

Multiple Random Oracles Are Better Than One

Jan Arpe and Elchanan Mossel

U.C. Berkeley [arpe,mossel]@stat.berkeley.edu

Abstract

We study the problem of learning k -juntas given access to examples drawn from a number of different product distributions. Thus we wish to learn a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ that depends on k (unknown) coordinates. While the best known algorithms for the general problem of learning a k -junta require running time of $n^k \text{poly}(n, 2^k)$, we show that given access to k different product distributions with biases separated by $\gamma > 0$, the functions may be learned in time $\text{poly}(n, 2^k, \gamma^{-k})$. More generally, given access to $t \leq k$ different product distributions, the functions may be learned in time $n^{k/t} \text{poly}(n, 2^k, \gamma^{-k})$. Our techniques involve novel results in Fourier analysis relating Fourier expansions with respect to different biases and a generalization of Russo's formula.

Keywords: learning juntas, PAC learning, biased product distributions, Fourier analysis of Boolean functions, Russo's formula

1 Introduction

1.1 Motivation

A k -junta is a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ that only depends on a subset of k variables x_{i_1}, \dots, x_{i_k} . Blum and Langley [6] proposed the problem of learning the class of k -juntas, which we refer to as the *junta learning problem*, as a clean and appealing model of learning in the presence of much irrelevant information. It is considered to be among the most important problems in computational learning theory to date [4, 23]. In addition to being an interesting class in itself, the importance of learning juntas is supported by its connections to learning decision trees and DNFs, see [23]. Mossel, O'Donnell, and Servedio [23] observed that junta learning is efficiently solvable in the membership query model and in the random walk model, whereas it is provably hard in the statistical query model. What lies in between is the uniform distribution PAC model for which [23] presented an algorithm with running time roughly $n^{0.7 \cdot k}$, being the currently best improvement upon a straightforward algorithm that runs in roughly n^k steps. For general distributions, no such improvement is known. The little progress on the junta learning problem in the PAC model to date might be considered evidence of the hardness of the problem in this model. At the same time, however, no lower bounds are available, either.

Apart from devising *fast* learning algorithms, another goal is often to have *low sample complexity* (i.e., a small number of examples needed to learn). Information-theoretically, $\Theta(k \log n + 2^k)$ examples are necessary and sufficient for learning k -juntas on n bits ([7, 26, 1]). The algorithm of [23], however, needs to draw roughly $n^{0.3k}$ examples in the worst case.

It thus seems reasonable to ask if we can find a natural extension of the PAC learning model under fixed distributions that admits junta learning algorithms that run in time $t(k) \cdot \text{poly}(n)$ for some function t that is independent of n and some polynomial that is independent of k . Moreover, such algorithms should ideally use $s(k) \cdot O(\log n)$ examples for some function s independent of n .

In this paper, we propose such a model: instead of giving the learner access to only one oracle, we study the setting in which a learner has access to *multiple* oracles that generate examples according to different

distributions. Although in this paper, we are mainly interested in learning from product distributions, we introduce the model in more generality since we believe that studying the learnability of other classes in this model, possibly under less restricted distributions, is a worthwhile goal for future research. In data mining and applied machine learning, researchers often depart from the assumption of having access to only one source of data in order to capture more realistic scenarios such as having multiple sources of different quality [10, 11], receiving partial information about tuples of examples [12], or observing sets of different attributes for the same examples [21]. We mention three possible real-world learning scenarios in which our model can be applied: e.g., the examples could be obtained as series of measurements in certain experimental setups, so that different oracles correspond to different setups, resulting in different distributions over the instance space. Or, examples could be sampled from disjoint populations in which the distributions of attributes differ significantly. Another application comes into mind when considering data generated by a *mixture* of distributions. After applying algorithms to tell the distributions apart (say, from unlabeled examples) [19, 28, 13], one could use algorithms designed for the model of learning from multiple distributions to finally learn the concept under consideration.

For our results on the junta learning problem, we consider *r-biased oracles* that generate examples $(\mathbf{x}, f(\mathbf{x}))$ according to *r-biased product distributions* μ_r on $\{-1, 1\}^n$ for biases $r \in (-1, 1)$. These are distributions such that every variable x_i independently takes on values -1 and $+1$ with probability $(1-r)/2$ and $(1+r)/2$, respectively (so that $\mathbb{E}_{\mu_r}[x_i] = r$).

As in the setting with one uniform distribution oracle [23] (this is the case $r = 0$), we show that the junta learning problem from multiple oracles reduces to the task of identifying at least one relevant variable. In general, a conceptual method to identify relevant variables is to find non-vanishing *Fourier coefficients* $\hat{f}(S, r)$, $S \subseteq [n]$, where r denotes the bias of the underlying distribution. The Fourier coefficient $\hat{f}(S, r)$ measures the correlation between the function value $f(\mathbf{x})$ and the function $\chi_S(\mathbf{x}, r) = \prod_{i \in S} (x_i - r)$ (see Section 2.3 for details). Most of the literature focuses on the case $r = 0$, in which $\chi_S(\mathbf{x}, 0)$ reduces to the parity of the variables indexed by S , and $\hat{f}(S, 0)$ is commonly denoted by $\hat{f}(S)$. The point is that whenever $\hat{f}(S, r) \neq 0$, then all variables x_i with $i \in S$ are relevant. If we pursue the search by starting with singletons S and then move on to higher levels, this method takes time about n^s if $\hat{f}(S, r) = 0$ for all S of size up to s . The question is how to proceed if such a situation occurs for some $s \in \omega(1)$. In [23] it is proposed to use a second approach based on the calculation of the coefficients of the polynomial representation of f over the two-element field and shown a trade-off between s and the degree of this polynomial. In a different direction, Atıcı and Servedio [3] enhance the uniform PAC model by a quantum subroutine to circumvent exhaustive search for non-zero Fourier coefficients. Our solution to give the learner access to several (classical) oracles can be considered as another (and maybe more realistic) alternative.

From a conceptual viewpoint, our main result shows that the junta learning problem is efficiently solvable in a *passive* learning model (as opposed to allowing the learner to actively ask membership queries) with *independent random examples* (as opposed to learning from, say, random walks, where examples are highly correlated).

1.2 Our Results

We solve the problem of vanishing Fourier coefficients up to high levels by considering Fourier coefficients with respect to multiple distributions: we show that if all Fourier coefficients $\hat{f}(S, r_i)$ of a k -junta f vanish up to level s with respect to t different biases r_1, \dots, r_t , then $s \cdot t < \deg(f)$, where

$$\deg(f) = \max \left\{ |S| \mid \hat{f}(S) \neq 0 \right\} \leq k$$

is the *degree* of f . Specifically, we prove

Theorem 1. *Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be non-constant function and $s, t \in \mathbb{N}$ be such that $s \cdot t \geq \deg(f)$. Let $r_1, \dots, r_t \in (-1, 1)$ be arbitrary pairwise different biases. Then there exists an $i \in [t]$ and a set $S \subseteq [n]$ with $1 \leq |S| \leq s$ such that $\hat{f}(S, r_i) \neq 0$.*

Letting $s = 1$ and $t = k$, Theorem 1 implies that there are at most $k - 1$ different biases r such that all r -biased first-level Fourier coefficients of f vanish. As a consequence, whenever a learner has access to k r -biased oracles for k pairwise distinct biases r , it suffices to consider, for each given bias r , only coefficients $\hat{f}(S, r)$ at all singletons S in order to find at least one relevant variable. The main technical issue we have to take care of is that Theorem 1 does not rule out the possibility that $|\hat{f}(S, r_i)|$ could be extremely small, so that it would require a large amount of examples to tell whether a coefficient is nonzero. To take this into account, we add the requirement that the biases are well separated, i.e., have pairwise distance at least $\gamma > 0$. In addition, we allow the running time to also depend on the inverse of the minimum distance of the biases to -1 or 1 since the degenerate cases $r = -1$ or $r = 1$ only produce the single example $(\bar{r}, f(\bar{r}))$, from which we cannot learn anything. Here \bar{r} denotes the vector with all n entries equal to r . Our main learning theory application of Theorem 1 (in the special case $s = 1$ and $t = k$) is:

Theorem 2. *Let $-1 + \alpha \leq r_1 < \dots < r_k \leq 1 - \alpha$ for some $\alpha > 0$ such that for all $i \in [k - 1]$, $r_{i+1} - r_i \geq \gamma > 0$. Then the class of k -juntas is exactly learnable with access to r_i -biased oracles, $i \in [k]$, from $m = \text{poly}(\log n, 2^k, (1/\gamma)^k, 1/\alpha, \log(1/\delta))$ examples in time $\text{poly}(m, n)$.*

Theorem 2 immediately follows from the following generalization which is based on the general case ($s \cdot t \geq k$) in Theorem 1. The trade-off between the number of r -biased oracles to which a learner has access and the level up to which the learner has to inspect the Fourier coefficients results in a trade-off between the number of oracles and the running time:

Theorem 3. *Let $k, s, t \in \mathbb{N}$ such that $s \cdot t \geq k$ and $-1 + \alpha \leq r_1 < \dots < r_t \leq 1 - \alpha$ for some $\alpha > 0$ such that for all $i \in [t - 1]$, $r_{i+1} - r_i \geq \gamma > 0$. Then the class of k -juntas is exactly learnable with access to r_i -biased oracles, $i \in [t]$, using $m = \text{poly}(\log n, 2^k, (1/\gamma)^k, (1/\alpha)^s, \log(1/\delta))$ examples and running in time $n^s \cdot \text{poly}(m, n)$.*

In other words, given access to t biased oracles with biases separated by $\gamma > 0$, the class of k -juntas is learnable in time $n^{k/t} \text{poly}(n, 2^k, \gamma^{-k})$. We should mention that we must have $\gamma \geq 2/t$ to be able to separate t biases, so that $\gamma^{-k} \geq (t/2)^k$. If $t = k$, the running time is thus at least polynomial in $2^{k \log k}$.

Theorems 2 and 3 are valid even if the biases are not known to the learner in advance. This follows since given the promise that the examples are generated according to r -biased product distributions, the learner can efficiently approximate these biases to within high accuracy (even from unlabeled examples) and working with such approximate biases is sufficient to recognize non-vanishing Fourier coefficients of the true biases (see Section 6).

It is observed in [23] that except for a set of measure zero of product distributions with bias vectors $\mathbf{r} = (r_1, \dots, r_n) \in [-1, 1]^n$ (i.e., $\mathbb{E}[x_i] = r_i$), every k -junta f has nonzero correlation with each of its relevant variables. They concluded that for each such vector of biases, k -juntas are learnable with confidence $1 - \delta$ in time $\text{poly}(2^k, n, \log(1/\delta))$. However, the correlations may become arbitrarily small, so that in order to identify nonzero correlations, these have to be approximated very precisely. As a consequence, the growth of the poly expression heavily depends on the bias vector \mathbf{r} . More precisely, the running time depends on $2^{c \cdot k}$, where the constant c depends on the choice of \mathbf{r} .

When we restrict the product distributions to r -biased distributions, we can improve from a set of measure zero of exceptional bias vectors to *finitely* many exceptional biases: for fixed k and arbitrary n , there are only finitely many *critical biases* $r \in (-1, 1)$ such that there exists a k -junta f with $\hat{f}(i, r) = 0$ for all $i \in [n]$. As an application, we show

Theorem 4. *Let $k \in \mathbb{N}$. Then for all but finitely many biases $r \in (-1, 1)$, there exists a function $t_r : \mathbb{N} \rightarrow \mathbb{N}$ such that k -juntas are exactly learnable under the r -biased distribution in time $t_r(k) \cdot \text{poly}(n, \log(1/\delta))$.*

Note that, unlike this rather non-constructive result, our algorithm for the “multiple-oracles model” works for k arbitrary and unknown biased product distributions.

1.3 Our methods

Denote by $\mathbb{E}_r[f]$ the expected value of $f(\mathbf{x})$ under the r -biased distribution ($r \in (-1, 1)$). Our main technical tool is a formula that connects the higher-order derivatives of $\mathbb{E}_r[f]$ with respect to r to the Fourier weights at certain levels of the Fourier spectrum. The formula is close in spirit to Russo’s well-known formula for monotone functions and generalizations thereof to arbitrary bounded functions on the hypercube.

Russo’s formula [24] states that for *monotone* Boolean functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$,

$$\frac{d}{dr} \mathbb{E}_r[f] = \mathbb{E}_r[f \cdot \sum_{i=1}^n x_i].$$

More generally, the following connection between the derivative of the expectation (with respect to the bias) and correlations between the function value and the variables is known (see Grimmett [15, Theorem 2.34]):

$$\frac{d}{dr} \mathbb{E}_r[f] = (1 - r^2)^{-1} \text{Cov}_{\mathbf{x} \sim \mu_r}[f(\mathbf{x}), \sum_{i=1}^n x_i] \quad (1)$$

(here, we have translated Grimmett’s notation to our setting, and Cov denotes the covariance). Since $\hat{f}(i, r) = \sigma^{-1} \text{Cov}_r[f, x_i - r] = \sigma^{-1} \text{Cov}_r[f, x_i]$ (see Section 2.3), (1) can be rewritten as

$$\frac{d}{dr} \mathbb{E}_r[f] = (1 - r^2)^{-1/2} \sum_{i=1}^n \hat{f}(i, r). \quad (2)$$

Define the *weight* $w_s(f, r)$ of the s -th r -biased Fourier level of f as the sum of all r -biased Fourier coefficients at level s , i.e.,

$$w_s(f, r) = \sum_{S \subseteq [n]: |S|=s} \hat{f}(S, r).$$

We use the following generalization of formula (2) which we attribute to folklore (and to the best of our knowledge, has not been published before).

Theorem 5 (Generalization of Russo’s Formula). *Let $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, $s \in [n]$, and $r_* \in (-1, 1)$. Then*

$$\left. \frac{d^s}{dr^s} \mathbb{E}_r[f] \right|_{r=r_*} = \frac{s!}{(1 - r_*^2)^{s/2}} \cdot w_s(f, r_*).$$

Theorem 5 follows from a similar statement for product distributions with arbitrary biases (see Proposition 1). The second ingredient to prove Theorem 1 is the observation that we can write

$$\mathbb{E}_r[f] = \sum_{t=0}^n w_t(f, 0) r^t \quad (3)$$

(see Section 2) and that this is a polynomial in r of degree at most $\deg(f)$. Moreover, this polynomial is constant (in r) if and only if f is constant. From Theorem 5, we obtain that if for some r_* , the Fourier coefficients $\hat{f}(S, r_*)$ vanish for all $S \subseteq [n]$ with $1 \leq |S| \leq s$, then $(d^t/dr^t) \mathbb{E}_r[f] \Big|_{r=r_*} = 0$ for all $t \in [s]$,

i.e., r_* is an s -fold root of the nonzero polynomial $(d/dr)\mathbb{E}_r[f]$, which is of degree at most $\deg(f) - 1$. Since there can be at most $(\deg(f) - 1)/s$ roots of multiplicity s , this proves Theorem 1. To the best of our knowledge, this is the first application of Theorem 5 in theoretical computer science. Let us remark further that we obtain the following relationship between Fourier weights with respect to different measures as a consequence of Theorem 5 and Equation (3):

$$w_s(f, r) = (1 - r^2)^{s/2} \sum_{t=s}^n \binom{t}{s} w_t(f, 0) r^{t-s}. \quad (4)$$

1.4 Related Work

If we restrict ourselves to subclasses of k -juntas $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ such as monotone or symmetric juntas (i.e., juntas invariant under permutations of the relevant variables), there do exist at least partially satisfying solutions to the junta learning problem: under the uniform distribution, monotone k -juntas are learnable in time $\text{poly}(n, 2^k)$ from $\text{poly}(\log n, 2^k)$ examples [23] and symmetric juntas are learnable in time $n^{O(k/\log k)} \text{poly}(n, 2^k)$ [20, 22]. Furthermore, results for other learning more general classes under fixed product distributions have been obtained [14, 16, 25, 9], including the polynomial time learnability of monotone $O(\log^2 n / \log^2 \log n)$ -juntas. Notably, also parity juntas, i.e., parities of subsets of at most k variables, are efficiently learnable from product distributions (even in the presence of attribute and classification noise), with the restriction that every variable has a non-zero bias [2].

Recently, Atıcı and Servedio [3] have studied the junta learning problem for the case that the learner has access to a uniform distribution PAC oracle plus a quantum oracle. They showed that k -juntas are learnable within accuracy ϵ from $O(\epsilon^{-1}k \log k)$ quantum examples and $O(2^k \log(1/\epsilon))$ classical (uniformly distributed) examples, both bounds being independent of n . Given this dramatic speed-up (which is impossible to achieve from classical queries only), we ask the more realistic question what can be done if we are given access to *multiple classical* oracles.

Interestingly, our results are obtained in terms of purely statistical evaluation of the given data, i.e., one can interpret the Fourier algorithm as a statistical query (SQ) algorithm with respect to several distributions. While in the original SQ model [18], in which queries are evaluated with respect to the uniform distribution on the input space, (parity) juntas are provably not efficiently learnable [5, 8, 23], our results show that such a lower bound is not valid if queries are evaluated with respect to several distributions.

1.5 Organization of this Paper

We introduce all necessary prerequisites in Section 2. In Section 3 we present the generalization of Russo’s formula. The reduction to identifying only one relevant variable is shown in Section 4. In Section 5, we prove Theorem 3 that addresses learnability via the s -th level Fourier algorithm from several oracles. Section 6 shows that the biases do not have to be known in advance. Finally, we prove Theorem 4 in Section 7.1 and propose open problems in Section 7.2.

2 Preliminaries

2.1 General Notation, Juntas, and Probability Theory

Let $\mathbb{N} = \{0, 1, 2, \dots\}$, and for $n \in \mathbb{N}$, let $[n] = \{1, \dots, n\}$. We use boldface letters such as \mathbf{r} , \mathbf{x} , and $\boldsymbol{\sigma}$ to denote (real) vectors of length n . The corresponding entries are denoted by r_i , x_i , σ_i , and so forth. For $\mathbf{x} \in \{-1, +1\}^n$ and $i \in [n]$, denote by $\mathbf{x}^{(i)}$ the vector \mathbf{x} with the sign of the i -th entry flipped.

Definition 1 (Relevant variables). Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$. For $i \in [n]$, the function f depends on variable x_i (equivalently, x_i is relevant to f) if there exists an $\mathbf{x} \in \{-1, 1\}^n$ such that $f(\mathbf{x}^{(i)}) \neq f(\mathbf{x})$.

Definition 2 (Junta). Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and $k \in [n]$. The function f is a k -junta if it depends on at most k variables.

Let x_1, \dots, x_n be independent random variables taking values -1 and $+1$ with $\mathbb{E}[x_i] = r_i \in [-1, 1]$. The value r_i is called the *bias* of x_i . Equivalently, $\Pr[x_i = -1] = (1 - r_i)/2$ and $\Pr[x_i = 1] = (1 + r_i)/2$. In this way, $\{-1, 1\}^n$ is equipped with the product measure $\mu_{\mathbf{r}}$, $\mathbf{r} = (r_1, \dots, r_n)$, given by

$$\mu_{\mathbf{r}}(\mathbf{x}) = \prod_{i=1}^n ((1 + r_i x_i)/2)$$

for $\mathbf{x} \in \{-1, 1\}^n$. For $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, we denote by $\mathbb{E}_{\mathbf{r}}[f]$ the expectation of f with respect to $\mu_{\mathbf{r}}$. Furthermore, for $f, g : \{-1, 1\}^n \rightarrow \mathbb{R}$, let

$$\text{Cov}_{\mathbf{r}}[f, g] = \mathbb{E}_{\mathbf{r}}[(f - \mathbb{E}_{\mathbf{r}}[f])(g - \mathbb{E}_{\mathbf{r}}[g])] = \mathbb{E}_{\mathbf{r}}[f \cdot g] - \mathbb{E}_{\mathbf{r}}[f] \cdot \mathbb{E}_{\mathbf{r}}[g]$$

denote the covariance of f and g with respect to $\mu_{\mathbf{r}}$. Denote by $\sigma_i = (1 - r_i^2)^{1/2}$ the standard deviation of x_i and let $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$. We will mostly be interested in the case that all biases r_i are equal. For $r \in [-1, 1]$, let $\bar{\mathbf{r}} = (r, \dots, r)$ be the vector that consists of n entries that are all equal to r . In this case, we write $\sigma = \sigma(r) = \sqrt{1 - r^2}$. We will frequently use that if $|r| \leq 1 - \alpha$ for some $\alpha > 0$, then $\sigma \geq \sqrt{\alpha}$. The measure $\mu_{\bar{\mathbf{r}}}$ is called the r -biased product distribution. We also write μ_r instead of $\mu_{\bar{\mathbf{r}}}$, \mathbb{E}_r instead of $\mathbb{E}_{\bar{\mathbf{r}}}$, etc.

2.2 Learning Theory

We introduce an extension of the classical PAC model [27]. Let $\mathcal{C} = \bigcup_{n \in \mathbb{N}} \mathcal{C}_n$ be a class of functions, where each \mathcal{C}_n contains some functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and let $\mathcal{M} = \bigcup_{n \in \mathbb{N}} \mathcal{M}_n$ be a class of input distributions, where each \mathcal{M}_n contains distributions on $\{-1, 1\}^n$. For $f \in \mathcal{C}_n$ and a distribution $\mu \in \mathcal{M}_n$, denote by $EX(f, \mu)$ an oracle that on request generates $\mathbf{x} \in \{-1, 1\}^n$ according to μ and returns the example $(\mathbf{x}, f(\mathbf{x}))$. For $r \in [-1, 1]$, we call $EX(f, \mu_r)$ an r -biased oracle. Let us first review the original PAC model. The class \mathcal{C} is PAC-learnable under distributions \mathcal{M} if there is an algorithm \mathcal{A} that for all $n \in \mathbb{N}$, all functions $f \in \mathcal{C}_n$, and all distributions $\mu \in \mathcal{M}_n$ on $\{-1, 1\}^n$, given $\delta, \epsilon > 0$ and access to $EX(f, \mu)$ but no further knowledge on f and μ , outputs a hypothesis $h : \{-1, 1\}^n \rightarrow \{-1, 1\}$ such that with probability at least $1 - \delta$ (taken over all random draws of the oracle), $\Pr_{\mathbf{x} \sim \mu}[h(\mathbf{x}) \neq f(\mathbf{x})] \leq \epsilon$. If \mathcal{M}_n is the class of all distributions on $\{-1, 1\}^n$, we say that \mathcal{C} is distribution-free PAC-learnable. If \mathcal{M}_n only contains the uniform distribution on $\{-1, 1\}^n$, we say that \mathcal{C} is uniform distribution PAC-learnable. If \mathcal{M}_n is the class of all r -biased product distributions μ_r on $\{-1, 1\}^n$, we say that \mathcal{C} is learnable from biased distributions. If \mathcal{A} even manages to output exactly f , (i.e., $\epsilon = 0$), we say that \mathcal{C} is exactly learnable.

The performance of a learning algorithm is measured by the number of examples it requests and by its running time, both of course depending on δ, ϵ, n , and possibly further parameters involved in the definition of the class \mathcal{C} .

Now we study what happens if, instead of having access to a single oracle $EX(f, \mu)$, we admit the learning algorithm to have access to multiple (pairwise different) oracles $EX(f, \mu_i)$, $\mu_i \in \mathcal{M}_n$ for $i \in [t]$. If we do not impose any restrictions other than being pairwise different on the distributions μ_i , then the learner does not gain any power since the distributions could be arbitrarily close to each other. Thus, we allow the running time to depend on the minimum distance γ between pairs of distributions (at this point, we leave open the choice of appropriate distance measures).

The notion of learnability is the same as above, except that we require that the hypothesis output by a learning algorithm has to satisfy with probability at least $1 - \delta$ that $\Pr_{x \sim \mu_i}[h(x) \neq f(x)] \leq \epsilon$ for all $i \in [t]$. In this case, we say that \mathcal{C} is *PAC-learnable from t oracles under distributions \mathcal{M} with separation γ* .

In the following, we motivate in which variants of this very general new learning model we are interested. Our goal is to find efficient learning algorithms for the class of k -juntas. More precisely, for a non-decreasing function $k : \mathbb{N} \rightarrow \mathbb{N}$, we want to learn the class $\mathcal{C} = \bigcup_{n \in \mathbb{N}} \mathcal{C}_n$, where \mathcal{C}_n consists of all $k(n)$ -juntas $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$. The fastest known (exact) learning algorithm for \mathcal{C} in the uniform distribution PAC-learning model runs in time $n^{0.7k} \text{poly}(n, 2^k, \log(1/\delta))$ [23]. Moreover, for $k \in \omega(1)$, there is *not any* explicit distribution μ for which \mathcal{C} is known to be PAC-learnable under μ in time $t(k) \cdot \text{poly}(n, \log(1/\delta))$ with an arbitrary function $t : \mathbb{N} \rightarrow \mathbb{N}$. It thus seems reasonable to ask if we can do any better if we are given access to more than one oracle with several *simple* distributions (possibly known to the learner). We will show that this is in fact the case if the distributions are biased product distributions μ_{r_i} with well-separated biases r_i , even without prior knowledge on the biases (except that each $|r_i|$ should be bounded away from 1). Consequently, we manage to learn efficiently in the model of *PAC-learning from multiple biased product distributions*. The separation of biases will be reflected in the dependence of the running time on $\gamma = \min_{i \neq j} |r_i - r_j|$.

2.3 Fourier Coefficients

For $\mathbf{t} \in \mathbb{R}^n$ and $S \subseteq [n]$, define $\mathbf{t}_S = \prod_{i \in S} t_i$. In particular, for $\mathbf{x} \in \{-1, 1\}^n$, \mathbf{x}_S is the *parity* of bits in \mathbf{x} indexed by S , and for $\mathbf{r} \in [-1, 1]^n$, $\mathbb{E}_{\mathbf{r}}[\mathbf{x}_S] = \prod_{i \in S} \mathbb{E}_{\mathbf{r}}[x_i] = \mathbf{r}_S$. For $i \in [n]$ and $\mathbf{r} \in (-1, 1)^n$, define $\chi_i(\mathbf{x}, \mathbf{r}) = (x_i - r_i)/\sigma_i$ and for $S \subseteq [n]$, let $\chi_S(\mathbf{x}, \mathbf{r}) = \prod_{i \in S} \chi_i(\mathbf{x}, \mathbf{r})$.

The measure $\mu_{\mathbf{r}}$ induces the inner product

$$\langle f, g \rangle_{\mathbf{r}} = \mathbb{E}_{\mathbf{r}}[f \cdot g] = \sum_{\mathbf{x} \in \{-1, 1\}^n} \mu_{\mathbf{r}}(\mathbf{x}) f(\mathbf{x}) g(\mathbf{x})$$

on $\mathbb{R}^{\{-1, 1\}^n}$. The associated norm is

$$\|f\|_{2, \mathbf{r}} = \langle f, f \rangle_{\mathbf{r}}^{1/2} = \mathbb{E}_{\mathbf{r}}[f^2]^{1/2}.$$

The functions $\chi_S = \chi_S(\cdot, \mathbf{r})$, $S \subseteq [n]$, form an orthonormal basis of this space with respect to $\langle \cdot, \cdot \rangle_{\mathbf{r}}$:

$$\langle \chi_S, \chi_S \rangle_{\mathbf{r}} = \mathbb{E}_{\mathbf{r}}[\chi_S^2] = \prod_{i \in S} \frac{\mathbb{E}_{\mathbf{r}}[(x_i - r_i)^2]}{\sigma_i^2} = 1,$$

and if $i \in S \setminus T$ for some sets $S, T \subseteq [n]$, then $\mathbb{E}_{\mathbf{r}}[\chi_S \chi_T] = \mathbb{E}_{\mathbf{r}}[\chi_i] \mathbb{E}_{\mathbf{r}}[\chi_{S \setminus \{i\}} \chi_T] = 0$ since $\mathbb{E}_{\mathbf{r}}[\chi_i] = 0$.

We can expand any function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ as a linear combination of the functions $\chi_S(\cdot, \mathbf{r})$, called the *Fourier expansion* of f with respect to $\mu_{\mathbf{r}}$:

$$f = \sum_{S \subseteq [n]} \langle f, \chi_S \rangle_{\mathbf{r}} \chi_S,$$

and we call $\hat{f}(S, \mathbf{r}) = \langle f, \chi_S \rangle_{\mathbf{r}}$ the *Fourier coefficient of f at S with respect to $\mu_{\mathbf{r}}$* . Note that $\chi_i(\cdot, \mathbf{r})$ is a linear function in x_i and thus $\chi_S(\cdot, \mathbf{r})$ is a multi-linear polynomial in the variables x_i , $i \in S$ (of degree $|S|$). Consequently, the Fourier expansion (with respect to any $\mathbf{r} \in (-1, 1)^n$) provides a representation of f as a real multi-linear polynomial of degree

$$\deg(f, \mathbf{r}) = \max\{k \in [n] \mid \exists S \subseteq [n] : |S| = k \wedge \hat{f}(S, \mathbf{r}) \neq 0\}.$$

Since this degree does actually not depend on \mathbf{r} (there is exactly one polynomial representation of f), we let $\deg(f) = \deg(f, \bar{0})$.

If $S = \{i\}$ is a singleton set, we also write $\hat{f}(i, r)$ instead of $\hat{f}(\{i\}, r)$. Note that for $S \neq \emptyset$, $\hat{f}(S, \mathbf{r}) = \text{Cov}_{\mathbf{r}}[f, \chi_S(\cdot, \mathbf{r})]$ since $\mathbb{E}_{\mathbf{r}}[\chi_S(\cdot, r)] = 0$. Put in another way,

$$\sigma^{|S|} \cdot \hat{f}(S, \mathbf{r}) = \text{Cov}_{\mathbf{r}} \left[f, \prod_{i \in S} (x_i - r_i) \right].$$

In case we consider the r -biased product measure for some $r \in (-1, 1)$, we call $\hat{f}(S, r) = \hat{f}(S, \bar{r})$ an r -biased Fourier coefficient. In particular, $\hat{f}(\emptyset, r) = \langle f, 1 \rangle_r = \mathbb{E}_r[f]$ (again using r as subscripts rather than \bar{r}). For the uniform measure $\mu_{\bar{r}}$ with $r = 0$, the Fourier expansion of f directly results in the representation of f as a real multilinear polynomial in canonical form (i.e., a linear combination of monomials \mathbf{x}_S) since $\chi_S(\mathbf{x}, 0) = \mathbf{x}_S$: $f(\mathbf{x}) = \sum_{S \subseteq [n]} \hat{f}(S, 0) \cdot \mathbf{x}_S$. Since for $\mathbf{r} \in [-1, 1]^n$, $\mathbb{E}_{\mathbf{r}}[\mathbf{x}_S] = \mathbf{r}_S$, we obtain

$$\mathbb{E}_{\mathbf{r}}[f] = \sum_{S \subseteq [n]} \hat{f}(S, 0) \mathbf{r}_S, \quad (5)$$

of which (3) is the special case $\mathbf{r} = \bar{r}$ for $r \in [-1, 1]$.

The *weight* of the i -th Fourier level of f with respect to $\mu_{\mathbf{r}}$ is defined to be

$$w_i(f, \mathbf{r}) = \sum_{S \subseteq [n]: |S|=i} \hat{f}(S, \mathbf{r}).$$

Lemma 1. *Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$. If $\sum_{i=1}^n w_i(f, 0) = 0$, then f is constant.*

Proof. We have

$$f(1^n) = \sum_{S \subseteq [n]} \hat{f}(S, 0) = \sum_{i=0}^n w_i(f, 0)$$

is either 1 or -1 . Thus, if $\sum_{i=1}^n w_i(f, 0) = 0$, then $|\hat{f}(\emptyset, 0)| = |w_0(f, 0)| = 1$, i.e., $f \equiv 1$ or $f \equiv -1$. \square

The connection between juntas and Fourier coefficients is given by the following characterization of relevant variables:

Lemma 2 ([2, 23]). *Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, $r \in (-1, 1)$, and $i \in [n]$. Then x_i is relevant to f if and only if there exists $S \subseteq [n]$ with $i \in S$ and $\hat{f}(S, r) \neq 0$.*

In particular, if $\hat{f}(S, r) \neq 0$ for some $S \subseteq [n]$ and some $r \in (-1, 1)$, then all variables x_i , $i \in S$, are relevant to f . Thus, one way to find relevant variables is to look for non-vanishing Fourier coefficients. Furthermore, if f is a k -junta, then $\hat{f}(S, r) = 0$ for all S with $|S| > k$, i.e., looking at coefficients up to level k is sufficient for finding all relevant variables.

2.4 Sampling Fourier Coefficients

To approximate biased Fourier coefficients, we will make use of the Hoeffding bound [17]:

Fact 1 (Hoeffding bound, [17]). *Let X_i , $i \in [m]$, be mutually independent random variables taking values in $[a, b]$, $a < b$. Then for any $\epsilon \in [0, 1]$,*

$$\Pr \left[\left| \sum_{i=1}^m X_i - \sum_{i=1}^m \mathbb{E}[X_i] \right| \geq \epsilon m \right] \leq 2 \exp \left(\frac{-2m\epsilon^2}{(b-a)^2} \right).$$

Lemma 3. Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, $r \in (-1, 1)$, $S \subseteq [n]$, and $\delta > 0$. Given access to $EX(f, r)$, we can estimate $\hat{f}(S, r)$ within accuracy $\epsilon > 0$ from $m = \text{poly}(2^{|S|}, (1/\sigma)^{|S|}, \log(1/\delta), 1/\epsilon)$ examples in time $O(m \cdot n)$ with confidence $1 - \delta$, provided that r is given exactly.

Proof. Draw $m = 2 \cdot \ln(2/\delta) \cdot (2^{|S|}/\epsilon)^2 \cdot (1/\sigma)^{2|S|}$ examples $(\mathbf{x}^t, f(\mathbf{x}^t))$ from $EX_r(f)$. Define $\Delta = (\max_{x_i \in \{-1, 1\}} |x_i - r|)^{|S|} = (1 + |r|)^{|S|} \leq 2^{|S|}$. Let $g(\mathbf{x}) = \sigma^{|S|} f(\mathbf{x}^t) \chi_S(\mathbf{x}^t, r) \in [-\Delta, \Delta]$. Then, by Fact 1,

$$\left| \frac{1}{m} \sum_{t=1}^m g(\mathbf{x}^t) - \sigma^{|S|} \hat{f}(S, r) \right| \leq \epsilon \sigma^{|S|}$$

with probability at least $1 - \delta$. \square

We will deal with the case that r is not exactly given in advance in Section 6. To distinguish the cases $\hat{f}(S, i) = 0$ and $\hat{f}(S, i) \neq 0$, we also need that a non-vanishing $\hat{f}(S, i)$ is not too small. For this, we will use the following (straightforward) lemma:

Lemma 4. Let $h \in \mathbb{R}[x]$ be a polynomial of degree d with leading coefficient b and roots $t_1, \dots, t_d \in \mathbb{C}$. Let $t \in \mathbb{R}$ and $\epsilon > 0$ such that $|t - \text{Re } t_i| \geq \epsilon$ for all $i \in [d]$. Then $|h(t)| \geq |b| \cdot \epsilon^d$.

Proof. Since $h(x) = b \cdot \prod_{i \in [d]} (x - t_i)$, $|h(t)| = |b| \cdot \prod_{i \in [d]} |t - t_i| \geq |b| \cdot \prod_{i \in [d]} |t - \text{Re } t_i| \geq |b| \cdot \epsilon^d$. \square

2.5 Derivatives

For a k -fold differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $S = \{i_1, \dots, i_k\} \subseteq [n]$ with pairwise different elements i_j , denote by $\frac{\partial^k}{\partial \mathbf{x}_S} f = \frac{\partial^k}{\partial x_{i_1} \dots \partial x_{i_k}} f$ the k -th order partial derivative with respect to x_{i_1}, \dots, x_{i_k} .

Lemma 5. Let $g \in \mathbb{R}[t_1, \dots, t_n]$ be a multilinear polynomial (i.e., all exponents are at most one) and define $h \in \mathbb{R}[t]$ by $h(t) = g(t, \dots, t)$. Then

$$\frac{d^k}{dt^k} h(t) = k! \cdot \sum_{S \subseteq [n]: |S|=k} \frac{\partial^k}{\partial t_S} g(t, \dots, t).$$

Proof. The easy way to see the claim is to simply apply the chain rule. For multi-linear polynomials, though, we can as well check the claim “by hand”: By linearity of the construction of h , it suffices to check the claim for the case that g is a monomial. Without loss of generality, assume that $g(t_1, \dots, t_n) = t_1 \dots t_\ell$. Let $S \subseteq [n]$ with $|S| = k$. If $S \not\subseteq [\ell]$, then clearly $(\partial^k / \partial t_S) g = 0 = (d^k / dt^k) h$. If $S \subseteq [\ell]$, then $(\partial^k / \partial t_S) g(t, \dots, t) = t^{\ell-k}$, so that

$$k! \cdot \sum_{S \subseteq [n]: |S|=k} \frac{\partial^k}{\partial t_S} g(t, \dots, t) = k! \cdot \binom{\ell}{k} t^{\ell-k} = \frac{\ell!}{(\ell-k)!} t^{\ell-k}.$$

On the other hand, $h(t) = t^\ell$ and thus $(d^k / dt^k) h(t) = \ell \cdot (\ell - 1) \cdot \dots \cdot (\ell - k + 1) \cdot t^{\ell-k} = \frac{\ell!}{(\ell-k)!} t^{\ell-k}$. \square

3 An Extension of Russo’s Formula to General Product Distributions and Higher Order Derivatives

In this section, we derive our connection between derivatives or $\mathbb{E}_{\mathbf{r}}[f]$ and Fourier levels. In particular, we prove Theorem 5 stated in Section 1.

Proposition 1. Let $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, $S \subseteq [n]$ with $|S| = k$, and $\mathbf{r}^* \in (-1, 1)^n$. Then

$$\frac{\partial^k}{\partial \mathbf{r}_S} \mathbb{E}_{\mathbf{r}}[f] \Big|_{\mathbf{r}=\mathbf{r}^*} = \prod_{i \in S} (1 - r_i^{*2})^{-1} \cdot \text{Cov}_{\mathbf{r}^*} \left[f, \prod_{i \in S} (x_i - r_i^*) \right] = \boldsymbol{\sigma}_S^{*-1} \cdot \hat{f}(S, \mathbf{r}^*).$$

Proof. Expanding f with respect to $\mu_{\mathbf{r}^*}$, we see that

$$\mathbb{E}_{\mathbf{r}}[f] = \sum_{S \subseteq [n]} \hat{f}(S, \mathbf{r}^*) \mathbb{E}_{\mathbf{r}}[\chi_S(\cdot, \mathbf{r}^*)] = \sum_{S \subseteq [n]} \hat{f}(S, \mathbf{r}^*) \prod_{i \in S} \frac{r_i - r_i^*}{\sigma_i^*}$$

is simply the Taylor expansion of the multi-linear polynomial $\mathbb{E}_{\mathbf{r}}[f]$, and the claim follows. \square

Putting together Proposition 1 and Equation (5), we obtain the relationship

$$\hat{f}(S, \mathbf{r}) = \boldsymbol{\sigma}_S \sum_{T \supseteq S} \hat{f}(T, 0) \mathbf{r}_{T \setminus S}. \quad (6)$$

Theorem 5 in the introduction now follows from Proposition 1 and Lemma 5, and (4) is a special case of (6).

4 Identifying One Relevant Variable Is Enough

In analogy to Proposition 6 in [23], we show that if we have an algorithm that identifies just *one* relevant variable of a non-constant k -junta f using $m = \text{poly}(\log n, 2^k, (1/\alpha)^k, \log(1/\delta))$ examples from $EX(f, r_1), \dots, EX(f, r_t)$ (where $\alpha > 0$ bounds away the biases r_i from -1 and 1) in time $n^{\beta k} \text{poly}(m, n)$, then we can construct an algorithm that identifies *all* relevant variables and outputs the truth table of f using $m' = t \cdot \text{poly}(\log n, 2^k, (1/\alpha)^k, \log(1/\delta))$ examples in time $n^{\beta k} \text{poly}(m', n)$ (for the same β , but a different polynomial):

Proposition 2. Let \mathcal{A} be an algorithm that, given access to $EX(f, r_1), \dots, EX(f, r_t)$ for some non-constant k -junta $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and some $r \in (-1 + \alpha, 1 - \alpha)$ ($\alpha > 0$) and given $\delta > 0$, outputs with probability at least $1 - \delta$ one relevant variable of f using $m = \text{poly}(\log n, 2^k, (1/\alpha)^k, \log(1/\delta))$ examples in time $n^{\beta k} \cdot \text{poly}(m, n)$. Then there is an algorithm \mathcal{B} that, for any k -junta $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, given access to $EX(f, r_1), \dots, EX(f, r_t)$ and $\delta > 0$, outputs with probability at least $1 - \delta$ all relevant variables and a truth table of f , using $m' = t \cdot \text{poly}(\log n, 2^k, (1/\alpha)^k, \log(1/\delta))$ examples in time $n^{\beta k} \text{poly}(m', n)$.

Proof. The proposition can be proved by an adaption of the proof of Proposition in [23], so we only point to the necessary modifications of the latter. First, if f is non-constant, then each output value $f(x)$ is drawn from $EX_{r_i}(f)$ with frequency at least $(\min\{(1 - r_i)/2, (1 + r_i)/2\})^k \geq (\alpha/2)^k$. Thus, the check for constancy with confidence δ requires $O((2/\alpha)^k \log(1/\delta))$ examples and $\text{poly}((2/\alpha)^k, n, \log(1/\delta))$ steps.

Next, for restrictions $f|_{\rho}$ of f fixing at most k variables, each simulation of a draw from $EX_{r_i}(f|_{\rho})$ requires the draw of $O((2\alpha)^k \log(m/\delta))$ from $EX_{r_i}(f)$.

Since \mathcal{A} is run at most $k2^k$ times with confidence $1 - \delta/(k2^k)$ each, it suffices to draw

$$O(m(2/\alpha)^k \log(mk2^k/\delta)) = m \log(m/\delta) \text{poly}(2^k, (1/\alpha)^k)$$

examples from each oracle $EX(f, r_i)$ (note that \mathcal{A} run on different restrictions may ask m examples from different oracles).

Finally, to read off a truth table of f from the examples, $\text{poly}((2/\alpha)^k, (1/\delta))$ examples (from any of the oracles) are again sufficient to ensure with probability $1 - \delta$ that every possible assignment of the relevant variables appears in the examples. The claim follows since $m = \text{poly}(\log n, 2^k, (1/\alpha)^k, \log(1/\delta))$. \square

5 Learning Relevant Variables via the s -th Level Fourier Algorithm

The goal of this section is to prove Theorem 3. For $s \in [n]$, let

$$\mathcal{R}_s(f) = \{r_0 \in (-1, 1) \mid r_0 = \operatorname{Re}(r) \text{ for some root } r \in \mathbb{C} \text{ of } \frac{d}{dr}\mathbb{E}_r[f] \text{ of multiplicity at least } s\}.$$

By Theorem 5, $\mathcal{R}_s(f)$ contains all $r \in (-1, 1)$ such that $w_1(f, r) = \dots = w_s(f, r) = 0$ and in particular all $r \in (-1, 1)$ for which $\hat{f}(S, r) = 0$ for all $S \subseteq [n]$ of size $1 \leq |S| \leq s$.

Lemma 6. *Let $f : \{-1, 1\}^n \rightarrow \{-1, +1\}$ be a non-constant k -junta, $s \in [k]$, and $r \in (-1, 1)$ such that $\operatorname{dist}(r, \mathcal{R}_s(f)) \geq \gamma > 0$. Then there exists $S \subseteq [n]$ with $1 \leq |S| \leq s$ such that $|\hat{f}(S, r)| \geq \sigma^s(\gamma/4)^k$. In particular, all variables x_i with $i \in S$ are relevant.*

Proof. Let $r_0 = r$. Let $g(r) = \mathbb{E}_r[f]$. By (3) and Lemma 1, g is a non-constant polynomial of degree $d = \deg(g) \leq \deg(f) \leq k$ with leading coefficient $w_d(f, 0)$. Let $t \geq 1$ be minimal with $(d^t/dr^t)g|_{r=r_0} \neq 0$. Since $r_0 \notin \mathcal{R}_s(f)$, $t < s$. Let $h = (d^t/dr^t)g$. Then h is a non-zero polynomial of degree $d - t \leq \deg(f) - t \leq k - t$. The highest coefficient of h is $b = d \cdot (d-1) \cdot \dots \cdot (d-t+1) \cdot w_d(f, 0)$. By Lemma 4, $|h(r_0)| \geq |b| \cdot \gamma^{d-t}$. Since w_d is a non-zero integer multiple of 2^{-k} , $|b| \geq \frac{d!}{(d-t)!} 2^{-k}$. By Theorem 5, $h(r_0) = t! \sigma^{-t} w_t(f, r_0)$, so that

$$|w_t(f, r_0)| = (t!)^{-1} \sigma^t |h(r_0)| \geq \binom{d}{t} 2^{-k} \sigma^t \gamma^{d-t}.$$

Hence there exists $S \subseteq [n]$ with $|S| = t$ such that

$$|\hat{f}(S)| \geq \binom{k}{t}^{-1} \binom{d}{t} 2^{-k} \sigma^t \gamma^{d-t} \geq \binom{k}{d}^{-1} (\gamma/2)^k \sigma^s \geq (\gamma/4)^k \sigma^s.$$

□

For the remainder of this section, we assume that a learning algorithm has exact knowledge of all biases. However, we will show in Section 6 that this assumption is not necessary.

Proposition 3. *There is an algorithm such that if $f : \{-1, 1\}^n \rightarrow \{-1, +1\}$ is a non-constant k -junta and $r \in (-1 + \alpha, 1 - \alpha)$ (for some $\alpha > 0$) is such that $\operatorname{dist}(r, \mathcal{R}_s(f)) \geq \gamma$ for some $\gamma > 0$, having access to the oracle $EX(f, r)$, for any $\delta > 0$ outputs at least one relevant variable of f with probability at least $1 - \delta$ using $m = \operatorname{poly}(\log n, 2^k, (1/\gamma)^k, (1/\alpha)^s, \log(1/\delta))$ examples and running in time $n^s \cdot \operatorname{poly}(m, n)$. Furthermore, for arbitrary $r \in [-1 + \alpha, 1 - \alpha]$, with probability at least $1 - \delta$, any variable output by the algorithm is relevant.*

Proof. By Lemma 6, there exists $S \subseteq [n]$ with $1 \leq |S| \leq s$ such that $|\hat{f}(S, r)| \geq \sigma^s(\gamma/4)^k \geq \alpha^{s/2}(\gamma/4)^k$. Thus, it suffices to estimate all coefficients $\hat{f}(S, r)$, $S \subseteq [n]$ with $1 \leq |S| \leq s$, within accuracy $\alpha^{s/2}(\gamma/4)^k/2$, each with confidence $1 - \delta \cdot n^{-s}$, to identify (with probability at least $1 - \delta$) at least one S such that $\hat{f}(S, r) \neq 0$ with confidence $1 - \delta$. This takes $\operatorname{poly}(2^{|S|}, (1/\alpha)^{|S|}, \log(n^s/\delta), (4/\gamma)^k(1/\alpha)^s)$ examples from the oracle $EX(f, r)$ by Lemma 3, and we can reuse the same examples to estimate all coefficients (since we use a union bound for the confidence). Overall, the number of examples used is

$$m = \operatorname{poly}(\log n, 2^k, (1/\gamma)^k, (1/\alpha)^s, \log(1/\delta)).$$

The algorithm outputs all variables x_i for which it finds a nonzero Fourier coefficient $\hat{f}(S)$ with $i \in S$. Since we have to check $\sum_{i=1}^s \binom{n}{i} = O(n^s)$ coefficients in the worst-case, the running time is bounded above by $n^s \cdot \operatorname{poly}(m, n)$.

For the second part of the claim, note that if $\hat{f}(S, r) = 0$ (and especially, if S contains an index i of some non-relevant variable), then the estimate for $|\hat{f}(S, r)|$ will with high probability be smaller than $\alpha^{s/2}(\gamma/4)^k/2$. \square

Theorem 6. *Let $s, t \in [k]$ such that $s \cdot t \geq k$, $\alpha, \gamma > 0$, and $-1 + \alpha \leq r_1 < \dots < r_t \leq 1 - \alpha$ with $r_{j+1} - r_j \geq \gamma$ for all $j \in [t - 1]$. Then there is an algorithm that, for any non-constant k -junta $f : \{-1, 1\}^n \rightarrow \{-1, +1\}$, given $\delta > 0$ and having access to the oracles $EX(f, r_1), \dots, EX(f, r_t)$, outputs a relevant variable of f with probability at least $1 - \delta$, using $m = \text{poly}(\log n, 2^k, (1/\gamma)^k, (1/\alpha)^s, \log(1/\delta))$ examples and running in time $n^s \cdot \text{poly}(m, n)$.*

Proof. Let $h(r) = w_1(f, r)/\sigma = (d/dr)\mathbb{E}_r[f]$. Since h is a nonzero polynomial of degree at most $\deg(f) - 1 \leq k - 1$ and since $s \cdot t \geq k$, h has less than t roots of multiplicity at least s . Consequently, there exists $j \in [t]$ such that $\text{dist}(r_j, \mathcal{R}_s(f)) \geq \gamma/2$. Running the algorithm from Proposition 3 for every single bias r_j , $j \in [t]$, (each time with confidence parameter δ/t , reusing the same examples) yields the claim. \square

Proof of Theorem 3. Theorem 6 shows that it is possible to identify at least one relevant variable from the claimed number of examples in time $n^s \cdot \text{poly}(m, n)$. By Proposition 2, the claim follows. \square

We note that since $h(r)$ is of degree at most $\deg(f) - 1$, it actually suffices to have $s \cdot t \geq d$ oracles if we are given the promise that $\deg(f) \leq d$.

6 Biases Unknown in Advance

The algorithms provided in Section 5 require that all biases r_i are precisely known to the learner. As one might expect, this assumption is not necessary since a learner can get good estimates of the biases from (unlabeled) random examples. The main technical issue is now to show that using good estimates r'_i still leads to sufficiently close approximations of the Fourier coefficients with respect to the true biases r_i . For this it suffices to show that $\chi_S(\cdot, r_i)$ and $\chi_S(\cdot, r'_i)$ are close in L^2 .

Lemma 7. *Let $\alpha, \gamma > 0$, $r, r' \in (-1, 1)$ such that $|r| \leq 1 - \alpha$ and $|r - r'| \leq \gamma$, $S \subseteq [n]$. Then*

$$\|\chi_S(\cdot, r') - \chi_S(\cdot, r)\|_{2,r} \leq \frac{|S| + 1}{\alpha^{1/2}\sigma^s} \gamma.$$

To prove Lemma 7, we will first compute, given $\mathbf{r}, \mathbf{r}' \in (-1, 1)^n$, the Fourier coefficients of $\chi_S(\cdot, \mathbf{r}')$ with respect to $\mu_{\mathbf{r}}$. Although we only need $\mathbf{r} = \bar{r}$ and $\mathbf{r}' = \bar{r}'$ for our applications, we state the result for general bias vectors \mathbf{r} and \mathbf{r}' since the proof does not simplify for the special case.

Lemma 8. *Let $\mathbf{r}, \mathbf{r}' \in (-1, 1)^n$ and $S, T \subseteq [n]$. Then*

$$\widehat{\chi_S(\cdot, \mathbf{r}')}_{(T, \mathbf{r})} = \langle \chi_S(\cdot, \mathbf{r}'), \chi_T(\cdot, \mathbf{r}) \rangle_{\mathbf{r}} = \begin{cases} 0 & \text{if } T \not\subseteq S \\ \frac{\sigma_T}{\sigma_S} (\mathbf{r} - \mathbf{r}')_{S \setminus T} & \text{if } T \subseteq S. \end{cases}$$

Proof. We have $\mathbb{E}_{\mathbf{r}}[\chi_i(\cdot, \mathbf{r}')] = (r_i - r'_i)/\sigma'_i$, $\mathbb{E}_{\mathbf{r}}[\chi_i(\cdot, \mathbf{r})] = 0$, and $\mathbb{E}_{\mathbf{r}}[\chi_i(\cdot, \mathbf{r}) \cdot \chi_i(\cdot, \mathbf{r}')] = \sigma_i/\sigma'_i$. The claim now follows from

$$\begin{aligned} \langle \chi_S(\cdot, \mathbf{r}'), \chi_T(\cdot, \mathbf{r}) \rangle_{\mathbf{r}} &= \mathbb{E}_{\mathbf{r}}[\chi_S(\cdot, \mathbf{r}') \cdot \chi_T(\cdot, \mathbf{r})] \\ &= \prod_{i \in S \setminus T} \mathbb{E}_{\mathbf{r}}[\chi_i(\cdot, \mathbf{r}')] \cdot \prod_{i \in T \setminus S} \mathbb{E}_{\mathbf{r}}[\chi_i(\cdot, \mathbf{r})] \cdot \prod_{i \in S \cap T} \mathbb{E}_{\mathbf{r}}[\chi_i(\cdot, \mathbf{r}) \cdot \chi_i(\cdot, \mathbf{r}')] . \end{aligned}$$

\square

Now we bound the L^2 -norm of the difference between $\chi_S(\cdot, r)$ and $\chi_S(\cdot, r')$. Here we do restrict ourselves to $r = \bar{r}$ and $r' = \bar{r}'$ to avoid an increase in technicality:

Proof of Lemma 7. By Parseval's equation,

$$\|\chi_S(\cdot, r) - \chi_S(\cdot, r')\|_{2,r}^2 = \sum_{T \subseteq [n]} \left(\widehat{\chi_S(\cdot, r')}(T, r) - \widehat{\chi_S(\cdot, r)}(T, r) \right)^2.$$

By Lemma 8 and since $\widehat{\chi_S(\cdot, r)}(T, r) = 0$ unless $T = S$, all summands for $T \not\subseteq S$ vanish. Furthermore, Lemma 8 states that for $T \subseteq S$, $\widehat{\chi_S(\cdot, r')}(T, r) = (\sigma^t/\sigma'^s)(r - r')^{s-t}$, where we let $s = |S|$ and $t = |T|$. Thus,

$$\begin{aligned} \|\chi_S(\cdot, r') - \chi_S(\cdot, r)\|_{2,r}^2 &\leq (\widehat{\chi_S(\cdot, r')}(S, r) - 1)^2 + \sum_{T \subsetneq S} \widehat{\chi_S(\cdot, r')}(T, r)^2 \\ &= (\sigma^s/\sigma'^s - 1)^2 + \sum_{t=0}^{s-1} \binom{s}{t} ((\sigma^t/\sigma'^s)\gamma^{s-t})^2 \\ &= (\sigma')^{-2s} [(\sigma^s - \sigma'^s)^2 + (\sigma^2 + \gamma^2)^s - \sigma^{2s}]. \end{aligned}$$

Now we use the following two facts:

Fact 2. For any $a, b \in [0, 1]$ with $|b - a| \leq \rho$, $|a^s - b^s| \leq s \cdot \rho$.

Proof. Let $a < b$. Then by convexity of the function $x \mapsto x^s$, $b^s \leq a^s + sb^{s-1}(b - a) \leq a^s + s\delta$. \square

Fact 3. If $|r' - r| \leq \gamma$, then $|\sigma' - \sigma| \leq \gamma/\sigma$.

Proof. Let $\sigma(r) = \sqrt{1 - r^2}$. The derivative of σ is $(d/dr)\sigma(r) = -\frac{r}{\sigma(r)}$. Since σ is concave, we have that for any δ such that $r, r + \delta \in (-1, 1)$, $\sigma(r + \delta) \leq \sigma(r) + (d/dr)\sigma(r)\delta = \sigma(r) - r\delta/\sigma(r)$. Since $|r| \leq 1$, the claim follows with $r' = r + \delta$, $|\delta| \leq \gamma$, $\sigma' = \sigma(r')$, and $\sigma = \sigma(r)$. \square

Let $\rho = \gamma\alpha^{-1/2}$. By Fact 3 and since $\sigma^2 = 1 - r^2 \geq 1 - r \geq \alpha$, $|\sigma' - \sigma| \leq \rho$. From Fact 2, we obtain $|\sigma'^s - \sigma^s| \leq s\rho$ and $(\sigma^2 + \gamma^2)^s - (\sigma^2)^s \leq s\gamma^2$. Consequently,

$$\sigma'^{2s} \|\chi_S(\cdot, r') - \chi_S(\cdot, r)\|_{2,r}^2 \leq (s\rho)^2 + s\gamma^2 = s^2\gamma^2/\alpha + s\gamma^2 \leq (s + 1)^2\gamma^2/\alpha.$$

This proves the lemma. \square

As a corollary, we obtain an estimate of how well $\langle f, \chi(\cdot, r') \rangle_r$ approximates $\hat{f}(S, r)$:

Corollary 1. Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, $\gamma > 0$, $r, r' \in (-1, 1)$ such that $|r' - r| \leq \gamma$, and $S \subseteq [n]$. Then

$$\left| \langle f, \chi_S(\cdot, r') \rangle_r - \hat{f}(S, r) \right| \leq \frac{|S| + 1}{\alpha^{1/2}\sigma'^{|S|}} \gamma.$$

Proof. By Cauchy-Schwartz,

$$\left| \langle f, \chi(\cdot, r') \rangle_r - \hat{f}(S, r) \right| = \left| \langle f, \chi_S(\cdot, r') - \chi_S(\cdot, r) \rangle_r \right| \leq \|f\|_{2,r} \|\chi_S(\cdot, r') - \chi_S(\cdot, r)\|_{2,r}.$$

The claim follows since $\|f\|_{2,r} = 1$. \square

Next we show how to closely approximate $\hat{f}(S, r)$ given no a priori knowledge on r :

Lemma 9. Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, $\alpha > 0$, $r \in [-1 + \alpha, 1 - \alpha]$, $S \subseteq [n]$, and $\delta > 0$. Given access to $EX(f, r)$, we can estimate $\hat{f}(S, r)$ within accuracy ϵ from $m = \text{poly}(2^{|S|}, (1/\alpha)^{|S|}, \log(1/\delta), 1/\epsilon)$ examples in time $O(m \cdot n)$ with confidence $1 - \delta$ without any a priori knowledge on r .

Proof. Let $\gamma = \alpha^{(|S|+1)/2} / (2(|S| + 1)) \leq \alpha^{1/2} \sigma^{|S|} / (2|S| + 1)$, so that, in particular, $\gamma \leq \alpha/2 \leq \sigma^2/2$ (note that we may assume $|S| \geq 1$ without loss of generality). First, we approximate r to within γ by requesting $m_1 = 8 \ln(4/\delta) / \gamma^2 = \text{poly}(|S|^2, (1/\alpha)^{|S|}, \log(1/\delta))$ examples $(\mathbf{x}^t, f(\mathbf{x}^t))$ from $EX(f, r)$ to compute $r' = (1/m_1) \sum_{t=1}^{m_1} x_i^t$. With probability at least $\delta/2$, $|r' - r| \leq \gamma$.

Now, letting $g(\mathbf{x}) = \sigma^{|S|} f(\mathbf{x}^t) \chi_S(\mathbf{x}^t, r')$, $\phi = (m_2 \sigma^{|S|})^{-1} \sum_{t=1}^{m_2} g(\mathbf{x}^t)$ approximates $\langle f, \chi_S(\cdot, r') \rangle_r$ within accuracy $\epsilon/2$ given $m_2 = \text{poly}(2^{|S|}, (1/\sigma')^{2|S|}, \log(1/\delta), 1/\epsilon)$ examples. Since

$$\sigma' \geq \sigma - \gamma/\sigma \geq \sigma/2 \geq \alpha^{1/2}/2$$

implies $(1/\sigma')^{2|S|} \leq (4/\alpha)^{|S|}$, m_2 is dominated by $\text{poly}(2^{|S|}, (1/\alpha)^{|S|}, \log(1/\delta), 1/\epsilon)$. Finally,

$$|\phi - \hat{f}(S, r)| \leq |\phi - \langle f, \chi_S(\cdot, r') \rangle_r| + |\langle f, \chi_S(\cdot, r') \rangle_r - \hat{f}(S, r)| \leq \epsilon/2 + (|S| + 1) \alpha^{-1/2} \sigma'^{-|S|} \gamma \leq \epsilon.$$

The total number of examples to be drawn is $\max\{m_1, m_2\}$, which is of the order indicated in the claim. \square

Using Lemma 9 in place of Lemma 3 shows that Proposition 3, Theorem 6, and finally also Theorems 2 and 3 even hold if the biases r_i are not known in advance (except for the bound $|r_i| \leq 1 - \alpha$).

7 Further Results and Open Problems

7.1 Learning in Polynomial Time for All But Finitely Many Biases

We have seen that for each k -junta f , there are at most $k - 1$ biases in $(-1, 1)$ for which $w_1(f, r) = 0$. Since for the r -biased product measure, $w_1(f, r)$ does not depend on *where* the relevant variables are hidden, it is not hard to see that there are at most $(k - 1) \cdot 2^{O(k^2)}$ biases for which there exists *some* k -junta f (for any n) with $w_1(f, r) = 0$. Let us call these biases *critical*. Let \mathcal{S}_k denote the set of biases $r \in (-1, 1)$ such that there exists a function $t_r : \mathbb{N} \rightarrow \mathbb{N}$ and a k -junta-learning algorithm that learns from $EX(f, r)$ in time $t_r(k) \cdot \text{poly}(n)$. Then \mathcal{S}_k is exactly the complement of the critical points. This is because the minimum distance between any two distinct critical points is a function of k only. This proves Theorem 4 stated in the introduction. Consequently, for each k , there are only finitely many biases for which junta-learning may not be feasible in time polynomial in n . The next step (left for future research) is to find lower bounds on $t_r(k)$.

Generalizing to arbitrary product distributions with bias vector $\mathbf{r} \in (-1, 1)^n$, we obtain that $w_1(f, \mathbf{r})$ is zero only for a set of biases of measure zero (since it is the zero set of a non-constant multi-linear polynomial). Considering the polynomials $\sigma \hat{f}(i)$ separately for each $i \in [n]$, we recover the statement of [23] that $\hat{f}(i, \mathbf{r}) = 0$ for all $i \in [n]$ only for a set of measure zero.

7.2 Open Problems

Next to the notoriously hard problem of designing more efficient algorithms for the junta learning problem under the uniform distribution, it would also constitute considerable progress to have, for any concretely given fixed bias $r \neq 0$, *some* algorithm improving over the n^k bound. Note that we have shown in Section 7.1 that for all but finitely many r , the degree-one algorithm works. However, it is not clear how to decide in general whether a given bias is critical. We believe that the relationship (3) between Fourier coefficients with respect to different biases could be useful to this end.

In a different direction, it seems worthwhile to further study our newly introduced model of learning from multiple oracles. Can we show positive results for other learning problems that appear to be hard in the

classical PAC setting? In particular, is there an efficient algorithm for learning DNFs or decision trees from multiple distributions? What general conditions on the distributions are required to make efficient learning possible? As the *number* of oracles obviously constitutes a significant resource parameter, it is natural to ask if polynomial time learning of juntas is also possible from $o(k)$ oracles (maybe at least for important subclasses).

7.3 Acknowledgement

Jan Arpe is supported by the Postdoc-Program of the German Academic Exchange Service (DAAD) and by NSF Career Award DMS 0548249 and BSF 2004105. Elchanan Mossel is supported by NSF Career Award DMS 0548249, BSF 2004105 and DOD ONR grant N0014-07-1-05-06.

References

- [1] Hussein Almuallim and Thomas G. Dietterich. Learning Boolean Concepts in the Presence of Many Irrelevant Features. *Artificial Intelligence*, 69(1-2):279–305, September 1994.
- [2] Jan Arpe and Rüdiger Reischuk. Learning Juntas in the Presence of Noise. *Theoretical Computer Science*, 384:2–21, 2007.
- [3] Alp Atıcı and Rocco A. Servedio. Quantum Algorithms for Learning and Testing Juntas. *Quantum Inf. Process.*, 6(5):323–348, October 2007.
- [4] Avrim Blum. Learning a Function of r Relevant Variables. In Bernhard Schölkopf and Manfred K. Warmuth, editors, *Computational Learning Theory and Kernel Machines, 16th Annual Conference on Computational Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003, Proceedings*, volume 2777 of *Lecture Notes in Computer Science*, pages 731–733. Springer, 2003.
- [5] Avrim Blum, Merrick Furst, Jeffrey C. Jackson, Michael Kearns, Yishay Mansour, and Steven Rudich. Weakly Learning DNF and Characterizing Statistical Query Learning Using Fourier Analysis. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 253–262, 1994.
- [6] Avrim Blum and Pat Langley. Selection of Relevant Features and Examples in Machine Learning. *Artificial Intelligence*, 97(1-2):245–271, December 1997.
- [7] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik-Chervonenkis Dimension. *J. ACM*, 36(4):929–965, October 1989.
- [8] Nader H. Bshouty and Vitaly Feldman. On Using Extended Statistical Queries to Avoid Membership Queries. *J. Mach. Learn. Res.*, 2(3):359–396, August 2002.
- [9] Nader H. Bshouty and Christino Tamon. On the Fourier Spectrum of Monotone Functions. *J. ACM*, 43(4):747–770, July 1996.
- [10] Koby Crammer, Michael Kearns, and Jennifer Wortman. Learning from Data of Variable Quality. In *Advances in Neural Information Processing Systems 18 [Neural Information Processing Systems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada]*. MIT Press, 2005.

- [11] Koby Crammer, Michael Kearns, and Jennifer Wortman. Learning from Multiple Sources. In *Advances in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 4-7, 2006 (NIPS '06)*, pages 321–328. MIT Press, 2006.
- [12] Daniel R. Dooley, Qi Zhang, Sally A. Goldman, and Robert A. Amar. Multiple-Instance Learning of Real-Valued Data. *J. Mach. Learn. Res.*, 3:651–678, December 2002.
- [13] Jon Feldman, Ryan O’Donnell, and Rocco A. Servedio. Learning Mixtures of Product Distributions over Discrete Domains. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 501–510. IEEE Press, 2005.
- [14] Merrick L. Furst, Jeffrey C. Jackson, and Sean W. Smith. Improved Learning of AC^0 Functions. In Leslie G. Valiant and Manfred K. Warmuth, editors, *Proceedings of the Fourth Annual Workshop on Computational Learning Theory (COLT 1991), Santa Cruz, California, USA*, pages 317–325. Morgan Kaufmann, 1991.
- [15] Geoffrey Grimmett. *Percolation*. Grundlehren Math. Wiss. Springer, 2nd edition, 1999.
- [16] Thomas R. Hancock and Yishay Mansour. Learning Monotone k - μ DNF Formulas on Product Distributions. In Leslie G. Valiant and Manfred K. Warmuth, editors, *Proceedings of the Fourth Annual Workshop on Computational Learning Theory (COLT 1991), Santa Cruz, California, USA*, pages 179–183. Morgan Kaufmann, 1991.
- [17] Wassily Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *J. Amer. Statist. Assoc.*, 58:13–30, 1963.
- [18] Michael Kearns. Efficient Noise-Tolerant Learning from Statistical Queries. *J. ACM*, 45(6):983–1006, November 1998.
- [19] Michael Kearns, Yishay Mansour, Dana Ron, Ronitt Rubinfeld, Robert E. Schapire, and Linda Sellie. On the Learnability of Discrete Distributions. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing, Montreal, Quebec, Canada (STOC '94)*, pages 273–282. ACM Press, 1994.
- [20] Mihail N. Kolountzakis, Evangelios Markakis, and Aranyak Mehta. Learning Symmetric Juntas in Time $n^{o(k)}$. In *Workshop on Interface between Harmonic Analysis and Number Theory, Marseille, 2005*, 2005. Available as Tech. Rep. at <http://arxiv.org/abs/math.CO/0504246v1>.
- [21] Charles X. Ling and Qiang Yang. Discovering Classification from Data of Multiple Sources. *Data Min. Knowl. Discov.*, 12(2-3):181–201, 2006.
- [22] Richard J. Lipton, Evangelos Markakis, Aranyak Mehta, and Nisheeth K. Vishnoi. On the Fourier Spectrum of Symmetric Boolean Functions with Applications to Learning Symmetric Juntas. In *20th Annual IEEE Conference on Computational Complexity (CCC '05)*, pages 112–119, 2005.
- [23] Elchanan Mossel, Ryan W. O’Donnell, and Rocco A. Servedio. Learning functions of k relevant variables. *J. Comput. System Sci.*, 69(3):421–434, November 2004.
- [24] Lucio Russo. On the Critical Percolation Probabilities. *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, 56:229–237, 1981.

- [25] Rocco A. Servedio. On learning monotone DNF under product distributions. *Inform. and Comput.*, 193(1):57–74, August 2004.
- [26] György Turán. Lower Bounds for PAC Learning with Queries. In Leslie G. Valiant and Manfred K. Warmuth, editors, *Proceedings of the Sixth Annual ACM Conference on Computational Learning Theory (COLT 1993), July 26-28, 1993, Santa Cruz, CA, USA*, pages 384–391. ACM, 1993.
- [27] Leslie G. Valiant. A Theory of the Learnable. *Commun. ACM*, 27(11):1134–1142, November 1984.
- [28] Santosh Vempala and Grant Wang. A Spectral Algorithm for Learning Mixtures of Distributions. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, page 113. IEEE Press, 2002.