Simple Algorithms for Smoothed Learning

Anish Jayant

November 19, 2025

1 Introduction

Classical perspectives on learning theory focus on arguably extreme data generation models of either i.i.d. or completely adversarial data. We review each of these settings and then introduce the *smoothed* online setting, which interpolates between the two.

Definition 1 (Offline Learning). *Let $F \subset \{\mathcal{Y}^{\mathcal{X}}\}$ be a function class and let $\mathcal{P}$ be distribution over $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. Nature's ground truth distribution $P \in \mathcal{P}$ is fixed and a training set $S \subset \mathcal{Z}^n$ is sampled* independently and identically *from P. A learner, represented as a map $\pi : \mathcal{Z}^n \to \Delta(\mathcal{D})$, selects a decision function $\hat{y} \in \mathcal{D}$ at random. If, as $n \to \infty$, the expected loss of π against P tends to the best possible performance of any predictor in $\mathcal{F}$, we say that $\mathcal{F}$ is* learnable in the offline setting.

The notion introduced in Def. 1 may be formalized as a game between the player and the environment as

$$\mathcal{V}^{\text{iid}}(\mathcal{F}, n) = \inf_{\pi} \sup_{\mathcal{P}} \left(\mathbb{E}\left[L(\hat{y})\right] - \inf_{f \in \mathcal{F}} L(f) \right), \tag{1}$$

where learnability implies $\lim_{n\to\infty} \mathcal{V}^{\text{iid}}(\mathcal{F}, n) = 0$. In the binary classification case, where $\mathcal{Y} = \{0, 1\}$ and we are interested in the 0-1 loss, the properties that $\mathcal{F}$ must obey to be properly learnable (i.e., $\mathcal{D} = \mathcal{F}$) are well-studied. The representation Eq. 1 may be bounded by symmetrization to the *Rademacher complexity* of $\mathcal{F}$ and ultimately controlled in terms of the combinatorial *VC dimension*, which is the largest cardinality n such that there exist $\{x_1, x_2, ..., x_d\} \in \mathcal{X}$ shattered by $\mathcal{F}$, i.e., $\{f(x_1), f(x_2), ..., f(x_d) | f \in \mathcal{F}\} = \{\pm 1\}^d$.

In stark contrast, the online learning setting makes no guarantee about the data generating process, which may be adversarial.

Definition 2 (Online Learning). *Let $\mathcal{F} \subset \{\mathcal{Y}^{\mathcal{X}}\}$ be a function class and $\mathcal{D} \subset \{\mathcal{Y}^{\mathcal{X}}\}$ be a decision space. Over a horizon T, for each $t \in [T]$ Nature selects data z_t and the learner decides randomized map $\pi_t \in \mathcal{D}$ to simultaneously. Then, the learner suffers loss $\ell(\hat{y}_t, z_t)$ and observes z_t. If, as $T \to \infty$, the* regret against *$\mathcal{F}$,*

$$\textit{Regret}(\mathcal{F}, T) = \sum_{i=1}^{T} \ell(\hat{y}_t, z_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^{n} \ell(f, z_t)$$

is sublinear, then $\mathcal{F}$ is learnable in the online setting.

In the former case, VC theory is known to completely classify learnability. A *function class* The minimax representation is even more informative in this case, exhibiting the sequential nature

$$\mathcal{V}^{\text{seq}}(\mathcal{F}, n) = \left\langle \inf_{q_t \in \Delta(\mathcal{D})} \sup_{z_t \in \mathcal{Z}} \mathbb{E}_{\hat{y}_t \sim q_t} \right\rangle_{t=1}^{T} \left[\frac{\text{Regret}(\mathcal{F}, T)}{T} \right], \tag{2}$$

where we use $\langle a_i \rangle_{i=1}^n = a_1 a_2 \dots a_n$ to compactly represent the n round game. Using a more wasteful *skolemization* rather than symmetrization, the quantity Eq. 2 may be characterized by sequential shattering via *Littlestone dimension*, which is strictly larger than the VC dimension and often infinite, even for simple classes.

Definition 3 (Smoothed online learning)**.** *For some base measure μ, Nature plays the game described in Def. 2 but suggests a distribution on $\mathcal{Z}$ with marginal over $\mathcal{X}$ bounded in Radon-Nikodym derivative (likelihood) by σ^{-1} against μ, for $\sigma \in (0,1] \cup \{0\}$.*

If we denote the allowable distributions $\tilde{\Delta} = \left\{ p \in \Delta(Z) : \frac{\mathrm{d}p}{\mathrm{d}\mu} \leq \frac{1}{\sigma} \right\}$, Def. 3 insists that the supremum in Eq. 2 operates over $\tilde{\Delta}$ instead of point-distributions over $\mathcal{Z}$. In the case where $\sigma = 1$, we recover the offline setting, and when $\sigma = 0$, we agree that $\tilde{\Delta} = \Delta(Z)$ and recover the online setting (the best possible moves for the second player in the game does *not* involve randomization). Under this restriction, can the sequential complexity be controlled once again by VC dimension? The seminal result of [4] shows that indeed i.i.d. complexity measures robustly describe this problem, and that adversarial characterizations are brittle.

We remark that philosophically interpreting the notion of smoothed learning is a difficult question itself: the adaptive smoothed setting has interesting statistical properties but is not *unique* in its interpolation between offline and online settings. The notation in this exposition is borrowed from [5], which initially proposed adding noise to Nature's choice in 2. Another promising direction in the recent [2], studies learnability for both the function class $\mathcal{F}$ and distribution class $\mathcal{U}$ in tandem. Clearly, a singleton $\mathcal{U} = \{\mu_0\}$ recovers the offline case and $\mathcal{U} = \Delta(\mathcal{X})$ being the set of all distributions over $\mathcal{X}$ recovers the online case. They are able to subsume known learnability results in traditional smoothed online learning and somewhat generalize both VC and Littlestone measures, but it is generally harder to characterize *rates* or design algorithms from this perspective. We focus on the formalism developed in Def. 3 in the rest of this survey, which has the comparatively richest algorithmic landscape.

2 Simple Algorithms

2.1 Hedge and Coupling

We present the first simple algorithm and analysis introduced by [4], which was the first to characterize learnability in the smoothed online setting using VC dimension. The main tool is an elementary coupling theorem.

Theorem 4 (Coupling)**.** *Let $\mathcal{D}$ be an adaptive sequence of σ-smooth distributions on $\mathcal{X}$. Then, for each $k > 0$, there is a coupling of Π such that $\left(X_1, Z_1^{(1)}, ..., Z_k^{(1)}, ..., X_t, Z_1^{(t)}, ..., Z_k^{(t)}\right) \sim \Pi$ satisfy*

1. *$X_1, ..., X_t$ is distributed according to $\mathcal{D}$.*
2. *$Z_i^{(j)}$ is uniformly and independently distributed on $\mathcal{X}$.*
3. *$\left\{ Z_i^{(j)} | j \geq t, i \in [k] \right\}$ is uniformly and independently distributed on $\mathcal{X}$ conditoined on $X_1, ..., X_{t-1}$.*
4. *With probability at least $1 - t(1-\sigma)^k \geq 1 - te^{-k\sigma}$, $\{X_1, ..., X_t\} \subset \left\{ Z_i^{(j)} | i \in [k], j \in [t] \right\}$.*

Proof For ease of exposition, we focus on $|\mathcal{X}| < \infty$ and the base measure μ as uniform over $\mathcal{X}$. We construct the coupling by first drawing $Z_1^{(i)}, ..., Z_k^{(i)}$ and constructing a set $S^{(i)}$ by including each $Z_s^{(i)}$ independently with probability $\sigma \cdot \frac{\mathrm{d}p_i}{\mathrm{d}\mu}(Z_s^{(i)}) \in [0, 1]$. If $S^{(i)}$ is non-empty, we select X_i as a uniformly random element in the set, otherwise, we draw $X_i \sim p_i$. We construct this sequence iteratively from $i = 1, 2..., t$, selecting p_i according to earlier realizations compliant with $\mathcal{D}$.

The properties (2) - (4) are readily checked from this construction, where the failure event is if any of the $S^{(i)}$ are empty. To confirm that the marginal distribution for each X_i matches p_i, we calculate for any $A \subset \mathcal{X}$,

$$\Pr[Z_s^{(i)} \in A | Z_s^{(i)} \in S] = \frac{\Pr[Z_s^{(i)} \in S, Z_s^{(i)} \in A]}{\Pr[Z_s^{(i)} \in A]} = \frac{\int_{x \in A} \mathrm{d}\mu(x) \cdot \sigma \frac{\mathrm{d}p_i}{\mathrm{d}\mu}(x)}{\int_{x \in \mathcal{X}} \mathrm{d}\mu(x) \cdot \sigma \frac{\mathrm{d}p_i}{\mathrm{d}\mu}(x)} = p_i(A).$$

It is apparent that the proof holds even when the base measure is arbitrary $\mu \in \Delta(\mathcal{X})$, but we focus on the uniform case in line with [4]. Also, $\mathcal{X}$ may be infinite with slight measure-theoretic care in defining likelihood ratios. ■

The main upshot of such a coupling is transferring from adaptive smooth stochastic processes to i.i.d. samples from a fixed measure, with respect to any bounded function class $\mathcal{G}$. It follows by a simple monotonicity argument relying on $g \geq 0$,

$$\mathop{\mathbb{E}}_{\mathcal{D}}\left[\sup_{g \in \mathcal{G}} \sum_{i=1}^{T} g(X_i)\right] \leq T^2(1-\sigma)^{\frac{\alpha}{\sigma}} + \mathop{\mathbb{E}}_{\mathcal{U}}\left[\sup_{g \in \mathcal{G}} \sum_{i \in [k], j \in [T]} g\left(Z_i^{(j)}\right)\right], \tag{3}$$

where $k = \alpha/\sigma$. From here, standard concentration bounds using the VC dimension d of $\mathcal{G}$, control

Lemma 5. *For any k and $\epsilon > \frac{120d\log(4e^2/\epsilon)}{Tk}$*

$$\mathop{\mathbb{E}}_{\mathcal{U}}\left[\sup_{g \in \mathcal{G}} \sum_{i \in [k], j \in [T]} g\left(Z_i^{(j)}\right)\right] \leq 72\sqrt{\epsilon T k d \log(1/\epsilon)} + Tk\epsilon. \tag{4}$$

With just these technical tools, we are ready to present the first simple algorithm!

Theorem 6. *Let $\mathcal{H}$ be a hypothesis class of VC dimension d. There is an algorithm $\mathcal{A}$ such that, for any adaptive sequence of σ-smooth distributions $\mathcal{D}$, it achieves regret*

$$\mathbb{E}\left[\mathit{Regret}\left(\mathcal{A}, \mathcal{D}\right)\right] \leq \tilde{\mathcal{O}}\left(\sqrt{Td\log\left(\frac{T}{\sigma d}\right)} + d\log\left(\frac{T}{\sigma d}\right)\right). \tag{5}$$

Proof Construct $\mathcal{H}' \subseteq \mathcal{H}$ which is an *ϵ-cover* of $\mathcal{H}$ with respect to the uniform distribution. That is, for each $h \in \mathcal{H}$, there is some $h'_h \in \mathcal{H}'$ such that $\Pr_\mu\left[h'_h(x) \neq h(x)\right] \leq \epsilon$. It is known that, when $\mathcal{H}$ has VC dimension $d < \infty$, such a class exists of size at most $(41/\epsilon)^d$. The algorithm that achieves regret bound in Eq. (5) is simply to run Hedge with experts $\mathcal{H}'$. Using the notation $\text{Regret}_{\mathcal{C}}(\cdot, \cdot)$ to express regret against the best hindsight predictor in $\mathcal{C}$,

$$\begin{aligned}\mathbb{E}\left[\text{Regret}_{\mathcal{H}}\left(\mathcal{A}, \mathcal{D}\right)\right] &\overset{(i)}{\leq} \mathbb{E}\left[\text{Regret}_{\mathcal{H}'}\left(\mathcal{A}, \mathcal{D}\right)\right] + \mathbb{E}\left[\max_{h \in \mathcal{H}} \min_{h' \in \mathcal{H}'} \sum_{t=1}^{T} 1\left[h(X_t) \neq h'(X_t)\right]\right] \\ &\overset{(ii)}{\leq} \mathcal{O}\left(\sqrt{Td\log(1/\epsilon)}\right) + \mathbb{E}\left[\sup_{g \in \mathcal{G}} \sum_{t=1}^{T} g(X_t)\right],\end{aligned}$$

where (i) follows by triangle inequality and (ii) follows from classical guarantees of Hedge: against the best (fixed) predictor of $\mathcal{H}'$ the regret scales as $O(\sqrt{T \ln |\mathcal{H}'|})$. We denote the class $\mathcal{G} = \{h \oplus h' : h \in \mathcal{H}, h' \in \mathcal{H}'(h)\}$ to be discrepancy indicators between $\mathcal{H}$ and its cover. By combining Equation 3 and Lemma 5 setting $\alpha = 10 \log(T)$,

$$\mathbb{E}\left[\text{Regret}_{\mathcal{H}}(\mathcal{A}, \mathcal{D})\right] \leq \mathcal{O}\left(\sqrt{Td\log(1/\epsilon)} + \sqrt{\frac{\epsilon}{\sigma}T\log(T)d\log(1/\epsilon)} + T\log(T)\frac{\epsilon}{\sigma}\right)$$

$$\overset{(iii)}{\leq} \tilde{\mathcal{O}}\left(\sqrt{Td\log\left(\frac{T}{\sigma d}\right)} + d\log\left(\frac{T}{\sigma d}\right)\right),$$

after balancing with $\epsilon = \mathcal{O}\left(\frac{d\sigma}{T\log(T)}\log\left(\frac{T\log T}{d\sigma}\right)\right)$ ∎

2.2 ERM and Decoupling

Empirical risk minimization is a wildly successful paradigm in offline binary classification, however, its online version Follow-The-Leader (FTL) is known to fail fantastically. Consider the setting where there are just two experts, and the initial loss is revealed $(1/2, 1/2)$. By switching the losses between $(1, 0)$ and $(0, 1)$ on each subsequent iteration, we may guarantee that FTL has loss as $T - O(1)$, while staying faithful to a single predictor achieves $T/2$ loss, leading to $\Omega(T)$ regret. This adversarial case exploits the brittleness of empirical minimization but can be sidestepped by techniques such as FTRL and FTPL, where additional structure is introduced through convex optimization or additional randomness, to achieve optimal rates.

Intuitively, this extreme adversarial behavior is impossible in the *smoothed* online setting, where the "surprise" of new data is bounded. Is it possible that empirical risk minimization may once again become viable, even optimal? We show that the error with respect to the squared loss of a function class $\mathcal{F} \subset \left\{[0,1]^{\mathcal{X}}\right\}$ may be bounded under the guarantee that $\mathbb{E}[y_t|x_t] = f^*(x_t)$ for some $f^* \in \mathcal{F}$. Consider the simple strategy where each prediction $\hat{f}_t$ is the empirical minimizer of the sequence viewed thus far,

$$\hat{f}_t \in \arg\min_{f \in \mathcal{F}} \sum_{s=1}^{t-1} (f(X_s) - Y_s)^2. \tag{6}$$

The main result is the following bound under the realizable and σ-smooth conditions:

Theorem 7 (Main result, [3][1])**.** *Let $\mathcal{F} \subset \left\{[-1,1]^{\mathcal{X}}\right\}$ be a function class and suppose that $(X_t, Y_t)_{t\in[T]}$ is a sequence of well-specified data such that the X_t are σ-smooth with respect to some base measure μ and Y_t are conditionally ν^2-subGaussian for some $\nu \geq 0$. If the learner chooses $\hat{f}_t$ as Eq. (6), then*

$$\mathbb{E}\left[\sum_{t=1}^{T}(\hat{f}_t(X_t) - f^*(X_t))^2\right] \leq \frac{20\log^3(T)}{\sigma}\sqrt{T(1+\nu)(1+\log\mathbb{E}_\mu\left[W_{2T\log(T)/\sigma}(256\cdot\mathcal{F})\right])}.$$

where $W_m(\mathcal{C})$ denotes the Will's functional on $\mathcal{C}$ projected onto m datapoints.

In the interest of concision, we present only the most beautiful and central step from the proof of Thm. 7 (in the author's opinion!), which precisely formalizes the "bounded surprise" intuition using the following simple lemma.

[1]We remark that the dependence on σ^{-1} can be improved to $\sqrt{\sigma^{-1}}$ by more delicate analysis, [1].

Lemma 8. *Let $(a_t)_{t\in\mathbb{N}}$ be a sequence of real numbers with $a_0 = 1$ and $a_t \in [0,1]$ for all $t > 0$. For $K > 0$ and $t \in \mathbb{N}$, let*

$$B_t(a, K) = \left\{ s < t : a_s \geq \frac{K}{s} \sum_{u<s} a_u \right\}.$$

Then, for any $\epsilon \in (0,1)$, it holds that $\frac{1}{T}|B_T(a,K)| \leq \epsilon$ as long as $K \geq \frac{2\log(T)}{\epsilon}$.

Using this lemma, it is possible to *decouple* the adversary and player using a *tangent sequence* of data instead, opening the gates for more standard concentration approaches.

Lemma 9 (Smooth Decoupling)**.** *Let $(X_t) \subset \mathcal{X}$ be a sequence of random variables and let $g_t : \mathcal{X} \to [0,1]$ be a sequence of random function adapted to filtration $(\mathcal{H}_t)_{t\geq 0}$ such that g_t is $\mathcal{H}_{t-1}$-measurable and $X_t|(\mathcal{H}_{t-1}, g_t)$ is σ-smooth with respect to some measure μ. Then, for a tangent sequence X'_s generated using the same filtration,*

$$\mathbb{E}\left[\sum_{t=1}^{T} g_t(X_t)\right] \leq \frac{\log^2(T)}{\sigma}\sqrt{2T \cdot \mathbb{E}\left[\sum_{t=1}^{T}\frac{1}{t}\sum_{s=1}^{t-1} g_t(X'_s)\right]}$$

Proof We first change the measure of adversarial data to the base measure μ.

$$\begin{aligned}
\mathop{\mathbb{E}}_{X_1,\ldots X_T}\left[\sum_{t=1}^{T} g_t(X_t)\right] &= \mathbb{E}\left[\sum_{t=1}^{T} \mathop{\mathbb{E}}_{X_t\sim p_t}[g_t(X_t)|g_t, \mathcal{H}_{t-1}]\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \mathop{\mathbb{E}}_{Z\sim\mu}\left[\frac{dp_t}{d\mu}(Z) g_t(Z)|g_t, \mathcal{H}_{t-1}\right]\right] \\
&= \mathbb{E}_Z \mathbb{E}_{g_t}\left[\sum_{t=1}^{T} \frac{dp_t}{d\mu}(Z) g_t(Z)\right].
\end{aligned}$$

We emphasize that Z is a single draw from μ and independent of $X_1, \ldots, X_T$. Using Lemma 8, we can split the summation by inspecting the sequence $a_t(Z) = \sigma \cdot \frac{dp_t}{d\mu}(Z) \in [0,1]$ and noting that $|B_T(a(Z), K)| \leq \epsilon T$ whenever $K \geq 2\log(T)/\epsilon$. Thus,

$$\begin{aligned}
\mathbb{E}_Z \mathbb{E}_{g_t}\left[\sum_{t=1}^{T} \frac{dp_t}{d\mu}(Z) g_t(Z)\right] &= \mathbb{E}_Z \mathbb{E}_{g_t}\left[\sum_{t=1}^{T} \frac{dp_t}{d\mu}(Z) g_t(Z) 1\left[t \in B_T(a(Z), K)\right]\right] \\
&\quad + \mathbb{E}_Z \mathbb{E}_{g_t}\left[\sum_{t=1}^{T} \frac{dp_t}{d\mu}(Z) g_t(Z) 1\left[t \notin B_T(a(Z), K)\right]\right] \\
&\overset{(i)}{\leq} \frac{1}{\sigma}\mathbb{E}_Z\left[\sum_{t=1}^{T} 1\left[t \in B_T(a(Z), K)\right]\right] + \mathbb{E}_Z \mathbb{E}_{g_t}\left[\sum_{t=1}^{T} K\tilde{p}_t(Z) g_t(Z) + \frac{K}{\sigma t}\right] \\
&\overset{(ii)}{\leq} \frac{\epsilon T}{\sigma} + \frac{K\log(T)}{\sigma} + K \cdot \mathbb{E}_{g_t}\left[\sum_{t=1}^{T}\frac{1}{t}\sum_{s=1}^{t-1} g_t(X'_s)\right]
\end{aligned}$$

where in (i) we used the boundedness of g_t and σ-smoothness and in (ii) we applied Lemma 8 where $\tilde{p}_t$ is shorthand for $\frac{1}{t}\sum_{s=1}^{t-1}\frac{dp_s}{d\mu}$. From the independence of Z from $X_1, \ldots, X_T$ emerges the tangent sequence X'_s. The statement is shown by balancing $\epsilon = \sqrt{\frac{2}{T} \cdot \mathbb{E}\left[\sum_{t=1}^{T}\frac{1}{t}\sum_{s=1}^{t-1} g_t(X'_s)\right]}$. ■

2.3 Hedge and ERM, Coupling and Decoupling!

Results in the previous section are highly specialized in that they study the square loss and realizability in the form of a guarantee that $Y_t = f(X_t) + \eta_t$ for η being ν^2-subGaussian. An astonishingly simple algorithm Cover achieves sublinear $\tilde{\mathcal{O}}\left(T^{2/3}\right)$ (oblivious) smoothed regret by running Hedge on a dynamically created *cover* of experts. Further, it holds in the *agnostic* setting, where there is no guarantee on the marginal distribution $Y_t|X_t$; the objective is simply to compete against $\mathcal{F}$. The algorithm is as follows:

Algorithm 1 Construction of the Cover algorithm

Input: horizon T, number of epochs $K \leq T$
Let $T_k = \lfloor k\frac{T}{K} \rfloor$ for $k \in \{0, \dots, K\}$
for $k \in [K]$ **do**
  Construct a minimal-size cover $S_k \subset \mathcal{F}$ such that for any $f \in \mathcal{F}$ there exists $g \in S_k$ with $f(x_s) = g(x_s)$ for $s \in [T_{k-1}]$
  For iterations $t \in (T_{k-1}, T_k]$, run any learning-with-expert-advice algorithm (e.g., Hedge) with expert set S_k
end for

Theorem 10. *Let $F \subset \left\{\{0,1\}^{\mathcal{X}}\right\}$ be a class of VC dimension d. Suppose that $(x_t)_{t\geq 1}$ is σ-smooth against base measure μ. Then Algorithm 1 run with $K = \mathcal{O}\left(\log T \cdot (T/d)^{1/3}\sigma^{-2/3}\right)$ makes prediction $\hat{y}_t$ satisfying*

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell_t(\hat{y}_t) - \inf_{f\in\mathcal{F}} \ell_t(f(x_t))\right] \leq C \log^2 T \left(\frac{dT^2}{\sigma}\right)^{1/3}.$$

We remark that in a clever recursive construction R-Cover [1] is able to achieve the optimal rate, though the discussion is slightly beyond the scope of this survey. It remains an open and very interesting problem (in the author's opinion!) if the form of R-Cover may be improved to be iterative rather than recursive and/or ERM oracle efficient.

Theorem 11. *Fix $T \geq 1$ and let $F \subset \left\{\{0,1\}^{\mathcal{X}}\right\}$ be as before. Suppose that $(x_t)_{t\geq 1}$ is σ-smooth with respect to μ. Then R-Cover achieves predictions $\hat{y}_t$ satisfying*

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell_t(\hat{y}_t) - \inf_{f\in\mathcal{F}} \ell_t(f(x_t))\right] \leq C \log^{5/2}(T)\sqrt{\frac{dT}{\sigma}},$$

which is the information-theoretically optimal *dependence on parameters.*

2.3.1 Oblivious Analysis of Theorem 10

We refine the hedging approach of [4] using the fundamental "bounded surprise" insight from [3]. Recall the proof follows in two parts: the analysis of Hedge over the cover, and the discretization error inherited by using the approximation. If we have nearly evenly sized epochs, so $\Delta T := \max_k (T_k - T_{k-1}) = \mathcal{O}(T/K)$, this first part may be bounded as

$$C \sum_{k\in[K]} \sqrt{(T_k - T_{k-1}) \log\left(T^d\right)} \leq \mathcal{O}\left(\sqrt{KdT\log T}\right),$$

where we have used Sauer-Shelah's Lemma to bound the number of experts over T iterations as $\mathcal{O}(T^d)$ and then Cauchy-Schwarz inequality to collate into a single expression.

We now address the second step. Consider the distortion of a cover created at t_0 at a time step t in the future through its ℓ_1 disagreement

$$\gamma_{t_0}(t) = \sup_{\substack{f,g\in\mathcal{F} \\ f(x_s)=g(x_s), s\in[t_0]}} \Pr\left[f(x_t) \neq g(x_t) | \mathcal{H}_{t-1}\right] = \sup_{\substack{f,g\in\mathcal{F} \\ f(x_s)=g(x_s), s\in[t_0]}} \Pr_{x\sim\mu_t}\left[f(x_t) \neq g(x_t)\right]$$

We argue that controlling this quantity is sufficient for regret bounds. Indeed, suppose that the optimal function in hindsight was $f^* \in \mathcal{F}$ and let $f_k \in S_k$ be the corresponding representative from the cover during epoch k. We have the expectation bound on discretization bounded as

$$\mathbb{E}\left[\sum_{k\in[K]}\sum_{t=T_{k-1}+1}^{T_k} \ell_t(f_k(x_t)) - \ell_t(f^*(x_t))\right] \leq \mathbb{E}\left[\sum_{k\in[K]}\sum_{t=T_{k-1}+1}^{T_k} \gamma_{T_{k-1}(t)}\right].$$

Using the surprise bound method, we can show epochs for which significant distortions occur are bounded by $\tilde{O}(1/(q\sigma))$.

Lemma 12. *Let $0 = T_0 < T_1 < \cdots < T_k = T$ define epochs and fix any parameter $q, \delta \in (0,1]$ and denote $w(T,\delta) = d\log\left(\frac{T}{\sigma}\log\frac{1}{\delta}\right) + \log\frac{T}{\delta} + 2$. Then, with probability at least $1-\delta$,*

$$\left|\left\{k \in [K] : \sum_{t=T_{k-1}+1}^{T_k} \gamma_{T_{k-1}}(t) \cdot 1\left[\gamma_{T_{k-1}}(t) \geq q\right] \geq w(T,\delta)\right\}\right| \leq C\frac{\log^2 T}{q\sigma}.$$

With the control endowed by the previous lemma, we may finish by direct calculation

$$\begin{aligned}
\sum_{k\in[K]}\sum_{t=T_{k-1}+1}^{T_k} \gamma_{T_{k-1}(t)} &\overset{(i)}{\leq} q_0 T + \sum_{l=0}^{l_0}\sum_{k\in[K]} \overbrace{\sum_{t=T_{k-1}+1}^{T_k} \gamma_{T_{k-1}}(t) 1\left(\gamma_{T_{k-1}}(t) \in \left[2^l q_0, 2^{l+1} q_0\right]\right)}^{\Gamma_{k,l}} \\
&\overset{(ii)}{\leq} q_0 T + (l_0+1)K \cdot w(T,\delta) + \sum_{l=1}^{l_0} 2^{l+1} q_0 \Delta T \left|\left\{k \in [K] : \Gamma_{k,l} \geq w(T,\delta)\right\}\right| \\
&\overset{(iii)}{\leq} q_0 T + (l_0+1)K \cdot w(T,\delta) + \sum_{l=0}^{l_0} 2^{l+1} q_0 \Delta T \cdot C\frac{\log^2 T}{2^l q_0 \sigma} \leq \frac{\log^3 T}{K\sigma} \cdot T
\end{aligned}$$

In (i) we partitioned the range $[0,1]$ of γ into dyadic intervals, in (ii) we separated $\Gamma_{k,l}$ into surprising and unsurprising events and applied pessimistic bounds, and in (iii) we finally applied Lemma 12 and took the summation. Putting the regret terms together, we get

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell_t(\hat{y}_t) - \inf_{f\in\mathcal{F}} \ell_t(f(x_t))\right] \leq \mathcal{O}\left(\sqrt{KdT\log T} + \frac{\log^3 T}{K\sigma} \cdot T\right),$$

and balancing K as before gives the desired $\tilde{\mathcal{O}}\left(T^{2/3}\right)$ rate.

3 ChatGPT Statement

ChatGPT or any other AI assistance was not used.

References

[1] Moïse Blanchard. Agnostic smoothed online learning. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, page 1997–2006, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400715105. doi: 10.1145/3717823.3718111. URL https://doi.org/10.1145/3717823.3718111.

[2] Moïse Blanchard and Samory Kpotufe. Distributionally-constrained adversaries in online learning, 2025. URL https://arxiv.org/abs/2506.10293.

[3] Adam Block, Alexander Rakhlin, and Abhishek Shetty. On the performance of empirical risk minimization with smoothed data, 2024. URL https://arxiv.org/abs/2402.14987.

[4] Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. *J. ACM*, 71(3), June 2024. ISSN 0004-5411. doi: 10.1145/3656638. URL https://doi.org/10.1145/3656638.

[5] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Stochastic, constrained, and smoothed adversaries. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL https://proceedings.neurips.cc/paper_files/paper/2011/file/692f93be8c7a41525c0baf2076aecfb4-Paper.pdf.