

MIT Open Access Articles

Lifting Uniform Learners via Distributional Decomposition

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Blanc, Guy, Lange, Jane, Malik, Ali and Tan, Li-Yang. 2023. "Lifting Uniform Learners via Distributional Decomposition."

As Published: <https://doi.org/10.1145/3564246.3585212>

Publisher: ACM|Proceedings of the 55th Annual ACM Symposium on Theory of Computing

Persistent URL: <https://hdl.handle.net/1721.1/151008>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Lifting Uniform Learners via Distributional Decomposition

Guy Blanc
Stanford University
Stanford, CA, USA

Ali Malik
Stanford University
Stanford, CA, USA

Jane Lange
MIT
Cambridge, MA, USA

Li-Yang Tan
Stanford University
Stanford, CA, USA

ABSTRACT

We show how any PAC learning algorithm that works under the uniform distribution can be transformed, in a blackbox fashion, into one that works under an arbitrary and unknown distribution \mathcal{D} . The efficiency of our transformation scales with the inherent complexity of \mathcal{D} , running in $\text{poly}(n, (md)^d)$ time for distributions over $\{\pm 1\}^n$ whose pmfs are computed by depth- d decision trees, where m is the sample complexity of the original algorithm. For monotone distributions our transformation uses only samples from \mathcal{D} , and for general ones it uses subcube conditioning samples.

A key technical ingredient is an algorithm which, given the aforementioned access to \mathcal{D} , produces an *optimal* decision tree decomposition of \mathcal{D} : an approximation of \mathcal{D} as a mixture of uniform distributions over disjoint subcubes. With this decomposition in hand, we run the uniform-distribution learner on each subcube and combine the hypotheses using the decision tree. This algorithmic decomposition lemma also yields new algorithms for learning decision tree distributions with runtimes that exponentially improve on the prior state of the art—results of independent interest in distribution learning.

CCS CONCEPTS

• **Theory of computation** → **Complexity theory and logic; Oracles and decision trees; Boolean function learning; Query learning; Semi-supervised learning.**

KEYWORDS

PAC Learning, Semi-supervised learning, Decision tree decomposition

ACM Reference Format:

Guy Blanc, Jane Lange, Ali Malik, and Li-Yang Tan. 2023. Lifting Uniform Learners via Distributional Decomposition. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC '23)*, June 20–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3564246.3585212>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '23, June 20–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9913-5/23/06...\$15.00

<https://doi.org/10.1145/3564246.3585212>

1 INTRODUCTION

A major strand of research in learning theory concerns the learning of an unknown target function $f : \{\pm 1\}^n \rightarrow \{0, 1\}$ with respect to an unknown source distribution \mathcal{D} . This line of work has both motivated and benefited from rich connections to the study of function complexity: underlying every new learning algorithm are new structural insights about the concept class of functions that f belongs to. Our work is motivated by the question of whether structural properties of the *source distribution* \mathcal{D} can be similarly leveraged for the design of efficient algorithms.

Distribution-free learning. In Valiant’s original distribution-free PAC learning model [54], the target function f is promised to belong to a *known*, and often “computationally simple” concept class \mathcal{C} , whereas no assumptions are made about the distribution \mathcal{D} .

Unfortunately, efficient algorithms have proven hard to come by even for simple concept classes. Consider, for instance, the class of polynomial-size DNF formulas—the motivating example that Valiant used throughout his paper. Despite decades of effort, the current fastest algorithm for this class runs in exponential time, $\exp(\tilde{O}(n^{1/3}))$ [41]. For other classes, even formal hardness results have been established. For example, under standard cryptographic assumptions it has been shown that there are no efficient algorithms for learning intersections of polynomially many halfspaces [42].

A major source of difficulty in the design of efficient learning algorithms in this model stems from the absence of assumptions about \mathcal{D} . Even if the concept class \mathcal{C} is assumed to be simple, this can be negated by the fact that the underlying distribution \mathcal{D} is arbitrarily complex. Indeed, this is precisely the idea underlying existing hardness results for distribution-free learning.

Distribution-specific learning. This leads us to the *distribution-specific* variant of PAC learning where \mathcal{D} is known rather than arbitrary [7, 45]. For the domain $\{\pm 1\}^n$, the most popular assumption takes \mathcal{D} to be the uniform distribution—an assumption that opens the door to many more positive results. For example, there is a quasipolynomial-time algorithm for learning polynomial-size DNF formulas under the uniform distribution [55]; if the learner is additionally given query access to f , there are even polynomial-time algorithms [30, 36]. Similarly, there is a quasipolynomial-time algorithm for learning intersections of polynomially many halfspaces under the uniform distribution [40]—a sharp contrast to the distribution-free setting where even learning intersections of *two* halfspaces in subexponential time is a longstanding open problem [39, 52]. Indeed, there is by now a suite of powerful techniques for the design of uniform-distribution learning algorithms, notably

“meta-algorithms” for learning an arbitrary concept class \mathcal{C} with performance guarantees that scale with its Fourier properties. Famous examples include the algorithms for learning (approximately) low-degree functions [44] and (approximately) sparse functions [43], as well as their agnostic variants [30, 37]; see Chapter §3 of [46] for an in-depth treatment.

While this setting has proved to be fertile grounds for the development of sophisticated techniques and fast algorithms, it is admittedly a stylized theoretical model, since real-world data distributions often exhibit correlations and are not uniform. For this reason, uniform-distribution learners have thus far been limited in their practical relevance.

This work. These two settings, distribution-free and uniform-distribution learning, correspond to two extremes in terms of distributional assumptions. Distribution-free algorithms make no assumptions about the distribution, but the design of such algorithms that are computationally efficient has proven correspondingly challenging. On the other hand, we have a wealth of techniques and positive results in the uniform-distribution setting, but the guarantees of such algorithms only hold under a strong, often unrealistic assumption on the data distribution. Due to these vast differences, research in these two settings has mostly proceeded independently with few known connections.

The motivation for our work is the search for fruitful middle grounds between these two extremes, where efficient algorithms can be obtained for expressive concept classes and where their guarantees hold for broad classes of distributions. More generally, we ask:

Can we interpolate between the two extremes via notions of distribution complexity, and design learning algorithms whose efficiency scale with the inherent complexity of \mathcal{D} ?

Relatedly, can the large body of existing results and techniques for learning under the uniform distribution be lifted, ideally in a blackbox fashion, to non-uniform but nevertheless “simple” distributions?

2 OUR RESULTS

We provide affirmative answers to these questions via a natural notion of distribution complexity:

DEFINITION 1 (DECISION TREE COMPLEXITY OF \mathcal{D}). *The decision tree complexity of a distribution \mathcal{D} over $\{\pm 1\}^n$, denoted $\text{depth}(\mathcal{D})$, is the smallest integer d such that its probability mass function (pmf) can be computed by a depth- d decision tree.*

Decision tree complexity interpolates between the uniform distribution on one extreme (depth 0) and arbitrary distributions on the other (depth n).

This notion, first considered by Feldman, O’Donnell, and Seredio [26], generalizes other well-studied notions of complexity of high-dimensional discrete distributions and is itself a special case of others. A d - *junta distribution* \mathcal{D} , introduced by Aliakbarpour, Blais, and Rubinfeld [3], is one for which there exists a set of d coordinates $J \subseteq [n]$ such that for every assignment $\rho \in \{\pm 1\}^J$, the conditional distribution $\mathcal{D}_{J \leftarrow \rho}$ is uniform. Every d -junta distribution has decision tree complexity d , but a depth- d decision tree

distribution can have junta complexity as large as 2^d . On the other hand, decision tree distributions are a special case of mixtures of subcubes [20], since a depth- d decision tree induces a partition of $\{\pm 1\}^n$ into 2^d many disjoint subcubes. Mixtures of subcubes are in turn a special case of mixtures of product distributions over discrete domains [22, 23, 26, 28].

2.1 First Main Result: Lifting Uniform-Distribution Learners

Our algorithms and analyses will use the notion of *influence* of variables. This notion is most commonly applied to boolean-valued functions, and in this work we extend its study to the pmfs of distributions:

DEFINITION 2 (INFLUENCE OF VARIABLES ON DISTRIBUTIONS). *Let \mathcal{D} be a distribution over $\{\pm 1\}^n$ and $f_{\mathcal{D}}(x) = 2^n \cdot \mathcal{D}(x)$ be its pmf scaled up by the domain size.¹ The influence of a coordinate $i \in [n]$ on a distribution \mathcal{D} over $\{\pm 1\}^n$ is the quantity*

$$\text{Inf}_i(f_{\mathcal{D}}) := \mathbb{E}_{x \sim \mathcal{U}^n} [|f_{\mathcal{D}}(x) - f_{\mathcal{D}}(x^{-i})|],$$

where \mathcal{U}^n denotes the uniform distribution over $\{\pm 1\}^n$ and x^{-i} denotes x with its i -th coordinate rerandomized. The total influence of \mathcal{D} is the quantity $\text{Inf}(f_{\mathcal{D}}) := \sum_{i=1}^n \text{Inf}_i(f_{\mathcal{D}})$.

With **Definitions 1** and **2** in hand we can now state our first main result. For clarity, we first state it assuming a unit-time oracle that computes the influences of variables:

THEOREM 1 (LIFTING UNIFORM-DISTRIBUTION LEARNERS; SEE **THEOREM 5 FOR THE FORMAL VERSION).** *For any concept class \mathcal{C} of functions $f : \{\pm 1\}^n \rightarrow \{0, 1\}$ closed under restrictions, assuming a unit-time influence oracle, if there is an algorithm for learning \mathcal{C} under the uniform distribution to accuracy ϵ with sample complexity m and running time $\text{poly}(n, m)$, there is an algorithm for learning \mathcal{C} under depth- d decision tree distributions with sample complexity and running time*

$$M = \text{poly}(n) \cdot \left(\frac{dm}{\epsilon} \right)^{O(d)}.$$

As an example setting of parameters, **Theorem 1** lifts a quasipolynomial-time uniform-distribution algorithm into one that still runs in quasipolynomial time, but now succeeds under any distribution with decision tree complexity $\text{polylog}(n)$. We note that such distributions are quite broad: the size of a depth- d tree can be as large as $2^d = \text{quasipoly}(n)$, corresponding to the mixture of that many subcubes.

The proof of **Theorem 1** readily extends to lift *agnostic* uniform-distribution learners [35, 38] to agnostic distribution-free ones. As mentioned in the introduction, many algorithms for learning under the uniform distribution are obtained through Fourier-analytic meta-algorithms. By **Theorem 1**, we now have analogues of these meta-algorithms for distributions with low decision tree complexity.

¹This 2^n normalisation factor makes the average value of $f_{\mathcal{D}}$ exactly 1, lending to a cleaner analysis.

Estimating influences efficiently. We have stated [Theorem 1](#) assuming that influences of variables on distributions can be computed exactly in unit time. In the body of the paper we show how these quantities can be approximated to sufficiently high accuracy given access to \mathcal{D} . For monotone distributions we show how this can be done using only samples from \mathcal{D} . For general distributions, we use the notion of *subcube conditioning samples*, proposed in [8, 18] and subsequently studied in [17, 21]: in this model, the algorithm specifies a subcube of $\{\pm 1\}^n$ and receives a draw $x \sim \mathcal{D}$ conditioned on x lying in the subcube. (For general distributions, estimating influences based only on samples is intractable as it can be easily seen to require $\Omega(\sqrt{2^n})$ many samples.)

A more general result. We obtain [Theorem 1](#) as a corollary of a more general result which shows how every uniform-distribution learner \mathcal{A} that is robust to *distributional noise* can be lifted in a way that the resulting runtime depends only on \mathcal{A} 's noise tolerance and not its sample complexity. [Theorem 1](#) follows since every m -sample algorithm is automatically tolerant to an $O(1/m)$ amount of distributional noise. We achieve better parameters for algorithms that are tolerant to higher amounts of noise.

2.2 Second Main Result: Learning Decision Tree Distributions

Our algorithm for [Theorem 1](#) proceeds in a two-stage manner: we first learn the decision tree structure of \mathcal{D} and then use the uniform-distribution learner to learn f restricted to each of the leaves of the tree. To carry out the first stage, we give an algorithm that learns the *optimal* decision tree decomposition of a distribution \mathcal{D} :

THEOREM 2 (LEARNING DECISION TREE DISTRIBUTIONS). *Let \mathcal{D} be a distribution that is representable by a depth- d decision tree. There is an algorithm that returns a depth- d tree representing a distribution \mathcal{D}' such that $\text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}') \leq \epsilon$, with high probability over the draw of samples, and with running time and sample complexity $\text{poly}(n) \cdot (d/\epsilon)^{O(d)}$. For monotone distributions, the algorithm only uses random samples from \mathcal{D} , and for general distributions, it uses subcube conditional samples.*

[Theorem 2](#) is a result of independent interest in distribution learning. Even setting aside properness, the performance guarantees of our algorithm improves, quite dramatically, the prior state of the art for learning decision tree distributions and circumvents existing hardness results. Aliakbarpour, Blais, and Rubinfeld [3] gave an $n^{O(k)}$ time algorithm for learning k -junta distributions, which implies an $n^{O(2^d)}$ time algorithm for learning depth- d decision tree distributions. Chen and Moitra [20] gave an $s^{s^3} \cdot n^{O(\log s)}$ time algorithm for learning the mixture of s subcubes, which implies a $2^{2^{O(d)}} \cdot n^d$ time algorithm for learning depth- d decision tree distributions. Finally, Feldman, O'Donnell, and Seredio [26] gave an $n^{O(m^3)}$ time algorithm for learning the mixture of m product distributions, which implies an $n^{2^{O(d)}}$ time algorithm for learning depth- d decision tree distributions. These runtimes all have a doubly-exponential dependence on d , whereas ours only depends exponentially on d . Note that the runtime of any algorithm must have at least an exponential dependence on d since that is the description length of a depth- d decision tree.

None of these prior algorithms are proper, and that fact that ours is is crucial to the application to [Theorem 1](#): the decision tree structure specifies a decomposition of $\{\pm 1\}^n$ into disjoint subcubes, and it is on these subcubes that we run our uniform-distribution learner.

Circumventing hardness results. A novel aspect of [Theorem 2](#) is that it sidesteps existing hardness results for learning decision tree distributions. [26] showed that the problem of learning depth- d decision tree distributions is as hard as that of learning depth- d decision tree functions under the uniform distribution. Despite significant efforts for over three decades, the current fastest algorithm for the latter problem runs in time $n^{O(d)}$ [24] and improving on this is a longstanding challenge of learning theory. (It contains as a special case the junta problem [13]—learning k -juntas in time better than $n^{O(k)}$ —itself already a notorious open problem.) [Theorem 2](#) shows that this barrier can be circumvented in two different ways: by giving the algorithm access to subcube conditional samples, and by considering monotone distributions.

Proof Overview of [Theorem 2](#). To describe the intuition behind [Theorem 2](#) and the role that influence plays in its proof, we begin by considering the following elementary equations:

$$\begin{aligned} \mathbb{E}_{b \sim \{\pm 1\}} [\text{Inf}((f_{\mathcal{D}})_{x_i=b})] &= \text{Inf}(f_{\mathcal{D}}) - \text{Inf}_i(f_{\mathcal{D}}) \quad (1) \\ 2 \cdot \text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{U}) &\leq \text{Inf}(f_{\mathcal{D}}), \quad (2) \end{aligned}$$

where $(f_{\mathcal{D}})_{x_i=b}$ denotes the restriction of $f_{\mathcal{D}}$ by fixing x_i to b . [Equation \(1\)](#), which follows from the definition of influence, says that the total influence of $f_{\mathcal{D}}$ drops by $\text{Inf}_i(f)$ when restricted by x_i . [Equation \(2\)](#), which is a consequence of the Efron–Stein inequality, says that the total influence of a distribution upper bounds its distance to uniformity.

Together, they suggest a simple and natural algorithm for learning decision tree distributions: build a decision tree hypothesis for \mathcal{D} greedily by iteratively querying the most influential variable. After sufficiently many stages, the total influence of the conditional distributions at most leaves will be close to zero (by [Equation \(1\)](#)), which in turn means that most leaves will be close to uniform ([Equation \(2\)](#)). Indeed, this intuition can be formalized using the techniques in this work to get an algorithm that learns depth- d decision tree distributions with depth- $O(d^2)$ decision tree hypotheses in time $2^{O(d^2)}$.

To obtain the improved parameters of [Theorem 2](#), we consider a generalization of this different algorithm: instead of splitting on the single most influential variable, we consider all $O(d)$ most influential ones as candidate splits. While this involves searching over more candidates at each split, we will show that at least one of the choices leads to a high accuracy hypothesis at depth d instead of $O(d^2)$, resulting in a smaller search space of $d^{O(d)}$ instead of $2^{O(d^2)}$.

Our approach is inspired by a recent algorithm of Blanc, Lange, Qiao, and Tan for properly learning decision tree functions under the uniform distribution [12]. Our analysis builds on and extends theirs to the setting of unsupervised learning, which poses a number of challenges that we have to overcome. First, while highly accurate estimates of influences can be easily obtained with membership queries to the function (in our case, the pmf of \mathcal{D}), subcube

conditioning samples provide more limited and coarse-grained information about \mathcal{D} . Our algorithms for estimating influences with subcube conditioning samples, and from samples alone for monotone distributions, could see further utility in other problems. Second, while it is easy to estimate how close a *function* is to a constant (relatedly, how close two functions are) via random sampling, the analogous problem of estimating the distance of *distribution* to uniformity is an intractable problem: for a distribution over $\{\pm 1\}^n$, the sample complexity of estimating its distance to uniformity is $\Theta(2^n/n)$ [53]. There are no known improvements using subcube conditioning samples (though [8, 17] give efficient uniformity *testers*), and the best known algorithm for monotone distributions uses $2^{n-\Theta(\sqrt{n}\log n)}$ samples [50]. We sidestep this barrier by showing how total influence—for which we provide efficient estimators—can be used as a good proxy for distance to uniformity.

2.3 Other Related Work

Conditional samples in distribution learning and testing. The subcube conditioning model falls within a recent line of work on the power of *conditional samples* in distribution learning and testing. In this more general model, which was independently introduced by Chakraborty, Fischer, Goldhirsh, and Matsliah [19] and Canonne, Ron, and Servedio [18], the algorithm can specify an arbitrary subset of the domain and receive a sample conditioned on falling within this subset. Since its introduction, a large number of works have designed conditional sample algorithms, in distribution learning and testing [2, 9, 16, 17, 21, 25, 27, 51] and beyond [1, 32, 33]. Our results add to this line of work, and further reinforce the message that conditional samples (indeed, even just subcube conditional samples) can be used circumvent sample complexity lower bounds in a variety of settings.

Other access models to distributions include queries to the pmf or cdf (the *evaluation oracle* model) [5, 15, 34, 48] and giving the algorithm probability revealing samples [47].

Semi-supervised learning. There is extensive research in the statistical machine learning literature on leveraging unlabeled examples to improve learning. Much of this work focuses on improving sample complexity or convergence rates of existing algorithms using additional unlabelled data [6, 31]. In contrast, our work is aimed at creating *computational and sample efficient* algorithms from existing ones that are only guaranteed to work under “nice” distributions. Results of this flavour have been studied in limited ways, e.g. transforming a 1d-uniform learning algorithm to one that works on any 1d continuous distribution [6].

The work of [10]. In [10], Blais, O’Donnell, and Wimmer gave an algorithm for performing polynomial regression under arbitrary product distributions over $\{\pm 1\}^n$. As an application, they showed how their algorithm can be lifted to *mixtures* of product distributions via the algorithm of [26] for learning mixtures of product distributions.

Like our work, this is also an example where algorithms for learning with respect to a “simple” distribution (the uniform distribution in our case and product distributions in [10]’s case) can be lifted to more complex ones (decision tree distributions in our case and mixtures of product distributions in [10]’s case), via a distribution

learning algorithm that decomposes the more complex one into simple ones (Theorem 2 in our case and [26]’s algorithm in [10]’s case). The quantitative details of our transformations are incomparable: our algorithm for learning depth- d decision tree distributions run in $\text{poly}(n) \cdot d^{O(d)}$ time, whereas [26]’s algorithm for learning the mixture of m product distributions run in $n^{O(m^3)}$ time. (Recall that a depth- d decision tree induces a mixture of as many as $m = 2^d$ product distributions.)

3 DISCUSSION AND FUTURE WORK

We view our work as part of two broader and potentially fruitful approaches to PAC learning. Much of the progress in the field thus far has been guided by the design of efficient learning algorithms for successively more expressive concept classes, as measured according to various notions of function complexity; this was the approach advocated in Valiant’s pioneering paper and other early works. For example, on one such axis we have small-width conjunctions as a special case of small juntas, which are in turn a special case of small-depth decision trees, which are in turn a special case of small-width DNFs, and so on, and the field seeks to design efficient algorithms for each of these classes. We believe that it is equally natural to make progress along a separate dimension, with respect to various notions of distribution complexity. The overall goal can then be cast as that of learning successively more expressive concept classes with respect to successively more expressive distributions.

Next, our algorithm is just one instantiation of a general two-stage approach to learning that is studied in the semi-supervised literature (see e.g. [4]): the first gathers information about the underlying distribution, and the second exploits this distributional information to learn the target function. It would be interesting to develop more computationally efficient examples of such a two-stage approach. More broadly, there should be much to be gained from using the insights of distribution learning to counter the difficulty of distribution-free PAC learning, the crux of which is the potential nastiness of the unknown distribution.

Finally, looking beyond PAC learning, a similar gulf exists between uniform-distribution and distribution-free *testing* of function properties. The original model of property testing was defined with respect to the uniform distribution [29, 49] and much of the ensuing research has focused on this setting, with the distribution-free variant receiving increasing attention in recent years. Can uniform-distribution testers be lifted generically to the distribution-free setting? A concrete avenue towards such a result would be via a variant of our distribution decomposition lemma that runs in sublinear time.

4 PRELIMINARIES

Notation. Given an input $x \in \{\pm 1\}^n$, coordinate $i \in [n]$, and setting $b \in \{\pm 1\}$, we use $x_{i=b}$ to refer to the input x with the i^{th} coordinate overwritten to take the value b . Similarly, given a sequence of (coordinate, value) pairs $\pi = \{(i_1, b_1), \dots, (i_k, b_k)\}$, we use x_π to represent x with the coordinates in π overwritten/inserted with their respective values.

Given a function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$, we denote the *restriction* $f_{i=b} : \{\pm 1\}^n \rightarrow \mathbb{R}$ to be the function that maps x to $f(x_{i=b})$. We define the restriction f_{π} analogously.

DEFINITION 3 (DECISION TREES (DT)). A *decision trees* $T : \{\pm 1\}^n \rightarrow \mathbb{R}$, is a binary tree whose internal nodes query a particular coordinate, and whose leaves are labelled by values. Each instance $x \in \{\pm 1\}^n$ follows a unique root-to-leaf path in T : at any internal node, it follows either the left or right branch depending on the value of the queried coordinate, until a leaf is reached and its value is returned.

The set of leaves $\ell \in \text{leaves}(T)$ therefore form a partition of $\{\pm 1\}^n$, with each leaf having $2^{n-|\ell|}$ elements, where $|\ell|$ is the depth of the leaf. Every leaf ℓ also corresponds to a sequence of (coordinates, value) pairs $\pi(\ell)$ that lead to the leaf. For a function f , will sometimes use the shorthand f_{ℓ} to mean the restriction $f_{\pi(\ell)}$.

DEFINITION 4 (DECISION TREE DISTRIBUTION). We say that a distribution $\mathcal{D} : \{\pm 1\}^n \rightarrow [0, 1]$ is representable by a depth- d DT, if its pmf is computable by a depth- d decision tree T . Specifically, each leaf ℓ is labelled by a value p_{ℓ} , so that $\mathcal{D}(x) = p_{\ell}$ for all $x \in \ell$. This means that the conditional distribution of all points that reach a leaf is uniform. Moreover, since \mathcal{D} is a distribution, we have: $\sum_{\ell \in \text{leaves}(T)} 2^{n-|\ell|} \cdot p_{\ell} = 1$.

For a given leaf ℓ , we will write $\mathcal{D}_{\ell} : \{\pm 1\}^{n-|\ell|} \rightarrow [0, 1]$ to represent the conditional distribution of \mathcal{D} at the leaf ℓ , so that for any $x \in \{\pm 1\}^{n-|\ell|}$, we have $\mathcal{D}_{\ell}(x) = \mathcal{D}(x_{\pi(\ell)}) / \Pr_{y \sim \mathcal{D}}[y \in \ell]$.

We will often scale up the pmfs of our distributions by the domain size, since it makes our analysis easier. As such, we also define the weighting function:

DEFINITION 5 (WEIGHTING FUNCTION OF DISTRIBUTION). Let \mathcal{E} be an arbitrary distribution over $\{\pm 1\}^m$. We define the weighting function:

$$f_{\mathcal{E}}(x) := 2^m \cdot \mathcal{E}(x).$$

DEFINITION 6 (MONOTONE DISTRIBUTION). We furthermore say that a distribution \mathcal{D} is monotone if its pmf is monotone: for $x, y \in \{\pm 1\}^n$, if $x_i \leq y_i$ for all $i \in [n]$, then $\mathcal{D}(x) \leq \mathcal{D}(y)$.

DEFINITION 7 (TV DISTANCE). For two distributions \mathcal{P}, \mathcal{Q} over a countable domain \mathcal{X} , we define the total variation distance:

$$\text{dist}_{\text{TV}}(\mathcal{P}, \mathcal{Q}) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathcal{P}(x) - \mathcal{Q}(x)| = \frac{1}{2} \|\mathcal{P} - \mathcal{Q}\|_1.$$

DEFINITION 8 (ℓ_1 INFLUENCE). For any function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$, the influence of the i -th variable on f is given by:

$$\text{Inf}_i(f) := \mathbb{E}_{x \sim \mathcal{U}^n} [|f(x) - f(x^{\oplus i})|],$$

where $x^{\oplus i}$ denotes x with the i -th coordinate re-randomised. Note that the influence of a function is defined with respect to the uniform distribution over its domain.

We further define the total influence as the sum of influences over all variables: $\text{Inf}(f) := \sum_{i=1}^n \text{Inf}_i(f)$.

FACT 4.1 (INFLUENCE \equiv CORRELATION FOR MONOTONE FUNCTIONS). Let $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ be a monotone function. Then

$$\text{Inf}_i(f) = \mathbb{E}_{x \sim \mathcal{U}^n} [f(x) \cdot x_i].$$

DEFINITION 9 (ℓ_1 VARIANCE). For any function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$,

$$\text{Var}^{(1)}(f) := \mathbb{E}_{x, y \sim \mathcal{U}^n} |f(x) - f(y)|.$$

We will also sometimes use a different definition of variance, given by the mean absolute deviation of f :

$$\text{Var}_{\mu}(f) := \mathbb{E}_{x \sim \mathcal{U}^n} |f(x) - \mathbb{E}[f]|.$$

These two definitions are equivalent, up to constant factors:

LEMMA 4.2. For a function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$,

$$\text{Var}_{\mu}(f) \leq \text{Var}^{(1)}(f) \leq 2 \text{Var}_{\mu}(f)$$

PROOF. The second part follows immediately from the triangle inequality and the first is an application of Jensen's:

$$\text{Var}_{\mu}(f) = \mathbb{E}_{x \sim \mathcal{U}} \left[\left| \mathbb{E}_{y \sim \mathcal{U}} [f(x) - f(y)] \right| \right] \leq \mathbb{E}_{x, y \sim \mathcal{U}} |f(x) - f(y)|. \quad \square$$

DEFINITION 10 (SENSITIVITY). For a function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ and $x \in \{\pm 1\}^n$, the sensitivity of f at x is defined to be

$$s(f, x) = \sum_{i=1}^n \mathbb{1}[f(x) \neq f(x^{\oplus i})].$$

Furthermore, the sensitivity of f is given by its maximum sensitivity over all points:

$$s(f) = \max_{x \in \{\pm 1\}^n} \{s(f, x)\}.$$

Note that the sensitivity of a decision tree is at most the depth of the decision tree, since any point can only be sensitive to the coordinates queried on its root-to-leaf path.

Useful inequalities. We present some useful inequalities for boolean functions.

LEMMA 4.3 (EFRON-STEIN). For any function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$:

$$\text{Var}^{(1)}(f) \leq \text{Inf}(f).$$

LEMMA 4.4 (TOTAL INFLUENCE AND SENSITIVITY). For any function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$:

$$\text{Inf}(f) \leq 2s(f) \cdot \text{Var}^{(1)}(f).$$

PROOF. Let $s = s(f)$ be the sensitivity of f and consider the set, $\text{snbr}(x) = \{i \in [n] \mid f(x) \neq f(x^{\oplus i})\}$. By assumption, $|\text{snbr}(x)| \leq s$. We define a coupling $(x, y) \sim \pi$ s.t. y is often in $\text{snbr}(x)$, but the marginal distributions $\pi(x)$ and $\pi(y)$ are still uniform. First, sample $x \sim \mathcal{U}$. Then, sample y given x as follows: for each $i \in \text{snbr}(x)$, let $y = x^{\oplus i}$ (i.e. flip the i -th coordinate of x) with probability $1/s$, and with the remaining $1 - |\text{snbr}(x)|/s$ probability, take $y = x$. It is easy to see that the marginal distribution over y is still uniform.

Unrolling the definition of influence, we have:

$$\begin{aligned}
\text{Inf}(f) &= \sum_{i=1}^n \mathbb{E}_{\mathbf{x} \sim \mathcal{U}} |f(\mathbf{x}) - f(\mathbf{x}^{\oplus i})| \\
&= \mathbb{E}_{\mathbf{x} \sim \mathcal{U}} \left[\sum_{i=1}^n |f(\mathbf{x}) - f(\mathbf{x}^{\oplus i})| \right] \\
&= \mathbb{E}_{\mathbf{x} \sim \mathcal{U}} \left[\sum_{i \in \text{snbr}(\mathbf{x})} |f(\mathbf{x}) - f(\mathbf{x}^{\oplus i})| \right] \\
&\quad \text{(only consider nonzero terms)} \\
&= \mathbb{E}_{\mathbf{x} \sim \mathcal{U}} \left[s \cdot \sum_{i \in \text{snbr}(\mathbf{x})} \frac{|f(\mathbf{x}) - f(\mathbf{x}^{\oplus i})|}{s} \right] \\
&= \mathbb{E}_{\mathbf{x} \sim \pi} \left[s \cdot \mathbb{E}_{\mathbf{y} \sim \pi(\cdot | \mathbf{x})} |f(\mathbf{x}) - f(\mathbf{y})| \right] \\
&\quad \text{(definition of coupling } \pi) \\
&= s \cdot \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \pi} |f(\mathbf{x}) - f(\mathbf{y})| \\
&\leq s \cdot \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \pi} |f(\mathbf{x}) - \mathbb{E}[f]| + s \cdot \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \pi} |f(\mathbf{y}) - \mathbb{E}[f]| \\
&\quad \text{(triangle inequality)} \\
&= 2s \cdot \text{Var}_{\mu}(f) \quad \text{(marginal distributions of } \pi \text{ are uniform)} \\
&\leq 2s \cdot \text{Var}^{(1)}(f). \quad \text{(Lemma 4.2)}
\end{aligned}$$

□

5 OUR ALGORITHMIC DECOMPOSITION LEMMA

Here we present an algorithm that constructs a decision tree of depth d for a distribution \mathcal{D} , and analyze its correctness and complexity. Throughout this section, we assume access to an oracle that gives the exact influences of variables in $f_{\mathcal{D}}$ or any of its restrictions. In the next sections we will show that the influences can be estimated from random examples for monotone distributions, and from subcube conditional examples for general distributions.

THEOREM 3 (LEARNING DECISION TREE DISTRIBUTIONS). *Let \mathcal{D} be a distribution that is representable by a depth- d decision tree. The algorithm BUILDDT returns a depth- d tree representing a distribution \mathcal{D}' such that $\text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}') \leq \varepsilon$ w.h.p.. Given access to a unit time influence oracle, its running time is $n \cdot (d/\varepsilon)^{O(d)}$.*

The algorithm BUILDDT is an exhaustive search over a subset of depth- d decision trees. We characterize this subset as follows:

DEFINITION 11 (EVERYWHERE τ -INFLUENTIAL). *Let T be a tree and v be an internal node with root variable $i(v)$. T is everywhere τ -influential with respect to some f if for every $v \in T$, we have $\text{Inf}_{i(v)}(f_v) \geq \tau$.*

Correctness. Here we show that under the oracle assumptions described above, BUILDDT returns a tree within TV distance ε . The proof will rely on the following fact, which relates TV distance to the uniform ℓ_1 error of the tree with respect to $f_{\mathcal{D}}$.

FACT 5.1 (TV DISTANCE = LABEL ERROR).

$$\begin{aligned}
\text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}') &= \frac{1}{2} \cdot \|\mathcal{D} - \mathcal{D}'\|_1 \\
&= 2^{-(n+1)} \cdot \|2^n \mathcal{D} - 2^n \mathcal{D}'\|_1 \\
&= 2^{-(n+1)} \cdot \|f_{\mathcal{D}} - T'\|_1.
\end{aligned}$$

First, we will show that BUILDDT outputs a decision tree T' with small average influence at the leaves. Then, we will show that this implies that the uniform ℓ_1 error of T' with respect to $f_{\mathcal{D}}$ is small. Correctness follows from the equivalence between $2^{-(n+1)} \|f_{\mathcal{D}} - T'\|_1$ and $\text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}')$.

The claim that BUILDDT outputs a decision tree T' with small average influence at the leaves extends a lemma from [12], instantiated here for the metric space \mathbb{R} equipped with the ℓ_1 -norm:

LEMMA 5.2 (THEOREM 5 OF [12]). *Let $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ be representable by a depth- d DT T . Then there exists T^* such that the following are satisfied:*

- (1) *The size and depth of T^* are at most the size and depth of T ,*
- (2) *T^* is everywhere τ -influential with respect to f ,*
- (3) *$2^{-n} \cdot \|f - T^*\|_1 \leq d\tau$.*

In our BUILDDT, we cannot compute $\|f_{\mathcal{D}} - T'\|_1$ and hence cannot search for trees that minimise this error. Instead, we find trees that minimise the expected total influence at the leaves. The following lemma relates these two values:

LEMMA 5.3 (EXPECTED TOTAL INFLUENCE AND ℓ_1 ERROR). *Let $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ be representable by a depth- d DT, and let T' be any other DT. Then:*

$$\mathbb{E}_{\ell \in T'} [\text{Inf}(f_{\ell})] \leq 4d \cdot 2^{-n} \|f - T'\|_1$$

PROOF. Since f is representable by a depth- d decision tree, its maximum sensitivity (and the sensitivity of each of its leaf restrictions) must be at most d . Therefore, for any leaf $\ell \in T'$, Lemma 4.4 asserts that $\text{Inf}(f_{\ell}) \leq 2d \cdot \text{Var}^{(1)}(f_{\ell})$. Moreover,

$$\begin{aligned}
\text{Var}^{(1)}(f_{\ell}) &= \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \mathcal{U}^n} |f_{\ell}(\mathbf{x}) - f_{\ell}(\mathbf{y})| \\
&= \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \mathcal{U}^n} |f_{\ell}(\mathbf{x}) - T'_{\ell} + T'_{\ell} - f_{\ell}(\mathbf{y})| \\
&\leq 2 \cdot \mathbb{E}_{\mathbf{x} \sim \mathcal{U}^n} |f_{\ell}(\mathbf{x}) - T'_{\ell}| \quad \text{(Triangle ineq.)} \\
&= 2 \cdot \mathbb{E}_{\mathbf{x} \sim \mathcal{U}^n} [|f(\mathbf{x}) - T'(\mathbf{x})| \mid \mathbf{x} \in \ell].
\end{aligned}$$

Therefore,

$$\begin{aligned}
\mathbb{E}_{\ell \in T'} [\text{Inf}(f_{\ell})] &\leq 2d \cdot \mathbb{E}_{\ell \in T'} [\text{Var}^{(1)}(f_{\ell})] \\
&\leq 4d \cdot \mathbb{E}_{\ell \in T'} \mathbb{E}_{\mathbf{x} \sim \mathcal{U}^n} [|f(\mathbf{x}) - T'(\mathbf{x})| \mid \mathbf{x} \in \ell] \\
&= 4d \cdot \mathbb{E}_{\mathbf{x} \sim \mathcal{U}^n} |f(\mathbf{x}) - T'(\mathbf{x})| \\
&= 4d \cdot 2^{-n} \cdot \|f(\mathbf{x}) - T'(\mathbf{x})\|_1. \quad \square
\end{aligned}$$

As a corollary of Lemma 5.2 and Lemma 5.3, we get our pruning lemma stated in terms of influences:

BUILDDT($\mathcal{D}, \pi, d, \tau$):

Input: Random examples from \mathcal{D} , restriction π , influence oracle for $(f_{\mathcal{D}})_{\pi}$, depth parameter d , influence parameter τ .

Output: A decision tree T that minimizes $\mathbb{E}_{\ell \in T} [\text{Inf}((f_{\mathcal{D}})_{\ell})]$ among all depth- d , everywhere τ -influential trees.

- (1) Let $S \subseteq [n]$ be the set of variables i such that $\text{Inf}_i((f_{\mathcal{D}})_{\pi}) \geq \tau$.
- (2) If S is empty or $d = 0$, return the leaf labeled with $2^{\lfloor \pi \rfloor} \cdot \Pr_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x}$ is consistent with $\pi]$.
- (3) Otherwise:
 - (a) For each $i \in S$, let T_i be the tree such that

$$\begin{aligned} \text{root}(T_i) &= x_i \\ \text{left-subtree}(T_i) &= \text{BUILDDT}(\mathcal{D}, \pi \cup \{x_i = -1\}, d - 1, \tau) \\ \text{right-subtree}(T_i) &= \text{BUILDDT}(\mathcal{D}, \pi \cup \{x_i = 1\}, d - 1, \tau) \end{aligned}$$

- (b) Return the tree among the T_i 's defined above that minimizes $\mathbb{E}_{\ell \in T_i} [\text{Inf}((f_{\mathcal{D}})_{\ell})]$.

Figure 1: BUILDDT recursively searches for the depth- d , everywhere τ -influential tree of minimal influence at the leaves.

COROLLARY 5.4 (PRUNING LEMMA WITH EXPECTED TOTAL INFLUENCE AT LEAVES). *Let $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ be representable by a depth- d DT. Then there exists T^* such that the following are satisfied:*

- (1) *The size and depth of T^* are at most the size and depth of T ,*
- (2) *T^* is everywhere τ -influential with respect to f ,*
- (3) *$\mathbb{E}_{\ell \in T^*} [\text{Inf}(f_{\ell})] \leq 4d^2\tau$.*

We now move to show that the tree output by BUILDDT satisfies $2^{-n} \cdot \|f_{\mathcal{D}} - T'\|_1 \leq \varepsilon$.

CLAIM 5.5. *Let \mathcal{D} be a distribution that is representable by a depth- d decision tree, and let T be the output of BUILDDT, with $\tau = \varepsilon/8d^2$. Then, with high probability, $2^{-n} \cdot \|f_{\mathcal{D}} - T\|_1 \leq \varepsilon$.*

PROOF. First, we claim that T minimizes $\mathbb{E}_{\ell \in T} [\text{Inf}((f_{\mathcal{D}})_{\ell})]$ among all depth- d , everywhere τ -influential trees. This claim holds by induction on d : since

$$\mathbb{E}_{\ell \in T} [\text{Inf}((f_{\mathcal{D}})_{\ell})] = \frac{1}{2} (\mathbb{E}_{\ell \in T_{\text{left}}} [\text{Inf}((f_{\mathcal{D}})_{\ell})] + \mathbb{E}_{\ell \in T_{\text{right}}} [\text{Inf}((f_{\mathcal{D}})_{\ell})]),$$

each candidate T_i minimizes influence among all depth- d , everywhere τ -influential trees with x_i at the root under the inductive assumption that T_{left} and T_{right} minimize influence for depth- $(d-1)$ trees. Then BUILDDT chooses the tree of smallest influence among all the candidate T_i 's, so it minimizes total influence at leaves among all trees in its search space of depth- d , τ -influential trees.

Since Corollary 5.4 establishes the existence of a tree with average total influence at leaves $\leq \varepsilon/2$, it follows that the influence T 's influence is also $\leq \varepsilon/2$. Furthermore, we may assume by standard Hoeffding bounds that with a sample size of $\text{poly}(2^d, 1/\varepsilon)$, each leaf's value estimate of $\mathbb{E}[f_{\mathcal{D}}(\mathbf{x}) \mid \mathbf{x}$ is consistent with $\ell] = 2^{\lfloor \ell \rfloor} \cdot \Pr_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x}$ is consistent with $\ell]$ is accurate to within $\pm \varepsilon/2$ w.h.p..

We can now show that $2^{-n} \cdot \|f_{\mathcal{D}} - T\|_1 \leq \varepsilon$. Throughout this section, $\mathbf{x} \in \ell$ will stand in as shorthand for

“ \mathbf{x} consistent with the restriction at ℓ ”.

$$\begin{aligned} 2^{-n} \cdot \|f_{\mathcal{D}} - T\|_1 &= 2^{-n} \cdot \sum_{\mathbf{x} \in \{\pm 1\}^n} |f_{\mathcal{D}}(\mathbf{x}) - T(\mathbf{x})| \\ &= 2^{-n} \cdot \sum_{\ell \in T} \sum_{\mathbf{x} \in \ell} |(f_{\mathcal{D}})_{\ell}(\mathbf{x}) - T_{\ell}(\mathbf{x})| \\ &\leq 2^{-n} \cdot \sum_{\ell \in T} \sum_{\mathbf{x} \in \ell} |(f_{\mathcal{D}})_{\ell}(\mathbf{x}) - \mathbb{E}[(f_{\mathcal{D}})_{\ell}]| + |\mathbb{E}[(f_{\mathcal{D}})_{\ell}] - T_{\ell}(\mathbf{x})| \\ &= 2^{-n} \cdot \left(\sum_{\ell \in T} 2^{n-|\ell|} \cdot \text{Var}_{\mu}((f_{\mathcal{D}})_{\ell}) \right. \\ &\quad \left. + \sum_{\ell \in T} 2^{n-|\ell|} \cdot |\mathbb{E}[(f_{\mathcal{D}})_{\ell}] - T_{\ell}| \right) \\ &\leq \mathbb{E}_{\ell \in T} \text{Inf}((f_{\mathcal{D}})_{\ell}) + \mathbb{E}_{\ell \in T} [|\mathbb{E}[(f_{\mathcal{D}})_{\ell}(\mathbf{x})] - T_{\ell}|] \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon. \end{aligned} \tag{Lemma 4.3}$$

□

We can now prove correctness and time bounds for BUILDDT.

PROOF OF THEOREM 3. Correctness follows by combining Claim 5.5 and Fact 5.1. To bound the running time, first we note that for all restrictions π , there are at most $d/\tau = 8d^3/\varepsilon$ variables of influence at least τ . This is because all restrictions of $f_{\mathcal{D}}$ are depth- d decision trees, and for any depth- d decision tree, the sum of all variable influences is at most d .

Then the number of recursive calls to BUILDDT is at most $(8d^3/\varepsilon)^d$. Each call at an internal node makes n calls to the unit-time influence oracle, and each call at a leaf processes $\text{poly}(2^d, 1/\varepsilon)$ samples to estimate the leaf label. Thus, the total running time is $n \cdot (d/\varepsilon)^{O(d)}$, as desired. □

6 ALGORITHMS FOR COMPUTING DISTRIBUTIONAL INFLUENCES

In Section 5, we assumed the ability to exactly compute influences of f and its restrictions in unit time. In this section, we show how to instead estimate the influences from samples, when the distribution is monotone or when we have access to subcube conditional samples. Just as in [12], our proof only requires estimates to be accurate to $\pm \min(\tau/4, \varepsilon/n)$. Letting $\text{INFEST}_i(f)$ denote such an estimate of $\text{Inf}_i(f)$, the pseudocode in Figure 1 and proof of Theorem 3 is modified as follows.

- (1) We modify S to include variables i such that $\text{INFEST}_i((f_{\mathcal{D}})_\pi) \geq 3\tau/4$. Since the estimate is accurate to $\pm\tau/4$, this is guaranteed to include all variables with influence $\geq \tau$, and furthermore will only include variables with influence at least $\tau/2$. Therefore, the total size of S is at most $\frac{d}{\tau/2}$, which is only a constant factor (of 2) larger than in the proof of Theorem 3 which assumed perfect influence oracles, and there does not affect asymptotics.
- (2) It returns the tree T_i that minimizes $\mathbb{E}_{\ell \in T_i} [\sum_{i \in [n]} \text{INFEST}_i((f_{\mathcal{D}})_\ell)]$. Since each estimate is accurate to ε/n , this estimate of total influence of $(f_{\mathcal{D}})_\ell$ will be accurate to $\pm\varepsilon$. Then, in Claim 5.5, rather than BUILDDT building a tree T minimizes $\mathbb{E}_{\ell \in T} [\text{Inf}((f_{\mathcal{D}})_\ell)]$ among all depth- d , everywhere τ -influential trees, T (roughly) minimizes $\mathbb{E}_{\ell \in T} [\sum_{i \in [n]} \text{INFEST}_i((f_{\mathcal{D}})_\ell)]$ among all depth- d , everywhere τ -influential trees. More formally, T will either be the best depth- d , everywhere τ -influential trees, or better, as we know its searches over all variables with $\text{INFEST}_i((f_{\mathcal{D}})_\pi) \geq 3\tau/4$ which is guaranteed to include variables with influence $\geq \tau$, but can include more variables. Finally, since the estimates of total influence are accurate to $\pm\varepsilon$, using INFEST rather than Inf can only incur at most 2ε additive error, which is a constant factor in the analysis.

Given any distribution \mathcal{D} over $\{\pm 1\}^n$, we need to estimate $\text{Inf}_i((f_{\mathcal{D}})_\ell)$ for any restriction ℓ of $f_{\mathcal{D}}$ and $i \in [n] - \ell$. In this section, we will instead show how to compute $\text{Inf}_i(f_{\mathcal{E}})$ for any distribution \mathcal{E} over $\{\pm 1\}^m$. We can then use our estimators with $\mathcal{E} = \mathcal{D}_\ell$ to obtain the necessary answers using the following fact:

FACT 6.1. *For any distribution \mathcal{D} over $\{\pm 1\}^n$, restriction ℓ of $\{\pm 1\}^n$, and $i \in [n] - \ell$:*

$$\text{Inf}_i((f_{\mathcal{D}})_\ell) = 2^{|\ell|} \cdot \Pr_{x \sim \mathcal{D}} [x \in \ell] \cdot \text{Inf}_i(f_{\mathcal{D}_\ell}).$$

PROOF. Let $w_\ell = \Pr_{x \sim \mathcal{D}} [x \in \ell]$. We have:

$$\begin{aligned} \text{Inf}_i((f_{\mathcal{D}})_\ell) &= \mathbb{E}_{x \sim \mathcal{U}^n} [|f_{\mathcal{D}}(x) - f_{\mathcal{D}}(x^{\sim i})| \mid x \in \ell] \\ &= 2^n \cdot \mathbb{E}_{x \sim \mathcal{U}^n} [|D(x) - D(x^{\sim i})| \mid x \in \ell] \\ &= 2^n \cdot w_\ell \cdot \mathbb{E}_{x \sim \mathcal{U}^n} \left[\left| \frac{D(x)}{w_\ell} - \frac{D(x^{\sim i})}{w_\ell} \right| \mid x \in \ell \right] \\ &= 2^n \cdot w_\ell \cdot \mathbb{E}_{y \sim \mathcal{U}^{n-|\ell}} |D_\ell(y) - D_\ell(y^{\sim i})| \\ &= 2^{|\ell|} \cdot w_\ell \cdot \mathbb{E}_{y \sim \mathcal{U}^{n-|\ell}} |f_{\mathcal{D}_\ell}(y) - f_{\mathcal{D}_\ell}(y^{\sim i})| \\ &= 2^{|\ell|} \cdot w_\ell \cdot \text{Inf}_i(f_{\mathcal{D}_\ell}). \quad \square \end{aligned}$$

Because of Fact 6.1, for any restriction ℓ of depth at most d , to estimate $\text{Inf}_i((f_{\mathcal{D}})_\ell)$ to accuracy $\pm\varepsilon$, it is sufficient to estimate $\text{Inf}_i(f_{\mathcal{D}_\ell})$ to accuracy $\pm\varepsilon/2^d$. This 2^d factor is dominated by the $d^{O(d)}$ term in Theorem 2, so we are free to do the later.

6.1 Monotone Distributions using Samples

Let \mathcal{E} be an arbitrary distribution over $\{\pm 1\}^m$. If \mathcal{E} is monotone, the influences of $f_{\mathcal{E}}$ can be efficiently computed directly from samples of \mathcal{E} , via an estimate of bias:

LEMMA 6.2 (ESTIMATING INFLUENCE USING BIAS). *If \mathcal{E} is monotone,*

$$\text{Inf}_i(f_{\mathcal{E}}) = \mathbb{E}_{x \sim \mathcal{E}} [x_i].$$

PROOF. Using Fact 4.1,

$$\begin{aligned} \text{Inf}_i(f_{\mathcal{E}}) &= \mathbb{E}_{x \sim \mathcal{U}^m} [f_{\mathcal{E}}(x) \cdot x_i] \\ &= \sum_{x \in \{\pm 1\}^m} 2^{-m} f_{\mathcal{E}}(x) \cdot x_i \\ &= \sum_{x \in \{\pm 1\}^m} \mathcal{E}(x) \cdot x_i \\ &= \mathbb{E}_{x \sim \mathcal{E}} [x_i]. \quad \square \end{aligned}$$

As a simple application of the above and Hoeffding's inequality, we obtain the following corollary.

COROLLARY 6.3 (ESTIMATING INFLUENCES OF MONOTONE DISTRIBUTIONS). *For any $\varepsilon, \delta > 0$, there is an efficient algorithm that given unknown monotone distribution \mathcal{E} , computes an estimate of $\text{Inf}_i(f_{\mathcal{E}})$ to accuracy $\pm\varepsilon$ with probability at least $1 - \delta$ using $O(\log(1/\delta)/\varepsilon^2)$ random samples from \mathcal{E} .*

Recall that for Theorem 3, we only need the influence estimates to be accurate to $\pm \text{poly}(2^{-d}, \tau, \varepsilon, 1/n)$. Setting $\tau = O(\varepsilon/d^2)$ and union bounding over $n \cdot (d/\varepsilon)^{O(d)}$ calls to the influence oracle gives a sample complexity of $\text{poly}(n, 2^d, \varepsilon, \log(1/\delta))$. The running time is still dominated by the number of recursive calls.

COROLLARY 6.4 (SAMPLE COMPLEXITY OF BUILDDT FOR MONOTONE DISTRIBUTIONS). *The algorithm $\text{BUILDDT}(\mathcal{D}, \emptyset, d, \frac{\varepsilon}{2d^2})$, given $\text{poly}(n, 2^d, 1/\varepsilon, \log(1/\delta))$ random examples from a monotone distribution \mathcal{D} , runs in $\text{poly}(n) \cdot (d/\varepsilon)^{O(d)} \cdot \log(1/\delta)$ time and outputs a distribution within TV distance ε of \mathcal{D} . The algorithm fails with probability at most δ .*

6.2 Beyond Monotone Distributions using Subcube Conditional Sampling

We also design an influence estimator for arbitrary distributions \mathcal{E} over $\{\pm 1\}^m$ using subcube conditional sampling.

PROPOSITION 6.5 (INFEST HAS LOW BIAS). *For any distribution \mathcal{E} over $\{\pm 1\}^m$, coordinate $i \in [m]$, and $\varepsilon > 0$,*

$$|\mathbb{E} [\text{INFEST}(\mathcal{E}, i, \varepsilon)] - \text{Inf}_i(f_{\mathcal{E}})| \leq \varepsilon$$

where INFEST is as defined in Figure 2.

Before proving Proposition 6.5, we note that it implies a high accuracy estimator.

INFEST($\mathcal{E}, i, \varepsilon$):

Input: A distribution \mathcal{E} over $\{\pm 1\}^m$, coordinate $i \in [m]$, and bias parameter ε .

Output: An estimate of $\text{Inf}_i(f_{\mathcal{E}})$ that has bias at most ε .

- (1) Sample a random $x \sim \mathcal{E}$ and define the subcube

$$S := \{x\} \cup \{x^{\oplus i}\}$$

where $x^{\oplus i}$ is x with the i^{th} bit flipped.

- (2) Take $\lceil 1/\varepsilon^2 \rceil$ independent samples from \mathcal{E} conditioned on the output being in S , and let p be the fraction of those samples that equal x .
- (3) Output $|p - (1 - p)|$.

Figure 2: Pseudocode for estimating the influence of a variable on a distribution's weighting function.

COROLLARY 6.6 (ESTIMATING INFLUENCES). *For any $\varepsilon, \delta > 0$, there is an efficient algorithm that given unknown distribution \mathcal{E} , computes an estimate of $\text{Inf}_i(f_{\mathcal{E}})$ to accuracy $\pm \varepsilon$ with probability at least $1 - \delta$ using $O(\log(1/\delta)/\varepsilon^4)$ subcube conditional samples from \mathcal{E} .*

PROOF. The algorithm outputs the mean of $O(\log(1/\delta)/\varepsilon^2)$ independent calls to $\text{INFEST}(\mathcal{E}, i, \varepsilon/2)$. Each call to $\text{INFEST}(\mathcal{E}, i, \varepsilon/2)$ gives an output bounded within $[0, 1]$. By Hoeffding's inequality, if **est** is the mean of $O(\log(1/\delta)/\varepsilon^2)$ independent calls to $\text{INFEST}(\mathcal{E}, i, \varepsilon/2)$,

$$\Pr[|\mathbf{est} - \mathbb{E}[\text{INFEST}(\mathcal{E}, i, \varepsilon/2)]| \geq \varepsilon/2] \leq \delta.$$

The result then follows from triangle inequality and [Proposition 6.5](#). \square

We prove that INFEST has low bias.

PROOF OF PROPOSITION 6.5. For each $x \in \{\pm 1\}^n$, let

$$p(x) := \frac{\mathcal{E}(x)}{\mathcal{E}(x) + \mathcal{E}(x^{\oplus i})}$$

be the relative weight of x in the subcube containing x and $x^{\oplus i}$. Then, we can rewrite the influence as:

$$\begin{aligned} \text{Inf}_i(f_{\mathcal{E}}) &= \sum_{x \in \{\pm 1\}^m} \frac{1}{2^m} |f_{\mathcal{E}}(x) - f_{\mathcal{E}}(x^{\oplus i})| \\ &= \sum_{x \in \{\pm 1\}^m} |\mathcal{E}(x) - \mathcal{E}(x^{\oplus i})| \quad (f_{\mathcal{E}}(x) = 2^m \mathcal{E}(x)) \\ &= \frac{1}{2} \cdot \sum_{x \in \{\pm 1\}^m} |\mathcal{E}(x) - \mathcal{E}(x^{\oplus i})| \\ &\quad (x^{\oplus i} = x^{\oplus i} \text{ wp } \frac{1}{2}, \text{ and otherwise } x^{\oplus i} = x) \\ &= \frac{1}{2} \cdot \sum_{x \in \{\pm 1\}^m} (\mathcal{E}(x) + \mathcal{E}(x^{\oplus i})) \cdot |p(x) - p(x^{\oplus i})|. \end{aligned}$$

(definition of $p(x)$)

Then, using the fact that $p(x) = 1 - p(x^{\oplus i})$, and distributing the $(\mathcal{E}(x) + \mathcal{E}(x^{\oplus i}))$ term, we can write

$$\begin{aligned} \text{Inf}_i(f_{\mathcal{E}}) &= \frac{1}{2} \sum_{x \in \{\pm 1\}^m} \mathcal{E}(x) \cdot |p(x) - (1 - p(x))| \\ &\quad + \mathcal{E}(x^{\oplus i}) \cdot |p(x^{\oplus i}) - (1 - p(x^{\oplus i}))| \\ &= \sum_{x \in \{\pm 1\}^m} \mathcal{E}(x) \cdot |p(x) - (1 - p(x))| \end{aligned}$$

where, in the last step, we used the fact that summing over $x \in \{\pm 1\}^m$ is equivalent to summing over $x^{\oplus i} \in \{\pm 1\}^m$. Let $\hat{p}(x)$ be the random variable for the estimate of $p(x)$ computed by INFEST step 2. We bound the bias of INFEST .

$$\begin{aligned} & \left| \mathbb{E}[\text{INFEST}(\mathcal{E}, i, \varepsilon)] - \text{Inf}_i(f_{\mathcal{E}}) \right| \\ &= \left| \mathbb{E}_{x \sim \mathcal{E}}[|2\hat{p}(x) - 1|] - \mathbb{E}_{x \sim \mathcal{E}}[|2p(x) - 1|] \right| \\ &\quad (2p - 1 = p - (1 - p)) \\ &= \left| \mathbb{E}_{x \sim \mathcal{E}}[|2\hat{p}(x) - 1| - |2p(x) - 1|] \right| \\ &\quad (\text{linearity of expectation}) \\ &\leq \mathbb{E}_{x \sim \mathcal{E}}[||2\hat{p}(x) - 1| - |2p(x) - 1||] \\ &\quad (\text{Jensen's inequality}) \\ &\leq 2 \left| \mathbb{E}_{x \sim \mathcal{E}}[|\hat{p}(x) - p(x)|] \right|. \end{aligned}$$

($||a| - |b|| \leq |a - b|$)

For any $x \in \{\pm 1\}^m$, $\hat{p}(x)$ is the average of $\lceil 1/\varepsilon^2 \rceil$ random variables each which is 1 with probability $p(x)$ and 0 otherwise. As a result, $\mathbb{E}[\hat{p}(x)] = p(x)$, and $\text{Var}[\hat{p}(x)] \leq \varepsilon^2/4$. Applying Jensen's inequality, we conclude

$$|\mathbb{E}[\text{INFEST}(\mathcal{E}, i, \varepsilon)] - \text{Inf}_i(f_{\mathcal{E}})| \leq 2\sqrt{\text{Var}[\hat{p}(x)]} \leq \varepsilon. \quad \square$$

7 LIFTING UNIFORM DISTRIBUTION LEARNERS: PROOF OF THEOREM 1

In this section, we show how to lift algorithms that learn over the uniform distribution to algorithms that learn with respect to arbitrary distributions, where the sample complexity and runtime scale with the decision tree complexity of the distribution. We first define our goal formally.

DEFINITION 12 (LEARNING WITH RESPECT TO A CLASS OF DISTRIBUTIONS). *For any concept class \mathcal{C} of functions $f : \{\pm 1\}^n \rightarrow \{0, 1\}$, $\varepsilon, \delta > 0$, set of distributions \mathcal{D} with support $\{\pm 1\}^n$, and $m, d \in \mathbb{N}$, we say that an algorithm $\mathcal{A}(\varepsilon, \delta)$ -learns \mathcal{C} for distributions \mathcal{D} using m samples if the following holds: For any $\mathcal{D} \in \mathcal{D}$ and any $f^* \in \mathcal{C}$, given m iid samples of the form $(x, f^*(x))$ where $x \sim \mathcal{D}$, \mathcal{A} outputs a hypothesis h satisfying*

$$\Pr_{x \sim \mathcal{D}}[f^*(x) \neq h(x)] \leq \varepsilon.$$

with probability at least $1 - \delta$.

Generally, we can think of δ as any fixed constant, as the success probability can always be boosted.

FACT 7.1 (BOOSTING SUCCESS PROBABILITY). *Given an algorithm \mathcal{A} that $(\epsilon, \frac{1}{2})$ -learns a concept \mathcal{C} using m samples, for any $\delta > 0$, we can construct an \mathcal{A}' that $(1.1\epsilon, \delta)$ -learns \mathcal{C} using $m \cdot \text{poly}(1/\epsilon, \log(1/\delta))$ samples.*

PROOF. By repeating \mathcal{A} $\log(1/\delta)$ times, we can guarantee that with probability at least $1 - \delta/2$, one of the returned hypotheses is ϵ -close to f^* . The accuracy of each of these hypothesis can be estimated to accuracy $\pm 0.05\epsilon$ using $O(\log(1/\delta)/\epsilon^2)$ random samples, and the most accurate one returned. With probability at least $1 - \delta$, that hypothesis will have at most 1.1ϵ error. \square

Our goal is to learn with respect to the class of all low-depth decision tree distributions. We use the following natural assumption on the concept class, which includes almost every concept class considered in the learning theory literature.

DEFINITION 13 (CLOSED UNDER RESTRICTION). *A concept class \mathcal{C} of functions $f : \{\pm 1\}^n \rightarrow \{0, 1\}$ is closed under restriction if, for any $f \in \mathcal{C}$, $i \in [n]$, and $b \in \{\pm 1\}$, the restriction $f_{i=b}$ is also in \mathcal{C} .*

We will use **Theorem 3** to first decompose \mathcal{D} into a mixture of nearly uniform distributions, and then run our learner on each of those distributions, as described in **Figure 3**.

LIFTLERNER(T, \mathcal{A}, S):

Input: A decision tree T , an algorithm for learning in the uniform distribution \mathcal{A} , and a random labeled sample S .

Output: A hypothesis

For each leaf $\ell \in T$ {

- (1) Let S_ℓ be the subset of points in S that reach ℓ .
- (2) Create a set S'_ℓ consisting of points in S_ℓ but where all coordinates queried on the root-to-leaf path for ℓ are rerandomized independently (this makes the marginal over the input uniform).
- (3) Use \mathcal{A} to learn a hypothesis, h_ℓ , with S'_ℓ as input.

}

Return the hypothesis that, when given an input x , first determines which leaf $\ell \in T$ that x follows and then outputs $h_\ell(x)$.

Figure 3: Pseudocode lifting a uniform distribution learner to one which succeeds on decision tree distributions. In this pseudocode, we assume that we have a decision tree representation which is close to the distribution, which can be accomplished using BUILDDT in Figure 1.

For our first result, we will assume that we already have a learner that succeeds on distributions that are sufficiently close to uniform.

DEFINITION 14 (ROBUST LEARNERS). *For any concept class \mathcal{C} of functions $f : \{\pm 1\}^n \rightarrow \{0, 1\}$ and algorithm \mathcal{A} , we say that*

$\mathcal{A}(\epsilon, \delta, c)$ -robustly learns \mathcal{C} using m samples under the uniform distribution if, for any $\eta > 0$ and the class of distributions

$\mathcal{D}_{\text{TV}, \eta} := \{\text{Distributions } \mathcal{D} \text{ over } \{\pm 1\}^n \text{ where } \text{dist}_{\text{TV}}(\mathcal{U}, \mathcal{D}) \leq \eta\}$, $\mathcal{A}(\epsilon + c\eta, \delta)$ -learns \mathcal{C} for the distributions in $\mathcal{D}_{\text{TV}, \eta}$ using m samples.

The study of robust learners is part of a long and fruitful line of work. In particular, every learner that is robust to nasty noise [14] meets our definition of robust learners.

Our result will also apply to learners that aren't explicitly robust. This is because every learner is robust for $c = O(m)$.

PROPOSITION 7.2. *For any concept class \mathcal{C} and algorithm \mathcal{A} , if $\mathcal{A}(\epsilon, \delta)$ -learns \mathcal{C} using m samples under the uniform distribution, then \mathcal{A} also $(\epsilon, \delta + \frac{1}{3}, 3m)$ -robustly learns \mathcal{C} using m samples.*

PROOF. Fix any $\eta > 0$. Our goal is to show that $\mathcal{A}(\epsilon + 3m\eta, \delta)$ -learns \mathcal{C} for distributions in $\mathcal{D}_{\text{distTV}, \eta}$. If $\eta \geq \frac{1}{3m}$, this is obviously true as any hypothesis has error ≤ 1 . We therefore need only consider $\eta < \frac{1}{3m}$. When \mathcal{A} receives a sample from \mathcal{U}^m , it returns a hypothesis with error $\leq \epsilon$ with probability at least $1 - \eta$. Instead, \mathcal{A} is receiving a sample from \mathcal{D}^m , where $\text{dist}_{\text{TV}}(\mathcal{U}, \mathcal{D}) < \frac{1}{3m}$. The success probability of any test given a sample from \mathcal{D}^m rather than \mathcal{U}^m can only differ by at most

$$\text{dist}_{\text{TV}}(\mathcal{U}^m, \mathcal{D}^m) \leq m \cdot \text{dist}_{\text{TV}}(\mathcal{U}, \mathcal{D}) < \frac{1}{3}.$$

Therefore, for $\eta < \frac{1}{3m}$, \mathcal{A} succeeds w.p at least $\delta + \frac{1}{3}$, as desired. \square

We now state the main result of this section.

THEOREM 4. *Choose any concept class \mathcal{C} of functions $\{\pm 1\}^n \rightarrow \{0, 1\}$ closed under restrictions, $\epsilon, \delta, c > 0$, $m, d \in \mathbb{N}$, and algorithm \mathcal{A} that $(\epsilon, \delta/(2 \cdot 2^d), c)$ -robustly learns \mathcal{C} using m samples under the uniform distribution. For any function $f^* \in \mathcal{C}$, distribution \mathcal{D} over $\{\pm 1\}^n$, depth- d decision tree $T : \{\pm 1\}^n \rightarrow \mathbb{R}$ computing the PMF of a distribution \mathcal{D}_T where*

$$\text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}_T) \leq \frac{\epsilon}{c},$$

and sample size of

$$M = m \cdot \text{poly}\left(2^d, \frac{1}{\epsilon}, \log\left(\frac{1}{\delta}\right)\right).$$

Let S a size- M iid sample of labeled points $(x, f^*(x))$ where $x \sim \mathcal{D}$. The output of **LIFTLERNER**(T, \mathcal{A}, S) is $O(\epsilon)$ -close to f^* w.r.t \mathcal{D} with probability at least $1 - \delta$.

By **Fact 7.1**, an algorithm with constant failure probability could be transformed into one with the failure probability required by **Theorem 4** with only a $\text{poly}(d, 1/\epsilon, \log(\delta))$ increase in the sample size. Therefore, would **Theorem 4** still holds when $\mathcal{A}(\epsilon, \frac{1}{2}, c)$ -robustly learns \mathcal{C} if **LIFTLERNER** applies **Fact 7.1** to boost the success probability of \mathcal{A} . Before proving **Theorem 4**, we show how it implies our main result.

THEOREM 5 (LIFTING UNIFORM-DISTRIBUTION LEARNERS, FORMAL VERSION OF THEOREM 1). *Choose any concept class \mathcal{C} of functions $\{\pm 1\}^n \rightarrow \{0, 1\}$ closed under restrictions, $\epsilon, c > 0$, $m, d \in \mathbb{N}$. If there is an efficient algorithm, \mathcal{A} , that $(\epsilon, \frac{1}{2})$ -learns \mathcal{C} using m samples for the uniform distribution, then for $M = \text{poly}(n) \cdot \left(\frac{dm}{\epsilon}\right)^{O(d)}$,*

- There is an algorithm that $(\varepsilon, \frac{1}{6})$ -learns \mathcal{C} using M samples for monotone distributions representable by a depth- d decision tree.
- There is an algorithm that uses M conditional subcube samples from \mathcal{D} and M random samples labeled by the target function that learns \mathcal{C} to ε -accuracy for arbitrary (not necessarily monotone) distributions representable by a depth- d decision tree.

In both cases, the algorithm runs in time $\text{poly}(n, M)$.

PROOF OF THEOREM 5 GIVEN THEOREM 4. Using THEOREM 3, we can learn the input distribution to TV-distance accuracy $\frac{\varepsilon}{3m}$ with a decision tree hypothesis. By PROPOSITION 7.2, $\mathcal{A}(\varepsilon, \frac{1}{2}, 3m)$ -robustly learns \mathcal{C} for the uniform distribution, which can be boosted to failure probability $O(2^{-d})$ using FACT 7.1. Applying THEOREM 4 gives the desired result, with the desired runtime following from PROPOSITION 7.5. \square

The remainder of this section is devoted to the proof of THEOREM 4. We'll use the following proposition.

PROPOSITION 7.3. For any distribution \mathcal{D} over $\{\pm 1\}^n$, depth- d decision tree T , $m \in \mathbb{N}$ and $p, \delta > 0$, as long as

$$M \geq O\left(\frac{d + m + \log(1/\delta)}{p}\right) \quad (3)$$

for a random sample of M points from \mathcal{D} , the probability there is a leaf $\ell \in T$ satisfying:

- (1) High weight: The probability a random sample from \mathcal{D} reaches ℓ is at least p ,
- (2) Few samples: The number of points in the size- M sample that reach ℓ is less than m

is at most δ .

PROOF. Fix a single leaf $\ell \in T$ with weight at least p . Then, the expected number of points that reach this leaf is $\mu \geq Mp$. As long as $\mu \geq 2m$, applying multiplicative Chernoff bounds,

$$\Pr[\text{Fewer than } m \text{ points reach } \ell] \leq \exp\left(-\frac{\mu}{8}\right).$$

Union bounding over all leaves 2^d , it is sufficient to choose an M where

$$2^d \exp\left(-\frac{\mu}{8}\right) \leq \delta.$$

This is satisfied for the M from EQUATION (3). \square

We'll also use that if we have a decision tree T that has learned the PMF to \mathcal{D} to high accuracy, \mathcal{D} restricted to the leaves of T is, on average over the leaves, close to uniform.

FACT 7.4 (LEMMA B.4 OF [11]). For any distribution \mathcal{D} and decision tree T computing the PMF for some distribution \mathcal{D}_T ,

$$\sum_{\text{leaves } \ell \in T} \Pr[\mathbf{x} \text{ reaches } \ell] \cdot \text{dist}_{\text{TV}}(\mathcal{D}_\ell, (\mathcal{D}_T)_\ell) \leq 2 \cdot \text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}_T).$$

PROOF OF THEOREM 4. First, we split up the accuracy of $h = \text{LIFTLEARNER}(T, \mathcal{A}, S)$ into the accuracy of the hypotheses h_ℓ

learned at each leaf ℓ :

$$\begin{aligned} \Pr_{\mathbf{x} \sim \mathcal{D}} [h(\mathbf{x}) \neq f^*(\mathbf{x})] \\ = \sum_{\ell \in T} \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \cdot \Pr_{\mathbf{x} \sim \mathcal{D}_\ell} [h_\ell(\mathbf{x}) \neq f^*(\mathbf{x})]. \end{aligned}$$

Each hypothesis h_ℓ is learned by running \mathcal{A} on the sample S'_ℓ . First, we argue that for all leaves ℓ with much of \mathcal{D} 's weight, will, whp, have $\geq m$ samples. Applying PROPOSITION 7.3 with probability at least $1 - \delta/2$, for all $\ell \in T$ with $\Pr_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x} \text{ reaches } \ell] \geq \frac{\varepsilon}{2^d}$, the size of S'_ℓ is at least m . Conditioning on that being true, we have that,

$$\begin{aligned} \Pr_{\mathbf{x} \sim \mathcal{D}} [h(\mathbf{x}) \neq f^*(\mathbf{x})] \\ \leq \sum_{\ell \in T} \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \cdot \Pr_{\mathbf{x} \sim \mathcal{D}_\ell} [h_\ell(\mathbf{x}) \neq f^*(\mathbf{x}) \mid |S'_\ell| \geq m] \\ + \sum_{\ell \in T} \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \cdot \mathbb{1}[\Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \leq p] \\ \leq \sum_{\ell \in T} \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \cdot \Pr_{\mathbf{x} \sim \mathcal{D}_\ell} [h_\ell(\mathbf{x}) \neq f^*(\mathbf{x}) \mid |S'_\ell| \geq m] + \varepsilon, \end{aligned}$$

where the last line uses that T has at most 2^d leaves and $p = \frac{\varepsilon}{2^d}$.

Each point in S'_ℓ is an iid sample with the input uniform over $\{\pm 1\}^n$ and labeled by the function f_ℓ^* . As \mathcal{C} is closed under restrictions, $f_\ell^* \in \mathcal{C}$. Therefore,

$$\begin{aligned} \Pr_S \left[\left[\Pr_{\mathbf{x} \sim \mathcal{D}_\ell} [h_\ell(\mathbf{x}) \neq f^*(\mathbf{x}) \mid |S'_\ell| \geq m] \right] \leq \varepsilon + c \cdot \text{dist}_{\text{TV}}(\mathcal{D}_\ell, \mathcal{U}) \right] \\ \geq 1 - \frac{\delta}{2 \cdot 2^d}. \end{aligned}$$

We union bound over all 2^d leaves ℓ and the earlier event that $|S_\ell| \geq m$ whenever $\Pr_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x} \text{ reaches } \ell] \geq p$, we have with probability at least $1 - \delta$,

$$\begin{aligned} \Pr_{\mathbf{x} \sim \mathcal{D}} [h(\mathbf{x}) \neq f^*(\mathbf{x})] \\ \leq \sum_{\ell \in T} \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \cdot (2\varepsilon + c \cdot \text{dist}_{\text{TV}}(\mathcal{D}_\ell, \mathcal{U})) \\ = 2\varepsilon + c \cdot \sum_{\ell \in T} \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \cdot \text{dist}_{\text{TV}}(\mathcal{D}_\ell, \mathcal{U}) \\ = 2\varepsilon + c \cdot \sum_{\ell \in T} \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \text{ reaches } \ell] \cdot \text{dist}_{\text{TV}}(\mathcal{D}_\ell, (\mathcal{D}_T)_\ell) \\ \leq 2\varepsilon + c \cdot 2 \cdot \text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}_T) \quad (\text{Fact 7.4}) \\ \leq 4\varepsilon. \quad (\text{dist}_{\text{TV}}(\mathcal{D}, \mathcal{D}_T) \leq \frac{\varepsilon}{2}) \end{aligned}$$

As desired, we have LIFTLEARNER learns an $O(\varepsilon)$ -accurate hypothesis w.h.p. \square

PROPOSITION 7.5 (RUNTIME OF LIFTLEARNER). Assuming unit-time calls to \mathcal{A} , given a depth- d decision tree T , $\text{LIFTLEARNER}(T, \mathcal{A}, S)$ runs in time $O(n2^d \cdot |S|)$.

PROOF. The number of leaves in T is at most 2^d . Since each point in S is in $\{\pm 1\}^n$, the entire sample takes $O(n \cdot |S|)$ bits to represent. For each leaf ℓ , it takes $O(n \cdot |S|)$ time to loop through this representation, find the points consistent with ℓ , and create the modified dataset S'_ℓ , and pass it into \mathcal{A} . Repeating this over all leaves can be done in $O(n2^d \cdot |S|)$ time. \square

Remark 1 (The agnostic setting). A popular variant of learning, as defined in Definition 12, is the *agnostic setting*. In this generalization of standard learning, rather than assume $f^* \in \mathcal{C}$, \mathcal{A} is required to output a hypothesis with error $\text{opt} + \epsilon$, where opt is the minimum error of a hypothesis in \mathcal{C} w.r.t f^* . It's straightforward to see that Theorem 4, and therefore theorem 5, extend to the agnostic setting, where if the uniform distribution learning \mathcal{A} succeeds in the agnostic setting, that learner is upgraded to one that succeeds for decision tree distributions in the agnostic setting. This is because, if there is an $f \in \mathcal{C}$ with error opt w.r.t f^* , then the average error of f over the leaves of a tree T w.r.t f^* will also be at most opt .

ACKNOWLEDGMENTS

We thank the STOC reviewers for their detailed feedback, especially for the references to the literature on semi-supervised learning.

Guy and Li-Yang are supported by NSF awards 1942123, 2211237, and 2224246. Jane is supported by NSF Award CCF-2006664. Ali is supported by a graduate fellowship award from Knight-Hennessy Scholars at Stanford University.

REFERENCES

- Jayadev Acharya, Clément L. Canonne, and Gautam Kamath. 2015. Adaptive estimation in weighted group testing. In *2015 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2116–2120.
- Jayadev Acharya, Clément L. Canonne, and Gautam Kamath. 2015. A Chasm Between Identity and Equivalence Testing with Conditional Queries. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* (2015), 449.
- Maryam Aliakbarpour, Eric Blais, and Ronitt Rubinfeld. 2016. Learning and Testing Junta Distributions. In *29th Annual Conference on Learning Theory (Proceedings of Machine Learning Research, Vol. 49)*, Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir (Eds.). PMLR, Columbia University, New York, New York, USA, 19–46. <https://proceedings.mlr.press/v49/aliakbarpour16.html>
- Maria-Florina Balcan and Avrim Blum. 2010. A Discriminative Model for Semi-Supervised Learning. *57, 3* (2010). <https://doi.org/10.1145/1706591.1706599>
- Tugkan Batu, Sanjoy Dasgupta, Ravi Kumar, and Ronitt Rubinfeld. 2005. The complexity of approximating the entropy. *SIAM J. Comput.* 35, 1 (2005), 132–150.
- Shai Ben-David, Tyler Lu, and David Pa. 2008. Does Unlabeled Data Provably Help? Worst-case Analysis of the Sample Complexity of Semi-Supervised Learning. In *Proceedings of the Twenty-First Annual Conference on Learning Theory*. <http://colt2008.cs.helsinki.fi/papers/92-Ben-David.pdf>
- Gyora M. Benedek and Alon Itai. 1991. Learnability with respect to fixed distributions. *Theoretical Computer Science* 86, 2 (1991), 377–389. [https://doi.org/10.1016/0304-3975\(91\)90026-X](https://doi.org/10.1016/0304-3975(91)90026-X)
- Rishiraj Bhattacharyya and Sourav Chakraborty. 2018. Property testing of joint distributions using conditional samples. *ACM Transactions on Computation Theory (TOCT)* 10, 4 (2018), 1–20.
- Eric Blais, Clément L. Canonne, and Tom Gur. 2019. Distribution testing lower bounds via reductions from communication complexity. *ACM Transactions on Computation Theory (TOCT)* 11, 2 (2019), 1–37.
- Eric Blais, Ryan O'Donnell, and Karl Wimmer. 2010. Polynomial regression under arbitrary product distributions. *Machine learning* 80, 2 (2010), 273–294.
- Guy Blanc, Jane Lange, Ali Malik, and Li-Yang Tan. 2022. Popular decision tree algorithms are provably noise tolerant. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*.
- Guy Blanc, Jane Lange, Mingda Qiao, and Li-Yang Tan. 2021. Properly learning decision trees in almost polynomial time. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*.
- Avrim Blum and Pat Langley. 1997. Selection of Relevant Features and Examples in Machine Learning. *Artificial Intelligence* 97, 1-2 (1997), 245–271.
- Nader H Bshouty, Nadav Eiron, and Eyal Kushilevitz. 2002. PAC learning with nasty noise. *Theoretical Computer Science* 288, 2 (2002), 255–275.
- Clément Canonne and Ronitt Rubinfeld. 2014. Testing probability distributions underlying aggregated data. In *International Colloquium on Automata, Languages, and Programming*. Springer, 283–295.
- Clément L. Canonne. 2015. Big data on the rise?. In *International Colloquium on Automata, Languages, and Programming*. Springer, 294–305.
- Clément L. Canonne, Xi Chen, Gautam Kamath, Amit Levi, and Erik Waingarten. 2021. Random restrictions of high dimensional distributions and uniformity testing with subcube conditioning. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 321–336.
- Clément L. Canonne, Dana Ron, and Rocco A. Servedio. 2015. Testing probability distributions using conditional samples. *SIAM J. Comput.* 44, 3 (2015), 540–616.
- Sourav Chakraborty, Eldar Fischer, Yonatan Goldhirsh, and Arie Matsliah. 2016. On the power of conditional samples in distribution testing. *SIAM J. Comput.* 45, 4 (2016), 1261–1296.
- Sitan Chen and Ankur Moitra. 2019. Beyond the low-degree algorithm: mixtures of subcubes and their applications. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC)*. 869–880.
- Xi Chen, Rajesh Jayaram, Amit Levi, and Erik Waingarten. 2021. Learning and testing junta distributions with sub cube conditioning. In *Conference on Learning Theory*. PMLR, 1060–1113.
- Mary Cryan. 1999. *Learning and approximation Algorithms for Problems motivated by evolutionary trees*. Ph. D. Dissertation. Department of Computer Science.
- Mary Cryan, Leslie Ann Goldberg, and Paul W. Goldberg. 2001. Evolutionary trees can be learned in polynomial time in the two-state general Markov model. *SIAM J. Comput.* 31, 2 (2001), 375–397.
- Andrzej Ehrenfeucht and David Haussler. 1989. Learning decision trees from random examples. *Information and Computation* 82, 3 (1989), 231–246.
- Moein Falahatgar, Ashkan Jafarpour, Alon Orlitsky, Venkatesh Pichapati, and Ananda Theertha Suresh. 2015. Faster algorithms for testing under conditional sampling. In *Conference on Learning Theory*. PMLR, 607–636.
- Jon Feldman, Ryan O'Donnell, and Rocco A. Servedio. 2008. Learning mixtures of product distributions over discrete domains. *SIAM J. Comput.* 37, 5 (2008), 1536–1564.
- Eldar Fischer, Oded Lachish, and Yadu Vasudev. 2019. Improving and extending the testing of distributions for shape-restricted properties. *Algorithmica* 81, 9 (2019), 3765–3802.
- Yoav Freund and Yishay Mansour. 1999. Estimating a mixture of two product distributions. In *Proceedings of the twelfth annual conference on Computational learning theory*. 53–62.
- Oded Goldreich, Shafi Goldwasser, and Dana Ron. 1998. Property testing and its connection to learning and approximation. *Journal of the ACM* 45 (1998), 653–750.
- Parikshit Gopalan, Adam Kalai, and Adam Klivans. 2008. Agnostically learning decision trees. In *Proceedings of the 40th ACM Symposium on Theory of Computing (STOC)*. 527–536.
- Christina Göpfert, Shai Ben-David, Olivier Bousquet, Sylvain Gelly, Ilya Tolstikhin, and Ruth Urner. 2019. When can unlabeled data improve the learning rate?. In *Proceedings of the Thirty-Second Conference on Learning Theory (Proceedings of Machine Learning Research, Vol. 99)*, Alina Beygelzimer and Daniel Hsu (Eds.). PMLR, 1500–1518. <https://proceedings.mlr.press/v99/gopfert19a.html>
- Themistoklis Gouleakis, Christos Tzamos, and Manolis Zampetakis. 2017. Faster sublinear algorithms using conditional sampling. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1743–1757.
- Themis Gouleakis, Christos Tzamos, and Manolis Zampetakis. 2018. Certified computation from unreliable datasets. In *Conference On Learning Theory*. PMLR, 3271–3294.
- Sudipto Guha, Andrew McGregor, and Suresh Venkatasubramanian. 2006. Streaming and sublinear approximation of entropy and information distances. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. 733–742.
- David Haussler. 1992. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and computation* 100, 1 (1992), 78–150.
- Jeffrey C. Jackson. 1997. An Efficient Membership-Query Algorithm for Learning DNF with Respect to the Uniform Distribution. *J. Comput. System Sci.* 55, 3 (1997), 414–440. <https://doi.org/10.1006/jcss.1997.1533>
- Adam Kalai, Adam Klivans, Yishay Mansour, and Rocco A. Servedio. 2008. Agnostically Learning Halfspaces. *SIAM J. Comput.* 37, 6 (2008), 1777–1805.
- Michael Kearns, Robert Schapire, and Linda Sellie. 1994. Toward efficient agnostic learning. *Machine Learning* 17, 2/3 (1994), 115–141.
- Subhash Khot and Rishi Saket. 2008. On hardness of learning intersection of two halfspaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*. 345–354.
- Adam R. Klivans, Ryan O'Donnell, and Rocco A. Servedio. 2004. Learning intersections and thresholds of halfspaces. *J. Comput. System Sci.* 68, 4 (2004), 808–840.
- Adam R. Klivans and Rocco A. Servedio. 2004. Learning DNF in time $2^{\tilde{O}(n^{1/3})}$. *J. Comput. System Sci.* 68, 2 (2004), 303–318. <https://doi.org/10.1016/j.jcss.2003.07.007> Special Issue on STOC 2001.
- Adam R. Klivans and Alexander A. Sherstov. 2009. Cryptographic hardness for learning intersections of halfspaces. *J. Comput. System Sci.* 75, 1 (2009), 2–12. <https://doi.org/10.1016/j.jcss.2008.07.008> Learning Theory 2006.
- Eyal Kushilevitz and Yishay Mansour. 1993. Learning Decision Trees Using the Fourier Spectrum. *SIAM J. Comput.* 22, 6 (Dec. 1993), 1331–1348.

- [44] Nathan Linial, Yishay Mansour, and Noam Nisan. 1993. Constant depth circuits, Fourier transform and learnability. *J. ACM* 40, 3 (1993), 607–620.
- [45] B. K. Natarajan. 1992. Probably Approximate Learning over Classes of Distributions. *SIAM J. Comput.* 21, 3 (1992), 438–449. <https://doi.org/10.1137/0221029> arXiv:<https://doi.org/10.1137/0221029>
- [46] Ryan O’Donnell. 2014. *Analysis of Boolean Functions*. Cambridge University Press.
- [47] Krzysztof Onak and Xiaorui Sun. 2018. Probability-revealing samples. In *International Conference on Artificial Intelligence and Statistics*. PMLR.
- [48] Ronitt Rubinfeld and Rocco A Servedio. 2009. Testing monotone high-dimensional distributions. *Random Structures & Algorithms* 34, 1 (2009), 24–44.
- [49] Ronitt Rubinfeld and Madhu Sudan. 1996. Robust characterizations of polynomials with applications to program testing. *SIAM J. Comput.* 25, 2 (1996), 252–271.
- [50] Ronitt Rubinfeld and Arsen Vasilyan. 2020. Monotone Probability Distributions over the Boolean Cube Can Be Learned with Sublinear Samples. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 151)*, Thomas Vidick (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 28:1–28:34. <https://doi.org/10.4230/LIPIcs.ITCS.2020.28>
- [51] Imdad S. B. Sardharwalla, Sergii Strelchuk, and Richard Jozsa. 2017. Quantum conditional query complexity. *Quantum Information & Computation* 17, 7-8 (2017), 541–567.
- [52] Alexander A Sherstov. 2013. The intersection of two halfspaces has high threshold degree. *SIAM J. Comput.* 42, 6 (2013), 2329–2374.
- [53] Gregory Valiant and Paul Valiant. 2011. The power of linear estimators. In *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 403–412.
- [54] Leslie Valiant. 1984. A theory of the learnable. *Commun. ACM* 27, 11 (1984), 1134–1142.
- [55] Karsten Verbeurgt. 1990. Learning DNF under the uniform distribution in quasipolynomial time. In *Proceedings of the 3rd Annual Workshop on Computational Learning Theory*. 314–326.

Received 2022-11-07; accepted 2023-02-06