
List Learning with Attribute Noise

Mahdi Cheraghchi

University of Michigan–Ann Arbor
mahdich@umich.edu

Elena Grigorescu

Purdue University
elena-g@purdue.edu

Brendan Juba

Washington University in St. Louis
bjuba@wustl.edu

Karl Wimmer

Duquesne University
wimmerk@duq.edu

Ning Xie

Florida International University
nxie@cis.fiu.edu

Abstract

We introduce and study the model of list learning with attribute noise. Learning with attribute noise was introduced by Shackelford and Volper (COLT, 1988) as a variant of PAC learning, in which the algorithm has access to noisy examples and uncorrupted labels, and the goal is to recover an accurate hypothesis. Sloan (COLT, 1988) and Goldman and Sloan (Algorithmica, 1995) discovered information-theoretic limits to learning in this model, which have impeded further progress. In this article we extend the model to that of list learning, drawing inspiration from the list-decoding model in coding theory, and its recent variant studied in the context of learning. On the positive side, we show that sparse conjunctions can be efficiently list learned under some assumptions on the underlying ground-truth distribution. On the negative side, our results show that even in the list-learning model, efficient learning of parities and majorities is not possible, regardless of the representation used.

1 INTRODUCTION

We study the attribute-noise PAC learning model, introduced by Shackelford and Volper (1988), in which learning must be achieved despite the presence of errors that corrupt the *attributes* of the data (instead of the *labels* of the data that are more commonly used in

the learning with error setting). The inherent difficulty in learning with attribute-noise has been formalized by Sloan (1988) and Goldman and Sloan (1995) by showing information-theoretic barriers: in the presence of attribute-noise, regardless of how much data is used, it is impossible to identify which representations are accurate. Historically, similar issues of identifiability were bypassed in coding theory by relaxing the notion of a solution to that of *list decoding* (Elias, 1957; Wozencraft, 1958); more recently, a similar notion of *list-learning* has been proposed to provide solutions in other learning settings where a correct solution simply cannot be identified from the given data (Balcan et al., 2008; Charikar et al., 2017; Diakonikolas et al., 2018; Karmalkar et al., 2019; Raghavendra and Yau, 2020). We further discuss this previous work in Section 2.2. Here we ask when and to what extent it is possible to overcome the non-identifiability barrier posed by attribute-noise by relaxing the solution to outputting a list of representations of Boolean functions.

In the attribute-noise model, the task is to learn a labeling function given labeled examples, where the examples may have corrupted entries. More formally, the algorithm has access to pairs $(\tilde{x}, c(x))$, where $x = (x_1, x_2, \dots, x_n) \in X$ is drawn independently from an unknown underlying distribution $\mathcal{D}$ over X , $c \in \mathcal{C}$ is an unknown labeling function from a concept class $\mathcal{C}$ over domain X , and $\tilde{x}$ is obtained from x by applying a noise vector $\rho = (\rho_1, \rho_2, \dots, \rho_n)$ from a noise distribution that affects the coordinates (a.k.a. attributes) of x . The goal is to output, with probability at least $1 - \delta$, a hypothesis c' that is $(1 - \epsilon)$ -accurate with respect to c over $\mathcal{D}$, namely $\Pr_{x \in \mathcal{D}}[c(x) = c'(x)] > 1 - \epsilon$. Hence, while in the standard PAC-learning model of Valiant (1984) the algorithm has access to $\tilde{x} = x$ (namely, actual examples from the input distribution), in the

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

A preliminary full version of this work is available as arXiv:2006.06850 (Cheraghchi et al., 2020).

attribute-noise version, the algorithm only has access to a noisy version of x , making the task of learning the labeling function significantly more difficult.

The attribute-noise model captures a setting in which one seeks an accurate model of dependencies in the “ground truth” process captured by $\mathcal{D}$ and c , in spite of errors in the recording of the data. For example, this formulation is appropriate for the task of formulating models in data-driven science; a small list of candidate functions in such a setting then corresponds to a list of possible hypotheses for further investigation. It stands in contrast to the (much easier) *label noise* model, which captures the task of making accurate predictions from the observed data while the observed data is generated from an unknown concept which may not match c . Indeed, if one is only interested in forecasting or building a device that works directly with the noisy data $\tilde{x}$ produced by given real-world sensors, such a setting may be captured by a suitable label-noise model. We stress that since accuracy in the attribute-noise model is assessed with respect to $\mathcal{D}$, which is never observed directly, the attribute-noise model is *not* captured by the label noise model, and is indeed much more challenging than the label noise model. We also note that the attribute-noise model can capture the classical *group testing* problem, see Section 2.1.

All previous work studies concept classes over Boolean attributes $x_i \in \{0, 1\}$ for all $i \in [n]$, and Boolean labeling functions $c : \{0, 1\}^n \rightarrow \{0, 1\}$. Specifically, Shackelford and Volper (1988) show that under *uniform random attribute-noise*, where the noise flips each coordinate independently with probability $p \in [0, 1]$, it is possible to learn k -DNF expressions and conjunctions efficiently, if the noise rate p is known to the algorithm. However, the knowledge of p is not necessary for efficient learning, as proved by Goldman and Sloan (1995). They further consider *product random attribute-noise* on conjunctions, where coordinates are affected independently by noise of possibly different rates p_i , and prove that if these rates are unknown, and if $p_i > 2\epsilon$ in each coordinate, then it is information-theoretically impossible to recover any $(1 - \epsilon)$ -accurate hypothesis. Hence, regardless of the running time of the algorithm, and the number of samples received, the algorithm is unable to output a good answer. On the other hand, if the noise rates are known, Decatur and Gennaro (1995) provide efficient algorithms for PAC-learning conjunctions and k -DNF formulas. More recently, Bshouty et al. (2003) studied the problem of learning in which noise distributions are unconstrained or unknown, but the examples are drawn from the uniform distribution.

We emphasize that the attribute-noise model is not

captured by noisy-PAC. Indeed, the celebrated results of Angluin and Laird (1987) show that learning in the noisy PAC model is information theoretically possible for any noise rate $\rho < 1/2$, and in fact k -CNF and k -DNFs can be learned efficiently in this high-noise regime. Again, this is in contrast with the attribute-noise setting where identifiability is not possible for unknown noise rate $\rho > 2\epsilon$ per coordinate (Goldman and Sloan, 1995). One can also view attribute-noise as an intermediate between noisy PAC and malicious noise, where the latter model assumes that $1 - \rho$ fraction of the output is correct, and the remaining ρ fraction may be completely irrelevant. Kearns and Li (1993) show that under this model in order to identify an ϵ -accurate hypothesis one must have $\rho < \epsilon/(1 + \epsilon)$.

Motivated by its applications in certain real-world machine learning scenarios, as well as its apparent difficulty, we revisit the learning with attribute-noise model and study it under product random attribute-noise, in which the noise rates are *not* known. We overcome the information-theoretic impossibility result of Sloan (1988); Goldman and Sloan (1995) by allowing the algorithm output a small *list* of labeling functions that contains one which is accurate. Thus, even if it is impossible to identify a single accurate function, we can hope to produce a small list of candidate hypotheses that contains an accurate one. Indeed, the proof of Sloan (1988); Goldman and Sloan (1995) follows from an explicit construction of two pairs $(\mathcal{D}_1, c_1, \mathcal{R}_1)$ and $(\mathcal{D}_2, c_2, \mathcal{R}_2)$ of distributions, distinct dictators as labeling functions, and product noise distributions, respectively. The two pairs of tuples lead to exactly the same observed distribution over the $n + 1$ bits received $(\tilde{x}, c(x))$, when $\nu > 2\epsilon$, where ν is an upper bound on the noise amount per attribute. In the list-learning model the algorithm is allowed to output both solutions. In fact, as in PAC learning, any $(1 - \epsilon)$ -accurate hypothesis with respect to the input distribution $\mathcal{D}_i$ is a valid solution to the learning problem, hence it is enough to outputs a small *net* of hypotheses that covers all the valid inputs, in the sense that for any valid input that could have resulted in the observed distribution, the list contains a hypothesis that is $(1 - \epsilon)$ -accurate with respect to that input. Our results provide some sufficient conditions under which efficient list learning is still possible despite the previous barriers. We also show strong lower bounds for most natural classes of Boolean functions.

1.1 The Model: List Learning with Attribute Noise

We denote an instance of the attribute-noise learning problem by a tuple $(\mathcal{D}, c, \mathcal{R})$, where $\mathcal{D}$ is the unknown distribution from which the algorithm receives noisy

samples, c is the labeling function, and $\mathcal{R}$ is the noise distribution. We will denote by $\tilde{\mathcal{D}}$ the observed distribution of $(\tilde{x}, c(x))$, where $\tilde{x} = x + \rho$, $x \leftarrow \mathcal{D}$ and $\rho \leftarrow \mathcal{R}$. We will often abuse notation and denote the marginal distribution on $\tilde{x}$ by $\tilde{\mathcal{D}}$ as well.

For an observed distribution $\tilde{\mathcal{D}}$, a *net* $\mathcal{H}$ (specifically, an ϵ -net) is a set of hypotheses such that for any tuple $(\mathcal{D}, c, \mathcal{R})$ that could have resulted in the observed distribution $\tilde{\mathcal{D}}$, there exists at least one $h \in \mathcal{H}$ that is a $(1 - \epsilon)$ -accurate solution with respect to c and $\mathcal{D}$.

Inspired by the list-decoding model in coding theory, we seek answers to the following general questions:

1. (Combinatorial): Does there exist a small net $\mathcal{H}$ for the attribute-noise learning problem with observed distribution $\tilde{\mathcal{D}}$?
2. (Algorithmic): Can a net for the attribute-noise learning problem with observed distribution $\tilde{\mathcal{D}}$ be computed efficiently?

We formalize these notions below, in the attribute-noise PAC-learning model, with product random noise.

Definition 1.1. (*List learning with random product attribute-noise*) Let $\mathcal{C}$ be a concept class containing Boolean functions $c : \{0, 1\}^n \rightarrow \{0, 1\}$, $\mathcal{D}$ a distribution over $\{0, 1\}^n$, let $\nu, \epsilon \in (0, 1)$, and $0 \leq p_1, \dots, p_n \leq \nu$. Let $\mathcal{R}$ be noise distribution defined as the product of n independent Bernoulli distribution with parameters p_i , $i \in [n]$.

1. (Combinatorial) $\mathcal{C}$ is said to be list-learnable with list size $\ell = \ell(\nu, \epsilon)$ if there exists a net $\mathcal{H}$ for the solutions of the attribute-noise learning problem with input distribution $\mathcal{D}$, such that $|\mathcal{H}| \leq \ell$.
2. (Algorithmic) $\mathcal{C}$ is said to be algorithmically list learnable if there exists a randomized algorithm outputting all $h \in \mathcal{H}$ with probability $1 - \delta$ in time proportional to ℓ .

1.2 Our Results

First, we show that the classes of parities and majorities are not amenable to efficient list learning, as every net for them has exponential size, regardless of the representation used for the net. More generally, we obtain our lower bound for any symmetric family of functions with sufficiently high *noise sensitivity*. (Recall that the noise sensitivity under ρ noise, $\text{NS}_\rho(f)$, is the probability the value of f changes when its inputs are corrupted by product noise of rate ρ .)

Theorem 1.2. (*Theorem 3.3, informal*) Let f be a symmetric function $f : \{0, 1\}^{n/2} \rightarrow \{0, 1\}$. Let $\mathcal{F}_f$

be the family of functions on n bits containing all functions f_S obtained by instantiating f on the set $S \subset [n]$ with $|S| = n/2$. Let $\rho > 0$. Suppose $\epsilon \leq (\frac{1}{2} - o(1))\text{NS}_{\rho/15}(f)$. Then if for every $f_S \in \mathcal{F}_f$ and distribution $\mathcal{D}$ on $\mathbf{x}$ there is an $h \in \mathcal{H}$ satisfying $\Pr_{\mathbf{x} \sim \mathcal{D}}[f_S(\mathbf{x}) \neq h(\mathbf{x})] < \epsilon$, then $|\mathcal{H}| > 2^{\Omega(n)}$.

Two immediate corollaries follow:

Corollary 1.3. Taking $f(x_1, x_2, \dots, x_{n/2}) = \sum_{i=1}^{n/2} x_i$, namely $f = \text{PARITY}_{n/2}$, in Theorem 1.2, the lower bound holds for any $\rho > 0$ and $\epsilon < \frac{1}{4} - o(1)$.

Corollary 1.4. Taking $f(x_1, x_2, \dots, x_{n/2}) = \text{MAJORITY}_{n/2}$, namely $f = \text{MAJORITY}_{n/2}$, in Theorem 1.2, the lower bound holds for any $\rho > 0$ and $\epsilon < \Omega(\sqrt{\rho})$.

We stress that since these lower bounds hold regardless of the representation used in the list, they give lower bounds for richer function classes that contain parities or majorities (respectively) as special cases, such as general linear threshold functions and so on. Of course, such a distinction between “proper” (representation-specific) and “improper” (representation-independent) solutions does not arise in coding theory, but is a common feature in learning theory. Improper learning is the main subject of interest in learning theory, but lower bounds against improper learning algorithms are usually much more challenging. The same holds here: it is generally much easier to argue that an exponential lower bound holds if the function is forced to be a parity function or a conjunction (see below), for example.

Our main results focus on conjunctions, for which we give a general lower bound, and an upper bound for a specific restriction on the input distribution on examples.

Theorem 1.5. (*Theorem A.2, informal*) Let $k > 0$ be an integer, $\epsilon > 0$, and let $\mathcal{C}_k$ be the set of all conjunctions over k bits out of n bits $f : \{0, 1\}^n \rightarrow \{0, 1\}$. If the attribute-noise is $\rho = \frac{1}{k} > 8\epsilon$, then there is an input distribution $\mathcal{D}$ such that list learning $\mathcal{C}_k$ under $\mathcal{D}$ with accuracy ϵ would require a list of size $|\mathcal{H}| > 2^{\Omega(k)}$.

Again, since this theorem is representation-independent, we obtain the same lower bound for any family of functions that can express the conjunctions on k out of n bits. Thus, even with $k = \Omega(n)$, we obtain lower bounds for decision trees, DNFs, s -CNFs, and so on. (By standard reductions, i.e., swapping 0 and 1, one can also obtain the same lower bound for s -DNFs.) Between Theorem 1.2 and the above, we have lower bounds for essentially all of the natural families of functions studied in learning theory, provided that the function depends

on $\omega(\log n)$ coordinates. (When $k = O(\log n)$, the problems are all open, see Section 4.)

Our main result is a sufficient assumption on the input distribution on examples that allows efficient list learning of sparse conjunctions under arbitrary probabilities of flipping individual attributes.

Theorem 1.6. (*Theorem B.1, informal*) *For any positive integer k, k' , and any real number $0 < \epsilon, \delta < 1$, $0 < \gamma \leq 1/2$, there exists a randomized algorithm which, with probability at least $1 - \delta$, list learns k -conjunctions with accuracy $1 - \epsilon$, with sample complexity $\text{poly}(k, \frac{1}{\epsilon}, \frac{1}{\gamma}, \log \frac{1}{\delta})$ and time complexity $\text{poly}(n, \frac{1}{\epsilon}, \log \frac{1}{\delta}, \frac{1}{\gamma}, (\frac{k}{\epsilon\gamma})^k)$ in the attribute-noise model with bit noise rate $0 \leq \nu_i < \frac{1}{2} - \gamma$ for every $1 \leq i \leq n$, under the assumption that the ground-truth distribution is k' -wise independent.*

We note that the trivial PAC learning algorithm that tries all monotone conjunctions of size at most k and outputs the best candidate works only for noise rate $\nu \leq \frac{\epsilon}{2k}$; we include a proof for completeness in the Appendix¹ C.

Theorem 1.6 shows that, even under high noise rates, list learning conjunctions with attribute noise is *fixed-parameter tractable*; meaning that the algorithm runs in polynomial time in n (with a universal constant exponent) for any fixed value of the parameter k . The result only requires a mild restriction of pairwise independence on the ground-truth distribution. This is a much milder assumption than being a product distribution (in turn, much milder than being uniform) which is commonplace in computational learning theory.

2 FURTHER RELATED WORK

We will briefly discuss the history and context of the models we use in this work, as well as a further application to the *group testing* problem.

2.1 Connection with Group Testing

As a further motivation for the study of conjunctions (or, equivalently, disjunctions) as a concept class, here we briefly discuss an application for the classical combinatorial group testing problem; cf. Du and Hwang (2000).

In combinatorial group testing, there is a population of n specimens among which up to k are defective, and the goal is to learn the subset S of the defectives via as few disjunctive tests as possible. Each disjunctive test

picks a group of the items and returns positive if and only if there is a defective among the picked group. Define the Boolean function $f(x_1, \dots, x_n) := \vee_{i \in S} x_i$, so the goal is now reduced to properly learning the disjunction f .

When a subset $T \subseteq [n]$ of the specimens are randomly pooled and the pooled sample is tested, the test result is positive if and only if at least one specimen in T is defective. Observe that in terms of f , the test outcome is $f(t_1, \dots, t_n)$ where $t_i = 1$ iff $i \in T$. Thus, effectively, learning disjunctions captures group testing (the distribution of attributes corresponds to the pooling design, for which it is known that product distributions lead to an essentially optimal number of tests (Du and Hwang, 2000, Chapter 4)).

Now, attribute noise captures the realistic consideration that each sample participating in each pool may register incorrectly with some probability due to such effects as dilution (it is worth noting that our algorithm in Theorem 1.6 actually outputs a *unique* superset of S of size not much larger than k).

As another related example, suppose a small group of people (call them *agents*) have interacted in a social gathering with a population of size n , among which k have a contagious disease. The agents are later tested for the infection, and from the results it is desired to identify the k infected individuals using contact tracing data. This problem has been considered by Cheraghchi et al. (2011). Again, we have a group testing instance, where each agent defines a pool of the individuals, with whom they have interacted in the event. One can assume that each agent has interacted with a random i.i.d. set of individuals (product distribution), and in each interaction the infection (if present) will be caught by the agent with some fixed but unknown probability (attributed noise). Again, the problem reduces to learning disjunctions with attribute noise.

2.2 Further Discussion of Related Work

The information theoretic lower bounds of Sloan (1988); Goldman and Sloan (1995) are analogous to the classical scenario in coding theory, in which, upon receiving a word corrupted by a high amount of noise, decoding becomes ambiguous. (We remark briefly that several authors, including Rubin (1976), Schuurmans and Greiner (1994), and Michael (2010) furthermore have introduced models of learning from data with omitted attributes, which is analogous to decoding from erasures, and is much easier.) As a result, Elias (1957) and Wozencraft (1958) extended the classical notion of unique decoding to that of *list-decoding*, where the algorithm is required to output a list of all possible messages that could have resulted in the re-

¹The appendices are available as supplementary material.

ceived one. A similar motivation prompted Balcan et al. (2008) to introduce the notion of *list-decodable learning* in the context of clustering, where their algorithm is required to output a small list that includes a “good” clustering, with high probability. Follow-up results by Charikar et al. (2017) use this framework in the context of learning from untrusted data when there is a minority fraction of “inliers” and so identifiability cannot hold. In the same vein, Diakonikolas et al. (2018) obtain algorithms for robust mean estimation, and learning mixtures of Gaussians. More recently, Karmalkar et al. (2019) and Raghavendra and Yau (2020) independently gave list-decodable linear regression algorithms for this minority-inlier setting. In all of these works, the difference is that there is guaranteed to be a fixed fraction of uncorrupted examples (whereas the corruption of the remaining examples is arbitrary). By contrast, in the attribute-noise model we study, with high probability *every* example has a non-negligible fraction of corrupted attributes, though conversely, the corruptions are stochastic and independent. Nevertheless, in spite of ours being a stochastic-noise model, we will see that the lack of clean examples still poses serious challenges, even for a list learner.

3 LOWER BOUNDS AND A CONSTRUCTIVE ALGORITHM

We now give a technical overview of our results.

3.1 The Lower Bounds, Theorems 1.2 and 1.5

The high-level idea of the lower-bound proofs is to explicitly construct a large set of labeling functions $c \in \mathcal{C}$ and initial input and noise distributions such that any function in the net can only be $(1 - \epsilon)$ -accurate for a small number of possible initial solutions $(\mathcal{D}, c, \mathcal{R})$, regardless of the representations used for the functions in the net. Hence, to cover an exponential number of such potential solutions a net has to have large size. The construction of the initial distributions exploits the idea that bits (x_{2i}, x_{2i+1}) that are ρ -correlated (meaning that x_{2i+1} takes the same value as x_{2i} w.p. $1 - \rho$, and takes the flipped value with probability ρ) appear identical to an observer when adding Bernoulli random noise ρ to one copy and no noise to the other copy.

For the proof, define the **noise operator at ρ on S** , denoted by $N_{S,\rho}(x)$, to be a random string such that $N_{S,\rho}(x)_i$ is a uniform random bit ρ -correlated with x_i if $i \in S$, and $N_{S,\rho}(x)_i = x_i$ with probability 1 for $i \notin S$. We further define the **noise sensitivity at ρ on S** to be $\text{NS}_{S,\rho}(f) = \Pr_{\mathbf{y} \sim \mathcal{U}_{n/2}}[f(\mathbf{y}) \neq f(N_{S,\rho}(\mathbf{y}))]$. These are related to the standard noise sensitivity constructions via $N_\rho(x) = N_{[n],\rho}(x)$, and

$$\text{NS}_\rho(f) = \Pr_{\mathbf{y} \sim \mathcal{U}_{n/2}}[f(\mathbf{y}) \neq f(N_\rho(\mathbf{y}))]:$$

Claim 3.1. *Let $S \subseteq [n]$ be a set such that $|S| = n/14$. For every symmetric Boolean function f on $n/2$ variables such that $\text{NS}_{S,\rho}(f) = 2^{-o(n)}$ for all S , $\text{NS}_{S,\rho}(f) \geq (1 - o(1))\text{NS}_{\rho/15}(f)$.*

We defer the proof of Claim 3.1 to Appendix A.

Before defining the functions in $\mathcal{F}$ for Theorem 1.2, we will make some notational conventions. For the sake of presentation we assume n is even.

For a string $x \in \{0, 1\}^n$, we may view it as the concatenation of pairs (x_{2i+1}, x_{2i+2}) , for $i = 0, 1, \dots, n/2 - 1$, and define two strings $x^0, x^1 \in \{0, 1\}^{n/2}$, by selecting the odd, respectively the even, indices of these pairs in order, namely $x^0 = x_1, x_3, \dots, x_{n-1}$ and $x^1 = x_2, x_4, \dots, x_n$. For $x \in \{0, 1\}^n$ and a string $z \in \{0, 1\}^{n/2}$, we define the hybrid string $x^z \in \{0, 1\}^{n/2}$ to be the string that for each $0 \leq i \leq n/2 - 1$ selects either x_{2i+1} if $z_i = 0$, or x_{2i+2} if $z_i = 1$, denoted by $x^z = (x_1^{z_1}, x_2^{z_2}, \dots, x_{n/2}^{z_{n/2}})$, where $x_i^{z_i} = x_{2i+1}$ if $z_i = 0$, and $x_i^{z_i} = x_{2i+2}$ if $z_i = 1$.

We now define the set of functions $\mathcal{F}$. For a symmetric function $f : \{0, 1\}^{n/2} \rightarrow \{0, 1\}$, such as *parity* or *majority*, and a string $z \in \{0, 1\}^{n/2}$, let $f^z : \{0, 1\}^n \rightarrow \{0, 1\}$ be the function $f^z(x) = f(x^z) = f(x_1^{z_1}, x_2^{z_2}, \dots, x_{n/2}^{z_{n/2}})$. Let $\mathcal{F} = \mathcal{F}(f) = \{f^z\}_{z \in \{0, 1\}^{n/2}}$.

Further, for $z \in \{0, 1\}^{n/2}$ let $\mathcal{D}^z$ be the distribution² on $\{0, 1\}^n$ defined by the following probability experiment:

- The coordinates in $\mathbf{x}^z$ are drawn independently and uniformly at random. That is, $\mathbf{x}^z \sim \mathcal{U}_{n/2}$, where $\mathcal{U}_{n/2}$ represents the uniform distribution on $\{0, 1\}^{n/2}$.
- The coordinates in $\mathbf{x}^{\bar{z}}$ are ρ -noisy copies of $\mathbf{x}^z$ —each bit $x_i^{\bar{z}}$ is a ρ -noisy copy of x_i^z .

We will show that if z is unknown, and we see labeled examples according to f^z under $\mathcal{D}^z$ with ρ -bounded attribute-noise, then list-learning to small accuracy requires an exponential size list. That is, for every set of functions $\mathcal{H}$ (our proposed net), the quantity $\max_{z \in \{0, 1\}^{n/2}} \min_{h \in \mathcal{H}} \Pr_{\mathbf{x} \sim \mathcal{D}^z}[f^z(x) \neq h(x)]$ is “large” if $|\mathcal{H}|$ is sub-exponential.

For f^z with respect to $\mathcal{D}^z$, given x , the attribute-noise $N_\rho^z(x)$ is as follows: we apply ρ -noise to each $x_i^{z_i}$, and no noise to $x_i^{\bar{z}_i}$. It follows that for *every* $\mathcal{D}^z$, the resulting distribution over the *labeled* examples is the same. We define $\mathcal{D}$ to be distribution³ on $\{0, 1\}^n$ such that, for each i , x_i^0 and x_i^1 are ρ -correlated uniformly

²Actually, $\mathcal{D}^z$ is the same distribution no matter what z is.

³Actually, this is the same as $\mathcal{D}^z$.

random bits, and the $n/2$ pairs (x_i^0, x_i^1) are chosen independently.

We now exploit the fact that totally symmetric functions with high “noise sensitivity” are often far apart.

Lemma 3.2. *Let $z, z' \in \{0, 1\}^{n/2}$ be strings such that $|z - z'| \geq n/14$. Then $\Pr_{\mathbf{x} \sim \mathcal{D}^z}[f^z(\mathbf{x}) \neq f^{z'}(\mathbf{x})] \geq (1 - o(1))\text{NS}_{\rho/15}(f)$.*

Proof. Define S to be the set of strings where z and z' differ.

$$\begin{aligned} \Pr_{\mathbf{x} \sim \mathcal{D}^z}[f^z(\mathbf{x}) \neq f^{z'}(\mathbf{x})] &= \Pr_{\mathbf{x} \sim \mathcal{D}^z}[f(\mathbf{x}^z) \neq f(\mathbf{x}^{z'})] \\ &= \Pr_{\mathbf{x} \sim \mathcal{D}^z}[f(\mathbf{x}^z) \neq f(N_{S,\rho}(\mathbf{x}^z))] \\ &= \Pr_{\mathbf{y} \sim \mathcal{U}_{n/2}}[f(\mathbf{y}) \neq f(N_{S,\rho}(\mathbf{y}))] \\ &= \text{NS}_{S,\rho}(f) \\ &\geq (1 - o(1))\text{NS}_{\rho/15}(f), \end{aligned}$$

where the last inequality follows by Claim 3.1. $\square$

Thus, any single member of the net can only be accurate for at most one of these far pairs, and so we must have a large net. We thus obtain a more specific version of Theorem 1.2.

Theorem 3.3. *Let $f : \{0, 1\}^{n/2} \rightarrow \{0, 1\}$ be a symmetric function, and $\rho > 0$. If $\epsilon \leq (\frac{1}{2} - o(1))\text{NS}_{\rho/15}(f)$ then, for family $\mathcal{F} = \{f^z\}_{z \in \{0, 1\}^{n/2}}$ of Boolean functions on n bits where the oracle produces examples with attribute-noise rate ρ , we have that any net $\mathcal{H}$ satisfying $\max_{z \in \{0, 1\}^{n/2}} \min_{h \in \mathcal{H}} \Pr_{\mathbf{x} \sim \mathcal{D}^z}[f^z(\mathbf{x}) \neq h(\mathbf{x})] < \epsilon$ must have $|\mathcal{H}| > 2^{\Omega(n)}$.*

Proof. By the triangle inequality, no function in the net can approximate both f^z and $f^{z'}$ for two strings z, z' where $|z - z'| \geq n/14$ (with respect to $\mathcal{D} = \mathcal{D}^z = \mathcal{D}^{z'}$) to within $(\frac{1}{2} - o(1))\text{NS}_{\rho/15}(f)$. Thus, any function in the net can cover at most $\binom{n}{n/14}$ such functions f^z with respect to $\mathcal{D}^z$. It follows that any net requires $2^{n/2} / \binom{n}{n/14} \geq 2^{n/14}$ functions (here we used that $\binom{n}{k} < (ne/k)^k$, with $k = n/14$). $\square$

The proof of Theorem 1.5 uses similar ideas, and appears in Appendix A.2. In this case we use distributions $\mathcal{D}^z$ over $\{0, 1\}^{2k}$ such that

- The coordinates in $\mathbf{x}^z$ are drawn independently at random with bias $1/k$. That is, $\mathbf{x}^z \sim \mu_{k,1/k}$, where $\mu_{n,p}$ denotes the p -biased distribution over $\{0, 1\}^n$.
- The coordinates in $\mathbf{x}^{\bar{z}}$ are ρ -noisy copies of $\mathbf{x}^z$ —each bit $x_i^{\bar{z}}$ is a ρ -noisy copy of x_i^z .

For f^z with respect to $\mathcal{D}^z$, given $\mathbf{x}$, the attribute-noise $N_\rho^z(\mathbf{x})$ is as follows: we apply ρ -noise to each x_i^z , and

no noise to $x_i^{\bar{z}}$. It follows again that for every $\mathcal{D}^z$, the resulting distribution over the labeled examples is the same.

We show a function in the net covers the most conjunctions by taking f to be 1 on $199k/100$ of these strings and 0 on the other $k/100$ since for every conjunction, a false 0 is roughly k times as costly as a false 1: To make the error less than $(1 - 1/k)^{k-1}(1/k)(1 - \rho)^{k-1}\rho \cdot (99k/100)$, there must be a function in the net that has no false 0’s and at most $99k/100$ false 1’s on these strings. (A function is covered if its bits correspond to those with ones.) There are $2^{99k/100}$ conjunctions covered, but 2^k conjunctions in total, so any net must have $2^{k/100}$ functions in it to achieve error below $(1 - 1/k)^{k-1}(1/k)(1 - \rho)^{k-1}\rho \cdot (99k/100)$. Taking $\rho = 1/k$, the error is at least $\rho/8$, so we need $\rho < 8\epsilon$ for a sub-exponential net.

3.2 The Conjunction Algorithm, Theorem 1.6

The essential difficulty in learning conjunctions under the attribute-noise model is that on the one hand, conjunctions are in general very sensitive to the attributes that appear in them; missing even one significant attribute incurs a large error. But, on the other hand, as illustrated in the lower bound, it is in general impossible to distinguish bits of the conjunction corrupted by noise in our examples from bits that would thus incur a serious error if they were included in the conjunction. Thus, we seek to find a small set of candidate coordinates and output all small subsets of these. Both the size of the set of candidates and the size of the conjunctions must be small to obtain a polynomial-size list. Proving that the algorithm does output a net for the solution space is the most difficult part of our arguments, the difficulty emerging from the fact that the accuracy of the solution is measured against the original unknown distribution rather than the observed distribution itself. The algorithm can only perform tests and optimize quantities using the corrupted examples, and we must then bound the distances from the unknown distribution.

We use the following notation: $\tilde{D}$: the observed distribution; D : the original distribution before applying the attribute-noise; $c = \bigwedge_{i \in c} \ell_i$: a conjunction⁴ of size at most k , where $c \subset [n]$, $|c| \leq k$ and ℓ_i is either x_i or $1 - x_i$; D_b (resp. $\tilde{D}_b$): the original (resp. observed) distribution conditioned on label c being b , for $b \in \{0, 1\}$; ν_i : the attribute-noise rate of bit i .

⁴We abuse notation here to let c denote both the conjunction and the set of variables in the conjunction. Furthermore, the conjunction over the empty set is understood to be 1.

Algorithm 1: Learning-Conjunction ($\tilde{E}X, k, \epsilon, \delta$)

input : Noisy example oracle $\tilde{E}X(c, D)$, integer k , error parameter ϵ , and confidence parameter δ
output: A list of conjunctions

```

1  $m := 32k^2/(\epsilon^5\gamma^2)$ 
2  $M := O(k^4 \log n \log(1/\delta)/(\epsilon^9\gamma^4))$ 
3  $\mathcal{M} \leftarrow M$  random labeled examples drawn from
   the noisy example oracle  $\tilde{E}X(c, D)$ 
4  $S \leftarrow$  Pairwise-Independence-Test ( $\mathcal{M}, \epsilon, \delta$ )
5 for  $i \leftarrow 1$  to  $n$  do
6   Use  $\mathcal{M}$  to estimate label sensitivity at the  $i^{\text{th}}$ 
   bit  $\widehat{LS}_i$ 
7   if  $\widehat{LS}_i < \epsilon\gamma/k$  then
8     remove  $i$  from  $S$ 
9 if  $|S| < m$  then
10  Output the list of conjunctions
    $\mathbf{0} \cup \{\wedge_{i \in c'} x_i\}_{c' \in \binom{S}{\leq k}}$ 
11 else
12  Output  $\mathbf{0}$ 
    
```

We call a bit $i \in [n]$ a *conjunction bit* if $i \in c$ and *non-conjunction bit* otherwise. Note that without loss of generality, we may assume that every candidate conjunction bit in S is biased towards 1, i.e. $E_{\tilde{D}}[x_i] \geq 1/2$ for every $i \in S$, as otherwise we simply replace x_i with $1 - x_i$ in our arguments.

The algorithm for list learning conjunctions under random product attribute-noise operates under the assumption that the attributes in the initial distribution on examples are pairwise independent.

Definition 3.4 (Non-uniform k -wise independence). *Let $P : \{0, 1\}^n \rightarrow \mathbb{R}^{\geq 0}$ be a distribution and k be a positive integer. P is said to be (non-uniform) k -wise independent if for any subset of k indices $\{i_1, \dots, i_k\} \subset [n]$ and for any $z_1 \dots z_k \in \{0, 1\}^k$,*

$$\Pr_P[X_{i_1} \cdots X_{i_k} = z_1 \cdots z_k] = \Pr_P[X_{i_1} = z_1] \times \cdots \times \Pr_P[X_{i_k} = z_k].$$

Claim 3.5. *For any positive integer k and any distribution $D : \{0, 1\}^n \rightarrow \mathbb{R}^{\geq 0}$, D is k -wise independent if and only if D is k -wise independent. In other words, attribute-noise does not change the k -wise independence of the underlying distribution.*

We defer the proof of this claim to Appendix D.

We first observe that since the bits of the actual conjunction must all take value 1 on label 1, and the noise is a product distribution, the bits of the actual con-

Algorithm 2: Pairwise-Independence-Test ($\mathcal{M}, \epsilon, \delta$)

input : M random labeled examples $\mathcal{M}$, error parameter ϵ , and confidence parameter δ
output: A subset $S \subset [n]$ of nearly pairwise independent bits under $\tilde{D}_1$

```

1  $S \leftarrow [n]$ 
2 for  $i \leftarrow 1$  to  $n$  do
3   Use positive examples in  $\mathcal{M}$  to empirically
   estimate  $\widehat{E}_{\tilde{D}_1}[x_i]$ 
4 for  $i \leftarrow 1$  to  $n - 1$  do
5   for  $j \leftarrow i + 1$  to  $n$  do
6     if  $i \notin S$  or  $j \notin S$  then
7       continue
8     if  $\widehat{E}_{\tilde{D}_1}[x_i] \leq 1/(8\epsilon m)$  or
        $\widehat{E}_{\tilde{D}_1}[x_j] \leq 1/(8\epsilon m)$  then
9       continue
10    Use sampled examples to empirically
    estimate  $\widehat{E}_{\tilde{D}_1}[x_i \cdot x_j]$ 
11    if  $|\widehat{E}_{\tilde{D}_1}[x_i] \cdot \widehat{E}_{\tilde{D}_1}[x_j] - \widehat{E}_{\tilde{D}_1}[x_i \cdot x_j]| >$ 
        $1/(8\epsilon m)$  then
12      Remove both  $i$  and  $j$  from  $S$ 
13 Output  $S$ 
    
```

junction in the noisy examples are fully independent when conditioned on label 1.

Learning conjunctions is easy when there is no attribute-noise because, if x_i is in the conjunction, then conditioned on label being 1, $\Pr[X_i = 1] = 1$ and this probability should be lower without the conditioning — unless variable x_i is almost surely being 1 under the distribution D . In other words, the expectation of a (relevant) conjunction bit should be *sensitive* to label change. This is also true under attribute-noise, although with lower sensitivity in general.

Definition 3.6. *The (observed) label sensitivity at bit i is defined by $LS_i = E_{\tilde{D}_1}[X_i] - E_{\tilde{D}_0}[X_i]$; that is, LS_i is the difference between expectation of x_i conditioned on label being 1 and the expectation of x_i conditioned on label being 0.*

The algorithm thus first identifies the subset of variables that are (at least) pairwise independent on label 1, and then eliminates from this surviving set the variables that are not too sensitive to the label. These eliminated variables could not have been significant bits of the conjunction: if there is no attribute-noise, the variables in the conjunction would be very sensitive to the label, since they would always take value 1 on label 1, and they would take value 0 on label 0 sig-

nificantly often. Indeed, we note the following simple fact: since attribute-noise does not change the labels of examples, the total mass of positive or negative examples are the same for D and $\tilde{D}$.

Fact 3.7. *For any underlying distribution D of the example oracle and any attribute-noise vector ν , $\Pr_D[c(x) = 1] = \Pr_{\tilde{D}}[c(x) = 1]$ and $\Pr_D[c(x) = 0] = \Pr_{\tilde{D}}[c(x) = 0]$.*

Now, either the function is nearly constant and so a constant function predicts the label sufficiently well, or else there is a bounded statistical distance between the distribution conditioned on label 1 and the original distribution, which is a mixture of the label 1 and label 0 distributions. We show that when the function is far from constant, there cannot be too many coordinates surviving. Intuitively, otherwise, the weight would allow us to distinguish the label 1 distribution from the original distribution beyond the statistical distance, due to Chebyshev’s inequality: the total weight would concentrate if there were many coordinates left. Thus we can afford to enumerate all small subsets of the surviving coordinates in this case.

In summary, our list-learning algorithm is described in Algorithm 1, in which we call Pairwise-Independence-Test (Algorithm 2, see Appendix B) as a subroutine to select the pairwise independent variables under distribution $\tilde{D}_1$. Essentially, for any pair of variables that each take both values with probability at least $1/(8\epsilon m)$, if the distance from independence is at least $1/(8\epsilon m)$, we filter out both members of the pair. Our analysis of Algorithm 1, sketched above, establishes Theorem 1.6; see Appendix B for the full details.

4 OPEN PROBLEMS

In our work we seek to develop the natural, yet difficult-to-analyze model of learning under attribute noise. While we prove several impossibility results and a sufficient condition for learning sparse conjunctions, our work leaves open a plethora of intriguing possibilities. We describe below a few important ones.

The first, most natural question is whether or not the pairwise-independence assumption is really needed for our algorithm:

Open Question 4.1. *Is the set of sparse conjunctions list-learnable under arbitrary product distributions of the attribute-noise?*

But, moreover, we note that our lower bounds do not rule out the possibility of obtaining polynomial-size lists for $O(\log n)$ -sparse functions in general. So it is still open whether or not natural function families with small numbers of relevant coordinates have efficient

list-learning algorithms, e.g.:

Open Question 4.2. *Is the set of sparse Boolean threshold functions list-learnable under arbitrary product distributions of the attribute-noise?*

Thus, in contrast to the usual theory of supervised learning, we do not have a characterization of which families of functions are (information-theoretically) learnable in terms of some parameter like the VC-dimension or Rademacher complexity in the attribute-noise list-learning setting:

Open Question 4.3. *What are necessary and sufficient conditions for families of Boolean functions to be list-learnable under the product distribution of the attribute-noise?*

Or, more generally:

Open Question 4.4. *What families of Boolean functions are list-learnable under general (not-necessarily independent product) noise distributions?*

Of course, one can ask both computational/algorithmic and statistical/combinatorial variants of these questions. But again, a central difficulty here is that the usual statistical techniques for estimating losses from data cannot be used directly to estimate losses from our corrupted data. Thus it seems that new tools may need to be developed to address these questions.

Acknowledgements

Mahdi Cheraghchi’s research was partially supported by the National Science Foundation under Grant No. CCF-2006455. Elena Grigorescu was supported in part by NSF awards CCF-1910659 and NSF CCF-1910411. Brendan Juba was supported by NSF awards CCF-1718380, IIS-1908287, and IIS-1939677, and was hosted by the Simons Institute for Theory of Computing during portions of this work. Ning Xie’s research was partially supported by grant ARO W911NF1910362.

References

- D. Angluin and P. D. Laird. Learning from noisy examples. *Mach. Learn.*, 2(4):343–370, 1987.
- M. Balcan, A. Blum, and S. Vempala. A discriminative framework for clustering via similarity functions. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 671–680, 2008.
- N. H. Bshouty, J. C. Jackson, and C. Tamon. Uniform-distribution attribute noise learnability. *Inf. Comput.*, 187(2):277–290, 2003.

- M. Charikar, J. Steinhardt, and G. Valiant. Learning from untrusted data. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 47–60, 2017.
- M. Cheraghchi, A. Hormati, A. Karbasi, and M. Vetterli. Group testing with probabilistic tests: Theory, design and application. *IEEE Transactions on Information Theory*, 57(10):7057–7067, 2011.
- M. Cheraghchi, E. Grigorescu, B. Juba, K. Wimmer, and N. Xie. List learning with attribute noise. *arXiv preprint arXiv:2006.06850*, 2020.
- S. E. Decatur and R. Gennaro. On learning from noisy and incomplete examples. In *Proceedings of the Eighth Annual Conference on Computational Learning Theory, COLT 1995, Santa Cruz, California, USA, July 5-8, 1995*, pages 353–360, 1995.
- I. Diakonikolas, D. M. Kane, and A. Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1047–1060, 2018.
- D.-Z. Du and F. Hwang. *Combinatorial Group Testing and its Applications*. World Scientific, second edition, 2000.
- P. Elias. List decoding for noisy channels. *Technical Report 335, Research Laboratory of Electronics, MIT*, 1957.
- S. A. Goldman and R. H. Sloan. Can PAC learning algorithms tolerate random attribute noise? *Algorithmica*, 14(1):70–84, 1995.
- S. Karmalkar, P. Kothari, and A. Klivans. List-decodable linear regression. In *Advances in Neural Information Processing Systems 32*, pages 7423–7432, 2019.
- M. J. Kearns and M. Li. Learning in the presence of malicious errors. *SIAM J. Comput.*, 22(4):807–837, 1993.
- L. Michael. Partial observability and learnability. *Artificial Intelligence*, 174(11):639–669, 2010.
- R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, USA, 2014. ISBN 1107038324.
- P. Raghavendra and M. Yau. List decodable learning via sum of squares. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms*, pages 161–180, 2020.
- D. B. Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.
- D. Schuurmans and R. Greiner. Learning default concepts. In *Proceedings of the Tenth Canadian Conference on Artificial Intelligence (AI’94)*, pages 99–106, 1994.
- G. Shackelford and D. Volper. Learning k -DNF with noise in the attributes. In *Proceedings of the First Annual Workshop on Computational Learning Theory, COLT ’88, Cambridge, MA, USA, August 3-5, 1988.*, pages 97–103, 1988.
- R. H. Sloan. Types of noise in data for concept learning. In *Proceedings of the First Annual Workshop on Computational Learning Theory, COLT ’88, Cambridge, MA, USA, August 3-5, 1988.*, pages 91–96, 1988.
- L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984.
- J. M. Wozenraft. List Decoding. *Quarterly Progress Report, Research Laboratory of Electronics, MIT*, 48:90–95, 1958.

Appendix

A LOWER BOUNDS

A.1 Noise Sensitivity Lower Bound for Some Symmetric Functions

In this section we show that some families of symmetric functions on subsets of half the bits are hard to improperly learn in an information-theoretic sense, and prove Theorem 1.2.

It can be easily checked that the distribution $\mathcal{D}$ has the following properties:

- For every $z \in \{0,1\}^{n/2}$ and a random string $\mathbf{x} \sim \mathcal{D}$, $\mathbf{x}^z$ is distributed as a uniformly random string over $\{0,1\}^{n/2}$.
- For every pair of strings $z, z' \in \{0,1\}^{n/2}$ and a random string $\mathbf{x} \sim \mathcal{D}$, the random strings $\mathbf{x}^z$ and $\mathbf{x}^{z'}$, restricted to the coordinates where z and z' disagree, are ρ -noisy copies of each other.
- To construct the distribution of $\mathbf{x}^{z'}$ from $\mathbf{x}^z$, one can apply ρ -noise to the coordinates of $\mathbf{x}^z$ in those coordinates where z and z' differ (and just read off the coordinates of $\mathbf{x}^z$ where they are the same).
- In fact, $\mathcal{D}^z$ is identical to $\mathcal{D}$ for every $z \in \{0,1\}^{n/2}$. However, the distribution of labeled examples $\langle \mathbf{x}, f^z(\mathbf{x}) \rangle$ where $\mathbf{x} \sim \mathcal{D}^z$ depends on z . The distribution of labeled examples after attribute noise $\langle N_\rho^z(\mathbf{x}), f^z(\mathbf{x}) \rangle$ is independent of z ; the marginal distribution on $N_\rho^z(\mathbf{x})$ is $\mathcal{D} = \mathcal{D}^z$.

A.1.1 Noise Sensitivity

For concreteness, recall that the **noise operator at ρ on S** is denoted by $N_{S,\rho}(x)$ is a random string such that $N_{S,\rho}(x)_i$ is a uniform random bit ρ -correlated with x_i if $i \in S$, and $N_{S,\rho}(x)_i = x_i$ with probability 1 for $i \notin S$ (cf. O'Donnell (2014)). We now prove Claim 3.1:

Proof of Claim 3.1. Note that, for every x , $N_{\rho/15}(x)$ is distributed as $N_{\mathbf{T},\rho}(x)$, where $\mathbf{T}$ is a set where each coordinate is included independently with probability $1/15$. It follows that

$$\begin{aligned} \text{NS}_{\rho/15}(f) &= \Pr_{\mathbf{y} \sim \mathcal{U}_{n/2}} [f(\mathbf{y}) \neq f(N_{\rho/15}(\mathbf{y}))] \\ &= \Pr_{\mathbf{y} \sim \mathcal{U}_{n/2}} [f(\mathbf{y}) \neq f(N_{\mathbf{T},\rho}(\mathbf{y}))] \\ &= \mathbb{E}_{\mathbf{T}} [\text{NS}_{\mathbf{T},\rho}(f)]. \end{aligned}$$

By a Chernoff bound, $\Pr[|\mathbf{T}| \leq n/14] \geq 1 - 2^{-\Omega(n)}$. Thus, for a set S such that $|S| = n/14$, we have

$$\begin{aligned} \text{NS}_{\rho/15}(f) &= \mathbb{E}_{\mathbf{T}} [\text{NS}_{\mathbf{T},\rho}(f)] \\ &= \mathbb{E}_{\mathbf{T}} [\text{NS}_{\mathbf{T},\rho}(f) \mid |\mathbf{T}| \leq n/14] \Pr_{\mathbf{T}}[|\mathbf{T}| \leq n/14] \\ &\quad + \mathbb{E}_{\mathbf{T}} [\text{NS}_{\mathbf{T},\rho}(f) \mid |\mathbf{T}| > n/14] \Pr_{\mathbf{T}}[|\mathbf{T}| > n/14] \\ &\leq \text{NS}_{S,\rho}(f) \Pr[|\mathbf{T}| \leq n/14] + 2^{-\Omega(n)} \\ &\leq \text{NS}_{S,\rho}(f)(1 + o(1)), \end{aligned}$$

where we used the fact that $\text{NS}_{S,\rho}$ is nondecreasing as $|S|$ increases. (Since we assumed that f is symmetric, only $|S|$ matters.) Dividing both sides by the $(1 + o(1))$ factor yields the claim. $\square$

Remark A.1. *The symmetric assumption can be relaxed by noting that the bound works for any function that is roughly balanced over the uniform distribution, since the noise sensitivity of such functions is $\Omega(\min\{\Pr[f(\mathbf{x}) = 0], \Pr[f(\mathbf{x}) = 1]\})$. Roughly speaking, this result asserts that we cannot learn with error smaller than the noise sensitivity.*

A.2 Maximum Sensitivity Lower Bound for Conjunctions

In this section we show a lower bound for improper list learning of conjunctions and by proving a more specific version of Theorem 1.5. We will use the same notation as in Section A.1.

Theorem A.2. *Let $k > 0$ be an integer, $\epsilon > 0$, and let $\mathcal{C}_k$ be the set of all conjunctions over k bits out of n bits $f : \{0, 1\}^n \rightarrow \{0, 1\}$. If the attribute noise is $\rho = \frac{1}{k} > 8\epsilon$, then any net $\mathcal{H}$ of functions satisfying*

$$\max_{z \in \{0, 1\}^{n/2}} \min_{h \in \mathcal{H}} \Pr_{\mathbf{x} \sim \mathcal{D}^z} [f^z(\mathbf{x}) \neq h(\mathbf{x})] < \epsilon$$

must have $|\mathcal{H}| > 2^{\Omega(k)}$.

Proof. Suppose that the distribution $\mathcal{D}^z$ over $\{0, 1\}^{2k}$ is such that

- The coordinates in $\mathbf{x}^z$ are drawn independently at random with bias $1/k$. That is, $\mathbf{x}^z \sim \mu_{k, 1/k}$, where $\mu_{n, p}$ denotes the p -biased distribution over $\{0, 1\}^n$.
- The coordinates in $\mathbf{x}^{\bar{z}}$ are ρ -noisy copies of $\mathbf{x}^z$; specifically, each bit $x_i^{\bar{z}}$ is a ρ -noisy copy of x_i^z .

We will show that if z is unknown, and we see labeled examples according to f^z under $\mathcal{D}^z$ with ρ -bounded attribute noise, then list-learning to small accuracy requires an exponential size list. That is, for every set of functions $\mathcal{H}$ (our proposed net), the quantity

$$\max_{z \in \{0, 1\}^{n/2}} \min_{h \in \mathcal{H}} \Pr_{\mathbf{x} \sim \mathcal{D}^z} [f^z(\mathbf{x}) \neq h(\mathbf{x})]$$

is “large” if $|\mathcal{H}|$ is sub-exponential in k .

For f^z with respect to $\mathcal{D}^z$, given x , the attribute noise $N_\rho^z(x)$ is as follows: we apply ρ -noise to each $x_i^{z_i}$, and no noise to $x_i^{\bar{z}_i}$. It follows that for *every* $\mathcal{D}^z$, the resulting distribution over the *labeled* examples is the same. We define $\mathcal{D}$ to be distribution⁵ on $\{0, 1\}^n$ such that, for each i , x_i^0 and x_i^1 are ρ -correlated random bits with bias $(1 - \rho)(1/k) + \rho(1 - 1/k)$, and the k pairs (x_i^0, x_i^1) are chosen independently. It can be easily checked that the distribution $\mathcal{D}$ has the following properties:

- For every $z \in \{0, 1\}^{n/2}$ and a random string $\mathbf{x} \sim \mathcal{D}$, $\mathbf{x}^z$ is distributed as a uniformly random string over $\{0, 1\}^{n/2}$.
- For every pair of strings $z, z' \in \{0, 1\}^{n/2}$ and a random string $\mathbf{x} \sim \mathcal{D}$, the random strings $\mathbf{x}^z$ and $\mathbf{x}^{z'}$, restricted to the coordinates where z and z' disagree, are ρ -noisy copies of each other.
- To construct $\mathbf{x}^{z'}$ from $\mathbf{x}^z$, one can apply ρ -noise to the coordinates of $\mathbf{x}^z$ in those coordinates where z and z' differ (and just read off the coordinates of $\mathbf{x}^z$ where they are the same).
- In fact, $\mathcal{D}^z$ is identical to $\mathcal{D}$ for every $z \in \{0, 1\}^{n/2}$. However, the distribution of labeled examples $\langle \mathbf{x}, f^z(\mathbf{x}) \rangle$ where $\mathbf{x} \sim \mathcal{D}^z$ depends on z . The distribution of labeled examples after attribute noise $\langle N_\rho^z(\mathbf{x}), f^z(\mathbf{x}) \rangle$ is independent of z ; the marginal distribution on $N_\rho^z(\mathbf{x})$ is $\mathcal{D} = \mathcal{D}^z$.

Unlike the uniform distribution case, when we consider the accuracy of a function in the net on a conjunction, the distribution under which we calculate the error depends on the conjunction. We compute the following quantities first:

- The probability of the all-0’s string in the true distribution is $(1 - 1/k)^k (1 - \rho)^k$; the all 0’s string in drawn in the conjunction bits, and no flips occur in the noisy version.

⁵Actually, this is the same as $\mathcal{D}^z$.

- The probability of a string of all-0's, except for $x_i^b = 1$ depends on the conjunction. If $z_i = b$ (x_i^b is in the conjunction), then the probability mass assigned is $(1 - 1/k)^{k-1}(1/k)(1 - \rho)^{k-1}\rho$. If $z_i = 1 - b$ (x_i^b is not in the conjunction), then the probability mass assigned is $(1 - 1/k)^k(1 - \rho)^{k-1}\rho$.

Consider the values of a function f on these standard basis strings.

- If $f(e_{i,b}) = 1$ ($x_i^b = 1$) and $z_i = b$ (x_i^b is in the conjunction), f incorrectly computes the conjunction. The contribution to the error is $(1 - 1/k)^{k-1}(1/k)(1 - \rho)^{k-1}\rho$.
- If $f(e_{i,1-b}) = 0$ ($x_i^b = 0$) and $z_i = b$ (x_i^b is in the conjunction), f incorrectly computes the conjunction. The contribution to the error is $(1 - 1/k)^k(1 - \rho)^{k-1}\rho$.

So for every conjunction, a false 0 is roughly k times as costly as a false 1. To make the error less than $(1 - 1/k)^{k-1}(1/k)(1 - \rho)^{k-1}\rho \cdot (99k/100)$, there must be a function in the net that has no false 0's and at most $99k/100$ false 1's on these strings. A function in the net covers the most conjunctions by taking f to be 1 on $k + 99k/100 = 199k/100$ of these strings and 0 on the other $k/100$. A function is covered if its bits are correspond to those with ones. There are $2^{99k/100}$ conjunctions covered, but 2^k conjunctions in total, so any net must have $2^{k/100}$ functions in it to achieve error below $(1 - 1/k)^{k-1}(1/k)(1 - \rho)^{k-1}\rho \cdot (99k/100)$. Taking $\rho = 1/k$, this is at least

$$\begin{aligned} & (1 - 1/k)^{k-1}(1/k)(1 - 1/k)^{k-1}(1/k) \cdot (99k/100) \\ &= 99(1 - 1/k)^{2k-2}/(100k) \\ &\geq 99/(100e^2k) \\ &\geq 1/(8k), \end{aligned}$$

so the error is at least $\rho/8$. We need $\rho < 8\epsilon$ for a sub-exponential size net. $\square$

B PROOF OF THEOREM 1.6

Our main theorem (the formal version of Theorem 1.6) is the following

Theorem B.1. *For any positive integer k and any real numbers $0 < \epsilon, \delta < 1$, $0 < \gamma \leq 1/2$, there exists a randomized algorithm which, with probability at least $1 - \delta$, list-learns k -conjunctions with accuracy $1 - \epsilon$, with sample complexity $\tilde{O}(k^4 \log(1/\delta)/(\epsilon^9 \gamma^4))$ and time complexity $\max\{\tilde{O}(n^2 k^4 \log(1/\delta)/(\epsilon^9 \gamma^4)), O((32k^2/\epsilon^5 \gamma^2)^k)\}$, in the attribute-noise model with bit noise rate $0 \leq \nu_i < \frac{1}{2} - \gamma$ for every $1 \leq i \leq n$, under the assumption that the ground-truth distribution is k' -wise independent for some $k' \geq 2$.*

In the rest of this section, we set $m := 32k^2/(\epsilon^5 \gamma^2)$. Also, by a simple application of Chernoff bound, if we draw $M := O(k^4 \log n \log(1/\delta)/(\epsilon^9 \gamma^4))$ random examples from the noisy example oracle $\tilde{E}X(c, D)$, then with probability at least $1 - \delta$, we can estimate quantities such as $E_{\tilde{D}_1}[x_i]$, $E_{\tilde{D}_1}[x_i \cdot x_j]$ with additive accuracy $O(1/(\epsilon m))$ for every $1 \leq i, j \leq n$. To ease exposition, from now on, we condition our arguments on this event happening.

Since every k' -wise independent distribution for $k' \geq 2$ is also pairwise independent, it is enough to prove Theorem B.1 for $k' = 2$. The algorithm for the pairwise independence test is given by Algorithm 2.

In the rest of this section, we use the notation $\hat{H}$ to denote the estimate of a quantity H using random examples sampled from the noisy example oracle $\tilde{E}X(c, D)$.

First of all, since we include the trivial functions $\mathbf{0}$ and $\mathbf{1}$ in the output list, our learning algorithm succeeds trivially whenever the target concept is ϵ -close to either $\mathbf{0}$ or $\mathbf{1}$. Therefore, from now on, we assume that $\epsilon \leq \Pr_D[c(x) = 1] \leq 1 - \epsilon$.

B.1 Conjunction Bits with Low Label-Sensitivity

The next lemma shows that using bits in S we can get a conjunction which approximates the target concept well.

Lemma B.2. *Let $c = \bigwedge_{i \in c} x_i$ be the target concept, and let c' be the set of bits obtained by removing from c the set of bits eliminated in Line 8 of Algorithm 1. Then conjunction c' is $\epsilon/2$ -close to c , i.e. $\Pr_D[c(x) \neq c'(x)] \leq \epsilon$.*

Proof. First note that eliminating non-conjunction bits can not worsen the performance of our learning algorithm, so we can focus on the effect of eliminating a conjunction bit from S in Line 8.

Since c' is a subset of c ,

$$\begin{aligned} \Pr_D[c(x) \neq c'(x)] &= \Pr_D[c'(x) = 1 \text{ and } \exists i \in c \setminus c' \text{ such that } x_i = 0] \\ &\leq \Pr_D[\exists i \in c \setminus c' \text{ such that } x_i = 0] \\ &\leq \sum_{i \in c \setminus c'} \Pr_D[x_i = 0]. \quad (\text{by union bound}) \end{aligned} \tag{1}$$

We can upper bound $\Pr_D[x_i = 0]$ for any $i \in c \setminus c'$ as

$$\begin{aligned} \Pr_D[x_i = 0] &= \Pr_D[c(x) = 0] \cdot \Pr_{D_0}[x_i = 0] + \Pr_D[c(x) = 1] \cdot \Pr_{D_1}[x_i = 0] \\ &= \Pr_D[c(x) = 0] \cdot \Pr_{D_0}[x_i = 0] \leq \Pr_{D_0}[x_i = 0]. \end{aligned}$$

On the other hand, in terms of quantities over the observed distribution $\tilde{D}$, we have

$$\begin{aligned} \Pr_{\tilde{D}_0}[x_i = 0] &= (1 - \nu_i) \Pr_{D_0}[x_i = 0] + \nu_i \Pr_{D_0}[x_i = 1] \\ &= (1 - \nu_i) \Pr_{D_0}[x_i = 0] + \nu_i (1 - \Pr_{D_0}[x_i = 0]) \\ &= (1 - 2\nu_i) \Pr_{D_0}[x_i = 0] + \nu_i, \end{aligned}$$

and

$$\Pr_{\tilde{D}_1}[x_i = 0] = (1 - \nu_i) \Pr_{D_1}[x_i = 0] + \nu_i \Pr_{D_1}[x_i = 1] = \nu_i \Pr_{D_1}[x_i = 1] \leq \nu_i.$$

Using $O(\log n \log(1/\delta) k^2 / \epsilon^3 \gamma^2) = o(M)$ random examples, we can, with probability at least $1 - \delta$, obtain $\Omega(\log n \log(1/\delta) k^2 / \epsilon^2 \gamma^2)$ random negative examples and $\Omega(\log n \log(1/\delta) k^2 / \epsilon^2 \gamma^2)$ random positive examples, and get an estimate of $\widehat{\text{LS}}_i$ with $|\widehat{\text{LS}}_i - \text{LS}_i| \leq \epsilon \gamma / (2k)$ for every $1 \leq i \leq n$. Since bit- i was eliminated from S , we

$$\text{LS}_i \leq \widehat{\text{LS}}_i + \epsilon \gamma / (2k) < 2\epsilon \gamma / k.$$

Combining this with bounds on $\Pr_{\tilde{D}_0}[x_i = 0]$ and $\Pr_{\tilde{D}_1}[x_i = 0]$, we have

$$2\epsilon \gamma / k > \text{LS}_i = \Pr_{\tilde{D}_0}[x_i = 0] - \Pr_{\tilde{D}_1}[x_i = 0] \geq (1 - 2\nu_i) \Pr_{D_0}[x_i = 0] > 2\gamma \Pr_{D_0}[x_i = 0],$$

where the last step follows from the fact that $\nu_i < \frac{1}{2} - \gamma$. Therefore we have $\Pr_{D_0}[x_i = 0] < \epsilon / k$.

Finally, plugging the above upper bound on $\Pr_{D_0}[x_i = 0]$ into inequality (1) completes the proof. $\square$

B.2 Pairwise Independent Bits

A simple but important observation is that, if the target concept conjunction is $c = \bigwedge_{i \in c} x_i$, then in the observed distribution $\tilde{D}_1$ of positive examples, the bits in c are totally independent. This is because, when restricting to bits in c , D_1 is supported on a single vector 1^k . After applying the (bit-wise independent) attribute noise, $\tilde{D}_1$ is a product distribution when restricting to bits in c .

As it is computationally expensive to check total independence among the conjunction bits on $\tilde{D}_1$, and pairwise independence suffices for our concentration argument, we check pairwise independence in Algorithm 2 by estimating the covariances between each pair of bits.

Lemma B.3. *With probability at least $1 - \delta$, the followings hold: the output S of Algorithm 2 includes every bit in c ; and conversely, every pair of bits X_i and X_j in S are close to being pairwise independent in the sense that $|\text{Cov}_{\tilde{D}_1}(X_i, X_j)| \leq 1/(4\epsilon m)$.*

Claim B.4. Let $D' : \{0, 1\}^n \rightarrow \mathbb{R}^{\geq 0}$ be a distribution and let $X \in \{0, 1\}^n$ be the random variable obtained from sampling according to D' . Then, for any $0 \leq \epsilon \leq 1/2$, if $\Pr[X_i = 1] \leq \epsilon$ for some $1 \leq i \leq n$, then $|\mathbf{Cov}(X_i, X_j)| \leq \epsilon$ for every $i \neq j$. The same bound holds when $\Pr[X_i = 0] \leq \epsilon$.

Proof. Let $p_0 = \Pr[X_i = 0 \wedge X_j = 0]$, $p_1 = \Pr[X_i = 0 \wedge X_j = 1]$, $p_2 = \Pr[X_i = 1 \wedge X_j = 0]$, and $p_3 = \Pr[X_i = 1 \wedge X_j = 1]$. Then $p_2 + p_3 = \Pr[X_i = 1] \leq \epsilon$ and $\mathbf{Cov}(X_i, X_j) = p_3 - (p_2 + p_3)(p_1 + p_3)$. Therefore, $\mathbf{Cov}(X_i, X_j) \geq -(p_2 + p_3)(p_1 + p_3) \geq -(p_2 + p_3) = -\epsilon$. On the other hand, $\mathbf{Cov}(X_i, X_j) \leq p_3 - p_3^2 \leq \epsilon - \epsilon^2 \leq \epsilon$, as $x - x^2$ is increasing for $0 \leq x \leq 1/2$.

The case of $\Pr[X_i = 0] \leq \epsilon$ follows directly from the identity $\mathbf{Cov}(1 - X_i, 1 - X_j) = \mathbf{Cov}(X_i, X_j)$. $\square$

Claim B.5. Let distribution D' and random variable X be the same as in Claim B.4. For any pair of distinct bits i and j , let $\widehat{\mathbf{Cov}}(X_i, X_j) := \widehat{E[X_i \cdot X_j]} - \widehat{E[X_i]} \cdot \widehat{E[X_j]}$ be the estimated covariance of X_i and X_j . Then the estimate error can be upper bounded as

$$|\widehat{\mathbf{Cov}}(X_i, X_j) - \mathbf{Cov}(X_i, X_j)| \leq |\widehat{E[X_i \cdot X_j]} - E[X_i \cdot X_j]| + 2|\widehat{E[X_i]} - E[X_i]| + 2|\widehat{E[X_j]} - E[X_j]|.$$

Proof. Let $\Delta X_i = \widehat{E[X_i]} - E[X_i]$ and $\Delta X_j = \widehat{E[X_j]} - E[X_j]$. Then we have

$$\begin{aligned} \left| \widehat{E[X_i]} \cdot \widehat{E[X_j]} - E[X_i] \cdot E[X_j] \right| &= |\Delta X_i E[X_j] + \Delta X_j E[X_i] + \Delta X_i \Delta X_j| \\ &\leq |\Delta X_i| (E[X_j] + |\Delta X_j|) + |\Delta X_j| (E[X_i] + |\Delta X_i|) \\ &\leq 2|\Delta X_i| + 2|\Delta X_j|, \end{aligned}$$

because both $\widehat{E[X_i]}$ and $E[X_i]$ are real numbers between 0 and 1. Now the bound in the claim follows directly from

$$\begin{aligned} \left| \widehat{\mathbf{Cov}}(X_i, X_j) - \mathbf{Cov}(X_i, X_j) \right| &= \left| \widehat{E[X_i \cdot X_j]} - \widehat{E[X_i]} \cdot \widehat{E[X_j]} - E[X_i \cdot X_j] + E[X_i] \cdot E[X_j] \right| \\ &\leq \left| \widehat{E[X_i \cdot X_j]} - E[X_i \cdot X_j] \right| + \left| \widehat{E[X_i]} \cdot \widehat{E[X_j]} - E[X_i] \cdot E[X_j] \right|. \end{aligned} \quad \square$$

Proof of Lemma B.3. As mentioned earlier, if we draw enough examples from the noisy example oracle, we can estimate quantities such as $E_{\bar{D}_1}[X_i]$ and $E_{\bar{D}_1}[X_i \cdot X_j]$ accurately enough. More specifically, using $O(\log(1/\delta) \log n (\epsilon m)^2 / \epsilon) = \tilde{O}(k^4 \log(1/\delta) / (\epsilon^9 \gamma^4))$ random samples, with probability at least $1 - \delta$, we have $|\widehat{E_{\bar{D}_1}}[X_i] - E_{\bar{D}_1}[X_i]| \leq 1/(48\epsilon m)$ for every $1 \leq i \leq n$ and $|\widehat{E_{\bar{D}_1}}[X_i \cdot X_j] - E_{\bar{D}_1}[X_i \cdot X_j]| \leq 1/(24\epsilon m)$ for every pair of distinct $1 \leq i, j \leq n$. Then for every pair of conjunction bits $i, j \in c$ or a pair of conjunction bit $i \in c$ and a non-conjunction bit $j \in [n] \setminus c$, we always have $\mathbf{Cov}_{\bar{D}_1}(X_i, X_j) = 0$. By Claim B.5, $|\widehat{\mathbf{Cov}}_{\bar{D}_1}(X_i, X_j)| \leq 1/(8\epsilon m)$, so any conjunction bit can never be removed from S in line 12 of Algorithm 2. On the other hand, by Claim B.4 and Claim B.5 and analogous calculations, for any pair of bits X_i and X_j that are in the output S of Algorithm 2, it must be the case that $|\mathbf{Cov}_{\bar{D}_1}(X_i, X_j)| \leq 1/(4\epsilon m)$. $\square$

B.3 Bounding the Size of S

Claim B.6. For every surviving bit X_i in S , we have $E_{\bar{D}_1}[X_i] - E_{\bar{D}}[X_i] > \epsilon^2 \gamma / (2k)$.

Proof. If x_i is in S , then by a similar argument as in the proof of Lemma B.2, $\text{LS}_i \geq \widehat{\text{LS}}_i - \epsilon \gamma / (2k) \geq \epsilon \gamma / (2k)$.

Now, by the definitions of $E_{\tilde{D}}[X_i]$ and $E_{\tilde{D}_1}[X_i]$,

$$\begin{aligned}
 E_{\tilde{D}_1}[X_i] - E_{\tilde{D}}[X_i] &= E_{\tilde{D}_1}[X_i] - (\Pr_{\tilde{D}}[c=0] \cdot E_{\tilde{D}_0}[X_i] + \Pr_{\tilde{D}}[c=1] \cdot E_{\tilde{D}_1}[X_i]) \\
 &= (1 - \Pr_{\tilde{D}}[c=1])(E_{\tilde{D}_1}[X_i] - E_{\tilde{D}_0}[X_i]) \\
 &\geq (1 - \Pr_{\tilde{D}}[c=1]) \frac{\epsilon\gamma}{4k} \\
 &> \epsilon \cdot \frac{\epsilon\gamma}{4k} \quad (\text{since } \Pr_{\tilde{D}}[c=1] = \Pr_{\tilde{D}}[c=1] \leq 1 - \epsilon) \\
 &= \frac{\epsilon^2\gamma}{2k}.
 \end{aligned}$$

□

Lemma B.7. *Suppose the size of S at line 9 in Algorithm 1 is at least m . Then the target concept c is ϵ -close to the all-zero function $\mathbf{0}$.*

Proof. Suppose $|S| \geq m$. Let $S' \subseteq S$ be any subset of S of size exactly m . Without loss of generality, assume that $S' = \{1, \dots, m\}$.

Let X and X^+ be the random variables obtained by sampling from $\{0, 1\}^n$ according to distributions $\tilde{D}$ and $\tilde{D}_1$ respectively. Let random variable $Z(X) := X_1 + \dots + X_m$ and $Z^+(X^+) := X_1^+ + \dots + X_m^+$.

Since D is pairwise independent, then by Claim 3.5, distribution $\tilde{D}$ is pairwise independent as well. Therefore,

$$\mathbf{Var}(Z) = \mathbf{Var}(X_1) + \dots + \mathbf{Var}(X_m) = \sum_{i=1}^m E_{\tilde{D}}[X_i](1 - E_{\tilde{D}}[X_i]) \leq \frac{m}{4}.$$

On the other hand, using the bound on covariances in Lemma B.3, we have

$$\mathbf{Var}(Z^+) = \sum_{i=1}^m \mathbf{Var}(X_i^+) + \sum_{i \neq j} \mathbf{Cov}(X_i^+, X_j^+) < \frac{m}{4} + m^2 \frac{1}{4\epsilon m} \leq \frac{m}{2\epsilon}.$$

Let $\bar{Z} = E_{\tilde{D}}[Z]$ and $\bar{Z}^+ = E_{\tilde{D}_1}[Z^+]$. Then by Claim B.6,

$$\Delta Z := \bar{Z}^+ - \bar{Z} > \frac{\epsilon^2\gamma m}{2k}.$$

Now, by setting $\Delta_1 = \sqrt{\frac{m}{2\epsilon}}$ and applying Chebyshev's inequality to Z , we have

$$\Pr_{\tilde{D}}[Z \geq \bar{Z} + \Delta_1] \leq \Pr[|Z - \bar{Z}| \geq \Delta_1] \leq \frac{\mathbf{Var}(Z)}{\Delta_1^2} \leq \epsilon/2.$$

Similarly, letting $\Delta_2 = \sqrt{\frac{2m}{\epsilon}}$ and applying Chebyshev's inequality to Z^+ yields

$$\Pr_{\tilde{D}_1}[Z^+ \leq \bar{Z}^+ - \Delta_2] \leq 1/4.$$

It is easily checked that $\Delta_1 + \Delta_2 < \frac{\epsilon^2\gamma m}{2k} < \Delta Z$. Therefore,

$$\begin{aligned}
 \epsilon/2 &\geq \Pr_{\tilde{D}}[Z(X) \geq \bar{Z} + \Delta_1] \geq \Pr_{\tilde{D}}[Z(X) \geq \bar{Z}^+ - \Delta_2] \\
 &\geq \Pr_{\tilde{D}}[Z(X) \geq \bar{Z}^+ - \Delta_2 \text{ and } X \text{ is a positive example}] \\
 &= \Pr_{\tilde{D}_1}[Z^+(X^+) \geq \bar{Z}^+ - \Delta_2] \Pr_{\tilde{D}}[c(X) = 1] \\
 &\geq (1 - \frac{1}{4}) \Pr_{\tilde{D}}[c(X) = 1],
 \end{aligned}$$

and hence

$$\Pr_{\tilde{D}}[c(X) = 1] = \Pr_D[c(X) = 1] \leq \frac{\epsilon/2}{1 - 1/4} = \frac{2}{3}\epsilon \leq \epsilon,$$

which completes the proof. $\square$

B.4 Putting Everything Together

Now we are ready to put everything together and prove the correctness of list-learning algorithm, i.e., Theorem B.1.

Proof of Theorem B.1. First of all, the claimed sample complexity of the learning algorithm follows directly from Lemma B.3, and the time complexity bound is due to the fact that we need to estimate, using the random examples, $\widehat{\text{Cov}}_{\tilde{D}_1}(X_i, X_j)$ for every pair $1 \leq i < j \leq n$, and that at the end we may need to output a list of $\binom{m}{\leq k}$ conjunctions.

Next, by Lemma B.3, every conjunction bit passes the Pairwise-Independence-Test and hence in S . Then, by Lemma B.2, filtering out low label-sensitive bits can cause at most an error of ϵ . That is, if we output all $\binom{m}{\leq k}$ conjunctions of size at most k from bits in S , at least one of these is ϵ -close to the target concept $c(x)$.

Finally, Lemma B.7 ensures that when the size of S is large, we can simply output the $\mathbf{0}$ function which is ϵ -close to c . $\square$

C THE TRIVIAL “BEST AGREEMENT” ALGORITHM (INFORMATION-THEORETIC-BOUND VERSION)

A naive algorithm for learning k -conjunctions with attribute noise is to try all $\sum_{i=0}^k 2^i \binom{n}{i} < (2n)^{k+1}$ conjunctions of size at most k and output the one that agrees with examples best.

Theorem C.1. *Given $0 < \epsilon < 1/2$ and assume the noise rate per coordinate is unknown and satisfies $\nu \leq \frac{\epsilon}{2k}$, the naive algorithm that outputs the k -conjunction with maximum agreement with the observed distribution runs in time $O(n^k)$ and with probability $1 - \delta$ outputs a conjunction that is $(1 - \epsilon)$ -close to the conjunction labeling the noisy examples.*

Proof. Let D be the underlying distribution and let $\boldsymbol{\nu} = (\nu_1, \dots, \nu_n)$ be the attribute noise vector with upper bound ν , i.e. $\nu_i \leq \nu$ for every $1 \leq i \leq n$. For ease of exposition, assume that $f(x) = x_1 \wedge \dots \wedge x_k$ is the target concept. For every $x \in \{0, 1\}^n$, let $\tilde{x} = x \oplus \mu$ be the vector obtained from x by adding the attribute noise μ specified by $\boldsymbol{\nu}$. Lastly, let $\hat{X}$ denote the set of noisy examples output by the oracle $\{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_m\}$. Define the *empirical disagreement* of a conjunction g on the sample by

$$\text{disagreement}(g)_{\hat{X}} = \frac{1}{m} \sum_{\tilde{x} \in \hat{X}} I_{g(\tilde{x}) \neq f(x)},$$

where $I_{g(\tilde{x}) \neq f(x)}$ is the indicator random variable of the event that $g(\tilde{x}) \neq f(x)$.

By a Hoeffding bound, it follows that

$$\Pr[|\text{disagreement}(g)_{\hat{X}} - \mathbb{E}_{x, \nu}[\text{disagreement}(g)_{\hat{X}}]| > t] \leq e^{-2mt^2}.$$

Let us calculate $\mathbb{E}_{x, \nu}[\text{disagreement}(g)_{\hat{X}}]$ first when $g = f$, and then when $\text{dist}_D(f, g) > \epsilon$. We will upper bound this quantity when $f = g$ and lower bound it when f and g are ϵ -far. We will show that the minimum disagreement among all ϵ -far functions g is larger than the disagreement of f on the observed set $\hat{X}$, with high probability. Therefore we output an ϵ -close conjunction with high probability $1 - \delta$.

Note that the example oracle generates an example in the following process: first draws a string x according to D , labels it as $f(x)$, then adds the attribute noise which transforms x into $\tilde{x}$. Therefore the example we

see is $(\tilde{x}, f(x))$. But f will predict the label as $f(\tilde{x})$. Hence, the probability that f makes a mistake, i.e., the disagreement between f and the example oracle is

$$\mathbb{E}_{x,\nu}[\text{disagreement}(g)_{\hat{X}}] = \Pr_{D,\nu}[f(x) \neq f(\tilde{x})] \leq \max_x \Pr_{\nu}[f(x) \neq f(\tilde{x})]. \quad (2)$$

Write $x|_{[k]}$ for the k -bit string obtained by projecting x onto index subset $[k]$. Clearly $f(x) = 1$ if and only if $x|_{[k]} = 1^k$. If $f(x) = 0$, then $\Pr_{\nu}[f(x) \neq f(\tilde{x})] = \Pr_{\nu}[\tilde{x}|_{[k]} = 1^k] = \prod_{i \in [k]: x_i=1} (1 - \nu_i) \cdot \prod_{i \in [k]: x_i=0} \nu_i \leq \prod_{i \in [k]} \nu_i \leq 1 - \prod_{i \in [k]} (1 - \nu_i)$, assuming $\nu < 1/2$.

On the other hand, when $f(x) = 1$, then

$$\Pr_{\nu}[f(x) \neq f(\tilde{x})] = \Pr_{\nu}[\tilde{x}|_{[k]} \neq 1^k] = 1 - \prod_{i \in [k]} (1 - \nu_i) \leq 1 - (1 - \nu)^k \leq k\nu.$$

Therefore, $\mathbb{E}_{x,\nu}[\text{disagreement}(g)_{\hat{X}}] \leq k\nu$.

Note that $\Pr_{\nu}[g(x) \neq g(\tilde{x})] \leq k\nu$ holds for any conjunction g of size at most k . Now for any k -conjunction g which is at distance ϵ from f under D , i.e. $\text{dist}_D(f, g) = \epsilon$, we have

$$\begin{aligned} \mathbb{E}_{x,\nu}[\text{disagreement}(g)_{\hat{X}}] &= \sum_x D(x) \Pr_{\nu}[f(x) \neq g(\tilde{x})] \\ &= \sum_{x: f(x)=g(x)} D(x) \Pr_{\nu}[g(x) \neq g(\tilde{x})] + \sum_{x: f(x) \neq g(x)} D(x) \Pr_{\nu}[g(x) = g(\tilde{x})] \\ &\geq \sum_{x: f(x) \neq g(x)} D(x) \Pr_{\nu}[g(x) = g(\tilde{x})] \geq (1 - k\nu) \text{dist}_D(f, g) = (1 - k\nu)\epsilon. \end{aligned}$$

By taking a union bound over all the $O(n^k)$ conjunctions that are ϵ -far from g , it follows that with probability $> 1 - n^k e^{-2mt^2}$ all these conjunctions g are such that

$$\text{disagreement}(g)_{\hat{X}} \geq (1 - k\nu)\epsilon - t.$$

By the above calculations it also follows that f itself satisfies

$$\text{disagreement}(f)_{\hat{X}} \leq k\nu + t.$$

It follows that if we assume that the maximum attribute noise is small enough, e.g. $\nu \leq \frac{\epsilon}{2k}$, $t = \epsilon/8$, $\epsilon < 1/2$ and $n^k e^{-2mt^2} < \delta/2$, then with probability $1 - \delta$ we output a conjunction that is ϵ -close to f , using $m = \Theta(\frac{1}{\epsilon^2}(\log \frac{1}{\delta} + k \log n))$ examples. □

D PROOF OF CLAIM 3.5

Proof. First of all, for any $1 \leq i \leq n$, if we let $p_i := \Pr_D[X_i = 1]$ and $\tilde{p}_i := \Pr_{\tilde{D}}[X_i = 1]$, then

$$\begin{pmatrix} 1 - \tilde{p}_i \\ \tilde{p}_i \end{pmatrix} = \begin{pmatrix} 1 - \nu_i & \nu_i \\ \nu_i & 1 - \nu_i \end{pmatrix} \begin{pmatrix} 1 - p_i \\ p_i \end{pmatrix}.$$

More generally, for any subset of k indices $\{i_1, \dots, i_k\} \subset [n]$,

$$\begin{pmatrix} \Pr_{\tilde{D}}[X_{i_1} \cdots X_{i_k} = 0^k] \\ \vdots \\ \Pr_{\tilde{D}}[X_{i_1} \cdots X_{i_k} = 1^k] \end{pmatrix} = \begin{pmatrix} 1 - \nu_{i_1} & \nu_{i_1} \\ \nu_{i_1} & 1 - \nu_{i_1} \end{pmatrix} \otimes \cdots \otimes \begin{pmatrix} 1 - \nu_{i_k} & \nu_{i_k} \\ \nu_{i_k} & 1 - \nu_{i_k} \end{pmatrix} \begin{pmatrix} \Pr_D[X_{i_1} \cdots X_{i_k} = 0^k] \\ \vdots \\ \Pr_D[X_{i_1} \cdots X_{i_k} = 1^k] \end{pmatrix},$$

where $\otimes$ stands for the Kronecker product of matrices. Now suppose that D is k -wise independent, then

$$\begin{pmatrix} \Pr_D[X_{i_1} \cdots X_{i_k} = 0^k] \\ \vdots \\ \Pr_D[X_{i_1} \cdots X_{i_k} = 1^k] \end{pmatrix} = \begin{pmatrix} 1 - p_{i_1} \\ p_{i_1} \end{pmatrix} \otimes \cdots \otimes \begin{pmatrix} 1 - p_{i_k} \\ p_{i_k} \end{pmatrix},$$

and it follows that

$$\begin{aligned} \begin{pmatrix} \Pr_{\tilde{D}}[X_{i_1} \cdots X_{i_k} = 0^k] \\ \vdots \\ \Pr_{\tilde{D}}[X_{i_1} \cdots X_{i_k} = 1^k] \end{pmatrix} &= \left(\begin{pmatrix} 1 - \nu_{i_1} & \nu_{i_1} \\ \nu_{i_1} & 1 - \nu_{i_1} \end{pmatrix} \begin{pmatrix} 1 - p_{i_1} \\ p_{i_1} \end{pmatrix} \right) \otimes \cdots \otimes \left(\begin{pmatrix} 1 - \nu_{i_k} & \nu_{i_k} \\ \nu_{i_k} & 1 - \nu_{i_k} \end{pmatrix} \begin{pmatrix} 1 - p_{i_k} \\ p_{i_k} \end{pmatrix} \right) \\ &= \begin{pmatrix} 1 - \tilde{p}_{i_1} \\ \tilde{p}_{i_1} \end{pmatrix} \otimes \cdots \otimes \begin{pmatrix} 1 - \tilde{p}_{i_k} \\ \tilde{p}_{i_k} \end{pmatrix}. \end{aligned}$$

That is, $\tilde{D}$ is also k -wise independent. The other direction follow from an identical argument by noting that matrix $\begin{pmatrix} 1 - \nu_i & \nu_i \\ \nu_i & 1 - \nu_i \end{pmatrix}$ is invertible — namely

$$\begin{pmatrix} 1 - \nu_i & \nu_i \\ \nu_i & 1 - \nu_i \end{pmatrix}^{-1} = \begin{pmatrix} \frac{1 - \nu_i}{1 - 2\nu_i} & -\frac{\nu_i}{1 - 2\nu_i} \\ -\frac{\nu_i}{1 - 2\nu_i} & \frac{1 - \nu_i}{1 - 2\nu_i} \end{pmatrix},$$

for every $0 \leq \nu_i < 1/2$. □