## A Additional Related Work

Beyond Worst-Case Online Learning [RST11] consider online learning where the adversary is constrained and build a framework based on minimax analysis and constrained sequential Rademacher complexity to analyze regret in these scenarios. These techniques have been applied to other constrained settings [ $\left.\mathrm{KAH}^{+} 19\right]$.

Smoothed Online Learning [GR17] consider smoothed online learning when looking at problems in online algorithm design. They prove that while optimizing parameterized greedy heuristics for combinatorial problems in presence of smoothing this problem can be learned with non-trivial sublinear regret. [CAK17] consider the same problem with an emphasis on the per-step runtime being logarithmic in $T$. [HRS20, HRS22] both study the notion of smoothed analysis with adaptive adversary and show that statistically the regret is bounded by $O(\sqrt{T d \log (1 / \sigma)})$.

Online Learning with Predictable Sequences Another line of work has focused on the future sequences being predictable given the past instances. [RS13b] incorporate additional information available in terms of an estimator for future instances. They achieve regret bounds depending on the path length of these estimators and can beat the worst-case $\Omega(\sqrt{T})$ if the estimators are accurate. [HM07] models predictability as knowing the first coordinate of loss vectors, which is revealed to the learner before he chooses actions. Some other work model predictability through hints which are additive estimate of loss vectors [HK10, RS13a, SL14, MY16]. [DHJ ${ }^{+}$17, BCKP20] considers settings where the learner has access to hints in form of vectors that are weakly correlated with the future instances and show exponential improvement in the regret in some cases. The literature on hints represents an active and growing subarea of online learning (see [BCKP20] and references within).

## B Oracle-Efficient Learning with Real-valued Functions

## B. 1 Coupling Lemma

Lemma B. 1 (Coupling, [HRS22]). Let $\mathscr{D}_{\sigma}$ be an adaptive sequence of $t$-smooth distributions on $\mathcal{X}$. Then, there is a coupling $\Pi$ such that $\left(x_{1}, z_{1,1}, \ldots, z_{1, K}, \ldots, x_{t}, z_{t, 1}, \ldots, z_{t, K}\right) \sim \Pi$ satisfy
a. $x_{1}, \ldots, x_{t}$ is distributed according $\mathscr{D}_{\sigma}$.
b. For every $j \leq t,\left\{z_{i, k}\right\}_{i \geq j, k \in[K]}$ are uniformly and independently distributed on $\mathcal{X}$, conditioned on $x_{1}, \ldots, x_{j-1}$.
c. With probability at least $1-t(1-\sigma)^{K},\left\{x_{1}, \cdots, x_{T}\right\} \subseteq\left\{z_{t, k}\right\}_{t=1: T, k=1: K}$.

## B. 2 Monotonicity of the Regularized Rademacher Complexity

Lemma 4.1 (Restated). Let $Z=\left\{z_{i}\right\}_{i \in[m]} \in \mathcal{X}^{m}$ be a set of unlabeled instances and $\Phi: \mathcal{H} \rightarrow \mathbb{R}$ be a mapping from the set of hypothesis to real values. Recall that the Rademacher complexity for set $Z$ regularized by $\Phi$ is defined as

$$
\mathfrak{R}(\Phi, Z)=\underset{\epsilon_{1: m} \stackrel{i i d}{\sim} \mathcal{U}( \pm 1)}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{i=1}^{m} \epsilon_{i} h\left(z_{i}\right)+\Phi(h)\right\}\right]
$$

Then for any dataset $z_{1: m} \in \mathcal{X}^{m}$ and any additional data point $x \in \mathcal{X}$, we have

$$
\mathfrak{R}\left(\Phi, z_{1: m}\right) \leq \mathfrak{R}\left(\Phi, z_{1: m} \cup\{x\}\right)
$$

Proof of Lemma 4.1. Using $\mathbb{E}\left[\sup _{\lambda} X_{\lambda}\right] \geq \sup _{\lambda} \mathbb{E}\left[X_{\lambda}\right]$, we have

$$
\begin{aligned}
\mathfrak{R}(\Phi, Z \cup\{x\}) & =\underset{\epsilon_{1: m+1}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{i=1}^{m} \epsilon_{i} h\left(z_{i}\right)+\epsilon_{m+1} h(x)+\Phi(h)\right\}\right] \\
& =\underset{\epsilon_{1: m}, \epsilon_{m+1}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{i=1}^{m} \epsilon_{i} h\left(z_{i}\right)+\epsilon_{m+1} h(x)+\Phi(h)\right\}\right] \\
& \geq \underset{\epsilon_{1: m}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{i=1}^{m} \epsilon_{i} h\left(z_{i}\right)+\underset{\epsilon_{m+1}}{\mathbb{E}}\left[\epsilon_{m+1} h(x)\right]+\Phi(h)\right\}\right] \\
& =\mathfrak{R}(\Phi, Z),
\end{aligned}
$$

as desired.

## B. 3 Notions for Real-Valued Functions

In this section we introduce the notions that will be useful in analyzing real-valued hypothesis classes, including pseudo dimension and covering numbers.
Definition B. 1 (Pseudo-dimension, [AB99]). For every $h \in \mathcal{H}$, let $B_{h}(x, y)=\operatorname{sgn}(h(x)-y)$ be the indicator of the region below or on the graph of $h$. The pseudo-dimension of hypothesis class $\mathcal{H}$ is defined as the VC dimension of the subgraph class $B_{\mathcal{H}}=\left\{B_{h}: h \in \mathcal{H}\right\}$.

We will see in the two following lemmas that pseudo dimension can be used to characterize the magnitude of covering numbers and Rademacher complexity.
Lemma B. $2\left(d_{L_{1}(\mathcal{U})}\right.$-Covering Number Bound, [AB99]). The $\epsilon$-covering number of $\mathcal{H}$ with respect to metric $d_{L_{1}(\mathcal{U})}$, denoted by $\mathcal{N}\left(\epsilon, \mathcal{H}, L_{1}(\mathcal{U}(\mathcal{X}))\right)$, is the cardinality of the smallest subset $\mathcal{H}^{\prime}$ of $\mathcal{H}$, such that for every $h \in \mathcal{H}$, there exists $h^{\prime} \in \mathcal{H}^{\prime}$ such that $d_{L_{1}(\mathcal{U})}\left(h, h^{\prime}\right) \leq \epsilon$, where $d_{L_{1}(\mathcal{U})}(f, g)=$ $\mathbb{E}_{\mathcal{U}}[|f-g|]$. If $d$ is the pseudo-dimension of $\mathcal{H}$, then for any $\epsilon>0$,

$$
\log \mathcal{N}\left(\epsilon, \mathcal{H}, L_{1}(\mathcal{U}(\mathcal{X}))\right) \in \widetilde{O}\left(d \log \left(\frac{1}{\epsilon}\right)\right)
$$

Lemma B. 3 (Rademacher Complexity Bound, [Bar06]). The Rademacher complexity of class $\mathcal{H}$ for a set of $n$ elements is upper bounded by $O(\sqrt{d n \log n})$, where $d$ is the pseudo dimension of $\mathcal{H}$.

## B. 4 Proof of Theorem 3.1

Theorem 3.1 (Restated). For any $\sigma$-smooth adversary $\mathscr{D}_{\sigma}$, Algorithm 1 has expected regret upper bounded by $\widetilde{O}(G \sqrt{T d / \sigma})$, where $\widetilde{O}$ hide factors that are polynomial in $\log (T)$ and $\log (1 / \sigma)$. Here $G$ is the Lipschitz constant of the loss and $d$ is the pseudodimension of class $\mathcal{H}$. Furthermore, the algorithm is oracle-efficient: at every round $t$, this algorithm uses two oracle calls with histories of length $\widetilde{O}(T / \sigma)$.

Proof. To prove Theorem 3.1 we use the following relaxation:

$$
\begin{align*}
\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right) & =2 G \underset{V^{(t)} \underset{\substack{i d}}{\mathbb{U}} \mathcal{U}(\mathcal{X})}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), V^{(t)}\right)\right]+2 G \beta(T-t) \\
& =2 G \underset{V^{(t)}, \mathcal{E}^{(t)}}{\mathbb{E}}\left[\sup _{\substack{ \\
h \in \mathcal{H}}}\left\{\sum_{\substack{i=t+1: T \\
k=1: K}} \epsilon_{i, k}^{(t)} h\left(v_{i, k}^{(t)}\right)-L^{\mathrm{r}}\left(h, s_{1: t}\right)\right\}\right]+2 G \beta(T-t), \tag{2}
\end{align*}
$$

where $K=100 \log T / \sigma$ and $\beta=10 T K(1-\sigma)^{K}$. We will show in Lemma B. 4 that the above relaxation is admissible. Therefore, Proposition 4.1 gives us the following upper bound on the expected regret:
$\mathbb{E}[\operatorname{REGRET}(T)] \leq \operatorname{Rel}_{T}(\mathcal{H} \mid \emptyset)+O(\sqrt{T})=2 G \underbrace{\underset{V^{(0)}, \mathcal{E}^{(0)}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{\substack{i=1: T \\ k=1: K}} \epsilon_{i, k}^{(0)} h\left(v_{i, k}^{(0)}\right)\right\}\right]}_{\text {(a) }}+2 G \beta T+O(\sqrt{T})$.
The first term (a) is the Rademacher complexity of the hypothesis class $\mathcal{H}$ with respect to the uniform distribution for sample size $T K$. By Lemma B.3, (a) $\leq O(\sqrt{d T K \log (T K)})$. For the second term, we have $\beta T \in o(1)$ because $\beta \lesssim T K e^{-\sigma K} \lesssim T^{-99} \log T / \sigma=o(1 / T)$. Plugging in $K=O(\log (T) / \sigma)$, we have the following bound:

$$
\mathbb{E}[\operatorname{REGRET}(T)] \leq O\left(G \sqrt{\frac{d T}{\sigma} \log T \log \left(\frac{T}{\sigma}\right)}\right) \subseteq \widetilde{O}(G \sqrt{d T / \sigma})
$$

where $\widetilde{O}$ hide factors that are polynomial in $\log (T)$ and $\log (1 / \sigma)$.

## B. 5 Admissibility of the Relaxation

Lemma B.4. The prediction rule $\mathscr{Q}=\left(\mathcal{Q}_{1}, \cdots, \mathcal{Q}_{T}\right)$ given by Algorithm 1 is admissible with respect to the relaxation defined in Equation (2).

Proof. Using the language of regularized Rademacher complexity, the above relaxation can be written as

$$
\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right)=2 G \underset{\left.V^{(t)}\right) \underset{\text { id }}{\sim} \mathcal{U}(\mathcal{X})}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), V^{(t)}\right)\right]+2 G \beta(T-t)
$$

where $L^{\mathrm{r}}\left(\cdot, s_{1: t}\right)=\sum_{i=1}^{t-1} l^{\mathrm{r}}\left(h\left(x_{i}\right), y_{i}\right)$. When $t=T$, the relaxation becomes

$$
\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: T}\right)=-2 G L^{\mathrm{r}}\left(h, s_{1: T}\right)=-\inf _{h \in \mathcal{H}} \sum_{i=1}^{T} l\left(h\left(x_{i}\right), y_{i}\right)
$$

thus it satisfies the second condition of Definition 4.1. For the first condition, we need to verify

$$
\begin{equation*}
\sup _{\mathcal{D}_{t} \in \mathfrak{D}_{t}} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \sup _{y_{t} \in \mathcal{Y}}\left\{\underset{\widehat{y_{t}} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1} \cup\left(x_{t}, y_{t}\right)\right)\right\} \leq \operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right) \tag{3}
\end{equation*}
$$

We first upper bound the LHS of Equation (3) by matching the randomness in $V^{(t)}$ and applying Jensen's inequality to the supremum function. For every fixed input $x_{t}$ and hint set $V^{(t)}$, we denote our prediction rule in Equation (1) with $\mathcal{Q}_{t}\left(V^{(t)}\right)$. This gives us

$$
\begin{align*}
& \sup _{\mathcal{D}_{t} \in \mathfrak{D}_{t}} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \sup _{y_{t} \in \mathcal{Y}}\left\{\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+\mathbf{R e l}_{T}\left(\mathcal{H} \mid s_{1: t-1} \cup\left(x_{t}, y_{t}\right)\right)\right\} \\
& =\sup _{\mathcal{D}_{t} \in \mathfrak{D}_{t}} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \sup _{y_{t} \in \mathcal{Y}} \underset{V^{(t)}}{\mathbb{i i d} \mathcal{\sim} \mathcal{U}(\mathcal{X})}\left[\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}\left(V^{(t)}\right)}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+2 G \cdot \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), V^{(t)}\right)\right]+2 G \beta(T(4) t) \\
& \leq \sup _{\mathcal{D}_{t} \in \mathfrak{D}_{t}} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{V^{(t)}}{\mathbb{E} \sim} \mathcal{U}^{\mathbb{E}} \mathcal{U}(\mathcal{X})\left[\sup _{y_{t} \in \mathcal{Y}}\left\{\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}\left(V^{(t)}\right)}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+2 G \cdot \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), V^{(t)}\right)\right\}\right]+2 G \beta(T-t) \tag{5}
\end{align*}
$$

In Equation (5), note that $\mathcal{Q}_{t}\left(V^{(t)}\right)$ is the same as the transductive prediction rule in [RSS12, Equation (25)], with $V^{(t)}$ being the set of unlabeled future instances and $s_{1: t-1}$ being the historical data with labels.

According to [RSS12, Lemma 12], for all input $x_{t}$ and unlabeled sequence $\mathscr{X}$ (which plays the role of $\left.x_{t+1: T}\right)$, the decision rule $\mathcal{Q}_{t}(\mathscr{X})$ satisfies

$$
\begin{gathered}
\sup _{y_{t} \in \mathcal{Y}}\left\{\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}(\mathscr{X})}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+2 G \underset{\mathcal{E}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{x \in \mathscr{X}} \epsilon_{x} h(x)-L^{\mathrm{r}}\left(h, s_{1: t-1} \cup\left(x_{t}, y_{t}\right)\right)\right\}\right]\right\} \\
\leq 2 G \underset{\mathcal{E}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{x \in \mathscr{X} \cup\left\{x_{t}\right\}} \epsilon_{x} h(x)-L^{\mathrm{r}}\left(h, s_{1: t-1}\right)\right\}\right]
\end{gathered}
$$

Therefore, if we choose the sequence $\mathscr{X}$ to be $V^{(t)}$, we obtain the following inequality which is written in the language of regularized Rademacher complexity:
$\sup _{y_{t} \in \mathcal{Y}}\left\{\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}\left(V^{(t)}\right)}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+2 G \cdot \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), V^{(t)}\right)\right\} \leq 2 G \cdot \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right)$.
By adding the expectations over $V^{(t)}$ and $x_{t}$ on both sides, we obtain the following upper bound:

$$
\begin{aligned}
(5) & \leq \sup _{\mathcal{D}_{t} \in \mathfrak{D}_{t}} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{V^{(t)} \stackrel{i d}{\sim} \mathcal{U}(\mathcal{X})}{\mathbb{E}}\left[2 G \cdot \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right)\right]+2 G \beta(T-t) \\
& \leq \underset{\left.V^{(t)}\right)}{\mathbb{E} \underset{\sim}{\sim} \mathcal{U}(\mathcal{X})}\left[\sup _{\mathcal{D}_{t} \in \mathfrak{D}_{t}} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} 2 G \cdot \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right)\right]+2 G \beta(T-t) .
\end{aligned}
$$

According to Lemma B.5, we can replace the $x_{t}$ sampled from the worst-case smooth distribution by $Z_{t}$ sampled independently from the uniform distribution, with the extra cost $\beta$. This gives

$$
\begin{align*}
(5) & \leq \underset{V^{(t)} \stackrel{\text { id }}{\sim} \mathcal{U}(\mathcal{X})}{\mathbb{E}} \underset{z_{t} \sim}{\underset{\text { in }}{\sim} \mathcal{U}(\mathcal{X})} \\
& =2 G \underset{V^{(t)} \underset{\mathbb{E}}{\mathbb{E}} \mathcal{U}(\mathcal{X})}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t-1)}\right)\right]+2 G \beta(T-t+1)  \tag{6}\\
& =\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right),
\end{align*}
$$

which is precisely the RHS of Equation (3).
Lemma B. 5 (Replacing Supremum by Expectation). For any $V^{(t)} \in \mathcal{X}^{K(T-t)}$, there exists a set of $K$ variables $Z_{t}=\left\{z_{t, k}\right\}_{k \in[K]}$, such that
$\sup _{\mathcal{D}_{t}^{\mathcal{X}} \in \Delta_{\sigma}(\mathcal{X})} \underset{x_{t} \sim \mathcal{D}_{t}^{\mathcal{X}}}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right)\right] \leq \underset{\substack{\text { id } \\ Z_{t} \sim \mathcal{U}(\mathcal{X})}}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right)\right]+\beta$.
Proof. To establish the monotonicity property, we need to show that the random instance $x_{t}$ drawn from a smooth distribution belongs to a set of uniform i.i.d. hints with high probability. This is where the coupling lemma comes in. For the smooth distribution $\mathcal{D}_{t} \in \Delta_{\sigma}(\mathcal{X})$ that achieves the supremum (assume the supremum is achievable), Lemma B. 1 shows the existence of a coupling $\Pi$ on $\left(x_{t}, z_{t, 1}, \cdots, z_{t, K}\right)$ such that $x_{t}$ is distributed according to $\mathcal{D}_{t}^{\mathcal{X}}$ and $Z_{t}=\left\{z_{t, k}\right\}_{k \in[K]}$ are uniformly and independently distributed. We thus have

$$
\begin{equation*}
\sup _{\mathcal{D}_{t} \in \Delta_{\sigma}(\mathcal{X})} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right)\right]=\underset{\Pi}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right)\right] . \tag{7}
\end{equation*}
$$

This joint distribution $\Pi$ has the property that event the $E_{t} \stackrel{\text { def }}{=}\left\{x_{t} \in Z_{t}\right\}$ happens with high probability. We now upper bound the expected value by conditioning on $E_{t}$ and $\bar{E}_{t}$ respectively.
Conditioned on $E_{t}$, we apply the monotonicity of regularized Rademacher complexity (Lemma 4.1) recursively and obtain

$$
\begin{equation*}
\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right) \leq \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right) \tag{8}
\end{equation*}
$$

Conditioned on $\bar{E}_{t}$, we skirt the monotonicity issue by directly using upper and lower bounds on the regularized Rademacher complexity. To be more precise, we use Lemma B. 6 in Appendix B. 6 to show that

$$
\begin{align*}
\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right) & \leq T K \leq T K+\left(\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right)+T\right) \\
& \leq \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right)+2 T K . \tag{9}
\end{align*}
$$

Finally, we expand the right hand side of Equation (7) by conditioning on $E_{t}$ and $\bar{E}_{t}$ respectively. Putting Equations (8) and (9) together, we obtain

$$
\begin{aligned}
& \text { (7) }=\operatorname{Pr}\left[E_{t}\right] \cdot \underset{\Pi}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right) \mid E_{t}\right] \\
& +\operatorname{Pr}\left[\bar{E}_{t}\right] \cdot \underset{\Pi}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup\left\{x_{t}\right\}\right) \mid \bar{E}_{t}\right] \\
& \leq \operatorname{Pr}\left[E_{t}\right] \cdot \underset{\Pi}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right) \mid E_{t}\right] \\
& +\operatorname{Pr}\left[\bar{E}_{t}\right] \cdot \underset{\Pi}{\mathbb{E}}\left[\Re\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right)+2 T K \mid \bar{E}_{t}\right] \\
& =\underset{\Pi}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right)\right]+\operatorname{Pr}\left[\bar{E}_{t}\right] \cdot 2 T K .
\end{aligned}
$$

Since $\Pi$ has uniform marginal distribution on $Z_{t}$, and that $\operatorname{Pr}\left[\bar{E}_{t}\right] \cdot 2 T K \leq(1-\sigma)^{K} \cdot 2 T K \leq \beta$, we further obtain

$$
(7) \leq \underset{Z_{t} \text { iid } \mathcal{U}(\mathcal{X})}{\mathbb{E}}\left[\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t-1}\right), V^{(t)} \cup Z_{t}\right)\right]+\beta
$$

thus completes the proof.

## B. 6 Upper and Lower Bounds on the Relaxation

Lemma B. 6 (Upper and Lower Bounds on the Relaxation). For all $t \in[T]$, all sequence $s_{1: T}$, and all instance set $Z$ of size no larger than $(T-t) K$,

$$
-\frac{T}{2} \leq \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), Z\right) \leq T K
$$

Proof. Let $I$ be the number of instances in set $Z$ and let $\mathcal{E}=\left(\epsilon_{1}, \cdots, \epsilon_{I}\right)$ be the random labels associated with them. By convexity of the supremum,

$$
\begin{aligned}
& \mathfrak{R}\left(-L_{t}, Z\right)=\underset{\mathcal{E} \sim}{\mathbb{E} \sim}\left[\operatorname{U}^{\mathrm{id}(\mathcal{Y})}\left[\sup _{h \in \mathcal{H}}\left\{\sum_{i=1}^{I} \epsilon_{i} h\left(z_{i}\right)-L^{\mathrm{r}}\left(h, s_{1: t}\right)\right\}\right] \geq \sup _{h \in \mathcal{H}}\left\{\underset{\mathcal{E}^{\mathrm{iid} \mathcal{U}(\mathcal{Y})}}{\mathbb{E}}\left[\sum_{i=1}^{I} \epsilon_{i} h\left(z_{i}\right)-L^{\mathrm{r}}\left(h, s_{1: t}\right)\right]\right\}\right. \\
& =\sup _{h \in \mathcal{H}} \sum_{i=1}^{t} \underbrace{\left(-l^{\mathrm{r}}\left(h\left(x_{t}\right), y_{t}\right)\right)}_{\geq-1} \geq-T \text {. }
\end{aligned}
$$

For the upper bound, we notice that $\forall \mathcal{E}, h$,

$$
\sum_{i=1}^{I} \epsilon_{i} h\left(z_{i}\right)-L^{\mathrm{r}}\left(h, s_{1: t}\right) \leq I+t \leq(T-t) K+t \leq T K
$$

So the $\mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), Z\right)$ also has an upper bound of $T K$.

## B. 7 Remark on the Requirement of Fresh Dataset

In order to beat the adaptive adversary, the learner needs to sample fresh random hints in each round. Otherwise, the adversary can enforce high regret by correlating future labels with the history. More precisely, we will see that the matching randomness argument in Inequality (4) uses the crucial fact that $V^{(t)}$ is a fresh dataset that is uniformly distributed independent of the interactions in the past. If $V^{(t)}$ is reused, then the adaptive adversary has the power to correlate $s_{1: t-1}$ with $V^{(t)}$ such that $V^{(t)}$ is no longer unbiased conditioned on the history. In this case, the algorithm fails to mimic the randomization in the relaxation, and the matching-randomness argument breaks down.
Another important property of the fresh self-generated hints $v_{i, k}^{(t)}$ is that they are identically distributed with the real hints $z_{i, k}$ in the coupling. Nevertheless, the analysis has to unite the fact that the learner can only access $v_{i, k}^{(t)} \mathrm{s}$, and the monotonicity property (lemma B.5) is based on $z_{i, k}$. This point is subtle because it is impossible for the self-generated hints to really tell the future (i.e., ensure $x_{t} \in\left\{v_{t, k}^{(t-1)}\right\}_{k \in[K]}$ ), since they are not controlled by the coupling $\Pi$. This issue is taken care of by Equation (6). We can see that it is sufficient for the uncoupled hints in $V^{(t-1)}$ to resemble the coupled hints $Z_{t}$ at distribution level. This distributional resemblence is not achievable if $V^{(t-1)}$ were not independent with the past.

## C Oracle-Efficient Online Binary Classification

## C. 1 Information Theoretic Lemmas

For two probability distributions $P$ and $Q$ over the same domain $\mathcal{X}$, let $\chi^{2}(P, Q)=$ $\sum_{x \in \mathcal{X}} P(x)^{2} / Q(x)-1$ be the $\chi^{2}$-divergence. The following lemma upper bounds the TV distance by the $\chi^{2}$-divergence; a proof could be found in [Tsy09, Chapter 2].
Lemma C. 1 (From TV to $\chi^{2}$ ). The following relations hold:

$$
\mathrm{TV}(P, Q) \leq \sqrt{\frac{1}{2} \log \left(1+\chi^{2}(Q, P)\right)} \leq \sqrt{\frac{\chi^{2}(Q, P)}{2}}
$$

The following statement is the well-known Ingster's $\chi^{2}$ method, and we refer to the excellent book [IS03] for a general treatment.
Lemma C. 2 (Ingster's $\chi^{2}$ method). For a mixture distribution $\mathbb{E}_{\theta \sim \pi}\left[Q_{\theta}\right]$ and a generic distribution $P$, the following identity holds:

$$
\chi^{2}\left(\underset{\theta \sim \pi}{\mathbb{E}}\left[Q_{\theta}\right], P\right)=\underset{\theta, \theta^{\prime} \sim \pi}{\mathbb{E}}\left[\underset{x \sim P}{\mathbb{E}}\left(\frac{Q_{\theta}(x) Q_{\theta^{\prime}}(x)}{P(x)^{2}}\right)\right]-1
$$

where $\theta^{\prime}$ is an independent copy of $\theta$.

## C. 2 Proof of Theorem 3.2

Theorem 3.2 (Restated). In the setting of binary classification with $\sigma$-smoothed adversaries, Algorithm 2 has regret that is at most $\widetilde{O}\left(\min \left\{\sqrt{d T \sigma^{-1 / 2}}, \sqrt{T(d|\mathcal{X}|)^{1 / 2}}\right\}\right)$. Furthermore, Algorithm 2 is a proper learning oracle-efficient algorithm: at every round $t$, this algorithm uses a single ERM oracle call on a history that is of length $t+O(T / \sqrt{\sigma})$ with high probability.

In the remainder of this section, we present a proof of the regret upper bound $\widetilde{O}\left(\sqrt{d T \sigma^{-1 / 2}}\right)$ in Theorem 3.2 when $\sigma \geq d /|\mathcal{X}|$. The proof of the other case $\sigma<d /|\mathcal{X}|$ is slightly different and will be presented in Appendix C.6.
To prove the upper bound, we will use the relaxation framework for analyzing the stability of FTPL. Our proof establishes the connection between the relaxation and FTPL frameworks, which we believe will be useful for the design of online algorithms more generally.

Writing $s=(x, y)$ and $L(h, s)=l(h(x), y)=-y h(x) / 2$, the relaxation is defined as

$$
\begin{equation*}
\operatorname{Re}_{T}\left(\mathcal{H} \mid s_{1: t}\right)=\underset{R^{(t+1)}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left(-\sum_{i=1}^{N^{(t+1)}} L\left(h, \widetilde{s}_{i}^{(t+1)}\right)-\sum_{\tau=1}^{t} L\left(h, s_{\tau}\right)\right)\right]+\eta(T-t) \tag{10}
\end{equation*}
$$

where

$$
\eta=\frac{1}{\sqrt{n \sigma}}+c \sqrt{\frac{d \log T}{n \sigma}}+\frac{n \sigma}{4 T^{2} \log T}+e^{-n / 8} \in \widetilde{O}\left(\sqrt{\frac{d}{n \sigma}}\right)
$$

with an absolute constant $c>0$ given in Lemma C. 5 later, and $R^{(t)}=\left(N^{(t)},\left\{\widetilde{s}_{i}\right\}_{i \in N^{(t)}}\right)$ is the fresh randomness generated at the beginning of time $t$, which is independent of $\left\{s_{\tau}\right\}_{\tau<t}$ generated by the adversary. The relaxation here is similar to Equation (2), where the key difference is a different generation process for the hint set and an additional term $\eta(T-t)$ to account for the stability.
Let $\mathcal{Q}_{t}$ be the distribution of the learner's action $h_{t} \in \mathcal{H}$ in Algorithm 2, then the relaxation in Equation (10) is admissible with respect to Algorithm 2 if the following two conditions hold:

$$
\begin{align*}
\sup _{\mathcal{D}_{t} \in \Delta_{\sigma}(\mathcal{X})} \underset{x_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \sup _{y_{t}}\left[\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]+\right. & \left.\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right)\right]
\end{align*} \leq \operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right), \quad \forall s_{1: t-1}(11),
$$

According to proposition 4.1, if both Inequalities (11) and (12) hold, the expected regret of Algorithm 2 will satisfy

$$
\begin{aligned}
\mathbb{E}[\operatorname{REGRET}(T)] & \leq \operatorname{Rel}_{T}(\mathcal{H} \mid \emptyset)+O(\sqrt{T})=\underset{R^{(1)}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left(-\sum_{i=1}^{N^{(1)}} L\left(h, \widetilde{s}_{i}^{(1)}\right)\right)\right]+\eta T+O(\sqrt{T}) \\
& \stackrel{(\mathrm{a})}{=} O\left(\underset{N^{(1)}}{\mathbb{E}}\left[\sqrt{d N^{(1)}}\right]+\eta T+\sqrt{T}\right) \stackrel{(\mathrm{b})}{=} O(\sqrt{d n}+\eta T+\sqrt{T}),
\end{aligned}
$$

and Theorem 3.2 follows from the choices $n=T / \sqrt{\sigma}$ and $\eta=O(\sqrt{d / n \sigma})$. In the above inequality, step (a) follows from random labels and the upper bound $O(\sqrt{n d})$ on the Rademacher complexity of $\mathcal{H}$ over $n$ points, and step (b) is due to Jensen's inequality and $\mathbb{E}\left[N^{(1)}\right]=n$.
Now it remains to verify Inequalities (11) and (12). It is not hard to verify Inequality (12): this follows from the fact that for any random variable $\lambda, \mathbb{E}\left[\sup _{\lambda} X_{\lambda}\right] \geq \sup _{\lambda} \mathbb{E}\left[X_{\lambda}\right]$ and $\mathbb{E}_{R}\left[L\left(h, \widetilde{s}_{i}\right)\right]=0$. The key technical difficulty is in the proof of Inequality (11). To overcome this challenge, we will show that the stability of learner's distribution $\mathcal{Q}_{t}$ implies the admissibility of the relaxation, where the stability is measured via

$$
\text { Stability }=\underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left(\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]\right)
$$

Note that here $s_{t} \sim \mathcal{D}_{t}$ denotes both the instance and its label and $\mathcal{D}_{t}$ 's marginal over $\mathcal{X}$ is $\sigma$-smooth.
The following lemma formalizes the discussion in Section 4.2 and shows that a small TV distance and modified generalization error suffice to ensure the stability of the algorithm, which in turn implies the admissibility of the relaxation. This result could be of independent interest. The proof can be found in Appendix C.3.
Lemma C. 3 (TV + Generalization $\Rightarrow$ Stability $\Rightarrow$ Admissibility). Let $\mathcal{Q}_{t}$ denote learner's distribution over $\mathcal{H}$ in Algorithm 2 at round $t, \mathcal{D}_{t}$ be adversary's distribution at time $t$ (given the history $\left.s_{1}, \cdots, s_{t-1}\right), s_{t} \sim \mathcal{D}_{t}$ be the realized adversarial instance at time $t$, and $s_{t}^{\prime}$ be an independent copy $s_{t}^{\prime} \sim \mathcal{D}_{t}$. It holds that

$$
\begin{aligned}
& \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left(\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]+\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right)\right)-\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right) \\
& \leq \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left(\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]\right)-\eta \\
& \leq \operatorname{TV}\left(\mathcal{Q}_{t}, \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\mathcal{Q}_{t+1}\right]\right)+\underbrace{\underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t} ; R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)-L\left(h_{t+1}, s_{t}\right)\right]-\eta .}_{\text {Modified Generalization Error }}
\end{aligned}
$$

Lemma C. 3 shows that, in order to prove the admissibility of the relaxation in Equation (10), it remains to upper bound the TV distance and the generalization error, respectively.
Our next lemma provides an upper bound on the TV distance between $\mathcal{Q}_{t}$ and the mixture distribution $\mathbb{E}_{s_{t} \sim \mathcal{D}_{t}}\left[\mathcal{Q}_{t+1}\right]$.
Lemma C. 4 (Upper Bound of TV Distance). Let $\mathcal{Q}^{t}$ be the distribution over $h_{t}$ at time $t$. We have

$$
\sup _{\mathcal{D}_{t} \in \Delta_{\sigma}(\mathcal{S})} \operatorname{TV}\left(\mathcal{Q}_{t}, \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\mathcal{Q}_{t+1}\right]\right) \leq \frac{1}{\sqrt{n \sigma}}
$$

The key ingredient in the proof of Lemma C. 4 is the Poissonization, which ensures the independence of the number of +1 (or -1 ) labels across instances, and enables us to write down the mixture distribution of inputs to the ERM oracle in a compact form. The proof of Lemma C. 4 is presented in Appendix C.4.
The following lemma upper bounds the modified generalization error for any smooth distribution $\mathcal{D}_{t}$. The proof of this lemma is based on the discussions in Section 4.4, and we formally present it in Appendix C.5.
Lemma C. 5 (Upper Bound of Generalization Error). Under the notations of Lemma C.3, it holds for an absolute constant $c>0$ (independent of $(n, d, T, \sigma)$ ) that

$$
\sup _{\mathcal{D}_{t} \in \Delta_{\sigma}(\mathcal{X})}\left\{{\underset{s}{t}, s_{t}^{\prime} \sim \mathcal{D}_{t} ; R^{(t+1)}}_{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)-L\left(h_{t+1}, s_{t}\right)\right]\right\} \leq c \sqrt{\frac{d \log T}{n \sigma}}+\frac{n \sigma}{4 T^{2} \log T}+e^{-n / 8}
$$

Now the claimed result of Theorem 3.2 when $\sigma \geq d /|\mathcal{X}|$ follows from Lemma C.3, Lemma C.4, and Lemma C.5.

## C. 3 Proof of Lemma C. 3

To prove this lemma, we first introduce some notations. For $t \in[T] \cup\{0\}$, let $r^{t} \in \mathbb{Z}^{\mathcal{X}}$ be the $|\mathcal{X}|$-dimensional random vector with $r^{t}(x)$ defined to be the difference between the number of +1 and -1 labels in the self-generated samples and the history up to time $t$ on instance $x$. Formally,

$$
r^{t}(x)=\sum_{i=1}^{N^{(t+1)}} \widetilde{y}_{i}^{(t+1)} \cdot \mathbf{1}\left(\widetilde{x}_{i}^{(t+1)}=x\right)+\sum_{\tau=1}^{t} y_{\tau} \cdot \mathbf{1}\left(x_{\tau}=x\right)
$$

Let $\mathcal{P}^{t}$ be the distribution of $r^{t}$. Also recall that $R^{(t)}=\left(N^{(t)},\left\{\widetilde{s}_{i}\right\}_{i \in N^{(t)}}\right)$ is the fresh randomness generated at the beginning of time $t$.
Using the definitions of $r^{t}, \mathcal{P}^{t}$ and $R^{(t)}$, the following chain of inequalities holds for any fixed $s_{t}$ :

$$
\begin{aligned}
& \underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}} {\left[L\left(h_{t}, s_{t}\right)\right]+\mathbf{R e l}_{T}\left(\mathcal{H} \mid s_{1: t}\right) } \\
& \stackrel{(\mathrm{a})}{=} \underset{\mathcal{P}^{t-1}}{\mathbb{E}} {\left[L\left(\operatorname{opt}\left(r^{t-1}\right), s_{t}\right)\right]-\underset{R^{(t+1)}}{\mathbb{E}}\left[\sum_{i=1}^{N^{(t+1)}} L\left(\operatorname{opt}\left(r^{t}\right), \widetilde{s}_{i}^{(t+1)}\right)+\sum_{\tau=1}^{t} L\left(\operatorname{opt}\left(r^{t}\right), s_{\tau}\right)\right]+\eta(T-t) } \\
&= \underset{\mathcal{P}^{t-1}}{\mathbb{E}}\left[L\left(\operatorname{opt}\left(r^{t-1}\right), s_{t}\right)\right]-\underset{\mathcal{P}^{t}}{\mathbb{E}}\left[L\left(\operatorname{opt}\left(r^{t}\right), s_{t}\right)\right]+\eta(T-t) \\
&-\underset{R^{(t+1)}}{\mathbb{E}}\left[\sum_{i=1}^{N^{(t+1)}} L\left(\operatorname{opt}\left(r^{t}\right), \widetilde{s}_{i}^{(t+1)}\right)+\sum_{\tau=1}^{t-1} L\left(\operatorname{opt}\left(r^{t}\right), s_{\tau}\right)\right] \\
& \leq \underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]+\eta(T-t)+\underset{R^{(t+1)}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left(-\sum_{i=1}^{N^{(t+1)}} L\left(h, \widetilde{s}_{i}^{(t+1)}\right)-\sum_{\tau=1}^{t-1} L\left(h, s_{\tau}\right)\right)\right] \\
&\left(\underset{h^{\prime}}{(\mathrm{b})} \underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]+\eta(T-t)+\underset{R^{(t)}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left(-\sum_{i=1}^{N^{(t)}} L\left(h, \widetilde{s}_{i}^{(t)}\right)-\sum_{\tau=1}^{t-1} L\left(h, s_{\tau}\right)\right)\right]\right. \\
&=\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]-\eta+\mathbf{R e l}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right),
\end{aligned}
$$

where (a) uses the definition of opt $\left(r^{t}\right)$, and (b) is due to the fact that $R^{(t+1)}$ is an independent copy of $R^{(t)}$ conditioned on $\left\{s_{\tau}\right\}_{\tau<t}$. This implies the first inequality of Lemma C.3.
For the second inequality, we further take the expectation with respect to $s_{t} \sim \mathcal{D}_{t}$, and note that $\mathcal{Q}_{t}$ and $\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right)$ are independent of $s_{t}$, while $\mathcal{Q}_{t+1}$ depends on $s_{t}$ :

$$
\begin{aligned}
& \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left(\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]+\mathbf{R e l}_{T}\left(\mathcal{H} \mid s_{1: t}\right)\right)-\mathbf{R e l}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right) \\
& \leq \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]-\eta \\
& \leq \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)\right] \\
& \quad+\underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)\right]-\underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]-\eta \\
& \left(\stackrel{\text { c }}{=} \underset{s_{t}^{\prime} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}^{\prime}\right)\right]-\underset{s_{t}^{\prime} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t+1} \sim \mathbb{E}_{s_{t} \sim \mathcal{D}_{t}}^{\mathbb{E}}\left[\mathcal{Q}_{t+1}\right]}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)\right]\right. \\
& \quad+\underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)\right]-\underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]-\eta \\
& \text { (d)} \\
& \leq \operatorname{TV}\left(\mathcal{Q}_{t}, \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\mathcal{Q}_{t+1}\right]\right)+\underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t} ; R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)-L\left(h_{t+1}, s_{t}\right)\right]-\eta,
\end{aligned}
$$

where (c) follows from the independence of $h_{t} \sim \mathcal{Q}_{t}$ and $\left(s_{t}, s_{t}^{\prime}\right)$, and (d) is due to $\mid \mathbb{E}_{X \sim P}[f(X)]-$ $\mathbb{E}_{X \sim Q}[f(X)] \mid \leq \operatorname{TV}(P, Q)$ for every measurable function $f$ with $\|f\|_{\infty} \leq 1$.

## C. 4 Upper Bounding TV Distance: Proof of Lemma C. 4

Using the notations in lemma C.3, since $h_{t}$ in Algorithm 2 only depends on the vector $r^{t-1}$, the ERM objective could be written as a quantity depending only on $r^{t-1}$ and $h \in \mathcal{H}$. We write $h_{t}=\operatorname{opt}_{\mathcal{H}, l}\left(r^{t-1}\right)$ in the sequel, and then opt $\mathcal{H}, l\left(r^{t-1}\right) \sim \mathcal{Q}_{t}$ as $r^{t-1} \sim \mathcal{P}^{t-1}$. Therefore, the data-processing inequality shows that

$$
\operatorname{TV}\left(\mathcal{Q}_{t}, \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\mathcal{Q}_{t+1}\right]\right) \leq \operatorname{TV}\left(\mathcal{P}^{t-1}, \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\mathcal{P}^{t}\right]\right)
$$

and is suffices to upper bound the TV distance $\operatorname{TV}\left(\mathcal{P}^{t-1}, \mathbb{E}_{s_{t} \sim \mathcal{D}_{t}}\left[\mathcal{P}^{t}\right]\right)$.
Let us first create a better understanding of the structures of the distributions $\mathcal{P}^{t-1}$ and $\mathcal{P}^{t}$. Without loss of generality we assume that $\mathcal{X}$ is discrete (the case of continuous $\mathcal{X}$ can be dealt by analyzing the appropriate Poisson point process). Let $n_{+}(x), n_{-}(x)$ be the numbers of +1 and -1 labels, respectively, given instance $x$ in the self-generated samples:

$$
n_{+}(x)=\sum_{i=1}^{N} \mathbf{1}\left(\widetilde{x}_{i}=x, \widetilde{y}_{i}=+1\right) \text { and } n_{-}(x)=\sum_{i=1}^{N} \mathbf{1}\left(\widetilde{x}_{i}=x, \widetilde{y}_{i}=-1\right) .
$$

As each $\widetilde{x}_{i}$ is uniformly distributed on $\mathcal{X}$ and $\widetilde{y}_{i} \sim \mathcal{U}(\{ \pm 1\})$, by the subsampling property of the Poisson distribution, the $2|\mathcal{X}|$ random variables $\left\{n_{ \pm}(x)\right\}_{x \in \mathcal{X}}$ are i.i.d. distributed as $\operatorname{Poi}(n / 2|\mathcal{X}|)$. This independence implied by the Poisson distribution plays a key role in the analysis. Moreover, $r^{0}(x)=n_{+}(x)-n_{-}(x)$, so $\mathcal{P}^{0}$ is determined by the joint distribution of $\left\{n_{ \pm}(x)\right\}_{x \in \mathcal{X}}$.
As we move to general $t$, note that the only contribution of the historic data $\left\{s_{\tau}\right\}_{\tau<t}$ to both $\mathcal{P}^{t-1}$ and $\mathcal{P}^{t}$ is a common translation independent of $\mathcal{P}^{0}$. Since the TV distance is translation invariant, it suffices to upper bound $\operatorname{TV}\left(\mathcal{P}^{0}, \mathbb{E}_{s_{1}}\left[\mathcal{P}^{1}\right]\right)$. Let $n_{ \pm}^{1}(x)=n_{ \pm}(x)+\mathbf{1}\left(x_{1}=x, y_{1}= \pm 1\right)$, it holds that $r^{1}(x)=n_{+}^{1}(x)-n_{-}^{1}(x)$. Consequently, let $P$ and $Q$ be the probability distributions of $\left\{n_{ \pm}(x)\right\}_{x \in \mathcal{X}}$ and $\left\{n_{ \pm}^{1}(x)\right\}_{x \in \mathcal{X}}$, respectively, the data-processing inequality implies that $\operatorname{TV}\left(\mathcal{P}^{0}, \mathbb{E}_{s_{1}}\left[\mathcal{P}^{1}\right]\right) \leq$ $\operatorname{TV}(P, \bar{Q})$.
As discussed above, the distribution $P$ is a product Poisson distribution:

$$
P\left(\left\{n_{ \pm}(x)\right\}\right)=\prod_{x \in \mathcal{X}} \prod_{y \in\{ \pm\}} \mathbb{P}\left(\operatorname{Poi}(n / 2|\mathcal{X}|)=n_{y}(x)\right) .
$$

As for the distribution $Q$, it could be obtained from $P$ in the following way: the smooth adversary draws $x^{\star} \sim \mathcal{D}$, independent of $\left\{n_{ \pm}(x)\right\}_{x \in \mathcal{X}} \sim P$, for some $\sigma$-smooth distribution $\mathcal{D} \in \Delta_{\sigma}(\mathcal{X})$. He then chooses a label $y^{\star}=y\left(x^{\star}\right) \in\{ \pm 1\}$ as a function of $x^{\star}$, and sets

$$
n_{y\left(x^{\star}\right)}^{1}\left(x^{\star}\right)=n_{y\left(x^{\star}\right)}\left(x^{\star}\right)+1, \quad \text { and } \quad n_{y}^{1}(x)=n_{y}(x), \quad \forall(x, y) \neq\left(x^{\star}, y\left(x^{\star}\right)\right)
$$

Consequently, given a $\sigma$-smooth distribution $\mathcal{D}$ and a labeling function $y: \mathcal{X} \rightarrow\{ \pm\}$ used by the adversary, the distribution $Q$ is a mixture distribution $Q=\mathbb{E}_{x^{\star} \sim \mathcal{D}^{\chi}}\left[Q_{x^{\star}}\right]$, with

$$
Q_{x^{\star}}\left(\left\{n_{ \pm}(x)\right\}\right)=\mathbb{P}\left(\operatorname{Poi}(n / 2|\mathcal{X}|)=n_{y\left(x^{\star}\right)}\left(x^{\star}\right)-1\right) \times \prod_{(x, y) \neq\left(x^{\star}, y\left(x^{\star}\right)\right)} \mathbb{P}\left(\operatorname{Poi}(n / 2|\mathcal{X}|)=n_{y}(x)\right)
$$

To upper bound the TV distance between a mixture distribution $Q$ and a base distribution $P$, we will rely on the smoothness properties of $\mathcal{D}$, in particular, that the probability of collision between two independent draws $x_{1}^{\star}, x_{2}^{\star} \sim \mathcal{D}$ is small. To formally address this, we make use of two technical lemmas, first to upperbound the TV distance in terms of the $\chi^{2}$ distance, and second to use the Ingster's method for bounding the $\chi^{2}$ distance between a mixture distribution and a base distribution. See Lemma C. 1 and Lemma C. 2 in the Appendix C. 1 for more details. Let $x_{1}^{\star}, x_{2}^{\star}$ be an arbitrary pair of instance. Using the closed-form expressions of distributions $P$ and $Q_{x^{\star}}$, it holds that

$$
\frac{Q_{x_{1}^{\star}}\left(\left\{n_{ \pm}(x)\right\}\right) Q_{x_{2}^{\star}}\left(\left\{n_{ \pm}(x)\right\}\right)}{P\left(\left\{n_{ \pm}(x)\right\}\right)^{2}}=\frac{2|\mathcal{X}| n_{y\left(x_{1}^{\star}\right)}\left(x_{1}^{\star}\right)}{n} \cdot \frac{2|\mathcal{X}| n_{y\left(x_{2}^{\star}\right)}\left(x_{2}^{\star}\right)}{n}
$$

Using the fact that $\left\{n_{ \pm}(x)\right\}_{x \in \mathcal{X}}$ are i.i.d. distributed as $\operatorname{Poi}(n / 2|\mathcal{X}|)$ under $P$, we have

$$
\underset{\left\{n_{ \pm}(x)\right\} \sim P}{\mathbb{E}}\left(\frac{Q_{x_{1}^{\star}}\left(\left\{n_{ \pm}(x)\right\}\right) Q_{x_{2}^{\star}}\left(\left\{n_{ \pm}(x)\right\}\right)}{P\left(\left\{n_{ \pm}(x)\right\}\right)^{2}}\right)=1+\frac{2|\mathcal{X}|}{n} \cdot \mathbf{1}\left(x_{1}^{\star}=x_{2}^{\star}\right) .
$$

Now using the aforementioned lemmas (Lemma C. 1 and Lemma C.2), we have

$$
\begin{aligned}
\operatorname{TV}(P, Q) & \leq \sqrt{\frac{\chi^{2}(Q, P)}{2}}=\sqrt{\frac{\chi^{2}\left(\mathbb { E } _ { x ^ { \star } \sim \mathcal { D } } \left[Q_{\left.\left.x^{\star}\right], P\right)}^{2}\right.\right.}{2}}=\sqrt{\frac{|\mathcal{X}|}{n}} \cdot \underset{x_{1}^{\star}, x_{2}^{\star} \sim \mathcal{D}}{\mathbb{E}}\left[\mathbf{1}\left(x_{1}^{\star}=x_{2}^{\star}\right)\right] \\
& =\sqrt{\frac{|\mathcal{X}|}{n} \sum_{x \in \mathcal{X}} \mathcal{D}(x)^{2}} \stackrel{(\mathrm{a})}{\leq} \sqrt{\frac{|\mathcal{X}|}{n} \sum_{x \in \mathcal{X}} \mathcal{D}(x) \cdot \frac{1}{\sigma|\mathcal{X}|}}=\frac{1}{\sqrt{\sigma n}}
\end{aligned}
$$

where (a) follows from the definition of a $\sigma$-smooth distribution. This completes the proof.

## C. 5 Upper Bounding Generalization Error: Proof of Lemma C. 5

In the proof of Lemma C.5, we shall need the following property of smooth distributions which is a slightly strengthened version of the coupling lemma in Lemma B.1. The proof of lemma C. 6 is presented in Appendix C.7.
Lemma C.6. Let $X_{1}, \cdots, X_{m} \sim Q$ and $P$ be another distribution with a bounded likelihood ratio: $d P / d Q \leq 1 / \sigma$. Then using external randomness $R$, there exists an index $I=I\left(X_{1}, \cdots, X_{m}, R\right) \in$ $[m]$ and a success event $E=E\left(X_{1}, \cdots, X_{m}, R\right)$ such that $\operatorname{Pr}\left[E^{c}\right] \leq(1-\sigma)^{m}$, and

$$
\left(X_{I} \mid\left(E, X_{\backslash I}\right)\right) \sim P
$$

Fix any realization of the Poissonized sample size $N \sim \operatorname{Poi}(n)$. Choose $m=4 \sigma^{-1} \log T$ in Lemma C.6, and without loss of generality assume that $N$ is an integral multiple of $m$. Since for any $\sigma$-smooth $\mathcal{D}_{t}$, it holds that

$$
\frac{\mathcal{D}_{t}(s)}{\mathcal{U}(\mathcal{X} \times\{ \pm 1\})(s)}=\frac{\mathcal{D}_{t}(x)}{\mathcal{U}(\mathcal{X})(x)} \cdot \frac{\mathcal{D}_{t}(y \mid x)}{\mathcal{U}(\{ \pm 1\})(y)} \leq \frac{2}{\sigma}
$$

the premise of Lemma C .6 holds with parameter $\sigma / 2$ for $P=\mathcal{D}_{t}, Q=\mathcal{U}(\mathcal{X} \times\{ \pm 1\})$. Consequently, dividing the self-generated samples $\widetilde{s}_{1}, \cdots, \widetilde{s}_{N}$ into $N / m$ groups each of size $m$, and running the procedure in Lemma C.6, we arrive at $N / m$ independent events $E_{1}, \cdots, E_{N / m}$, each with probability
at least $1-(1-\sigma / 2)^{m} \geq 1-T^{-2}$. Moreover, conditioned on each $E_{j}$, we can pick an element $u_{j} \in\left\{\widetilde{s}_{(j-1) m+1}, \cdots, \widetilde{s}_{j m}\right\}$ such that

$$
\left(u_{j} \mid\left(E_{j},\left\{\widetilde{s}_{(j-1) m+1}, \cdots, \widetilde{s}_{j m}\right\} \backslash\left\{u_{j}\right\}\right)\right) \sim \mathcal{D}_{t}
$$

For notational simplicity we denote the set of unpicked samples $\left\{\widetilde{s}_{(j-1) m+1}, \cdots, \widetilde{s}_{j m}\right\} \backslash\left\{u_{j}\right\}$ by $v_{j}$. As a result, thanks to the mutual independence of different groups and $s_{t} \sim \mathcal{D}_{t}$ conditioned on $s_{1: t-1}$ (note that we draw fresh randomness at every round), for $E \triangleq \cap_{j \in[N / m]} E_{j}$ we have

$$
\left(u_{1}, \cdots, u_{N / m}, s_{t}\right) \mid\left(E, s_{1: t-1}, v_{1}, \cdots, v_{N / m}\right) \stackrel{\mathrm{iid}}{\sim} \mathcal{D}_{t}
$$

Consequently, for each $j \in[N / m]$ we have

$$
\begin{aligned}
& \underset{s_{t} \sim \mathcal{D}_{t}, R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right) \mid E\right] \\
& =\underset{s_{t} \sim \mathcal{D}_{t}, \widetilde{s}_{1}, \cdots, \widetilde{s}_{N}}{\mathbb{E}}\left[L\left(\operatorname{opt}\left(\widetilde{s}_{1}, \cdots, \widetilde{s}_{N}, s_{1: t-1}, s_{t}\right), s_{t}\right) \mid E\right] \\
& =\underset{v, s_{1: t-1} \mid E}{\mathbb{E}}\left(\underset{s_{t}, u_{1}, \cdots, u_{N / m}}{\mathbb{E}}\left[L\left(\operatorname{opt}\left(s_{1: t-1}, v, u_{1}, \cdots, u_{N / m}, s_{t}\right), s_{t}\right) \mid\left(E, s_{1: t-1}, v\right)\right]\right) \\
& \stackrel{(\text { a) }}{=} \underset{v, s_{1: t-1} \mid E}{\mathbb{E}}\left(\underset{s_{t}, u_{1}, \cdots, u_{N / m}}{\mathbb{E}}\left[L\left(\operatorname{opt}\left(s_{1: t-1}, v, u_{1}, \cdots, u_{j-1}, s_{t}, u_{j+1}, \cdots, u_{N / m}, u_{j}\right), u_{j}\right) \mid\left(E, s_{1: t-1}, v\right)\right]\right) \\
& \stackrel{(\mathrm{b})}{=} \underset{v, s_{1: t-1} \mid E}{\mathbb{E}}\left(\underset{s_{t}, u_{1}, \cdots, u_{N / m}}{\mathbb{E}}\left[L\left(\operatorname{opt}\left(s_{1: t-1}, v, u_{1}, \cdots, u_{N / m}, s_{t}\right), u_{j}\right) \mid\left(E, s_{1: t-1}, v\right)\right]\right) \\
& =\underset{s_{t} \sim \mathcal{D}_{t}, R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, u_{j}\right) \mid E\right],
\end{aligned}
$$

where (a) follows from the conditional iid (and thus exchangeable) property of ( $u_{1}, \cdots, u_{N / m}, s_{t}$ ) after the conditioning, and (b) is due to the invariance of the ERM output after any permutation of the inputs. On the other hand, if $s_{t}^{\prime}, u_{1}^{\prime}, \cdots, u_{N / m}^{\prime}$ are independent copies of $s_{t} \sim \mathcal{D}_{t}$, by independence it is clear that

$$
\underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}, R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right) \mid E\right]=\underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}, R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, u_{j}^{\prime}\right) \mid E\right], \quad \forall j \in[N / m] .
$$

Consequently, using the shorthand $u_{0}=s_{t}, u_{0}^{\prime}=s_{t}^{\prime}$, we have

$$
\begin{aligned}
& s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}, R^{(t+1)} \\
& =\frac{1}{N / m+1}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)-L\left(h_{t+1}, s_{t}\right) \mid E\right] \\
& \leq \frac{1}{N / m+s_{t} \sim \mathcal{D}_{t}, R^{(t+1)}} \mathbb{E}\left[\sum_{j=0}^{N / m}\left(L\left(h_{t+1}, u_{j}^{\prime}\right)-L\left(h_{t+1}, u_{j}\right)\right) \mid E\right] \\
& \leq \frac{1}{u_{0}, \cdots, u_{N / m}, u_{0}^{\prime}, \cdots, u_{N / m}^{\prime} \sim \mathcal{D}_{t}}\left[\sup _{h \in \mathcal{H}}^{N / m}\left(L\left(h, u_{j}^{\prime}\right)-L\left(h, u_{j}\right)\right)\right] \\
& \leq \frac{2}{N / m+1} \underset{u_{0}, \cdots, u_{N / m} \sim \mathcal{D}_{t}}{\mathbb{E}} \underset{\epsilon_{1} \ldots \ldots \epsilon_{N / m}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}^{N / m} \sum_{j=0}^{N / m} \epsilon_{j} h\left(u_{j}\right)\right] \leq c_{0} \sqrt{\frac{d}{N / m+1}},
\end{aligned}
$$

where the last inequality is due to the classical $O(\sqrt{d / n})$ upper bound on the Rademacher complexity, and $c_{0}>0$ in an absolute constant. Note that the union bound gives

$$
\operatorname{Pr}\left[E^{c}\right] \leq \sum_{j=1}^{N / m} \operatorname{Pr}\left[E_{j}^{c}\right] \leq \frac{N}{m T^{2}}
$$

the law of total expectation gives

$$
\begin{aligned}
& \underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}, R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)-L\left(h_{t+1}, s_{t}\right)\right] \\
& \leq \underset{s_{t}, s_{t}^{\prime} \sim \mathcal{D}_{t}, R^{(t+1)}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}^{\prime}\right)-L\left(h_{t+1}, s_{t}\right) \mid E\right]+\operatorname{Pr}\left[E^{c}\right] \leq c_{0} \sqrt{\frac{d}{N / m+1}}+\frac{N}{m T^{2}} .
\end{aligned}
$$

Finally, plugging the choice of $m=4 \sigma^{-1} \log T$, taking the expectation of $N \sim \operatorname{Poi}(n)$, and using $\operatorname{Pr}[N>n / 2] \geq 1-e^{-n / 8}$ in the above inequality completes the proof of Lemma C.5.

## C. 6 Proof of Theorem 3.2 for Small Domain

In this section we complete the proof of the $O\left(\sqrt{T(d|\mathcal{X}|)^{1 / 2}}\right)$ upper bound in Theorem 3.2 when $\sigma<d /|\mathcal{X}|$ (and thus $n=T \sqrt{|\mathcal{X}| / d}$ ). The proof is still through the same relaxation in Equation (10), though we will choose a different parameter $\eta$ and prove a slightly modified version of Lemma C.3:
Lemma C. 7 (Expected TV $\Rightarrow$ Admissibility). Let $\mathcal{Q}_{t}$ denote learner's distribution over $\mathcal{H}$ in Algorithm 2 at round $t$, and $s_{t} \sim \mathcal{D}_{t}$ be the conditional distribution of $s_{t}$ given the history $s_{1}, \cdots, s_{t-1}$. It holds that

$$
\underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left(\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]+\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right)\right)-\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right) \leq \underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\operatorname{TV}\left(\mathcal{Q}_{t}, \mathcal{Q}_{t+1}\right)\right]-\eta
$$

Proof of Lemma C.7. The analysis is similar to the proof of Lemma C.3. In fact, an intermediate step of Lemma C. 3 gives

$$
\underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]+\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right)-\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t-1}\right) \leq \underset{h_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[L\left(h_{t}, s_{t}\right)\right]-\underset{h_{t+1} \sim \mathcal{Q}_{t+1}}{\mathbb{E}}\left[L\left(h_{t+1}, s_{t}\right)\right]-\eta .
$$

Now using $\left|\mathbb{E}_{X \sim P}[f(X)]-\mathbb{E}_{X \sim Q}[f(X)]\right| \leq \operatorname{TV}(P, Q)$ for every measurable function $f$ with $\|f\|_{\infty} \leq 1$, the RHS is further upper bounded by $\operatorname{TV}\left(\mathcal{Q}_{t}, \mathcal{Q}_{t+1}\right)-\eta$. The proof of Lemma C. 7 is completed by taking the expectation over $s_{t} \sim \mathcal{D}_{t}$.

Note that in Lemma C.7, the expectation is outside the TV distance and no smaller than the TV distance when the mixture distribution is inside the expectation compared with Lemma C.3. We can simply upper bound this expected TV distance, with the worst case choice of $s_{t}$ and apply the data processing inequality, i.e.,

$$
\underset{s_{t} \sim \mathcal{D}_{t}}{\mathbb{E}}\left[\operatorname{TV}\left(\mathcal{Q}_{t}, \mathcal{Q}_{t+1}\right)\right] \leq \sup _{s_{t}} \operatorname{TV}\left(\mathcal{P}^{t-1}, \mathcal{P}^{t}\right)
$$

Using the similar independence property of Poissonization in Appendix C.4, the target TV distance is at most $\operatorname{TV}(P, Q)$, where $P \sim \operatorname{Poi}(n / 2|\mathcal{X}|)$, and $Q$ is a right-translation of $P$ by one. Consequently,

$$
\mathrm{TV}(P, Q) \leq \sqrt{\frac{\chi^{2}(Q, P)}{2}}=\sqrt{\frac{1}{2}\left(\underset{X \sim P}{\mathbb{E}}\left[\left(\frac{X}{n / 2|\mathcal{X}|}\right)^{2}\right]-1\right)}=\sqrt{\frac{|\mathcal{X}|}{n}}
$$

so the choice of $\eta=\sqrt{|\mathcal{X}| / n}$ and Lemma C. 7 again makes the relaxation in Equation (10) admissible, and we complete the proof of Theorem 3.2.

## C. 7 Proof of Lemma C. 6

The proof is essentially similar to [BDGR22, Lemma 12], and we include it here for completeness. For each $i \in[m]$, compute the value $p_{i}=\sigma \frac{d P}{d Q}\left(X_{i}\right)$, which lies in $[0,1]$ due to the likelihood ratio upper bound. Now we draw an independent Bernoulli random variable $Y_{i} \sim \operatorname{Bern}\left(p_{i}\right)$, and define the random index $I$ and success event $E$ as follows:

$$
\begin{aligned}
E & \triangleq \cup_{i=1}^{m}\left\{Y_{i}=1\right\} \\
I & \triangleq \text { a uniformly random element of }\left\{i \in[m]: Y_{i}=1\right\}
\end{aligned}
$$

Note that $Y_{1}, \cdots, Y_{m}$ are mutually independent, and for each $i \in[m]$,

$$
\operatorname{Pr}\left[Y_{i}=1\right]=\underset{X_{i} \sim Q}{\mathbb{E}}\left[p_{i}\right]=\underset{X_{i} \sim Q}{\mathbb{E}}\left[\sigma \frac{d P}{d Q}\left(X_{i}\right)\right]=\sigma,
$$

we conclude that $\operatorname{Pr}[E]=1-(1-\sigma)^{m}$. For the second statement, we denote by $r_{i}$ the external randomness used in drawing $Y_{i} \sim \operatorname{Bern}\left(p_{i}\right)$, and by $r$ the external randomness used in the definition of $I$. Then for any measurable set $A \subseteq \mathcal{X}$,

$$
\begin{aligned}
& \operatorname{Pr}\left[X_{I} \in A \mid E, X_{\backslash I}\right] \\
& =\sum_{i, r_{\backslash i}, r} \operatorname{Pr}\left[X_{I} \in A \mid E, X_{\backslash I}, I=i, r_{\backslash i}, r\right] \cdot \operatorname{Pr}\left[I=i, r_{\backslash i}, r \mid E, X_{\backslash I}\right] \\
& =\sum_{i, r_{\backslash i}, r} \operatorname{Pr}\left[X_{i} \in A \mid E, X_{\backslash i}, I=i, r_{\backslash i}, r\right] \cdot \operatorname{Pr}\left[I=i, r_{\backslash i}, r \mid E, X_{\backslash I}\right] \\
& \stackrel{(\text { a) }}{=} \sum_{i, r_{\backslash i}, r} \operatorname{Pr}\left[X_{i} \in A \mid Y_{i}=1, X_{\backslash i}, r_{\backslash i}, r\right] \cdot \operatorname{Pr}\left[I=i, r_{\backslash i}, r \mid E, X_{\backslash I}\right] \\
& \stackrel{(\mathrm{b})}{=} \sum_{i, r_{\backslash i}, r} \operatorname{Pr}\left[X_{i} \in A \mid Y_{i}=1\right] \cdot \operatorname{Pr}\left[I=i, r_{\backslash i}, r \mid E, X_{\backslash I}\right] \\
& \stackrel{(\mathrm{c})}{=} \sum_{i, r_{\backslash i}, r} P(A) \cdot \operatorname{Pr}\left[I=i, r_{\backslash i}, r \mid E, X_{\backslash I}\right] \\
& =P(A),
\end{aligned}
$$

where (a) is due to the event $\left\{E, I=i, X_{\backslash i}, r_{\backslash i}, r\right\}$ is the same as $\left\{Y_{i}=1, X_{\backslash i}, r_{\backslash i}, r\right\}$ as long as the former event $\left\{E, I=i, X_{\backslash i}, r_{\backslash i}, r\right\}$ is non-empty (note that empty events do not contribute to the sum), (b) follows from the mutual independence of $\left(X_{i}, r_{i}, Y_{i}\right)_{i \in[m]}$ and $r$, (c) is due to

$$
\operatorname{Pr}\left[X_{i} \in A \mid Y_{i}=1\right]=\frac{\operatorname{Pr}\left[X_{i} \in A, Y_{i}=1\right]}{\operatorname{Pr}\left[Y_{i}=1\right]}=\frac{1}{\sigma} \underset{X_{i} \sim Q}{\mathbb{E}}\left[\mathbf{1}\left(X_{i} \in A\right) \sigma \frac{d P}{d Q}\left(X_{i}\right)\right]=P(A)
$$

The above identity shows that the conditional distribution of $X_{I}$ conditioned on $\left(E, X_{\backslash I}\right)$ is always $P$, as desired.

## D Unknown Smoothness Parameters

Suppose we have upper and lower bounds $\sigma_{\max }$ and $\sigma_{\min }$ on the exact value of $\sigma$, i.e., $\sigma_{\min } \leq \sigma \leq$ $\sigma_{\max }$. In this section, we introduce a meta algorithm that uses a geometric doubling approach to incorporate knowledge of $\sigma_{\max }$ and $\sigma_{\min }$ into the algorithms introduced in Section 3.
We start by constructing $\log \left(\sigma_{\max } / \sigma_{\min }\right)$ experts, where each expert $i$ runs a local version of our algorithm (can be either Algorithm 1 for the real-valued case or Algorithm 2 for the binary case) with parameter $\sigma_{i}=2^{i} \cdot \sigma_{\min }$. We then run Hedge on these experts. Note that the parameter $i^{\star}$ of the best expert satisfies $\frac{\sigma}{2} \leq \sigma_{i^{\star}} \leq \sigma$, so the expected regret of this expert matches the expected regret of the same algorithm running on true $\sigma$ up to a constant factor. Therefore, the expected regret of this meta algorithm is comparable to the bound in Theorem 3.1 and 3.2, with an additive term of order at most $O\left(\sqrt{T \log \log \left(\sigma_{\max } / \sigma_{\min }\right)}\right)$. The number of oracle calls also blows up only by $\log \left(\sigma_{\max } / \sigma_{\min }\right)$ per round. This could potentially be improved using a more aggressive step size for the Hedge meta algorithm.

## E Proof of Lower Bounds (Theorem E. 1 and Theorem 5.1)

Theorem E. 1 (Limitations of Algorithms). For any choice of the parameter $n$ in Algorithm 2 (or Algorithm 1), there exists a $\sigma$-smoothed online learning instance such that Algorithm 2 (or Algorithm 1) suffers from at least $\Omega\left(\min \left\{T, \sqrt{d T \sigma^{-1 / 2}}, \sqrt{T(d|\mathcal{X}|)^{1 / 2}}\right\}\right)$ expected regret.

## E. 1 Proof of Theorem E. 1

This section proves the regret lower bounds for Algorithm 2 and Algorithm 1 stated in Theorem E.1. We split the analysis into two subsections, and in each subsection we prove a large regret both when the sample size parameter $n$ is large and small.

## E.1.1 Lower Bound Analysis for Algorithm 2

We shall only prove the regret lower bound $\Omega\left(\sqrt{d T \sigma^{-1 / 2}}\right)$ under the assumption $\sigma \geq$ $\max \left\{d /|\mathcal{X}|,(d / T)^{2}\right\}$, for a smaller $\sigma$ only makes the worst-case regret larger, and the other lower bounds follow from this case by taking $\sigma=d /|\mathcal{X}|$ and $\sigma=(d / T)^{2}$, respectively. We split the analysis into two cases depending on the choice of parameter $n$.

Case I: Large $n$. When $n$ is large, or more specifically, when $n \geq T / \sqrt{\sigma}$, consider the behavior of Algorithm 2 on the following instance. Consider any domain $\mathcal{X}$ where $|\mathcal{X}|$ is an integral multiple of $d$, and partition $\mathcal{X}=\cup_{j=1}^{d} \mathcal{X}_{j}$ into $d$ sets $\left\{\mathcal{X}_{j}\right\}_{j \in[d]}$ with an equal size. Consider the following hypothesis class:

$$
\mathcal{H}=\left\{h: \mathcal{X} \rightarrow\{ \pm 1\} \mid h \text { is a constant on } \mathcal{X}_{i}, \forall i \in[d]\right\} .
$$

Clearly $\mathcal{H}$ has VC dimension $d$. The adversary chooses a hypothesis $h^{\star} \in \mathcal{H}$ uniformly at random, and sets $x_{t}$ to be uniformly distributed on $\mathcal{X}$. As for the label $y_{t}$, the adversary sets $y_{t}=h^{\star}\left(x_{t}\right)$. This adversary is 1 -smooth, and the best expert in $\mathcal{H}$ incurs a zero loss under this realizable setting. We claim that for each of the first $\min \{T, c \sqrt{n d}\}$ time steps, for an absolute constant $c>0$ sufficiently small, Algorithm 2 makes a mistake with $\Omega(1)$ probability. Summing over these steps, the expected regret of Algorithm 2 is then $\Omega(\min \{T, c \sqrt{n d}\})$, which gives Theorem E. 1 by our assumption $n \geq T / \sqrt{\sigma}$.
To prove this claim, we need the following lemma.
Lemma E. 2 (Minimum Error on Hallucinated Samples). For $N \sim \operatorname{Poi}(n)$ hallucinated samples $\left(x_{1}, y_{1}\right), \cdots,\left(x_{N}, y_{N}\right)$, if $n \geq d$, it holds that

$$
\mathbb{P}\left(\sum_{i=1}^{N} y_{i} \cdot \mathbf{1}\left(x_{i} \in \mathcal{X}_{j}\right) \geq \sqrt{\frac{n}{d}}\right)=\Omega(1), \quad \forall j \in[d] .
$$

Proof. For $j \in[d]$, let $n_{j,+}, n_{j,-}$ denote the number of hallucinate samples $\left(x_{i}, y_{i}\right)$ with $x_{i} \in \mathcal{X}_{j}$ and $y_{i}= \pm 1$, respectively. By the Poisson subsampling property, $\left\{n_{j, \pm}\right\}_{j \in[d]}$ are mutually independent $\operatorname{Poi}(n /(2 d))$ random variables. By definition of $\mathcal{H}$, we have

$$
n_{j,+}-n_{j,-}=\sum_{i=1}^{N} y_{i} \cdot \mathbf{1}\left(x_{i} \in \mathcal{X}_{j}\right)
$$

Consequently, the quantity of interest is $n_{j,+}-n_{j,-}$. As $n / d \geq 1$, by the Poisson tail property, both events $n_{j,+} \geq n /(2 d)+\sqrt{n / d} / 2$ and $n_{j,-} \leq n /(2 d)-\sqrt{n / d} / 2$ happen with $\Omega(1)$ probability, and their independence gives the claimed result.

Since $(d / T)^{2} \leq \sigma \leq 1$, we have $T \geq d$ and thus $n \geq T / \sqrt{\sigma} \geq d$, the premise of Lemma E. 2 holds. Consequently, at each time step $t \leq \min \{T, c \sqrt{n d}\}$ with $x_{t} \in \mathcal{X}_{j}$, with $\Omega(1)$ probability there are at least $\sqrt{n / d}$ net positive labels in the hallucinated samples, while the learner has only observed at most $\alpha c \sqrt{n / d}$ labels in the history with probability at least $1-1 / \alpha$, by Markov's inequality. By choosing constants $c>0$ small and $\alpha>0$ large, the perturbed leader will predict +1 depending only on the hallucination, and this prediction is independent of the choice of $h^{\star}$ and thus incurs an error with probability $1 / 2$. This proves the claim that before time $\min \{T, c \sqrt{n d}\}$, there is always $\Omega(1)$ probability of error.

Case II: Small $n$. Now we turn to the scenario where $n<T / \sqrt{\sigma}$. Consider the following learning instance: choose $\mathcal{X}_{0} \subseteq \mathcal{X}$ with $\left|\mathcal{X}_{0}\right|=\sigma|\mathcal{X}| \geq d$, the adversary always chooses $x_{t} \sim \mathcal{U}\left(\mathcal{X}_{0}\right)$, which is $\sigma$-smooth. Assuming that $\left|\mathcal{X}_{0}\right|$ is an integral multiple of $d$, we partition $\mathcal{X}_{0}=\cup_{j=1}^{d} \mathcal{X}_{j}$ into $d$ subsets with equal size. Condition on each $\mathcal{X}_{j}$, consider an alternating label sequence:

$$
\left(y_{t}: x_{t} \in \mathcal{X}_{j}\right)_{t=1}^{T}=(+1,-1,+1,-1, \cdots)
$$

The hypothesis class $\mathcal{H}$ consists of $2^{d}$ functions:

$$
\mathcal{H}=\left\{h: \mathcal{X} \rightarrow\{ \pm 1\} \mid h \text { is a constant on } \mathcal{X}_{j}, \forall j \in[d], \text { and } h(x) \equiv 1, \forall x \in \mathcal{X} \backslash \mathcal{X}_{0}\right\} .
$$

Clearly $\mathcal{H}$ has VC dimension $d$, and the best hypothesis in $\mathcal{H}$ incurs a cumulative loss $T / 2$.
Now we examine the performance of Algorithm 2. Let $r_{j}$ be the difference between the number of +1 and -1 labels in the hallucinated samples with feature in $\mathcal{X}_{j}$, similar to the proof of Lemma E. 2 we have $r_{j}=n_{j,+}-n_{j,-}$ for independent Poisson random variables $n_{j,+}, n_{j,-} \sim \operatorname{Poi}(n \sigma / 2 d)$. Suppose that ties are broken by always predicting -1 when calling the ERM oracle, we observe that Algorithm 2 always makes a mistake when $x \in \mathcal{X}_{j}$ and $r_{j}=0$ - this is the same counterexample where Follow-The-Leader (FTL) makes a mistake at every step. Moreover, when $r_{j} \neq 0$, Algorithm 2 makes $T / 2$ mistakes, same as the best expert in $\mathcal{H}$. Consequently, the expected regret of Algorithm 2 is at least $T \cdot \mathbb{P}\left(r_{j}=0\right)$, where

$$
\begin{aligned}
\mathbb{P}\left(r_{j}=0\right) & =\underset{N \sim \operatorname{Poi}(n \sigma / d)}{\mathbb{E}}\left[\mathbb{P}\left(\operatorname{Bin}\left(N, \frac{1}{2}\right)=\frac{N}{2}\right)\right]=\underset{N \sim \operatorname{Poi}(n \sigma / d)}{\mathbb{E}}\left[\Omega\left(\frac{\mathbf{1}(N \text { is even })}{\sqrt{N+1}}\right)\right] \\
& \stackrel{(\text { a) }}{=} \Omega\left(\frac{\mathbb{P}_{N \sim \operatorname{Poi}(n \sigma / d)}(N \text { is even })}{\sqrt{n \sigma / d+1}}\right) \stackrel{(\text { b) }}{=} \Omega\left(\frac{1}{\sqrt{n \sigma / d+1}}\right)=\Omega\left(\min \left\{1, \sqrt{\frac{d}{n \sigma}}\right\}\right) .
\end{aligned}
$$

In the above display, (a) follows from the conditional Jensen's inequality, and (b) is due to

$$
\mathbb{P}_{N \sim \operatorname{Poi}(\lambda)}(N \text { is even })=\sum_{k=0}^{\infty} e^{-\lambda} \frac{\lambda^{2 k}}{(2 k)!}=e^{-\lambda} \cdot \frac{e^{\lambda}+e^{-\lambda}}{2} \geq \frac{1}{2}
$$

This leads to the claimed regret lower bound in Theorem E.1.

## E.1.2 Lower Bound Analysis for Algorithm 1

Similar to the lower bound analysis for Algorithm 2, we also split into the cases where $n$ is large and $n$ is small, respectively. Recall that for Algorithm 1, the parameter $n$ is the number of random draws from the uniform distribution for each future time. Our current version of Algorithm 1 sets the parameter to be $n=K=\Theta(\log T / \sigma)$.

Case I: Large $n$. We first focus on the case where $n \geq 1 / \sqrt{\sigma}$. Consider the same construction of $\mathcal{X}, \mathcal{H}$ and the adversary in Case I of Section E.1.1.

The regret analysis is essentially the same as Section E.1.1. For every $t \leq T / 2$, the learner in Algorithm 1 essentially generates $(T-t) n \geq T /(4 \sqrt{\sigma})$ uniformly random samples (with replacement) in $\mathcal{X}$. A similar analysis to Lemma E. 2 shows that for each $j \in[d]$, with $\Omega(1)$ probability there are $\Omega\left(\sqrt{T /\left(d \sigma^{1 / 2}\right)}\right)$ more +1 labels than -1 labels within $\mathcal{X}_{j}$ in the hallucinated samples. Consequently, for $t \leq \min \left\{T / 2, \Omega\left((d T)^{1 / 2} \sigma^{-1 / 4}\right)\right\}$, two calls of the ERM oracle in Algorithm 1 will return the same hypothesis, and the learner's prediction is always +1 . Similar to Section E.1.1, these time steps lead to an $\Omega\left(\min \left\{T,(d T)^{1 / 2} \sigma^{-1 / 4}\right\}\right)$ regret.

Case II: Small $n$. Next we turn to the case where $n<1 / \sqrt{\sigma}$.
Consider the same construction as Case II in Appendix E.1.1. For the performance of Algorithm 1, let $r_{j}$ be the difference between the number of +1 and -1 labels in the hallucinated samples with input in $\mathcal{X}_{j}$. One can check that if $r_{j} \neq 0$, the learner makes half of the mistakes along the alternating sequence; if $r_{j}=0$, the fraction of mistakes becomes $3 / 4$ (Algorithm 1 cyclically predicts a wrong label and makes a random guess). Consequently, the expected regret of Algorithm 1 is lower bounded by $\Omega\left(T \cdot \mathbb{P}\left(r_{j}=0\right)\right)$. To compute the probability $\mathbb{P}\left(r_{j}=0\right)$, note that $r_{j}=2 M-N$, with $N \sim \operatorname{Bin}(n(T-t), \sigma / d)$ being the number of observations in $\mathcal{X}_{j}$ in the hallucinated data, and $M \mid N \sim \operatorname{Bin}(N, 1 / 2)$. Using a similar argument to Section E.1.1, this probability is lower bounded by $\Omega(\min \{1, \sqrt{d /(n \sigma)}\})$, as desired.

## E. 2 Proof of Theorem 5.1

Theorem 5.1 (Restated). For $1 / \sigma \geq d$, any proper algorithm which only has access to the ERM oracle and achieves a regret $o\left(\min \left\{T, \sqrt{T(d / \sigma)^{1 / 2}}\right\}\right)$ for any $\sigma$-smoothed online learning problem must have an $\omega(\sqrt{d / \sigma})$ total running time.

We make the following remarks.

1. First, although the lower bound of the regret and running time in Theorem 5.1 does not match the counterparts of Algorithm 2 in Theorem 3.2, the upper and lower bounds share the same $\Theta\left(\sigma^{-1 / 4}\right)$ dependence on $\sigma$. This suggests that the improvement from $\Theta\left(\sigma^{-1 / 2}\right)$ to $\Theta\left(\sigma^{-1 / 4}\right)$ thanks to Poissonization is not superfluous and might be fundamental. We also conjecture that for all efficient algorithms with runtime poly $(T, d, 1 / \sigma)$, the $\Theta\left(\sigma^{-1 / 4}\right)$ dependence is the best one can hope for in the regret of such algorithms, as opposed to the $\Theta(\sqrt{\log (1 / \sigma)})$ dependence in the statistical regret.
2. Second, Theorem 5.1 shows a poly $(d, 1 / \sigma)$ computational lower bound to achieve the statistical regret $\widetilde{O}(\sqrt{T d \log (1 / \sigma)})$, while the $\varepsilon$-net $\operatorname{argument}$ in [HRS22] requires a poly $\left(\sigma^{-d}\right)$ computational time. One may wonder whether this exponential dependence on $d$ is in fact unavoidable, and this is a missing feature not covered in [HK16]. This motivates the open question in Section 5:

Open Question. For $d / \sigma \gg T^{2}$ in the smoothed setting, does any algorithm achieving o $o(T)$ regret require $\Omega\left(\operatorname{poly}\left(T, 2^{d}, 1 / \sigma\right)\right)$ computational time given access to the ERM oracle?

The proof of Theorem 5.1 uses a similar idea to [HK16]. There are two lower bound arguments in [HK16]: one reduces the problem to the Aldous' problem, and the other is based on an explicit construction of the hard instance. Although both arguments could work for our problem, we adopt the latter which corresponds to Theorem 25 of [HK16]. In the sequel, we will always take the domain size $|\mathcal{X}|=1 / \sigma$ so that the smooth adversary becomes the usual adaptive adversary. In the next subsections, we first prove the theorem for the simpler case $d=1$, and then generalize our argument for any VC dimension $d$.

## E.2.1 The case $d=1$.

We first show how the argument in [HK16] proves the claimed $\omega(\sqrt{|\mathcal{X}|})$ computational lower bound when $T=\sqrt{|\mathcal{X}|}=\sqrt{1 / \sigma}$. Assuming that $N \triangleq \sqrt{|\mathcal{X}|}$ is an integer, we partition the domain $\mathcal{X}$ into disjoint subsets $\mathcal{X}_{1}, \cdots, \mathcal{X}_{N}$, each of size $N$. For each $x \in \mathcal{X}$, we associate two independent Rademacher variables $\varepsilon(x)$ and $\varepsilon^{\star}(x)$, and they are mutually independent across different $x \in \mathcal{X}$. For each $i \in[N]$, the adversary chooses $x_{i}^{\star} \sim \mathcal{U}\left(\mathcal{X}_{i}\right)$, and sets the hypothesis class $\mathcal{H}=\left\{h_{x}\right\}_{x \in \mathcal{X}}$ with

$$
h_{x}\left(x^{\prime}\right)= \begin{cases}\varepsilon^{\star}\left(x^{\prime}\right) & \text { if } x=x_{i}^{\star}, x^{\prime}=x_{j}^{\star}, \text { and } i \geq j \\ \varepsilon\left(x^{\prime}\right) & \text { otherwise }\end{cases}
$$

At each time $t \in[N]$, the adversary sets $x_{t}=x_{t}^{\star}$, and $y_{t}=h_{x_{N}^{\star}}\left(x_{t}\right)=\varepsilon^{\star}\left(x_{t}^{\star}\right)$. Under this setting, [HK16] proved the following lower bound.
Theorem E. 3 (Theorem 25 of [HK16], restated). Given access to the ERM oracle, any proper algorithm achieving an expected regret at most $N / 4$ requires $\Omega(N)=\Omega(\sqrt{|\mathcal{X}|})$ running time.

Here by running time, we assume that each oracle call takes unit time, and maintaining each element in the input $\left\{\left(x_{i}, y_{i}\right)\right\}_{i \in I}$ to the oracle also takes unit time. We also sketch the proof idea of Theorem E. 3 for completeness: the crucial observation is that, when the learner feeds the input $\left\{\left(x_{i}, y_{i}\right)\right\}_{i \in I}$ to the ERM oracle, the oracle can always return some $h \in\left\{h_{0}, h_{x_{1}^{\star}}, \cdots, h_{x_{j}^{\star}}\right\}$, where $h_{0}$ is any hypothesis in $\mathcal{H} \backslash\left\{h_{x_{1}^{\star}}, \cdots, h_{x_{N}^{\star}}\right\}$, and $j \in[N]$ is the largest index such that $x_{j}^{\star} \in\left\{x_{i}\right\}_{i \in I}$. See Lemma 27 of [HK16] for a proof. Therefore, the label $y_{t}=\varepsilon^{\star}\left(x_{t}^{\star}\right)$ at time $t$ will look random to the learner unless the learner has seen a function $h_{x_{s}^{\star}}$ for some $s \geq t$. By the above observation, this occurs only if the learner has set one (or more) of $\left\{x_{s}^{\star}\right\}_{s>t}$ as the input to the ERM oracle, but this requires one to find a random element in a size- $N$ set and thus take $\Omega(N)$ time (note that a proper algorithm only observes $\left\{x_{1}^{\star}, \cdots, x_{t-1}^{\star}\right\}$ at time $\left.t\right)$. Consequently, with $o(N)$ running time, the learner suffers from an $\Omega(N)$ loss with high probability, while the best expert incurs zero loss - giving the $\Omega(N)$ regret.
Since the restriction of $\mathcal{H}$ on any two elements $\left\{x, x^{\prime}\right\}$ with $x<x^{\prime}$ could only be one of the three possibilities: $\left\{\left(\varepsilon(x), \varepsilon\left(x^{\prime}\right)\right),\left(\varepsilon^{\star}(x), \varepsilon\left(x^{\prime}\right)\right),\left(\varepsilon^{\star}(x), \varepsilon^{\star}\left(x^{\prime}\right)\right)\right\}$, the VC dimension of $\mathcal{H}$ is 1 . Therefore, Theorem E. 3 gives a valid proof of Theorem 5.1 when $d=1$ and $T=\sqrt{1 / \sigma}$. For $T<\sqrt{1 / \sigma}$, the
above construction still gives the $\Omega(T)$ regret lower bound given $o(\sqrt{|\mathcal{X}|})$ computational time. For general $T>\sqrt{1 / \sigma}$, we make the following modification to the adversary: partition the time horizon $[T]$ into $N$ intervals $T_{1}, \cdots, T_{N}$, each of length $T / N$. For each $i \in[N]$ and $t \in T_{i}$, the adversary sets $x_{t}=x_{i}^{\star}$, and

$$
y_{t}= \begin{cases}h_{x_{N}^{\star}}\left(x_{t}\right) & \text { with probability } \frac{1}{2}+\delta, \\ -h_{x_{N}^{\star}}^{\star}\left(x_{t}\right) & \text { with probability } \frac{1}{2}-\delta\end{cases}
$$

Consequently, the best expert $h_{x_{N}^{\star}}$ incurs an expected cumulative loss $(1 / 2-\delta) T$. Meanwhile, as long as the learner cannot distinguish the distributions $\operatorname{Bern}(1 / 2+\delta)^{\otimes(T / N)}$ and $\operatorname{Bern}(1 / 2-\delta)^{\otimes(T / N)}$, she is not able to estimate $\varepsilon^{\star}\left(x_{i}^{\star}\right)$ based on labels $\left\{y_{t}\right\}_{t \in T_{i}}$ in the $i$-th interval. This condition is fulfilled when $\delta \asymp \sqrt{N / T}$. In addition, a similar argument for Theorem E. 3 shows that with an $o(N)$ computational time, the learner cannot predict future $x_{s}^{\star}$ either. Therefore, any proper learner with $o(N)=o(\sqrt{|\mathcal{X}|})$ computational time must incur a regret $\Omega(\delta T)=\Omega\left(\sqrt{T|\mathcal{X}|^{1 / 2}}\right)$, which is precisely the statement of Theorem 5.1 for $d=1$.

## E.2.2 General $d$.

In this section we lift the hypothesis construction for $d=1$ to general $d$. Since $1 / \sigma \geq d$, we assume that $1 /(\sigma d)$ is an integer. Partition $\mathcal{X}=\cup_{j=1}^{d} \mathcal{X}_{j}$ each of size $|\mathcal{X}| / d$, we apply the hypothesis class $\mathcal{H}$ in the previous section to each $\mathcal{X}_{j}$, and set the entire hypothesis class as

$$
\mathcal{H}_{d}=\left\{h=\left(h_{1}, \cdots, h_{d}\right) \in \mathcal{H}^{d}:\left.h\right|_{\mathcal{X}_{j}}=h_{j}, \forall j \in[d]\right\} .
$$

Clearly the VC dimension of $\mathcal{H}_{d}$ is $d$. The adversary is constructed as follows: partition $[T]$ into $d$ sub-intervals $T_{1}, \cdots, T_{d}$, each of size $T / d$. For the $i$-th sub-interval, we run the subroutine in the previous section independently on $\mathcal{X}_{i}$. Now suppose that the total runtime is $o(\sqrt{d|\mathcal{X}|})$, then for at least half of the sub-intervals, the runtime during each such interval is $o(\sqrt{|\mathcal{X}| / d})$. By the lower bound for $d=1$, the expected regret during each such sub-interval is

$$
\Omega\left(\min \left\{\frac{T}{d}, \sqrt{\frac{T}{d} \cdot\left(\frac{|\mathcal{X}|}{d}\right)^{1 / 2}}\right\}\right)=\Omega\left(\min \left\{\frac{T}{d}, \sqrt{T \cdot\left(\frac{|\mathcal{X}|}{d^{3}}\right)^{1 / 2}}\right\}\right)
$$

Summing over at least $d / 2$ such independent sub-problems, the total regret lower bound is then $\Omega\left(\min \left\{T, \sqrt{T(d|\mathcal{X}|)^{1 / 2}}\right\}\right)$, establishing the claim of Theorem 5.1.

## F Statistical Upper Bound for Real-valued Labels

In this section, we present a statistical upper bound achieved by a computationally inefficient algorithm. The $\mathscr{Q}$ be the algorithm that runs Hedge on a finite subset $\mathcal{H}^{\prime}$ on $\mathcal{H}$, where $\mathcal{H}^{\prime}$ is a $\epsilon$-cover of $\mathcal{H}$ with respect to the uniform distribution $\mathcal{U}(\mathcal{X})$. The regret upper bound of this algorithm is bounded as follows.

Theorem F. 1 (Statistical Upper Bound). For any $\sigma$-smooth adversary $\mathscr{D}_{\sigma}$, the algorithm $\mathscr{Q}$ described above has regret upper bound

$$
\mathbb{E}\left[\operatorname{RegREt}\left(T, \mathscr{D}_{\sigma}, \mathscr{Q}\right)\right] \in \widetilde{O}\left(\sqrt{T d \log \left(\frac{T}{d \sigma}\right)}+G d \log \left(\frac{T}{d \sigma}\right)\right)
$$

Proof. Let $\mathcal{H}^{\prime}$ be the smallest $\epsilon$-cover of $\mathcal{H}$ with respect to the uniform distribution, i.e., for any $h \in \mathcal{H}$, there exists a proxy $h^{\prime} \in \mathcal{H}^{\prime}$ such that $\mathbb{E}_{x \sim \mathcal{U}(\mathcal{X})}\left[\left|h(x)-h^{\prime}(x)\right|\right] \leq \epsilon$. By lemma B.2, the size of $\mathcal{H}$ can be upper bounded in terms of the pseudo dimension $d$ :

$$
\log \left(\left|\mathcal{H}^{\prime}\right|\right)=\log \mathcal{N}\left(\epsilon, \mathcal{H}, L_{1}(\mathcal{U}(\mathcal{X}))\right) \leq O\left(d \log \left(\frac{1}{\epsilon}\right)\right)
$$

Based on the net $\mathcal{H}^{\prime}$, we also define function class $\mathcal{G}$ as follows.

$$
\mathcal{G}=\left\{g_{h, h^{\prime}}(x)=\left|h(x)-h^{\prime}(x)\right|: h \in \mathcal{H}, h^{\prime} \in \mathcal{H}^{\prime} \text { is its proxy. }\right\}
$$

Letting $L\left(h, s_{1: T}\right)=\sum_{t=1}^{T} l\left(h\left(x_{t}\right), y_{t}\right)$ for all $h \in \mathcal{H}$, we now consider the following regret decomposition:

$$
\begin{aligned}
\mathbb{E}[\operatorname{REGRET}(T)] & =\mathbb{E}\left[\sum_{t=1}^{T} l\left(\widehat{y}_{t}, y_{t}\right)-\inf _{h \in \mathcal{H}} L\left(h, s_{1: T}\right)\right] \\
& =\mathbb{E}\left[\sum_{t=1}^{T} l\left(\widehat{y}_{t}, y_{t}\right)-\inf _{h^{\prime} \in \mathcal{H}^{\prime}} L\left(h^{\prime}, s_{1: T}\right)\right]+\mathbb{E}\left[\inf _{h^{\prime} \in \mathcal{H}^{\prime}} L\left(h^{\prime}, s_{1: T}\right)-\inf _{h \in \mathcal{H}} L\left(h, s_{1: T}\right)\right]
\end{aligned}
$$

Note that the first term is precisely the regret of Hedge on the cover $\mathcal{H}^{\prime}$. It is thus bounded by

$$
\mathbb{E}\left[\sum_{t=1}^{T} l\left(\widehat{y}_{t}, y_{t}\right)-\inf _{h^{\prime} \in \mathcal{H}^{\prime}} L\left(h^{\prime}, s_{1: T}\right)\right] \leq O\left(\sqrt{T \log \left|\mathcal{H}^{\prime}\right|}\right) \in O\left(\sqrt{T d \log \left(\frac{1}{\epsilon}\right)}\right)
$$

As for the second term, we reformulate it in terms of class $\mathcal{G}$ :

$$
\begin{align*}
& \mathbb{E}\left[\inf _{h^{\prime} \in \mathcal{H}^{\prime}} L\left(h^{\prime}, s_{1: T}\right)-\inf _{h \in \mathcal{H}} L\left(h, s_{1: T}\right)\right]=\mathbb{E}\left[\sup _{h \in \mathcal{H}} \inf _{h^{\prime} \in \mathcal{H}^{\prime}} \sum_{t=1}^{T} l\left(h^{\prime}\left(x_{t}\right), y_{t}\right)-l\left(h\left(x_{t}\right), y_{t}\right)\right] \\
& \stackrel{\text { (a) }}{\leq} \mathbb{E}\left[\sup _{h \in \mathcal{H}} \inf _{h^{\prime} \in \mathcal{H}^{\prime}} \sum_{t=1}^{T} G\left|h\left(x_{t}\right)-h\left(x_{t}\right)\right|\right]=G \cdot \underset{\mathscr{D}}{\mathbb{E}}\left[\sup _{g \in \mathcal{G}} \sum_{t=1}^{T} g\left(x_{t}\right)\right] \tag{13}
\end{align*}
$$

where (a) is because the loss function $l$ has Lipschitz constant $G$. Analogous to [HRS22, Claim 3.4], we apply the coupling argument in Lemma B. 1 to replace the adaptive sequence $x_{t} \mathrm{~s}$ by $z_{t, k} \mathrm{~s}$ that are sampled independently from the uniform distribution. Thus we obtain

$$
\underset{\mathscr{D}}{\mathbb{E}}\left[\sup _{g \in \mathcal{G}} \sum_{t=1}^{T} g\left(x_{t}\right)\right] \leq T^{2}(1-\sigma)^{K}+\underset{\mathcal{U}(\mathcal{X})}{\mathbb{E}}\left[\sup _{g \in \mathcal{G}} \sum_{t=1}^{T} \sum_{i=1}^{K} g\left(z_{t, k}\right)\right]
$$

The expected supremum can be further bounded in terms of the magnitude of $\mathcal{G}$ (i.e., $\epsilon$ ) as well as the pseudo dimension of the original hypothesis class $\mathcal{H}$. Using the bound in Lemma F.2, and together with Equation (13), we obtain

$$
\mathbb{E}[\operatorname{Regret}(T)] \leq \widetilde{O}\left(\sqrt{T d \log \left(\frac{1}{\epsilon}\right)}+G\left(T^{2}(1-\sigma)^{K}+T K \epsilon+\sqrt{T K \epsilon d \log \left(\frac{1}{\epsilon}\right)}\right)\right)
$$

In order to satisfy the condition on $n$ in lemma F. 2 and to make the failure probability of the coupling $\operatorname{argument}$ sufficiently small, we take $\alpha=10 \log (T), K=\frac{\alpha}{\sigma}, \epsilon=\Theta\left(\frac{d \sigma}{T \log (T)} \log \left(\frac{T \log (T)}{d \sigma}\right)\right)$. With this choice of parameters, we have $T^{2}(1-\sigma)^{K}=o(1)$ and

$$
\begin{aligned}
\mathbb{E}[\operatorname{RegRET}(T)] & \leq O\left(\sqrt{T d \log \left(\frac{1}{\epsilon}\right)}+G\left(\frac{T \log (T)}{\sigma} \epsilon+\sqrt{\frac{T \log (T)}{\sigma} \epsilon d \log \left(\frac{1}{\epsilon}\right)}\right)\right) \\
& \leq \widetilde{O}\left(\sqrt{T d \log \left(\frac{T}{d \sigma}\right)}+G d \log \left(\frac{T}{d \sigma}\right)\right)
\end{aligned}
$$

as desired.

Lemma F. 2 (Concentration for the expected value of supremum). When $n \geq \Omega\left(\frac{d}{\epsilon} \log \left(\frac{1}{\epsilon}\right)\right)$, we have

$$
\underset{\substack{\text { idd } \\ x_{1: n} \sim \mathcal{U}(\mathcal{X})}}{\mathbb{E}}\left[\sup _{g \in \mathcal{G}} \sum_{i=1}^{n} g\left(x_{i}\right)\right] \leq O\left(n \epsilon+\sqrt{n \epsilon d \log \left(\frac{1}{\epsilon}\right)}\right) .
$$

Proof. We will use the bound on expected values of suprema of empirical processes in [GK06, Theorem 3.1]. To apply their result, the first step is to establish a bound on the $L_{2}(P)$-covering number of class $\mathcal{G}$. Let $P_{n}=\frac{1}{n} \sum_{i=1}^{n} \delta_{x_{i}}$ be the empirical distribution based on independent samples $x_{1}, \cdots, x_{n}$. A similar argument to [BKP97, Lemma 2] gives us

$$
\mathcal{N}\left(\epsilon, \mathcal{G}, L_{2}\left(P_{n}\right)\right) \leq \mathcal{N}\left(\frac{\epsilon}{2}, \mathcal{H}, L_{2}\left(P_{n}\right)\right)^{2} .
$$

Thus we obtain

$$
\log \mathcal{N}\left(\epsilon, \mathcal{G}, L_{2}\left(P_{n}\right)\right) \leq 2 \log \mathcal{N}\left(\frac{\epsilon}{2}, \mathcal{H}, L_{2}\left(P_{n}\right)\right) \leq 2 \log \mathcal{M}\left(\frac{\epsilon}{2}, \mathcal{H}, L_{2}\left(P_{n}\right)\right) \leq O\left(d \log \left(\frac{1}{\epsilon}\right)\right)
$$

where $\mathcal{M}$ denotes the packing number and the last inequality is due to [Bar06, Theorem 3.1]. Therefore, for the function $H(x)=O(d \log x)$, we can guarantee that for any $\epsilon>1$,

$$
\log \mathcal{N}\left(\epsilon, \mathcal{G}, L_{2}\left(P_{n}\right)\right) \leq H(1 / \epsilon)
$$

satisfying the condition of [GK06, Theorem 3.1]. Therefore, when $n \geq \Omega\left(\frac{H(1 / \epsilon)}{\epsilon}\right)=$ $\Omega\left(\frac{d}{\epsilon} \log \left(\frac{1}{\epsilon}\right)\right)$, [GK06] gives us

$$
\underset{\mathcal{U}}{\mathbb{E}}\left[\sup _{g \in \mathcal{G}} \sum_{i=1}^{n}\left(g\left(x_{t}\right)-\mathbb{E}\left[g\left(x_{t}\right)\right]\right)\right] \leq O(\sqrt{n \epsilon H(1 / \epsilon)})=O\left(\sqrt{n \epsilon d \log \left(\frac{1}{\epsilon}\right)}\right)
$$

Finally, since $\mathbb{E}_{\mathcal{U}} g(x) \leq \epsilon$ for any $g \in \mathcal{G}$, we obtain

$$
\mathbb{E}\left[\sup _{g \in \mathcal{G}} \sum_{i=1}^{n} g\left(x_{t}\right)\right] \leq O\left(n \epsilon+\sqrt{n \epsilon d \log \left(\frac{1}{\epsilon}\right)}\right)
$$

and the proof is complete.

## G Proof of the Admissible Relaxation Framework

Proposition 4.1 (Restated). In the smoothed online learning setting, let $\mathscr{Q}=\left(\mathcal{Q}_{1}, \cdots, \mathcal{Q}_{T}\right)$ be an algorithm that is admissible with respect to relaxations $\operatorname{Rel}_{T}(\mathcal{H})$, then the following bound on the expected regret holds regardless of the strategies $\mathscr{D}_{\sigma}$ of the adversary:

$$
\mathbb{E}\left[\operatorname{RegREt}\left(T, \mathscr{Q}, \mathscr{D}_{\sigma}\right)\right] \leq \operatorname{Rel}_{T}(\mathcal{H} \mid \emptyset)+O(\sqrt{T})
$$

Proof of proposition 4.1. To prove this lemma we break the expected regret into two parts:

$$
\mathbb{E}[\operatorname{REGRET}(T)]=\underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}\left[\sum_{t=1}^{T} \underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]-\inf _{h \in \mathcal{H}} L\left(h, s_{1: T}\right)\right]+\underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}\left[\sum_{t=1}^{T} l\left(\widehat{y}_{t}, y_{t}\right)-\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]\right]
$$

For the first part, we use an inductive argument to show that

$$
\begin{equation*}
\underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}\left[\sum_{t=1}^{T} \underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]-\inf _{h \in \mathcal{H}} \sum_{t=1}^{T} l\left(h\left(x_{t}\right), y_{t}\right)\right] \leq \operatorname{Rel}_{T}(\mathcal{H} \mid \emptyset) \tag{14}
\end{equation*}
$$

According to the definition of admissibility, we have

$$
\begin{aligned}
& \underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}[\sum_{t=1}^{T} \underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right] \underbrace{-\inf _{h \in \mathcal{H}} \sum_{t=1}^{T} l\left(h\left(x_{t}\right), y_{t}\right)}_{\leq \operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: T}\right) \text { by 2nd condition of admissibility }}] \\
& \leq \underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}[\mathbb{E}[\sum_{t=1}^{T-1} \underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+\underbrace{\underset{x_{T} \sim \mathcal{D}_{T}}{\mathbb{E}}\left[\underset{\widehat{y}_{T} \sim \mathcal{Q}_{T}}{\mathbb{E}}\left[l\left(\widehat{y}_{T}, y_{T}\right)\right]+\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: T}\right)\right]}_{\leq \operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: T-1}\right) \text { by 1st condition of admissibility }} \mid s_{1: T-1}]] \\
& \leq \underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}\left[\mathbb{E}\left[\sum_{t=1}^{T-1} \underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: T-1}\right) \mid s_{1: T-1}\right]\right] \\
& =\underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}\left[\sum_{t=1}^{T-1} \underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]+\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: T-1}\right)\right],
\end{aligned}
$$

where the last step uses the tower property of conditional expectations.
Repeat this process for $(T-1)$ times and note that $\boldsymbol{\operatorname { R e l }}_{T}(\mathcal{H} \mid \emptyset)$ is a constant that does not dependent on $\mathscr{D}$ proves Equation (14).

Since the second part is the expected sum of a martingale difference sequence, we apply the AzumaHoeffding inequality and obtain

$$
\begin{equation*}
\underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}\left[\sum_{t=1}^{T} l\left(\widehat{y}_{t}, y_{t}\right)-\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l\left(\widehat{y}_{t}, y_{t}\right)\right]\right] \leq \int_{0}^{\infty} \exp \left(-\frac{2 t^{2}}{T}\right) d t \in O(\sqrt{T}) \tag{15}
\end{equation*}
$$

Combining Equation (14) and Equation (15) completes the proof.

## H Transductive Learning with $K$ Hints

## H. 1 Model

In the traditional transductive setting, the adversary releases the sequence of unlabeled instances $\left\{x_{t}\right\}_{t=1}^{T}$ to the learner before the game starts. We generalize this setting and introduce a $K$-hint version of transductive learning. In this setting, the exact sequence of instances is replaced with a sequence of $K$ hints per time step such that the set of hints at each time step includes the instance at that time step. More formally, before the interaction starts, the adversary releases $T$ sets (multisets) of size $K$ to the learner. We denote these sets by $\left\{Z_{t}=\left\{z_{t, 1}, \cdots, z_{t, K}\right\}\right\}_{t=1}^{T}$. On releasing these sets, the adversary promises to always pick $\mathcal{D}_{t}$ supported only on the elements of $Z_{t}$. The regret of a learner with prediction rules $\mathscr{Q}$ on the adaptive sequence $\mathscr{D}=\left(\mathcal{D}_{1}, \cdots, \mathcal{D}_{T}\right)$ following above restrictions is defined as:

$$
\mathbb{E}[\operatorname{REGRET}(T, \mathscr{D}, \mathscr{Q})]=\underset{\mathscr{D}, \mathscr{Q}}{\mathbb{E}}\left[\sum_{t=1}^{T} l\left(\widehat{y}_{t}, y_{t}\right)-\inf _{h \in \mathcal{H}} \sum_{t=1}^{T} l\left(h\left(x_{t}\right), y_{t}\right)\right] .
$$

## H. 2 Efficient Algorithm for Transductive Online Learning with $K$ Hints

We will show an oracle-efficient regret upper bound of $O(\sqrt{d T K})$ by constructing an oracle-efficient algorithm based on the random playout technique. We consider the optimization oracle defined in Definition 2.3 with the loss functions specified by $l^{\mathrm{r}}(\hat{y}, y)=\frac{1}{2 G} l(\hat{y}, y)$ and $l^{\mathrm{b}}(\hat{y}, y)=\mathbf{1}\{\widehat{y} \neq y\}-\frac{1}{2}$.

Similar to Algorithm 1, at each time step $t$, our algorithm applies the offline optimization oracle to two input sequences: One where the real history $s_{1: t-1}$ is mixed with two copies ${ }^{5}$ of randomly labeled set of all hints corresponding to future time steps and the current instance is labled +1 , and another, where the current instance is labeled -1 .
More specifically, with $\mathcal{E}^{(t)}=\left\{\epsilon_{i, k}^{(t)}\right\}_{i=t+1: T, k=1: K}$ denoting the set of random labels and $S^{(t)}=$ $\left(Z_{t+1: T}, \mathcal{E}^{(t)}\right)$ denoting the set of hints labeled by $\mathcal{E}^{(t)}$, we consider

$$
\begin{equation*}
\widehat{y_{t}}=\mathrm{OPT}\left(s_{1: t-1} ; S^{(t)} \cup S^{(t)} \cup\left\{\left(x_{t},-1\right)\right\}\right)-\mathrm{OPT}\left(s_{1: t-1} ; S^{(t)} \cup S^{(t)} \cup\left\{\left(x_{t},+1\right)\right\}\right) \tag{16}
\end{equation*}
$$

See Algorithm 3 for a formal description of the algorithm.

```
Algorithm 3: Oracle-Efficient Online Transductive Learning with \(K\) Hints
Input: \(T, K,\left\{Z_{t}\right\}_{t=1}^{T}\)
for \(t \leftarrow 1\) to \(T\) do
    Receive \(x_{t}\). Assert that \(x_{t} \in Z_{t}\)
    for \(i=t+1, \cdots, T ; k=1, \cdots, K\) do
            Draw new \(\epsilon_{i, k}^{(t)} \sim \mathcal{U}(\{-1,+1\})\).
        end
        \(S_{1}^{(t)} \leftarrow\left\{\left(z_{i, k}^{(t)}, \epsilon_{i, k}^{(t)}\right),\left(z_{i, k}^{(t)}, \epsilon_{i, k}^{(t)}\right)\right\}_{\substack{i=t+1: T \\ k=1: K}} \quad\) // Two copies of each tuple
        \(S_{2}^{(t)} \leftarrow\left\{\left(x_{\tau}, y_{\tau}\right)\right\}_{\tau=1}^{t-1}\)
        \(p_{t} \leftarrow \frac{1}{2}+\frac{\mathrm{oPT}_{\mathcal{H}, l}\left(S_{t}^{(1)} \cup S_{t}^{(2)} \cup\left\{\left(x_{t},+1\right)\right\}\right)-\mathrm{OPT}_{\mathcal{H}, l}\left(S_{t}^{(1)} \cup S_{t}^{(2)} \cup\left\{\left(x_{t},-1\right)\right\}\right)}{2}\).
        With probability \(p_{t}\), predict \(\widehat{y}_{t}=-1\); otherwise predict \(\widehat{y}_{t}=+1\)
        Receive \(y_{t}\), suffer loss \(l\left(\widehat{y}_{t}, y_{t}\right)\).
end
```

Theorem H. 1 (Regret Bound for Efficient $K$-Hint Transductive Learning). In the setting of transductive learning with K-hints, Algorithm 3 achieve expected regret bound of $O(\sqrt{d T K})$. The algorithm can be implemented using two calls to the optimization oracle per round.

The proof of Theorem H. 1 follows a similar approach to that in [RSS12] and uses the admissible relaxation framework. This proof is also very similar to the proof of Theorem 3.1 for the case of smoothed adversaries. We include the proof of this Theorem for completeness.
Specifically, we will show that the algorithm is admissible with respect to the following relaxation:

$$
\left.\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right)=\underset{\mathcal{E}^{(t)}}{\mathbb{E}}\left[\sup _{h \in \mathcal{H}}\left\{2 G \sum_{\substack{i=t+1: T \\ k=1: K}} \epsilon_{i, k}^{(t)} h\left(z_{i, k}\right)-\sum_{i=1}^{t} l\left(h\left(x_{i}\right), y_{i}\right)\right)\right\}\right], \quad t=0, \cdots, T
$$

Using the language of regularized Rademacher complexity introduced in Section 4.3, the relaxation at the end of time step $t$ can be written as the Rademacher complexity for the union of future hints, regularized by the past total loss. That is,

$$
\begin{equation*}
\operatorname{Rel}_{T}\left(\mathcal{H} \mid s_{1: t}\right)=2 G \cdot \mathfrak{R}\left(-L^{\mathrm{r}}\left(\cdot, s_{1: t}\right), Z_{t+1: T}\right) \tag{17}
\end{equation*}
$$

where $L^{\mathrm{r}}\left(h, s_{1: t}\right)=\sum_{i=1}^{t} l^{\mathrm{r}}\left(h\left(x_{i}\right), y_{i}\right)=\frac{1}{2 G} \sum_{i=1}^{t} l\left(h\left(x_{i}\right), y_{i}\right)$ for $h \in \mathcal{H}$.
To use the relaxation framework and Proposition 4.1, it suffices to establish two claims: 1) the relaxation in Equation (17) is admissible in the $K$-hint setting, 2) the value of this relaxation at the beginning of the game is not too large.
For the second claim, we notice that $\operatorname{Rel}_{T}(\mathcal{H} \mid \emptyset)$ is equal to the unregularized Rademacher complexity for the dataset that includes all the hints. Since there are at most $T K$ hints, the Rademacher

[^0]complexity is at most $\widetilde{O}(\sqrt{d T K})$ according to Lemma B.3. That's where we get the extra $\sqrt{K}$ in the bounds compared to the standard transductive setting.
The first claim is the more technically interesting one. For admissibility, here we will focus on proving the following bound
\[

$$
\begin{equation*}
\sup _{x_{t} \in Z_{t}} \underbrace{\sup _{y_{t} \in \mathcal{Y}}\left\{\underset{\widehat{y}_{t} \sim \mathcal{Q}_{t}}{\mathbb{E}}\left[l^{\mathrm{r}}\left(\widehat{y}_{t}, y_{t}\right)\right]+\mathfrak{R}\left(-L_{t}^{\mathrm{r}}, Z_{t+1: T)}\right)\right\}}_{\text {(a) }} \leq \sup _{x_{t} \in Z_{t}} \mathfrak{R}\left(-L_{t-1}^{\mathrm{r}}, Z_{t+1: T} \cup\left\{x_{t}\right\}\right) \leq \mathfrak{R}\left(-L_{t-1}^{\mathrm{r}}, Z_{t: T}\right) \tag{18}
\end{equation*}
$$

\]

where $L_{t}^{\mathrm{r}}(h)$ abbreviates for $L^{\mathrm{r}}\left(h, s_{1: t}\right)$.
Let us consider the L.H.S of the above inequality and note that for any fixed $x_{t}$, the term (a) captures the standard transductive learning setting with $Z_{t+1: T}$ being the set of unlabeled instances for the future. In this case, the convexity of loss function $l^{\mathrm{r}}$ together with the min-max theorem can be used to show that the learner's strategy $\mathcal{Q}_{t}$, which makes the two values inside the supremum over $\mathcal{Y}$ equalize as $y_{t}$ takes value -1 and +1 , is indeed the optimal strategy. At a high level, this technique which is also used by [RSS12], gives rise to the algorithm in Equation (16) and proves the first inequality in inequality (18). We refer the readers to [RSS12, Lemma 12] for more details about the proof.
The second transition in Inequality (18) can be established using the fact that regularized Rademacher complexity is monotone in the dataset, as shown in Lemma 4.1. Therefore, we have proved that the relaxation given by Equation (17) is admissible, and the final regret upper bound follows from Proposition 4.1.


[^0]:    ${ }^{5}$ We use two copies to scale the loss appropriately.

