

Degree- d Chow Parameters Robustly Determine Degree- d PTFs (and Algorithmic Applications)

Ilias Diakonikolas*

University of Southern California
diakonik@usc.edu

Daniel M. Kane†

University of California, San Diego
dakane@cs.ucsd.edu

November 7, 2018

Abstract

The degree- d Chow parameters of a Boolean function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ are its degree at most d Fourier coefficients. It is well-known that degree- d Chow parameters uniquely characterize degree- d polynomial threshold functions (PTFs) within the space of all bounded functions. In this paper, we prove a robust version of this theorem: For f any Boolean degree- d PTF and g any bounded function, if the degree- d Chow parameters of f are close to the degree- d Chow parameters of g in ℓ_2 -norm, then f is close to g in ℓ_1 -distance. Notably, our bound relating the two distances is completely independent of the dimension n . That is, we show that Boolean degree- d PTFs are *robustly identifiable* from their degree- d Chow parameters. Results of this form had been shown for the $d = 1$ case [OS11, DDFS14], but no non-trivial bound was previously known for $d > 1$.

Our robust identifiability result gives the following algorithmic applications: First, we show that Boolean degree- d PTFs can be efficiently approximately reconstructed from approximations to their degree- d Chow parameters. This immediately implies that degree- d PTFs are efficiently learnable in the uniform distribution d -RFA model [BDD98]. As a byproduct of our approach, we also obtain the first low integer-weight approximations of degree- d PTFs, for $d > 1$. As our second application, our robust identifiability result gives the first efficient algorithm, with dimension-independent error guarantees, for malicious learning of Boolean degree- d PTFs under the uniform distribution.

The proof of our robust identifiability result involves several new technical ingredients, including the following structural result for degree- d multivariate polynomials with very poor anti-concentration: If p is a degree- d polynomial where $p(x)$ is *very* close to 0 on a *large* number of points in $\{-1, 1\}^n$, then there exists a degree- d hypersurface that exactly passes through *almost all* of these points. We leverage this structural result to show that if the degree- d Chow distance between f and g is small, then we can find many degree- d polynomials that vanish on their disagreement region, and in particular enough that forces the ℓ_1 -distance between f and g to also be small. To implement this proof strategy, we require additional technical ideas. In particular, in the $d = 2$ case we show that for any large vector space of degree-2 polynomials with a large number of common zeroes, there exists a linear function that vanishes on almost all of these zeroes. The degree- d degree generalization of this statement is significantly more complex, and can be viewed as an effective version of Hilbert's Basis Theorem for our setting.

*Supported by NSF Award CCF-1652862 (CAREER) and a Sloan Research Fellowship.

†Supported by NSF Award CCF-1553288 (CAREER) and a Sloan Research Fellowship.

1 Introduction

This paper is concerned with the problem of reconstructing degree- d polynomial threshold functions over the Boolean hypercube from their degree at most d Fourier coefficients. Before we state our main contributions, we provide some background and motivation for this work.

1.1 Background

A *degree- d polynomial threshold function (PTF)* is a Boolean function of the form $f(x) = \text{sign}(p(x))$, where $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ is a degree- d polynomial with real coefficients. The function $\text{sign}(z)$ takes value 1 for $z \geq 0$ and -1 otherwise. For degree $d = 1$, degree- d PTFs are referred to as Linear Threshold Functions (LTFs) or Boolean Halfspaces. Degree- d PTFs are a fundamental family of Boolean functions that have been extensively studied in a number of scientific disciplines for at least six decades [Ros58, Cho61, MTT61, Der65, MP68, SRK95]. During the past decade, there has been renewed research attention on degree- d PTFs from several perspectives of theoretical computer science, including complexity theory [Ser07, RS10, DGJ⁺10, DKN10, DS13, MZ13, Kan14, DDS14, DS14, GKM15, Kan17] and computational learning theory [DHK⁺10, DOSW11, OS11, FGRW12, DDFS14, DRST14, HKM14, Dan15, Dan16, DDS17, ABL17, DKS18, BGS18].

In this work, we study the problem of efficiently approximating degree- d PTFs over $\{-1, 1\}^n$ from their Fourier coefficients of degree at most d , which we will call degree- d Chow parameters:

Definition 1. Given any function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, its *degree- d Chow Parameters* are its Fourier coefficients of degree at most d , i.e., $\widehat{f}(S) = \mathbf{E}_{x \sim \mathcal{U}_n}[f(x)\chi_S(x)]$, for all $S \subseteq [n]$ with $|S| \leq d$, where \mathcal{U}_n denotes the uniform distribution on $\{-1, 1\}^n$. We say that the *degree- d Chow vector* of f is $\vec{\chi}_f^d = (\widehat{f}(S))_{S \subseteq [n], |S| \leq d}$.

In general, if the degree d is relatively small, the degree- d Chow parameters of a Boolean function f provide limited information about the function. Perhaps surprisingly, this is not the case if we know that f is an LTF or, more generally, a degree- d PTF. A classical result of C.K. Chow [Cho61] shows that Boolean LTFs are uniquely specified by their degree-1 Chow parameters. Chow's result was later generalized by Bruck [Bru90] to the class of degree- d PTFs. Specifically, [Cho61, Bru90] establish the following:

Degree- d Chow's Theorem: For any $d \geq 1$, if $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ is any degree- d PTF and $g : \{-1, 1\}^n \rightarrow [-1, 1]$ is any bounded function such that $\vec{\chi}_f^d = \vec{\chi}_g^d$, then $f(x) = g(x)$ for all $x \in \{-1, 1\}^n$.

Unfortunately, the original proof of Chow's theorem [Cho61] (and its straightforward generalization to the degree- d case [Bru90]) is non-constructive, and in particular does not suggest any algorithm to reconstruct (even approximately) a degree- d PTF from its degree- d Chow parameters. This naturally suggests the following computational problem:

Degree- d Chow Parameters Problem: Given (approximations to) the degree- d Chow parameters of an unknown degree- d PTF over $\{-1, 1\}^n$, output an (approximate) representation of f as $\text{sign}(p(x))$, where $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ is a degree- d polynomial.

The (degree-1) Chow Parameters problem has a rich history and has been extensively studied since the 1960s. Specifically, researchers in various communities have been interested in finding an efficient algorithm for the problem, including electrical engineering [Elg60, MTK62, Win64, MTB67, Kas63, Win63, KW65, Der65, Win69, Bau73, Hur73], game theory and voting theory [Lap72, Pen46, Ban65, DS79, EL89, TZ92, Fre97, Lee03, Car04, FM04, TT06, APL07, LW98, Lee02a, Lee02b, dKKZ10, Kur12, KN12], and computational learning theory [BDJ⁺98, BDD98, Gol06, Ser07, OS11, DDFS14]. More recently, Chow's theorem and the Chow parameters have played an important role in various aspects of complexity theory (see, e.g., [CHIS10, KW16]). The reader is referred to [OS11] for a detailed summary of previous work.

The first provably efficient algorithm for the degree-1 Chow parameters problem was obtained by O’Donnell and Servedio [OS11]: they gave a polynomial time algorithm that, given sufficiently accurate approximations to the degree-1 Chow parameters of an unknown LTF f , it outputs the weights-based representation of an LTF h that is close to f in (normalized) Hamming distance. In subsequent work, De, Diakonikolas, Feldman, and Servedio [DDFS14] gave a significantly faster algorithm for the (degree-1 Chow parameters) problem. (See Section 1.3 for a more detailed description of prior work.) The degree- d Chow parameters problem for $d > 1$ has remained open. Prior to this work, no non-trivial upper bound was known, even for special cases of degree-2 PTFs.

To facilitate the subsequent discussion, we introduce some basic notation. For $f, g : \{-1, 1\}^n \rightarrow \mathbb{R}$ the *distance* between f and g is $\text{dist}(f, g) \stackrel{\text{def}}{=} \mathbf{E}_x[|f(x) - g(x)|]$, where the underlying distribution will be the uniform distribution on $\{-1, 1\}^n$, unless explicitly stated otherwise. If $\text{dist}(f, g) \leq \epsilon$, we say that f and g are ϵ -close. (Note that if f, g are Boolean-valued, then $\text{dist}(f, g) = 2\Pr_x[f(x) \neq g(x)]$.) The degree- d Chow parameters also naturally induce a distance measure between functions over the Boolean hypercube:

Definition 2. Let $f, g : \{-1, 1\}^n \rightarrow \mathbb{R}$. We define the *degree- d Chow distance* between f and g to be $\text{Chow}_d(f, g) \stackrel{\text{def}}{=} \|\vec{\chi}_f^d - \vec{\chi}_g^d\|_2$, i.e., the Euclidean distance between their degree- d Chow vectors.

(A useful equivalent reformulation is $\text{Chow}_d(f, g) = \sup_{p \text{ degree at most } d, \|p\|_2=1} \mathbf{E}_x[p(x)(f(x) - g(x))]$, where the supremum is taken over all normalized multilinear polynomials of degree at most d .) Using this terminology, the degree- d Chow’s theorem can be rephrased as follows: If f is a degree- d PTF and g is a bounded function such that $\text{Chow}_d(f, g) = 0$, then $\text{dist}(f, g) = 0$. An immediate question that arises when thinking about this problem is to what extent is the degree- d Chow’s theorem *robust*: In particular, if $\text{Chow}_d(f, g)$ is small, does this necessarily imply that f and g are close? Or equivalently, suppose that $\text{dist}(f, g) = \epsilon > 0$, where ϵ is a small universal constant. Is it the case that $\text{Chow}_d(f, g)$ cannot be too small? We note that Chow’s original argument does not establish any non-trivial robustness.

For the case of LTFs ($d = 1$), a sequence of works [BDJ⁺98, Gol06, Ser07, OS11, DDFS14] established robust versions of the degree-1 Chow’s theorem with varying quantitative guarantees. In particular, for f, g with $\text{dist}(f, g) = \epsilon$, [OS11, DDFS14] showed lower bounds on $\text{Chow}_1(f, g)$ that only depend on ϵ and are independent of n . Similarly to the algorithmic version of the problem, the existence of a robust version of the degree- d Chow’s theorem for $d > 1$ was one of the main open questions in [DDFS14] and has remained unresolved. Prior to this work, no non-trivial bound was known, even for degree-2 PTFs.

It turns out that the *robustness* question discussed above — a purely structural question — is intimately related to the *algorithmic* question of approximately reconstructing a degree- d PTF from its degree- d Chow parameters. It should be noted that both previous works that provide efficient algorithms [OS11, DDFS14] for the $d = 1$ case establish robust versions of Chow’s theorem and crucially use them for the analysis of their algorithms. The connection between robustness and computationally efficient reconstruction was made explicit in [DDFS14] (see also [TTV08]), where it was established that a sufficiently robust version of the degree- d Chow’s theorem suffices to obtain an efficient approximation algorithm for the problem (see Theorem 37 in Section 4.1 for a precise quantitative version).

1.2 Our Results

The main contribution of this paper is a robust version of the degree- d Chow’s theorem that is completely independent of n . Specifically, we prove the following:

Theorem 3 (Main Result). *There exists a function $\kappa : \mathbb{R} \times \mathbb{N} \rightarrow \mathbb{R}$ such that the following holds: Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be any degree- d PTF and $g : \{-1, 1\}^n \rightarrow [-1, 1]$ be an arbitrary bounded function. If $\text{Chow}_d(f, g) \leq \kappa(\epsilon, d)$, then $\text{dist}(f, g) \leq \epsilon$.*

Some comments are in order: The main conceptual message of Theorem 3 is that the function $1/\kappa(\epsilon, d)$ is *independent* of the number of variables n . Prior to our work, no structural result of this form was known with a sub-exponential dependence on n , even for restricted classes of degree-2 PTFs and $\epsilon = 0.49$.

We note that the growth rate of the function $1/\kappa(\epsilon, d)$ established by our current proof is very large. Specifically, $1/\kappa(\epsilon, d)$ grows like Ackermann($d + O(1), 1/\epsilon$). We believe that the right dependence is quasi-polynomial in $1/\epsilon$ for constant d ¹, though proving such an improved bound would require additional ideas. The correct dependence on d is less clear, but ought to be at least doubly exponential.

Theorem 3 is a natural structural result on the Fourier structure of degree- d PTFs that we believe is of independent interest. Below, we describe a number of algorithmic and structural applications of Theorem 3.

Algorithmic and Structural Applications. Our first algorithmic application is an efficient algorithm for the degree- d Chow parameters problem. Combined with known algorithmic machinery [TTV08, DDFS14], Theorem 3 yields the following:

Theorem 4 (Reconstruction of Degree- d PTFs from Degree- d Chow Parameters). *There is an algorithm that on input ϵ, δ , and a vector $\vec{\alpha}$ satisfying $\|\vec{\alpha} - \vec{\chi}_f^d\|_2 \leq \kappa(\epsilon, d)$, for an unknown degree- d PTF f , has the following behavior: it runs in time $\tilde{O}(n^{2d}) \cdot \text{poly}(1/\kappa(\epsilon, d)) \cdot \log(1/\delta)$ and outputs a vector $(H_S)_{S \subseteq [n], |S| \leq d}$, such that with probability at least $1 - \delta$, the degree- d PTF $h(x) = \text{sign}(\sum_S H_S \prod_{i \in S} x_i)$ satisfies $\Pr_x[f(x) \neq h(x)] \leq \epsilon$.*

In words, we obtain an algorithm for the degree- d Chow parameters problem that, for any constant accuracy ϵ , runs in time $\tilde{O}_d(n^{2d})$. As an immediate corollary of Theorem 4, we obtain an algorithm with similar running time for learning degree- d PTFs in the uniform distribution d -RFA model of Ben-David and Dichterman [BDD98]. In this learning model, the learner can only observe a desired subset of coordinates of each unlabeled example of size at most d . See Section 4.1 for a detailed statement.

The algorithm of Theorem 4 can be shown to output a degree- d PTF with integer weights whose sum of squares is at most $n^d \cdot \text{poly}(1/\kappa(\epsilon, d))$. Hence, we obtain the first non-trivial bounds on approximating arbitrary degree- d PTFs using degree- d PTFs with small integer weights.

Theorem 5 (Low Integer-Weight Approximation for Degree- d PTFs). *Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be a degree- d PTF. There is a degree- d PTF $h(x) = \text{sign}(\sum_{S \subseteq [n], |S| \leq d} H_S \prod_{i \in S} x_i)$ such that $\Pr_x[f(x) \neq h(x)] \leq \epsilon$ and the weights H_S are integers that satisfy $\sum_S H_S^2 = O(n^d) \cdot \text{poly}(1/\kappa(\epsilon, d))$.*

A number of previous works [Ser07, DS13, DDFS14] obtained low integer-weight approximators to LTFs, culminating in the near-optimal² integer weight bound of $O(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))}$ [DDFS14]. For $d > 1$, no non-trivial bound was known prior to our work. We note that [DSTW14] gave low integer-weight approximators for degree- d PTFs, but the degree of the approximating PTF is $(1/\epsilon)^{\Omega(d)}$, as opposed to d .

Our main structural result also has algorithmic implications for the problem of learning Boolean degree- d PTFs in the malicious learning model of Valiant, Kearns and Li [Val85, KL93]. The *malicious noise model* is a generalization of the PAC model in which an adversary can arbitrarily corrupt a small constant fraction of both the unlabeled data points and their labels. Using the machinery of [DKS18], we obtain an algorithm that learns Boolean degree- d PTFs in the presence of a small constant fraction of corrupted data:

Theorem 6 (Learning Boolean Low-Degree PTFs with Nasty Noise). *There is a polynomial-time algorithm for learning Boolean degree- d PTFs in the presence of malicious noise with respect to the uniform distribution on $\{-1, 1\}^n$. Specifically, if $\kappa(\epsilon, d)$ is the noise rate, the algorithm runs in $\text{poly}(n^d, 1/\epsilon)$ time and outputs a hypothesis degree- d PTF $h(x)$ that with high probability satisfies $\Pr_x[h(x) \neq f(x)] \leq \epsilon$, where f is the unknown target PTF.*

¹A quasi-polynomial lower bound is known for $d = 1$ [DDFS14].

²A construction of [Hås94] implies a lower bound of $\max\{n^{1/2}, (1/\epsilon)^{\Omega(\log \log(1/\epsilon))}\}$.

We note that the algorithm establishing Theorem 6 was given in [DKS18]. Our Theorem 3 is the missing technical ingredient to prove correctness of this algorithm for the setting of the uniform distribution on the hypercube. See Section 4.2 for a detailed explanation.

1.3 Related and Prior Work

In this section, we review some relevant prior work on the degree-1 version of the Chow parameters problem [Gol06, Ser07, OS11, DDFS14]. Goldberg [Gol06] showed that for f an n -variable LTF and g any Boolean function, if $\text{dist}(f, g) = \epsilon$ then $\text{Chow}_1(f, g) \geq (\epsilon/n)^{O(\log(n/\epsilon)\log(1/\epsilon))}$. In the same setting, [Ser07] obtained a lower bound of $\text{Chow}_1(f, g) \geq 1/(\text{poly}(n) \cdot 2^{\tilde{O}(1/\epsilon^2)})$ and [OS11] obtained the bound $\text{Chow}_1(f, g) \geq 2^{-\tilde{O}(1/\epsilon^2)}$. Finally, [DDFS14] improved the latter lower bound to $\text{Chow}_1(f, g) \geq \epsilon^{O(\log^2(1/\epsilon))}$, which is the best known bound to date and qualitatively nearly-matches an upper bound of $\epsilon^{O(\log \log(1/\epsilon))}$. Building on their structural result, [OS11] gave an algorithm for the degree-1 Chow parameters problem that finds an ϵ -approximator to the unknown LTF in time $\text{poly}(n) \cdot 2^{2^{\tilde{O}(1/\epsilon^2)}}$. [DDFS14] gave a new algorithm for the problem that, combined with their structural result, was shown to run in time $\text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))}$. The algorithm of [DDFS14] straightforwardly generalizes to the degree- d case, but its analysis hinges on a robust version of the degree- d Chow’s theorem, which we prove in this work.

1.4 Our Techniques

In this section, we provide an overview of our techniques that lead to the proof of Theorem 3 in tandem with a comparison to prior work. We start by reviewing previous approaches that give robust versions of the degree-1 Chow’s theorem. Let f be an LTF and g a Boolean function such that $\text{dist}(f, g) = \epsilon$, $\text{Chow}_1(f, g) = \delta$; we would like to show that δ cannot be too small (as a function of ϵ and, potentially, the dimension n). When one tries to robustify the original proof of Chow [Cho61], one finds that the argument goes through unless the LTF $f(x) = \text{sign}(L(x))$ has $|L(x)|$ very close to 0 on almost all the points where f and g differ. In other words, if $L(x) = w \cdot x - \theta$ is *anti-concentrated* around the origin, i.e., the fraction of points $x \in \{-1, 1\}^n$ such that $|L(x)|$ is very close to 0 is small, then Chow’s argument can be naturally extended. Unfortunately, this is not always the case: it is quite possible that $|L(x)|$ is very close to 0 for a significant fraction of points $x \in \{-1, 1\}^n$, which makes it challenging to robustify Chow’s argument.

Two approaches have been proposed to circumvent the above obstacle. The idea in [OS11] (implicit in [Ser07]) is to approximate an arbitrary LTF by an LTF with “good anti-concentration”. For this idea to work, it is crucial that the normalized Hamming distance between f (the original LTF) and its approximator f' to be very small compared to the anti-concentration radius. While such an approach was shown to be feasible for the degree-1 case, we do not know if it is possible to extend even to the case of degree-2 PTFs.

On the other hand, Goldberg [Gol06] and [DDFS14] (that builds on and substantially strengthens [Gol06]) uses a more direct geometric view of the problem. Roughly speaking, it is shown in [Gol06, DDFS14] that if $L(x)$ has “very poor” anti-concentration, then the linear function $L(x)$ satisfies certain important structural properties. More specifically, suppose that for some moderately large ϵ and very small δ that $\Pr_x[|L(x)| < \delta] > \epsilon$. Then there exists a linear polynomial $L'(x)$ so that all but a tiny fraction of the points x in the disagreement region between f and g satisfy $L'(x) = 0$ — as opposed to $L(x)$, which is very close to 0. By slightly modifying g , we can reduce to the case where *all* of the discrepancies lie on the hyperplane $L'(x) = 0$. This allows us to renormalize L , by taking it modulo L' , and potentially find a second linear function on which nearly all of the discrepancies lie. Repeating this process, we can eventually find a large number of linear functions so that nearly all of the disagreements between f and g lie on the intersection of the corresponding hyperplanes. However, given enough such functions, there will no longer be enough points for this to be the case, yielding a contradiction.

At a high-level, our approach for the degree- d case is a generalization of [Gol06, DDFS14]. Firstly, we note that the robustification of the degree-1 Chow’s result still works for the degree- d case, unless the degree- d polynomial p defining our degree- d PTF $f(x) = \text{sign}(p(x))$ has the same kind of very poor anti-concentration as before (see Claim 15). We will next need to show that this implies that almost all of the disagreements lie on a degree- d polynomial hypersurface (Proposition 14). While [Gol06, DDFS14] accomplish this for $d = 1$ by a careful analysis of the vectors perpendicular to the discrepancy points (in order to get nearly optimal quantitative bounds), our techniques for the degree- d case are less accurate. We first need some sort of general anti-concentration result — which we obtain via a combination of a regularity lemma [DSTW14] and an invariance principle [MOO10] — to show that this kind of bad anti-concentration implies that there exists a small set of coordinates $S \subseteq [n]$, so that upon fixing the variables in S , most of the disagreements between f, g reduce to polynomials in the remaining coordinates with small ℓ_2 -norm (see Corollary 20). We then show that by scaling p , we can get a polynomial that is in some sense “nearly integral”, and establish that the integer part must vanish on almost all of the points where $|p(x)|$ is extremely close to 0 (see Fact 21 and the paragraph preceding it for a more detailed overview).

It turns out that the much more challenging part of our proof is to generalize the iteration of the above result. We have established that if p has very bad anti-concentration, then almost all of its near zeroes lie on the zero-set of a degree- d polynomial. By iterating this fact, it is not hard to show that if f and g have very small degree- d Chow distance, we can find a whole sequence q_1, q_2, \dots of linearly independent, degree- d multilinear polynomials so that almost all disagreements of f and g lie on the joint zero set of the q_i ’s. When $d = 1$, completing the proof from this point would be easy, since at most a 2^{-m} -fraction of points in $\{-1, 1\}^n$ can lie on the joint zeroes of m linear polynomials (see Fact 24).

Unfortunately, even for degree-2 polynomials, this statement is false. For example, consider the sequence of polynomials $q_i = (x_0 + 1)(x_i + 1)$. These are linearly independent, however half of all points (those with $x_0 = -1$) are joint zeroes. This can happen because almost all of our zeroes lie on a hyperplane. In fact, for the degree-2 case, we can show (see Proposition 22) that this is essentially the only thing that can go wrong. In particular, we prove that given sufficiently many linearly independent degree-2 polynomials, we can replace them with a *single* degree-1 polynomial without losing too many disagreements in the zero set. Iterating this result, until we have enough linear polynomials, yields a contradiction as before.

The higher degree case runs along the same lines, however the recursion becomes somewhat more complicated. At each stage of the process, we maintain an *ideal* I of polynomials so that almost all of the disagreements of f and g lie on zeroes of I . We show that if the original degree- d Chow distance was small enough, we can always add another degree- d polynomial to I . From here what we need can be seen as a robustification of Hilbert’s Basis Theorem for our setting. The Hilbert Basis Theorem [Hil90] (see, e.g., [CLO07]) says that, starting with an ideal, if one repeatedly adds new polynomials, this process must eventually terminate (perhaps with I being the unit ideal). Unfortunately, the number of rounds of this iterative process is unbounded, and even under reasonable restrictions, will still depend on the number of variables n . What we establish here is that (1) If the added polynomials are all degree at most d , and (2) If we are allowed to throw sets of negligible mass out of the associated variety, then we can actually obtain an upper bound on the number of steps. To achieve this, we show that we can replace sufficiently many degree- k polynomials by a *single* degree- $(k - 1)$ polynomial without losing too much probability mass. The basic idea of the proof is to use a degree- d version of the Littlewood-Offord lemma [MNV16] to show that if there are many degree- k polynomials with a large number of joint zeroes, then either (1) there is a degree- $(k - 1)$ polynomial which vanishes on almost all of them, or (2) there is a small set of coordinates so that all of the polynomials depend on only these coordinates (and thus the dimension of the space they span is bounded). The above ingredients suffice in order to obtain a contradiction in the degree- d case.

1.5 Organization

The structure of this paper is as follows: In Section 2, we introduce the mathematical background required for our results. Section 3 contains the proof of our main structural result (Theorem 3). In Section 4, we present our algorithmic and structural applications. In Section 5, we conclude with a few open problems.

2 Preliminaries

Notation. We start by establishing basic notation. For $n \in \mathbb{Z}_+$, we write $[n]$ to denote $\{1, 2, \dots, n\}$. We write $\mathbf{E}[X]$ and $\mathbf{Var}[X]$ to denote expectation and variance of a random variable X , where the underlying distribution will be the uniform distribution \mathcal{U}_n on $\{-1, 1\}^n$, unless explicitly stated otherwise. For $x \in \{-1, 1\}^n$ and $S \subseteq [n]$ we write x_S to denote $(x_i)_{i \in S}$. For a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ and $q \geq 1$, we denote by $\|f\|_q$ its l_q -norm, i.e., $\|f\|_q \stackrel{\text{def}}{=} \mathbf{E}_x[|p(x)|^q]^{1/q}$. For Boolean functions $f, g : \{-1, 1\}^n \rightarrow \mathbb{R}$ the distance between f and g , denoted $\text{dist}(f, g)$, is defined by $\text{dist}(f, g) \stackrel{\text{def}}{=} \mathbf{E}_x[|f(x) - g(x)|]$. If $\text{dist}(f, g) \leq \epsilon$, we say that f and g are ϵ -close. Note that if f, g are Boolean-valued, i.e., take values in $\{\pm 1\}$, then $\text{dist}(f, g) = 2\Pr_x[f(x) \neq g(x)]$. The *disagreement region* between f and g is defined by $D(f, g) \stackrel{\text{def}}{=} \{x \in \{-1, 1\}^n : f(x) \neq g(x)\}$. For a multilinear polynomial $P : \mathbb{R}^n \rightarrow \mathbb{R}$ with $P(x) = \sum_{S \subseteq [n]} P_S \prod_{i \in S} x_i$, we denote by $\text{supp}(P) = \{S \subseteq [n] : P_S \neq 0\}$ and we call $|\text{supp}(P)|$ the support size of P .

Fourier Analysis and Influences. We consider functions $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ and we think of the inputs x as being distributed according to the uniform distribution \mathcal{U}_n . The set of such functions forms a 2^n -dimensional inner product space with inner product given by $\langle f, g \rangle = \mathbf{E}_x[f(x)g(x)]$. The set of functions $(\chi_S)_{S \subseteq [n]}$ defined by $\chi_S(x) = \prod_{i \in S} x_i$ forms a complete orthonormal basis for this space. Given a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ we define its *Fourier coefficients* by $\hat{f}(S) \stackrel{\text{def}}{=} \mathbf{E}_x[f(x)\chi_S(x)]$, and we have that $f(x) = \sum_{S \subseteq [n]} \hat{f}(S)\chi_S(x)$. As a consequence of orthonormality, we have *Plancherel's identity* $\langle f, g \rangle = \sum_{S \subseteq [n]} \hat{f}(S)\hat{g}(S)$, which has as a special case *Parseval's identity*, $\mathbf{E}_x[f(x)^2] = \sum_{S \subseteq [n]} \hat{f}(S)^2$. From this it follows that for every $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ we have $\sum_S \hat{f}(S)^2 = 1$. The expectation and the variance of $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ can be expressed in terms of the Fourier coefficients of f by $\mathbf{E}[f] = \mathbf{E}_x[f(x)] = \hat{f}(\emptyset)$ and $\mathbf{Var}[f] = \mathbf{Var}_x[f(x)] = \sum_{\emptyset \neq S \subseteq [n]} \hat{f}(S)^2$. The *influence* of variable i on $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ is $\text{Inf}_i(f) \stackrel{\text{def}}{=} \sum_{S \ni i} \hat{f}(S)^2$ and the *total influence* of f is $\text{Inf}(f) = \sum_{i=1}^n \text{Inf}_i(f) = \sum_{S \subseteq [n]} |S| \hat{f}(S)^2$.

Useful Probability Bounds. We will need the following well-known concentration bound for degree- d polynomials, a simple corollary of hypercontractivity (see, e.g., Theorem 9.23 in [O'D14]):

Fact 7. *Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a degree- d multilinear polynomial. For any $t > e^d$, we have that $\Pr_x[|p(x)| \geq t\|p\|_2] \leq \exp(-\Omega(t^{2/d}))$.*

We say that a polynomial $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ is τ -regular if $\max_{i \in [n]} \text{Inf}_i(p) \leq \tau \cdot \text{Inf}(p)$. Our second technical fact is that regular polynomials over the hypercube are anti-concentrated. This follows by combining the invariance principle [MOO10] and Gaussian anti-concentration [CW01] (see Claim 4.2 in [DSTW14] for an explicit reference):

Claim 8. *Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a τ -regular degree- d multilinear polynomial. Then it holds that $\Pr_x[|p(x)| \leq \tau\|p\|_2] \leq O(d\tau^{1/(8d)})$.*

Our proof makes essential use of a degree- d version [MNV16] of the classical Littlewood-Offord lemma [LO43, Erd45]. To state it, we need the following definition:

Definition 9 ([RV13, MNV16]). [Rank of Multilinear Polynomials] For a degree- d multilinear polynomial on n variables $P(x) = \sum_{S \subseteq [n], |S| \leq d} P_S \cdot \chi_S(x)$, the *rank* of P , denoted by $\text{rank}(P)$, is the largest integer r such that there exist disjoint sets $S_1, \dots, S_r \subseteq [n]$ of size d with $P_{S_j} \neq 0$ for $j \in [r]$.

Theorem 10 ([MNV16], Theorem 1.6). *For any degree- d multilinear polynomial P in n variables of rank $r \geq 2$ and any point $t \in \mathbb{R}$ we have that $\Pr_x[P(x) = t] \leq \frac{d^{O(d^2)} \cdot (\log r)^{O(d \log d)}}{\sqrt{r}}$.*

We note that the above is not quite the theorem from [MNV16]. They show that if there are r disjoint sets of coordinates S_i so that the $|P_{S_j}|$ are all at least 1, then the probability that $|P(x) - t| < 1$ is bounded by $\frac{d^{O(d^2)} \cdot (\log r)^{O(d \log d)}}{\sqrt{r}}$. Our result follows easily from this by replacing P by $m \cdot P$ for m an integer larger than any of the $1/|P_{S_j}|$'s. Applying [MNV16]'s original result to bound the probability that $|m \cdot P(x) - mt| < 1$ gives Theorem 10. We will make use of Theorem 10 in the form of the following corollary:

Corollary 11. *For any degree- d multilinear polynomial P , if $\Pr_x[P(x) = t] > \eta$, then there exists a set S of at most $d^{O(d^2)} \eta^{-3}$ coordinates so that every non-zero degree- d term in P has at least one variable of S in it.*

Proof. By the contrapositive of Theorem 10, the rank of P must be at most $d^{O(d^2)} \eta^{-3}$. Let S_1, S_2, \dots, S_r be a maximal set of disjoint subsets of $[n]$ with $|S_i| = d$ and $P_{S_i} \neq 0$. We claim that $S = \bigcup S_i$ suffices. This is because any other $S' \subset [n]$ of size d with $P_{S'} \neq 0$ must intersect some S_i , and thus intersect S . Furthermore, we have that $|S| = d \cdot r \leq d^{O(d^2)} \eta^{-3}$. \square

3 Main Structural Result: Proof of Theorem 3

In this section, we prove Theorem 3. Let f be any Boolean degree- d PTF and $g : \{-1, 1\}^n \rightarrow [-1, 1]$ be any bounded function. We will show that if $\text{Chow}_d(f, g)$ is sufficiently small, then $\text{dist}(f, g)$ is small.

The structure of this section is as follows: In Section 3.1, we reduce to the case that g is *Boolean-valued*. In Section 3.2, we develop our first major technical tool. In particular, we prove a generalization of a structural result in [Gol06, DDFS14], showing that if f and g have abnormally small degree- d Chow distance, then there must be some degree- d polynomial that exactly vanishes on almost all of their points of disagreement (Proposition 14). We then proceed to apply Proposition 14 in order to obtain a contradiction for the assumption that f and g have degree- d Chow distance too small relative to their ℓ_1 -distance. In Section 3.3, we start by showing this in the degree-2 case. In Section 3.4, we generalize to higher degrees.

3.1 Reduction to the Case that g is Boolean-valued

We begin by showing that it suffices to prove Theorem 3 for the special case that the function g is Boolean-valued, as opposed to bounded.

The idea of the proof is fairly simple: Let g be a $[-1, 1]$ -valued function. We can randomly round g to a Boolean-valued function in such a way as to maintain its expected degree- d Chow parameters and distance from f . Furthermore, it is not hard to show that if the dimension n is large (which can be achieved by introducing new irrelevant variables if necessary), the errors are, on average, quite small. That is, for any f, g we can find a Boolean-valued function g_0 with $\text{dist}(f, g) \approx \text{dist}(f, g_0)$ and $\text{Chow}_d(f, g) \approx \text{Chow}_d(f, g_0)$. Therefore, if we have a theorem relating distance to degree- d Chow distance for f and g_0 , essentially the same statement applies to f and g . More formally, we have:

Lemma 12. *Suppose that for some $\eta, \delta, d > 0$ that for any n and any degree- d PTF f and Boolean function g in n variables with $\text{dist}(f, g) > \eta$ have $\text{Chow}_d(f, g) \geq \delta$. Then for those same η, d, δ , f any degree- d PTF in any number of variables and g any function valued in $[-1, 1]$ with $\text{dist}(f, g) > \eta$, we must have $\text{Chow}_d(f, g) \geq \delta$.*

Proof. Assume for sake of contradiction that we have a degree- d PTF, f , and a bounded function $g : \{-1, 1\}^n \rightarrow [-1, 1]$ with $\text{dist}(f, g) > \eta$ and $\text{Chow}_d(f, g) < \delta$. By introducing irrelevant variables if necessary, we can make the dimension n as large as we like. Define a random rounding g_0 of g as follows: $g_0(x) = 1$ with probability $(g(x) + 1)/2$ and $g_0(x) = -1$ otherwise. Furthermore, the $g_0(x)$'s are independent of each other. Notice that $\text{dist}(f, g) = \mathbf{E}[\text{dist}(f, g_0)]$. Furthermore, since $\text{dist}(f, g_0)$ is the average of 2^n independent and bounded random variables, we have that $\text{dist}(f, g_0) = \text{dist}(f, g) + O(2^{-n/3})$ with high probability. Similarly, for any subset S of coordinates we have that $\mathbf{E}_{g_0}[\mathbf{E}_x[g_0(x)\chi_S(x)]] = \mathbf{E}_x[g(x)\chi_S(x)]$, where again it is an average of 2^n independent, bounded random variables. Thus, with high probability over the choice of g_0 , every degree at most d Fourier coefficient of g_0 is within $2^{-n/3}$ of the corresponding parameter of g . Thus, with high probability over the choice of g_0 we have that $\text{dist}(f, g_0) = \text{dist}(f, g) + O(2^{-n/3})$ and $\text{Chow}_d(f, g_0) = \text{Chow}_d(f, g) + O(n^d 2^{-n/3})$. For n sufficiently large, this implies that we have a Boolean g_0 so that $\text{dist}(f, g_0) > \eta$ and $\text{Chow}_d(f, g_0) < \delta$, which is a contradiction. \square

Lemma 12 implies that, in order to establish Theorem 3, it suffices to prove the following:

Theorem 13. *There exists a function $\kappa : \mathbb{R} \times \mathbb{N} \rightarrow \mathbb{R}$ such that the following holds: Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be any degree- d PTF and $g : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be an arbitrary Boolean-valued function. If $\text{Chow}_d(f, g) \leq \kappa(\epsilon, d)$, then $\text{dist}(f, g) \leq \epsilon$.*

In the main part of this section, we prove Theorem 13.

3.2 Existence of Degree- d Polynomial that Vanishes in the Disagreement Region

Our first result shows that there exists a polynomial that captures almost all the disagreement region, in the sense that it vanishes on most such points. Formally, we have:

Proposition 14. *Fix $0 < \delta < \eta < 1$ such that $\eta = 1/\text{poly}_d(\log \log(1/\delta))$ is sufficiently large. Let $f(x) = \text{sign}(p(x))$ be a degree- d PTF and $g : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be a Boolean function such that $\text{Chow}_d(f, g) \leq \delta$. Then there exists a degree- d multilinear polynomial $r : \mathbb{R}^n \rightarrow \mathbb{R}$ such that the following holds:*

- (i) $\Pr_x[(f(x) \neq g(x)) \cap (r(x) \neq 0)] < \eta$.
- (ii) $\mathbf{E}_x[p(x)r(x)] \neq 0$, i.e., r non-trivially correlates with p .

Note that although δ here will be much much smaller than η , our application of this lemma will only rely on the fact that η goes to 0 with δ for fixed d .

Proof. Our proof proceeds in several stages. We begin by noting that, by a generalization of Chow's original argument, if the degree- d Chow distance between f and g is small, then nearly all of the discrepancies between $f = \text{sign}(p(x))$ and g have $|p(x)|$ very small (see Claim 15). The intuition is that if the polynomial p is regular, then this cannot happen for very many points (by Claim 8). In order to reduce to this regular case, we apply the regularity lemma of [DSTW14]. This will provide us with a small set S of coordinates (independent of n) so that for most of our points of disagreement, after fixing the coordinates in S , our polynomial must be reduced to one with small ℓ_2 norm. From here we use basic techniques in Diophantine approximation theory to approximate a multiple of p by a polynomial r that is in some sense integral. This allows us to show that on all but a small number of disagreements, not only does the restriction of r (obtained by fixing the coordinates in S) have small ℓ_2 norm, but the restriction must in fact be 0. This r will thus satisfy the necessary requirements of our proposition.

We now proceed with the detailed proof. Suppose that $f(x) = \text{sign}(p(x))$, where $p : \mathbb{R}^n \rightarrow \mathbb{R}$ is a degree- d multilinear polynomial that without loss of generality satisfies $\|p\|_2 = 1$. Let $D = D(f, g)$ be the

disagreement region between f and g . If $\Pr_x[D] < \eta$, the proposition follows trivially by taking $r \equiv p$. We will hence assume that $\Pr_x[D] \geq \eta$.

We show that the upper bound on the degree- d Chow distance implies that all but a few elements of D have $|p(x)|$ very small. More concretely, we have the following claim:

Claim 15. *Let $D' = \{x \in D : |p(x)| > \delta/\eta\}$. Then, $\Pr_x[D'] \leq \eta/2$.*

Proof. The proof follows by a simple argument very similar to the original proof of Chow's theorem. Let $p(x) = \sum_{S \subseteq [n], |S| \leq d} p_S \chi_S(x)$ and note that $\|p\|_2^2 = \sum_{S \subseteq [n], |S| \leq d} p_S^2 = 1$. We can write

$$\mathbf{E}_x[|(f(x) - g(x)) \cdot p(x)|] = \mathbf{E}_x[(f(x) - g(x)) \cdot p(x)] = \sum_{S \subseteq [n], |S| \leq d} p_S \cdot (\widehat{f}(S) - \widehat{g}(S)) \leq \|p\|_2 \cdot \text{Chow}_d(f, g) \leq \delta, \quad (1)$$

where the first equality uses the fact that $(f(x) - g(x)) \cdot p(x)$ is non-negative for all $x \in \{-1, 1\}^n$, the second equality is Plancherel's identity, and the inequality is Cauchy-Schwarz. Then, we can write

$$\mathbf{E}_x[|(f(x) - g(x)) \cdot p(x)|] \geq \Pr_x[D'] \cdot 2 \cdot \min_{x \in D'} |p(x)| > \Pr_x[D'] \cdot (2\delta/\eta). \quad (2)$$

Combining (1) and (2), we get that $\Pr_x[D'] \leq \eta/2$, completing the proof of Claim 15. \square

For $S \subseteq [n]$, we can partition the coordinates as $x = (x_S, x')$, where $x' = x_{[n] \setminus S}$ and rewrite the degree- d multilinear polynomial $p(x)$ as $p_{x_S}(x') = p(x_S, x')$. For any fixed assignment to x_S , we will view $p_{x_S}(x')$ as a degree- d multilinear polynomial in x' .

We will require the following structural lemma showing that there exists a set $S \subseteq [n]$ of coordinates, whose size is independent of the dimension n , such that for at least $1 - \eta/2$ fraction of points $x \in \{-1, 1\}^n$, the restricted polynomial p_{x_S} has ℓ_2 -norm not much larger than $|p(x)|$:

Lemma 16. *There exists $S \subseteq [n]$ of size $|S| \leq 2^{\text{poly}_d(1/\eta)}$ such that*

$$\Pr_x[\|p_{x_S}\|_2 > \text{poly}_d(1/\eta) \cdot |p(x)|] < \eta/2.$$

Proof. We start by applying a regularity lemma to the polynomial p . We will use the following statement:

Theorem 17 ([DSTW14]). *Let $f(x) = \text{sign}(p(x))$, where $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ is a degree- d multilinear polynomial and $0 < \tau < 1$. Then f is equivalent to a decision tree \mathcal{T} , of depth $\text{depth}(d, \tau) := \frac{1}{\tau} \cdot (d \log \frac{1}{\tau})^{O(d)}$ with variables at the internal nodes and a degree- d PTF $f_\rho = \text{sign}(p_\rho)$ at each leaf ρ , with the following property: with probability at least $1 - \tau$, a random path³ from the root reaches a leaf ρ such that either (i) p_ρ is τ -regular, or (ii) $\text{Var}[p_\rho] < \tau \cdot \mathbf{E}[p_\rho^2]$*

Remark 18. The above statement is not explicit in [DSTW14], but easily follows from their proof. Specifically, Theorem 17 follows mutatis mutandis by taking the parameter β in their Lemma 3.5 to satisfy $1/\tau = \Theta(\log(1/\beta)^d)$ and noting that the restriction ρ in its statement satisfies their Definition 3.6.

We call a leaf *Good* if it satisfies (i) or (ii) in the statement of Theorem 17 above. By an application of Theorem 17 for the polynomial p defining our degree- d PTF f and $\tau \stackrel{\text{def}}{=} \Theta(\eta/d)^{8d}$ (with the implied constant sufficiently small), we obtain a decision tree of depth $\text{depth}(d, \tau) = (d/\eta)^{O(d)}$ such that with probability at least $1 - \tau$ a random path in the tree leads to a good leaf, i.e., $\Pr_x[\rho \text{ is Good}] \geq 1 - \tau$. We show the following claim:

Claim 19. *We have that $\Pr_x[|p(x)| > \text{poly}_d(\eta) \cdot \|p_\rho\|_2] > 1 - \eta/4$.*

³A random path corresponds to the standard uniform random walk on the tree.

Proof. Note that the restriction ρ defining the set of variables that are fixed in the path from the root to the corresponding leaf of the tree depends on the input x . By Theorem 17, $\Pr_x[\rho \text{ is Good}] \geq 1 - \tau$. We condition on this event and analyze each case separately.

Consider a restriction ρ satisfying (i). In this case, since p_ρ is τ -regular the polynomial $p_\rho(x')$ is anti-concentrated. Specifically, Claim 8 gives that $\Pr_{x'}[|p_\rho(x')| \leq \tau \cdot \|p_\rho\|_2] \leq \eta/8$.

Consider a restriction ρ satisfying (ii). In this case, the constant term of p_ρ is very large and it is very unlikely that the non-constant term dominates. Let $p_\rho(x') = p'(\rho) + q(\rho, x')$, where x' are the variables not fixed by ρ . Then $p'(\rho)$ is the constant term and we note that $\mathbf{Var}[p_\rho] = \|q\|_2^2$ and $\mathbf{E}[p_\rho^2] = (p'(\rho))^2 + \|q\|_2^2$. Since $\mathbf{Var}[p_\rho] < \tau \cdot \mathbf{E}[p_\rho^2]$, it follows that $\|q\|_2^2 \leq \tau \cdot (p'(\rho))^2$. Hence, by the concentration bound of Theorem 7, we obtain $\Pr_{x'}[|q(\rho, x')| \geq |p'(\rho)|/4] \leq \eta/4$. If the latter event fails to hold, we get that $|p(x)| = |p_\rho(x')| \geq (3/4)|p'(\rho)| \geq (1/2)\|p_\rho\|_2$, as desired.

Combining the above, we get that $\Pr_x[|p(x)| \leq \tau \cdot \|p_\rho\|_2 \mid \rho \text{ is Good}] \leq \eta/8$. We therefore obtain

$$\Pr_x[|p(x)| \leq \tau \cdot \|p_\rho\|_2] \leq \Pr_x[\rho \text{ is not Good}] + \Pr_x[|p(x)| \leq \tau \cdot \|p_\rho\|_2 \mid \rho \text{ is Good}] \leq \tau + \eta/8 < \eta/4.$$

This completes the proof of Claim 19. \square

To complete the proof of the lemma, we let S be the set of all coordinates appearing in the decision tree. We note that $|S| \ll 2^{\text{depth}(\mathcal{T})} = 2^{\text{poly}_d(1/\eta)}$. What remains to show is that $\|p_{x_S}\|_2$ is small for almost all $x \in D$. We do this by showing that with high probability $\|p_{x_S}\|_2$ is not much larger than $\|p_\rho\|_2$.

Recall that S is the set of variables that appear in the decision tree, hence contains the variables fixed by any restriction ρ in any root to leaf path. Note that for any such restriction ρ , we have that

$$\mathbf{E}_{x_S}[\|p_{x_S}\|_2^2 \mid \rho] = \mathbf{E}_{x_S}[\mathbf{E}_{x'}[p(x_S, x')^2 \mid \rho]] = \mathbf{E}_x[p(x)^2 \mid \rho] = \|p_\rho\|_2^2.$$

Therefore, by Markov's inequality, it follows that

$$\Pr_x[\|p_{x_S}\|_2^2 > (4/\eta) \cdot \|p_\rho\|_2^2] \leq \eta/4. \quad (3)$$

Inequality (3), Claim 19, and a union bound give that

$$\Pr_x[\|p_{x_S}\|_2 > \text{poly}_d(1/\eta) \cdot |p(x)|] < \eta/2.$$

This completes the proof of Lemma 16. \square

Combining Claim 15 and Lemma 16, we get that that for all but $\eta \cdot 2^n$ points in D we have that

$$\delta/\eta > |p(x)| > \text{poly}_d(\eta) \cdot \|p_{x_S}\|_2.$$

Therefore, we have shown the following:

Corollary 20. *There exists a set of coordinates $S \subset [n]$ of size $|S| \leq 2^{\text{poly}_d(1/\eta)}$ such that the set*

$$\tilde{D} \stackrel{\text{def}}{=} \{x \in D : \|p_{x_S}\|_2 = O(\delta/\text{poly}_d(\eta))\}$$

has $\Pr_x[\tilde{D}] > \Pr_x[D] - \eta$.

We will construct a degree- d multilinear polynomial r that correlates with p and is such that $r(x) = 0$ for all $x \in \tilde{D}$. This will complete the proof of Proposition 14, since $\Pr_x[(f(x) \neq g(x)) \cap (r(x) \neq 0)] \leq \Pr_x[D \setminus \tilde{D}] < \eta$. To do so, we will leverage the structural information provided by Corollary 20. Let S be the set of coordinates satisfying Corollary 20. We rewrite $p(x) = p_{x_S}(x')$ and note that, for each fixed x_S , $p_{x_S}(x')$ is a degree- d multilinear polynomial in x' . We can view p as a degree- d multilinear polynomial

in x_S that returns degree- d multilinear polynomials on x' . Let v_0, v_1, \dots, v_R , where $R = O(|S|^d)$, be an orthonormal basis of the image $\text{Im}(p_{x_S})$. Then we can write

$$p_{x_S}(x') = \sum_{0 \leq j \leq R, T \subseteq S, |T| \leq d} \alpha_{T,j} \cdot \chi_T(x) \cdot v_j(x'). \quad (4)$$

We note that the functions $\{\chi_T(x) \cdot v_j(x')\}_{T \subseteq S, 0 \leq j \leq R}$ are orthonormal.

The idea of the rest of the proof is as follows. We note that if the $\alpha_{T,j}$'s were all integers, we could use $r = p$ and would be done. This is because *any* $p_{x_S}(x')$ would be an integer linear combination of the v_j 's, and thus would have integer ℓ_2 -norm. On the other hand, for $x \in \tilde{D}$, it must be the case that p_{x_S} has small ℓ_2 -norm. The only way that these can simultaneously hold is if $p_{x_S} = 0$. In particular, this would imply that $p_{x_S} = 0$ (and thus that $p(x) = 0$) for all $x \in \tilde{D}$.

In order to prove this result for more general p , we show that some multiple of p can be *approximated* by a polynomial where all of the $\alpha_{T,j}$'s are integers. To do so, we will need to make use of the following fact from the theory of Diophantine approximation:

Fact 21. *Let $N \in \mathbb{Z}_+$ and $w \in \mathbb{R}^N$. For any $\gamma \in (0, 1/2)$ there exists $t \in [1, O(1/\gamma)^N]$ such that*

$$t \cdot w \in (\mathbb{Z} + (-\gamma, \gamma))^N.$$

Proof. Let $\gamma > 1/m$ for some integer $m = O(1/\gamma)$. We show in fact that t can be taken to be an integer.

Partition $[0, 1]^N$ into m^N subcubes of side length $1/m$ in each dimension. For each integer k from 0 to m^N , sort $kw \pmod{1}$ into the appropriate subcube. Since there are $m^N + 1$ values of k and only m^N subcubes, by the pigeonhole principle, there must be $k \neq k'$ so that $kw \pmod{1}$ and $k'w \pmod{1}$ fall in the same subcube. Without loss of generality, $k \geq k'$, and we let $t = k - k'$. It follows that the coordinates of $tw = kw - k'w$ are all in $\mathbb{Z} + [-1/m, 1/m] \subset \mathbb{Z} + (-\gamma, \gamma)$. \square

We apply Fact 21 to the vector $(\alpha_{T,j})$ defined by the coefficients of the polynomial p in (4). We have that $N = O(|S|^{2d})$ and we set $\gamma \stackrel{\text{def}}{=} 1/N^2$. It follows that there exists $t < 2^{2^{\text{poly}_d(1/\eta)}}$ such that all $t \cdot \alpha_{T,j}$ are within an additive $1/N^2$ of being integers. We can thus write

$$t \cdot p(x_S, x') = r(x_S, x') + e(x_S, x'),$$

where r and e are degree- d multilinear polynomials

$$r_{x_S}(x') = \sum_{0 \leq j \leq R, T \subseteq S, |T| \leq d} \beta_{T,j} \cdot \chi_T(x) \cdot v_j(x'),$$

and

$$e_{x_S}(x') = \sum_{0 \leq j \leq R, T \subseteq S, |T| \leq d} \gamma_{T,j} \cdot \chi_T(x) \cdot v_j(x'),$$

where $\beta_{T,j} \in \mathbb{Z}$ and $|\gamma_{T,j}| < 1/N^2$ for all T, j . Note that this implies that the sum of the $|\gamma_{T,j}|$ is at most $1/N$. It follows that for any value of $x_S \in \{\pm 1\}^S$ it holds $\|e_{x_S}\|_2 < 1/N$. On the other hand, for any such x_S , r_{x_S} is an integer linear combination of $v_j(x')$, and thus $\|r_{x_S}\|_2^2 \in \mathbb{Z}$. However, for $x \in \tilde{D}$, we have that

$$\|r_{x_S}\|_2 \leq \|t \cdot p_{x_S}\|_2 + \|e_{x_S}\|_2 \leq t \cdot O(\delta / \text{poly}_d(\eta)) + 1/N.$$

Therefore, if $O(t \cdot \delta / \text{poly}_d(\eta)) < 1/2$ (which holds for $1/\eta$ a sufficiently small polynomial in $\log \log(1/\delta)$), we have that for all $x \in \tilde{D}$ it holds

$$\|r_{x_S}\|_2 < 1.$$

However, since $\|r_{x_S}\|_2^2$ is an integer, this can only hold if $r_{x_S} = 0$. Therefore, for all $x \in \widetilde{D}$, we have that $r(x) = 0$, establishing (i).

We can now prove part (ii) of Proposition 14, i.e., that $\mathbf{E}_x[p(x)r(x)] \neq 0$. From the definition of r , we have that $(1/t)r(x) = p(x) - e(x)/t$. We note that e_{x_S} always has ℓ_2 -norm at most $1/N$. It follows that $\|e(x)/t\|_2 < 1/2$.

Therefore,

$$\mathbf{E}_x[p(x)(1/t)r(x)] = \mathbf{E}_x[p^2(x)] - \mathbf{E}_x[p(x)e(x)/t] \geq 1 - 1/2 > 0 ,$$

where we used the assumption that $\mathbf{E}_x[p^2(x)] = \|p\|_2^2 = 1$ and the Cauchy-Schwartz inequality. This establishes (ii), and completes the proof of Proposition 14. \square

3.3 Warm-Up: Completing the Proof for Degree-2 PTFs

In this subsection, we complete the proof of Theorem 3 for degree-2 polynomial threshold functions.

The high-level idea of the proof is as follows: Given a pair f, g contradicting our desired statement, by iteratively applying Proposition 14 we can find many degree-2 polynomials in which almost all of the disagreements between f and g vanish. We would like to use this fact in order to reach a contradiction, by showing that there will be not enough points in the disagreement region. Unfortunately, this does not necessarily suffice, as it is possible to have many degree-2 polynomials that have joint zeroes at a substantial number of points. However, we show (see Proposition 22) that given enough such degree-2 polynomials, we can find a single *linear* polynomial that vanishes on almost all of these discrepancies. From there, we can restrict to the hyperplane defined by this polynomial yielding a new pair of functions with greater discrepancy and repeat this process.

We now proceed with the formal proof. We start by showing that for any subspace of degree-2 polynomials whose dimension is sufficiently large, there exists a non-trivial linear function that vanishes on almost all of its zero set.

Proposition 22. *Fix $0 < \eta < 1$. There exist positive integers $N_0 = N_0(\eta) = \Theta(1/\eta^3)$ and $N_1 = N_1(\eta) = \Theta(1/\eta^7)$ such that the following holds: Let V be any subspace of degree-2 n -variable polynomials of dimension $\dim(V) > N_1$ and $Z(V) \stackrel{\text{def}}{=} \{x \in \{-1, 1\}^n : p(x) = 0 \text{ for all } p \in V\}$. Then there exists a linear function $L : \mathbb{R}^n \rightarrow \mathbb{R}$, not identically zero, with $|\text{supp}(L)| \leq N_0$ such that $\Pr_x[Z(V) \cap L(x) \neq 0] \leq \eta$.*

Proof. The proof proceeds in two stages: First, via Corollary 11, we show that there is a small set S of coordinates so that all degree-2 terms of any polynomial in V have a variable in S . Next, we consider the set of points in $\{\pm 1\}^S$ which define sub-cubes with a reasonably large number of joint zeroes. We show that either this set is contained in a linear subspace or that $\dim(V)$ must be bounded.

We can assume that $\Pr_x[Z(V)] > \eta$, otherwise there is nothing to prove. Therefore, for any $p \in V$ it holds that $\Pr_x[p(x) = 0] > \eta$. This implies the existence of structure on the coefficients of p , which can be formalized using Theorem 10, a degree- d version of the classical Littlewood-Offord lemma [LO43, Erd45].

We start with the following claim:

Claim 23. *There exists a set S of at most $N_0 = \Theta(\eta^{-3})$ coordinates such that for all $q \in V$ each degree-2 term of q has a coordinate in S .*

Proof. We let p be a generic element of V . In particular, p_T should be non-zero for all T for where $q_T \neq 0$ for any $q \in V$. Applying Corollary 11 to p , we find a set of $O(\eta^{-3})$ coordinates so that each non-vanishing degree-2 term of p contains a variable in S . Since p is generic, this implies that every non-vanishing degree-2 term of every q in V also has a variable in S . This completes our proof. \square

So far we have shown that any polynomial $p \in V$ is of the form

$$p(x) = p_{\text{junta}}(x_S) + p_{\text{cross}}(x_S, x'),$$

where $p_{\text{junta}}(x_S)$ is a degree-2 polynomial in x_S and $p_{\text{cross}}(x_S, x')$ is linear in x' .

Let V' be the subspace of V consisting of polynomials with *no* non-zero terms using only coordinates in S . That is, any polynomial $p \in V'$ is of the form:

$$p(x) = p_{\text{cross}}(x_S, x').$$

Note that

$$m \stackrel{\text{def}}{=} \dim(V') \geq \dim(V) - |S|^2 > N_1 - |S|^2 = \Omega(\eta^{-7}),$$

where we used that $\dim(V) \geq N_1$ and $|S| \leq N_0$. Let $Z(V') \stackrel{\text{def}}{=} \{x \in \{-1, 1\}^n : p(x) = 0 \text{ for all } p \in V'\}$. Since $V' \subseteq V$, it follows that $Z(V) \subseteq Z(V')$. Therefore, $\Pr_x[Z(V')] \geq \Pr_x[Z(V)] > \eta$.

We can view each $p \in V'$ as an affine linear function from x_S to the set of linear functions on x' . That is, for any fixed x_S , the space $V'_{x_S} = \{p(x_S, x'), p \in V'\}$ is a subspace of linear functions on x' .

We will need the following basic fact (see, e.g., Lemma 1 in [Gol06]):

Fact 24. *Let W be a subspace of linear functions on x . Then $\Pr_x[L(x) = 0 \text{ for all } L \in W] \leq 2^{-\dim(W)}$.*

Let

$$G \stackrel{\text{def}}{=} \left\{ x_S \in \{-1, 1\}^{|S|} \text{ such that } \Pr_{x'}[p(x_S, x') = 0 \text{ for all } p \in V'] > \eta \right\}.$$

Let p_1, \dots, p_m be a basis of V' . Since, for any fixed x_S , each $p_i(x_S, x')$ is a linear function in x' , by Fact 24 we get that for any fixed $x_S \in G$ we have

$$\dim(V'_{x_S}) = \dim(\text{span}\{p_i(x_S, x'), i \in [m]\}) < \log_2(1/\eta).$$

We now establish the following claim:

Claim 25. *The set G lies in a proper affine linear subspace of $\{-1, 1\}^{|S|}$.*

Proof. For the sake of contradiction, suppose that there exist $y_1, \dots, y_R \in G$, with $R = |S| + 1$, whose affine span is $\{-1, 1\}^{|S|}$. Then every $x_S \in \{-1, 1\}^{|S|}$ can be written as an affine linear combination of the y_i 's. Since the $p_i(x_S, x')$'s are linear functions for every fixed x_S , it follows that for every fixed p_i , $p_i(x_S, x') \in \text{span}(\{p_i(y_j, x'), j \in [R]\})$. Therefore, for all $x_S \in \{-1, 1\}^{|S|}$ and p_i we have that

$$p_i(x_S, x') \in \text{span}(\{p_i(y_j, x'), j \in [R], i \in [m]\}) =: U.$$

Note that $\dim(U) \leq R \cdot \log(1/\eta) = (|S| + 1) \log(1/\eta)$. Since all $p_i(x_S, x')$'s are affine linear functions from x_S to U , it follows that

$$\dim(V') \leq (|S| + 1) \dim(U) \leq (|S| + 1)^2 \log(1/\eta) = \tilde{O}(\eta^{-6}),$$

which leads to the desired contradiction, since $m = \Omega(\eta^{-7})$. This completes the proof of Claim 25. \square

By Claim 25, it follows that there is a non-zero linear function $L : \mathbb{R}^{|S|} \rightarrow \mathbb{R}$ such that $L(x_S) = 0$ for all $x_S \in G$. We can trivially extend L to \mathbb{R}^n by adding zero coefficients for the remaining coordinates, i.e., $|\text{supp}(L)| \leq |S| \leq N_0$. It remains to argue that

$$\Pr_x [Z(V) \cap (L(x) \neq 0)] \leq \eta.$$

Indeed, we can write

$$\begin{aligned}
\Pr_x [Z(V) \cap (L(x) \neq 0)] &\leq \Pr_x [Z(V') \cap (L(x) \neq 0)] \\
&\leq \Pr_x [Z(V') \cap (x_S \in \bar{G})] \\
&\leq \Pr_{x=(x_S, x')} [Z(V') \mid x_S \in \bar{G}] \\
&\leq \eta,
\end{aligned}$$

where the first inequality holds since $Z(V) \subseteq Z(V')$, the second follows from the definition of L and the last inequality from the definition of G . This completes the proof of Proposition 22. \square

We now have the necessary ingredients to complete the proof of our main result for $d = 2$. We prove the following proposition:

Proposition 26. *Fix $0 < \epsilon, \delta < 1$. Let $f(x) = \text{sign}(p(x))$ be a degree-2 PTF and $g : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be a Boolean function such that $\text{dist}(f, g) \geq \epsilon$ and $\text{Chow}_2(f, g) \leq \delta$. Then, for all $m \geq 1$, there exists a degree-2 PTF f'_m and a Boolean-valued function g'_m such that $\text{dist}(f'_m, g'_m) \geq (3/2)^m \cdot \epsilon$ and $\text{Chow}_2(f'_m, g'_m) \leq \tilde{O}((\log^*)^{(m)}(1/\delta)^{-1/7})$, so long as the latter term is less than $(3/2)^m \cdot \epsilon$.*

The $d = 2$ case of Theorem 3 follows immediately from Proposition 26. Indeed, if $1/\epsilon = 2^{o(\log^{**}(1/\delta))}$, then we obtain a contradiction by setting $m = \Omega(\log(1/\epsilon))$ in Proposition 26.

In the rest of this section, we give the proof of Proposition 26. The proof follows by induction on m , where the base case is given in the following lemma:

Lemma 27. *Let $f(x) = \text{sign}(p(x))$ be a degree-2 PTF and $g : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be such that $\text{dist}(f, g) \geq \epsilon$ and $\text{Chow}_2(f, g) \leq \delta$. There exists a degree-2 PTF f' and a Boolean function g' such that $\text{dist}(f', g') \geq (3/2) \cdot \epsilon$ and $\text{Chow}_2(f', g') \leq \tilde{O}(\log^*(1/\delta)^{-1/7})$.*

The basic idea of the proof is quite simple. By repeatedly applying Proposition 14, we find many quadratic polynomials which vanish on almost all of the points of disagreement between f and g . From there, we apply Proposition 22 to produce a single linear polynomial which vanishes on almost all disagreements. Restricting ourselves to this hyperplane, we can set one of our variables as a linear function of the others, and reduce to an $(n - 1)$ -dimensional cube without substantially changing the number of disagreements.

One issue with the above strategy is that we need to ensure that the repeated applications of Proposition 14 produce linearly independent polynomials. This can be achieved by modifying p to be orthogonal to the previously found polynomials (and modifying f and g appropriately) so that the correlation condition in Proposition 14 implies that the polynomial produced will be new.

We begin this program by showing that we can find a single linear polynomial that vanishes on almost all of the disagreements:

Lemma 28. *Fix $0 < \delta < 1$ and η a sufficiently large multiple of $(\log^*(1/\delta))^{-1/7}$. Let $f(x) = \text{sign}(p(x))$ be a degree-2 PTF and $g : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be such that $\text{Chow}_2(f, g) \leq \delta$. Then there exists a non-trivial linear function $L : \mathbb{R}^n \rightarrow \mathbb{R}$ with $|\text{supp}(L)| \leq O(\log^*(1/\delta))$ such that $\Pr_x[(f(x) \neq g(x)) \cap L(x) \neq 0] \leq \eta$.*

Proof. Let $D = D(f, g)$ be the disagreement region between f and g . For a sequence of polynomials $\{p_i\}_{i \in [m]}$, we will denote by $Z_m = \{x \in \{-1, 1\}^n : p_i(x) = 0 \text{ for all } i \in [m]\}$ and by \bar{Z}_m its complement. We start by establishing the following claim:

Claim 29. *For all $m \geq 1$, there exist linearly independent degree-2 polynomials p_1, \dots, p_m with $\eta_0 = \delta$ and $\eta_m = 1/\text{poly}(\log \log(1/\eta_{m-1}))$ such that $\Pr_x[D \cap \bar{Z}_m] \leq \eta_m$.*

Proof. By induction on m . The base case, ($m = 0$) is trivial. For the induction step, suppose that p_1, \dots, p_m and η_m exist such that the claim holds. We will prove that there exists p_1, \dots, p_{m+1} and η_{m+1} satisfying the claim statement. Consider the polynomial p' defined by $p' = p - \text{proj}(p, \text{span}(p_1, \dots, p_m))$. Now note that the degree-2 PTF $f'(x) = \text{sign}(p'(x))$ satisfies $f'(x) = f(x)$ for all $x \in Z_m$. Let

$$g'(x) = \begin{cases} f'(x), & x \in \bar{Z}_m \\ g(x), & x \in Z_m \end{cases}$$

Note that $f'(x) - g'(x) = f(x) - g(x)$, $x \in Z_m$, and $f'(x) - g'(x) = 0$, otherwise. This means that $f - g = f' - g'$ except on η_m -mass of D . Therefore, $\text{Chow}_2(f', g') = \text{Chow}_2(f, g) + \tilde{O}(\eta_m) = \tilde{O}(\delta + \eta_m)$ and $\text{dist}(f', g') \geq \epsilon - \eta_m \approx \epsilon$.

Therefore, we can apply Proposition 14 to the pair f', g' . We thus obtain that there exists a polynomial p_{m+1} that correlates with p' such that on all but $\eta_{m+1} = 1/\text{poly}(\log \log(1/\eta_m))$ mass of $D(f', g')$ we have $p_{m+1}(x) = 0$.

Since $D(f', g')$ and $D(f, g)$ are within η_m mass of each other and $\eta_m \ll \eta_{m+1}$, we obtain that all but at most $O(\eta_{m+1}) \cdot 2^n$ points of $D(f, g)$ satisfy that $p_1(x) = \dots = p_{m+1}(x) = 0$.

By the conditions of Proposition 14, we have that $\mathbf{E}_x[p_{m+1}(x)p'(x)] \neq 0$. But by the definition of p' , we have that $\mathbf{E}_x[p_i(x)p'(x)] = 0$ for all $i \leq m$. Thus, $p_{m+1} \notin \text{span}(\{p_i, i \in [m]\})$ and the inductive step is complete. This completes the proof of Claim 29. \square

We are now ready to complete the proof of the lemma. We select N such that $\log^{[O(N^7)]}(1/\delta) > N$. By Claim 29, we obtain a set of linearly independent polynomials p_1, \dots, p_{N^7} such that all but $(1/N) \cdot 2^n$ points of $D(f, g)$ satisfy $p_1(x) = \dots = p_{N^7}(x) = 0$. By Proposition 22, there exists a non-zero linear form $L : \mathbb{R}^n \rightarrow \mathbb{R}$ with support $O(N^7)$ such that for all but $O(1/N) \cdot 2^n$ points x of $D(f, g)$ it holds $L(x) = 0$. Note that the probability mass of points where $L(x) = 0$ is $O(1/\log^*(1/\delta)^{1/7})$. This completes the proof of Lemma 28. \square

We are now prepared to prove Lemma 27.

Proof of Lemma 27. By Lemma 28, there exists a linear form L , not identically zero, with support of size $O(\log^*(1/\delta))$ such that all but $O(1/\log^*(1/\delta)^{1/7})$ mass of points x in $D(f, g)$ has $L(x) = 0$.

Since in most of the disagreement region $D(f, g)$ we have $L(x) = 0$, it follows that, for each such point x , we can express one of the variables in $L(x)$ as a function of the others. Specifically, we can write $L(x_1, \dots, x_n) = x_1 - L'(x_2, \dots, x_n)$.

We now consider the function $g_0(x)$ which equals $f(x)$ if $L(x) \neq 0$ and $g(x)$ otherwise. Note that $\text{dist}(f, g_0) \geq 3\epsilon/4$ and $\text{Chow}_2(f, g_0) \leq \tilde{O}(\log^*(1/\delta)^{-1/7})$.

Recall that $f(x) = \text{sign}(p(x))$. We consider the function f' which is a degree-2 PTF over x_2, \dots, x_n defined as follows: $f'(x_2, \dots, x_n) = \text{sign}(p(L'(x_2, \dots, x_n), x_2, \dots, x_n))$ and the Boolean function g' defined by $g'(x_2, \dots, x_n) = f'(x_2, \dots, x_n)$ if $L'(x_2, \dots, x_n) \notin \{\pm 1\}$; and $g'(x_2, \dots, x_n) = g(L'(x_2, \dots, x_n), x_2, \dots, x_n)$ otherwise. Then we have that $\text{dist}(f', g') = 2\text{dist}(f, g_0) \geq 3\epsilon/2$ and $\text{Chow}_2(f', g') \leq \text{Chow}_2(f, g_0) \leq \tilde{O}(\log^*(1/\delta)^{-1/7})$. This completes the proof of Lemma 27. \square

The proof of Proposition 26 is now immediate. Given the existence of f'_m and g'_m , we obtain f'_{m+1} and g'_{m+1} by applying Lemma 27.

3.4 Completing the Proof for Degree- d PTFs

The high-level approach of our degree- d generalization is similar as in the degree-2 case. First, we make use of Proposition 14, showing that if we have a small degree- d Chow distance, then all but a tiny number of discrepancies must lie on the zero set of a degree- d polynomial. From there, we can iterate this to show

that our discrepancies must lie on the intersection of *many* degree- d polynomials. We would ideally like to be able to generalize Lemma 28 and deduce that once there are enough of them, we will be forced to lie on a linear subspace, but this is too ambitious in general. Instead, we will use many degree- d polynomials to force our discrepancies to lie on the zero set of a single degree- $(d - 1)$ polynomial. From there, we will go back to gathering degree- d polynomials until we have enough to force a second degree- $(d - 1)$ polynomial; and eventually we will have enough of those to force a degree- $(d - 2)$ polynomial, and so on. To keep track of the general state of this recursion, we will need to maintain an *ideal* of polynomials which vanish on our discrepancies.

We start by recalling the definition of an ideal:

Definition 30. An ideal (in $\mathbb{R}[x_1, \dots, x_n]$) is a set I of polynomials so that

1. For any $p, q \in I$, we have $p + q \in I$.
2. For any $p \in I$ and $a \in \mathbb{R}[x_1, \dots, x_n]$, we have that $a \cdot p \in I$.

Note that an ideal I is also a vector subspace of the space of all polynomials.

Notation. We will use the following notation in the rest of the proof. For polynomials p_1, p_2, \dots, p_m , we consider the ideal $I = \{\sum_{i=1}^m a_i p_i : a_i \in \mathbb{R}[x_1, \dots, x_n]\}$, which we will denote by (p_1, p_2, \dots, p_m) . For two ideals, I and J , we have that the set of polynomials $\{p + q : p \in I, q \in J\}$ is another ideal, which we will denote $I + J$. We say that two polynomials p and q are congruent modulo I if $p - q \in I$. We note that this is an equivalence relation on polynomials.

We will also care about the points of the hypercube on which an ideal vanishes.

Definition 31. For a set \mathcal{P} of polynomials, let $Z(\mathcal{P})$ denote the set $\{x \in \{\pm 1\}^n : p(x) = 0 \text{ for all } p \in \mathcal{P}\}$.

Of particular interest is the case where \mathcal{P} is an ideal I . In particular, we note that $Z((p_1, \dots, p_m)) = \{x \in \{\pm 1\}^n : p_i(x) = 0 \text{ for all } 1 \leq i \leq m\}$ and that $Z(I + J) = Z(I) \cap Z(J)$.

We now proceed to generalize Proposition 22 to the degree- d setting. It turns out that this generalization is somewhat more complicated. We prove:

Proposition 32. *For every positive integer d and every $\eta > 0$ there exists an $m = \eta^{-d^{O(d)}}$ so that the following holds: For I an ideal containing $x_i^2 - 1$ for all $i \in [n]$, and V a vector subspace of the space of degree at most d polynomials with $I \cap V = \{0\}$, and $\dim(V) > m$, there exists a degree at most $d - 1$ polynomial $P \notin I$ so that $\Pr[Z(I) \cap Z(V) \setminus Z(P)] \leq \eta$.*

Proof. Firstly, we note that we may assume that $\Pr[Z(I) \cap Z(V)] > \eta$, or the result is trivial (for example, we can take $P = 1$, as 1 cannot be in I or I would have to contain V). Secondly, let V have a basis p_1, p_2, \dots, p_t . We note that if we find another set of degree-at-most- d polynomials q_1, q_2, \dots, q_t with $p_i \equiv q_i \pmod{I}$ and let $V' = \text{span}(\{q_i\})$, then $Z(I) \cap Z(V) = Z(I) \cap Z(V')$, $I \cap V' = \{0\}$ and $\dim(V') = \dim(V)$, and thus proving our statement for V is equivalent to proving it for V' . We call such q_i an equivalent basis.

Our proof proceeds by proving the following lemma:

Lemma 33. *Given any ideal I and vector space V and any integer $0 \leq k \leq d$ either there exists a set S_k of $O(\eta^{-d^{2(d-k)}})$ many coordinates and an equivalent basis q_i of V so that every non-zero monomial of every q_i contains at most k coordinates not in S_k or there exists a degree less than d polynomial $P \notin I$ so that $\Pr[Z(I) \cap Z(V) \setminus Z(P)] \leq \eta$.*

We note that the $k = 0$ case of this lemma immediately implies our result, as either such a P exists, or (up to taking an equivalent basis) V is a space of polynomials in $|S_0|$ variables, and thus has dimension at most $\sum_{i=0}^d \binom{|S_0|}{i}$. If m is larger than this, we are done. It remains to prove this lemma.

Proof. The proof proceeds by backwards induction on k and is reminiscent of the proof of Proposition 22. In particular, given an S_k , we think of our polynomials as being degree- d polynomials from $\{\pm 1\}^{S_k}$ to degree- k polynomials in the remaining variables. We consider the set of x in $\{\pm 1\}^{S_k}$ with a substantial number of points in $Z(V) \cap Z(I)$. By Corollary 11, we note that for each such x_S , there is a small set of remaining coordinates T_x so that all of the degree k terms in $p_{x_{S_k}}$ have a variable in T . By an appropriate dimensionality argument, we show that in fact we can use only a single T for all good x . From there, we show that it is either the case that some polynomial not in I vanishes on all good x , or that the generators of V can be reduced modulo I to have no terms of degree more than $k - 1$ in variables outside of $S_k \cup T$.

We proceed by backwards induction on k . In particular, for $k = d$, the result is trivial. Otherwise, assume that it holds for a given value of k . By the inductive hypothesis, either there exists a polynomial P , or there exists an appropriate set of coordinates S_k . By replacing V by an appropriate equivalent basis, we may assume that the p_i are multilinear and have all non-zero terms have at most k coordinates not in S_k . Define

$$G := \{x \in \{\pm 1\}^{S_k} : \Pr(y \in Z(I) \cap Z(V) | y_{S_k} = x) > \eta\}.$$

Note that all but an η -mass of the points of $Z(I) \cap Z(V)$ have $x_{S_k} \in G$.

Next, for $x \in \{\pm 1\}^{S_k}$, define $M(x)$ to be the vector of degree at most $(d - k)$ -monomials in x . Note that if $M(x_0)$ is a linear combination of some $M(x_i)$, then for any polynomial p of degree at most $d - k$, $p(x)$ will be the same linear combination of the $p(x_i)$. Let x_1, x_2, \dots, x_s be points in G so that $M(x_1), M(x_2), \dots, M(x_s)$ spans all of the $M(x)$ for all $x \in G$. Note that we can take s to be at most the dimension of the range of M , which is $O(|S_k|^d)$.

Let q be a generic element of V . Let q_i be the degree- k polynomial obtained by setting of the S_k -coordinates of the input of q to those of x_i . We note that $\Pr(q_i(y) = 0) \geq \eta$. By Corollary 11, this implies that there is a set T_i of at most $2^{O(d^2)}\eta^{-3}$ coordinates so that every non-zero, degree- k monomial in q_i has at least one coordinate in T_i . Since q_i is a generic linear combination of the p_j , this must mean that each p_j when restricted to x_i must have all its degree- k terms having a coordinate in T_i .

Let $S = S_k \cup \bigcup_{i=1}^s T_i$. We note that $|S| = O(\eta^{-d^2(d-k+1)})$. We will attempt to use S for S_{k-1} .

Rewrite each $p_\ell(x)$ as $r_\ell(x) + \sum m_{\ell j}(x_{S^c})c_{\ell j}(x_S)$, where all the monomials in r_ℓ have at most $k - 1$ coordinates not in S , each $m_{\ell j}$ is a monomial of degree- k with only coordinates not in S , and $c_{\ell j}$ a polynomial with coordinates in S . Note that for any x_i that $c_{\ell j}(x_i) = 0$. For any $x \in G$ $c_{\ell j}(x)$ is a linear combination of the $c_{\ell j}(x_i)$ and is thus also zero. We now split into two cases:

Case 1: $c_{\ell j} \in I$ for all ℓ and j .

In this case, $p_\ell \equiv r_\ell \pmod{I}$. Thus, taking the r_ℓ as the equivalent basis, it contains no monomial with more than $k - 1$ coordinates not in S , so we are done.

Case 2: Some $c_{\ell j} \notin I$.

We note that $c_{\ell j}$ is a degree less than d polynomial that is not in I . Furthermore, by the above, it vanishes on G . Thus, it vanishes on all but $\eta 2^n$ points of $Z(I) \cap Z(V)$. Thus, we can take $P = c_{\ell j}$.

This completes the inductive step, and proves our lemma. □

□

In order to properly analyze the process of iteratively applying the above proposition, some work needs to be done in order to find the correct inductive statement. The following is the one that works conveniently:

Proposition 34. *For any integers $d \geq k \geq 0$ there exists a function $h_{d,k} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ so that $\lim_{\delta \rightarrow 0} h_{d,k}(\delta) = 0$ and so that if f is any degree- d PTF, $g : \{-1, 1\}^n \rightarrow \text{bits}$ is any boolean function, and I any ideal containing all $x_i^2 - 1$ so that $f(x) \neq g(x)$ only for $x \in Z(I)$ and where $\text{Chow}_d(f, g) < \delta$, then there exists a degree at most k polynomial $P \notin I$ so that $P(x) = 0$ for all but an $h_{d,k}(\delta)$ -mass of the points where $f(x) \neq g(x)$.*

Before we prove this proposition, we note why it immediately implies Theorem 3. In particular, taking $k = 0$ and $I = (x_1^2 - 1, \dots, x_n^2 - 1)$, this says that if $\text{Chow}_d(f, g) < \delta$, there exists a *constant* function $P \neq 0$, so that all but an $h_{d,0}(\delta)$ -mass of the disagreements of f and g have $P = 0$ (which never happens). In particular, this means that

$$\text{Chow}_d(f, g) < \delta \Rightarrow \text{dist}(f, g) < h_{d,0}(\delta).$$

Proof. We proceed once again by backwards induction on k . We start with the base case of $k = d$. This will follow essentially from Proposition 14. The idea is simple. Let $f(x) = \text{sign}(p(x))$. Let I_d be the vector space of all degree-at-most- d polynomials in I . Let p' be the part of p perpendicular to I_d , and let $f'(x) = \text{sign}(p'(x))$. Notice that since $p' \equiv p \pmod{I}$, we have that $f(x) = f'(x)$ for $x \in Z(I)$. Let

$$g'(x) = \begin{cases} f'(x) & \text{if } x \notin Z(I) \\ g(x) & \text{otherwise} \end{cases}.$$

Note that $f(x) - g(x) = f'(x) - g'(x)$ and so the discrepancy sets are the same as are the Chow distances. Thus, it suffices to prove our result for f' and g' . However, we know from Proposition 14 that for $h_{d,d}(\delta) = \eta = 1/\text{poly}_d(\log \log(1/\delta))$, that there exists a polynomial P of degree at most d with P not perpendicular to p' so that all but an η -mass of the discrepancy of f' and g' lies in $Z(P)$. However, p' is orthogonal to all of I_d , and therefore P cannot be in I . This completes the proof when $k = d$.

For the inductive step, assume that there is such a function $h_{d,k}$ for some given value of k . Note that for any triple f, g, I satisfying our hypotheses, there is a P of degree at most k where all but an $h_{d,k}(\delta)$ -mass of the discrepancies lie on $Z(P)$. If we define $g'(x)$ to be $g(x)$ when $P(x) = 0$ and $f(x)$ otherwise, note that all of the discrepancies of f and g' lie on $Z(I + (P))$. Furthermore since g and g' differ in hamming weight $h_{d,k}(\delta)$, we have that

$$\text{Chow}_d(f, g') \leq \text{Chow}_d(f, g) + \text{Chow}_d(g, g') \leq \delta + O_d(h_{d,k}(\delta))^{1/3} := h'_{d,k}(\delta).$$

We are now able to define $h_{d,k-1}$. In particular, let $m_d(\eta)$ be the implied constant from Proposition 32. We claim that we can take $h_{d,k-1}$ to be less than η whenever $h'_{d,k}$ iterated on δ $m_k(\eta^2)$ times is less than η^2 . In particular, suppose that this holds for δ and we have that $\text{Chow}_d(f, g) < \delta$. We would like to claim that we can find an appropriate P of degree less than k with all but an η mass of the discrepancies between f and g lying on $Z(P)$.

Let $g = g_0$. We note that all discrepancies between f and g lie on $Z(I)$. By the above, there exists a degree- d polynomial $P_1 \notin I$ and a g_1 with $\text{dist}(g, g_1) < h_{d,k}(\delta)$, $\text{Chow}_d(f, g_1) < \delta_1 = h'_{d,k}(\delta)$ and all discrepancies between f and g_1 lying in $Z(I_1)$ where $I_1 = I + (P_1)$. Iterating this, we can find g_t with $\text{dist}(g_{t-1}, g_t) < h_{d,k}(\delta_{t-1})$, $\text{Chow}_d(f, g_t) < \delta_t = h'_{d,k}(\delta_{t-1})$ and all discrepancies between f and g_t lying in $Z(I_t)$, where $I_1 = I + (P_1, P_2, \dots, P_t)$. Take $t = m_d(\eta)$, so that $\delta_t < \eta^2$. Note that since $P_i \notin I + (P_1, \dots, P_{i-1})$, it must be the case that $V = \text{span}(P_1, \dots, P_t)$ has trivial intersection with I . Therefore, applying Proposition 32 to I and V , we find that there is a degree less than k polynomial P so that all but an η^2 -mass of the discrepancies between f and g_t (which in turn is all but an η^2 -mass of the discrepancies between f and g) lie on $Z(P)$.

This completes our inductive step and the proof. \square

4 Algorithmic Applications of Theorem 3

Our main structural result, Theorem 3, together with machinery developed in [TTV08, DDFS14, DKS18] yields the first efficient algorithms for the degree- d Chow parameters problem and for learning Boolean degree- d PTFs with malicious noise. As a corollary of the former result, we also obtain the existence of low integer-weight approximations to degree- d PTFs. In this section, we describe these applications and explain how they are obtained by combining Theorem 3 with prior work.

4.1 Degree- d Chow Parameters Problem and Low Integer-Weight Approximation

We start by proving Theorem 4. To prove our theorem, we need an efficient algorithm that starts with (approximations to) the degree- d Chow parameters of our degree- d PTF f and computes the coefficients of a degree- d multilinear polynomial that approximately sign represents f . This can be done by known techniques, as follows from prior work [TTV08, DDFS14].

We first need to define the notion of a projection:

Definition 35. For $a \in \mathbb{R}$, we denote its projection to $[-1, 1]$ by $P_1(a)$. That is, $P_1(a) = a$ if $|a| \leq 1$ and $P_1(a) = \text{sign}(a)$, otherwise.

Our algorithm will make essential use of a variant of degree- d PTFs, which we call degree- d *polynomial bounded functions* (PBFs):

Definition 36. A function $g : \{-1, 1\}^n \rightarrow [-1, 1]$ is referred to as a degree- d *polynomial bounded function* (PBF) if there exists a degree- d multilinear polynomial $Q : \{-1, 1\}^n \rightarrow \mathbb{R}$ with $Q(x) = \sum_{S \subseteq [n], |S| \leq d} Q_S \chi_S(x)$ such that $g(x) = P_1(Q(x))$. The vector of coefficients $(Q_S)_S$ is said to represent g .

The following result shows that there exists an efficient algorithm which, given approximations to the degree- d Chow parameters of an arbitrary Boolean-valued function, outputs a degree- d PBF with approximately these degree- d Chow parameters:

Theorem 37 (Degree- d Chow Reconstruction). *There exists a randomized algorithm `Chow-d-Reconstruct` that for every Boolean function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, given $\xi > 0, \delta > 0$ and a vector $\vec{\alpha} = (\alpha_S)_{S \subseteq [n], |S| \leq d}$ such that $\|\vec{\chi}_f^d - \vec{\alpha}\|_2 \leq \xi$, with probability at least $1 - \delta$, outputs a degree- d PBF g such that $\|\vec{\chi}_f^d - \vec{\chi}_g^d\|_2 \leq 6\xi$. The algorithm runs in time $\tilde{O}(n^{2d})\text{poly}(1/\xi) \log(1/\delta)$. Further, g is represented by a weight-vector λv , where $\lambda \in \mathbb{R}$ and v is an integer vector with $\|v\|_2^2 = O(n^d)\text{poly}(1/\xi)$.*

We remark that the condition on the weight vector v given by Theorem 37 is the key for the proof of Theorem 5. We note that Theorem 37 follows directly either from a more general in [TTV08] (specifically, their Theorem 3.1; also see Theorem 16 of [DDS17]), or from a straightforward generalization of the $d = 1$ algorithm for the same problem in [DDFS14].

The algorithm establishing Theorem 37 is a simple iterative algorithm running for $O(1/\xi^2)$ iterations, where in each iteration it estimates the degree- d Chow parameters of the current hypothesis h' within ℓ_2 error $O(\xi)$. Here, each h' is a degree- d PBF, i.e., function of the form $h'(x) = P_1(\frac{\xi}{2} \cdot \sum_{S \subseteq [n], |S| \leq d} H_S \chi_S(x))$, where the coefficients H_S are integers whose absolute values sum to $O(1/\xi^2)$.

Proof of Theorem 4. Given a vector $\vec{\alpha}$ such that $\Delta := \|\vec{\alpha} - \vec{\chi}_f^d\|_2 \leq \kappa(\epsilon, d)$, where f is the unknown degree- d PTF to be learned, we proceed as follows: To construct the desired hypothesis f^* , we run algorithm `Chow-d-Reconstruct` (from Theorem 37) on input $\vec{\alpha}$. The algorithm runs in time $\text{poly}(1/\Delta) \cdot \tilde{O}(n^{2d}) \cdot \log(1/\delta)$ and outputs a degree- d PBF g such that with probability at least $1 - \delta$ we have $\text{Chow}_d(f, g) \leq 6\Delta \leq 6\kappa(\epsilon, d)$. By Theorem 3, we get that with probability at least $1 - \delta$ we have $\text{dist}(f, g) \leq \epsilon/2$. (By setting the constants appropriately in the definition of $\kappa(\epsilon, d)$ above, we can guarantee that the conclusion of Theorem 3 is $\text{dist}(f, g) \leq \epsilon/2$.) Writing the degree- d PBF g as $g(x) = P_1(Q(x))$, we now claim that $f^*(x) = \text{sign}(Q(x))$ has $\text{dist}(f, f^*) \leq \epsilon$. This holds because, for each input $x \in \{-1, 1\}^n$, the contribution that x makes to $\text{dist}(f, f^*)$ is at most twice the contribution x makes to $\text{dist}(f, g)$. This completes the proof of Theorem 4. \square

As a simple corollary, we obtain Theorem 5.

Proof of Theorem 5. Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be an arbitrary degree- d PTF. We apply Theorem 4, for $\delta = 1/3$, and consider the degree- d PTF, f^* , output by the algorithm. Note that the weights v_i defining f^* are identical to the weights of the PBF g output by the algorithm `Chow-d-Reconstruct`. It follows from (the proof of) Theorem 4 that these weights are integers that satisfy $\sum_S v_S^2 = O(n^d \cdot \text{poly}(1/\Delta))$, where $\Delta = \Omega(\kappa(\epsilon, d))$, and the proof is complete. \square

Learning PTFs in the RFA Model. Ben-David and Dichterman [BDD98] introduced the “Restricted Focus of Attention” (RFA) learning framework to model the phenomenon of a learner having incomplete access to examples. We focus here on the uniform-distribution “ d -RFA” model. In this setting each time the learner is to receive a labeled example, it first specifies a set $J \subseteq [n]$ of at most d indices; then an n -bit string x is drawn from the uniform distribution over $\{-1, 1\}^n$ and the learner is given $(x_J, f(x))$. So for each labeled example, the learner is only shown the bits of the example indexed by J along with the label.

Note that learning in the d -RFA model is closely related to the degree- d Chow parameters problem. Indeed, a learning algorithm in the d -RFA model can only use the examples to estimate the degree- d Chow parameters of the unknown target concept f to any desired accuracy. This connection was established in Birkendorf et al. [BDJ⁺98] who asked the $d = 1$ version of this question, i.e., whether LTFs can be learned in the uniform distribution 1-RFA model. For $d = 1$, the structural results of Goldberg [Gol06] and Servedio [Ser07] established information-theoretic upper bounds on the sample complexity of the problem. The algorithmic results of [OS11, DDFS14] giving algorithms for the degree-1 Chow parameters problem immediately imply efficient algorithms for learning LTFs in the uniform distribution 1-RFA model.

As a direct consequence of Theorem 4, we obtain the first efficient learning algorithm for learning degree- d PTFs in the uniform distribution d -RFA model:

Theorem 38. *There is an algorithm which performs $\tilde{O}(n^{2d}) \cdot \text{poly}(1/\kappa(\epsilon, d)) \cdot \log(\frac{1}{\delta})$ bit-operations and properly learns degree- d PTFs to accuracy ϵ and confidence $1 - \delta$ in the uniform distribution d -RFA model.*

Prior to our work, even for $d = 2$, no sub-exponential in n upper bound was known for this problem, even with respect to sample complexity only.

4.2 Malicious Learning of Boolean Degree- d PTFs

In this section, we sketch how Theorem 6 is obtained using Theorem 3. The results of this section follow directly from the recent work [DKS18]. We provide a brief description here for the sake of completeness.

We focus on the efficient learnability of Boolean low-degree PTFs in the presence of *nasty noise*, a model of corruptions that strengthens malicious noise and has recently received renewed attention motivated by robust high-dimensional statistics [DKK⁺16]. In the *nasty noise model* [BEK02], an omniscient adversary can arbitrarily corrupt a small constant fraction of both the unlabeled data points and their labels. Formally, we have the following definition:

Definition 39 (Nasty Noise Learning Model). Let \mathcal{C} be a class of Boolean-valued functions over \mathbb{R}^n , D a distribution over \mathbb{R}^n , and f an unknown target concept $f \in \mathcal{C}$. For $0 < \epsilon < 1/2$, we say that a set T of m labeled examples is an ϵ -corrupted set of examples from \mathcal{C} if it is obtained using the following procedure: First, we draw a set $S = \{(x^{(i)}, y_i)\}$ of m labeled examples, $1 \leq i \leq m$, where for each i we have that $x^{(i)} \sim D$, $y_i = f(x^{(i)})$, and the $x^{(i)}$'s are independent. Then an omniscient adversary, upon inspecting the set S , is allowed to remove an ϵ -fraction of the examples and replace these examples by the same number of arbitrary examples of its choice. The modified set of labeled examples is the ϵ -corrupted set T . A learning algorithm in the nasty noise model is given as input an ϵ -corrupted set of examples from \mathcal{C} and its goal is to output a hypothesis h such that with high probability the error $\Pr_{x \sim D}[h(x) \neq f(x)]$ is small.

It should be noted that the nasty model generalizes a number of well-studied noise models, including the malicious noise model [Val85, KL93] and the agnostic (adversarial label noise) model [Hau92, KSS94].

A first qualitative goal is to design polynomial-time learning algorithms that can tolerate nasty noise of *constant* rate, i.e., we want to achieve error guarantees that are *independent of the dimension*. In recent work, [DKS18] obtained such an algorithm for degree- d PTFs over the Gaussian distribution (and under well-behaved continuous distributions).

The robust learning algorithm of [DKS18] for degree- d PTFs works in two steps: (1) Start by robustly approximating the degree- d Chow parameters of our function, and (2) Use the approximate degree- d Chow parameters from Step (1) to find a proper hypothesis that is close to the target concept.

Step (2) uses known algorithmic techniques [TTV08, DDFS14] (in particular, Theorem 37) to efficiently find an accurate proper hypothesis with approximately these degree- d Chow parameters. The correctness of Step (2) leverages the fact that *approximations* to the degree- d Chow parameters information-theoretically approximately determine our function. While such a structural result is easy to show for continuous well-behaved distributions, it was very challenging to prove for the uniform distribution on the hypercube. Our Theorem 4 provides the required structural result for degree- d PTFs over $\{-1, 1\}^n$, hence allowing the [DKS18] algorithmic approach to go through. This gives Theorem 6.

5 Conclusions and Open Problems

In this paper, we showed that the degree- d Chow parameters of a degree- d PTF f robustly determine f with respect to ℓ_1 -distance. As a corollary, we obtained a number of algorithmic and structural applications.

The main qualitative message of our result is that the relation between degree- d Chow distance and ℓ_1 -distance is independent of the dimension n . On the other hand, our quantitative bounds can almost certainly be improved. The obvious open problem is to prove a nearly tight relation between the two metrics. More specifically, what is the best possible function $\kappa(\epsilon, d)$ in the statement of Theorem 3? It is known that for $d = 1$ the answer is between $\epsilon^{\Omega(\log \log(1/\epsilon))}$ and $\epsilon^{O(\log^2(1/\epsilon))}$. For general d , standard bounds on approximating arbitrary degree- d PTFs by degree- d PTFs with small integer weights imply that $\kappa(\epsilon, d)$ cannot be larger than $\epsilon^{\Omega_d(\log(1/\epsilon)^{d-1})}$. We believe that obtaining such improved bounds requires new ideas.

References

- [ABL17] P. Awasthi, M. F. Balcan, and P. M. Long. The power of localization for efficiently learning linear separators with noise. *J. ACM*, 63(6):50:1–50:27, 2017.
- [APL07] H. Aziz, M. Paterson, and D. Leech. Efficient algorithm for designing weighted voting games. In *IEEE Intl. Multitopic Conf.*, pages 1–6, 2007.
- [Ban65] J. Banzhaf. Weighted voting doesn’t work: A mathematical analysis. *Rutgers Law Review*, 19:317–343, 1965.
- [Bau73] C. R. Baugh. Chow parameters in pseudothreshold logic. In *SWAT (FOCS)*, pages 49–55, 1973.
- [BDD98] S. Ben-David and E. Dichterman. Learning with restricted focus of attention. *Journal of Computer and System Sciences*, 56(3):277–298, 1998.
- [BDJ⁺98] A. Birkendorf, E. Dichterman, J. Jackson, N. Klasner, and H.U. Simon. On restricted-focus-of-attention learnability of Boolean functions. *Machine Learning*, 30:89–123, 1998.
- [BEK02] N. Bshouty, N. Eiron, and E. Kushilevitz. PAC Learning with Nasty Noise. *Theoretical Computer Science*, 288(2):255–275, 2002.

- [BGS18] A. Bhattacharyya, S. Ghoshal, and R. Saket. Hardness of learning noisy halfspaces using polynomial thresholds. In *Conference On Learning Theory, COLT 2018*, pages 876–917, 2018.
- [Bru90] J. Bruck. Harmonic analysis of polynomial threshold functions. *SIAM Journal on Discrete Mathematics*, 3(2):168–177, 1990.
- [Car04] F. Carreras. On the design of voting games. *Mathematical Methods of Operations Research*, 59(3):503–515, 2004.
- [CHIS10] M. Cheraghchi, J. Håstad, M. Isaksson, and O. Svensson. Approximating Linear Threshold Predicates. In *13th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems — APPROX 2010*, pages 110–123, 2010.
- [Cho61] C.K. Chow. On the characterization of threshold functions. In *Proceedings of the Symposium on Switching Circuit Theory and Logical Design (FOCS)*, pages 34–38, 1961.
- [CLO07] D. A. Cox, J. Little, and D. O’Shea. *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra, 3/e (Undergraduate Texts in Mathematics)*. Springer-Verlag, Berlin, Heidelberg, 2007.
- [CW01] A. Carbery and J. Wright. Distributional and L^q norm inequalities for polynomials over convex bodies in R^n . *Mathematical Research Letters*, 8(3):233–248, 2001.
- [Dan15] A. Daniely. A PTAS for agnostically learning halfspaces. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015*, pages 484–502, 2015.
- [Dan16] A. Daniely. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the 48th Annual Symposium on Theory of Computing, STOC 2016*, pages 105–117, 2016.
- [DDFS14] A. De, I. Diakonikolas, V. Feldman, and R. A. Servedio. Nearly optimal solutions for the chow parameters problem and low-weight approximation of halfspaces. *J. ACM*, 61(2):11:1–11:36, 2014.
- [DDS14] A. De, I. Diakonikolas, and R. A. Servedio. Deterministic approximate counting for juntas of degree-2 polynomial threshold functions. In *IEEE 29th Conference on Computational Complexity, CCC 2014*, pages 229–240, 2014.
- [DDS17] A. De, I. Diakonikolas, and R. A. Servedio. The inverse shapley value problem. *Games and Economic Behavior*, 105:122–147, 2017.
- [Der65] M. Dertouzos. *Threshold Logic: A Synthesis Approach*. MIT Press, Cambridge, MA, 1965.
- [DGJ⁺10] I. Diakonikolas, P. Gopalan, R. Jaiswal, R. Servedio, and E. Viola. Bounded independence fools halfspaces. *SIAM J. on Comput.*, 39(8):3441–3462, 2010.
- [DHK⁺10] I. Diakonikolas, P. Harsha, A. Klivans, R. Meka, P. Raghavendra, R. A. Servedio, and L. Y. Tan. Bounding the average sensitivity and noise sensitivity of polynomial threshold functions. In *STOC*, pages 533–542, 2010.
- [DKK⁺16] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Robust estimators in high dimensions without the computational intractability. In *Proceedings of FOCS’16*, pages 655–664, 2016.

- [dKKZ10] B. de Keijzer, T. Klos, and Y. Zhang. Enumeration and exact design of weighted voting games. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems : volume 1 - Volume 1*, AAMAS '10, pages 391–398, 2010.
- [DKN10] I. Diakonikolas, D. M. Kane, and J. Nelson. Bounded independence fools degree-2 threshold functions. In *FOCS*, pages 11–20, 2010.
- [DKS18] I. Diakonikolas, D. M. Kane, and A. Stewart. Learning geometric concepts with nasty noise. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 1061–1073, 2018.
- [DOSW11] I. Diakonikolas, R. O’Donnell, R. A. Servedio, and Y. Wu. Hardness results for agnostically learning low-degree polynomial threshold functions. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pages 1590–1606, 2011.
- [DRST14] I. Diakonikolas, P. Raghavendra, R. A. Servedio, and L. Y. Tan. Average sensitivity and noise sensitivity of polynomial threshold functions. *SIAM J. Comput.*, 43(1):231–253, 2014.
- [DS79] P. Dubey and L.S. Shapley. Mathematical properties of the Banzhaf power index. *Mathematics of Operations Research*, 4:99–131, 1979.
- [DS13] I. Diakonikolas and R. A. Servedio. Improved approximation of linear threshold functions. *Computational Complexity*, 22(3):623–677, 2013.
- [DS14] A. De and R. A. Servedio. Efficient deterministic approximate counting for low-degree polynomial threshold functions. In *Symposium on Theory of Computing, STOC 2014*, pages 832–841, 2014.
- [DSTW14] I. Diakonikolas, R. A. Servedio, L. Y. Tan, and A. Wan. A regularity lemma and low-weight approximators for low-degree polynomial threshold functions. *Theory of Computing*, 10:27–53, 2014.
- [EL89] E. Einy and E. Lehrer. Regular simple games. *International Journal of Game Theory*, 18:195–207, 1989.
- [Elg60] C.C. Elgot. Truth functions realizable by single threshold organs. In *Proceedings of the Symposium on Switching Circuit Theory and Logical Design (FOCS)*, pages 225–245, 1960.
- [Erd45] P. Erdős. On a lemma of Littlewood and offord. *Bull. Amer. Math. Soc.*, 51:898–902, 1945.
- [FGRW12] V. Feldman, V. Guruswami, P. Raghavendra, and Y. Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM J. Comput.*, 41(6):1558–1590, 2012.
- [FM04] D. Felsenthal and M. Machover. A priori voting power: what is it all about? *Political Studies Review*, 2(1):1–23, 2004.
- [Fre97] J. Freixas. Different ways to represent weighted majority games. *Top (Journal of the Spanish Society of Statistics and Operations Research)*, 5(2):201–212, 1997.
- [GKM15] P. Gopalan, D. M. Kane, and R. Meka. Pseudorandomness via the discrete fourier transform. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*, pages 903–922, 2015.

- [Gol06] P. Goldberg. A Bound on the Precision Required to Estimate a Boolean Perceptron from its Average Satisfying Assignment. *SIAM Journal on Discrete Mathematics*, 20:328–343, 2006.
- [Hås94] J. Håstad. On the size of weights for threshold gates. *SIAM Journal on Discrete Mathematics*, 7(3):484–492, 1994.
- [Hau92] D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100:78–150, 1992.
- [Hil90] D. Hilbert. Ueber die theorie der algebraischen formen. *Mathematische Annalen*, 36(4):473–534, Dec 1890.
- [HKM14] P. Harsha, A. R. Klivans, and R. Meka. Bounding the sensitivity of polynomial threshold functions. *Theory of Computing*, 10:1–26, 2014.
- [Hur73] S.L. Hurst. The application of Chow Parameters and Rademacher-Walsh matrices in the synthesis of binary functions. *The Computer Journal*, 16:165–173, 1973.
- [Kan14] D. M. Kane. The correct exponent for the gotsman-linial conjecture. *Computational Complexity*, 23(2):151–175, 2014.
- [Kan17] D. M. Kane. A structure theorem for poorly anticoncentrated polynomials of gaussians and applications to the study of polynomial threshold functions. *Ann. Probab.*, 45(3):1612–1679, 05 2017. Preliminary version in FOCS’12.
- [Kas63] P. Kaszerman. A geometric test-synthesis procedure for a threshold device. *Information and Control*, 6(4):381–398, 1963.
- [KL93] M. Kearns and M. Li. Learning in the presence of malicious errors. *SIAM Journal on Computing*, 22(4):807–837, 1993.
- [KN12] S. Kurz and S. Napel. Heuristic and exact solutions to the inverse power index problem for small voting bodies. Available as arxiv report <http://arxiv.org/abs/1202.6245>, 2012.
- [KSS94] M. Kearns, R. Schapire, and L. Sellie. Toward Efficient Agnostic Learning. *Machine Learning*, 17(2/3):115–141, 1994.
- [Kur12] S. Kurz. On the inverse power index problem. *Optimization*, 61(8):989–1011, 2012.
- [KW65] K.R. Kaplan and R.O. Winder. Chebyshev approximation and threshold functions. *IEEE Trans. Electronic Computers*, EC-14:315–325, 1965.
- [KW16] D. M. Kane and R. Williams. Super-linear gate and super-quadratic wire lower bounds for depth-two and depth-three threshold circuits. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC ’16, pages 633–643, 2016.
- [Lap72] E. Lapidot. The counting vector of a simple game. *Proceedings of the AMS*, 31:228–231, 1972.
- [Lee02a] D. Leech. Designing the voting system for the eu council of ministers. *Public Choice*, 113:437–464, 2002.
- [Lee02b] D. Leech. Voting power in the governance of the international monetary fund. *Annals of Operations Research*, 109:375–397, 2002.

- [Lee03] D. Leech. Power indices as an aid to institutional design: the generalised apportionment problem. In M. Holler, H.Kliemt, D. Schmidtchen, and M. Streit, editors, *Yearbook on New Political Economy*, 2003.
- [LO43] J. E. Littlewood and A. C. Offord. On the number of real roots of a random algebraic equation. *III. Rec. Math. [Mat. Sbornik] N.S.*, 12:277–286, 1943.
- [LW98] A. Laruelle and M. Widgren. Is the allocation of voting power among eu states fair? *Public Choice*, 94:317–339, 1998.
- [MNV16] R. Meka, O. Nguyen, and V. Vu. Anti-concentration for polynomials of independent random variables. *Theory of Computing*, 12(1):1–17, 2016.
- [MOO10] E. Mossel, R. O’Donnell, and K. K. Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, 171:295–341, 2010.
- [MP68] M. Minsky and S. Papert. *Perceptrons: an introduction to computational geometry*. MIT Press, Cambridge, MA, 1968.
- [MTB67] S. Muroga, T. Tsuboi, and C.R. Baugh. Enumeration of threshold functions of eight variables. Technical Report 245, Univ. of Illinois, Urbana, 1967.
- [MTK62] S. Muroga, I. Toda, and M. Kondo. Majority decision functions of up to six variables. *Math. Comput.*, 16:459–472, 1962.
- [MTT61] S. Muroga, I. Toda, and S. Takasu. Theory of majority switching elements. *J. Franklin Institute*, 271:376–418, 1961.
- [MZ13] R. Meka and D. Zuckerman. Pseudorandom generators for polynomial threshold functions. *SIAM J. Comput.*, 42(3):1275–1301, 2013.
- [O’D14] R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, New York, NY, USA, 2014.
- [OS11] R. O’Donnell and R. Servedio. The Chow Parameters Problem. *SIAM J. on Comput.*, 40(1):165–199, 2011.
- [Pen46] L.S. Penrose. The elementary statistics of majority voting. *Journal of the Royal Statistical Society*, 109(1):53–57, 1946.
- [Ros58] F. Rosenblatt. The Perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65:386–407, 1958.
- [RS10] Y. Rabani and A. Shpilka. Explicit construction of a small epsilon-net for linear threshold functions. *SIAM J. Comput.*, 39(8):3501–3520, 2010.
- [RV13] A. A. Razborov and E. Viola. Real advantage. *TOCT*, 5(4):17:1–17:8, 2013.
- [Ser07] R. Servedio. Every linear threshold function has a low-weight approximator. *Comput. Complexity*, 16(2):180–209, 2007.
- [SRK95] K.-Y. Siu, V.P. Roychowdhury, and T. Kailath. *Discrete Neural Computation: A Theoretical Foundation*. Prentice-Hall, Englewood Cliffs, NJ, 1995.

- [TT06] K. Takamiya and A. Tanaka. Computational complexity in the design of voting games. Technical Report 653, The Institute of Social and Economic Research, Osaka University, 2006.
- [TTV08] L. Trevisan, M. Tulsiani, and S. Vadhan. Regularity, Boosting and Efficiently Simulating every High Entropy Distribution . Technical Report 103, ECCC, 2008. Conference version in Proc. CCC 2009.
- [TZ92] A. Taylor and W. Zwicker. A Characterization of Weighted Voting. *Proceedings of the AMS*, 115(4):1089–1094, 1992.
- [Val85] L. Valiant. Learning disjunctions of conjunctions. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 560–566, 1985.
- [Win63] R.O. Winder. Threshold logic in artificial intelligence. *Artificial Intelligence*, IEEE Publication S-142:107–128, 1963.
- [Win64] R.O. Winder. Threshold functions through $n = 7$. Technical Report 7, Air Force Cambridge Research Laboratories, 1964.
- [Win69] R.O. Winder. Threshold gate approximations based on chow parameters. *IEEE Transactions on Computers*, pages 372–375, 1969.