

PAC learning with stable and private predictions

Vitaly Feldman*

Google Research

Yuval Dagan†

MIT

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We study binary classification algorithms for which the prediction on any point is not too sensitive to individual examples in the dataset. Specifically, we consider the notions of uniform stability (Bousquet and Elisseeff, 2001) and prediction privacy (Dwork and Feldman, 2018). Previous work on these notions shows how they can be achieved in the standard PAC model via simple aggregation of models trained on disjoint subsets of data. Unfortunately, this approach leads to a significant overhead in terms of sample complexity. Here we demonstrate several general approaches to stable and private prediction that either eliminate or significantly reduce the overhead. Specifically, we demonstrate that for any class C of VC dimension d there exists a γ -uniformly stable algorithm for learning C with excess error α using $\tilde{O}(d/(\alpha\gamma) + d/\alpha^2)$ samples. We also show that this bound is nearly tight. For ε -differentially private prediction we give two new algorithms: one using $\tilde{O}(d/(\alpha^2\varepsilon))$ samples and another one using $\tilde{O}(d^2/(\alpha\varepsilon) + d/\alpha^2)$ samples. The best previously known bounds for these problems are $O(d/(\alpha^2\gamma))$ and $O(d/(\alpha^3\varepsilon))$, respectively.

1. Introduction

For a domain X and $Y = \{0, 1\}$, let \mathcal{A} be a learning algorithm that given a dataset $S \in (X \times Y)^n$ and a data-point $x \in X$ outputs a randomized prediction $\mathcal{A}^S(x)$. (Equivalently, one can think of algorithm \mathcal{A} as returning a randomized classifier $\mathcal{A}^S : X \rightarrow \{0, 1\}$). We consider algorithms \mathcal{A} that for every x, y , $\Pr[\mathcal{A}^S(x) = y]$ is not sensitive to individual examples in S . More formally, we say that \mathcal{A} is an (ε, δ) -differentially private prediction algorithm (Dwork and Feldman, 2018) if for any pair of datasets S and S' that differ in a single example and every $(x, y) \in X \times Y$ we have

$$\Pr_{\mathcal{A}}[\mathcal{A}^S(x) = y] \leq e^\varepsilon \cdot \Pr_{\mathcal{A}}[\mathcal{A}^{S'}(x) = y] + \delta.$$

The setting when $\varepsilon = 0$ is equivalent to the δ -uniform stability¹(Bousquet and Elisseeff, 2002).

Stability is a classical approach for understanding and analysis of generalization bounds (Rogers and Wagner, 1978; Devroye and Wagner, 1979; Bousquet and Elisseeff, 2002; Shalev-Shwartz et al., 2010; Feldman and Vondrák, 2019). In practice, stability-inducing methods such as Bagging

* Now at Apple. Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing.

† Part of the work was done while the author was at Google Research.

1. Alternatively, uniform stability can be considered for algorithms that output a model $h : X \rightarrow [0, 1]$ that gives the confidence level at x . It is easy to see (Elisseeff et al., 2005), that our discussion and results extend to this setting by defining $h(x) := \mathbf{E}[\mathcal{A}^S(x)]$.

(Breiman, 1996) and regularization (Bousquet and Elisseeff, 2002; Shalev-Shwartz et al., 2010) are used to improve accuracy and robustness to outliers.

Prediction privacy was defined by Dwork and Feldman (2018) to study privacy-preserving learning in the setting where the description of the model trained on sensitive data is not accessible to the (potentially adversarial) users. Instead, the users have access to the model through an interface (or API) which is provided by the data curator. A query to this interface is an input point $x \in X$ and, in response, the user receives a prediction $y \in Y$ on x of the model. Thus the definition of prediction privacy is exactly the differential privacy of the output of the curator on a single query measured with respect to the training data.

Many of the existing ML systems, such as query prediction in online search and credit scores, are deployed in this way. At the same time, as demonstrated in a growing number of works, even such black-box exposure of a learned model presents significant privacy risks. For example, for models trained using typical deep learning algorithms, such access allows to infer whether a point x is present in the training dataset with high accuracy (Shokri et al., 2017; Long et al., 2017; Truex et al., 2018) and to complete sentences used in the training data (Carlini et al., 2019). Thus understanding of prediction privacy is a crucial step towards ensuring that such systems preserve the privacy of the users who contributed the training data. Bassily et al. (2018) consider the same notion while focusing on algorithms for answering multiple queries to the curator by aggregating multiple non-private learning algorithms.

For comparison, the standard setting of privacy-preserving learning aims to ensure that the model learned from the data is produced in a differentially private way. Thus this approach preserves privacy even when a potential adversary has complete access to the description of the predictive model. The downside of this strong guarantee is that for some learning problems, achieving the guarantee is known to have substantial additional costs, both in terms of sample complexity and computation. In particular, such guarantee may be an overkill in the setting where only the predictive interface needs to be exposed.

Another application of algorithms with private predictions is for labeling public unlabeled data. Training on data labeled in this way gives a differentially private learning algorithm (Hamm et al., 2016; Papernot et al., 2017, 2018; Bassily et al., 2018). It is also easy to see that these notions of stability ensure some level of protection against targeted data poisoning attacks (Biggio et al., 2012) in which an attacker can add examples to the dataset with the goal of changing the prediction on a point they choose.

Given the significant benefits of these notions of stability and privacy it is natural to ask what is the overhead of ensuring them. We will address this question in the classical context of learning a class of Boolean functions in the (agnostic) PAC learning framework (Valiant, 1984; Haussler, 1992; Kearns et al., 1994). Namely, for a class H of Boolean functions over X , we say that \mathcal{A} is an (α, β) -agnostic PAC learning algorithm for H if for every distribution D over $X \times Y$, given a dataset S of i.i.d. examples from D , \mathcal{A} outputs a hypothesis h that with probability at least $1 - \beta$ over the choice of S satisfies:

$$\Pr_{(x,y) \sim D, \mathcal{A}} [\mathcal{A}^S(x) \neq y] \leq \inf_{f \in H} \Pr_{(x,y) \sim D} [f(x) \neq y] + \alpha.$$

Namely, the *excess error* is at most α . In the realizable setting it is additionally known that there exists $f \in H$ such that for all (x, y) in the support of D , $y = f(x)$.

A simple, general, and well-known approach to improve stability is via averaging of models trained on disjoint subsets of data. Alternatively, one can pick a random subset of the dataset and run

the algorithm on that subset. Clearly, this technique improves stability by a factor k at the expense of using k times more data. It was used in an influential work of [Shalev-Shwartz et al. \(2010\)](#) to demonstrate that learnability of H implies existence of a uniformly stable algorithm for learning H and, to the best of our knowledge, this is the only known technique for learning an arbitrary class H of finite VC dimension with uniform stability. An unfortunate disadvantage of this technique is that it substantially increases the sample complexity. In particular, it uses $O(d/(\gamma\alpha^2))$ samples to learn a class of VC dimension d with uniform stability γ and excess error α (the dependence on β is logarithmic and thus we set it to a small fixed constant). By the same argument, in the realizable case $O(d/(\gamma\alpha))$ samples suffice.

A somewhat more careful aggregation and analysis are needed to achieve privacy since naive averaging does not improve the ε parameter. In the realizable case $\tilde{O}(d/(\alpha\varepsilon))$ samples are known to suffice (and this is tight), but in the agnostic case the best known bound was $\tilde{O}(d/(\alpha^3\varepsilon))$ ([Dwork and Feldman, 2018](#); [Bassily et al., 2018](#)). An additional factor of α results from the need to add some noise in the aggregation step while ensuring that it contributes at most α to the excess error.

For comparison, we note that ensuring that the entire model is produced with differential privacy cannot, in general, be done using a number of samples polynomial in d . For example, threshold functions on a line have VC dimension of 1 but require an infinite number of samples to learn with privacy ([Feldman and Xiao, 2015](#); [Alon et al., 2019](#)). The only general approach for PAC learning with differential privacy is the technique of [Kasiviswanathan et al. \(2011\)](#) that is based on the exponential mechanism ([McSherry and Talwar, 2007](#)). This approach leads to sample complexity of $O(\log(|H|)/(\alpha\varepsilon) + d/\alpha^2)$. In this bound $\log(|H|)$ can be replaced with the *representation dimension* ([Beimel et al., 2013a](#)) of H which can sometimes be lower than $\log(|H|)$.

1.1. Our contribution

We describe several new algorithms for agnostic PAC learning of arbitrary VC classes that significantly reduce the overheads of achieving stability and privacy of predictions. We first show a simple and natural algorithm that is γ -uniformly stable and is a (α, β) -agnostic PAC learner with the nearly optimal sample complexity of $\tilde{O}(d/(\gamma\alpha) + d/\alpha^2)$.

Theorem 1 *For every class H of VC dimension d and $\alpha, \beta, \gamma \in (0, 1)$, there exists a γ -uniformly stable (α, β) -agnostic PAC learning algorithm with sample complexity of*

$$n \leq O\left(\frac{d \log(1/\alpha)}{\gamma\alpha} + \frac{d + \log(1/\beta)}{\alpha^2}\right).$$

This bound is tight up to poly-logarithmic factors. First, a lower bound of $n = \Omega((d + \log(1/\beta))/\alpha^2)$ holds for any (α, β) -agnostic PAC algorithm (not necessarily stable). Secondly, we prove that $n = \Omega(d/(\gamma\alpha))$ is required even for the realizable setting. The proof of the lower bound is based on a similar bound for private prediction algorithms ([Dwork and Feldman, 2018](#), Theorem 2.2) and appears in Section D.

One way to interpret this bound is that α -uniform stability can be achieved essentially “for free”, that is, using asymptotically the same number of samples as is necessary for learning. In contrast, known general approaches for achieving stability do not give any non-trivial guarantees without increasing the sample complexity by at least a factor of $1/\gamma$ ([Shalev-Shwartz et al., 2010](#)).

This uniformly stable algorithm easily implies existence of an $(\varepsilon, 0)$ -differentially private prediction learning algorithm with sample complexity of $\tilde{O}(d/(\varepsilon\alpha^2))$. This step can be done using the aggregation technique in (Dwork and Feldman, 2018) but we give an even simpler conversion.

Corollary 2 *For every class H of VC dimension d and $\alpha, \beta, \varepsilon \in (0, 1)$, there exists an (α, β) -agnostic PAC learning algorithm with ε -differentially private prediction that has sample complexity of $\tilde{O}(d/\varepsilon\alpha^2)$.*

This bound implies that for $\varepsilon = \Omega(1)$, ε -differentially private prediction can be achieved without increasing the sample complexity (asymptotically).

One limitation of the bound in Corollary 2 is that it only improves on learning with differential privacy when α is sufficiently large. Specifically, the sample complexity of ε -DP (α, β) -agnostic PAC learning is $\tilde{O}(r/(\alpha\varepsilon) + d/\alpha^2)$, where r is the representation dimension of H (Beimel et al., 2013a). Thus Corollary 2 gives a better bound than samples complexity of ε -DP learning only when $\alpha \geq d/r$ (for most known classes of functions H , $r = \Theta(\log(|H|))$ (Feldman and Xiao, 2015)). The only known result in which learning with ε -private prediction requires fewer samples than ε -DP learning for all α is the algorithm for learning thresholds on a line with sample complexity of $\tilde{O}(1/(\varepsilon\alpha) + 1/\alpha^2)$ (Dwork and Feldman, 2018). Note that the VC dimension of this class is 1 while the representation dimension is $\Theta(\log |X|)$.

Our main technical results is the first general bound that for a number of classes improves on ε -DP (α, β) -agnostic PAC learning in all parameter regimes. Specifically, we describe an algorithm with sample complexity $\tilde{O}(d^2/(\varepsilon\alpha) + d/\alpha^2)$:

Theorem 3 *For every class H of VC dimension d and $\alpha, \beta, \varepsilon \in (0, 1)$, there exists an (α, β) -agnostic PAC learning algorithm with ε -differentially private prediction that has sample complexity of*

$$n \leq \tilde{O} \left(\frac{d^2}{\varepsilon\alpha} + \frac{d}{\alpha^2} \right).$$

For comparison, ε -DP learning of linear threshold functions over B^d , where $B \subseteq \mathbb{R}$ has sample complexity of $\Theta(d^2 \log(|B|)/(\varepsilon\alpha) + d/\alpha^2)$ (Feldman and Xiao, 2015). Thus the first term in our bound is better by a factor of $\log(|B|)$. Another example is the class consisting of (the indicator function of) lines on a plane over a finite field $GF(p)$. Its VC dimension is 1 while its representation dimension is $\Theta(\log p)$ (Feldman and Xiao, 2015). Thus our bound is better by a factor of $\log p$. Although these gaps are may appear to be small, we believe that there is an important conceptual difference between bounds in terms of size of H and those depending on the VC dimension, as apparent also in the classical learning theory. In particular, our bounds apply to infinite hypothesis classes.

The best known lower bound for the problem is $\Omega(d/(\varepsilon\alpha) + (d + \log(1/\beta))/\alpha^2)$ (Dwork and Feldman, 2018, Theorem 2.2). Thus the bound in Theorem 3 is tight with respect to ε and α , however, there is a gap in the dependence on d . Closing this gap is an interesting open problem.

Overview of the techniques: The algorithm we use to prove Theorem 1 combines two standard tools. First, we pick a random subset T of S of size $\tilde{O}(d/\alpha)$, or approximately a γ fraction of S . Classical results in VC theory imply that T is an α -net for H with high probability. We can therefore use it to define an α -cover H_T of H that has size of at most $(1/\alpha)^{O(d)}$. We then use the exponential mechanism over H_T as in (Kasiviswanathan et al., 2011) to sample a hypothesis

$h \sim H_T$. Finally, we predict according to h , outputting $h(x)$. A simple argument implies stability: firstly, the random choice of the subset T ensures that each example has small influence on the choice H_T and secondly, the exponential mechanism ensures that each example has small influence on the final output function. We remark that similar uses of an exponential mechanism over a data-dependent cover have appeared in prior and subsequent work on differentially private learning (Chaudhuri and Hsu, 2011; Beimel et al., 2013b, 2015; Bassily et al., 2019, 2020).

The above stable algorithm can be converted to a differentially private prediction algorithm (Corollary 2). In order to do so, we simulate it with $\gamma = \varepsilon\alpha$ and, in addition, randomly flip the output label with probability α . This additional noise increases the error by at most α . At the same time it ensures that both 0 and 1 are output with probability at least α . In particular, the additive guarantee of $\varepsilon\alpha$ implies a multiplicative guarantee of at most $(\alpha + \varepsilon\alpha)/\alpha = 1 + \varepsilon \leq e^\varepsilon$.

Next, we describe the second differentially private algorithm, from Theorem 3. The first step also uses an exponential mechanism over a cover defined using a random subset of S . The function output by this step is used to re-label S , and the result is used as a training set for the private prediction algorithm from (Dwork and Feldman, 2018) that assumes a realizable setting. It provides the final output. The proof relies on a substantially more delicate analysis of the differential privacy of the exponential mechanism when the set of functions it is applied to changes, combined with the effect of privacy amplification by subsampling (Kasiviswanathan et al., 2011). Although the final step of the algorithm and some elements of the proof are borrowed from the general technique of (Beimel et al., 2015), new techniques and ingredients are required to derive the final bound. A more detailed outline of the analysis can be found in Section 4.1.

We remark that an alternative way to derive Corollary 2 is to use the general relabeling approach for converting a realizable PAC learning algorithm to an agnostic one (Beimel et al., 2015). The algorithm and the analysis resulting from this approach are more involved than our proof of Corollary 2 (we use elements of this approach in the proof of Theorem 3). This approach has been communicated to us by Nissim and Stemmer (2018) and has been the starting point for this work. The details of this technique can be found in an (independent) work of Nandi and Bassily (2020) which applies this result to answering multiple prediction queries with privacy.

2. Preliminaries

Notation: We denote by X the domain, and by $Y = \{0, 1\}$ the label set. We use H to denote the underlying hypothesis class of functions from X to Y and d denotes the VC dimension of H . The dataset is denoted by $S = ((x_i, y_i))_{i=1}^n$, and the underlying distribution is denoted by D . Define $[n] = \{1, \dots, n\}$. For any $I \subseteq [n]$, denote $S_I = ((x_i, y_i))_{i \in I}$. Given a hypothesis $h: X \rightarrow Y$, denote the expected zero-one error of h by $L_D(h) := \Pr_{(x,y) \sim D}[h(x) \neq y]$ and the empirical error of h by $L_S(h) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{h(x_i) \neq y_i}$. We use C to denote positive universal constants.

PAC learning for randomized algorithms: We consider learning algorithms that receive a training set $S \in (X \times Y)^n$ and $x \in X$ and output a random prediction $\mathcal{A}^S(x)$, where $\hat{y} \in Y$. Given S , define the *population* and *training* losses of \mathcal{A}^S (respectively) by

$$L_D(\mathcal{A}^S) := \Pr_{(x,y) \sim D, \mathcal{A}}[\mathcal{A}^S(x) \neq y]; \quad L_S(\mathcal{A}^S) := \frac{1}{n} \sum_{i \in [n]} \Pr_{\mathcal{A}}[\mathcal{A}^S(x_i) \neq y_i].$$

Given $\alpha, \beta \in (0, 1)$, we say that \mathcal{A} is (α, β) -agnostic PAC learner for H if for every distribution D over $X \times Y$,

$$\Pr_{S \sim D^n} \left[L_D(\mathcal{A}^S) \leq \inf_{h \in H} L_D(h) + \alpha \right] \geq 1 - \beta.$$

Uniform stability and prediction privacy: A γ -uniformly stable (or, γ -stable for brevity) learner is a learner whose prediction probabilities at every point change by at most an additive γ when one example in S is changed. Formally, for any $S, S' \in (X \times Y)^n$ which differ in at most one example and any $x \in X$ and $y \in Y$,

$$\Pr_{\mathcal{A}}[\mathcal{A}^S(x) = y] \leq \Pr_{\mathcal{A}}[\mathcal{A}^{S'}(x) = y] + \gamma.$$

The notion of differentially private prediction (Dwork and Feldman, 2018) is an application of the definition of differential privacy (Dwork et al., 2006) to learning in the setting where the only output that is exposed to (potentially adversarial) users is predictions on their points. Formally, given $\varepsilon, \delta \geq 0$, we say that an algorithm \mathcal{A} gives (ε, δ) -private prediction if

$$\forall S, S', x, y : \Pr_{\mathcal{A}}[\mathcal{A}^S(x) = y] \leq e^\varepsilon \Pr_{\mathcal{A}}[\mathcal{A}^{S'}(x) = y] + \delta.$$

Postprocessing guarantees of differential privacy imply that any learning algorithm that outputs a predictor with (ε, δ) -differential privacy also gives (ε, δ) -private prediction. We say that an algorithm gives ε -private prediction when $\delta = 0$. Note that any ε -private algorithm is $(e^\varepsilon - 1)$ -stable and $(0, \delta)$ -private prediction is exactly the same as δ -stability.

2.1. Technical preliminaries

Nets for VC classes: Fix $\alpha \in (0, 1)$, a hypothesis class H and a distribution D over $X \times Y$. Given a subset $A \subseteq X$, we say that A is an α -net for H with respect to D if it satisfies the following: any $h, h' \in H$ that satisfy $h(x) = h'(x)$ for all $x \in A$, also satisfy $\Pr_{(x,y) \sim D}[h(x) \neq h'(x)] \leq \alpha$. We say that A is an α -net for H with respect to S if it is an α -net with respect to the uniform distribution over S . The following is a fundamental theorem in machine learning (see, e.g. Shalev-Shwartz and Ben-David, 2014, Section 28.3).

Lemma 4 *Let D be a distribution over $X \times Y$, fix $\alpha, \beta \in (0, 1)$, and let U be a set of $n' = O((d \log(1/\alpha) + \log(1/\beta))/\alpha)$ i.i.d. samples from D . Then, with probability at least $1 - \beta$, U is an α -net for H with respect to D . Furthermore, this holds also if the samples are selected without replacement (namely, re-sampling from D until one gets n' distinct elements).*

Note that this theorem is usually stated for i.i.d. samples with replacement. However, since repetitions do not matter for the definition of ε nets, one can only gain from sampling without replacement.

The class H_T of all possible labelings of a subset $T \subseteq X \times Y$: Given a subset $T \subseteq X \times Y$, we will create a hypothesis class H_T in the following manner: define an equivalence relation \sim_T over H , by $h \sim_T g$ if $h(x) = g(x)$ for all $(x, y) \in T$. Then, H_T will contain one representative from each equivalence class (chosen arbitrarily).

The growth function and the Sauer-Shelah lemma: Given a hypothesis class H , the *growth function of H* is defined as $\tau_n(H) := \sup_{A \subseteq X: |A|=n} |H_A|$. The well-known Sauer-Shelah lemma states (see, e.g., [Shalev-Shwartz and Ben-David, 2014](#), Lemma 6.10):

Lemma 5 *For any hypothesis class H of VC dimension d and $n > d + 1$,*

$$\tau_n(H) \leq \sum_{i=0}^d \binom{n}{i} \leq (em/d)^d.$$

Uniform convergence bounds: The following is a standard uniform bound on the estimation error of a hypothesis class H (e.g. [Shalev-Shwartz and Ben-David, 2014](#), Theorem 6.8):

Lemma 6 *Fix $\alpha, \beta \in (0, 1)$ and assume that $n \geq \Omega((d + \log(1/\beta))/\alpha^2)$. Then,*

$$\Pr_{S \sim D^n} [\forall h \in H, |L_D(h) - L_S(h)| \leq \alpha] \geq 1 - \beta.$$

The exponential mechanism: The exponential mechanism is a well-known ε -differentially private algorithm for selecting a candidate that approximately maximizes some objective that has low sensitivity ([McSherry and Talwar, 2007](#)). Following ([Kasiviswanathan et al., 2011](#)) we apply it to a hypothesis class H , dataset S and privacy parameter ε . Namely $\mathcal{A}_{H,\varepsilon}^{\text{exp}}(S)$ samples $h \in H$ with probability proportional to $\exp(-nL_S(h)\varepsilon/2)$. It is ε -differentially private (in the usual sense of ([Dwork et al., 2006](#))) and, in particular, it gives ε -private predictions. The utility guarantees for the expected function can for example be found in ([Steinke and Ullman, 2017](#)).

Lemma 7 *The exponential mechanism $\mathcal{A}_{H,\varepsilon}^{\text{exp}}$ gives ε -private predictions. Additionally, for any $\alpha \in (0, 1)$ and $S \in (X \times Y)^n$, if $n \geq 2 \ln(|H|)/(\varepsilon\alpha)$ then*

$$L_S(\mathcal{A}_{H,\varepsilon}^{\text{exp}}(S)) \leq \inf_{h \in H} L_S(h) + \alpha.$$

Privacy amplification by subsampling: The following lemma is an adaptation of the standard privacy amplification-by-sampling technique ([Kasiviswanathan et al., 2011](#)) to private prediction.

Lemma 8 *Let \mathcal{A}' be an algorithm operating on a sample of size $n' = \eta n$ for $\eta \in (0, 1)$ and let \mathcal{A} be the following algorithm that receives a sample of size n :*

1. *Select a uniformly random subset $T \subseteq S$ of size n' .*
2. *Run \mathcal{A}' on T .*

For any $\varepsilon \in (0, 1)$, if \mathcal{A}' gives ε -private prediction then \mathcal{A} gives $2\varepsilon\eta$ -private prediction.

Private prediction in the realizable setting: [Dwork and Feldman \(2018, Theorem 4.1\)](#) have shown that $\tilde{O}(d/(\varepsilon\alpha))$ samples suffice to learn with private prediction in the realizable setting (that is when there exists $h \in H$ with zero error).

Lemma 9 *For any $\varepsilon, \alpha, \beta$ and H , there exists an (α, β) -PAC learner with ε -private prediction $\mathcal{A}_{\varepsilon,\alpha,H}^R$ for distributions realizable by H with sample complexity of $n = \tilde{O}(d/(\alpha\varepsilon))$.*

3. Uniform Stability of PAC Learning

In this section we describe the uniformly stable PAC learning algorithm, prove a (nearly) matching lower bound on its sample complexity and derive the corollary for private prediction.

For convenience, we start by restating Theorem 1.

Theorem 10 (Thm. 1 restated) *For every class H of VC dimension d and $\alpha, \beta, \gamma \in (0, 1)$, there exists a γ -uniformly stable (α, β) -agnostic PAC learning algorithm with sample complexity of*

$$n = O\left(\frac{d \log(d/\alpha)}{\gamma\alpha} + \frac{d + \log(1/\beta)}{\alpha^2}\right).$$

We outline the algorithm and its analysis below. The details of the proof appear in Section B.

Proof outline: The algorithm, that receives as input $S \in (X \times Y)^n$ and $x \in X$, consists of two steps:

- Randomly select a subset $T \subseteq S$ of size $|T| = \gamma n/2$, and create the hypothesis class H_T , as defined in Section 2.1 (H_T contains all classifications of H on T).
- Run the exponential mechanism (from Section 2.1) on the set of hypotheses H_T evaluated on the set of examples S with privacy parameter $\gamma/4$, to draw a hypothesis $\hat{h} \sim \mathcal{A}_{H_T, \gamma/4}^{\text{exp}}(S)$. Output $\hat{h}(x)$.

The following properties are useful for the proof:

1. With high probability,

$$\inf_{h \in H_T} L_S(h) \leq \inf_{h \in H} L_S(h) + \alpha/3. \quad (1)$$

This follows from the fact that $T = \tilde{\Omega}(d/\alpha)$, hence it is an $\alpha/3$ net for H , with respect to the uniform distribution over S (Lemma 4).

2. From Sauer-Shelah lemma (Lemma 5), the cardinality of H_T is upper-bounded by $|H_T| \leq O(n^d)$.
3. The exponential mechanism in the second step is $\gamma/4$ -private, hence its stability parameter is $(e^{\gamma/4} - 1) \leq \gamma/2$. By Lemma 7, $|S| = O(\log |H_T| / (\gamma\alpha)) = \tilde{O}(d/(\gamma\alpha))$ suffices to ensure that

$$\mathbf{E}_{\hat{h} \sim \mathcal{A}_{H_T, \gamma/4}^{\text{exp}}(S)} [L_S(\hat{h})] \leq \inf_{h \in H_T} L_S(h) + \alpha/3. \quad (2)$$

The algorithm has excess error of at most α : combining Eq. (1) and Eq. (2), we get that with high probability over the choice of T , $\mathbf{E}_{\hat{h}} [L_S(\hat{h})] \leq \inf_{h \in H} L_S(h) + 2\alpha/3$, and from uniform convergence (Lemma 6) we obtain that $\mathbf{E}_{\hat{h}} [L_D(\hat{h})] \leq \inf_{h \in H} L_D(h) + \alpha$, with high probability over the choice of S and T . Next, we claim that it is γ -stable: assume that S and S' are two training sets that differ in one example. This example has a probability of at most $\gamma/2$ to appear in T and affect the creation of H_T . Additionally, the exponential mechanism is $\gamma/2$ stable. Composing the stability bounds from both steps we obtain that the algorithm is γ -stable as desired.

The tightness of the upper bound is implied by the following lower bound.

Theorem 11 *Let H be a class of functions of VC dimension d . Fix $\alpha \leq 1/4$. Let \mathcal{A} be a PAC learning algorithm that is γ -stable and for any realizable distribution D over $X \times Y$, $\mathbf{E}_{S \sim D^n} [L_D(\mathcal{A}^S)] \leq \alpha$. Then, $n \geq (d - 1)/(8\gamma\alpha)$.*

The proof follows the same structure as the proof of Theorem 2.2 in (Dwork and Feldman, 2018) and appears in Appendix D.

Prediction privacy via stability: We now observe that any $\alpha\varepsilon/2$ -stable (α, β) -PAC learning algorithm can be converted to a $(2\alpha, \beta)$ -PAC learning algorithm with ε -private prediction simply by flipping the predictions with probability α .

More formally, the algorithm consists of the following two steps:

1. Run the γ -stable algorithm from Theorem 1 with $\gamma = \varepsilon\alpha/2$ and approximation parameter α , and let y' be the output prediction.
2. With probability $1 - \alpha$ predict y' and otherwise predict $1 - y'$.

Let \mathcal{A} denote the algorithm. Its error is trivially bounded by 2α . Privacy follows from definition, using the facts that \mathcal{A} is $\varepsilon\alpha/2$ -stable and that for any $S \in (X \times Y)^n$, $x \in X$ and $y \in Y$, $\Pr[\mathcal{A}(S, x) = y] \geq \alpha$. Thus we obtain Corollary 2.

4. PAC Learning with Prediction Privacy

In this section we describe our main technical result for private prediction. Our algorithm is the first general algorithm that can PAC learn an arbitrary class H of VC dimension d with the (nearly) optimal dependence of the sample complexity on ε and α . However the dependence on the dimension d is quadratic. More formally, we prove the following upper bound on the sample complexity of the problem.

Theorem 12 (Theorem 3 restated) *For every class H of VC dimension d and $\alpha, \beta, \varepsilon \in (0, 1)$, there exists an (α, β) -agnostic PAC learning algorithm with ε -differentially private prediction that has sample complexity of*

$$n \leq \tilde{O} \left(\frac{d^2}{\varepsilon\alpha} + \frac{d}{\alpha^2} \right).$$

We present the algorithm and the outline of its analysis below. Additional technical details are given in Section C.

4.1. Overview of the algorithm and its analysis

The algorithm that achieves the claimed sample complexity is described below (with slight simplifications for the clarity of exposition):

- The first steps are similar to those of the stable algorithm: we draw a random subset $T \subseteq S$ and then select, using the exponential mechanism, a hypothesis $\hat{h} \in H_T$. The only difference is that T is now selected to be of size $|T| = O(\varepsilon|S|)$ (rather than $O(\gamma|S|)$) and the privacy parameter of the exponential mechanism is now set to be $\min(\varepsilon/2, d/(\alpha|S|))$ (rather than $\gamma/4$).

- In the second step, we “privatize” \hat{h} as follows:
 1. Let $S_{\hat{h}}$ denote the set of examples obtained by labeling the points in S with \hat{h} , namely, $S_{\hat{h}} := \{(x, \hat{h}(x)) : (x, y) \in S\}$.
 2. Let $\varphi(\hat{h})$ denote the PAC learner from Lemma 9, with sample $S_{\hat{h}}$, $d/(\alpha|S|)$ prediction privacy and excess error $O(\alpha)$. The final prediction is drawn according to the distribution of $\varphi(\hat{h})(x)$.

Let \mathcal{A} denote the algorithm and we first establish that it is (α, β) -PAC. As in the analysis of the stable algorithm, $\mathbf{E}[L_D(\hat{h})] \leq \inf_{h \in H} L_D(h) + O(\alpha)$. By the guarantees of the learner from Lemma 9, $\Pr_{(x,y) \sim D, \varphi}[\varphi(\hat{h})(x) \neq \hat{h}(x)] \leq O(\alpha)$. Combining the above two inequalities:

$$\mathbf{E}[L_D(\mathcal{A}^S)] = \mathbf{E}[L_D(\varphi(\hat{h}))] \leq L_D(\hat{h}) + O(\alpha) \leq \inf_{h \in H} L_D(h) + O(\alpha).$$

Next, we explain why the algorithm gives ε -private prediction. Denote by $h_{S,T}(x)$ the random prediction given a training set S and a fixed subset T : $h_{S,T}(x)$ is obtained by first selecting \hat{h} from H_T using the exponential mechanism and then outputting a random value drawn according to the distribution of $\varphi(\hat{h})(x)$. For the analysis, we extend the definition of $h_{S,T}(x)$ to sets T which are not necessarily subsets of S . Our goal is to show that for any $x \in X$ and $y \in Y$,

$$\Pr_{T \subseteq S, h_{S,T}(x)}[h_{S,T}(x) = y] \leq e^\varepsilon \Pr_{T' \subseteq S', h_{S',T'}(x)}[h_{S',T'}(x) = y], \quad (3)$$

where the probability is both over the selections of the subsets T and T' and over the random selections of $h_{S,T}(x)$ and $h_{S',T'}(x)$. We split Eq. (3) in two, first comparing $h_{S,T}$ with $h_{S,T'}$ and then $h_{S,T'}$ with $h_{S',T'}$: for any $x \in X$ and $y \in Y$,

$$\Pr_{T \subseteq S, h_{S,T}(x)}[h_{S,T}(x) = y] \leq e^{\varepsilon/2} \Pr_{T' \subseteq S', h_{S,T'}(x)}[h_{S,T'}(x) = y]; \text{ and} \quad (4)$$

$$\Pr_{T' \subseteq S', h_{S,T'}(x)}[h_{S,T'}(x) = y] \leq e^{\varepsilon/2} \Pr_{T' \subseteq S', h_{S',T'}(x)}[h_{S',T'}(x) = y]. \quad (5)$$

Eq. (5) follows from the fact that for any fixed T' , $h_{S,T'}(x)$ is the composition of two private algorithms: first, the exponential mechanism selects \hat{h} privately, and then $\varphi(\hat{h})(x)$ is a private prediction given any fixed \hat{h} .

Next, we sketch the proof of Eq. (4). The terms in both sides of the inequality can be calculated as follows: first T and T' are randomly drawn as subsets of S and S' , respectively. Then, the random predictions $h_{S,T}(x)$ and $h_{S,T'}(x)$ are made. Due to this structure, we can use *privacy amplification by sub-sampling* (Lemma 8)² to compare $h_{S,T}(x)$ with $h_{S,T'}(x)$. In particular, we prove that $h_{S,T}(x)$ is $O(1)$ -private as a function of T , and since $|T| = O(\varepsilon|S|)$, the sub-sampling boosts the privacy by a factor of $O(\varepsilon)$ and Eq. (4) follows. Hence, we are left with proving that $h_{S,T}(x)$ is $O(1)$ -private as a function of T :

$$\Pr_{h_{S,T}(x)}[h_{S,T}(x) = y] \leq e^{O(1)} \Pr_{h_{S,T'}(x)}[h_{S,T'}(x) = y], \quad (6)$$

for (almost) any T and T' that differ on one element.

² To prove Eq. (4), one cannot use Lemma 8 as stated. However, the proof in our case is analogous.

To prove Eq. (6), we fix T and T' that differ in one element and create a matching between the elements of H_T and the elements of $H_{T'}$. This matching is either one-to-one, one-to-two or two-to-one and for any matched pair $h \in H_T$ and $h' \in H_{T'}$, the following properties hold:

- $|L_S(h) - L_S(h')| \leq \alpha/d$.
- For any $x \in X$ and $y \in Y$, $\Pr[\varphi(h)(x) = y] \leq e^{O(1)} \Pr[\varphi(h')(x) = y]$, where the probability is over the random selections of $\varphi(h)(x)$ and $\varphi(h')(x)$.

These two properties imply that $\Pr[h_{S,T}(x) = y]$ is within constant factors of $\Pr[h_{S,T'}(x) = y]$: the first property ensures that the probability to select h from H_T via the exponential mechanism is within a constant factor from the probability to select h' from $H_{T'}$ (for any matched h and h'). The second property states that $\Pr[\varphi(h)(x) = y]$ and $\Pr[\varphi(h')(x) = y]$ are within a constant factor from each other. Hence, Eq. (6) follows from these two properties, and it remains to describe why they hold.

- First, we define the aforementioned matching between H_T and $H_{T'}$: $h \in H_T$ is matched with $h' \in H_{T'}$ if and only if $h(x) = h'(x)$ for all $(x, y) \in T \cap T'$.
- We apply Lemma 4 on η -nets, substituting $U = T \cap T'$ and $\eta = \alpha/d$. We obtain that $T \cap T'$ is an α/d -net for H with respect to the uniform distribution over S (with high probability over a random selection of T). Any pair of matched hypotheses satisfy $h(x) = h'(x)$ for all $(x, y) \in T \cap T'$, hence, by definition of η -nets, they satisfy

$$|\{(x, y) \in S : h(x) \neq h'(x)\}| \leq |S|\alpha/d. \quad (7)$$

- The first desired property, namely $|L_S(h) - L_S(h')| \leq \alpha/d$, follows immediately from Eq. (7).
- For the second property, recall that $\varphi(h)(x)$ is the output of a learner with $d/(\alpha|S|)$ -private prediction, trained on $\{(x, h(x)) : (x, y) \in S\}$. Eq. (7) implies that the training sets used to train $\varphi(h)(x)$ and $\varphi(h')(x)$ differ on at most $\alpha|S|/d$ examples. Applying the privacy guarantee $\alpha|S|/d$ times (referred to as *group privacy*), one obtains that $\Pr[\varphi(h)(x) = y] \leq (e^{d/(\alpha|S|)})^{\alpha|S|/d} \Pr[\varphi(h')(x) = y] = e \Pr[\varphi(h')(x) = y]$, as required. The second property holds, and the proof follows.

5. Conclusions and open problems

Our work investigates the sample complexity of agnostic PAC learning with two natural and practically-motivated constraints: uniform stability and prediction privacy. We demonstrate the first general techniques that go beyond the naive averaging that has been used in a number of theoretical and practical works (Breiman, 1996; Shalev-Shwartz et al., 2010; Hamm et al., 2016; Papernot et al., 2017, 2018; Dwork and Feldman, 2018; Bassily et al., 2018). Our results settle the sample complexity of uniformly stable PAC learning (up to logarithmic factors). At the same time, achieving prediction privacy appears to be a significantly more challenging task and there remains a gap between our best upper and lower bounds. Specifically, the best known lower bound for the problem is $\tilde{\Omega}(d/(\varepsilon\alpha) + d/\alpha^2)$ (Dwork and Feldman, 2018), whereas our upper bound is $\tilde{O}(\min\{d/(\varepsilon\alpha^2), d^2/(\varepsilon\alpha)\} + d/\alpha^2)$. Closing this gap is a natural open problem.

Another avenue for future research is the computational complexity of learning with private or stable predictions. Our algorithms rely on the exponential mechanism which cannot be implemented efficiently in general and is usually more computationally expensive than empirical risk minimization (used in the averaging-based approaches). Thus designing general techniques whose computational complexity is comparable to that of ERM is a natural open problem.

The focus of this work is the classical distribution-independent PAC model for binary classification. Sample complexity of learning in this framework is well-understood and it has also been used in many prior works on the sample complexity of privacy-preserving learning. This makes it particularly suitable for investigating the sample complexity of private and stable prediction. At the same time, other learning models are likely to be more suitable for understanding of specific practical applications. Thus it is important to investigate the cost of stability and private prediction in other learning models. One example of such cost has been recently demonstrated in the context of learning from long-tailed data distributions (Feldman, 2019). Namely, it was demonstrated that uniformly stable (or prediction private) algorithms are inherently sub-optimal for such data distributions.

Acknowledgments

We are grateful to Kobbi Nissim and Uri Stemmer for pointing out that the general approach for converting a realizable PAC learning algorithm to an agnostic one from (Beimel et al., 2015) can be used to prove Corollary 2.

References

- Noga Alon, Roi Livni, Maryanthe Malliaris, and Shay Moran. Private PAC learning implies finite littlestone dimension. In *STOC*, pages 852–860, 2019.
- Raef Bassily, Om Thakkar, and Abhradeep Thakurta. Model-agnostic private learning via stability. *CoRR*, abs/1803.05101, 2018. URL <http://arxiv.org/abs/1803.05101>.
- Raef Bassily, Shay Moran, and Noga Alon. Limits of private learning with access to public data. In *Advances in Neural Information Processing Systems*, pages 10342–10352, 2019.
- Raef Bassily, Albert Cheu, Shay Moran, Aleksandar Nikolov, Jonathan Ullman, and Zhiwei Steven Wu. Private query release assisted by public data. *arXiv preprint arXiv:2004.10941*, 2020.
- Amos Beimel, Kobbi Nissim, and Uri Stemmer. Characterizing the sample complexity of private learners. In *ITCS*, pages 97–110, 2013a.
- Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. In *APPROX-RANDOM*, pages 363–378, 2013b.
- Amos Beimel, Kobbi Nissim, and Uri Stemmer. Learning privately with labeled and unlabeled examples. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 461–477. Society for Industrial and Applied Mathematics, 2015.
- Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In *ICML*, pages 1467–1474, 2012.

- Olivier Bousquet and André Elisseeff. Stability and generalization. *JMLR*, 2:499–526, 2002.
- Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- Nicholas Carlini, Chang Liu, Jernej Kos, Úlfar Erlingsson, and Dawn Song. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *Usenix Security (to appear)*, 2019.
- Kamalika Chaudhuri and Daniel Hsu. Sample complexity bounds for differentially private learning. In *COLT*, pages 155–186, 2011.
- Luc Devroye and Terry J. Wagner. Distribution-free inequalities for the deleted and holdout error estimates. *IEEE Trans. Information Theory*, 25(2):202–207, 1979.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284, 2006.
- Cynthia Dwork and Vitaly Feldman. Privacy-preserving prediction. *CoRR*, abs/1803.10266, 2018. URL <http://arxiv.org/abs/1803.10266>. Extended abstract in COLT 2018.
- André Elisseeff, Theodoros Evgeniou, and Massimiliano Pontil. Stability of randomized learning algorithms. *Journal of Machine Learning Research*, 6:55–79, 2005. URL <http://www.jmlr.org/papers/v6/elisseeff05a.html>.
- Vitaly Feldman. Does learning require memorization? A short tale about a long tail. *CoRR*, abs/1906.05271, 2019. URL <http://arxiv.org/abs/1906.05271>.
- Vitaly Feldman and Jan Vondrák. High probability generalization bounds for uniformly stable algorithms with nearly optimal rate. *CoRR*, abs/1902.10710, 2019. URL <http://arxiv.org/abs/1902.10710>.
- Vitaly Feldman and David Xiao. Sample complexity bounds on differentially private learning via communication complexity. *SIAM J. Comput.*, 44(6):1740–1764, 2015.
- Jihun Hamm, Yingjun Cao, and Mikhail Belkin. Learning privately from multiparty data. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *ICML*, volume 48 of *Proceedings of Machine Learning Research*, pages 555–563, 2016. URL <http://proceedings.mlr.press/v48/hamm16.html>.
- D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100(1):78–150, 1992. ISSN 0890-5401.
- Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, June 2011.
- M. Kearns, R. Schapire, and L. Sellie. Toward efficient agnostic learning. *Machine Learning*, 17(2-3):115–141, 1994.
- Yunhui Long, Vincent Bindschaedler, and Carl A. Gunter. Towards measuring membership privacy. *CoRR*, abs/1712.09136, 2017. URL <http://arxiv.org/abs/1712.09136>.

- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *FOCS*, pages 94–103, 2007.
- Anupama Nandi and Raef Bassily. Privately answering classification queries in the agnostic pac model. In *Algorithmic Learning Theory*, pages 687–703, 2020.
- Kobbi Nissim and Uri Stemmer. Personal communication, 2018.
- Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian J. Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
- Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson. Scalable private learning with PATE. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rkZB1XbRZ>.
- W. H. Rogers and T. J. Wagner. A finite sample distribution-free performance bound for local discrimination rules. *The Annals of Statistics*, 6(3):506–514, 1978. URL <http://www.jstor.org/stable/2958555>.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670, 2010.
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy, SP 2017*, pages 3–18, 2017.
- Thomas Steinke and Jonathan Ullman. Subgaussian tail bounds via stability arguments. *arXiv preprint arXiv:1701.03493*, 2017. URL <https://arxiv.org/abs/1701.03493>.
- Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Lei Yu, and Wenqi Wei. Towards demystifying membership inference attacks. *arXiv preprint arXiv:1807.09173*, 2018.
- L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

Appendix A. Notations for the proof

A different notation for randomized classifiers: Recall that \mathcal{A}_S can be viewed as a randomized classifier, namely, $\mathcal{A}_S(x)$ is a random variable. In the formal proof, we use a different way to describe such classifiers. Instead, a randomized classifier is represented by $h: X \rightarrow [0, 1]$, where $h(x)$ represents the probability to classify x as 1. This notation is convenient for describing mixtures of hypotheses: if h satisfies that for all x , $h(x)$ is a mixture of $h_1(x), \dots, h_N(x)$ with fixed probabilities p_1, \dots, p_N , then according to the introduced notation, $h = \sum_{i=1}^N p_i h_i$. Further, the expected population loss equals $L_D(h) = \mathbf{E}_{(x,y) \sim D}[|h(x) - y|]$.

Comparing hypotheses ($h \preceq \lambda g + \kappa$): For two randomized hypotheses $h, g: X \rightarrow [0, 1]$ and $\lambda, \kappa \geq 0$, we write $h \preceq \lambda g + \kappa$ if for any $x \in X$ and $y \in \{0, 1\}$, $|h(x) - y| \leq \lambda|g(x) - y| + \kappa$. Note that in this notation, the condition of (ε, δ) -private prediction is equivalent to

$$\mathcal{A}(S) \preceq e^\varepsilon \cdot \mathcal{A}(S') + \delta.$$

Appendix B. Proof of Theorem 1

The algorithm receives as an input a training set S and a point $x \in X$, and outputs $y \in \{0, 1\}$ as a prediction for x defined below:

1. Select a uniformly random subset $I \subseteq [n]$ of size n' , where $n' \in [n]$ is a parameter to be defined later (n' should be thought of as $n' \approx \gamma n$).
2. Create the hypothesis class H_{S_I} containing all possible labelings of S_I , as defined in Section 2.1.
3. Execute the exponential mechanism $\mathcal{A}_{H_{S_I}, \gamma}^{\text{exp}}(S)$ with privacy parameter γ on the sample S to randomly select a hypothesis $h_{S, S_I} \in H_{S_I}$, and output $h_{S, S_I}(x)$.

We proceed with the formal definition. First, for any $S \in (X \times Y)^n$ and $T \in (X \times Y)^{n'}$, define $h_{S, T} = \mathcal{A}_{H_T, \gamma}^{\text{exp}}(S)$, viewing $h_{S, T}$ as a function from $X \rightarrow [0, 1]$. Note that we do not require $T \subseteq S$ in this definition. Next, define

$$\mathcal{A}^{\text{sta}}(S) := \frac{1}{\binom{n}{n'}} \sum_{I \subseteq [n]: |I|=n'} h_{S, S_I}.$$

The final prediction of the algorithm given $x \in X$ is $\mathcal{A}^{\text{sta}}(S)(x)$.

We are ready to state the main lemma:

Lemma 13 *Let $N_{\alpha, \beta, d}^{\text{net}}$ denote the smallest number n' that suffices for Lemma 4 to hold, given parameters α, β and d (the minimal size required for a random set to be an α -net with probability $1 - \beta$). Let $N_{k, \varepsilon, \alpha}^{\text{exp}}$ denote the sample complexity required for ε -differentially private exponential mechanism to be α -approximate given a hypothesis class of size k . Assume that*

$$n' \geq N_{\alpha, \alpha, d}^{\text{net}}; \quad n \geq N_{\tau_{n'}(H), \gamma, \alpha}^{\text{exp}}; \quad n \geq n'/\gamma.$$

Then

$$L_S(\mathcal{A}^{\text{sta}}(S)) - \inf_{h \in H} L_S(h) \leq 3\alpha$$

and \mathcal{A}^{sta} is 3γ -stable.

First, we prove Theorem 1 using Lemma 13, and then we prove this lemma.

Proof [Proof of Theorem 1] Given α, β and γ , we will show an algorithm with the desired sample complexity. It suffices to find an algorithm which is $O(\gamma)$ -stable and $(O(\alpha), O(\beta))$ -agnostic PAC. To do so, we apply the above algorithm \mathcal{A}^{sta} with parameter $n' = N_{\alpha, \alpha, d}^{\text{net}}$, and then select n as the smallest integer which satisfies the conditions of Lemma 13.

We apply Lemma 13 to get a bound on the empirical error of $\mathcal{A}^\gamma(S)$ and then apply uniform convergence to generalize to the data distribution. In particular, from Lemma 6, with probability $1 - \beta$ over S ,

$$L_D(\mathcal{A}^{\text{sta}}(S)) \leq L_S(\mathcal{A}^{\text{sta}}(S)) + \alpha \leq \inf_{h \in H} L_S(h) + 4\alpha \leq \inf_{h \in H} L_D(h) + 5\alpha.$$

■

Lastly, we prove Lemma 13:

Proof [Proof of Lemma 13] First we prove 3α -approximation. From the condition that $n \geq N_{\tau_{n'}(h), \varepsilon, \alpha}^{\text{exp}}$ for any I ,

$$L_S(h_{S, S_I}) \leq \min_{h \in H_{S_I}} L_S(h) + \alpha.$$

Since $n' \geq N_{\alpha, \alpha, d}^{\text{net}}$, with probability at least $1 - \alpha$ over the choice of I :

$$\min_{h \in H_{S_I}} L_S(h) \leq \inf_{h \in H} L_S(h) + \alpha.$$

Thus,

$$L_S(\mathcal{A}^{\text{sta}}(S)) = \frac{1}{\binom{n}{n'}} \sum_{I \subseteq [n]: |I|=n'} L_S(h_{S, S_I}) \leq \alpha + \frac{1}{\binom{n}{n'}} \sum_{I \subseteq [n]: |I|=n'} \min_{h \in H_{S_I}} L_S(h) \leq \inf_{h \in H} L_S(h) + 3\alpha.$$

Now we prove the 3γ -stability. Let S be a sample and let S' be obtained from S by replacing one sample. Without loss of generality we can assume that S' is obtained by removing (x_1, y_1) from S and adding (x'_1, y'_1) . Our goal is to show that $\mathcal{A}^{\text{sta}}(S) \preceq \mathcal{A}^{\text{sta}}(S') + 3\gamma$.

For any T , Lemma 7 implies that

$$h_{S, T} \preceq e^\gamma h_{S', T} \preceq (1 + 2\gamma) h_{S', T} \preceq h_{S', T} + 2\gamma,$$

using the inequality $e^x \leq 1 + 2x$ for $x \leq 1$. To conclude the proof, using the fact that $n'/n \leq \gamma$,

$$\begin{aligned} \mathcal{A}^{\text{sta}}(S) &= \frac{1}{\binom{n}{n'}} \sum_{I \subseteq [n]: |I|=n'} h_{S, S_I} \preceq \frac{1}{\binom{n}{n'}} \sum_{I \subseteq [n]: |I|=n'} h_{S', S_I} + 2\gamma \\ &= \frac{1}{\binom{n}{n'}} \sum_{I \subseteq [n] \setminus \{1\}: |I|=n'} h_{S', S_I} + \sum_{I \subseteq [n]: |I|=n', 1 \in I} h_{S', S_I} + 2\gamma \\ &\preceq \mathcal{A}^{\text{sta}}(S') + \frac{n'}{n} + 2\gamma \preceq \mathcal{A}^{\text{sta}}(S') + 3\gamma. \end{aligned}$$

■

Remark 14 *It is possible to slightly improve the bound in Theorem 1, replacing $\log(d/\alpha)$ with $\log(1/\alpha)$. In order to do so, one has to replace H_{S_I} with an α -net, namely, a set G_{S_I} of hypotheses from H which satisfies:*

$$\forall h \in H, \exists g \in G_{S_I}, \text{ s.t. } |\{(x, y) \in S_I: h(x) \neq g(x)\}| \leq \alpha |I|.$$

There exists such a set of cardinality $(1/\alpha)^{O(d)}$. Using this set, one can relax the requirement $n \geq N_{\tau_{n'}(H), \gamma, \alpha}^{\text{exp}}$ from Lemma 13 to $n \geq N_{1/\alpha^{O(d)}, \gamma, \alpha}^{\text{exp}}$, and the improved bound would follow.

Appendix C. Proof of Theorem 3

We will start by describing an algorithm which gives $(O(\varepsilon), O(\varepsilon\alpha))$ -private prediction (namely, (ε, δ) -privacy with $\delta = O(\varepsilon\alpha)$, see Section 2). We then convert it to an algorithm with $O(\varepsilon)$ -private prediction as in the proof of Corollary 2.

The algorithm receives a training set S and a point $x \in X$, and outputs a label y for x . It is defined as follows:

1. Select a subset $I \subseteq [n]$ of size $|I| = n'$ uniformly at random (n' is a parameter to be defined later, and should be thought of as $n' \approx \varepsilon n$).
2. Create the hypothesis class H_{S_I} , as defined in Section 2.1 (H_{S_I} contains all the different prediction patterns of hypotheses from H on S_I).
3. Draw a random $h \in H_{S_I}$ using the exponential mechanism with loss evaluated on S and privacy parameter $2/(\eta n)$, where $\eta > 0$ is a parameter to be defined later (should be thought of as ε/d).
4. Define the dataset S_h as the predictions of h on S , $S_h := ((x_i, h(x_i)))_{i=1}^n$ (where (x_i, y_i) is the i 'th sample from H). Recall that $\mathcal{A}_{\varepsilon, \alpha, H}^R$ denotes the realizable private learner from Lemma 9 with privacy ε and approximation error α (we will ensure that the sample size is sufficiently large to achieve these guarantees). Let $\varphi_S(h) := \mathcal{A}_{1/(\eta n), \alpha, H}^R(S_h)$ denote the learned hypothesis given the sample S_h .
5. Output $\varphi_S(h)(x)$.

Viewed as a deterministic function the hypothesis output by this algorithm is described as follows. First, for any $S \in (X \times Y)^n$ and $T \in (X \times Y)^{n'}$ (not necessarily $T \subseteq S$), define:

$$h_{S,T} = \frac{1}{Z} \sum_{h \in H_T} \lambda_h \varphi_S(h), \quad \text{where } \lambda_h = \exp(-L_S(h)/\eta) \quad \text{and} \quad Z = \sum_{h \in H_{S_I}} \lambda_h. \quad (8)$$

Then, for any S , define

$$h_S := \frac{1}{\binom{n}{n'}} \sum_{I \subseteq [n]: |I|=n'} h_{S, S_I}. \quad (9)$$

We now summarize the conditions that need to be satisfied to prove the main result.

Lemma 15 *For $\alpha, \beta, \varepsilon \in (0, 1)$ and a positive integer d , define the following values:*

- Let $N_{\alpha, \beta, d}^{\text{net}}$ be the smallest sample size such that for any distribution D and any hypothesis class H of VC dimension d , a random i.i.d. sample of size $N_{\alpha, \beta, d}^{\text{net}}$ is an α -net for H with probability $1 - \beta$. Lemma 4 gives that $N_{\alpha, \beta, d}^{\text{net}} = O((d \log(1/\alpha) + \log(1/\beta))/\alpha)$.
- Let $N_{\varepsilon, \alpha, \beta, d}^R$ denote the smallest sample size such that with probability $1 - \beta$ over the choice of S :

$$\forall h \in H, \quad \mathbf{E}_{(x,y) \sim D} [|h(x) - \varphi_S(h)(x)|] \leq \alpha, \quad (10)$$

where $\varphi_S(h)$ is obtained from h by applying the PAC learner with private prediction for the realizable case. As we explain in Remark 16, $N_{\varepsilon, \alpha, \beta, d}^R$ is essentially the same as the sample complexity stated in Lemma 9, and equals $\tilde{O}(d/(\varepsilon\alpha))$.

- Let $N_{\varepsilon,\alpha,k}^{\text{exp}}$ denote the sample size required for the ε -differentially private exponential mechanism to be α -approximate for any hypothesis class of size k . Lemma 7 implies that $N_{\varepsilon,\alpha,k}^{\text{exp}} \leq 2 \ln(k)/(\varepsilon\alpha)$.
- Let $N_{\alpha,\beta,d}^G$ denote the sample size required to ensure that for any hypothesis class H of VC dimension d , with probability $1 - \beta$ over the choice of the training set S :

$$\forall h \in H, |L_D(h) - L_S(h)| \leq \alpha.$$

Standard uniform convergence bounds (Lemma 6) give $N_{\alpha,\beta,d}^G = \Theta((d + \log(1/\beta))/\alpha^2)$.

Recall that $\tau_n(H)$ is the growth function of H , defined in Section 2.1. Let n' and η be parameters such that the following holds:

$$n \geq N_{\alpha,\beta,d}^G; \quad n' \geq N_{\alpha,\alpha,d}^{\text{net}}; \quad n' \geq N_{\eta,\alpha,d}^{\text{net}} + 1; \quad n \geq N_{2/(\eta n),\alpha,\tau_n(H)}^{\text{exp}}; \quad n \geq N_{1/(\eta n),\alpha,d}^R; \quad n' \leq \varepsilon n; \quad \varepsilon \geq 1/(\eta n).$$

Then, the hypothesis h_S defined in Eq. (9) above gives $(O(\varepsilon), O(\varepsilon\alpha))$ -private predictions. Additionally, for any set S ,

$$L_D(h_S) \leq \inf_{h \in H} L_D(h) + O(\alpha),$$

with probability $1 - \beta$.

Remark 16 We briefly explain why $N_{\varepsilon,\alpha,\beta,d}^R = \tilde{O}(d/(\varepsilon\alpha))$. In Theorem 4.1 in (Dwork and Feldman, 2018) shows that this sample complexity is sufficient for the existence of a PAC learning algorithm with ε -DP prediction for the realizable setting. This implies that for any fixed h , with probability $1 - \beta$ over the choice of S labeled by h , $\mathbf{E}_{(x,y) \sim D}[|h(x) - \varphi_S(h)(x)|] \leq \alpha$. In our case, we apply the algorithm to S_h , where h depends on S . Therefore we need the learning result to hold for all possible ways to label S by $h \in H$. Namely, that with probability $\geq 1 - \beta$, for all $h \in H$, $\mathbf{E}_{(x,y) \sim D}[|h(x) - \varphi_S(h)(x)|] \leq \alpha$.

The proof in (Dwork and Feldman, 2018) relies on having a (non-private) (α, β) -PAC learner with sample complexity $O((d + \log(1/\beta))/\alpha^2)$. In order for the stronger statement in Eq. (10) to hold, all we need is a non-private (α, β) -PAC learner that succeeds for all $h \in H$ at the same time. Namely, that with probability $\geq 1 - \beta$ over a random sample S of size $\tilde{\Theta}(d/\alpha)$ for all $h \in H$ the algorithm, given S_h outputs g that satisfies $\Pr_{x \sim D}[h(x) \neq g(x)] \leq \alpha$. For classes of VC dimension d , uniform convergence results imply that any learning algorithm that outputs g that is consistent with S_h will have this property for sample size $m = O((d + \log(1/\beta))/\alpha^2)$. This is the same bound as the one used in (Dwork and Feldman, 2018) and thus we get $N_{\varepsilon,\alpha,\beta,d}^R = \tilde{O}(d/(\varepsilon\alpha))$.

Theorem 3 follows from Lemma 15 by setting $n = \tilde{\Theta}(d^2/(\varepsilon\alpha) + d/\alpha^2)$, $n' = \tilde{\Theta}(d^2/\alpha)$, $\eta = \tilde{\Theta}(\alpha/d)$. To convert the prediction from $(\varepsilon, \varepsilon\alpha)$ -private to ε -private we flip the value of the final output h_S with probability α (as in the proof of Corollary 2). The conditions in Lemma 15 have the following roles:

- $n \geq N_{\alpha,\beta,d}^G$ is required for generalization (guaranteeing that the error relative to D is within α of the error on S).
- $n' \geq N_{\alpha,\alpha,d}^{\text{net}}$ is required to ensure that the sub-sampled set T is an α -net for H with probability $1 - \alpha$. This ensures that the hypothesis class H_T contains an α -optimal hypothesis.

- $n' \geq N_{\eta, \alpha, d}^{\text{net}} + 1$ guarantees that $T \cap T'$ is an $\eta \approx \alpha/d$ -net.
- $n \geq N_{2/(\eta n), \alpha, \tau_n(H)}^{\text{exp}}$ guarantees that the executed exponential mechanism is $2/(\eta n)$ -differentially private and α -approximate. The condition $\varepsilon \geq 1/(\eta n)$ guarantees that the exponential mechanism is $O(\varepsilon)$ -differentially private.
- $n \geq N_{1/(\eta n), \alpha, \beta, d}^R$ guarantees that the PAC learning algorithm for the realizable case is $1/(\eta n)$ -private and (α, β) -approximate.
- $n' \leq \varepsilon n$ is required for privacy amplification by sub-sampling from $O(1)$ -privacy to $O(\varepsilon)$ -privacy.

We now prove that the algorithm we defined is an $(O(\alpha), O(\beta))$ -PAC learning algorithm.

Lemma 17 *Under the conditions of Lemma 15 it holds that $L_D(h_S) \leq \inf_{h \in H} L_D(h) + 4\alpha$, with probability at least $1 - \beta$.*

Proof Assume that I is a uniformly random subset of $[n]$ of size n' . Since $n' \geq N_{\alpha, \alpha, d}^{\text{net}}$, with probability at least $1 - \alpha$ over the choice of I :

$$\min_{h \in H_{S_I}} L_S(h) \leq \inf_{h \in H} L_S(h) + \alpha.$$

In particular,

$$\mathbf{E}_I \left[\min_{h \in H_{S_I}} L_S(h) \right] \leq \inf_{h \in H} L_S(h) + 2\alpha,$$

where I is a uniformly random subset.

Next, let h'_{S, S_I} denote the output of the exponential mechanism on H_{S_I} , namely

$$h'_{S, S_I} = \frac{1}{Z} \sum_{h \in H_{S_I}} \lambda_h h.$$

Since the exponential mechanism is defined as α -approximate, it holds that $L_S(h'_{S, S_I}) \leq \min_{h \in H_{S_I}} L_S(h) + \alpha$.

Recall the assumption that $n = \Omega((d + \log(1/\beta))/\alpha^2)$, or, equivalently, $n \geq N_{\alpha, \beta, d}^G$. This, combining with the previous two steps, imply that with probability $1 - \beta$:

$$\mathbf{E}_I[L_D(h'_{S, S_I})] \leq \mathbf{E}_I[L_S(h'_{S, S_I})] + \alpha \leq \mathbf{E}_I \left[\min_{h \in H_{S_I}} L_S(h) \right] + 2\alpha \leq \inf_{h \in H} L_S(h) + 4\alpha \leq \inf_{h \in H} L_D(h) + 5\alpha. \quad (11)$$

To go from h'_{S, S_I} to h_{S, S_I} , note that $h_{S, S_I} = \varphi_S(h'_{S, S_I})$. From $n \geq N_{\varepsilon, \alpha, \beta, d}^R$ it follows that with probability $\geq 1 - \beta$: $L_D(\varphi_S(h)) \leq L_D(h) + \alpha$ for all $h \in H$. In particular, $L_D(h_{S, S_I}) \leq L_D(h'_{S, S_I}) + \alpha$. This and Eq. (11) imply that with probability $\geq 1 - 2\beta$,

$$L_D(h_S) = \mathbf{E}_I[L_D(h_{S, S_I})] \leq \mathbf{E}_I[L_D(h'_{S, S_I})] + \alpha \leq \inf_{h \in H} L_D(h) + 6\alpha.$$

■

The rest of this section is dedicated to proving that the algorithm is $(O(\varepsilon), O(\varepsilon\alpha))$ -private. Fix S and let S' be a set which is obtained from S by replacing one sample. Without loss of generality, we can assume that S' is obtained from S by replacing (x_1, y_1) with (x'_1, y'_1) . Our goal is to show that

$$h_S = \frac{1}{\binom{n}{n'}} \sum_{|I|=n'} h_{S,S_I} \preceq e^{O(\varepsilon)} h_{S'} + O(\varepsilon\alpha).$$

We will divide the summation over subsets $|I| = n'$ into three terms. Let \mathcal{I} denote the set of all such subsets, and we will partition it to \mathcal{I}' , \mathcal{I}'' and \mathcal{I}''' as follows:

- \mathcal{I}' : all subsets satisfying $1 \notin I$, namely, those that do not correspond to the changed point (x_1, y_1) .
- \mathcal{I}'' : those that satisfy $1 \in I$, and that samples with indices in $I \setminus \{1\}$ are *not* an η -net for H with respect to the uniform distribution over $S_{\{2, \dots, n\}}$. There are not many such sets: from the requirements $|I \setminus \{1\}| = n' - 1 \geq N_{\eta, \alpha, d}^{\text{net}}$ and $n' \leq \varepsilon n$, it follows that $|\mathcal{I}''| \leq \alpha |\mathcal{I} \setminus \mathcal{I}'| = \alpha |\mathcal{I}| n' / n \leq \alpha \varepsilon |\mathcal{I}|$.
- \mathcal{I}''' : the remaining sets. With respect to these, we will be able to perform the matching discussed in the proof outline.

Next, we will show prediction privacy of h_{S,S_I} with respect to I , for $I \in \mathcal{I}'''$:

Lemma 18 *Let $I \in \mathcal{I}'''$. Fix $i \in [n] \setminus I$, and let $I' = (I \setminus \{1\}) \cup \{i\}$. Then,*

$$h_{S,S_I} \preceq 4e^6 h_{S,S_{I'}} = 4e^6 h_{S,S_{I'}}. \quad (12)$$

Proof Note that the equality in Eq. (12) follows from the fact that $1 \notin I'$ and that S and S' differ only on the first sample. Hence, we prove the inequality. We will create a matching between H_{S_I} and $H_{S_{I'}}$ as follows: for any $h \in H_{S_I}$, let P_h denote the set of all hypotheses $g \in H_{S_{I'}}$ that satisfy $h(x_j) = g(x_j)$ for all $j \in I \setminus \{1\}$. Intuitively, P_h is the set of hypotheses matched with h , and it holds that $1 \leq |P_h| \leq 2$.

We will fix $h \in H_{S_I}$ and $g \in P_h$. Since $I \in \mathcal{I}'''$, $I \setminus \{1\}$ is an η -net with respect to $S_{\{2, \dots, n\}}$. From definition of η nets (Section 2.1), since h and g agree on $I \setminus \{1\}$, there are at most $\eta(n-1)$ indices $j \in [n] \setminus \{1\}$ for which $h(x_j) \neq g(x_j)$. Hence, the number of samples $(x, y) \in S$ for which $h(x) \neq g(x)$ is at most $\eta(n-1) + 1 \leq 2\eta n$ (we used that $1 \leq 1/\varepsilon \leq \eta n$). This implies the following properties:

- It immediately follows that $|L_S(h) - L_S(g)| \leq 2\eta$.
- Let λ_h and λ_g denote the coefficients of h and g in the exponential weight mechanism corresponding to the definitions of h_{S,S_I} and $h_{S,S_{I'}}$, respectively (as defined in Eq. (8)). We will bound λ_h in terms of λ_g :

$$\lambda_h = e^{-L_S(h)/\eta} \leq e^{-L_S(g)/\eta+2} = \lambda_g e^2.$$

- Let Z_I and $Z_{I'}$ denote the values of Z in the definitions of h_{S,S_I} and $h_{S,S_{I'}}$, respectively. It holds that

$$Z_I = \sum_{h \in H_{S_I}} \lambda_h \leq e^2 \sum_{h \in H_{S_I}} \sum_{g \in P_h} \lambda_g \leq 2e^2 \sum_{g \in H_{S_{I'}}} \lambda_g = 2e^2 Z_{I'},$$

using the fact that for any $g \in H_{S_{I'}}$, there exist at most two hypotheses $h \in H_{S_I}$ for which $g \in P_h$. Similarly, one can derive $Z_{I'} \leq 2e^2 Z_I$.

- We will compare $\varphi_S(h)$ with $\varphi_S(g)$. Recall that $\varphi_S(h)$ is the output of a $1/(\eta n)$ -private realizable learner on the training set S_h (and similarly for $\varphi_S(g)$ and S_g). Since h and g predict differently on at most $2\eta n$ samples from S , S_h and S_g differ on at most $2\eta n$ entries, hence $\varphi_S(h) \preceq (e^{1/(\eta n)})^{2\eta n} \varphi_S(g) = e^2 \varphi_S(g)$.

Combining the above inequalities, one obtains that

$$h_{S,S_I} = \frac{1}{Z_I} \sum_{h \in H_{S_I}} \lambda_h \varphi_S(h) \preceq \frac{2e^6}{Z_{I'}} \sum_{h \in H_{S_I}} \sum_{g \in P_h} \lambda_g \varphi_S(g) \preceq \frac{4e^6}{Z_{I'}} \sum_{g \in H_{S_{I'}}} \lambda_g \varphi_S(g) = 4e^6 h_{S,S_{I'}}.$$

■

Using an adaptation of *amplification by sub-sampling* (discussed in Section 2.1) and using Lemma 18, one can bound the summation over $I \in \mathcal{I}'''$:

$$\sum_{I \in \mathcal{I}'''} \frac{1}{\binom{n}{n'}} h_{S,S_I} \preceq \sum_{I \in \mathcal{I}'''} \sum_{i \in [n] \setminus I} \frac{4e^6}{(n-n') \binom{n}{n'}} h_{S,S'_{(I \setminus \{i\}) \cup \{i\}}} \preceq \sum_{J \in \mathcal{I}} \frac{4e^6 n'}{(n-n') \binom{n}{n'}} h_{S,S'_J} \preceq \sum_{J \in \mathcal{I}} \frac{O(\varepsilon)}{\binom{n}{n'}} h_{S,S'_J},$$

where we used Lemma 18 and the facts that $n' \leq \varepsilon n$ and $n - n' \geq n - \varepsilon n \geq n/2$. This implies that

$$\begin{aligned} h_S &= \sum_{I \in \mathcal{I}'} \frac{1}{\binom{n}{n'}} h_{S,S_I} + \sum_{I \in \mathcal{I}''} \frac{1}{\binom{n}{n'}} h_{S,S_I} + \sum_{I \in \mathcal{I}'''} \frac{1}{\binom{n}{n'}} h_{S,S_I} \\ &\preceq \frac{1}{\binom{n}{n'}} \sum_{J \in \mathcal{I}} h_{S,S'_J} + O(\varepsilon \alpha) + \frac{O(\varepsilon)}{\binom{n}{n'}} \sum_{J \in \mathcal{I}} h_{S,S'_J}, \end{aligned} \quad (13)$$

using the derived bound $|\mathcal{I}''| \leq \varepsilon \alpha |\mathcal{I}|$. To conclude the proof, it suffices to bound h_{S,S'_J} in terms of h_{S',S'_J} .

Lemma 19 Fix a subset $T \subseteq X$ of size n' . Then, $h_{S,T} \preceq e^{3\varepsilon} h_{S',T}$.

Proof For any $h \in H_T$, let $\lambda_{h,S}$ and $\lambda_{h,S'}$ denote the values of λ_h in the definition of $h_{S,T}$ and $h_{S',T}$, respectively. Similarly, define Z_S and $Z_{S'}$. Note the following:

$$\lambda_{h,S} = e^{-L_S(h)/\eta} \leq e^{-L_{S'}(h)/\eta + 1/(\eta n)} = \lambda_{h,S'} e^{1/(\eta n)} \leq e^\varepsilon \lambda_{h,S'},$$

where we used the condition $1/(\eta n) \leq \varepsilon$. The same calculation implies that $\lambda_{h,S'} \leq e^\varepsilon \lambda_{h,S}$, hence $Z_S \geq e^{-\varepsilon} Z_{S'}$. Next, since $\mathcal{A}_{1/(\eta n), \alpha, H}^R$ has $1/(\eta n)$ -private prediction, it holds that $\varphi_S(h) \preceq e^{1/(\eta n)} \varphi_{S'}(h) \preceq e^\varepsilon \varphi_{S'}(h)$. The proof concludes, by plugging $\lambda_{h,S} \leq e^\varepsilon \lambda_{h,S'}$, $Z_S \geq e^{-\varepsilon} Z_{S'}$ and $\varphi_S(h) \preceq e^\varepsilon \varphi_{S'}(h)$ in the definition of $h_{S,T}$. ■

The proof of Lemma 15 now follows from Lemma 19 and Eq. (13):

$$h_S \preceq (1+O(\varepsilon)) \frac{1}{\binom{n}{n'}} \sum_{J \in \mathcal{I}} h_{S,S'_J} + O(\varepsilon \alpha) \preceq (1+O(\varepsilon)) e^{3\varepsilon} \frac{1}{\binom{n}{n'}} \sum_{J \in \mathcal{I}} h_{S',S'_J} + O(\varepsilon \alpha) = (1+O(\varepsilon)) e^{3\varepsilon} h_{S'} + O(\varepsilon \alpha).$$

Appendix D. Proof of Theorem 11

By the definition of the VC dimension, there exists a subset $X' \subseteq X$ of size $|X'| = d$ such that the restriction of H to X' contains all binary functions over X' , and assume without loss of generality that $X' = \{1, \dots, d\}$.

Define a distribution $D_{X'}$ over X' as follows: for any $k \in \{1, \dots, d-1\}$, $D_{X'}(k) = 4\alpha/(d-1)$ and $D_{X'}(d) = 1 - 4\alpha$. For any hypothesis $h: X' \rightarrow \{0, 1\}$, let D_h be the following distribution over pairs $(x, y) \in X' \times Y$: first, $x \sim D_{X'}$ is drawn, then $y = h(x)$. Let D_h^n denote the distribution over n i.i.d. samples from D_h .

Given $S = ((x_i, y_i))_{i=1}^n$, denote by $S^{\text{flip}(k)}$ the set obtained from S by flipping all values y_i corresponding to $x_i = k$, namely $S^{\text{flip}(k)} = ((x_i, y'_i))_{i=1}^n$ where $y'_i = 1 - y_i$ if $x_i = k$ and otherwise $y'_i = y_i$. Let $\#_S(i) := |\{i \in [n]: x_i = i\}|$ denote the number of times that k appears in S .

Drawing $h: X' \rightarrow \{0, 1\}$ uniformly at random, we obtain the following:

$$\begin{aligned} \alpha &\geq \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} [L_{D_h}(\mathcal{A}(S))] = \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} \left[\sum_{k=1}^d D_{X'}(k) |h(k) - \mathcal{A}(S)(k)| \right] \\ &\geq \frac{4\alpha}{d-1} \sum_{k=1}^{d-1} \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} [|h(k) - \mathcal{A}(S)(k)|] \\ &= \frac{4\alpha}{d-1} \sum_{k=1}^{d-1} \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} \left[\frac{1}{2} |h(k) - \mathcal{A}(S)(k)| + \frac{1}{2} |(1-h(k)) - \mathcal{A}(S^{\text{flip}(k)})(k)| \right] \end{aligned} \quad (14)$$

$$\geq \frac{2\alpha}{d-1} \sum_{k=1}^{d-1} \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} \left[|(2h(k) - 1) + \mathcal{A}(S^{\text{flip}(k)})(k) - \mathcal{A}(S)(k)| \right] \quad (15)$$

$$\geq \frac{2\alpha}{d-1} \sum_{k=1}^{d-1} \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} \left[|2h(k) - 1| - |\mathcal{A}(S)(k) - \mathcal{A}(S^{\text{flip}(k)})(k)| \right] \quad (16)$$

$$\begin{aligned} &= \frac{2\alpha}{d-1} \sum_{k=1}^{d-1} \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} \left[1 - |\mathcal{A}(S)(k) - \mathcal{A}(S^{\text{flip}(k)})(k)| \right] \\ &\geq \frac{2\alpha}{d-1} \sum_{k=1}^{d-1} (1 - \gamma \mathbf{E}_h \mathbf{E}_{S \sim D_h^n} [\#_S(k)]) \end{aligned} \quad (17)$$

$$= 2\alpha(1 - \gamma \frac{4\alpha}{d-1} n),$$

where Eq. (14) follows from a change of measure argument, Eq. (15) and Eq. (16) follow from the triangle inequality and Eq. (17) follows from the fact that \mathcal{A} is γ -stable. We obtain that $1 - 4\gamma\alpha n/(d-1) \leq 1/2$ which implies that $n \geq (d-1)/(8\gamma\alpha)$ as required.