H}_{t-1}^{\mathcal{X}}$ -measurable laws.

We will denote a generic class of adversaries as \mathcal{C} .

Performance Metrics The two principal quantities of interest are the number of mistakes made by the learner, and the number of times it has abstained. We will denote these as

$$M_T := \sum_{t \leq T} \mathbb{1}\{\hat{Y}_t \notin \{\perp, Y_t\}\}, \quad \text{and} \quad A_T := \sum_{t \leq T} \mathbb{1}\{\hat{Y}_t = \perp\}.$$

As previously discussed, the performance of a learner is measured in terms of regret with respect to the best-in-hindsight abstaining classifier from \mathcal{F} that makes no mistakes, that is

$$f^* \in \arg \min_{f \in \mathcal{F}} \sum_{t \leq T} \mathbb{1}\{f(X_t) = \perp\} \quad \text{s.t.} \quad \sum_{t \leq T} \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} = 0.$$

Note that such an f^* is always realised, since the class is finite, and since it contains the all abstaining classifier. Let $A_T^* := \sum_{t \leq T} \mathbb{1}\{f^*(X_t) = \perp\}$ denote the value of the minimum above. The principal metrics of interest to us are the *abstention regret* $A_T - A_T^*$, and the *total mistakes* M_T .

Solution Concept The two performance metrics naturally involve a tradeoff - for instance, making some mistakes may allow a learner to drastically reduce its abstention regret to the point that it is negative. We pursue the trade-off between the worst possible behaviour of either regret.

Definition (Regret Achievability) For functions $\varphi, \psi : \mathbb{N}^2 \rightarrow \mathbb{R}$, we say that *expected regret bounds of (φ, ψ) are achievable against a class of adversaries \mathcal{C} if there exists a learner such that for every adversary in \mathcal{C} , $\mathbb{E}[A_T - A_T^*] \leq \varphi(T, N)$ and $\mathbb{E}[M_T] \leq \psi(T, N)$.*

As is common, we are interested in the growth rates of achievable bounds with T . We thus define

Definition (Achievable rates) we say that *asymptotic expected-regret rates of $(\alpha, \mu) \in [0, 1]^2$ are achievable against a class of adversaries \mathcal{C} if an expected regret bound of (φ, ψ) can be achieved against it for functions φ, ψ , said to be witnesses for the rate, such that*

$$\limsup_{T \rightarrow \infty} \frac{\log \varphi(T, N)}{\log T} \leq \alpha \quad \text{and} \quad \limsup_{T \rightarrow \infty} \frac{\log \psi(T, N)}{\log T} \leq \mu.$$

Notice that if (α, μ) is an achievable rate, so is (α', δ') for $\alpha' \geq \alpha, \delta' \geq \delta$. As a result, the lower boundary of the set of achievable rates is well defined, and we will refer to this as the *Pareto frontier of achievable rates*. This is equivalently characterised by the function $\underline{\alpha}(\mu) := \inf\{\alpha : (\alpha, \mu) \text{ is an achievable rate}\}$. This is well defined since $\forall \mu, (1, \mu)$ is achievable by always abstaining.

3 The Adversarial Case

We begin with the adversarial case. The scheme, called the ‘versioned uniform explorer’ (VUE) is described below, and we discuss both the motivation of the scheme, and its analysis.

The main idea underlying VUE is that any function f that is observed to make a mistake on an instance X_t (due to the learner abstaining on this instance) can be removed from future consideration, since we are only trying to match the behaviour of the competitor f^* , and clearly $f \neq f^*$ as it has made a mistake. This motivates setting up a ‘version space,’

$$\mathcal{V}_t := \left\{ f : \sum_{s < t} \mathbb{1}\{Z_s \neq *, f(X_s) \notin \{\perp, Y_s\}\} = 0 \right\},$$

the set of functions that are consistent with the observations made up to time t . Notice that $f^* \in \mathcal{V}_t$ for all t . Given \mathcal{V}_t , we can restrict to playing an action in the set $\mathcal{Y}_t := \{f(X_t) : f \in \mathcal{V}_t\} - f^*(X_t)$ lies in this set, and thus any action outside of it can be eliminated. Of course, if \mathcal{Y}_t is a singleton, then it contains $f^*(X_t)$, and we can just play it.

Next, since we are incentivised to minimise the total

number of abstentions, it behooves us to play non-abstaining actions whenever possible. However, this puts us in a bind, since feedback is produced only when we play an abstaining action. Taking inspiration from [CLS05], we abstain at a rate p by tossing a biased ‘exploratory coin’, C_t , abstaining when $C_t = 1$, and otherwise playing any non-abstaining action in $\widehat{\mathcal{Y}}_t$. Clearly, such a strategy can incur at most pT excess abstention regret in expectation. Mistakes made by this strategy are controlled via the following ‘adversarial law of large numbers’ (ALLN).

Lemma 1. Let $\{\mathcal{F}_t\}_{t=1}^\infty$ be any filtration, and $\{U_t\}_{t=1}^\infty, \{B_t\}_{t=1}^\infty$ be $\{\mathcal{F}_t\}$ -adapted binary processes, such that $B_t \sim \text{Bern}(p)$, $p < 1/2$ is jointly independent of \mathcal{F}_{t-1}, U_t for each t . Let $W_t = \sum_{s \leq t} U_s$,

Algorithm 1 VUE

- 1: **Inputs:** \mathcal{F} , Exploration rate p .
 - 2: **Initialise:** $\mathcal{V}_1 \leftarrow \mathcal{F}$.
 - 3: **for** $t \in [1 : T]$ **do**
 - 4: $\widehat{\mathcal{Y}}_t \leftarrow \{f(X_t) : f \in \mathcal{V}_t\}$.
 - 5: **if** $|\widehat{\mathcal{Y}}_t| = 1$ **then**
 - 6: $\widehat{Y}_t \leftarrow f(X_t)$ for any $f \in \mathcal{V}_t$.
 - 7: $\mathcal{V}_{t+1} \leftarrow \mathcal{V}_t$.
 - 8: **else**
 - 9: Sample $C_t \sim \text{Bern}(p)$.
 - 10: **if** $C_t = 1$ **then**
 - 11: Set $\widehat{Y}_t = \perp$, observe Y_t .
 - 12: $\mathcal{U}_t \leftarrow \{f : f(X_t) \in \{\perp, Y_t\}\}$
 - 13: $\mathcal{V}_{t+1} = \mathcal{V}_t \cap \mathcal{U}_t$.
 - 14: **else**
 - 15: Pick $\widehat{Y}_t \in \widehat{\mathcal{Y}}_t \setminus \{\perp\}$.
 - 16: $\mathcal{V}_{t+1} \leftarrow \mathcal{V}_t$.
-

and $\widetilde{W}_t = \sum_{s \leq t} U_s B_s$. For any $\delta \in (0, 1/\sqrt{\epsilon})$,

$$\mathbb{P} \left(\exists t : \widetilde{W}_t \leq 1, W_t > \frac{8 \log(1/\delta)}{p} \right) \leq \delta.$$

The above is argued in §A using a self-normalised martingale tail inequality [HRMS20]. We note that this self-normalisation is critical, and without this techniques such as Freedman’s inequality yield an extraneous \sqrt{T} factor in the bounds that is untenable for our purposes. The same argument, along with the shaping technique of Howard et al. [HRMS18] yields a Bernstein-type law of iterated logarithms that controls $|W_t - \widetilde{W}_t/p|$ at a level $\widetilde{O}(1/p + \sqrt{\widetilde{W}_t/p \log \log t})$, which should be useful more broadly. This full version (presented in §A) further shows that the ‘Bernoulli-sampler’ [BY20; Alo+21] offers a continuous approximation in the sense of Ben-Eliezer & Yogev [BY20], but with the error for sets of low incidence flattened as expected due to Bernstein’s inequality.

For our purposes, the point of Lemma 1 is to allow us to argue that no matter what the adversary does, if we uniformly abstain at a rate p , then we will ‘catch’ any mistake-prone function before it makes $O(1/p)$ mistakes. Exploiting a union bound, this in turn means that with high probability, any such function will fall out of the version space \mathcal{V}_t before it has incurred much more than $\log N/p$ mistakes. Since the label produced by Algorithm 1 must equal $f(X_t)$ for *some* f in the version space, we can infer that the number of mistakes the learner makes is at most the number of times any function in the version space is wrong. Using the Lemma yields a bound of $\widetilde{O}(1/p)$ on the number of mistakes that any functions in the version space can have ever made, and since there are only N possible functions, in total the number of mistakes the learner can make is bounded as $\widetilde{O}(N/p)$. More formally, the argument, presented in §B, argues this for a single function $f \in \mathcal{F}$ by instantiating the lemma with $\mathcal{F}_t = \sigma(\mathcal{F}_t = \sigma(\mathcal{H}_t^{\mathcal{X}}))$, $B_t = C_t$, and $U_t^f := 1\{f(X_t) \notin \{\perp, Y_t\}\}$. The resulting \widetilde{W}_t^f is the number of mistakes f is *observed* to have made, and $f \in \mathcal{V}_t$ if and only if $\widetilde{W}_t^f = 0^2$. Along with a use of Bernstein’s inequality to control A_T this yields the result below.

Theorem 2. *Algorithm 1 instantiated with $p < 1/2$, and run against an adaptive adversary, attains the following with probability at least $1 - \delta$ over the randomness of the learner and the adversary:*

$$M_T \leq \frac{9N \log(2N/\delta)}{p}$$

$$A_T - A_T^* \leq pT + \sqrt{2p(1-p)T \log(2/\delta)} + 2 \log(2/\delta).$$

In particular, taking $p = \sqrt{N/T}$ yields the symmetric regret bound

$$\max(M_T, A_T - A_T^*) \lesssim \sqrt{NT} \log(N/\delta).$$

We conclude with a few remarks.

Achievable rates: Taking $\delta = 1/T$, and varying p in $(\log T/T, 1]$ gives the rates attainable by VUE

Corollary 3. *All rates (α, μ) such that $\alpha > 0$, $\alpha + \mu > 1$ are achievable against adaptive adversaries.*

These rates are tight - as expressed in Corollary 6, rates such that $\alpha + \mu < 1$ are not achievable even against stochastic adversaries. The Pareto frontier is therefore the line $\alpha + \mu = 1$.

Dependence on N : It should be noted that the dependence on the number of functions, N , in Thm. 2 is polynomial, as opposed to the more typical logarithmic dependence on the same in online classification. The problem of characterising this dependence appears to be subtle, and we do not resolve the same. In the following section, we explore schemes that improve this aspect, but at a cost - §4 yields logarithmic dependence against stochastic adversaries, while §5 gives a scheme that has a logarithmic dependence against adaptive adversaries, but worse dependence with T .

It is worth stating that the analysis above is tight for Algorithm 1 - consider the domain $\mathcal{X} = [1 : N]$, and the class $\mathcal{F} = \{f_t : t \in [0 : N]\}$ such that $f_t(x) = \perp$ if $x \leq t$ and $= 1$ if $x > t$. Now consider an adversary that chooses a t^* in advance, and presents the contexts $1 T/N$ times, $2 T/N$ times and so on, labelling contexts smaller than t^* as 0, and contexts larger than t^* as 1. Notice that in each case, there is exactly one function in \mathcal{V}_t that does not abstain. The scheme above incurs $\Omega(pT(1 - t^*/N))$ excess abstention, and $\Omega(t^*/p)$ mistakes, and linearly large t^* form a tight example. Of course, this is not a lower bound on this problem, and the question of the optimal dependence on N remains open.

²This argument only needs control for the case $\widetilde{W}_t = 0$. The ≤ 1 in Lemma 1 is exploited in §5.2.

Hedge-Type Schemes The natural approach of proceeding by weighing the cost of abstention versus a mistake, and running a hedge-type scheme on an importance-estimate of the resulting loss does not lead to tight rates - the scheme MIXED-LOSS-PROD of §5 pursues precisely this strategy, and the worse case symmetric regret bounds that standard analyses lead to scale as $T^{2/3}$ instead of as $T^{1/2}$ as for VUE (Cor. 8). This may be due to the fact certain-error prone classifiers in \mathcal{F} may have very low abstention rates, and thus overall large weight, and it is unclear how to eliminate this behaviour.

4 The Stochastic Case

This section argues that the regret bounds of Thm. 2 can be improved to behave logarithmically in N in the stochastic setting. There are a couple of issues with Algorithm 1 that impede a better analysis in the stochastic case. The first, and obvious, one is that how \hat{Y}_t is chosen is not specified. More subtly, the fact that the scheme insists on playing non-abstaining actions whenever possible makes it difficult to control the number of mistakes without a polynomial dependence on N .

We sidestep these issues in Algorithm 2 by maintaining a law π_t on functions in \mathcal{V}_t that only depends on \mathcal{H}_{t-1}^Ω , and predicting by setting $\hat{Y}_t = f(X_t)$ for $f_t \sim \pi_t$. Notice that playing this way it is possible that we abstain on X_t even if the exploratory coin comes up tails. We control mistakes by arguing that very error-prone functions are all quickly eliminated (due to the stochasticity), and using the property that π_t does not depend on X_t to limit the mistakes incurred up to such a time. Abstention control follows by choosing π according to a strategy that favours f s with small overall abstention rate over the history. In Algorithm 2, we use a version of the PROD scheme of [CMS07] to set weights, analysed with shrinking decision sets. The following is shown along these lines in §C.

Theorem 4. *Algorithm 2, run against stochastic adversaries with $\eta = p$, attains the regret bounds*

$$\mathbb{E}[M_T] \leq 8 \frac{\log T \log(NT)}{p}, \quad \text{and} \quad \mathbb{E}[A_T - A_T^*] \leq pT + \frac{\log N}{p}.$$

We note that VUE-PROD also enjoys favourable bounds in the adversarial case - mistakes are bounded as $\tilde{O}(N/p)$, and abstention regret as in the above result. This is in contrast to simpler follow-the-versioned-leader type schemes that also satisfy similar bounds as Thm. 4 in the stochastic case. Also note that the above cannot attain rates such that $\alpha \leq 1/2$, an inefficiency introduced due to the conditional independence of π_t and X_t .

Finally, we show a lower bound. The statement equates stochastic adversaries with their laws.

Theorem 5. *If \mathcal{F} contains two functions f_1, f_2 such that there exists a point x for which $f_1(x) = \perp \neq f_2(x)$, then for every $\gamma \in [0, 1/2]$, there exists a pair of laws P_1^γ, P_2^γ such that any learner that attains $\mathbb{E}_{P_1^\gamma}[A_T - A_T^*] = K$ must incur $\mathbb{E}_{P_2^\gamma}[M_T] \geq \gamma(e^{-2\gamma K} T - K)$.*

Thus, if a (φ, ψ) regret bound with $\sup \frac{\varphi}{T} < \frac{1}{2e^2}$ is achievable, then $\varphi \cdot \psi = \Omega(T)$. Indeed, using the above with $\gamma = 1/\varphi(T, N)$, gives $\mathbb{E}_{P_1}[A_T - A_T^*] = K \leq \varphi(T, N)$, and so $\psi(T, N) \geq \mathbb{E}_{P_2}[M_T] \geq \frac{T}{\varphi(T, N)} e^{-2K/\varphi(T, N)} - 1$. This proves the following.

Corollary 6. *If $(\alpha, \mu) \in [0, 1]^2$ is such that $\alpha + \mu < 1$, then an (α, μ) regret rate is not achievable against stochastic adversaries, and, a fortiori, against adaptive adversaries.*

5 Reducing the dependence of regret bounds on N in the adversarial case

This section concentrates on improving the N -dependence of regret bounds in the adversarial case via two avenues. The first improves this dependence to $\log(N)$ by running PROD with a weighted loss, but at the cost of increasing T dependence. This holds greatest relevance when T is bounded as a polynomial of N , which is of interest because N can be quite large even in reasonable settings - e.g., a discretisation of d -dimensional hyperplanes induces $N = \exp(Cd)$. The second approach

Algorithm 2 VUE-PROD

- 1: **Inputs:** \mathcal{F}, p , Learning rate η .
 - 2: **Initialise:** $\mathcal{V}_1 \leftarrow \mathcal{F}, \forall f, w_1^f \leftarrow 1$.
 - 3: **for** $t \in [1 : T]$ **do**
 - 4: Sample $f_t \sim \pi_t = \frac{w_t^f \mathbb{1}\{f \in \mathcal{V}_t\}}{\sum_{f \in \mathcal{V}_t} w_t^f}$.
 - 5: Toss $C_t \sim \text{Bern}(p)$.
 - 6: $\hat{Y}_t \leftarrow \begin{cases} \perp & C_t = 1 \\ f_t(X_t) & C_t = 0 \end{cases}$.
 - 7: $\mathcal{V}_{t+1} \leftarrow \mathcal{V}_t$.
 - 8: **if** $C_t = 1$ **then**
 - 9: $\mathcal{U}_t \leftarrow \{f : f(X_t) \in \{\perp, Y_t\}\}$
 - 10: $\mathcal{V}_{t+1} = \mathcal{V}_t \cap \mathcal{U}_t$.
 - 11: **for** $f \in \mathcal{V}_{t+1}$ **do**
 - 12: $a_t^f \leftarrow \mathbb{1}\{f(X_t) = \perp\}$
 - 13: $w_{t+1}^f \leftarrow w_t^f \cdot (1 - \eta a_t^f)$.
-

considers the case when the set of possible contexts, i.e. \mathcal{X} is not too large. While in this case, N can be as large as $(|\mathcal{Y}| + 1)^{|\mathcal{X}|}$, we show bounds depending only linearly on $|\mathcal{X}|$.

5.1 Weighted PROD

We continue the uniform exploration, but play according to the PROD method, with the loss

$$\ell_t^f := C_t \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} + \lambda \mathbb{1}\{f(X_t) = \perp\},$$

where λ both trades-off the relative costs of mistakes and abstentions, in the vein of the fixed cost Chow loss, and accounts for the sub-sampling of the mistake loss.

The analysis of this scheme, presented in §D exploits the quadratic bound of PROD due to [CMS07] to control the sum $\mathbb{E}[pM_T + \lambda(A_T - pT)]$ by $\min_g \log N/\eta + \sum \eta(\ell_t^g)^2$, where the expectation is only over the coins C_t , and the $-pT$ term is due to the extra abstentions due to the exploratory coin. The key observation is that since f^* makes no mistakes, $\sum (\ell_t^{f^*})^2 = \lambda^2 A_T^*$, and so taking $g = f^*$, and exploiting the weight allows us to separately control the regrets in terms of A_T^* .

Theorem 7. *Algorithm 3, when run against adaptive adversaries with $\eta = 1/2$, $\lambda \leq p$, attains*

$$\mathbb{E}[M_T] \leq \frac{2 \log N}{p} + \frac{2\lambda}{p} \mathbb{E}[A_T^*], \quad \text{and} \quad \mathbb{E}[A_T - A_T^*] \leq pT + \frac{2 \log N}{\lambda}.$$

5.1.1 Rates

Theorems 4 and 7 show regret bounds with logarithmic dependence in N . The following concept separates rates attainable with this advantageous property from those with worse N -dependence.

Definition (Logarithmically Achievable Rates) *We say that rates (α, μ) are logarithmically achievable against adversaries from a class \mathcal{C} if there exists a learner that attains a (ψ, φ) -regret against such adversaries for ψ, φ that witness the rate (α, μ) , and satisfy that for every fixed T , $\max(\varphi(T, N), \psi(T, N)) = O(\text{polylog}(N))$ as $N \rightarrow \infty$.*

Since $A_T^* \leq T$, choosing $p = T^{-u}$, $\lambda = T^{-(u+v)}$ in MIXED-LOSS-PROD for any $(u, v) \in [0, 1]^2$, $u + v \leq 1$ allows us to attain rates of the form $(\alpha, \mu) = (\max(1 - u, u + v), 1 - v)$. Notice that for any fixed v , the smallest α so attainable is $1^{+v/2}$. This shows

Corollary 8. *Any rate (α, μ) such that $\alpha + \mu/2 > 1$ is logarithmically achievable against adaptive adversaries.*

The following figure illustrates the worst case achievable rate regions in the three cases considered.




Figure 1: Left shows rates achievable against adaptive adversaries. Middle and right show logarithmically achievable rates against stochastic and adaptive adversaries respectively.

Adaptive Rates Observe that if $A_T^* \asymp T^{\alpha^*}$ for some $\alpha^* < 1$, then nominally, the achievable rates can be improved. Indeed, with the parametrisation $p = T^{-u}$, $\lambda = T^{-u+v}$, we may attain rates of the form $(\alpha, \mu) = (\max(1 - u, u + v), \max(u, \alpha^* - v))$. Further, a given mistake rate μ can be attained by setting $u = \mu$, and $\alpha^* - v \leq \mu$. With these constraints, the smallest abstention rate attainable is

$$\tilde{\alpha}(\mu; \alpha_*) = \max(1 - \mu, (1 + (\alpha^* - \mu)_+)/2),$$

achieved by setting $v = (\alpha_* - \mu)_+$, $u = \min(1 - (\alpha_* - \mu)_+, 2\mu)/2$. Such rates can in fact be attained adaptively, without prior knowledge of α^* . The main bottleneck here is that the quantity A_T^* is not observable. However, every function g that is never *observed* to make a mistake satisfies $\sum (\ell_t^g)^2 = \lambda^2 \sum \mathbb{1}\{g(X_t) = \perp\}$, and such functions are identifiable given \mathcal{H}_t^Σ . Let

$$B_t^* := \min \sum_{s \leq t} \mathbb{1}\{g(X_s) = \perp\} \quad \text{s.t.} \quad \sum_{s \leq t} C_s \mathbb{1}\{g(X_s) \notin \{\perp, Y_t\}\} = 0.$$

Note that B_t^* grows monotonically, and is always smaller than $A_t^* = \sum_{s \leq t} \mathbb{1}\{f^*(X_t) = \perp\}$. We show the following in §D.1 via a scheme that adaptively sets p, λ according to B_t^* .

Theorem 9. *For any $\alpha^*, \mu, \varepsilon \in (0, 1]$, Algorithm 4 attains, without prior knowledge of α^* , any rate of the form $(\tilde{\alpha}(\mu, \alpha^*) + \varepsilon, \mu + \varepsilon)$ against adaptive adversaries that induce $A_T^* \leq T^{\alpha^*}$ almost surely.*

The rates $\tilde{\alpha}$ essentially interpolate between the second and third panels of Fig. 1. Concretely the region achieved consists of the intersection of the regions $\{\alpha > 1/2\}$, $\{\alpha + \mu > 1\}$ and $\{2\alpha + \mu > 1 + \alpha^*\}$, with the last set being active only when $\alpha^* \geq 1/2$.

5.2 A $|\mathcal{X}|$ -dependent analysis of VUE

We give an alternate mistake analyse for VUE over finite domains. The analysis is slightly stronger: let $\mathbf{y} \in ([1 : K] \cup \{\perp\})^{|\mathcal{F}|}$ be indexed by elements of \mathcal{F} , with the ‘fth’ entry \mathbf{y}_f represents a value that f might take. Consider the resulting partition of $\{\mathcal{X}_{\mathbf{y}}\}_{\mathbf{y} \in ([1:K] \cup \{\perp\})^{\mathcal{F}}}$, where each part $\mathcal{X}_{\mathbf{y}} \subset \mathcal{X}$ contains points that have the same pattern of function values, that is $\mathcal{X}_{\mathbf{y}} = \{x : \forall f \in \mathcal{F}, f(x) = \mathbf{y}_f\}$. The following argument can be run unchanged by replacing single x s in the following by all x s in one $\mathcal{X}_{\mathbf{y}}$. That is, we may replace $|\mathcal{X}|$ in the following Theorem 10 with $|\{\mathcal{X}_{\mathbf{y}}\}|$. For simplicity, we present the argument for $|\mathcal{X}|$ only.

Denote $\hat{\mathcal{Y}}_t^x := \{f(x) : f \in \mathcal{V}_t\}$. Notice that after the first time t such that $X_t = x$, $\hat{\mathcal{Y}}_t = \perp$, we will remove from the version space all classifiers that did not abstain or output the correct classification at time t . Thus if we define $y^x \in [1 : K]$ to be Y_t , then for all subsequent times, $\hat{\mathcal{Y}}_t^x \subset \{\perp, y^x\}$. As a result, if we observe two mistakes at any given x , then we cannot make any more mistakes at a subsequent time t' with $X_{t'} = x$, because the only remaining decision in $\hat{\mathcal{Y}}_{t'}^x$ must be \perp .

We may now proceed in much the same way as §3 - instantiate $U_t^x = \mathbb{1}\{X_t = x, \hat{\mathcal{Y}}_t \notin \{\perp, Y_t\}\}$, $B_t = C_t$, and union bound over the x s. Then $|\hat{\mathcal{Y}}_t^x| \geq 2$ if and only if $\tilde{W}_t^x \leq 1$, and, invoking Lemma 1, up to such a time at most $W_t^x = O(\log |\mathcal{X}|/p)$ mistakes may be made on instances such that $X_t = x$. But then totting up, we make at most $O(|\mathcal{X}| \log |\mathcal{X}|/p)$ mistakes, as encapsulated below

Theorem 10. *Algorithm 1 instantiated with $p \leq 1/2$ and run against an adaptive adversary, attains the following with probability at least $1 - \delta$ over the randomness of the learner and the adversary:*

$$M_T \leq \frac{9|\mathcal{X}| \log(2|\mathcal{X}|/\delta)}{p}$$

$$A_T - A_T^* \leq pT + \sqrt{2p(1-p)T \log(2/\delta)} + 2 \log(2/\delta).$$

Along with the bound itself, the above result makes a couple of points regarding the characterisation of N -dependence of the regrets in online selective classification. Firstly, it suggests that efficient analyses, and possibly schemes, must incorporate the structure of \mathcal{X} ; and secondly it shows that constructions that attempt to show superlogarithmic in N lower bounds must have both N and $|\mathcal{X}|$ large, and thus typical strategies placing a very rich class on a small domain will not be effective.

6 Experiments

We evaluate the performance of Algorithm 2 on two tasks - CIFAR 10 [KH09], and GAS [Ver+12] - see §E for details of implementation, and here for the relevant code. The former represents a setting where an expert can be adaptively invoked, which we treat by providing the true labels of the classes upon abstention. The second case is more explicitly an adaptive feature selection task - the GAS dataset has features from 16 sensors, and we train one model, g , on all of this data, while the selective classification task operates on data from the first 8 sensors only, and receives the output of g when abstaining. The standard accuracies of the model classes we implement are $\sim 90\%$ on CIFAR-10,

and $\sim 77\%$ on GAS. In both cases, a training set is used to learn a parameterized family of selective classifiers, $f_{\mu,t}$. The hyperparameters (μ, t) provide control over various levels of accuracy and abstention. For training, we leverage a recent method [GKS21] that yields such a parameterisation, which is discretised to get $N = 600$ of these functions to form our class \mathcal{F} . We then sequentially classify the test datasets of each of the tasks.

One subtlety with the setting is that none of the selective classifiers in \mathcal{F} actually make no mistakes. To avoid the trivialities emerging from this, we relax the versioning condition to only drop classifiers that are seen to make mistakes on at least $\varepsilon N_t + \sqrt{2\varepsilon N_t}$ mistakes at time t , where N_t is the number of times feedback was received up to time t , and the second term handles noise. Additionally, if it turns out that all functions in \mathcal{V}_t are wrong on a particular observed instance, we ignore this feedback (since such an error is unavoidable). Such variations of ‘relaxed versioning’ are natural ideas when extending the present problem to the one where the competitor may be allowed to make non-zero mistakes, although its analysis is beyond the scope of this paper. The scheme’s viability in this extended setting with only simple modifications indicates the practicality of such strategies.

Below, we take the competitor to be the function that makes the fewest mistakes, denoted as M_T^* . If there is more than one such function, we take the one that makes the fewest abstentions to get A_T^* . We measure *excess mistakes* $M_T - M_T^*$ and *excess abstentions* $A_T - A_T^*$ with respect to this competitor.

Behaviour of regrets with the length of the game Fig. 2 presents the excess mistakes as a fraction of T for the two datasets, i.e. $(M_T - M_T^*)/T$, as T , is varied. The learners are all instantiated with the exploration rate $p = 1/\sqrt{T}$. We observe that the excess abstentions are negative (or near-zero) over this range (see Fig. 4 in §E). Therefore we do not plot these below (the orange line is MMEA, see below). We note that the relative mistakes stay below $\sqrt{2 \log N/T}$, bearing out the theory.

Achievable Operating Points of Mistakes and Abstentions Fig. 3 shows the mistake and abstention rates attainable by varying p and ε , while holding T fixed at 500 (which is large enough to show long-run structure, but small enough allow fast experimentation). Concretely, we vary these linearly for 20 values of $p \in [0.015, 0.285]$, and 10 values of $\varepsilon \in [0.001, 0.046]$. The resulting values represent operating points that can be attained by a choice of p, ε . The same plot includes lines that represent the operating points when the scheme is run with $\varepsilon = 0.001$, the smallest value we take. Note that in practice, the best choices of ε, p may be data dependent, and choosing them in an online way is an interesting open problem (also see §E.6).

The Price of Being Online We characterise this in two ways beyond the excess mistakes.

- In Fig. 2, we also plot the ‘mistake-matched excess abstention’ (MMEA). This is defined as follows - if the scheme concludes with having made M_T mistakes, we find, in hindsight, the classifier that minimises the number of abstentions, subject to making at most M_T mistakes. The MMEA is the excess abstention of the learner over those of this relaxed competitor, and represents how many fewer abstentions a batch learner would make if allowed to make as many mistakes as the online learner. Notice that this MMEA remains well controlled in Fig. 2, and appears to scale as $O(\sqrt{T})$.
- In Fig. 3, we also plot the post-hoc operating points of the classifiers in \mathcal{F} as black triangles. This amounts to plotting the optimal abstentions amongst classifiers that make at most m mistakes, varying m .³ We note that the red operating points of the scheme get close to the black frontier, illustrating that the inefficiency due to being online is limited. As the time-behaviour of MMEA in Fig. 2 illustrates, the inefficiency is expected to grow sublinearly with T , and to thus vanish under amortisation.

7 Discussion

Online selective classification offers a primitive that has relevance to both safety-critical and resource-limited settings. In the paper, we highlighted the role of long-term abstentions in such situations, and studied this problem under the feedback limitation that labels are only provided when the system abstains, which is the only time high-complexity evaluation would be invoked in a selective classification system. When working with a finite class of model, we identified a simple scheme that provides a tight (in terms of T) trade-off between mistakes and excess abstentions against adaptive adversaries. We further discussed two schemes that improve upon the dependence of the same on the size of the model class - tightly against stochastic adversaries, and at the cost of some rate

³Observe that the MMEA corresponds to the horizontal distance between a red-point with m mistakes, and the left-most black point with y -coordinate under m .




Figure 2: $M_T - M_T^*$, and MMEA as fractions of T , as the number of rounds T is varied for CIFAR-10 (left) and GAS (right). The plots are averaged over 100 runs, and one-standard-deviation error regions are drawn.

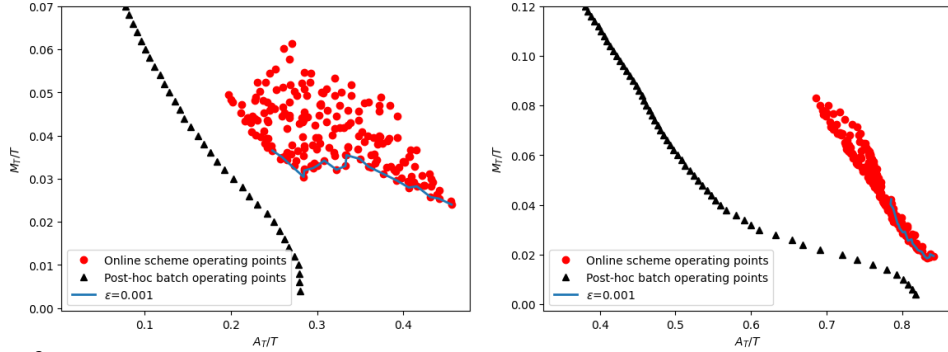


Figure 3: Operating points for our scheme as ε and p are varied are represented as red dots (for CIFAR-10 in the left, and GAS in the right). The black triangles represent operating points obtained by batch learning with the benefit of full feedback. The blue lines interpolate points obtained by varying p for $\varepsilon = 0.001$. Points are averaged over 200 runs. Note that the values are raw mistakes and abstentions, and not regrets.

performance against adaptive adversaries. Together, these schemes and analyses provide some basic foundations for the problem when competing against no-mistake models. Additionally, we carried out empirical studies that validate the scheme in the stochastic case, and demonstrate that with minor modifications, the scheme is resilient to the situation where no selective classifier in the model class is mistake-free. A number of interesting questions remain open, and we discuss a few of these below.

Perhaps the most basic question left open by the above study is how the minimax regrets against adaptive adversaries depend on N . Along with being a basic scientific question, this issue has implications for whether the results can be extended to infinite classes. Indeed, under assumptions of bounded combinatorial dimensions, the VUE-PROD and MIXED-LOSS-PROD schemes can be extended to infinite model classes, but the basic technique to do so yields trivial bounds for VUE due to the linear dependence on N . If this dependence could be improved to logarithmic, the extension to model classes with finite (multiclass versions of) Littlestone dimension would be immediate.

A practically relevant and theoretically interesting direction is online SC but where the competitor can make non-zero mistakes. This can be set up in at least two ways - either an error parameter ε is given to the learner, which must ensure that both notions of regret are small against competitors that make at most εT mistakes; or, no explicit error parameter is specified, and the learner is required to compete against the least mistake-prone model in a given set (similarly to §6). Both settings raise new challenges, since one must relax the notion of versioning used in the above work for related schema to be viable. The latter setting raises a further issue of how one can adapt to the mistake rate of the competitor. Also of practical relevance is the case where abstentions are not equally penalised, but have some variable cost. Here too, one can study variants of signalling regarding whether the cost of abstention is available before or only after an abstaining decision is made.

Finally, we observe that while tight, the random exploration technique is somewhat unsatisfying, and practically a context-adapted abstention strategy is likely to offer meaningful advantages over it. In analogy with the exploration in label-efficient prediction, one direction towards exploring context-aware methods is to study more concrete structured situations, such as linear models with noisy feedback that are popular in the investigation of online selective sampling.

Acknowledgements Our thanks to Tianrui Chen for helpful discussions.

Funding Disclosure This research was supported by the Army Research Office Grant W911NF2110246, the National Science Foundation grants CCF-2007350 and CCF-1955981, ARM Research Inc, and the Hariri Data Science Faculty Fellowship Grants. Additional revenues related to this work: AC was a visiting researcher at Google when this work was completed.

References

- [ACDK15] Noga Alon, Nicolo Cesa-Bianchi, Ofer Dekel, and Tomer Koren. “Online learning with feedback graphs: Beyond bandits”. In: *Conference on Learning Theory*. PMLR. 2015, pp. 23–35.
- [Alo+21] Noga Alon, Omri Ben-Eliezer, Yuval Dagan, Shay Moran, Moni Naor, and Eylon Yogev. “Adversarial Laws of Large Numbers and Optimal Regret in Online Classification”. In: *arXiv preprint arXiv:2101.09054* (2021).
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [BW08] Peter L Bartlett and Marten Wegkamp. “Classification with a reject option using a hinge loss”. In: *Journal of Machine Learning Research* 9.Aug (2008), pp. 1823–1840.
- [BY20] Omri Ben-Eliezer and Eylon Yogev. “The adversarial robustness of sampling”. In: *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 2020, pp. 49–62.
- [CDM16] Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. “Learning with rejection”. In: *International Conference on Algorithmic Learning Theory*. Springer. 2016, pp. 67–82.
- [Cho57] C Chow. “An optimum character recognition system using decision functions”. In: *IRE Transactions on Electronic Computers* EC-6.4 (1957), pp. 247–254.
- [Cho70] C Chow. “On optimum recognition error and reject tradeoff”. In: *IEEE Transactions on Information Theory* 16.1 (1970), pp. 41–46.
- [CLS05] Nicolo Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. “Minimizing regret with label efficient prediction”. In: *IEEE Transactions on Information Theory* 51.6 (2005), pp. 2152–2162.
- [CMS07] Nicolo Cesa-Bianchi, Yishay Mansour, and Gilles Stoltz. “Improved second-order bounds for prediction with expert advice”. In: *Machine Learning* 66.2 (2007), pp. 321–352.
- [Cor+18] Corinna Cortes, Giulia DeSalvo, Claudio Gentile, Mehryar Mohri, and Scott Yang. “Online learning with abstention”. In: *international conference on machine learning*. PMLR. 2018, pp. 1059–1067.
- [DZ13] Erik D Demaine and Morteza Zadimoghaddam. “Learning Disjunctions: Near-Optimal Trade-off between Mistakes and “I Don’t Knows””. In: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2013, pp. 1369–1379.
- [EW10] Ran El-Yaniv and Yair Wiener. “On the foundations of noise-free selective classification”. In: *Journal of Machine Learning Research* 11.May (2010), pp. 1605–1641.
- [GE19] Yonatan Geifman and Ran El-Yaniv. “SelectiveNet: A Deep Neural Network with an Integrated Reject Option”. In: *International Conference on Machine Learning*. 2019, pp. 2151–2159.
- [GKS21] Aditya Gangrade, Anil Kag, and Venkatesh Saligrama. “Selective Classification via One-Sided Prediction”. In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2021, pp. 2179–2187.
- [HRMS18] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. “Time-uniform, nonparametric, nonasymptotic confidence sequences”. In: *arXiv preprint arXiv:1810.08240* (2018).
- [HRMS20] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. “Time-uniform Chernoff bounds via nonnegative supermartingales”. In: *Probability Surveys* 17 (2020), pp. 257–317.

- [HW06] Radu Herbei and Marten Wegkamp. “Classification with reject option”. In: *The Canadian Journal of Statistics/La Revue Canadienne de Statistique* (2006), pp. 709–721.
- [HZRS16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.
- [Ide19] Yerlan Idelbayev. *Proper ResNet Implementation for CIFAR10/CIFAR100 in pytorch*. Accessed on 2020-2-28. 2019. URL: https://github.com/akamaster/pytorch_resnet_cifar10.
- [KH09] Alex Krizhevsky and Geoffrey Hinton. “Learning multiple layers of features from tiny images”. In: (2009).
- [KKM12] Adam Tauman Kalai, Varun Kanade, and Yishay Mansour. “Reliable agnostic learning”. In: *Journal of Computer and System Sciences* 78.5 (2012), pp. 1481–1495.
- [Lei14] Jing Lei. “Classification with confidence”. In: *Biometrika* 101.4 (Oct. 2014), pp. 755–769. ISSN: 0006-3444. DOI: [10.1093/biomet/asu038](https://doi.org/10.1093/biomet/asu038).
- [LLWS11] Lihong Li, Michael L Littman, Thomas J Walsh, and Alexander L Strehl. “Knows what it knows: a framework for self-aware learning”. In: *Machine learning* 82.3 (2011), pp. 399–443.
- [LS20] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [MS11] Shie Mannor and Ohad Shamir. “From Bandits to Experts: On the Value of Side-Observations”. In: *Advances in Neural Information Processing Systems*. Ed. by J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger. Vol. 24. Curran Associates, Inc., 2011. URL: <https://proceedings.neurips.cc/paper/2011/file/e1e32e235eee1f970470a3a6658dfdd5-Paper.pdf>.
- [NS17] Feng Nan and Venkatesh Saligrama. “Adaptive classification for prediction under a budget”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 4727–4737.
- [NZ20] Gergely Neu and Nikita Zhivotovskiy. “Fast rates for online prediction with abstention”. In: *Conference on Learning Theory*. PMLR. 2020, pp. 3030–3048.
- [OR94] Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.
- [RST15a] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. “Online learning via sequential complexities.” In: *J. Mach. Learn. Res.* 16.1 (2015), pp. 155–186.
- [RST15b] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. “Sequential complexities and uniform martingale laws of large numbers”. In: *Probability Theory and Related Fields* 161.1-2 (2015), pp. 111–153.
- [SS11] István Szita and Csaba Szepesvári. “Agnostic KWIK learning and efficient approximate reinforcement learning”. In: *Proceedings of the 24th Annual Conference on Learning Theory*. JMLR Workshop and Conference Proceedings. 2011, pp. 739–772.
- [SZB10] Amin Sayedi, Morteza Zadimoghaddam, and Avrim Blum. “Trading off mistakes and don’t-know predictions”. In: (2010).
- [Ver+12] Alexander Vergara, Shankar Vembu, Tuba Ayhan, Margaret A Ryan, Margie L Homer, and Ramón Huerta. “Chemical gas sensor drift compensation using classifier ensembles”. In: *Sensors and Actuators B: Chemical* 166 (2012), pp. 320–329.
- [WE11] Yair Wiener and Ran El-Yaniv. “Agnostic selective classification”. In: *Advances in neural information processing systems*. 2011, pp. 1665–1673.
- [Xu+14] Zhixiang (Eddie) Xu, Matt J. Kusner, Kilian Q. Weinberger, Minmin Chen, and Olivier Chapelle. “Classifier Cascades and Trees for Minimizing Feature Evaluation Cost”. In: *Journal of Machine Learning Research* 15 (2014), pp. 2113–2144. URL: <http://jmlr.org/papers/v15/xu14a.html>.
- [ZC16] Chicheng Zhang and Kamalika Chaudhuri. “The extended littlestone’s dimension for learning with mistakes and abstentions”. In: *Conference on Learning Theory*. PMLR. 2016, pp. 1584–1616.

Checklist

1. For all authors...

- (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] The main contributions of this paper are the setting and formulation, and the schema and their analysis. These are all accurately described in the abstract and introduction.
 - (b) Did you describe the limitations of your work? [Yes] See §3
 - (c) Did you discuss any potential negative societal impacts of your work? [No] As a primarily theoretical paper concerned with efficiency, our work does not have any meaningful negative impacts on society beyond the larger issues of how it is employed.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We have considered the potential societal impacts, and conformed to the ethical conduct guidelines
2. If you are including theoretical results...
- (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Main experimental details are discussed in §6 and §E, and code to reproduce the same is made available at <https://github.com/anilkagak2/Online-Selective-Classification>
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See §6 and §E
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See §6
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See §E
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A] We used publicly available datasets and libraries.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A] The datasets are made publicly available, and do not contain human data - The CIFAR-10 dataset contains low-resolution images of inanimate objects, while the GAS dataset consists of chemical sensor readings.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] See above.
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A An Adversarial Anytime Uniform Law of Large Numbers For Probing Binary Sequences

A.1 Proofs of Lemma 1

We begin with a simple lemma that underlies the remaining argument. Below, κ is chosen so that $\kappa''(0) = 1$.

Lemma 11. *Let $\mathcal{F}_t, U_t, B_t, W_t, \widetilde{W}_t$ be as in Lemma 1. Let $\bar{p} = 1 - p$. Then for any $\eta \in \mathbb{R}$, the process*

$$\xi_t^\eta := \exp\left(\eta(W_t - \widetilde{W}_t/p) - \kappa(\eta)V_t\right)$$

is a non-negative, \mathcal{F}_t -adapted martingale, where

$$\begin{aligned} V_t &= \frac{\bar{p}}{p}W_t, \\ \kappa(\eta) &= \frac{p}{\bar{p}}\log\left(pe^{-\eta\bar{p}/p} + \bar{p}e^\eta\right). \end{aligned}$$

Proof. The nonnegativity of ξ_t^η is trivial, and it is \mathcal{F}_t -adapted since it is a deterministic function of the adapted processes W_t, \widetilde{W}_t . We need to argue that ξ is a martingale. To this end, observe that since $W_t = \sum_{s < t} U_s, \widetilde{W}_t = \sum_{s < t} U_s B_s$,

$$\xi_t^\eta = \xi_{t-1}^\eta \cdot \exp\left(\eta U_t(1 - B_t/p - \bar{p}\kappa(\eta)/p)\right).$$

Due to the independence of B_t from $\sigma(U_t, \mathcal{F}_{t-1})$, we have

$$\begin{aligned} &\mathbb{E}[\exp(\eta U_t(1 - B_t/p)) | \mathcal{F}_{t-1}, U_t] \\ &= \left(pe^{-\eta U_t \bar{p}/p} + \bar{p}e^{\eta U_t}\right) \\ &\stackrel{*}{=} \left(pe^{-\eta \bar{p}/p} + \bar{p}e^\eta\right)^{U_t} = \exp\left(\frac{\bar{p}}{p}U_t \kappa(\eta)\right), \end{aligned}$$

where the equality marked $*$ exploits the fact that U_t is $\{0, 1\}$ -valued. Rearranging, we have

$$\mathbb{E}\left[\exp\left(\eta U_t(1 - B_t/p) - \frac{\bar{p}}{p}U_t \kappa(\eta)\right) \middle| \mathcal{F}_{t-1}, U_t\right] = 1,$$

and exploiting the tower rule, we conclude that

$$\mathbb{E}[\xi_t^\eta | \mathcal{F}_{t-1}] = \xi_{t-1}^\eta \mathbb{E}\left[\mathbb{E}\left[\exp\left(\eta U_t(1 - B_t/p) - \frac{\bar{p}}{p}U_t \kappa(\eta)\right) \middle| \mathcal{F}_{t-1}, U_t\right] \middle| \mathcal{F}_{t-1}\right] = \xi_{t-1}^\eta. \quad \square$$

The following argument heavily exploits the techniques of Howard et al. [HRMS20], and assumes familiarity with the same. It also exploits the property that only the upper tail of Δ_t is being controlled, although this is extended in the following section.

Proof of Lemma 1. We define the deviation of W_t from \widetilde{W}_t as

$$\Delta_t := W_t - \frac{\widetilde{W}_t}{p}.$$

Notice that $\Delta_0 = 1$. As a result of the above lemma, Δ_t is a 1-sub- κ process with the associated variance process V_t , in the sense of Definition 1 of Howard et al. [HRMS20]. In particular, since κ is the (normalised) cumulant generating function of a centred Bernoulli random variable taking values $\{-\bar{p}/p, 1\}$, the process is sub-binary. Further, since $p < 1/2, \bar{p}/p > 1$, and thus the process is sub-gamma, with the scale parameter $c = 0$. [HRMS20, §3.1, and Prop.2].

We can thus invoke the line-crossing inequality of Corollary 1, part c) of Howard et al., instantiated with $c = 0$ to find that for any $x, m > 0$

$$\mathbb{P}(\exists t : \Delta_t \geq x + \mathfrak{s}(x/m)(V_t - m)) \leq \exp\left(-\frac{x^2}{2m}\right),$$

where [HRMS20, Table 2]

$$\mathfrak{s}(x/m) = \frac{x}{2m}.$$

Plugging these in, we observe that

$$\mathbb{P}\left(\exists t : \Delta_t \geq \frac{x}{2} + \frac{x}{2m}V_t\right) \leq \exp\left(-\frac{x^2}{2m}\right).$$

Now notice that if $V_t \geq m$, then $x/2 + (x/2m)V_t \leq (x/m)V_t$. Therefore, we can conclude that

$$\mathbb{P}\left(\exists t : \Delta_t \geq \frac{x}{m}V_t, V_t \geq m\right) \leq \exp\left(-\frac{x^2}{2m}\right),$$

and substituting $V_t = \frac{\bar{p}}{p}W_t, \Delta_t = W_t - \widetilde{W}_t/p$,

$$\mathbb{P}\left(\exists t : \frac{\widetilde{W}_t}{p} \leq \frac{mp - x\bar{p}}{pm}W_t, W_t \geq \frac{pm}{\bar{p}}\right) \leq \exp\left(-\frac{x^2}{2m}\right).$$

Now, if we choose $m = \frac{\bar{p}}{p}(x + 1/p)$ it follows that

$$\forall W_t \geq \frac{p}{\bar{p}}m, \frac{pm - x\bar{p}}{mp}W_t \geq \frac{1}{p},$$

and thus

$$\mathbb{P}\left(\exists t : \frac{\widetilde{W}_t}{p} \leq \frac{1}{p}, W_t \geq \frac{1}{p} + x\right) \leq \exp\left(-\frac{px^2}{2(1/p + x)\bar{p}}\right),$$

and choosing $x \geq 1/p$ further ensures that

$$\mathbb{P}\left(\exists t : \widetilde{W}_t \leq 1, W_t \geq 2x\right) \leq \exp\left(-\frac{px}{4\bar{p}}\right).$$

Now, setting $x = \max\left(\frac{1}{p}, \frac{4\bar{p}}{p} \log(1/\delta)\right)$ leaves us with

$$\mathbb{P}\left(\exists t : \widetilde{W}_t \leq 1, W_t \geq \max\left(\frac{2}{p}, \frac{8\bar{p}}{p} \log(1/\delta)\right)\right) \leq \delta.$$

The conclusion follows on observing since $p < 1/2, 8\bar{p} \geq 4$, and thus, for $\log(1/\delta) \geq 1/2, \frac{2}{p} \leq 8\bar{p} \log(1/\delta)$. \square

A.2 An improved ALLN via a Self-Normalised Law of Iterated Logarithms

The line-crossing inequalities we utilised in the previous subsection can be stitched together, by picking an exponentially increasing set of x s, and optimising the m s at each, to yield a curve crossing inequality, which in effect determines a curve that the deviations are unlikely to cross. We use the results of Howard et al. [HRMS18] that produce non-asymptotic constructions.

For our purposes, note that the processes Δ_t and $-\Delta_t$ are both sub-Gamma with variance process V_t , with the scale parameters $c_+ = 0$ and $c_- = \frac{1}{3} \cdot \frac{1-2p}{p}$ respectively. The former property is useful for controlling the upper deviations of Δ_t , and the latter for the lower deviations. Note that since the scale parameter c_+ is 0, the upper tails in the following can be improved, but for ease of presentation we will just set $c = |c_+| = c_-$ in the following.

Using Theorem 1 of Howard et al. [HRMS18] twice - for Δ_t and $-\Delta_t$, and instantiating it with $\eta = e, h(k) = \frac{\pi^2 k^2}{6}$ yields that for the sub-gamma process Δ_t with scale parameter $\leq c$, and variance process V_t , and any constant $m > 0$, and for the functions

$$S_{m,\delta}(v) = 2\sqrt{v\ell_{m,\delta}(v)} + c\ell_{m,\delta}(v),$$

$$\ell_{m,\delta}(v) = \log \frac{\pi^2}{6} + 2 \log \log \frac{v}{m} + \log \frac{2}{\delta},$$

the following bound holds true

$$\mathbb{P}(\exists t : |\Delta_t| \geq \mathcal{S}_{m,\delta}(\max(V_t, m)) \leq \delta.$$

The curve $\mathcal{S}(\max(V_t, m))$ can be simplified upon observing that

$$\{\exists t : V_t \geq m, |\Delta_t| \geq \mathcal{S}_{m,\delta}(V_t)\} \subset \{\exists t : |\Delta_t| \geq \mathcal{S}_{m,\delta}(\max(V_t, m))\}.$$

With the above in hand, set $m = \bar{p}/p$, so that $W_t \geq 1 \iff V_t \geq m$, and observe that $\log(\pi^2/6) < 1$. The following bound is immediate upon recalling that $V_t = \frac{\bar{p}}{p}W_t, \Delta = W_t - \widetilde{W}_t/p$.

Theorem 12. *In the setting of Lemma 1,*

$$\mathbb{P}\left(\exists t : W_t \geq 1, |W_t - \widetilde{W}_t/p| \geq 2\sqrt{\frac{\bar{p}W_t}{p} \left(\log \frac{2e}{\delta} + 2 \log \log W_t\right)} + \frac{\log 2e/\delta + 2 \log \log W_t}{3p}\right) \leq \delta.$$

Technically, the $\log \log W_t$ is not always defined in the above. This should be read as $\log(\max(1, \log W_t))$ to handle edge cases - alternately, it can be handled by replacing $W_t \geq 1$ by $W_t \geq 3 > e$ in the above.⁴

Notice that the bound above has the correct form when taking into account the behaviour of binomial tails, which \widetilde{W}_t behaves like. Indeed, if W is some natural number valued random variable, and $\widetilde{W}|W \sim \text{Bin}(W, p)$, then Bernstein's inequality [BLM13, Ch. 2] states that

$$\mathbb{P}\left(|W - \widetilde{W}/p| \geq C\sqrt{\bar{p}\frac{W}{p} \log(2/\delta)} + C \log(2/\delta)\right) \leq \delta,$$

which entirely parallels the form of the above theorem, barring the $\log \log W_t$ blowup due to the uniformity over time.

The above analysis was inspired by studying the recent work of Ben-Eliezer and Yogev [BY20], on adversarial sketching - their goal was to maintain an estimate of the incidence of a process within a given set (and more generally, within sets in a given system) while using limited memory, and they analysed a similar sampling approach, showing via an application of Freedman's inequality that [BY20, Lemma 4.1]

$$\mathbb{P}\left(|W_T - \widetilde{W}_T/p| \geq C\sqrt{\frac{T}{p} \log(2/\delta)} + C \frac{\log(2/\delta)}{p}\right) \leq \delta.$$

This essentially amounts to using the crude bound $W_T \leq T$. The same paper, in Theorem 1.4 and associated lemmata argues that the Reservoir Sampler [BY20, §2] of size $\sim pT$ controls deviations uniformly over time at scale $\sqrt{\frac{T}{p} \log \frac{\log T}{\delta}}$, and it was asserted that the Bernoulli Sampler cannot attain such a 'continuous robustness' [BY20, §1]. The above result improves upon this in a few ways - firstly, the result applies to the simpler Bernoulli sampler, and improves the deviation control to $O(\sqrt{W_t})$ instead of $O(\sqrt{T})$. This has the further advantage that if one is concerned with the number of samples queried along with the memory, the Bernoulli sampler only queries $\sim pT$ times with high probability, while the reservoir sampler queries about $pT \log T$ times. Secondly, it shows that the Bernoulli sampler *does* offer continuous robustness, but up to a flattening of the deviation control for sets of small incidence (small W_t). Ben-Eliezer & Yogev show a number of applications of such bounds to sketching, and Alon et al. have recently applied this to tightly characterise the regret in online classification [Alo+21], using techniques of Rakhlin et al. [RST15a; RST15b]. We believe that self-normalised bounds as above can contribute to showing adaptive versions of these results.

B Analysis of VUE Against Adaptive Adversaries

This section serves to show Theorems 2 and 10. We will analyse the excess abstention, and the mistakes separately. Both deviations are controlled with probability $1 - \delta/2$, and so a union bound completes the argument. The excess abstention control is common to both, and exploits Bernstein's inequality.

⁴In a similar vein of edge-cases, if $W_t < 1 \implies W_t = 0$, then $0 \leq \widetilde{W}_t \leq W_t = 0$, and thus the bound extends to all possible values of W_t .

Proof of excess abstention bound. Notice that the procedure only abstains if $C_t = 1$ or if $\widehat{Y}_t = \{\perp\}$. In the latter case, the competitor also abstains, and thus no excess abstention is incurred. Therefore, the net excess abstention is bounded as $A_T - A_T^* \leq \sum C_t$. Now, $\sum C_t$ is a Binomial random variable with parameters T, p . By Bernstein's inequality [BLM13, Ch. 2],

$$\mathbb{P}\left(\sum C_t \geq pT + 2\sqrt{p(1-p)T \log(2/\delta)} + 2\log(2/\delta)\right) \leq \frac{\delta}{2}. \quad \square$$

We move on to bounding mistakes in a N -dependent way.

Proof of mistake bound from Theorem 2. As in the main text, consider the filtration $\{\mathcal{F}_t\} = \{\sigma(\mathcal{H}_t^{\mathcal{A}})\}$, $U_t^f := \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\}$, and consider the processes $W_t^f = \sum_{s < t} U_s^f$, $B_t = C_t$, $\widetilde{W}_t^f = U_t^f C_t$. Note that since $N \geq 2$, $\frac{\delta}{2N} \leq \frac{1}{4} \leq \frac{1}{\sqrt{e}}$.

Note that for every f , U_t^f and C_t satisfy the requirements of Lemma 1, since C_t is tossed independently of $\mathcal{H}_{t-1}^{\mathcal{A}}$. Therefore, we may invoke Lemma 1 to find that

$$\mathbb{P}\left(\exists t : \widetilde{W}_t^f = 0, W_t^f \geq \frac{8}{p} \log(2N/\delta)\right) \leq \frac{\delta}{2N},$$

and applying a union bound over $f \in \mathcal{F}$, we conclude that

$$\mathbb{P}\left(\exists t, f : \widetilde{W}_t^f = 0, W_t^f \geq \frac{8}{p} \log(2N/\delta)\right) \leq \frac{\delta}{2},$$

Notice that if \widetilde{W}_{t-1}^f is non-zero, then $f \notin \mathcal{V}_t$ since we've seen it make a mistake prior to the time t . Now define the stopping times $\tau_f := \max\{t : f \in \mathcal{V}_t\} = \max\{t : \widetilde{W}_{t-1}^f = 0\}$. We observe that

$$\begin{aligned} M_T &= \sum_t \mathbb{1}\{\widehat{Y}_t \notin \{\perp, Y_t\}\} \leq \sum_t \mathbb{1}\{\exists f \in \mathcal{V}_t : f(X_t) \notin \{\perp, Y_t\}\} \\ &\leq \sum_f \sum_t \mathbb{1}\{f \in \mathcal{V}_t, f(X_t) \notin \{\perp, Y_t\}\} \\ &= \sum_f \sum_t \mathbb{1}\{t \leq \tau_f\} U_t^f. \end{aligned}$$

Next, define the event

$$E := \left\{ \exists t, f : f \in \mathcal{V}_t, W_{t-1}^f \geq 8 \log(2N/\delta)/p \right\}.$$

Since $f \in \mathcal{V}_t \iff \widetilde{W}_{t-1}^f = 0 \iff t \leq \tau_f$. Also recall that $W_{t-1}^f = \sum_{s < t} \mathbb{1}\{f(X_s) \notin \{\perp, Y_s\}\}$. Therefore, given E^c ,

$$\sum_t \mathbb{1}\{t \leq \tau_f, f(X_t) \notin \{\perp, Y_t\}\} \leq 8 \frac{\log(2N/\delta)}{p} + 1,$$

since on E^c , $t \leq \tau_f \implies \widetilde{W}_{t-1}^f = 0 \implies \sum_{s < t} U_s^f \leq \frac{8 \log(2N/\delta)}{p}$, and the additional 1 arises since E^c does not control behaviour at τ_f . We conclude that given E^c , we have

$$M_T \leq \sum_f 9 \frac{\log(2N/\delta)}{p} = 9 \frac{N \log(2N/\delta)}{p}.$$

But E occurs with probability at most $\delta/2$, and we have shown that

$$\mathbb{P}\left(M_T > \frac{9N \log(2N/\delta)}{p}\right) \leq \frac{\delta}{2}. \quad \square$$

As discussed in §5, the \mathcal{X} -dependent argument proceeds similarly.

Proof of mistake bound from Theorem 10. Again, consider the filtration $\{\mathcal{F}_t\} = \{\sigma(\mathcal{H}_t^\mathcal{X})\}$. Define $\widehat{\mathcal{Y}}_t^x = \{f(x) : f \in \mathcal{V}_t\}$, and the process $U_t^x := \mathbb{1}\{X_t = x, \widehat{Y}_t \notin \{\perp, Y_t\}\}$, and consider the processes $W_t^x = \sum_{s < t} U_s^x, B_t = C_t, \widetilde{W}_t^x = U_t^x C_t$. Again, since $|\mathcal{X}| \geq 2, \frac{\delta}{2|\mathcal{X}|} \leq \frac{1}{4} \leq \frac{1}{\sqrt{e}}$.

Invoking Lemma 1, since C_t is tossed independently of $\mathcal{H}_{t-1}^\mathcal{X}$, we find that

$$\mathbb{P}\left(\exists t : \widetilde{W}_t^x \leq 1, W_t^x \geq \frac{8}{p} \log(2|\mathcal{X}|/\delta)\right) \leq \frac{\delta}{2|\mathcal{X}|},$$

and applying a union bound over $x \in \mathcal{X}$, we conclude that

$$\mathbb{P}\left(\exists t, x : \widetilde{W}_t^x \leq 1, W_t^x \geq \frac{8}{p} \log(2|\mathcal{X}|/\delta)\right) \leq \frac{\delta}{2},$$

Now, from the argument in the main text, $U_t^x \geq 0 \implies |\widehat{\mathcal{Y}}_t^x| \geq 2 \iff W_{t-1}^x \leq 1$. So, define the stopping times

$$\tau_x := \max\{t : |\widehat{\mathcal{Y}}_t^x| \geq 2\} = \max\{t : W_{t-1}^x \leq 1\}.$$

We have that

$$\begin{aligned} M_T &= \sum_t \mathbb{1}\{\widehat{Y}_t \notin \{\perp, Y_t\}\} \\ &= \sum_x \sum_t \mathbb{1}\{|\widehat{\mathcal{Y}}_t^x| \geq 2\} U_t^x \\ &= \sum_x \sum_t \mathbb{1}\{t \leq \tau_x\} U_t^x. \end{aligned}$$

Defining the event

$$E := \{\exists t, x : t \leq \tau_x, W_{t-1}^x \geq 8 \log(2|\mathcal{X}|/p)\},$$

we again observe that given E^c ,

$$\sum_t \mathbb{1}\{t \leq \tau_x\} U_t^x \leq 1 + 8 \frac{\log(2N/\delta)}{p},$$

since on $E^c, t \leq \tau_x \iff \widetilde{W}_{t-1}^x \leq 1 \implies \sum_{s < t-1} U_s^x \leq \frac{8 \log(2|\mathcal{X}|/\delta)}{p}$. We thus conclude that

$$M_T \leq \sum_x 9 \frac{\log(2|\mathcal{X}|/\delta)}{p} = \frac{9|\mathcal{X}| \log(2|\mathcal{X}|/\delta)}{p}.$$

But E occurs with probability at most $\delta/2$, and we have shown that

$$\mathbb{P}\left(M_T > \frac{9|\mathcal{X}| \log(2|\mathcal{X}|/\delta)}{p}\right) \leq \frac{\delta}{2}. \quad \square$$

C Stochastic Adversaries

This section contains proofs omitted from §4, and further provides a sample-and-commit based scheme that also attains tight performance in the stochastic case.

C.1 Performance of VUE-PROD

This section constitutes a proof of Theorem 4. We begin by controlling the excess abstentions.

Proof of excess abstention bound. We begin by analysing the PROD algorithm for the setting where decision sets may shrink with time. For succinctness, denote $a_t^f = \mathbb{1}\{f(X_t) = \perp\}$, $A_t^f := \sum_{s \leq t} a_s^f$.

Lemma 13. *Let π_t^f be as in Algorithm 2. If $\eta \leq 1/2$, then for any $g \in \mathcal{V}_T$, it holds that*

$$\sum_{t,f} \pi_t^f a_t^f \leq \frac{\log N}{\eta} + A_T^g + \eta \sum_{t \leq T} (a_t^g)^2.$$

Proof. We follow the standard analysis of PROD, updated slightly to account for versioning. Consider the potential $W_t := \sum_{f \in \mathcal{V}_t} w_t^f$, where recall that $w_t^f = \prod_{s < t} (1 - \eta a_s^f)$. Since the weights are always non-negative, for any $g \in \mathcal{V}_T$, we have that

$$W_{T+1} \geq \prod_{t \leq T} (1 - \eta a_t^g).$$

Therefore, we have the lower bound

$$\log \frac{W_{T+1}}{W_1} \geq -\log N + \sum \log(1 - \eta a_t^g) \geq -\log N - \sum \eta a_t^g - \sum (\eta a_t^g)^2,$$

which exploits the fact that for $z \leq 1/2$, $\log(1 - z) \geq -z - z^2$.

To upper bound the same quantity, notice that for any t ,

$$W_{t+1} = \sum_{f \in \mathcal{V}_{t+1}} w_{t+1}^f \leq \sum_{f \in \mathcal{V}_t} w_t^f (1 - \eta a_t^f) = W_t \left(1 - \eta \sum_f \pi_t^f a_t^f \right),$$

which again exploits that weights are non-negative, and that \mathcal{V}_t is a non-increasing sequence of sets. Taking ratios and bounding $\log(1 - z)$ by $-z$, and finally summing over $t = 1 : T$, we have

$$\log \frac{W_{T+1}}{W_1} = \sum_t \log \frac{W_{t+1}}{W_t} \leq -\eta \sum_t \sum_f \pi_t^f a_t^f.$$

Rearranging the inequality obtained by sandwiching $\log \frac{W_{T+1}}{W_1}$ yields the bound. \square

Note that the above lemma holds generically, for any loss $\ell_t^f \leq 1$, and any sequence of shrinking decision sets. We will exploit this fact later.

For our purposes, observe that since a_t^f is an indicator, $(a_t^f)^2 = a_t^f$. Thus, using Lemma 13 for $g = f^* \in \mathcal{V}_T$,

$$\sum_{t,f} \pi_t^f a_t^f \leq \frac{\log N}{\eta} + A_T^* + \eta A_T^*.$$

Now, the total abstention incurred by the learner is

$$A_T = \sum \mathbb{1}\{C_t = 1\} + \mathbb{1}\{C_t = 0, f_t(X_t) = \perp\}.$$

Exploiting the independence of the exploratory coin, we find that

$$\mathbb{E}[A_T] = pT + (1 - p)\mathbb{E}\left[\sum_{t,f} \pi_t^f a_t^f\right].$$

Invoking the above bound on $\sum_{t,f} \pi_t^f a_t^f$ and rearranging then yields that

$$\mathbb{E}[A_T] \leq pT + \frac{(1 - p) \log N}{\eta} + (1 - p)\mathbb{E}[A_T^*] + \eta(1 - p)\mathbb{E}[A_T^*].$$

Now, if $\eta = p$, then $\eta(1 - p) - p = -p^2 < 0$, and then exploiting that $A_T^* \geq 0$ yields the bound

$$\mathbb{E}[A_T - A_T^*] \leq pT + \frac{\log N}{p}. \quad \square$$

This leaves the mistake control. The argument we present critically relies on the law π_t^f being chosen independently of X_t , given \mathcal{H}_{t-1}^Σ . This is ultimately a source of inefficiency - for instance, if π_t^f were allowed to depend also on X_t , then we could enforce that non-abstaining actions are not played when $C_t = 0$, and drop the second $\log(N)/p$ term from the excess abstention bound. However, we were unable to show mistake control with only logarithmic dependence on N in this situation.

Proof of mistake bound. The mistake control proceeds by partitioning the class \mathcal{F} according to the mistake rates of individual \mathcal{F} s and arguing that whole groups of these are simultaneously, and quickly, eliminated from the version space without incurring too many mistakes. This fundamentally exploits the stochasticity of the setting.

To this end, define

$$\begin{aligned}\mathcal{F}_\zeta &:= \{f \in \mathcal{F} : 2^{-\zeta} \leq P(f(X_t) \notin \{?, Y_t\}) \leq 2^{1-\zeta}\} \\ \bar{\mathcal{F}}_\zeta &:= \{f \in \mathcal{F} : P(f(X_t) \notin \{?, Y_t\}) \leq 2^{-\zeta}\}.\end{aligned}$$

In the following, ζ_0 is a parameter for the purposes of analysis, that will be chosen later. Notice that $\mathcal{F} = \bigcup_{\zeta \leq \zeta_0} \mathcal{F}_\zeta \cup \bar{\mathcal{F}}_{\zeta_0}$.

We'll argue that all $f \in \mathcal{F}_\zeta$ are eliminated quickly (for small ζ). For this, it is useful to define the stopping times

$$\tau_\zeta := \max\{t : \exists f \in \mathcal{F}_\zeta \cap \mathcal{V}_t\}.$$

Notice that for any $f \in \mathcal{F}_\zeta$,

$$P(C_t = 1, f(X_t) \notin \{\perp, Y_t\}) \geq 2^{-\zeta} p.$$

As a consequence of this and the union bound, we have the following tail inequality.

Lemma 14. For any $\delta \in (0, 1)$,

$$\mathbb{P}(\exists \zeta \leq \zeta_0 : \tau_\zeta > \sigma_{\delta, \zeta_0}(\zeta)) \leq \delta,$$

where

$$\sigma_{\delta, \zeta_0}(\zeta) := \frac{2^\zeta}{p} \log(\zeta_0 N / \delta).$$

With this in hand, notice that

$$\begin{aligned}M_T &= \sum_t \sum_f \mathbb{1}\{f_t = f\} \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} \\ &= \sum_t \sum_{\zeta \leq \zeta_0} \sum_{f \in \mathcal{F}_\zeta} \mathbb{1}\{f_t = f\} \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} + \sum_t \sum_{f \in \bar{\mathcal{F}}_{\zeta_0}} \mathbb{1}\{f_t = f\} \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\}.\end{aligned}$$

Next, we observe that

$$\begin{aligned}\mathbb{E} \left[\sum_{f \in \mathcal{F}_\zeta} \mathbb{1}\{f_t = f\} \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} \middle| \mathcal{H}_{t-1}^\Sigma \right] &= \sum_{f \in \mathcal{F}_\zeta} \pi_t^f P(f(X_t) \notin \{\perp, Y_t\}) \\ &\leq 2^{1-\zeta} \pi_t(f_t \in \mathcal{F}_\zeta) \\ &\leq 2^{1-\zeta} \mathbb{1}\{t \leq \tau_\zeta\},\end{aligned}$$

where the first equality is because π_t^f is predictable given \mathcal{H}_{t-1}^Σ , the second uses the definition of \mathcal{F}_ζ , and the final inequality is because π_t is a distribution that is supported on \mathcal{V}_t , and thus has total mass at most 1, and mass 0 when $\mathcal{F}_\zeta \cap \mathcal{V}_t = \emptyset$. In much the same way, also notice that

$$\mathbb{E} \left[\sum_{f \in \bar{\mathcal{F}}_{\zeta_0}} \mathbb{1}\{f_t = f, f(X_t) \notin \{\perp, Y_t\}\} \middle| \mathcal{H}_{t-1}^\Sigma \right] \leq 2^{-\zeta_0}.$$

Exploiting both the linearity of expectations and the tower rule,

$$\begin{aligned}\mathbb{E}[M_T] &\leq \sum_t \sum_{\zeta \leq \zeta_0} 2^{1-\zeta} P(\tau_\zeta \geq t) + 2^{-\zeta_0} T \\ &\leq \sum_{\zeta \leq \zeta_0} \left(2^{1-\zeta} \sum_{t \leq \sigma_{\delta, \zeta_0}(\zeta)} 1 + \sum_{t > \sigma_{\delta, \zeta_0}(\zeta)} \delta \right) + 2^{-\zeta_0} T \\ &\leq 2\zeta_0 \frac{\log(\zeta_0 N / \delta)}{p} + 2\delta T + 2^{-\zeta_0} T.\end{aligned}$$

Now set $\zeta_0 = \lfloor \log T \rfloor$, $\delta = 1/T$. Since $\zeta_0 N / \delta \leq N^2 T^2$, we find that

$$\mathbb{E}[M_T] \leq 4 \frac{\log T \log(NT)}{p} + 4,$$

and finally since $p \leq 1$, $\frac{4}{p} \geq 4$, leading to the claimed bound (for $T \geq 3$). \square

C.2 Lower Bound

Proof of Theorem 5. Without loss of generality, assume $f_2(x) = 1$. Recall that $f_1(x) = \perp$. We describe the two adversaries -

- P_1^γ is supported on $\{(x, 1)\}$, so that for each time $X_t = x$, and the label $Y_t = 1$.
- P_2^γ is supported on $\{(x, 1), (x, 2)\}$ such that for each time $X_t = x$, while the label is drawn iid from the law $Y_t = \begin{cases} 1 & \text{w.p. } 1 - \gamma \\ 2 & \text{w.p. } \gamma \end{cases}$.

Notice that against P_1^γ , the competitor is f_2 , which attains $A_T^{(P_1^\gamma)} = 0$, while against P_2^γ , the competitor is f_1 , which attains $A_T^{(P_2^\gamma)} = T$. Observe further that since $\gamma < 1/2$, if any learner does not play \perp , it is advantageous for it to play 1 and never play 2.⁵ We thus lose no generality in assuming that the learner's actions lie in $\{\perp, 1\}$. Now, run two coupled versions of the learner, so that if these observe the same Z_t s, they produce identical actions. Feed the first of these data generated from P_1^γ , and the second of these data generated from P_2^γ .

Let η_1 be the (random) number of abstentions that the first version of the learner makes - this means that it must have played 1 $T - \eta_1$ times. Denote the number of mistakes that the second version of the learner makes as η_2 . Given η_1 , the second version gets exactly the same sequence as the first with probability $(1 - \gamma)_1^{\eta_1}$ - indeed, due to the coupling, they first abstain together, and then receive the same label with probability $1 - \gamma$. Conditioned on this, they again abstain together, and then receive the same label with probability $1 - \gamma$ and so on, η_1 times. This means that, given η_1 , and the event that they get the same sequence, the second version of the learner plays $T - \eta_1$ '1' actions. Since each of these is wrong with probability γ , independently and identically,

$$\mathbb{E}[\eta_2 | \eta_1] \geq (1 - \gamma)^{\eta_1} \gamma (T - \eta_1).$$

Notice that $(1 - \gamma)^{\eta_1}$ is a convex function of η_1 . Thus, $\mathbb{E}[(1 - \gamma)^{\eta_1}] \geq (1 - \gamma)^{\mathbb{E}[\eta_1]} = (1 - \gamma)^K$. Further, $\mathbb{E}[-(1 - \gamma)^{\eta_1} \eta_1] \geq \mathbb{E}[-\eta_1] = -K$, and finally, for $\gamma \leq 1/2$, $(1 - \gamma) \geq e^{-2\gamma}$. It follows that

$$\mathbb{E}[\eta_2] \geq (1 - \gamma)^K \gamma T - \gamma K = \gamma(e^{-2\gamma} T - K). \quad \square$$

While here, let us also comment that the proof of Corollary 6 is mildly incomplete, since the argument requires that $\varphi \geq 2$. If instead $\varphi < 2$, then notice that setting $\gamma = 1/2$ in the above, and using that $\mathbb{E}[\eta_2] \geq \gamma((1 - \gamma)^K T - K)$, we have $\psi \geq 2^{-\varphi} \frac{T}{2} - \frac{2}{2} \geq \frac{T}{8} - 1$, which grows linearly with T .

D Analysis of MIXED-LOSS-PROD Against Adaptive Adversaries

This section provides a proof of Theorem 7, and describes an adaptive variant of the same scheme, based on a doubling trick, that serves to show Theorem 9.

Proof of Theorem 7. Recall that the scheme runs PROD with the loss

$$\ell_t^f := \mathbb{1}\{C_t = 1\} \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} + \lambda \mathbb{1}\{f(X_t) = \perp\}.$$

We first observe that repeating the proof of Lemma 13 with a_t^f replaced by ℓ_t^f gives us that for any $g \in \mathcal{V}_T$,

$$\sum_{t,f} \pi_t^f \ell_t^f \leq \frac{\log N}{\eta} + \sum_t \ell_t^g + \eta \sum (\ell_t^g)^2. \quad (1)$$

⁵More formally, given any learner, we can create the better—in expectation—learner that abstains when the given one does, and predicts 1 when the given one plays something other than \perp .

Note that this relation holds given the context and label processes. For $g = f^* \in \mathcal{V}_T$, we observe that $\ell_t^{f^*} = \lambda \mathbb{1}\{f^*(X_t) = \perp\}$, since by definition f^* makes no mistakes. Instantiating the above with f^* , and noting $\sum \mathbb{1}\{f^*(X_t) = \perp\} = A_T^*$, we conclude that

$$\sum_{t,f} \pi_t^f \ell_t^f \leq \frac{\log N}{\eta} + \lambda A_T^* + \eta \lambda^2 A_T^*. \quad (2)$$

We proceed to characterise the mistakes and abstentions that the learner makes in terms of $\sum_{t,f} \ell_t^f$. To this end, notice that

$$M_T = \sum_{t,f} \mathbb{1}\{f_t = f\} \cdot \mathbb{1}\{C_t = 0\} \cdot \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\}.$$

As a result, integrating over the randomness of the algorithm, but not over the contexts or labels, we find that

$$\begin{aligned} \mathbb{E}[M_T] &= \mathbb{E} \left[\sum_{t,f} \mathbb{E}[\mathbb{1}\{f_t = f\} \mathbb{1}\{C_t = 0\} \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} | \mathcal{H}_{t-1}^{\mathfrak{A}}, X_t, Y_t] \right] \\ &= \sum_{t,f} \mathbb{E} \left[\pi_t^f (1-p) \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} \right]. \end{aligned}$$

But, observe that

$$\begin{aligned} \mathbb{E}[\pi_t^f \ell_t^f] &= \mathbb{E} \left[\mathbb{E}[\pi_t^f C_t \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} + \lambda \pi_t^f \mathbb{1}\{f(X_t) = Y_t\} | \mathcal{H}_{t-1}^{\mathfrak{A}}] \right] \\ &= \mathbb{E}[p \pi_t^f \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\}] + \lambda \mathbb{E}[\pi_t^f \mathbb{1}\{f(X_t) = \perp\}]. \end{aligned}$$

Therefore,

$$\mathbb{E}[M_T] = \sum_{t,f} \mathbb{E} \left[\frac{(1-p)}{p} \left(\pi_t^f \ell_t^f - \pi_t^f \lambda \mathbb{1}\{f(X_t) = \perp\} \right) \right]. \quad (3)$$

Further, notice that

$$A_T = \sum_t \mathbb{1}\{C_t = 1\} + \sum_{t,f} \mathbb{1}\{C_t = 0\} \mathbb{1}\{f_t = f\} \mathbb{1}\{f(X_t) = \perp\},$$

and thus,

$$\mathbb{E}[A_T] = \mathbb{E} \left[pT + (1-p) \sum_{t,f} \pi_t^f \mathbb{1}\{f(X_t) = \perp\} \right].$$

Moving the negative terms in (3) to the left hand side, and exploiting the above, we find that

$$\mathbb{E}[M_T] + \frac{\lambda}{p} \mathbb{E}[A_T - pT] = \frac{1-p}{p} \mathbb{E} \left[\sum_{t,f} \pi_t^f \ell_t^f \right],$$

where we note that both the terms $\mathbb{E}[M_T]$ and $\mathbb{E}[A_T - pT]$ are non-negative.

Exploiting the inequality 2 and the above relation, we conclude that

$$\mathbb{E}[M_T] + \mathbb{E} \left[\frac{\lambda}{p} (A_T - pT) \right] \leq \mathbb{E} \left[\frac{1-p}{p} \left(\frac{\log N}{\eta} + \lambda A_T^* + \eta \lambda^2 A_T^* \right) \right]. \quad (4)$$

The required bounds are now forthcoming. Dropping the M_T term in the left hand side of (4), and pushing the constants N, η, p, λ through the expectations,

$$\begin{aligned} \frac{\lambda}{p} \mathbb{E}[A_T - pT] &\leq \frac{(1-p) \log N}{p\eta} + \frac{(1-p)\lambda}{p} \mathbb{E}[A_T^*] + \frac{\eta(1-p)\lambda^2}{p} \mathbb{E}[A_T^*] \\ \iff \mathbb{E}[A_T - pT] &\leq \frac{(1-p) \log N}{\eta\lambda} + (1-p) \mathbb{E}[A_T^*] + \eta\lambda(1-p) \mathbb{E}[A_T^*] \\ \iff \mathbb{E}[A_T - A_T^*] &\leq pT + \frac{\log N}{\eta\lambda} + (\eta\lambda - p) \mathbb{E}[A_T^*]. \end{aligned}$$

Taking $\eta = 1/2, \lambda \leq p$, observe that the last term is negative (since $A_T^* \geq 0$). Thus, making these substitutions and dropping the final term gives the required excess abstention control.

In a similar way, dropping the $\mathbb{E}[A_T - pT]$ term in (4) gives

$$\mathbb{E}[M_T] \leq \frac{\log N}{p\eta} + \frac{\lambda(1 + \eta\lambda)}{p} \mathbb{E}[A_T^*].$$

The claim follows on setting $\eta = 1/2$, and observing that $\eta\lambda \leq 1$. \square

D.1 Adapting Rates for small A_T^*

D.1.1 Deriving the form of $\tilde{\alpha}$

We first describe a derivation of the form of $\tilde{\alpha}$. As noted, the relevant parametrisation is $p = T^{-u}, \lambda = T^{-(u+v)}$, for $u, v \geq 0$. This, with the bounds of the previous section gives the control

$$\begin{aligned} \mathbb{E}[M_T] &\leq 2T^u \log N + T^{\alpha^* - v} \\ \mathbb{E}[A_T - A_T^*] &\leq T^{1-u} + 2T^{u+v} \log N + T^{\alpha^* - u - v}. \end{aligned}$$

Notice that $\alpha^* - u - v \leq 1 - u - v \leq 1 - u$, since $\alpha^* \leq 1, v \geq 0$. Thus, we have the rate bounds

$$\begin{aligned} \mu &= \max(u, \alpha^* - v) \\ \alpha &= \max(1 - u, u + v) \end{aligned}$$

Deriving the optimal α attainable for a fixed μ then amounts to the following convex program

$$\begin{aligned} \min \max(1 - u, u + v) \\ \text{s.t. } 0 \leq u \leq \mu \\ \max(0, \alpha^* - \mu) \leq v \end{aligned}$$

Notice that the objective is a non-decreasing function of v , so the optimal choice of the same is $(\alpha^* - \mu)_+$, the smallest value it may take. This leaves us with trying to minimise $\max(1 - u, u + (\alpha^* - \mu)_+)$ for $0 \leq u \leq \mu$. The unconstrained minimum of this function occurs at $u_0 = \frac{1 - (\alpha^* - \mu)_+}{2}$, which is feasible if $\mu \geq u_0$. If on the other hand $\mu < u_0$, then the max-affine function is in the decreasing branch $1 - u$, and the optimal choice of u is just μ . Thus, the optimum is achieved at

$$\begin{aligned} v &= (\alpha^* - \mu)_+ \\ u &= \begin{cases} \frac{1 - (\alpha^* - \mu)_+}{2} & 1 - (\alpha^* - \mu)_+ \leq 2\mu \\ \mu & 1 - (\alpha^* - \mu)_+ > 2\mu \end{cases} = \frac{\min(1 - (\alpha^* - \mu)_+, 2\mu)}{2}. \end{aligned}$$

Correspondingly, $\tilde{\alpha}$ takes the form

$$\tilde{\alpha}(\mu; \alpha^*) = \begin{cases} \frac{1 + (\alpha^* - \mu)_+}{2} & 1 - (\alpha^* - \mu)_+ \leq 2\mu \\ \max(1 - \mu, \mu + (\alpha^* - \mu)_+) & 1 - (\alpha^* - \mu)_+ > 2\mu \end{cases}.$$

But,

$$1 - (\alpha^* - \mu)_+ > 2\mu \iff 1 - \mu \geq \mu + (\alpha^* - \mu)_+,$$

and therefore

$$\tilde{\alpha}(\mu; \alpha^*) = \begin{cases} \frac{1 + (\alpha^* - \mu)_+}{2} & 1 - (\alpha^* - \mu)_+ \leq 2\mu \\ 1 - \mu & 1 - (\alpha^* - \mu)_+ > 2\mu \end{cases} = \max\left(1 - \mu, \frac{1 + (\alpha^* - \mu)_+}{2}\right).$$

D.1.2 Adaptive Scheme and Proofs

We start by recalling the definition of B_t^*

$$B_t^* = \min_{f \in \mathcal{V}_t} \sum_{s \leq t} \mathbb{1}\{f(X_s) = \perp\}.$$

We will also use the term

$$\beta_t^* := \frac{\log B_t^*}{\log T}.$$

For the remainder of this section, let $\kappa := \frac{\lambda}{p}$. Recall that the optimal behaviour is attained by setting $p = T^{-u}$, $\kappa = T^{-v}$, where

$$u = \frac{\min(1 - (\alpha^* - \mu)_+, 2\mu)}{2}$$

$$v = (\alpha^* - \mu)_+.$$

Algorithm 4 essentially constitutes a doubling trick by setting p and κ in phases, which are indexed by non-negative integers, n . The scheme is parametrised by a scale parameter, θ .

- We begin in the zeroth phase, with $\kappa = 1$, $p = T^{-\min(1, 2\mu)/2}$. This phase ends when β^* first exceeds μ , at which point the first phase begins.
- At the beginning of each phase, we re-initialise the scheme.
- For $n \geq 1$, the n th phase ends when (the reinitialised) β^* first exceeds $\mu + n\theta$.
- Each time the n th phase ends, we restart the scheme, with $\kappa = T^{-(n+1)\theta}$, $p = T^{-\min(1-(n+1)\theta, 2\mu)/2}$.

Since the scheme is restarted in each phase, we may analyse each phase separately. Note that if $A_T \leq T^{\alpha^*}$ almost surely, then the index of the largest phase is at most $n^* = \lfloor \frac{(\alpha^* - \mu)_+}{\theta} \rfloor$ phases, since $\beta_t^* \leq \alpha^*$ always. For convenience, we set T_n to be the length of the n th phase. Times t_n correspond to rounds within the n th phase, and $M_{T_n}^n, A_{T_n}^n$ are the number of mistakes and abstentions incurred by the learner in the n th phase, while $A_{T_n}^{*,n}$ is the number of abstentions incurred by f^* in the n th phase.

Consider the behaviour in the n th phase. Let g_n be the function that minimises $\sum_{s_n \leq T_n} \mathbb{1}\{g(X_t) = \perp\}$, subject to $\sum_{s_n \leq T_n} C_t \mathbb{1}\{g(X_t) \notin \{\perp, Y_t\}\} = 0$, and set the value of this optimum to $B_{T_n}^{*,n}$. By exploiting inequality (1) instantiated with g_n , and setting $\eta = 1/2$, we may infer that

$$\sum_{t_n \leq T_n} \pi_{t_n}^f \ell_{t_n}^f \leq 2 \log N + p_n \kappa_n B_{T_n}^{*,n} + \frac{p_n^2 \kappa_n^2}{2} B_{T_n}^{*,n}.$$

As a result, reiterating the previous analysis over the n th phase, the number of mistakes and abstentions incurred in this phase

$$\mathbb{E}[M_{T_n}^n] \leq \frac{2 \log N}{p_n} + 2\mathbb{E}[\kappa_n B_{T_n}^{*,n}]$$

$$\mathbb{E}[A_{T_n}^n - B_{T_n}^{*,n}] \leq \mathbb{E}[p_n T_n + 2 \frac{\log N}{\kappa_n p_n}]$$

Further, notice that in each phase, $B_{T_n}^{*,n} \leq T^{\mu+(n+1)\theta}$, $\kappa_n = T^{-n\theta}$, $p_n = T^{-\min(1-n\theta, 2\mu)/2}$. Substituting these into the above bounds, we have

$$\mathbb{E}[M_{T_n}^n] \leq 2T^{\min(1-n\theta, 2\mu)/2} \log N + 2T^{\mu+\theta} \leq 4T^{\mu+\theta} \log N$$

$$\mathbb{E}[A_{T_n}^n - B_{T_n}^{*,n}] \leq T^{-\min(1-n\theta, 2\mu)/2} \mathbb{E}[T_n] + T^{n\theta + \min(1-n\theta, 2\mu)/2} \log N$$

But then, summing over the phases,

$$\mathbb{E}[M_T] = \sum_{0 \leq n \leq n^*} \mathbb{E}[M_{T_n}^n]$$

$$\leq 4T^\mu \log N \cdot (n^* + 1)T^\theta$$

$$\leq 4T^\mu \log N \cdot \frac{T^\theta}{\theta}.$$

Further,

$$\begin{aligned}
\mathbb{E}[A_T - A_T^*] &= \mathbb{E}\left[\sum_{n \leq n^*} A_{T_n}^n - A_{T_n}^{*,n}\right] \\
&\leq \mathbb{E}\left[\sum_{0 \leq n \leq n^*} A_{T_n}^n - B_{T_n}^{*,n}\right] \\
&\leq \mathbb{E}\left[\sum_{0 \leq n \leq n^*} T^{-\min(\mu, 1-n\theta/2)} T_n\right] + \log N \sum_{0 \leq n \leq n^*} T^{n\theta + \min(1-n\theta/2, \mu)} \\
&\leq \left(\sum_{n=0}^{n^*} T^{1-\min(\mu, 1-n\theta/2)} + \sum_{n=0} T^{n\theta + \min(1-n\theta/2, \mu)}\right) \log N.
\end{aligned}$$

To simplify the above, let $n_0 = \lfloor \frac{1-2\mu}{\theta} \rfloor$, so that $\min(\mu, \frac{1-n\theta}{2}) = \mu$ for $n \leq n_0$. Notice that n_0 may be bigger or smaller than n^* . We can then write the bound as

$$\frac{\mathbb{E}[A_T - T_T^*]}{\log N} \leq \sum_{n=0}^{\min(n^*, n_0)} T^{1-\mu} + \sum_{n=\min(n^*, n_0)+1}^{n^*} T^{\frac{1+n\theta}{2}} + \sum_{n=0}^{\min(n^*, n_0)} T^{n\theta + \mu} + \sum_{n=\min(n^*, n_0)+1}^{n^*} T^{\frac{1+n\theta}{2}},$$

where we interpret $\sum_{n=i}^j = 0$ for $i > j$. This can further be simplified to

$$\frac{\mathbb{E}[A_T - A_T^*]}{\log N} \leq \min(n^* + 1, n_0 + 1) T^{1-\mu} + \frac{T^\mu}{T^\theta - 1} T^{(\min(n^*, n_0)+1)\theta} + 2\mathbf{1}\{n_0 < n^*\} \frac{T^{\frac{1+(n^*+1)\theta}{2}}}{T^{\theta/2} - 1}.$$

If we further assume that θ is chosen so that $T^{\theta/2} \geq 2$, we can lower bound $T^{\theta/2} - 1 \geq T^{\theta/2}/2$, $T^\theta - 1 \geq T^\theta/2$ which gives the bound

$$\frac{\mathbb{E}[A_T - A_T^*]}{4 \log N} \leq (\min(n_0, n^*) + 1) \left(T^{1-\mu} + T^{\mu + \min(n_0, n^*)\theta} + \mathbf{1}\{n_0 < n^*\} T^{(1+n^*\theta)/2} \right),$$

from which we can derive the rate control

$$\alpha \leq \zeta(\mu, n_0, n^*, \theta) = \max(1 - \mu, \mu + \min(n_0, n^*)\theta, \mathbf{1}\{n_0 < n^*\}(1 + n^*\theta)/2)$$

The exact statement of the theorem is now straightforward to prove

Proof of Theorem 9. We run the above procedure with $\theta = \frac{2 \ln 2}{\log T}$. Notice that $T^{\theta/2} \geq 2$, and that $T^\theta/\theta \leq \frac{2}{\ln 2} \log T \leq T^\varepsilon$ for large enough T . Therefore, mistakes are controlled at $O(T^{\mu+\varepsilon})$.

Further, for the abstention control, again $\min(n^*, n_0) + 1 \leq n_0 + 1 \leq \frac{1}{\theta} = \frac{\log T}{2 \ln 2}$. Recall the abstention rate bound ζ above. It suffices to argue that $\zeta \leq \tilde{\alpha} + \theta$, since $T^\theta = 4 = O(1)$.

To this end, first notice that

$$n_0 < n^* \iff \lfloor \frac{1-2\mu}{\theta} \rfloor < \lfloor \frac{(\alpha^* - \mu)_+}{\theta} \rfloor \implies 1 - 2\mu < (\alpha^* - \mu)_+.$$

In this case,

$$\begin{aligned}
\zeta &= \max\left(1 - \mu, \mu + n_0\theta, \frac{1 + n^*\theta}{2}\right) \\
&\leq \max\left(1 - \mu, \mu + \frac{(1-2\mu)}{\theta} \cdot \theta, \frac{1 + \frac{(\alpha^* - \mu)_+}{\theta} \cdot \theta}{2}\right) \\
&= \max\left(1 - \mu, \frac{1 + (\alpha^* - \mu)_+}{2}\right) \\
&= \tilde{\alpha}(\mu; \alpha^*).
\end{aligned}$$

Algorithm 4 ADAPTIVE-MIXED-LOSS-PROD

1: **Inputs:** \mathcal{F} , Time T , Mistake rate μ , Scale θ .
2: **Initialise:** $n \leftarrow 0$; $n_{\max} \leftarrow \lceil 1/\theta \rceil$; $\forall f \in \mathcal{F}, w_1^f \leftarrow 1$; $\forall n \leq n_{\max}, \tau_n \leftarrow T$.
3: **for** $t \in [1 : T]$ **do**
4: $u \leftarrow \min(1 - n\theta, 2\mu)/2, v \leftarrow n\theta$
5: $p \leftarrow T^{-u}, \lambda \leftarrow T^{-(u+v)}$.
6: **Sample** $f_t \sim \pi_t = w_t^f / \sum w_t^f$.
7: **Toss** $C_t \sim \text{Bern}(p)$.
8: **if** $C_t = 1$ **then**
9: $\hat{Y}_t \leftarrow \perp$
10: **else**
11: $\hat{Y}_t \leftarrow f_t(X_t)$
12: $\forall f \in \mathcal{F}$, evaluate
$$\ell_t^f = C_t \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} + \lambda \mathbb{1}\{f(X_t) = \perp\}$$

13: $w_{t+1}^f \leftarrow w_t^f (1 - \eta \ell_t^f)$.
14: **Compute**
$$B^* = \min_{g \in \mathcal{F}} \sum_{\tau_n < s \leq t} \mathbb{1}\{g(X_s) = \perp\}$$

$$\text{s.t. } \sum_{\tau_n < s \leq t} C_s \mathbb{1}\{g(X_s) \notin \{\perp, Y_t\}\} = 0.$$

15: **if** $\log B^* \geq (\mu + n\theta) \log T$ **then**
16: $n \leftarrow n + 1$
17: $\tau_{n+1} \leftarrow t$
18: $\forall f \in \mathcal{F}, w_{t+1}^f \leftarrow 1$.

On the other hand, if $n_0 \geq n^*$ then we have that

$$\frac{(\alpha^* - \mu)_+}{\theta} - 1 \leq \frac{(1 - 2\mu)}{\theta} \iff \mu \leq \frac{1 + \theta - (\alpha^* - \mu)_+}{2}.$$

As a result, in this case,

$$\begin{aligned} \zeta &\leq \max(1 - \mu, \mu + n^*\theta) \\ &\leq \max(1 - \mu, \mu + (\alpha^* - \mu)_+) \\ &\leq \max\left(1 - \mu, \frac{1 + (\alpha^* - \mu)_+ + \theta}{2}\right) \\ &\leq \tilde{\alpha}(\mu; \alpha^*) + \theta/2 \end{aligned} \quad \square$$

E Details of Experiments.

N.B. Code required to reproduce the experiments is provided at <https://github.com/anilkagak2/Online-Selective-Classification>.

E.1 Dataset Details

GAS [Ver+12] dataset is a 6-way classification task based on the 16 chemical sensors data. These sensors are used to discriminate 6 gases at various levels of concentrations. The data consists of these sensor readings for over a period of 36 months divided into 10 batches. There are 13,910 data points in this dataset. We use the first 7 batches as training set and the remaining 3 batches as test set. This split results in train and test sets with 9546 and 4364 data points respectively. The gas task contains data from 16 sensors (each of which gives 8 numbers). The standard error attained by the class we use (see below) on this is $\approx 87\%$. For the selective classification task, we use only the data from

the first 8 sensors (and thus only 64 out of 128 features). The standard error attainable for this is $\approx 67\%$. Importantly, for the GAS task, the selective classification setting we study only demands matching the performance of the best classifier with the full 16-sensor data, and thus supervision for the 8-sensor function is according to this best function. To be more concrete, denote the training data as $\{(X_i^1, X_i^2, Y_i)\}$, where X^1 and X^2 are the features from the first and second 8 sensors respectively, and Y is the label. We train a classifier g on this whole dataset. Then we produce the labelled dataset $\{(X_i^1, g(X_i^1, X_i^2))\}$, and train selective classifiers on this dataset. The online problem then takes the test dataset, and gives to the learner only the X^1 features from it. If the learner abstains, then the label $Y_t = g(X_t^1, X_t^2)$ is given to the learner.

CIFAR-10 [KH09] dataset is a popular image recognition dataset that consists of 32×32 pixels RGB images of 10 classes. It contains 50,000 training and 10,000 test images. We use standard data augmentations (shifting, mirroring and mean-std gaussian normalisation) for preprocessing the datasets. The best standard error attainable for this task by the models we use (see below) is $\approx 90\%$. This experiment is more straightforward to describe- selective classifiers are trained on the whole dataset. For the online problem, the test image is supplied to the learner, and if it abstains, then the true label of that image is provided as feedback.

E.2 Training Experts

[GKS21] proposed a scheme to train classifiers with an in-built abstention option. This scheme provides a loss function, which takes a single hyper-parameter μ , and is trained as a minimax program using gradient ascent-descent. The scheme then uses the outputs of this training with a second hyper-parameter t to provide classification or abstention decisions. Therefore, the scheme utilises two hyper-parameters (μ, t) to control the classification accuracy and abstentions.

We trained selective classifiers using this scheme. As per their recommendation, we used 30 values of μ with 10 values equally spaced in $[0.01, 1]$ and remaining 20 values in the $[1, 16]$. For the threshold parameter t , we used 20 equally spaced values in $[0.2, 0.95]$. The minimax program was run with the learning rates $(10^{-4}, 10^{-6})$ for the descent and ascent respectively. Notice that the resulting set of classifiers have $20 \times 30 = 600$ functions.

Note that classification on CIFAR-10 is a relatively difficult task than GAS. Hence, we used a simpler 3-layer fully connected neural network architecture for the GAS dataset, and a Resnet32 architecture [Ide19; HZRS16] for the CIFAR-10 dataset.

E.3 Algorithm implementation, Hyper-parameters, Compute requirements

We implemented Algorithm 5 (which relaxes the versioning in 2) using Python constructs. It has three hyper-parameters: (a) T denoting the number of rounds, (b) the exploration rate p , and (c) ε controlling the mistake tolerance. For each run, the test data points were randomly permuted, and the first T of them were presented to the algorithm.

There are two main departures from the scheme in the main text. Firstly, rather than only using feedback gained when $C_t = 1$, the version space is refined whenever $\hat{Y}_t = \perp$, allowing faster learning. Secondly, the versioning is relaxed as already described, to only exclude functions that make too many mistakes, as determined by ε .

An important implementation detail is that for very small ε , the version space may get empty before the run concludes. This is particularly relevant for small values of ε . As a simple fix, we modify the versioning rule so that if the version space were to become empty at the end of a round, it is not updated (and, indeed, the state of the scheme is retained, see below).

Since our experiments are CPU compute bounded, we used a machine with two Intel Xeon 2.60 GHz CPUs providing 40 cores. Both the regret-with-varying-time experiments took about 1 hour compute time, and the operating point experiments took nearly 5 hours each.

E.4 Regret Behaviour as Time-horizon in Varied.

We use the hyperparameter $\varepsilon = 0.01$. For the sake of efficiency, we use the adaptive scheme Algorithm 6 that adapts to the time horizon, that instead varies p with the number of rounds as $p_t = \min(0.1, \frac{1}{\sqrt{t}})$, $\eta_t = p_t$. This adaptation strategy is a standard way to handle varying horizons,

Algorithm 5 VUE-PROD-RELAXED

```
1: Inputs:  $\mathcal{F}$ , Exploration rate  $p$ , Learning rate  $\eta$ , Tolerance  $\varepsilon$ .
2: Initialise:  $\mathcal{V}_1 \leftarrow \mathcal{F}; \forall t, \mathcal{U}_t \leftarrow \emptyset; \forall f \in \mathcal{F}, w_1^f \leftarrow 1, o_0^f \leftarrow 0; \text{Ctr}_0 \leftarrow 0$ .
3: for  $t \in [1 : T]$  do
4:   Sample  $f_t \sim \pi_t = \frac{w_t^f \mathbb{1}\{f \in \mathcal{V}_t\}}{\sum_{f \in \mathcal{V}_t} w_t^f}$ .
5:   Toss  $C_t \sim \text{Bern}(p)$ .
6:    $\hat{Y}_t \leftarrow \begin{cases} \perp & C_t = 1 \\ f_t(X_t) & C_t = 0 \end{cases}$ .
7:   if  $\hat{Y}_t = \perp$  then ▷ Refine the version space if the exploratory coin is heads
8:      $\text{Ctr}_t \leftarrow \text{Ctr}_{t-1} + 1$ .
9:     for  $f \in \mathcal{V}_t$  do
10:       $o_t^f \leftarrow o_{t-1}^f + \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\}$ 
11:      if  $o_t^f \leq \varepsilon \text{Ctr}_t + \sqrt{2\varepsilon \text{Ctr}_t}$  then ▷ Retain all  $f$ s that have error rate  $< \varepsilon$  w.h.p.
12:         $\mathcal{U}_t \leftarrow \mathcal{U}_t \cup \{f\}$ .
13:       $\mathcal{V}_{t+1} = \mathcal{V}_t \cap \mathcal{U}_t$ .
14:   else
15:      $\mathcal{V}_{t+1} \leftarrow \mathcal{V}_t$ .
16:      $\forall f \in \mathcal{V}_{t+1}, o_t^f \leftarrow o_{t-1}^f$ 
17:      $\text{Ctr}_t \leftarrow \text{Ctr}_{t-1}$ .
18:   if  $\mathcal{V}_{t+1} \neq \emptyset$  then ▷ Penalise Abstentions if the version space is non-empty
19:     for  $f \in \mathcal{V}_{t+1}$  do
20:        $a_t^f \leftarrow \mathbb{1}\{f(X_t) = \perp\}$ 
21:        $w_{t+1}^f \leftarrow w_t^f \cdot (1 - \eta a_t^f)$ .
22:   else ▷  $\mathcal{V}_{t+1} = \emptyset$ , and so revert the state
23:      $\mathcal{V}_{t+1} \leftarrow \mathcal{V}_t$ .
24:      $\text{Ctr}_t \leftarrow \text{Ctr}_{t-1}$ .
25:     for  $f \in \mathcal{V}_{t+1}$  do
26:        $o_t^f \leftarrow o_{t-1}^f$ .
27:        $w_{t+1}^f \leftarrow w_t^f$ .
```

and the observations obtained via this represent (and slightly overestimate) the regrets for when Algorithm 5 is run with $p = \eta = \frac{1}{\sqrt{T}}$. A major advantage is that this significantly increases the efficiency of the procedure, since instead of re-starting the experiment for each time horizon, we can now run for one single time horizon, and obtain representative values of regret at smaller horizons by recording the values at checkpoints corresponding to these. In the plots, we ran for $T = 4000$, and checkpointed every 250 rounds.

E.4.1 Excess Abstention Behaviour

As noted in the main text, the excess abstention regret for both datasets is negative. This remains consistent with the theory, and likely arises since these datasets are, of course, not the worst case distributions. The excess abstentions regret are plotted below.

E.5 Achievable Operating Points of Mistakes and Abstentions

We use Algorithm 5, instantiated with $T = 500$, and always choosing $\eta = p$. The particular values of p, ε that are scanned are, as listed in the main text, 20 equally spaced values of p in the range $[0.015, 0.285]$, and 10 equally spaced values of ε in the range $[0.001, 0.046]$, giving in total 200 values of (p, ε) pairs that are scanned over.

The post-hoc batch operating points are obtained as follows: We first find the largest value of the number of mistakes that are attained by the online learner for some choice of (p, ε) . Call this M . The values attained were $M_{\text{CIFAR}} = 50$ and $M_{\text{GAS}} = 120$. Then, we then instantiated the set $\mathcal{M}_{\text{CIFAR}} = \{2, 3, \dots, 50\}$, and for $\mathcal{M}_{\text{GAS}} = \{2, 7, \dots, 117\}$. The density was chosen lower for GAS

Algorithm 6 VUE-PROD-RELAXED-TIME-ADAPTED

```
1: Inputs:  $\mathcal{F}$ , Tolerance  $\varepsilon$ .
2: Initialise:  $\mathcal{V}_1 \leftarrow \mathcal{F}; \forall t, \mathcal{U}_t \leftarrow \emptyset; \forall f \in \mathcal{F}, w_1^f \leftarrow 1, o_0^f \leftarrow 0; \text{Ctr}_0 \leftarrow 0$ .
3: for  $t \in [1 : T]$  do
4:    $p_t \leftarrow \min(0.1, 1/\sqrt{t})$ .
5:    $\eta_t \leftarrow p_t$ .
6:   Sample  $f_t \sim \pi_t = \frac{w_t^f \mathbb{1}\{f \in \mathcal{V}_t\}}{\sum_{f \in \mathcal{V}_t} w_t^f}$ .
7:   Toss  $C_t \sim \text{Bern}(p_t)$ .
8:    $\widehat{Y}_t \leftarrow \begin{cases} \perp & C_t = 1 \\ f_t(X_t) & C_t = 0 \end{cases}$ .
9:   if  $\widehat{Y}_t = \perp$  then ▷ Refine the version space if the exploratory coin is heads
10:      $\text{Ctr}_t \leftarrow \text{Ctr}_{t-1} + 1$ .
11:     for  $f \in \mathcal{V}_t$  do
12:        $o_t^f \leftarrow o_{t-1}^f + \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\}$ 
13:       if  $o_t^f \leq \varepsilon \text{Ctr}_t + \sqrt{2\varepsilon \text{Ctr}_t}$  then ▷ Retain all  $f$ s that have error rate  $< \varepsilon$  w.h.p.
14:          $\mathcal{U}_t \leftarrow \mathcal{U}_t \cup \{f\}$ .
15:        $\mathcal{V}_{t+1} = \mathcal{V}_t \cap \mathcal{U}_t$ .
16:     else
17:        $\mathcal{V}_{t+1} \leftarrow \mathcal{V}_t$ .
18:        $\forall f \in \mathcal{V}_{t+1}, o_t^f \leftarrow o_{t-1}^f$ 
19:        $\text{Ctr}_t \leftarrow \text{Ctr}_{t-1}$ .
20:     if  $\mathcal{V}_{t+1} \neq \emptyset$  then ▷ Penalise Abstentions if the version space is non-empty
21:       for  $f \in \mathcal{V}_{t+1}$  do
22:          $a_t^f \leftarrow \mathbb{1}\{f(X_t) = \perp\}$ 
23:          $w_{t+1}^f \leftarrow w_t^f \cdot (1 - \eta_t a_t^f)$ .
24:       else ▷  $\mathcal{V}_{t+1} = \emptyset$ , and so revert the state
25:          $\mathcal{V}_{t+1} \leftarrow \mathcal{V}_t$ .
26:          $\text{Ctr}_t \leftarrow \text{Ctr}_{t-1}$ .
27:         for  $f \in \mathcal{V}_{t+1}$  do
28:            $o_t^f \leftarrow o_{t-1}^f$ .
29:            $w_{t+1}^f \leftarrow w_t^f$ .
```

for visual pleasantness. Finally, for each $m \in \mathcal{M}_*$, we run the post-hoc optimisation

$$a(m) := \min_{f \in \mathcal{F}} \sum_t \mathbb{1}\{f(X_t) = \perp\} \quad \text{s.t.} \quad \sum_t \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} \leq m.$$

The resulting points $(a(m), m)$ are plotted as black triangles.

Definition of MMEA As stated in the main text, the mistake matched competitor is defined as follows: suppose that the scheme makes M mistakes and A abstentions over a stream. If the following program is feasible, then we define

$$A^*(m) = \min_{f \in \mathcal{F}} \sum \mathbb{1}\{f(X_t) = \perp\} \quad \text{s.t.} \quad \sum \mathbb{1}\{f(X_t) \notin \{\perp, Y_t\}\} \leq M.$$

If not, then we take $A^*(M)$ to be the abstentions made by the least mistake f , which is the competitor in the rest of the section. Then we define

$$\text{MMEA} = A - A^*(M).$$

E.6 Sensitivity of the scheme to hyperparameters

Working in the setting of Figure 2, we show how the excess mistake and abstention regrets vary at $T = 4000$ (the final point) as ε is varied in Figure 5. As expected, the excess mistakes increase roughly linearly with large ε , but the data reflects subtle non-monotonicities in the same. The variation in abstentions is, as expected, essentially opposite to that of the mistakes.




Figure 4: Excess abstention regret, normalised by T , in the setting of Figure 2 for CIFAR-10 (left) and GAS (right). The plots are averaged over 100 runs, and one-standard-deviation error bars are drawn. Notice that the values are negative for GAS, and strongly dominated by the MMEA for CIFAR.




Figure 5: Sensitivity with ε of the excess mistakes (left) and excess abstention (right) regrets at $T = 4000$ for CIFAR (top) and GAS (bottom) datasets. Points are averaged over 100 runs, and one-standard-deviation error bars are included.

Similarly, in Figure 6, we show the operating points that can be achieved by varying ε for a fixed p , and by varying p for a fixed ε . We observe first that the variation with ε for a fixed p is relatively regular, with larger ε increasing mistakes but decreasing abstentions at roughly the same rate, up to small variations. On the other, the behaviour with increasing p for a fixed ε is much more subtle, and indicates that a sweet-spot of the coin-based exploration rate exists for each tolerance level.

Together, these plots indicate that the optimal tuning of ε and p together can be subtle, and exploring how one can execute the same in an online way is an interesting open problem.




Figure 6: Illustration of how operating points achieved by the scheme vary as p is changed for fixed values of ϵ (left) and as ϵ is changed for fixed values of p (right), in the CIFAR (top) and GAS (bottom) datasets. The sets of ϵ s and p s marking the traces is reduced with respect to Figure 3 for the sake of legibility. The arrow denotes the direction of increasing the varied parameter.