A Limitation of the PAC-Bayes Framework

Roi Livni
Department of Electrical Engineering
Tel-Aviv University
Israel
rlivni@tauex.tau.ac.il

Shay Moran
Department of Mathematics
Technion, Haifa
Israel
shaymoran1@gmail.com

Abstract

PAC-Bayes is a useful framework for deriving generalization bounds which was introduced by McAllester ('98). This framework has the flexibility of deriving distribution- and algorithm-dependent bounds, which are often tighter than VC-related uniform convergence bounds. In this manuscript we present a limitation for the PAC-Bayes framework. We demonstrate an easy learning task which is not amenable to a PAC-Bayes analysis.

Specifically, we consider the task of linear classification in 1D; it is well-known that this task is learnable using just $O(\log(1/\delta)/\epsilon)$ examples. On the other hand, we show that this fact can not be proved using a PAC-Bayes analysis: for any algorithm that learns 1-dimensional linear classifiers there exists a (realizable) distribution for which the PAC-Bayes bound is arbitrarily large.

1 Introduction

The classical setting of supervised binary classification considers *learning algorithms* that receive (binary) labelled examples and are required to output a *predictor* or a *classifier* that predicts the label of new and unseen examples. Within this setting, Probably Approximately Correct (PAC) generalization bounds quantify the success of an algorithm to approximately predict with high probability. The PAC-Bayes framework, introduced in [23, 34] and further developed in [22, 21, 30], provides PAC-flavored bounds to Bayesian algorithms that produce *Gibbs-classifiers* (also called *stochastic-classifiers*). These are classifiers that, instead of outputting a single classifier, output a probability distribution over the family of classifiers. Their performance is measured by the expected success of prediction where expectation is taken with respect to both sampled data and sampled classifier.

A PAC-Bayes generalization bound relates the generalization error of the algorithm to a KL distance between the stochastic output classifier and some *prior distribution* P. In more detail, the generalization bound is comprised of two terms: first, the empirical error of the output Gibbs-classifier, and second, the KL distance between the output Gibbs classifier and some arbitrary (but sample-independent) prior distribution. This standard bound captures a basic intuition that a good learner needs to balance between bias, manifested in the form of a prior, and fitting the data, which is measured by the empirical loss. A natural task is then, to try and characterize the potential as well as limitations of such Gibbs-learners that are amenable to PAC-Bayes analysis. As far as the potential, several past results established the strength and utility of this framework (e.g. [33, 31, 19, 12, 18]).

In this work we focus on the complementary task, and present the first limitation result showing that there are classes that are learnable, even in the strong distribution-independent setting of PAC, but do not admit any algorithm that is amenable to a non-vacuous PAC-Bayes analysis. We stress that this is true even if we exploit the bound to its fullest and allow any algorithm and any possible, potentially distribution-dependent, prior.

34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada.

More concretely, we consider the class of 1-dimensional thresholds, i.e. the class of linear classifiers over the real line. It is a well known fact that this class is learnable and enjoys highly optimistic sample complexity. Perhaps surprisingly, though, we show that any Gibbs-classifier that learns the class of thresholds, must output posteriors from an unbounded set. We emphasize that the result is provided even for priors that depend on the data distribution.

From a technical perspective our proof exploits and expands a technique that was recently introduced by Alon et al. [1] to establish limitations on differentially-private PAC learning algorithms. The argument here follow similar lines, and we believe that these similarities in fact highlight a potentially powerful method to derive further limitation results, especially in the context of stability.

2 Preliminaries

2.1 Problem Setup

We consider the standard setting of binary classification. Let $\mathcal{X}$ denote the domain and $\mathcal{Y} = \{\pm 1\}$ the label space. We study learning algorithms that observe as input a sample S of labelled examples drawn independently from an unknown target distribution D, supported on $\mathcal{X} \times \mathcal{Y}$. The output of the algorithm is an hypothesis $h : \mathcal{X} \to \mathcal{Y}$, and its goal is to minimize the $0/1$-loss, which is defined by:

$$\mathcal{L}_D(h) = \mathop{\mathbb{E}}_{(x,y)\sim D}\big[\mathbf{1}[h(x) \neq y]\big].$$

We will focus on the setting where the distribution D is *realizable* with respect to a fixed hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ which is known in advance. That is, it is assumed that there exists $h \in \mathcal{H}$ such that: $\mathcal{L}_D(h) = 0$. Let $S = \langle (x_1, y_1), \ldots, (x_m, y_m) \rangle \in (\mathcal{X} \times \mathcal{Y})^m$ be a sample of labelled examples. The empirical error $\mathcal{L}_S$ with respect to S is defined by

$$\mathcal{L}_S(h) = \frac{1}{m} \sum_{i=1}^{m} \mathbf{1}[h(x) \neq y].$$

We will use the following notation: for a sample $S = \langle (x_1, y_1), \ldots (x_m, y_m) \rangle$, let $\underline{S}$ denote the underlying set of unlabeled examples $\underline{S} = \{x_i : i \leq m\}$.

The Class of Thresholds. For $k \in \mathbb{N}$ let $h_k : \mathbb{N} \to \{\pm 1\}$ denote the *threshold function*

$$h_k(x) = \begin{cases} -1 & x \leq k \\ +1 & x > k. \end{cases}$$

The class of thresholds $\mathcal{H}_{\mathbb{N}}$ is the class $\mathcal{H}_{\mathbb{N}} := \{h_k : k \in \mathbb{N}\}$ over the domain $\mathcal{X}_{\mathbb{N}} := \mathbb{N}$. Similarly, for a finite $n \in \mathbb{N}$ let $\mathcal{H}_n$ denote the class of all thresholds restricted to the domain $\mathcal{X}_n := [n] = \{1, \ldots, n\}$. Note that S is realizable with respect to $\mathcal{H}_{\mathbb{N}}$ if and only if either (i) $y_i = +1$ for all $i \leq m$, or (ii) there exists $1 \leq j \leq m$ such that $y_i = -1$ if and only if $x_i \leq x_j$.

A basic fact in statistical learning is that $\mathcal{H}_{\mathbb{N}}$ is PAC-learnable. That is, there exists an algorithm A such that for every realizable distribution D, if A is given a sample of size $O(\frac{\log 1/\delta}{\epsilon})$ examples drawn from D, then with probability at least $1 - \delta$, the output hypothesis h_S satisfies $\mathcal{L}_D(h_S) \leq \epsilon$. In fact, any algorithm A which returns an hypothesis $h_k \in \mathcal{H}_{\mathbb{N}}$ which is consistent with the input sample, will satisfy the above guarantee. Such algorithms are called empirical risk minimizers (ERMs). We stress that the above sample complexity bound is *independent* of the domain size. In particular it applies to $\mathcal{H}_n$ for every n, as well as to the infinite class $\mathcal{H}_{\mathbb{N}}$. For further reading, we refer to text books on the subject, such as [32, 24].

2.2 PAC-Bayes Bounds

PAC Bayes bounds are concerned with *stochastic-classifiers*, or *Gibbs-classifiers*. A Gibbs-classifier is defined by a distribution Q over hypotheses. The distribution Q is sometimes referred to as a *posterior*. The loss of a Gibbs-classifier with respect to a distribution D is given by the expected loss over the drawn hypothesis and test point, namely:

$$\mathcal{L}_D(Q) = \mathop{\mathbb{E}}_{h\sim Q,(x,y)\sim D}[\mathbf{1}\big[h(x) \neq y\big]].$$

A key advantage of the PAC-Bayes framework is its flexibility of deriving generalization bounds that do not depend on an hypothesis class. Instead, they provide bounds that depend on the KL distance between the output posterior and a fixed prior P. Recall that the KL divergence between a distribution P and a distribution Q is defined as follows[1]:

$$\mathrm{KL}\left(P\|Q\right) = \mathop{\mathbb{E}}_{x\sim P}\Big[\log\frac{P(x)}{Q(x)}\Big].$$

Then, the classical PAC-Bayes bound asserts the following:

Theorem 1 (PAC-Bayes Generalization Bound [23])**.** *Let D be a distribution over examples, let P be a* prior distribution *over hypothesis, and let $\delta > 0$. Denote by S a sample of size m drawn independently from D. Then, the following event occurs with probability at least $1-\delta$: for every* posterior distribution *Q,*

$$\mathcal{L}_D(Q) \leq \mathcal{L}_S(Q) + O\left(\sqrt{\frac{\mathrm{KL}\left(Q\|P\right) + \ln\sqrt{m}/\delta}{m}}\right).$$

The above bound relates the generalization error to the KL divergence between the posterior and the prior. Remarkably, the prior distribution P can be chosen as a function of the target distribution D, allowing to obtain distribution-dependent generalization bounds.

Since this pioneer work of McAllester [22], many variations on the PAC-Bayes bounds have been proposed. Notably, Seeger et al. [31] and Catoni [8] provided bounds that are known to converge at rate $1/m$ in the realizable case (see also [16] for an up-to-date survey). We note that our constructions are all provided in the realizable setting, hence readily apply.

3 Main Result

We next present the main result in this manuscript. Proofs are provided in the full version [20]. The statements use the following function $\Phi(m,\gamma,n)$, which is defined for $m,n>1$ and $\gamma\in(0,1)$:

$$\Phi(m,\gamma,n) = \frac{\log^{(m)}(n)}{(\frac{10m}{\gamma})^{3m}}.$$

Here, $\log^{(k)}(x)$ denotes the iterated logarithm, i.e.

$$\log^{(k)}(x) = \underbrace{\log(\log\ldots(\log(x)))}_{k \text{ times}}.$$

An important observation is that $\lim_{n\to\infty}\Phi(m,\gamma,n)=\infty$ for every fixed m and γ.

Theorem 2 (Main Result)**.** *Let $n,m>1$ be integers, and let $\gamma\in(0,1)$. Consider the class $\mathcal{H}_n$ of thresholds over the domain $\mathcal{X}_n=[n]$. Then, for any learning algorithm A which is defined on samples of size m, there exists a realizable distribution $D=D_A$ such that for any prior P the following event occurs with probability at least $1/16$ over the input sample $S\sim D^m$,*

$$\mathrm{KL}\left(Q_S\|P\right) = \tilde{\Omega}\left(\frac{\gamma^2}{m^2}\log\Big(\frac{\Phi(m,\gamma,n)}{m}\Big)\right) \quad \text{or} \quad \mathcal{L}_D(Q_S) > 1/2-\gamma-\frac{m}{\Phi(m,\gamma,n)},$$

where Q_S denotes the posterior outputted by A.

To demonstrate how this result implies a limitation of the PAC-Bayes framework, pick $\gamma=1/4$ and consider any algorithm A which learns thresholds over the natural numbers $\mathcal{X}_{\mathbb{N}}=\mathbb{N}$ with confidence $1-\delta\geq 99/100$, error $\epsilon<1/2-\gamma=1/4$, and m examples[2]. Since $\Phi(m,1/4,n)$ tends to infinity with n for any fixed m, the above result implies the existence of a realizable distribution D_n supported on $X_n\subseteq\mathbb{N}$ such that the PAC-Bayes bound with respect to any possible prior P will produce vacuous bounds. We summarize it in the following corollary.

[1] We use here the standard convention that if $P(\{x : Q(x)=0\})>0$ then $\mathrm{KL}\left(P\|Q\right)=\infty$.

[2] We note in passing that any Empirical Risk Minimizer learns thresholds with these parameters using <50 examples.

Corollary 1 (PAC-learnability of Linear classifiers cannot be explained by PAC-Bayes)**.** *Let $\mathcal{H}_{\mathbb{N}}$ denote the class of thresholds over $\mathcal{X}_{\mathbb{N}} = \mathbb{N}$ and let $m > 0$. Then, for every algorithm A that maps inputs sample S of size m to output posteriors Q_S and for every arbitrarily large $N > 0$ there exists a realizable distribution D such that, for any prior P, with probability at least $1/16$ over $S \sim D^m$ on of the following holds:*

$$\mathrm{KL}\left(Q_S \| P\right) > N \qquad \text{or,} \qquad \mathcal{L}_D(Q_S) > 1/4.$$

A different interpretation of Theorem 2 is that in order to derive meaningful PAC-Bayes generalization bounds for PAC-learning thresholds over a finite domain X_n, the sample complexity must grow to infinity with the domain size n (it is at least $\Omega(\log^\star(n))$). In contrast, the true sample complexity of this problem is $O(\log(1/\delta)/\epsilon)$ which is independent of n.

4 Technical Overview

A common approach of proving impossibility results in computer science (and in machine learning in particular) exploits a Minmax principle, whereby one specifies a fixed hard distribution over inputs, and establishes the desired impossibility result for any algorithm with respect to random inputs from that distribution. As an example, consider the "No-Free-Lunch Theorem" which establishes that the VC dimension lower bounds the sample complexity of PAC-learning a class $\mathcal{H}$. Here, one fixes the distribution to be uniform over a shattered set of size $d = \mathsf{VC}(H)$, and argues that every learning algorithm must observe $\Omega(d)$ examples. (See e.g. Theorem 5.1 in [32].)

Such "Minmax" proofs establish a stronger assertion: they apply even to algorithms that "know" the input-distribution. For example, the No-Free-Lunch Theorem applies even to learning algorithms that are designed given the knowledge that the marginal distribution is uniform over some shattered set.

Interestingly, such an approach is bound to fail in proving Theorem 2. The reason is that if the marginal distribution $D_{\mathcal{X}}$ over $\mathcal{X}_n$ is fixed, then one can pick an $\epsilon/2$-cover[3] $\mathcal{C}_n \subseteq \mathcal{H}_n$ of size $|\mathcal{C}_n| = O(1/\epsilon)$, and use any Empirical Risk Minimizer for $\mathcal{C}_n$. Then, by picking the prior distribution P to be uniform over $\mathcal{C}_n$, one obtains a PAC-Bayes bound which scales with the entropy $H(P) = \log|\mathcal{C}_n| = O(\log(1/\epsilon))$, and yields a $\mathsf{poly}(1/\epsilon, \log(1/\delta))$ generalization bound, which is independent of n. In other words, in the context of Theorem 2, there is no single distribution which is "hard" for all algorithms.

Thus, to overcome this difficulty one must come up with a "method" which assigns to any given algorithm A a "hard" distribution $D = D_A$, which witnesses Theorem 2 with respect to A. The challenge is that A is an arbitrary algorithm; e.g. it may be improper[4] or add different sorts of noise to its output classifier.

The method we use in the proof of Theorem 2 exploits Ramsey Theory. In a nutshell, Ramsey Theory provides powerful tools which allow to detect, for any learning algorithm, a large *homogeneous* set such that the behavior of A on inputs from the homogeneous set is highly regular. Then, we consider the uniform distribution over the homogeneous set to establish Theorem 2.

We note that similar applications of Ramsey Theory in proving lower bounds in computer science date back to the 80's [25]. For more recent usages see e.g. [7, 10, 9, 1]. Our proof closely follows the argument of Alon et al. [1], which establishes an impossibility result for learning $\mathcal{H}_n$ by differentially-private algorithms.

Technical Comparison with the Work by Alon et al. [1]. For readers who are familiar with the work of [1], let us summarize the main differences between the two proofs. The main challenge in extending the technique from [1] to prove Theorem 2 is that PAC-Bayes bounds are only required to hold for *typical samples*. This is unlike the notion of differential-privacy (which was the focus of [1]) that is defined with respect to *all samples*. Thus, establishing a lower bound in the context of differential privacy is easier: one only needs to demonstrate a single sample for which privacy is breached. However, to prove Theorem 2 one has to demonstrate that the lower bound applies to many samples. Concretely, this affects the following parts of the proof:

[3] I.e. $\mathcal{C}_n$ satisfies that $(\forall h \in \mathcal{H}_n)(\exists c \in \mathcal{C}_n) : \Pr_{x \sim D_X}\left(c(x) \neq h(x)\right) \leq \epsilon/2$.

[4] I.e. A may output hypotheses which are not thresholds, or Gibbs-classifiers supported on hypotheses which are not thresholds.

(i) The Ramsey argument in the current manuscript (Lemma 1) is more complex: to overcome the above difficulty we needed to modify the coloring and the overall construction is more convoluted.

(ii) Once Ramsey Theorem is applied and the homogeneous subset $R_n \subseteq X_n$ is derived, one still needs to derive a lower bound on the PAC-Bayes quantity. This requires a technical argument (Lemma 2), which is tailored to the definition of PAC-Bayes. Again, this lemma is more complicated than the corresponding lemma in [1].

(iii) Even with Lemma 1 and Lemma 2 in hand, the remaining derivation of Theorem 2 still requires a careful analysis which involves defining several "bad" events and bounding their probabilities. Again, this is all a consequence of that the PAC-Bayes quantity is an "average-case" complexity measure.

4.1 Proof Sketch and Key Definitions

The proof of Theorem 2 consists of two steps: (i) detecting a hard distribution $D = D_A$ which witnesses Theorem 2 with respect to the assumed algorithm A, and (ii) establishing the conclusion of Theorem 2 given the hard distribution D. The first part is combinatorial (exploits Ramsey Theory), and the second part is more information-theoretic. For the purpose of exposition, we focus in this technical overview, on a specific algorithm $\mathcal{A}$. This will make the introduction of the key definitions and presentation of the main technical tools more accessible.

The algorithm $\mathcal{A}$. Let $S = \langle (x_1, y_1), \ldots, (x_m, y_m) \rangle$ be an input sample. The algorithm $\mathcal{A}$ outputs the posterior distribution Q_S which is defined as follows: let $h_{x_i} = \mathbf{1}[x > x_i] - \mathbf{1}[x \leq x_i]$ denote the threshold corresponding to the i'th input example. The posterior Q_S is supported on $\{h_{x_i}\}_{i=1}^m$, and to each h_{x_i} it assigns a probability according to a decreasing function of its empirical risk. (So, hypotheses with lower risk are more probable.) The specific choice of the decreasing function does not matter, but for concreteness let us pick the function $\exp(-x)$. Thus,

$$Q_S(h_{x_i}) \propto \exp\bigl(-\mathcal{L}_S(h_{x_i})\bigr). \tag{1}$$

While one can directly prove that the above algorithm does not admit a PAC-Bayes analysis, we provide here an argument which follows the lines of the general case. We start by explaining the key property of *Homogeneity*, which allows to detect the hard distribution.

4.1.1 Detecting a Hard Distribution: Homogeneity

The first step in the proof of Theorem 2 takes the given algorithm and identifies a large subset of the domain on which its behavior is *Homogeneous*. In particular, we will soon see that the algorithm $\mathcal{A}$ is *Homogeneous* on the entire domain $\mathcal{X}_n$. In order to define Homogeneity, we use the following equivalence relation between samples:

Definition 1 (Equivalent Samples)**.** *Let $S = \langle (x_1, y_1), \ldots, (x_m, y_m) \rangle$ and $S' = \langle (x'_1, y'_1), \ldots, (x'_m, y'_m) \rangle$ be two samples. We say that S and S' are equivalent if for all $i, j \leq m$ the following holds.*

1. *$x_i \leq x_j \iff x'_i \leq x'_j$, and*
2. *$y_i = y'_i$.*

For example, $\langle (1,-), (5,+), (8,+) \rangle$ and $\langle (10,-), (70,+), (100,+) \rangle$ are equivalent, but $\langle (3,-), (6,+), (4,+) \rangle$ is not equivalent to them (because of Item 1). For a point $x \in \mathcal{X}_n$ let $\mathsf{pos}(x; S)$ denote the number of examples in S that are less than or equal to x:

$$\mathsf{pos}(x; S) = \Bigl| \{ x_i \in \underline{S} : x_i \leq x \} \Bigr|. \tag{2}$$

For a sample $S = \langle (x_1, y_1), \ldots, (x_m, y_m) \rangle$ let $\pi(S)$ denote the *order-type* of S:

$$\pi(S) = (\mathsf{pos}(x_1; S), \mathsf{pos}(x_2; S), \ldots, \mathsf{pos}(x_m; S)). \tag{3}$$

So, the samples $\langle (1,-), (5,+), (8,+) \rangle$ and $\langle (10,-), (70,+), (100,+) \rangle$ have order-type $\pi = (1,2,3)$, whereas $\langle (3,-), (6,+), (4,+) \rangle$ has order-type $\pi = (1,3,2)$.

Note that S, S' are equivalent if and only if they have the same labels-vectors and the same order-type. Thus, we encode the equivalence class of a sample by the pair $(\pi, \bar{y})$, where π denotes its order-type and $\bar{y} = (y_1 \dots y_m)$ denotes its labels-vector. The pair (π, y) is called the *equivalence-type of* S.

We claim that $\mathcal{A}$ satisfies the following property of *Homogeneity*:

Property 1 (Homogeneity)**.** *The algorithm $\mathcal{A}$ possesses the following property: for every two equivalent samples S, S' and every $x, x' \in \mathcal{X}_n$ such that* $\mathsf{pos}(x, S) = \mathsf{pos}(x', S')$*,*

$$\Pr_{h \sim Q_S}[h(x) = 1] = \Pr_{h' \sim Q_{S'}}[h'(x') = 1],$$

where $Q_S, Q_{S'}$ denote the Gibbs-classifier outputted by $\mathcal{A}$ on the samples S, S'.

In short, Homogeneity means that the probability $h \sim Q_S$ satisfies $h(x) = 1$ depends only on $\mathsf{pos}(x, S)$ and on the equivalence-type of S. To see that $\mathcal{A}$ is indeed homogeneous, let S, S' be equivalent samples and let $Q_S, Q_{S'}$ denote the corresponding Gibbs-classifiers outputted by $\mathcal{A}$. Then, for every x, x' such that $\mathsf{pos}(x, S) = \mathsf{pos}(x', S')$, Equation (1) yields that:

$$\Pr_{h \sim Q_S}\big[h(x) = +1\big] = \sum_{x_i < x} Q_S(h_{x_i}) = \sum_{x'_i < x'} Q_{S'}(h_{x'_i}) = \Pr_{h' \sim Q_{S'}}\big[h'(x') = +1\big],$$

where in the second transition we used that $Q_S(h_{x_i}) = Q_{S'}(h_{x'_i})$ for every $i \leq m$ (because S, S' are equivalent), and that $x_i \leq x \iff x'_i \leq x'$, for every i (because $\mathsf{pos}(x, S) = \mathsf{pos}(x', S')$).

The General Case: Approximate Homogeneity. Before we continue to define the hard distribution for algorithm A, let us discuss how the proof of Theorem 2 handles arbitrary algorithms that are not necessarily homogeneous.

The general case complicates the argument in two ways. First, the notion of Homogeneity is relaxed to an approximate variant which is defined next. Here, an order type π is called a *permutation* if $\pi(i) \neq \pi(j)$ for every distinct $i, j \leq m$. (Indeed, in this case $\pi = (\pi(x_1) \dots \pi(x_m))$ is a permutation of $1 \dots m$.) Note that the order type of $S = \langle (x_1, y_1) \dots (x_m, y_m)) \rangle$ is a permutation if and only if all the points in S are distinct (i.e. $x_i \neq x_j$ for all $i \neq j$).

Definition 2 (Approximate Homogeneity)**.** *An algorithm $\mathcal{B}$ is γ-approximately m-homogeneous if the following holds: let S, S' be two equivalent samples of length m whose order-type is a permutation, and let $x \notin \underline{S}, x' \notin \underline{S'}$ such that* $\mathsf{pos}(x, S) = \mathsf{pos}(x', S')$*. Then,*

$$|Q_S(x) - Q_{S'}(x')| \leq \frac{\gamma}{5m}, \tag{4}$$

where $Q_S, Q_{S'}$ denote the Gibbs-classifier outputted by $\mathcal{B}$ on the samples S, S'.

Second, we need to identify a sufficiently large subdomain on which the assumed algorithm is approximately homogeneous. This is achieved by the next lemma, which is based on a Ramsey argument.

Lemma 1 (Large Approximately Homogeneous Sets)**.** *Let $m, n > 1$ and let $\mathcal{B}$ be an algorithm that is defined over input samples of size m over $\mathcal{X}_n$. Then, there is $\mathcal{X}' \subseteq \mathcal{X}_n$ of size $|\mathcal{X}'| \geq \Phi(m, \gamma, n)$ such that the restriction of $\mathcal{B}$ to input samples from $\mathcal{X}'$ is γ-approximate m-homogeneous.*

We prove Lemma 1 in the full version [20]. For the rest of this exposition we rely on Property 1 as it simplifies the presentation of the main ideas.

The Hard Distribution D. We are now ready to finish the first step and define the "hard" distribution D. Define D to be uniform over examples (x, y) such that $y = h_{n/2}(x)$. So, each drawn example (x, y) satisfies that x is uniform in $\mathcal{X}_n$ and $y = -1$ if and only if $x \leq n/2$. In the general case, D will be defined in the same way with respect to the detected homogeneous subdomain.

4.1.2 Hard Distribution $\implies$ Lower Bound: Sensitivity

We next outline the second step of the proof, which establishes Theorem 2 using the hard distribution D. Specifically, we show that for a sample $S \sim D^m$,

$$\mathrm{KL}\left(Q_S \| P\right) = \tilde{\Omega}\left(\frac{1}{m^2} \log(|\mathcal{X}_n|)\right),$$

with a constant probability bounded away from zero. (In the general case $|X_n|$ is replaced by $\Phi(m, \gamma, n)$ – the size of the homogeneous set.)

Sensitive Indices. We begin with describing the key property of homogeneous learners. Let $(\pi, \bar{y})$ denote the equivalence-type of the input sample S. By homogeneity (Property 1), there is a list of numbers $p_0, \ldots, p_m$, which depends only on the order-type $(\pi, \bar{y})$, such that $\Pr_{h\sim Q_S}[h(x) = 1] = p_i$ for every $x \in \mathcal{X}_n$, where $i = \mathsf{pos}(x, S)$. The crucial observation is that there exists an index $i \leq m'$ which is *sensitive* in the sense that

$$p_i - p_{i-1} \geq \frac{1}{m}. \tag{5}$$

Indeed, consider x_j such that $h_{x_j} = \arg\min_k \mathcal{L}_S(h_{x_k})$, and let $i = \mathsf{pos}(x_j, S)$. Then,

$$p_i - p_{i-1} = \frac{\mathcal{L}_S(h_{x_j})}{\sum_{i' \leq m} \mathcal{L}_S(h_{x_{i'}})} \geq \frac{1}{m}.$$

In the general case we show that any homogeneous algorithm that learns $\mathcal{H}_n$ satisfies Equation (5) for *typical* samples (see the full version [20]). The intuition is that any algorithm that learns the distribution D must output a Gibbs-classifier Q_S such that for typical points x, if $x > n/2$ then $\Pr_{h\sim Q_S}[h(x) = 1] \approx 1$, and if $x \leq n/2$ then $\Pr_{h\sim Q_S}[h(x) = 1] \approx 0$. Thus, when traversing all x's from 1 up to n there must be a jump between p_{i-1} and p_i for some i.

From Sensitive Indices to a Lower Bound on the KL-divergence. How do sensitive indices imply a lower bound on PAC-Bayes? This is the most technical part of the proof. The crux of it is a connection between sensitivity and the KL-divergence which we discuss next. Consider a sensitive index i and let x_j be the input example such that $\mathsf{pos}(x_j, S) = i$. For $\hat{x} \in \mathcal{X}_n$, let $S_{\hat{x}}$ denote the sample obtained by replacing x_j with $\hat{x}$:

$$S_{\hat{x}} = \langle (x_1, y_1), \ldots, (x_{j-1}, y_{j-1}), (\hat{x}_j, y_j), (x_{j+1}, y_{j+1}) \ldots (x_m, y_m).\rangle,$$

and let $Q_{\hat{x}} := Q_{S_{\hat{x}}}$ denote the posterior outputted by $\mathcal{A}$ given the sample $S_{\hat{x}}$. Consider the set $I \subseteq \mathcal{X}_n$ of all points $\hat{x}$ such that $S_{\hat{x}}$ is equivalent to S. Equation (5) implies that that for every $x, \hat{x} \in I$,

$$\Pr_{h\sim Q_{\hat{x}}}[h(x) = 1] = \begin{cases} p_{i-1} & x < \hat{x}, \\ p_i & x > \hat{x}. \end{cases}$$

Combined with the fact that $p_i - p_{i-1} \geq 1/m$, this implies a lower bound on KL-divergence between an arbitrary prior P and $Q_{\hat{x}}$ for most $\hat{x} \in I$. This is summarized in the following lemma:

Lemma 2 (Sensitivity Lemma)**.** *Let I be a linearly ordered set and let $\{Q_{\hat{x}}\}_{\hat{x}\in I}$ be a family of posteriors supported on $\{\pm 1\}^I$. Suppose there are $q_1 < q_2 \in [0, 1]$ such that for every $x, \hat{x} \in I$:*

$$x < \hat{x} \implies \Pr_{h\sim Q_{\hat{x}}}[h(x) = 1] \leq q_1 + \frac{q_2 - q_1}{4},$$

$$x > \hat{x} \implies \Pr_{h\sim Q_{\hat{x}}}[h(x) = 1] \geq q_2 - \frac{q_2 - q_1}{4}.$$

Then, for every prior distribution P, if $\hat{x} \in I$ is drawn uniformly at random, then the following event occurs with probability at least $1/4$:

$$\mathrm{KL}\left(Q_{\hat{x}} \| P\right) = \Omega\Big((q_2 - q_1)^2 \frac{\log|I|}{\log\log|I|}\Big).$$

The sensitivity lemma tells us that in the above situation, the KL divergence between $Q_{\hat{x}}$ and any prior P, for a random choice $\hat{x}$, scales in terms of two quantities: the distance between the two values, $q_2 - q_1$, and the size of I.

The proof of Lemma 2 is provided in the full version [20]. In a nutshell, the strategy is to bound from below $\mathrm{KL}\left(Q^r_{\hat{x}} \| P^r\right)$, where r is sufficiently small; the desired lower bound then follows from the chain rule, $\mathrm{KL}\left(Q_{\hat{x}} \| P\right) = \frac{1}{r}\mathrm{KL}\left(Q^r_{\hat{x}} \| P^r\right)$. Obtaining the lower bound with respect to the r-fold products is the crux of the proof. In short, we will exhibit events $E_{\hat{x}}$ such that $Q^r_{\hat{x}}(E_{\hat{x}}) \geq \frac{1}{2}$ for every $\hat{x} \in I$, but $P^r(E_{\hat{x}})$ is tiny for $\frac{|I|}{4}$ of the $\hat{x}$'s. This implies a lower bound on $\mathrm{KL}\left(Q^r_{\hat{x}} \| P^r\right)$ since

$$\mathrm{KL}\left(Q^r_{\hat{x}} \| P^r\right) \geq \mathrm{KL}\left(Q^r_{\hat{x}}(E_{\hat{x}}) \| P^r(E_{\hat{x}})\right),$$

by the data-processing inequality.

Wrapping Up. We now continue in deriving a lower bound for $\mathcal{A}$. Consider an input sample $S \sim D^m$. In order to apply Lemma 2, fix any equivalence-type (π, y) with a sensitive index i and let x_j be such that $\mathsf{pos}(x_j; S) = i$. The key step is to condition the random sample S on (π, y) as well as on $\{x_t\}_{t=1}^m \setminus \{x_j\}$ – all sample points besides the sensitive point x_j. Thus, only x_j is remained to be drawn in order to fully specify S. Note then, that by symmetry $\hat{x}$ is uniformly distributed in a set $I \subseteq \mathcal{X}_n$, and plugging $q_1 := p_i, q_2 := p_{i-1}$ in Lemma 2 yields that for any prior distribution P:

$$\mathrm{KL}\left(Q_S \| P\right) \geq \tilde{\Omega}\left(\frac{1}{m^2}\log(|I|)\right),$$

with probability at least $1/4$. Note that we are not quite done since the size $|I|$ is a random variable which depends on the type $(\pi, \bar{y})$ and the sample points $\{x_k\}_{k \neq j}$. However, the distribution of $|I|$ can be analyzed by elementary tools. In particular, we show that $|I| \geq \Omega(|\mathcal{X}_n|/m^2)$ with high enough probability, which yields the desired lower bound on the PAC-Bayes quantity. (In the general case $|\mathcal{X}_n|$ is replaced by the size of the homogeneous set.)

5 Discussion

In this work we presented a limitation for the PAC-Bayes framework by showing that PAC-learnability of one-dimensional thresholds can not be established using PAC-Bayes.

Perhaps the biggest caveat of our result is the mild dependence of the bound on the size of the domain in Theorem 2. In fact, Theorem 2 does not exclude the possibility of PAC-learning thresholds over $\mathcal{X}_n$ with sample complexity that scale with $O(\log^* n)$ such that the PAC-Bayes bound vanishes. It would be interesting to explore this possibility; one promising direction is to borrow ideas from the differential privacy literature: [3] and [5] designed a private learning algorithm for thresholds with sample complexity $\exp(\log^* n)$; this bound was later improved by [17] to $\tilde{O}((\log^* n)^2)$. Also, [6] showed that finite Littlestone dimension is sufficient for private learnability, and it would be interesting to extend these results to the context of PAC-Bayes. Let us note that in the context of *pure* differential privacy, the connection between PAC-Bayes analysis and privacy has been established in [13].

Non-uniform learning bounds Another aspect is the implication of our work to learning algorithms beyond the uniform PAC setting. Indeed, many successful and practical algorithms exhibit sample complexity that depends on the target-distribution. E.g.,the k-Nearest-Neighbor algorithm eventually learns any target-distribution (with a distribution-dependent rate). The first point we address in this context concerns *interpolating algorithms*. These are learners that achieve zero (or close to zero) training error (i.e. they interpolate the training set). Examples of such algorithms include kernel machines, boosting, random forests, as well as deep neural networks [4, 29]. PAC-Bayes analysis has been utilized in this context, for example, to provide margin-dependent generalization guarantees for kernel machines [19]. It is therefore natural to ask whether our lower bound has implications in this context. As a simple case-study, consider the 1-Nearest-Neighbour. Observe that this algorithm forms a proper and consistent learner for the class of 1-dimensional thresholds[5], and therefore enjoys a very fast learning rate. On the other hand, our result implies that for any algorithm (including as 1-Nearest-Neighbor) that is amenable to PAC-Bayes analysis, there is a distribution realizable by thresholds on which it has high population error. Thus, no algorithm with a PAC-Bayes generalization bound can match the performance of nearest-neighbour with respect to such distributions.

Finally, this work also relates to a recent attempt to explain generalization through the implicit bias of learning algorithms: it is commonly argued that the generalization performance of algorithms can be explained by an implicit algorithmic bias. Building upon the flexibility of providing distribution-dependent generalization bounds, the PAC-Bayes framework has seen a resurgence of interest in this context towards explaining generalization in large-scale modern-time practical algorithms [27, 28, 12, 13, 2]. Indeed PAC-Bayes bounds seem to provide non-vacuous bounds in several relevant domains [18, 13]. Nevertheless, the work here shows that any algorithm that can learn 1D thresholds is necessarily not biased, in the PAC-Bayes sense, towards a (possibly distribution-dependent) prior. We mention that recently, [11] showed that SGD's generalization performance indeed cannot be attributed to some implicit bias of the algorithm that governs the generalization.

[5] Indeed, given any realizable sample it will output the threshold which maximizes the margin.

Broader Impact

There are no foreseen ethical or societal consequences for the research presented herein.

Acknowledgments and Disclosure of Funding

R.L is supported by an ISF grant no. 2188/20 and partially funded by an unrestricted gift from Google. Any opinions, findings, and conclusions or recommendations expressed in this work are those of the author(s) and do not necessarily reflect the views of Google. S.M is supported by the Israel Science Foundation (grant No. 1225/20), by an Azrieli Faculty Fellowship, and by a grant from the United States - Israel Binational Science Foundation (BSF). Part of this work was done while the author was at Google Research.

References

[1] N. Alon, R. Livni, M. Malliaris, and S. Moran. Private pac learning implies finite littlestone dimension. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 852–860, 2019.

[2] S. Arora, R. Ge, B. Neyshabur, and Y. Zhang. Stronger generalization bounds for deep nets via a compression approach. volume 80 of *Proceedings of Machine Learning Research*, pages 254–263. PMLR, 10–15 Jul 2018. URL `http://proceedings.mlr.press/v80/arora18b.html`.

[3] A. Beimel, K. Nissim, and U. Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. *Theory of Computing*, 12(1):1–61, 2016.

[4] M. Belkin, D. J. Hsu, and P. Mitra. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In *Advances in neural information processing systems*, pages 2300–2311, 2018.

[5] M. Bun, K. Nissim, U. Stemmer, and S. Vadhan. Differentially private release and learning of threshold functions. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 634–649. IEEE, 2015.

[6] M. Bun, R. Livni, and S. Moran. An equivalence between private classification and online prediction. *arXiv preprint arXiv:2003.00563*, 2020.

[7] M. M. Bun. *New Separations in the Complexity of Differential Privacy*. PhD thesis, Harvard University, Graduate School of Arts & Sciences, 2016.

[8] O. Catoni. Pac-bayesian supervised classification: The thermodynamics of statistical learning. *stat*, 1050:3, 2007.

[9] A. Cohen, A. Hassidim, H. Kaplan, Y. Mansour, and S. Moran. Learning to screen. In *Advances in Neural Information Processing Systems 32*, 2019. URL `http://papers.nips.cc/paper/9067-learning-to-screen`.

[10] J. R. Correa, P. Dütting, F. A. Fischer, and K. Schewior. Prophet inequalities for I.I.D. random variables from an unknown distribution. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, pages 3–17. ACM, 2019. URL `https://doi.org/10.1145/3328526.3329627`.

[11] A. Dauber, M. Feder, T. Koren, and R. Livni. Can implicit bias explain generalization? stochastic convex optimization as a case study. *arXiv preprint arXiv:2003.06152*, 2020.

[12] G. K. Dziugaite and D. M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI*. AUAI Press, 2017. URL `http://auai.org/uai2017/proceedings/papers/173.pdf`.

[13] G. K. Dziugaite and D. M. Roy. Data-dependent pac-bayes priors via differential privacy. In *Advances in Neural Information Processing Systems*, pages 8430–8441, 2018.

[14] P. Erdos and R. Rado. Combinatorial theorems on classifications of subsets of a given set. *Proceedings of the London mathematical Society*, 3(1):417–439, 1952.

[15] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey theory*, volume 20. John Wiley & Sons, 1990.

[16] B. Guedj and J. Shawe-Taylor. A primer on pac-bayesian learning. In *ICML 2019-Thirty-sixth International Conference on Machine Learning*, 2019.

[17] H. Kaplan, K. Ligett, Y. Mansour, M. Naor, and U. Stemmer. Privately learning thresholds: Closing the exponential gap. In *Conference on Learning Theory, COLT 2020, 9-12 July 2020*, volume 125 of *Proceedings of Machine Learning Research*, pages 2263–2285. PMLR, 2020. URL `http://proceedings.mlr.press/v125/kaplan20a.html`.

[18] J. Langford and R. Caruana. (not) bounding the true error. In *Advances in Neural Information Processing Systems*, pages 809–816, 2002.

[19] J. Langford and J. Shawe-Taylor. Pac-bayes & margins. In *Advances in neural information processing systems*, pages 439–446, 2003.

[20] R. Livni and S. Moran. A limitation of the pac-bayes framework. *CoRR*, abs/2006.13508, 2020. URL `https://arxiv.org/abs/2006.13508`.

[21] D. McAllester. Simplified pac-bayesian margin bounds. In *Learning theory and Kernel machines*, pages 203–215. Springer, 2003.

[22] D. A. McAllester. Pac-bayesian model averaging. In *Proceedings of the twelfth annual conference on Computational learning theory*, pages 164–170, 1999.

[23] D. A. McAllester. Some pac-bayesian theorems. *Machine Learning*, 37(3):355–363, 1999.

[24] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. MIT press, 2018.

[25] S. Moran, M. Snir, and U. Manber. Applications of ramsey's theorem to decision tree complexity. *Journal of the ACM (JACM)*, 32(4):938–949, 1985.

[26] D. Mubayi and A. Suk. A survey of hypergraph ramsey problems. *arXiv preprint arXiv:1707.04229*, 2017.

[27] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 5947–5956, 2017.

[28] B. Neyshabur, S. Bhojanapalli, and N. Srebro. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. In *International Conference on Learning Representations*, 2018.

[29] R. Salakhotdinov. Deep learning tutorial at the simons institute, berkeley. 2017. URL `https://simons.berkeley.edu/talks/ruslan-salakhutdinov-01-26-2017-1`.

[30] M. Seeger. Pac-bayesian generalisation error bounds for gaussian process classification. *Journal of machine learning research*, 3(Oct):233–269, 2002.

[31] M. Seeger, J. Langford, and N. Megiddo. An improved predictive accuracy bound for averaging classifiers. In *Proceedings of the 18th International Conference on Machine Learning*, number CONF, pages 290–297, 2001.

[32] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

[33] J. Shawe-Taylor and D. Hardoon. Pac-bayes analysis of maximum entropy classification. In *Artificial Intelligence and Statistics*, pages 480–487, 2009.

[34] J. Shawe-Taylor and R. C. Williamson. A pac analysis of a bayesian estimator. In *Proceedings of the tenth annual conference on Computational learning theory*, pages 2–9, 1997.

A Proofs

A.1 Proof of Theorem 2

Let A be an algorithm as in the premise of Theorem 2. That is, A receives as input a labeled sample S of length m and outputs a posterior Q_S. By Lemma 1, there exists $\mathcal{X}' \subseteq \mathcal{X}_n$ of size

$|\mathcal{X}'| = k \geq \Phi(m, \gamma, n)$ such that the restriction of A to inputs from $\mathcal{X}'$ is γ-approximate m-homogeneous. Without loss of generality, assume that $\mathcal{X}' = \mathcal{X}_k$ consists of the first k points in $\mathcal{X}_n$ and that k is an even number.

By the definition of approximate homogeneity (Definition 2) it follows that for every equivalence type $(\pi, \bar{y})$, where π is a permutation, there is a list $(p_i^{(\pi,\bar{y})})_{i=0}^m \in [0,1]^{m+1}$ such that for every sample $S \in (\mathcal{X}_k \times \{0,1\})^m$ whose type is $(\pi, \bar{y})$ and and every $x \in \mathcal{X}_k \setminus S$:

$$\left|Q_S(x) - p_i^{(\pi,\bar{y})}\right| \leq \frac{\frac{\gamma}{5m}}{2} = \frac{\gamma}{10m},$$

where $\mathsf{pos}(x, S) = i$. For the rest of the proof fix D to be the distribution over examples (x, y) such that x is drawn uniformly from $\mathcal{X}_k$ and $y = -1$ if and only if $x \leq k/2$. The underlying property we will require is summarized in the following claim:

Claim 1. *Let $(\pi, \bar{y})$ be an equivalence-type, where π is a permutation. Then, one of the following holds: either there exists a sensitive index $0 \leq i \leq m$ such that*

$$|p_i^{(\pi,\bar{y})} - p_{i-1}^{(\pi,\bar{y})}| \geq \frac{\gamma}{2m}, \tag{6}$$

or else,

$$\mathcal{L}_D(Q_S) > \frac{1}{2} - \gamma - \frac{m}{k}$$

with probability 1 over $S \sim D^m(\cdot|(\pi, \bar{y}))$.

The proof of Claim 1 is deterred to Appendix A.1.1.

With Claim 1 in hand we proceed with the proof of Theorem 2. Let S be a sample and let $(\pi, \bar{y})$ denote its equivalence-type. Define an interval $I(S) \subseteq \mathcal{X}_k$ as follows.

- if π is not a permutation then $I(S) = \emptyset$.
- If $(\pi, \bar{y})$ does not have a sensitive index that satisfies Equation (6) then $I(S) = \emptyset$.
- Finally, if π is a permutation and $(\pi, \bar{y})$ has a sensitive index i then set[6]

$$I(S) = \begin{cases} (x_j^-, x_j^+) & \frac{k}{2} \notin (x_j^-, x_j^+), \\ (x_j^-, \frac{k}{2}] & \frac{k}{2} \in (x_j^-, x_j^+) \text{ and } y_j = -1, \\ (\frac{k}{2}, x_j^+) & \frac{k}{2} \in (x_j^-, x_j^+) \text{ and } y_j = +1, \end{cases}$$

  where x_j is such that $\mathsf{pos}(x_j; S) = i$, and $x_j^- = \max(\{x_t : x_t < x_j\} \cup \{0\})$ and $x_j^+ = \min(\{x_t : x_t > x_j\} \cup \{k+1\})$.

We next define two events which will be used to finish the proof. First, consider the event that the drawn sample S satisfies either[7]

$$\mathrm{KL}\left(Q_S \| P\right) = \Omega\Big(\frac{\gamma^2}{m^2} \frac{\log|I(S)|}{\log\log|I(S)|}\Big), \tag{7}$$

or

$$\mathcal{L}(Q_S) \geq \frac{1}{2} - \gamma - \frac{m}{k}, \tag{8}$$

We show that this event occurs with probability at least $1/4$:

Claim 2. *Define E_1 to be the event*

$$E_1 = \Big\{S \in (\mathcal{X}_k \times \{\pm 1\})^m : S \textit{ satisfies Equation (7) or Equation (8)}\Big\}.$$

Then, E_1 occurs with probability at least $1/4$ over $S \sim D^m$.

[6] For concreteness, let i be the minimal sensitive index.

[7] We use here the convention, that $\frac{\log x}{\log\log x} = -\infty$ for $x \leq 2$. Alternatively, one can assume that Equation (7) holds vacuously if $|I(S)| = 0$

The proof of Claim 2 is deterred to Appendix A.1.2. The second event we consider is that the drawn sample S satisfies either Equation (8) or

$$|I(S)| \geq \frac{\Phi(m,\gamma,n)}{8(m+1)^2}. \tag{9}$$

We show that this event occurs with probability at least $7/8$:

Claim 3. *Define E_2 to be the event*

$$E_2 = \{S : S \textit{ satisfies Equation } (8) \textit{ or Equation } (9)\}.$$

Then E_2 occurs with probability at least $7/8$ over $S \sim D^m$

The proof of Claim 3 is deterred to Appendix A.1.3. With Claims 2 and 3 in hand, the proof of Theorem 2 is completed as follows. First, a union bound implies that the event $E_1 \cap E_2$ occurs with probability at least $1/16$. That is, with probability at least $1/16$ either Equation (8) holds and we are done, or else, if Equation (8) doesn't hold, then both Equations (7) and (9) hold simultaneously, which yields that

$$\begin{aligned}
\mathrm{KL}\left(Q_S \| P\right) &\geq \Omega\Big(\frac{\gamma^2}{m^2}\frac{\log|I(S)|}{\log\log|I(S)|}\Big) && \text{(By Equation (7))} \\
&\geq \Omega\Big(\frac{\gamma^2}{m^2}\frac{\log\frac{\Phi(m,\gamma,n)}{8(m+1)^2}}{\log\log\frac{\Phi(m,\gamma,n)}{8(m+1)^2}}\Big). && \text{(By Equation (9))}
\end{aligned}$$

This concludes the proof of Theorem 2.

□

We are thus left with proving Claims 1 to 3.

A.1.1 Proof of Claim 1

Let $(\pi, \bar{y})$ be an equivalence-type such that π is a permutation. Assume that

$$\mathcal{L}(Q_S) < \frac{1}{2} - \gamma - \frac{m}{k} \tag{10}$$

occurs with a positive probability over $S \sim D^m(\cdot|\pi, \bar{y})$. We first show that there is i such that

$$|p_i^{\pi,\bar{y}} - p_0^{\pi,\bar{y}}| > \gamma/2. \tag{11}$$

Indeed, assume the contrary and fix a sample S with type $(\pi, \bar{y})$ which satisfies Equation (10). Recall that A is homogeneous, hence for every $x \notin \underline{S}$,

$$|Q_S(x) - p_i^{\pi,\bar{y}}| < \frac{\gamma}{10m},$$

where $i = \mathsf{pos}(x, S)$. On the other hand, since Equation (11) is not met by any i, it follows that for every $x \notin \underline{S}$:

$$\begin{aligned}
|Q_S(x) - p_0^{\pi,\bar{y}}| &= |Q_S(x) - p_i^{\pi,\bar{y}} + p_i^{\pi,\bar{y}} - p_0^{\pi,\bar{y}}| \\
&\leq |Q_S(x) - p_i^{\pi,\bar{y}}| + |p_i^{\pi,\bar{y}} - p_0^{\pi,\bar{y}}| \\
&\leq \frac{\gamma}{10m} + \frac{\gamma}{2} \\
&\leq \gamma.
\end{aligned}$$

Thus, $\Pr_{h\sim Q_S}[h(x) = 1] \in [p_0^{\pi,\bar{y}} - \gamma, p_0^{\pi,\bar{y}} + \gamma]$, for every $x \in \mathcal{X}_k \setminus \underline{S}$. Now, since $\Pr_{(x,y)\sim D}[y = 1] = 1/2$ it follows that

$$\mathcal{L}_D(Q_S) \geq \frac{1}{2} - \gamma - \frac{m}{k}.$$

Indeed, for every $x \notin S$, if $x \leq k/2$ then $h \sim Q_S$ errs on x with probability at least $q_1 = p_0^{\pi,\bar{y}} - \gamma$, and if $x > k/2$ then $h \sim Q_S$ errs on x with probability at least $q_2 = 1 - (p_0^{\pi,\bar{y}} + \gamma)$. Thus, the

expected loss of $h \sim Q_S$ conditioned on $x \notin \underline{S}$ is at least $\frac{q_1+q_2}{2} = 1/2 - \gamma$, and the above inequality follows by taking into account that $h \sim Q_S$ may have zero error on the m points in S.

Finally, let i be some index that satisfy Equation (11), then because $0 \leq i \leq m$ we obtain via telescoping that there must be some $i' \leq i$, such that

$$|p_{i'}^{\pi,\bar{y}} - p_{i'-1}^{\pi,\bar{y}}| \geq \frac{\gamma}{2m}.$$

□

A.1.2 Proof of Claim 2

Proof of Claim 2. It is enough to show that E_1 occurs with probability at least $1/4$ over $S \sim D^m(\cdot|\pi,\bar{y})$ for every fixed equivalence-type $(\pi,\bar{y})$. Indeed, by summing over all equivalence types, the law of total probability then implies that E_1 occurs with probability at least $1/4$ over $S \sim D^m$.

Fix an equivalence-type $(\pi,\bar{y})$. We may assume that π is a permutation and that $(\pi,\bar{y})$ has a sensitive index i (or else Equation (7) trivially holds by the definition of $I(S)$ and we are done). If Equation (8) holds with probability at least $1/4$ then also E occurs with probability at least $1/4$ and we are done. Thus, assume that Equation (8) holds with probability less than $1/4$. It suffices to show that Equation (7) holds with probability at least $1/4$. By Claim 1, there is a sensitive index i such that

$$|p_i^{(\pi,\bar{y})} - p_{i-1}^{(\pi,\bar{y})}| > \frac{\gamma}{2m}.$$

Let x_j in S be such that $\mathsf{pos}(x_j;S) = i$. It will be convenient to consider the following (slightly convoluted) process of sampling a pair of (correlated) samples from $D^m(\cdot|\pi,\bar{y})$:

1. Sample $T = \langle (x_1,y_1)\ldots(x_m,y_m)\rangle \sim D^m(\cdot|\pi,\bar{y})$.
2. Resample only the sensitive point x_j while keeping all other points fixed, as well as the equivalence type $(\pi,\bar{y})$. Let $\hat{x}$ denote the newly sampled point and let $T_{\hat{x}}$ denote the sample obtained by replacing x_j by $\hat{x}$.
3. Set $S = T_{\hat{x}}$

Note that both T and S are drawn from $D^m(\cdot|\pi,\bar{y})$ and that $I(T) = I(S)$ always. Since the marginal distribution of D is uniform over $\mathcal{X}_k$, by symmetry it follows that the point $\hat{x}$ drawn in Step 2 is uniform in the interval $I(T) = I(S)$. Our next step is to apply Lemma 2 on the family of distributions $\{Q_{T_{\hat{x}}}\}_{\hat{x}\in I(T)}$. Towards this end, we first fix T and show that the premise of Lemma 2 is satisfied, with $I = I(T)$, $q_1 = p_{i-1}^{(\pi,\bar{y})}$ and $q_2 = p_i^{(\pi,\bar{y})}$.[8] Indeed, by homogeneity it follows that for each $x \in I(T)$, if $x < \hat{x}$

$$\begin{aligned}\left|\Pr_{h\sim Q_{\hat{x}}}[h(x)=1] - p_{i-1}^{\pi,\bar{y}}\right| &\leq \frac{\gamma}{10m} \\ &< \frac{|p_i^{(\pi,\bar{y})} - p_{i-1}^{(\pi,\bar{y})}|}{4}, \qquad \text{(because } i \text{ is sensitive)}\end{aligned}$$

and similarly if $x \geq \hat{x}$:

$$\left|\Pr_{h\sim Q_{\hat{x}}}[h(x)=1] - p_i^{(\pi,\bar{y})}\right| < \frac{|p_i^{(\pi,\bar{y})} - p_{i-1}^{(\pi,\bar{y})}|}{4}.$$

[8] Here we assume without loss of generality that $p_{i-1}^{(\pi,\bar{y})} < p_i^{(\pi,\bar{y})}$. If the reverse inequality holds then the argument follows by applying Lemma 2 with respect to the reverse linear order over $I(T)$.

Thus, applying Lemma 2 on the family of distributions $\{Q_{T_{\hat{x}}}\}_{\hat{x}\in I(T)}$ yields that for every T sampled in Step 1, the following holds with probability at least $1/4$ over sampling $\hat{x}$:

$$\begin{aligned}\mathrm{KL}\left(Q_S\|P\right) &= \mathrm{KL}\left(Q_{T_{\hat{x}}}\|P\right)\\ &\geq \Omega\Big(\big(p_{i-1}^{(\pi,\bar{y})}-p_i^{(\pi,\bar{y})}\big)^2\frac{\log(|I(T)|)}{\log\log|I(T)|}\Big)\\ &\geq \Omega\Big(\frac{\gamma^2}{m^2}\frac{\log|I(T)|}{\log\log|I(T)|}\Big)\\ &= \Omega\Big(\frac{\gamma^2}{m^2}\frac{\log|I(S)|}{\log\log|I(S)|}\Big).\end{aligned}$$

Note that the above holds for any fixed T. Taking expectation over T it follows that with probability at least $1/4$ over $S\sim D(\cdot|(\pi,\bar{y}))$,

$$\mathrm{KL}\left(Q_S\|P\right)\geq\Omega\Big(\frac{\gamma^2}{m^2}\frac{\log|I(S)|}{\log\log|I(S)|}\Big).$$

As discussed, taking expectation over the equivalence type concludes the proof. □

A.1.3 Proof of Claim 3

Proof of Claim 3. Consider $S\sim D^m$ where $S=\langle(x_1,y_1),\ldots(x_m,y_m)\rangle$. We claim that with probability at least $7/8$, every two unlabeled examples x_i,x_j with $i\neq j$ are at distance at least $\frac{k}{8(m+1)^2}$ from each other and from $k/2$. Indeed, fix any distinct $x',x''\in\{x_1,\ldots,x_m,k/2\}$. Recall that the distribution D satisfies that $x_1,\ldots,x_m$ are sampled uniformly and ind. from $\mathcal{X}_k$. Thus, the probability that $0\leq x'-x''<\frac{k}{8(m+1)^2}$ is at most $\frac{1}{8(m+1)^2}$. A union bound over all possible $\binom{m+1}{2}$ pairs implies that that the following holds with probability at least $\frac{7}{8}$ over $S\sim D^m$:

$$\Big(\forall\text{ distinct }x',x''\in\Big\{x_1,\ldots,x_m,\frac{k}{2}\Big\}\Big):|x'-x''|\geq\frac{k}{8(m+1)^2}. \tag{12}$$

We will now show that the latter event implies E_2. Let S be a sample satisfying Equation (12). In particular, $x_i\neq x_j$ for every distinct $i,j\leq m$ and so the order-type $\pi=\pi(S)$ is a permutation. Now, if S satisfies Equation (8) then $S\in E_2$ and we are done. Else, by Claim 1 there exists a sensitive index that satisfies Equation (6) and therefore $I(S)=(x',x'')$, where x',x'' are distinct points in $\{x_1\ldots,x_m,k/2\}$. Thus,

$$|I(S)|\geq\frac{k}{8(m+1)^2},$$

and Equation (9) holds, which also gives $S\in E_2$. Thus, every S which satisfies Equation (12) is in E_2 and so E_2 occurs with probability at least $7/8$. □

A.2 Proof of Lemma 1

We next prove Lemma 1 which establishes the existence of a "largish" homogeneous set with respect to an arbitrary algorithm A.

Notation. Recall from Equation (2) the definition of $\mathsf{pos}(x,S)$ which was defined for a sample S and a point x. It will be convenient to extend this definition to sets: for $R\subseteq\mathcal{X}_n$ and $x\in\mathcal{X}_n$ define $\mathsf{pos}(x,R)=|\{x'\in R:x'\leq x\}|$.

From Sets to Samples. Let $(\pi,\bar{y})$ be an equivalence-type whose order-type is a permutation and let $D=\{x_1<\ldots<x_m\}\subseteq\mathcal{X}_n$ be a set of m points. Denote by $D^{\pi,\bar{y}}=\langle(x_{i_j},y_{i_j})\rangle_{j=1}^m$ the sample obtained by ordering and labeling the elements of D such that $D^{\pi,\bar{y}}$ has type $(\pi,\bar{y})$; that is, $D^{\pi,\bar{y}}$ is defined such that for every $j\leq m$,

$$\pi(j)=\mathsf{pos}(x_{i_j},D^{\pi,\bar{y}})=\mathsf{pos}(x_{i_j},D)\ \text{ and }\ \bar{y}=(y_1,\ldots,y_m). \tag{13}$$

A Coloring. We define a coloring over subsets $D \subseteq \mathcal{X}_n$ of size $|D| = m+1$. Let $D = \{x_0 < x_1 < \ldots < x_m\}$ be a $(m+1)$-subset of $\mathcal{X}_n$. The coloring assigned to D is

$$c(D) = \left\{(p_0^{\pi,\bar{y}}, \ldots, p_m^{\pi,\bar{y}}) : (\pi, \bar{y}) \text{ is an equivalence-type s.t. } \pi \text{ is a permutation}\right\},$$

where each $p_i^{\pi,\bar{y}}$ is defined as follows: let $D_{-i} = D \setminus \{x_i\}$. For each equivalence type $(\pi, \bar{y})$ such that π is a permutation consider the sample $D_{-i}^{\pi,\bar{y}}$ (see Equation (13)), and define $p_i^{\pi,\bar{y}}$ to be the fraction of the form $\frac{t \cdot \gamma}{10m}$ for $t \in \mathbb{N}$ which is closest to

$$\Pr_{h \sim Q_{-i}^{\pi,\bar{y}}}[h(x_i) = 1],$$

where $Q_{-i}^{\pi,\bar{y}}$ is the stochastic classifier obtained by applying A on $D_{-i}^{\pi,\bar{y}}$.

Since the total number of equivalence-types whose order-type is a permutation is at most $m! \cdot 2^m$, it follows that the total number of colors is at most $m! \cdot 2^m \cdot \lceil \frac{10m}{\gamma} + 1 \rceil^{(m+1)} \leq (\frac{100m}{\gamma})^{2m}$.

Ramsey. We next apply Ramsey Theorem to derive a large $\mathcal{X}' \subseteq \mathcal{X}_n$ such that every subset $D \subseteq \mathcal{X}_n$ of size $m+1$ has the same color. Later we will argue that A is γ-approximately homogeneous with respect to $\mathcal{X}'$ which will finish the proof.

We will use the following quantitative version of Ramsey Theorem due to [14] (see also the book [15], or Theorem 10.1 in the survey by [26]). Here, the *tower function* $\mathbf{twr}_k(x)$ is defined by the recursion

$$\mathbf{twr}^{(i)}x = \begin{cases} x & i = 1, \\ 2^{\mathbf{twr}(i-1)(x)} & i > 1. \end{cases}$$

Theorem 3 (Ramsey Theorem [14])**.** *Let $s > t \geq 2$ and q be integers, and let*

$$N \geq \mathbf{twr}_t(3sq \log q).$$

Then, for every coloring of the subsets of size t of a universe of size N using q colors there is a homogeneous subset[9] *of size s.*

Stated differently, Theorem 3 guarantees the existence of a homogeneous subset of size

$$\frac{\log^{(t-1)}(N)}{3q \log q}. \tag{14}$$

Thus, by plugging $q := (\frac{10m}{\gamma})^{2m}, t := m+1, N := n$ in Equation (14) we get a homogeneous set $\mathcal{X}' \subseteq \mathcal{X}_n$ of size

$$|\mathcal{X}'| \geq \frac{\log^{(m)}(n)}{3(\frac{10m}{\gamma})^{2m} \cdot 2m \log(\frac{10m}{\gamma})} \geq \frac{\log^{(m)}(n)}{(\frac{10m}{\gamma})^{3m}}.$$

Wrapping-up. It remains to show that A is γ-approximately homogeneous with respect to $\mathcal{X}'$. By the construction of $\mathcal{X}'$ there exist a specific color

$$L = \{(p_i^{\pi,\bar{y}})_{i=0}^m : (\pi, \bar{y}) \text{ is an equivalence-type s.t. } \pi \text{ is a permutation}\}$$

such that $c(D) = L$ for every $D = \{x_0 < \ldots < x_m\} \subseteq \mathcal{X}'$. We need to show that for every pair of equivalent samples S', S'' whose order-type is a permutation and for every $x \in \mathcal{X}' \setminus \underline{S}, x' \in \mathcal{X}' \setminus \underline{S'}$ such that $\mathsf{pos}(x, S) = \mathsf{pos}(x', S')$:

$$\left| \Pr_{h \sim Q_S}[h(x) = 1] - \Pr_{h' \sim Q_{S'}}[h'(x') = 1] \right| \leq \frac{\gamma}{5m}.$$

Let $(\pi, \bar{y})$ be an equivalence-type such that π is a permutation, let S be any sample whose equivalence-type is $(\pi, \bar{y})$, and let $x \in \mathcal{X}' \setminus \bar{S}$. Consider the set $D = \{x_j : j \leq m\} \cup \{x\}$ and set $i = \mathsf{pos}(x, S)$. By the definition of $D_{-i}^{\pi,\bar{y}}$, we have $D_{-i}^{\pi,\bar{y}} = S$ and hence by the definition of $p_i^{\pi,\bar{y}}$ we have

$$\left| \Pr_{h \sim Q_S}[h(x) = 1] - p_i^{\pi,\bar{y}} \right| = \left| \Pr_{h \sim Q_{-i}^{\pi,\bar{y}}}[h(x) = 1] - p_i^{\pi,\bar{y}} \right| \leq \frac{\gamma}{10m}.$$

[9] A subset of the universe is homogeneous if all of its t-subsets have the same color.

Since the latter holds for every sample S whose order type is $(\pi, \bar{y})$ and every $x \notin \bar{S}$, it follows that for every pair of samples S, S' whose order-type is $(\pi, \bar{y})$ and every $x \in \mathcal{X}' \setminus \underline{S}, x' \in \mathcal{X}' \setminus \underline{S'}$ such that $\mathsf{pos}(x, S) = \mathsf{pos}(x', S')$:

$$\Big|\Pr_{h\sim Q_S}[h(x)=1] - \Pr_{h'\sim Q_{S'}}[h'(x')=1]\Big| \leq$$
$$\Big|\Pr_{h\sim Q_S}[h(x)=1] - p_i^{\pi,\bar{y}}\Big| + \Big|\Pr_{h\sim Q_S}[h(x)=1] - p_i^{\pi,\bar{y}}\Big| \leq \frac{\gamma}{10m} + \frac{\gamma}{10m} = \frac{\gamma}{5m},$$

where $i := \mathsf{pos}(x, S) = \mathsf{pos}(x', S')$. This finishes the proof. □

A.3 Proof of Lemma 2

Notation. We will assume without loss of generality that $I = \{1, 2, 3, ..., |I|\}$. Also, to simplify the presentation, we will assume that $|I|$ is a power of 2, i.e. $|I| = 2^b$ for some $b \in \mathbb{N}$. (Removing this assumption is straight-forward, but complicates some of the notation.)

Overview. Let P be an arbitrary prior supported on $\{\pm 1\}^I$. Our goal is to show that at least $|I|/4$ of all $\hat{x}$'s in I satisfy

$$\mathrm{KL}\left(Q_{\hat{x}} \| P\right) \geq \Omega\Big((q_2 - q_1)^2 \frac{\log|I|}{\log\log|I|}\Big) = \Omega\Big((q_2 - q_1)^2 \frac{b}{\log(b)}\Big).$$

The proof strategy is to bound from below $\mathrm{KL}\left(Q_{\hat{x}}^m \| P^m\right)$, where m is sufficiently small; the desired lower bound then follows from the chain rule:

$$\mathrm{KL}\left(Q_{\hat{x}} \| P\right) = \frac{1}{m}\mathrm{KL}\left(Q_{\hat{x}}^m \| P^m\right).$$

Obtaining the lower bound with respect to the m-fold products is the crux of the proof. In a nutshell, we will exhibit events $E_{\hat{x}}$ such that for every $\hat{x} \in I$, $Q_{\hat{x}}^m(E_{\hat{x}}) \geq 1/2$, , but for $|I|/4$ of the $\hat{x}$'s, $P^m(E_{\hat{x}})$ is tiny. This implies a lower bound on $\mathrm{KL}\left(Q_{\hat{x}}^m \| P^m\right)$ since

$$\mathrm{KL}\left(Q_{\hat{x}}^m \| P^m\right) \geq \mathrm{KL}\left(Q_{\hat{x}}^m(E_{\hat{x}}) \| P^m(E_{\hat{x}})\right),$$

by the data-processing inequality.

Construction of The Events $E_{\hat{x}}$. For every Gibbs-classifier $Q \in \{Q_{\hat{x}} : \hat{x} \in I\} \cup \{P\}$ define its *rounded-hypothesis* $\mathbf{h}_Q : X \to \{\pm 1\}$ as follows:

$$\mathbf{h}_Q(x) = \begin{cases} -1 & \mathbb{E}_{h\sim Q_{\hat{x}}}[h(x)] \leq \frac{q_1+q_2}{2}, \\ +1 & \mathbb{E}_{h\sim Q_{\hat{x}}}[h(x)] > \frac{q_1+q_2}{2}. \end{cases}$$

To simplify notation, let $\mathbf{h}_{\hat{x}} = \mathbf{h}_{Q_{\hat{x}}}$. Note that by the assumption of Lemma 2:

$$\mathbf{h}_{\hat{x}}(x) = \begin{cases} -1 & x < \hat{x}, \\ +1 & x > \hat{x}. \end{cases} \tag{15}$$

In words, each $\mathbf{h}_{\hat{x}}$ is a threshold with a sign-change either right before $\hat{x}$ or right after it. Next, given $\mathbf{h} : I \to \{\pm 1\}$, consider the following iterative process which applies binary-search on $\mathbf{h}$ towards detecting a pair of subsequent coordinates which contain a sign-change.

Binary-Search

Input: $\mathbf{h} : I \to \{\pm 1\}$.

1. Set $I_0 = [a_0, b_0]$, where $a_0 = 0, b_0 = |I| = 2^b$.
2. For $j = 0, \ldots$
   (a) If $|I_j| \leq 2$ then output I_j.
   (b) Query the coordinate $\mathbf{h}(m_j)$, where $m_j = \frac{a_j+b_j}{2}$.
   (c) If $\mathbf{h}(m_j) = +1$ then set $a_{j+1} = a_j, b_{j+1} = m_j$,
   (d) Else, set $a_{j+1} = m_j + 1, b_{j+1} = b_j$.

The following observations follow from the standard analysis of binary-search.

1. The process ends after $b-1$ iterations and each of the points m_j queried in Item (b) are even numbers.
2. If the process is applied on a threshold $\mathbf{h}$ which changes sign from $-$ to $+$ between x and $x+1$ then the output interval I_{out} is $\{x, x+1\}$. Thus, by Equation (15), if we apply this process on $\mathbf{h} = \mathbf{h}_{\hat{x}}$ then $\hat{x} \in I_{out}$.

Given a sequence of hypotheses $h_1, \ldots, h_m : I \to \{\pm 1\}$, define the *empirical rounded-hypothesis* $\mathbf{h}_{h_{1:m}}$ by:

$$\mathbf{h}_{h_{1:m}}(x) = \begin{cases} -1 & \frac{1}{m}\sum_{i=1}^{m} \mathbf{1}[h_i(x) = 1] \leq \frac{q_1+q_2}{2}, \\ +1 & \frac{1}{m}\sum_{i=1}^{m} \mathbf{1}[h_i(x) = 1] > \frac{q_1+q_2}{2}. \end{cases}$$

Consider $h_1, \ldots, h_m \sim Q_{\hat{x}}$ for an odd $\hat{x} \in I$. The following claim shows that with high probability, applying the binary search on $\mathbf{h}_{h_{1:m}}$ yields an output interval I_{out} such that $\hat{x} \in I_{out}$.

Claim 4. *Let $\hat{x} \leq 2^b$ be an odd number. Let J_{out} denote the interval outputted by applying the binary search on $\mathbf{h}_{\hat{x}}$ and let I_{out} denote the interval outputted by applying the binary search on $\mathbf{h}_{h_{1:m}}$, where $h_1, \ldots h_m \sim Q_{\hat{x}}$ are drawn independently. Then,*

$$\Pr_{h_1 \ldots h_m \sim Q_{\hat{x}}^m}[I_{out} \neq J_{out}] \leq b \cdot \exp\Big(-\frac{m}{2}(q_2 - q_1)^2\Big).$$

In particular, if $m = \frac{2(\ln(b)+2)}{(q_2-q_1)^2}$ then $\Pr[\hat{x} \notin I_{out}] \leq \frac{1}{2}$.

Proof. Let $x_1, \ldots x_2, \ldots, x_{b-1}$ be the coordinates queried by the binary search on J_{out}. We will show that with high probability $\mathbf{h}_{h_{1:m}}(x_i) = \mathbf{h}_{\hat{x}}(x_i)$ for every i, which implies that $J_{out} = I_{out}$. Let $i \leq b-1$ and define

$$\mu_i = \mathop{\mathbb{E}}_{h \sim Q_{\hat{x}}}[\mathbf{1}[h(x_i) = +1]] = \Pr_{h \sim Q_{\hat{x}}}[h(x_i) = +1].$$

Note that $\hat{x} \neq x_i$ (because x_i is even and $\hat{x}$ is odd). Therefore, by the assumption of Lemma 2:

$$\mu_i \begin{cases} \leq \frac{q_2+q_1}{2} - \frac{q_2-q_1}{4} & x_i < \hat{x}, \\ \geq \frac{q_2+q_1}{2} + \frac{q_2-q_1}{4} & x_i > \hat{x}. \end{cases}$$

Hence, by a Chernoff bound:

$$\begin{aligned} \Pr_{h_1 \ldots h_m}[\mathbf{h}_{h_{1:m}}(x_i) \neq \mathbf{h}_{\hat{x}}(x_i)] &\leq \Pr_{h_1 \ldots h_m}\Big[\frac{1}{m}\sum_{j=1}^{m} \mathbf{1}[h_j(x_i) = 1] \geq \mu_i + \frac{q_2 - q_1}{4}\Big] \\ &\leq \exp\Big(-\frac{m}{2}(q_2 - q_1)^2\Big) \qquad \text{(Chernoff Bound)} \end{aligned}$$

Thus, by taking a union bound over all $i \leq b-1$ it follows that $\mathbf{h}_{h_{1:m}}(x) = \mathbf{h}_{\hat{x}}(x)$ for every $i \leq b-1$ with probability at least $1 - \log(|I|) \cdot \exp(-\frac{m}{2}(q_2 - q_1)^2)$. In particular, with the above probability we have that $J_{out} = I_{out}$.

Lastly, assume $m = \frac{2(\ln(b)+2)}{(q_2-q_1)^2}$. Then, $b \cdot \exp(-\frac{m}{2}(q_2 - q_1)^2) \leq 1/2$, and therefore $\Pr[J_{out} = I_{out}] \geq \frac{1}{2}$. Since $\mathbf{h}_{\hat{x}}$ is a threshold which changes sign either right before $\hat{x}$ or right after $\hat{x}$, it follows that $\hat{x} \in J_{out}$, and therefore $\Pr[\hat{x} \in I_{out}] \geq 1/2$. □

We are now ready to define the events $E_{\hat{x}}$. Set $m = \frac{2(\ln(b)+2)}{(q_2-q_1)^2}$, according to Claim 4, and let $E_{\hat{x}}$ denote the event that $\hat{x} \in I_{out}$. That is, $E_{\hat{x}}$ is the set of all sequences $h_1, \ldots h_m$ such that $\hat{x} \in I_{out}$, where I_{out} is the interval outputted by the binary-search on $\mathbf{h}_{h_{1:m}}$. Thus, Claim 4 says that $Q_{\hat{x}}^m(E_{\hat{x}}) \geq \frac{2}{3}$ for an odd $\hat{x}$.

Bounding the KL-divergence. We next use the events $E_{\hat{x}}$ to lower bound $\mathrm{KL}\left(Q_{\hat{x}} \| P\right)$:

$$\begin{aligned}
\mathrm{KL}\left(Q_{\hat{x}} \| P\right) &= \frac{1}{m}\mathrm{KL}\left(Q_{\hat{x}}^m \| P^m\right) && \text{(Chain Rule)} \\
&\geq \frac{1}{m}\mathrm{KL}\left(Q_{\hat{x}}^m(E_{\hat{x}}) \| P^m(E_{\hat{x}})\right) && \text{(Data Processing Ineq.)} \\
&\geq \frac{1}{m}\left(-\frac{2}{3}\log\Big(\frac{2}{3}\Big) - \frac{1}{3}\log\Big(\frac{1}{3}\Big) - \frac{2}{3}\log\big(P^m(E_{\hat{x}})\big)\right) \\
&\geq \frac{-\log\big(P^m(E_{\hat{x}})\big) - 1}{2m}
\end{aligned}$$

Therefore, to lower bound $\mathrm{KL}\left(Q_{\hat{x}} \| P\right)$ it suffices to shows that $P^m(E_{\hat{x}})$ is small. We next establish this for $1/4$ of the $\hat{x}$'s in I. Note that whenever $\hat{x}_1, \hat{x}_2 \in I$ are odd and distinct then $E_{\hat{x}_1} \cap E_{\hat{x}_2} = \emptyset$. Indeed, this follows since the outputted interval I_{out} is of size ≤ 2 and hence contains at most one odd number. Thus,

$$\sum_{\hat{x} \text{ is odd}} P^m(E_{\hat{x}}) \leq 1.$$

In particular, since there are 2^{b-1} odd numbers in I, at least $1/2$ of them must satisfy $P^m(E_{\hat{x}}) \leq \frac{1}{2^{b-2}}$. Taken together we obtain that at least $1/4$ of all $\hat{x} \in I$ satisfy:

$$\begin{aligned}
\mathrm{KL}\left(Q_{\hat{x}} \| P\right) &\geq \frac{b-2-1}{2m} \\
&= \frac{b-1}{2\frac{2(\ln(b)+2)}{(q_2-q_1)^2}} = \Omega\Big((q_2-q_1)^2\frac{b}{\log(b)}\Big),
\end{aligned}$$

which finishes the proof of Lemma 2

□