Extremal properties of polynomial threshold functions

Ryan O'Donnell[*]
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
odonnell@cs.cmu.edu

Rocco A. Servedio[†]
Department of Computer Science
Columbia University
New York, NY 10025
rocco@cs.columbia.edu

June 5, 2007

Abstract

In this paper we give new extremal bounds on polynomial threshold function (PTF) representations of Boolean functions. Our results include the following:

- Almost every Boolean function has PTF degree at most $\frac{n}{2} + O(\sqrt{n \log n})$. Together with results of Anthony and Alon, this establishes a conjecture of Wang and Williams [32] and Aspnes, Beigel, Furst, and Rudich [4] up to lower order terms.

- Every Boolean function has PTF density at most $(1 - \frac{1}{O(n)})2^n$. This improves a result of Gotsman [14].

- Every Boolean function has weak PTF density at most $o(1)2^n$. This gives a negative answer to a question posed by Saks [28].

- PTF degree $\lfloor \log_2 m \rfloor + 1$ is necessary and sufficient for Boolean functions with sparsity m. This answers a question of Beigel [7].

We also give new extremal bounds on polynomials which approximate Boolean functions in the ℓ_∞ norm.

1 Introduction

A broad research goal in computational complexity is to understand the properties of various representation schemes for Boolean functions. Many representation schemes have been studied, such as DNF and CNF formulas, decision trees, branching programs, the Fourier representation (i.e. polynomials over the reals), polynomials over GF_2, monotone span programs, and so on.

Our main focus in this paper is on Boolean functions represented as *polynomial threshold functions*. Given a Boolean function $f : \{+1, -1\}^n \to \{+1, -1\}$, a polynomial threshold function (PTF) for f is an n-variable real polynomial p such that $\operatorname{sgn}(p(x)) = f(x)$ for all $x \in \{+1, -1\}^n$. (Alternatively, we sometimes say that such a polynomial p *sign-represents* f.)

Polynomial threshold functions play an important role in theoretical computer science. They are very useful in structural complexity theory; the Beigel *et al.* [8] proof that PP is closed under intersection uses clever constructions of polynomial threshold functions, and many oracle results have been obtained using PTFs, e.g. [4, 6, 13, 31]. Polynomial threshold functions can be viewed

---

	Strong PTF degree		Weak PTF degree	
	lower bound	upper bound	lower bound	upper bound
every function	n	n	n	n
almost every function	$\frac{n}{2}$	$\frac{n}{2} + O(\sqrt{n \log n})$ (†)	$\frac{n}{2} - O(\sqrt{n \log n})$ (†)	$\frac{n}{2}$

Table 1: Best bounds to date on strong and weak PTF degrees of n-variable Boolean functions. Lower bounds for "every function" mean that some function has this as a lower bound. Entries marked with (†) are new bounds proved in this paper.

	Strong PTF density		Weak PTF density	
	lower bound	upper bound	lower bound	upper bound
every function	$.11\ 2^n$	$(1 - \frac{1}{O(n)})2^n$ (†)	$2^{n/2}$ (†)	$o(1)2^n$ (†)
almost every function	$.11\ 2^n$	$(1 - \frac{1}{O(n)})2^n$ (‡)	$\frac{1}{2\sqrt{n}}2^{n/2}$ (†)	$\frac{2}{n}2^n$ (†)

Table 2: Best bounds to date on strong and weak PTF densities of n-variable Boolean functions. Lower bounds for "every function" mean that some function has this as a lower bound. Entries marked with (†) are new bounds proved in this paper. For (‡), we in fact show that *every* set of $(1 - \frac{1}{O(n)})2^n$ monomials can serve as a PTF support for almost every Boolean function.

as threshold-of-parity circuits and as such have been studied by researchers in circuit complexity [10, 11] and learning theory [18]. More recently, upper bounds on polynomial threshold function degree have been used to obtain learning algorithms for various classes of Boolean circuits [20, 19, 25]. Finally, polynomial threshold functions are an inherently interesting intermediate model of computation between purely algebraic models such as Fourier or GF_2 polynomials and purely combinatorial models such as decision trees or logic circuits. See Saks [28] for an extensive survey on polynomial threshold functions.

The two most basic complexity measures for a polynomial threshold function are its degree and its density (number of nonzero monomials). The *PTF degree* of a Boolean function f is the minimum degree over all polynomials p which sign-represent f, and the *PTF density* of f is the minimum density over all polynomials p which sign-represent f. Note that without loss of generality we may take any sign-representing polynomial to be multilinear, and hence every Boolean function has PTF degree at most n and PTF density at most 2^n.

Aspnes *et al.* [4] introduced a useful variant on polynomial threshold representations, namely, weak polynomial threshold representations. Given a Boolean function $f : \{+1, -1\}^n \rightarrow \{+1, -1\}$ we say the n-variable polynomial p is a *weak polynomial threshold representation* of f (alternatively, *p weakly sign-represents f*) if $p(x)$ is not identically 0 on $\{+1, -1\}^n$ and $\mathrm{sgn}(p(x)) = f(x)$ for all $x \in \{+1, -1\}^n$ such that $p(x) \neq 0$. The "Theorem of the Alternative" [4] shows that weak polynomial threshold representations are intimately connected to the usual threshold representations (see Theorem 3), and thus the study of *weak PTF degree* and *weak PTF density*, defined in analogy with PTF degree and PTF density, is of interest.

1.1 Previous Work

Prior to our work many authors have studied extremal properties of polynomial threshold functions. Here we touch briefly on the most relevant previous results (see Saks [28] for a detailed treatment).

In a famous result Minsky and Papert [23] proved upper and lower bounds of n for the PTF degree of the n-variable parity function. Aspnes *et al.* [4] proved upper and lower bounds of n for the weak PTF degree of parity as well. Both Aspnes *et al.* and Wang and Williams [32] conjectured that almost every n-variable Boolean function has PTF degree either $\lfloor n/2 \rfloor$ or $\lceil n/2 \rceil$. Toward this conjecture, Anthony [3] and Alon [1] used a counting argument to show that almost every Boolean function has PTF degree at least $n/2$. For the upper bound Razborov and Rudich [27] used a counting argument to show that almost every Boolean function has PTF degree at most $\frac{19}{20}n$, and Alon [1] used results of Gotsman [14] to show that almost every Boolean function has PTF degree at most $.89n$.

Less was known for PTF density. Saks [28] noted that results of Cover [12] imply that almost every Boolean function has PTF density at least $(.11)2^n$. Gotsman [14] proved that every Boolean function has PTF density at most $2^n - 2^{n/2}$. Aspnes *et al.* proved that every Boolean function has weak PTF density at most $\frac{1}{2}2^n$. Saks [28] has asked whether almost every Boolean function (i) has PTF density at most $(1-\epsilon)2^n$ for some $\epsilon > 0$, (ii) has weak PTF density at most $(\frac{1}{2} - \epsilon)2^n$ for some $\epsilon > 0$.

1.2 Our Results

We give many new extremal results on the degree and density of polynomial threshold functions. These results, which are summarized in Tables 1 and 2, improve on previous bounds and answer several of the questions described above. In addition to the results shown in Tables 1 and 2, we also prove a tight bound on the PTF degree of sparse Boolean functions, answering a question posed by Beigel [7].

We also give some new extremal bounds on *approximate* polynomial representations of Boolean functions. Approximating polynomials have multiple applications in complexity theory and quantum computation and have been studied by many authors, see e.g. [26, 24, 2, 5]. For $\epsilon \in [0, 1)$ we say that an n-variable real polynomial p is an ϵ-approximating polynomial for f if $|p(x) - f(x)| \le \epsilon$ for all $x \in \{+1, -1\}$, i.e. if p is an ϵ-approximator for f in the ℓ_∞ norm. It is easy to see that an approximating polynomial satisfies a stricter condition than a strong PTF, and in many cases our bounds for strong PTFs follow directly from bounds which we establish for ϵ-approximating polynomials. Our results on approximating polynomials are proved in Sections 3 through 5 and are summarized in Table 3.

1.3 Organization of the paper

In Section 2 we give some necessary background on strong and weak threshold representations and approximating polynomials, tail bounds, and Fourier analysis. Section 3 gives our new upper bound on PTF density and approximating polynomial density for all Boolean functions. Our upper bounds on PTF density and degree and approximating polynomial density and degree for almost all Boolean functions are in Section 4. Section 5 gives lower bounds on density and degree of approximating polynomials for almost all Boolean functions. In Section 6 we give new upper and lower bounds on weak PTF density for all and almost all Boolean functions. Finally, we prove a tight bound on the PTF degree of sparse Boolean functions in Section 7. We close in Section 8 with suggestions for future work and a conjecture.

| | ϵ-approximating degree | | ϵ-approximating density | |
	lower bound	upper bound	lower bound	upper bound
every function	n	n	$.11\ 2^n$	$(1 - \Omega(\frac{\epsilon^2}{n}))2^n$ (†)
a.e. function	$\frac{n}{2} + \Omega\left(\sqrt{n\log\frac{1}{\epsilon}}\right)$ (†)	$\frac{n}{2} + O\left(\sqrt{n\log\frac{n}{\epsilon}}\right)$ (†)	$.11\ 2^n$	$(1 - \Omega(\frac{\epsilon^2}{n}))2^n$ (‡)

Table 3: Best bounds to date on density and degree for ϵ-approximating polynomials for n-variable Boolean functions. Lower bounds for "every function" mean that some function has this as a lower bound. Entries marked with (†) are new bounds proved in this paper — for the range of ϵ for which they hold, please see the relevant theorems. For (‡), we in fact show that *every* set of $(1 - \Omega(\frac{\epsilon^2}{n}))\,2^n$ monomials can serve as an ϵ-approximating polynomial support for almost every Boolean function.

2 Preliminaries

2.1 Sign-representing and approximating polynomials

In this section we formally define weak and strong polynomial threshold functions (PTFs) and ϵ-approximating polynomials.

Definition 1 *Let $f\colon \{+1,-1\}^n \to \{+1,-1\}$ be a Boolean function and let $p\colon \mathbf{R}^n \to \mathbf{R}$ be a real multilinear polynomial.*

- *We say that p (strongly) sign-represents f, or p is a (strong) PTF for f, if $p(x) \neq 0$ for all $x \in \{+1,-1\}^n$ and $sgn(p(x)) = f(x)$ for all $x \in \{+1,-1\}^n$.*

- *We say that p weakly sign-represents f, or p is a weak PTF for f, if $p(x) \neq 0$ for at least one $x \in \{+1,-1\}^n$ and $sgn(p(x)) = f(x)$ for all $x \in \{+1,-1\}^n$ such that $p(x) \neq 0$.*

- *For $\epsilon \in [0,1)$ we say that p is an ϵ-approximating polynomial for f if $|p(x) - f(x)| \leq \epsilon$ for all $x \in \{+1,-1\}^n$.*

Since $x^2 = 1$ for all $x \in \{+1,-1\}$, the assumption that p is a multilinear polynomial is without loss of generality. Any multilinear p can be written as a linear combination of all 2^n multilinear monomials over $x_1, \ldots, x_n$; we let $\mathcal{M}$ denote this set of 2^n monomials. The *support* of p is the set of monomials which have non-zero coefficients in p, and the *density* of p is the number of monomials in the support.

It is easy to see that for $\epsilon \in [0,1)$ any ϵ-approximating polynomial for f is also a strong PTF for f, and that any strong PTF for f is also a weak PTF for f. Every Boolean function has a 0-approximating polynomial, i.e., a multilinear polynomial which agrees exactly with f on all inputs in $\{+1,-1\}^n$. This is precisely the *Fourier* representation of f as described in Section 2.3.

The two most important complexity measures for PTFs and approximating polynomials are *degree* and *density*.

Definition 2 *Given a Boolean function f, we say the* strong PTF *(respectively,* weak PTF, ϵ-approximating*) degree of f is the minimum degree over all polynomials which strongly sign-represent (respectively, weakly sign-represent, ϵ-approximate) f. We similarly define the* strong PTF, weak PTF, *and ϵ-approximating density of f.*

It follows from the above discussion that for each of these three notions (strong, weak, and ϵ-approximating) every Boolean function has degree at most n and has density at most 2^n.

We will use the so-called "Theorem of the Alternative" of Aspnes *et al.* [4] which relates weak and strong representations. This theorem follows immediately from the theorems of the alternative used for proving linear programming duality (e.g., Farkas's Lemma, the Stiemke Transposition Theorem). See [4, 25, 28] for more details.

Theorem 3 *Let $\mathcal{S}$ be any set of monomials over $x_1, \ldots, x_n$ and let f be any Boolean function. Then exactly one of the following statements is true:*

1. *f has a strong representation with support in $\mathcal{S}$;*

2. *f has a weak representation with suport in $\mathcal{M} - \mathcal{S}$.*

2.2 Concentration bounds

The following three concentration bounds will be useful for us. The first is quite standard, the second and third somewhat less so:

Chernoff bound: *Let $X_1, \ldots, X_m$ be independent ± 1 random variables. Let $\overline{X} = \frac{1}{m} \sum_{i=1}^{m} X_i$, and let $\mu = \mathbf{E}[\overline{X}]$. Then for all $\epsilon > 0$ we have*

$$\Pr[|\overline{X} - \mu| \geq \epsilon] \leq 2 \exp\left(-\frac{1}{2}\epsilon^2 m\right).$$

Hoeffding bound [17]: *Let $X_1, \ldots, X_m$ be independent random variables with common mean μ and bounded deviance from the mean, $|X_i - \mu| \leq c$. Let $\sigma^2 = \frac{1}{m} \sum_{i=1}^{m} Var[X_i]$, and let $X = \sum_{i=1}^{m} X_i$. Then for each $0 < t < cm$, we have:*

$$\Pr\left[|X - \mu m| \geq t\right] \leq \begin{cases} 2\left(\frac{3\sigma^2 m}{ct}\right)^{t/c} & \text{when } 3\sigma^2 m/ct < 1, \\ 2\exp\left(-\frac{t^2}{4\sigma^2 m}\right) & \text{when } 3\sigma^2 m/ct \geq 1. \end{cases}$$

This inequality also holds in the scenario where $X_1, \ldots, X_m$ are chosen without replacement from a fixed population $\{\alpha_1, \ldots, \alpha_N\}$, and μ and σ^2 denote $\frac{1}{N} \sum_{i=1}^{N} \alpha_i$ and $\frac{1}{N} \sum_{i=1}^{N} (\alpha_i - \mu)^2$, respectively.

Talagrand's Deviation Inequality [30] (see [21, Sec. 3.2]): for the statement here) *Let $v_1, \ldots v_m$ be fixed vectors in $\mathbf{R}^d$, let $X_1, \ldots, X_m$ be independent ± 1 random variables, and let $X = \sum_{i=1}^{m} X_i v_i$. Then for all $\epsilon > 0$ we have*

$$\Pr\left[\|X\| \geq \mathbf{E}[\|X\|] + \epsilon\right] \leq \exp\left(-\frac{\epsilon^2}{8\sigma^2}\right),$$

where

$$\sigma^2 = \max_{\|u\| \leq 1} \sum_{i=1}^{m} \langle u, v_i \rangle^2.$$

2.3 Fourier background

We view Boolean functions as maps $\{+1, -1\}^n \to \{+1, -1\}$. We consider the vector space V of all real-valued functions on $\{+1, -1\}^n$ endowed with inner product $\langle \cdot, \cdot \rangle$ defined by

$$\langle f, g \rangle = \mathbf{E}[f(x)g(x)],$$

where the expectation is over a uniform choice of $x \in \{+1, -1\}^n$. For $S \subseteq [n]$ we write x_S to denote $\prod_{i \in S} x_i$. As is well known, the collection of functions $\{x_S\}_{S \subseteq [n]}$ forms an orthonormal basis for V. We denote $\langle f(x), x_S \rangle$ by $\hat{f}(S)$ and hence for any function f,

$$f(x) = \sum_{S \subseteq [n]} \hat{f}(S) x_S.$$

This 0-approximating polynomial is known as the *Fourier representation* of f. Thus the Fourier coefficient $\hat{f}(S)$ is precisely the coefficient of x_S in the (unique) multilinear polynomial for f.

We will write $\|\hat{f}\|_p$ for the quantity $\left(\sum_{S \subseteq [n]} |\hat{f}(S)|^p \right)^{1/p}$. We also write $\|\hat{f}\|_\infty$ for $\max_S |\hat{f}(S)|$. An easy consequence of the orthonormality of $\{x_S\}$ is Plancherel's identity: for any $f, g : \{+1, -1\}^n \to \mathbf{R}$,

$$\mathbf{E}[fg] = \sum_{S \subseteq [n]} \hat{f}(S)\hat{g}(S).$$

As a special case we have Parseval's identity: for any $f : \{+1, -1\}^n \to \mathbb{R}$,

$$\|\hat{f}\|_2^2 = 2^{-n} \sum_{x \in \{+1, -1\}^n} f(x)^2.$$

In particular, all Boolean functions $f : \{+1, -1\}^n \to \{+1, -1\}$ have $\|\hat{f}\|_2 = 1$.

For $\mathcal{S} \subseteq 2^{[n]}$ define $f_{\mathcal{S}}$ to be the real-valued multilinear polynomial given by

$$f_{\mathcal{S}}(x) = \sum_{S \in \mathcal{S}} \hat{f}(S) x_S,$$

so $f_{\mathcal{S}}$ is obtained by zeroing the Fourier coefficients of all monomials x_T such that $T \notin \mathcal{S}$. We will often use the following simple fact:

Fact 4 *Let $f : \{+1, -1\}^n \to \{+1, -1\}$ be any Boolean function. Suppose that $\mathcal{S} \subseteq [n]$ is such that $\sum_{S \notin \mathcal{S}} |\hat{f}(S)| < \epsilon$. Then $f_{\mathcal{S}}$ is an ϵ-approximating polynomial for f.*

We conclude this section by analyzing the Fourier coefficients of a randomly chosen Boolean function $f : \{+1, -1\}^n \to \{+1, -1\}$. Let f be constructed by choosing each of its 2^n values independently and uniformly from $\{+1, -1\}$. It is easy to see that each Fourier coefficient $\hat{f}(S)$ is distributed as $2^{-n}B(\pm 1, 2^n)$, where $B(\pm 1, 2^n)$ denotes the binomial random variable given by adding 2^n independent uniformly random ± 1 values. Hence $\mathbf{E}[\hat{f}(S)] = 0$ and $\mathbf{E}[\hat{f}(S)^2] = 2^{-n}$. Furthermore, a Chernoff bound tells us that $\Pr_f[|\hat{f}(S)| \geq c\sqrt{n}2^{-n/2}] < 2\exp(-\frac{1}{2}c^2 n)$. Taking (for example) $c = 2$, a union bound over all 2^n sets S lets us conclude:

Proposition 5 *All but a 2^{-n} fraction of Boolean functions $f : \{+1, -1\}^n \to \{+1, -1\}$ satisfy $\|\hat{f}\|_\infty < 2\sqrt{n}\, 2^{-n/2}$.*

Gotsman noted this fact in [14].

The Fourier coefficients of a random function are not independent random variables. Nevertheless, we can prove that for any fixed set of Fourier coefficients $\mathcal{S} \subseteq 2^{[n]}$, for a random Boolean function f the value $\|\hat{f}_{\mathcal{S}}\|_2^2$ (which equals the L_2 Fourier weight of f on $\mathcal{S}$) is tightly concentrated around its expectation, $|\mathcal{S}|2^{-n}$.

Proposition 6 *Let $\mathcal{S} \subseteq 2^{[n]}$ be a collection of $\mu 2^n$ monomials. Then for a random Boolean function f, for any $0 < \epsilon \le 1$ we have that*

$$\Pr_f \left[\sum_{S \in \mathcal{S}} \hat{f}(S)^2 \ge \mu + \epsilon \right] \le \exp\left(-\frac{\epsilon^2}{72} 2^n \right).$$

Proof: In Talagrand's Deviation Inequality we take $m = 2^n$ and consider

$$X = \sum_{x \in \{+1,-1\}^n} f(x) v_x,$$

where the $f(x)$'s play the role of the random ± 1 bits X_i, and the vectors $v_x \in \mathbf{R}^{\mathcal{S}}$ are defined by

$$v_x = (2^{-n} x_S)_{S \in \mathcal{S}}.$$

By definition, the random vector $X \in \mathbf{R}^{\mathcal{S}}$ consists of the $\mathcal{S}$-Fourier coefficients of f. In particular, $\|X\|^2 = \sum_{S \in \mathcal{S}} \hat{f}(S)^2$ and $\mathbf{E}[\|X\|^2] = \mu$.

We now compute the quantity σ^2. Let A be the $|\mathcal{S}| \times 2^n$ matrix given by placing the column vectors v_x's side by side. Then for any vector u we have that $\sum_x \langle u, v_x \rangle^2 = \|u^\top A\|^2$. Hence

$$\sigma^2 = \max_{\|u\| \le 1} \sum_{x \in \{+1,-1\}^n} \langle u, v_x \rangle^2$$

is equal to the square of the largest singular value of A. This in turn equals the largest magnitude among the eigenvalues of $AA^\top$. But the rows of A are orthogonal, so $AA^\top$ is equal to 2^{-n} times the $|\mathcal{S}| \times |\mathcal{S}|$ identity matrix. We conclude that $\sigma^2 = 2^{-n}$, and therefore Talagrand's Deviation Inequality tells us that

$$\Pr\big[\|X\| \ge \mathbf{E}[\|X\|] + \delta\big] \le \exp\left(-\frac{\delta^2}{8} 2^n \right).$$

So except with probability at most $\exp\left(-\frac{\delta^2}{8} 2^n\right)$ we have

$$\|X\| < \mathbf{E}[\|X\|] + \delta \quad \Rightarrow \quad \|X\|^2 < (\mathbf{E}[\|X\|] + \delta)^2$$
$$\Rightarrow \quad \|X\|^2 < \mathbf{E}[\|X\|]^2 + 3\delta \le \mathbf{E}[\|X\|^2] + 3\delta = \mu + 3\delta,$$

where the second step used $\mathbf{E}[\|X\|] \le 1$ (which holds because $\|X\|^2 \le 1$ always). Taking $\delta = \epsilon/3$ completes the proof. $\square$

3   A new upper bound for PTF density

We first study the maximum PTF density of any Boolean function. As noted earlier, for any $f : \{+1, -1\}^n \to \{+1, -1\}$ the PTF density of f is at most 2^n. Gotsman [14] obtained a slightly better bound of $2^n - 2^{n/2} + 1$. The proof is straightforward: Let $\mathcal{T}$ denote the set of $2^{n/2} - 1$

monomials on which f has Fourier coefficients of smallest magnitude. Since $\|\hat{f}\|_2 = 1$, we have $\|\hat{f}\|_1 \le 2^{n/2}$, and hence the sum of the magnitudes of the smallest $2^{n/2}$ Fourier coefficients is at most 1. Thus $\sum_{S\in\mathcal{T}} |\hat{f}(S)| < 1$, so $\text{sgn}(f_{\mathcal{M}-\mathcal{T}})$ sign-represents f by Fact 4.

In this section we improve this upper bound to $(1 - \frac{1}{O(n)})2^n$. Indeed, we show that for all f and sufficiently large ϵ, f has an ϵ-approximating polynomial of density $(1 - \Omega(\frac{\epsilon^2}{n}))2^n$:

Theorem 7 *Let $f\colon \{+1,-1\}^n \to \{+1,-1\}$ be any Boolean function. Then for every $\epsilon > n2^{-n/4}$, f has an ϵ-approximating polynomial of density at most $(1-\Omega(\frac{\epsilon^2}{n}))2^n$. In particular, f has a PTF of density at most $(1 - \frac{1}{O(n)})2^n$.*

Proof: Let $L = \|\hat{f}\|_1$. Bruck and Smolensky [11] gave a randomized construction showing that f has PTF density at most $\lceil 2nL^2 \rceil$. It is easy to see that their proof generalizes to give an ϵ-approximating polynomial of density at most $\lceil \frac{2nL^2}{\epsilon^2} \rceil$. Hence if $L \le (\epsilon/2\sqrt{n})2^{n/2}$ then f has an ϵ-approximating polynomial of density at most $\frac{1}{2}2^n$; consequently we assume $L > (\epsilon/2\sqrt{n})2^{n/2}$. Since $\epsilon > n2^{-n/4}$, we conclude that $L > 2n/\epsilon$ with room to spare.

Since L is large, f must have a very large number of very small Fourier coefficients. The basic idea of the proof is that if we throw out a few monomials with very small Fourier coefficients, the function's values are unlikely to change by more than an additive ϵ.

To this end, let $\mathcal{T}$ be the set of monomials on which f has its Fourier coefficients of smallest magnitude, where the cutoff is selected so that:

$$\sum_{S\notin\mathcal{T}} |\hat{f}(S)| \in [n/\epsilon - 2, n/\epsilon - 1). \tag{1}$$

Note that this makes sense since $L > 2n/\epsilon$.

Since the average value of $|\hat{f}(S)|$ on $\mathcal{M} - \mathcal{T}$ must be at least the overall average, namely $2^{-n} \sum_{S\in\mathcal{M}} |\hat{f}(S)| = 2^{-n}L$, we conclude $(n/\epsilon-1)/|\mathcal{M}-\mathcal{T}| \ge 2^{-n}L$, whence $|\mathcal{T}| \ge (1-(n/\epsilon-1)/L)2^n$. Let N denote $|\mathcal{T}|$, the number of small Fourier coefficients. Using $L > 2n/\epsilon$, we conclude that $N > \frac{1}{2}2^n$.

We now randomly select $m = (\epsilon^2/Cn)2^n$ of the monomials in $\mathcal{T}$ without replacement; call the resulting set of monomials $\mathcal{S}$. (Here $C > 0$ is a large absolute constant to be determined later.) We will prove that for every fixed $x \in \{+1,-1\}^n$, the probability (over the choice of S) that $|\sum_{S\in\mathcal{S}} \hat{f}(S)x_S| > \epsilon$ is at most 3^{-n}. By a union bound, it will follow that for some particular collection $\mathcal{S}$, the polynomial

$$f_{\mathcal{M}-\mathcal{S}}(x) = \sum_{S\notin\mathcal{S}} \hat{f}(S)x_S$$

is within ϵ of $f(x)$ for *every* $x \in \{+1,-1\}^n$; i.e., $f_{\mathcal{M}-\mathcal{S}}$ is an ϵ-approximating polynomial for f. This will prove the theorem, since $f_{\mathcal{M}-\mathcal{S}}$ has density $2^n - m$, as desired.

Fix x, and let $(\alpha_1, \ldots, \alpha_N)$ denote the list of numbers $(\hat{f}(S)x_S)_{S\in\mathcal{T}}$. We have:

$$\left| \sum_{i=1}^{N} \alpha_i \right| = \left| \sum_{S\subseteq[n]} \hat{f}(S)x_S - \sum_{S\notin\mathcal{T}} \hat{f}(S)x_S \right|$$

$$\le \left| \sum_{S\subseteq[n]} \hat{f}(S)x_S \right| + \sum_{S\notin\mathcal{T}} |\hat{f}(S)x_S|$$

$$< 1 + (n/\epsilon - 1) = n/\epsilon.$$

Write $\mu = \frac{1}{N} \sum_{i=1}^{N} \alpha_i$, so $|\mu| < n/\epsilon N$. Now we bound σ^2:

$$
\begin{aligned}
\sigma^2 &:= \frac{1}{N} \sum_{i=1}^{N} (\alpha_i - \mu)^2 \\
&\leq \frac{1}{N} \sum_{i=1}^{N} 2(\alpha_i^2 + \mu^2) \\
&= 2\mu^2 + \frac{2}{N} \sum_{i=1}^{N} \alpha_i^2 \\
&\leq 2\mu^2 + 2/N \\
&< 3/N
\end{aligned}
$$

where the next to last inequality is by Parseval's identity and the last is since $\mu^2 < (n/\epsilon N)^2 < 2^{n/2}/N^2 < 1/2N$, since $\epsilon > n2^{-n/4}$ and $N > \frac{1}{2}2^n$.

Finally, we have that $|\alpha_i| \leq 1/(n/\epsilon - 2)$ and hence $|\alpha_i - \mu| \leq 2\epsilon/n$ for all $1 \leq i \leq N$. The second of these inequalities follows from the first since $1/(n/\epsilon - 2) < 1.5\epsilon/n$, and $|\mu| < n/\epsilon N < .5\epsilon/n$ (using $\epsilon > n2^{-n/4}$, $N > \frac{1}{2}2^n$). To see the first inequality, note that otherwise we would have $|\hat{f}(S)| > 1/(n/\epsilon - 2)$ for all $S \notin \mathcal{T}$, hence $|\mathcal{M} - \mathcal{T}| < (n/\epsilon - 2)^2$ because $\sum_{S \notin \mathcal{T}} \hat{f}(S)^2 \leq 1$ by Parseval. But by (1) and Cauchy-Schwarz we have:

$$
\begin{aligned}
n/\epsilon - 2 &\leq \sum_{S \notin \mathcal{T}} |\hat{f}(S)| \\
&\leq \sqrt{|\mathcal{M} - \mathcal{T}|} \sqrt{\sum_{S \notin \mathcal{T}} \hat{f}(S)^2} \\
&\leq \sqrt{|\mathcal{M} - \mathcal{T}|},
\end{aligned}
$$

which is a contradiction.

Now we consider selecting m of the α_i's at random. Let X denote the sum of the selected numbers. Our goal is to show that $|X| > \epsilon$ with probability at most 3^{-n}. By Hoeffding's bound, with $t = \epsilon/2$ and $c = 2\epsilon/n$, we have:

$$
\begin{aligned}
\Pr\left[|X - \mu m| \geq t\right] &\leq 2(3\sigma^2 m/ct)^{t/c} \\
\Rightarrow \Pr\left[|X| \geq |\mu m| + t\right] &\leq 2(3\sigma^2 m/ct)^{t/c} \\
\Rightarrow \Pr\left[|X| \geq (n/\epsilon N)(\epsilon^2/Cn)2^n + \epsilon/2\right] &\leq 2\left[(9/C)(2^n/N)\right]^{n/4} \\
\Rightarrow \Pr\left[|X| \geq (2/C)\epsilon + \epsilon/2\right] &\leq (18/C)^{n/4} \\
\Rightarrow \Pr\left[|X| \geq \epsilon\right] &\leq 3^{-n},
\end{aligned}
$$

Here the third inequality follows from the previously established bounds $|\mu| < n/\epsilon N$, $\sigma^2 \leq 3/N$ and the substitution $m = (\epsilon^2/Cn)2^n$. The fourth inequality follows from $N > \frac{1}{2}2^n$, and the final inequality follows by taking C to be a large enough constant. $\qquad\square$

4 Upper bounds on density and degree for almost all functions

In the previous section we showed that every Boolean function has PTF density at most $(1 - \frac{1}{O(n)})2^n$. In this section we show that *every* subset of $(1 - \frac{1}{O(n)})2^n$ monomials can serve

as a polynomial threshold support for almost every Boolean function. Indeed, every subset of $(1 - \Omega(\frac{\epsilon^2}{n}))2^n$ monomials can serve as the support of an ϵ-approximating polynomial for almost every Boolean function. Precisely:

Theorem 8 *Let $\mathcal{S} \subseteq 2^{[n]}$ be any collection of subsets of $[n]$ such that $|\mathcal{S}| \geq (1 - \frac{\epsilon^2}{6n})2^n$. Then for all but a 2^{-n} fraction of Boolean functions f on n bits, there is a ϵ-approximating polynomial p for f whose support is contained in $\mathcal{S}$.*

An interesting special case of Theorem 8 occurs when we take $\epsilon = 1/n$ and $\mathcal{S}$ to be the $(1 - \frac{1}{6n^3})2^n$ smallest subsets of $2^{[n]}$. By the Chernoff bound we then have that $|S| \leq \frac{n}{2} + O(\sqrt{n \log n})$ for all $S \in \mathcal{S}$. We thus obtain the following corollary:

Corollary 9 *Almost all Boolean functions have $\frac{1}{n}$-approximating polynomial degree (and hence also PTF degree) at most $\frac{n}{2} + O(\sqrt{n \log n})$.*

As noted earlier, Anthony and also Alon have used a counting argument to show that almost every Boolean function has PTF degree at least $\lfloor n/2 \rfloor$. Together with this lower bound, our upper bound answers in the affirmative a conjecture of Wang and Williams [32] and Aspnes *et al.* [4] up to lower order terms. (They conjectured that almost all Boolean functions have PTF degree exactly $n/2$.) We have been informed that a PTF degree upper bound of $n/2 + O(\sqrt{n \log n})$ for almost every function has also been independently proved by Samorodnitsky [29].

Using the Theorem of the Alternative, Aspnes *et al.* gave a simple proof that for any n-bit Boolean function f, the sum of the strong degree of f and the weak degree of $f \cdot \mathsf{PARITY}_n$ is exactly n (Lemma 2.5 of [4]). Hence Corollary 9 also implies that almost all Boolean functions have weak degree at least $n/2 - O(\sqrt{n \log n})$.

4.1 Proof of Theorem 8

Let $f : \{+1, -1\}^n \to \{+1, -1\}$ be a randomly chosen Boolean function. In the sequel, all probabilities are taken over this choice of f. To motivate our proof of Theorem 8, consider the earlier proof of Alon and Gotsman (see [14, 28]); they show the weaker PTF upper bound of $2^n - \frac{1}{2\sqrt{n}}2^{n/2}$ by combining Proposition 5 and Fact 4. (This gives the PTF degree upper bound of $.89n$ noted in Section 1.1 by taking $\mathcal{S}$ to be the set of monomials of lowest degree.) Alon and Gotsman's argument uses a "worst-case" assumption about the magnitude of the sum of the omitted Fourier coefficients. If the Fourier coefficients of the random function f were not just binomially distributed but were *independent* random variables, then we could use standard tail inequalities on sums of independent random variables to obtain a stronger bound. However the Fourier coefficients are not at all independent, so this direct approach does not seem to work.

We get around this by showing that in fact the error term $\sum_{S \notin \mathcal{S}} \hat{f}(S) x_S$ can be viewed as a sum of independent random variables. These new independent variables no longer correspond to the individual Fourier coefficients $\hat{f}(S)$; nevertheless, we can exactly characterize the variance of the sum of these new random variables, and this enables us to push the argument through.

We now proceed with the proof. For $z \in \{+1, -1\}^n$ let $\delta_z : \{+1, -1\}^n \to \mathbb{R}$ be the function

$$\delta_z(x) = \begin{cases} 1 & \text{if } x = z \\ 0 & \text{otherwise.} \end{cases}$$

The Fourier representation of δ_z is easily seen to be

$$\delta_z(x) \;=\; \frac{(1+z_1x_1)(1+z_2x_2)\cdots(1+z_nx_n)}{2^n}$$

$$=\; \frac{1}{2^n}\sum_{S\subseteq[n]} z_Sx_S.$$

Consequently any function $f:\{+1,-1\}^n \to \mathbb{R}$ may be written as:

$$f(x) = \sum_{z\in\{+1,-1\}^n} f(z)\frac{1}{2^n}\sum_{S\subseteq[n]} z_Sx_S.$$

For any $\mathcal{S}\subseteq 2^{[n]}$ we thus have

$$f_{\mathcal{S}}(x) = \frac{1}{2^n}\sum_{z\in\{+1,-1\}^n} f(z)\sum_{S\in\mathcal{S}} z_Sx_S. \tag{2}$$

Let $\delta_{\mathcal{S},z}(x) = \sum_{S\in\mathcal{S}} z_Sx_S$. It is clear that $\delta_{\mathcal{S},x}(x) = |\mathcal{S}|$ for any $x\in\{+1,-1\}^n$. We now claim:

Lemma 10 *For any $x\in\{+1,-1\}^n$, we have $\sum_{z\neq x}\delta_{\mathcal{S},z}(x)^2 = 2^n|\mathcal{S}| - |\mathcal{S}|^2$.*

Proof:

$$\sum_{z\neq x}\delta_{\mathcal{S},z}(x)^2 \;=\; \sum_{z\in\{+1,-1\}^n} \delta_{\mathcal{S},z}(x)^2 - \delta_{\mathcal{S},x}(x)^2$$

$$=\; \sum_{z\in\{+1,-1\}^n} \delta_{\mathcal{S},z}(x)^2 - |\mathcal{S}|^2$$

$$=\; \sum_{z\in\{+1,-1\}^n} \delta_{\mathcal{S},x}(z)^2 - |\mathcal{S}|^2 \tag{3}$$

$$=\; 2^n\sum_{S\subseteq[n]} \hat{\delta}_{\mathcal{S},x}(S)^2 - |\mathcal{S}|^2 \tag{4}$$

$$=\; 2^n|\mathcal{S}| - |\mathcal{S}|^2, \tag{5}$$

where (3) is because $\delta_{\mathcal{S},z}(x) = \delta_{\mathcal{S},x}(z)$, (4) is Parseval's identity, and (5) follows because $\delta_{\mathcal{S},x}$ has exactly $|\mathcal{S}|$ nonzero Fourier coefficients, each of magnitude exactly 1. $\qquad\square$

To prove Theorem 8, fix any $\mathcal{S}\subseteq 2^{[n]}$ with $|\mathcal{S}| \geq (1-\frac{\epsilon^2}{6n})2^n$. Fix any $x\in\{+1,-1\}^n$. We claim that for a random Boolean function f, with probability at least $1-4^{-n}$ we have

$$\left|\sum_{z\neq x} f(z)\delta_{\mathcal{S},z}(x)\right| < \epsilon|\mathcal{S}|. \tag{6}$$

Given this, a union bound lets us conclude that for all but a 2^{-n} fraction of functions f, inequality (6) holds for *every* $x\in\{+1,-1\}^n$. Since we have

$$2^n f_{\mathcal{S}}(x) = \sum_{z\in\{+1,-1\}^n} f(z)\delta_{\mathcal{S},z}(x) = |\mathcal{S}|f(x) + \sum_{z\neq x} f(z)\delta_{\mathcal{S},z}(x)$$

11

this implies that

$$
\begin{aligned}
|2^n f_{\mathcal{S}}(x) - |\mathcal{S}| f(x)| &= \left| \sum_{z \neq x} f(z) \delta_{\mathcal{S},z}(x) \right| \\
&\leq \epsilon |\mathcal{S}|,
\end{aligned}
$$

for all x, and hence $(2^n/|\mathcal{S}|) f_{\mathcal{S}}$ is a ϵ-approximating polynomial for f which is supported on $\mathcal{S}$.

To see the claim, note that since each $f(z)$ is an independent random ± 1 value, we may view the sum over $z \neq x$ in (6) as a sum of $2^n - 1$ independent random variables, where the zth random variable takes values $\pm \delta_{\mathcal{S},z}(x)$ each with probability $1/2$. From Lemma 10 we know that the sum of the squares of $\delta_{\mathcal{S},z}(x)$ is precisely $2^n |\mathcal{S}| - |\mathcal{S}|^2$, and hence the variance of the sum of these $2^n - 1$ random variables is precisely $\sigma^2 = \frac{2^n |\mathcal{S}| - |\mathcal{S}|^2}{2^n - 1}$. We can bound each random variable's deviance from the mean 0 by noting that $|\delta_{\mathcal{S},z}(x)| \leq 2^n - |\mathcal{S}|$ for all $z \neq x$ (this holds since by adding $\sum_{S \notin \mathcal{S}} z_S x_S$ to $\delta_{\mathcal{S},z}(x)$ we would get $\sum_{S \subseteq [n]} z_S x_S$ which is 0). Hence by Hoeffding's bound, with $m = 2^n - 1$, $t = \epsilon |\mathcal{S}|$, and $c = 2^n - |\mathcal{S}|$, we have

$$
\begin{aligned}
\Pr \left[\left| \sum_{z \neq x} f(z) \delta_{\mathcal{S},z}(x) \right| \geq \epsilon |\mathcal{S}| \right] &\leq 2 \exp\left(-\frac{\epsilon^2 |\mathcal{S}|}{4(2^n - |\mathcal{S}|)} \right) \qquad (7) \\
&\leq 2 \exp\left(-\frac{6n(1 - \epsilon^2/6n)}{4} \right) \\
&\leq 4^{-n},
\end{aligned}
$$

as desired, where the second-last step uses $|\mathcal{S}| \geq (1 - \epsilon^2/6n) 2^n$. $\qquad$ (Theorem 8) $\blacksquare$

5 Lower bounds for ϵ-approximating density and degree

J. Håstad communicated to us a proof that if $\mathcal{S} \subseteq 2^{[n]}$ contains noticeably less than a $1 - 4\epsilon$ fraction of the monomials in $\mathcal{M}$, then only a tiny fraction of Boolean functions can have an ϵ-approximating polynomial supported on $\mathcal{S}$. (Note the contrast with Theorem 8.)

Theorem 11 *(J. Håstad) Given $0 < \epsilon < \frac{1}{5}$, let $\mathcal{S}$ be any set of $(1 - 5\epsilon) 2^n$ monomials. Then the fraction of Boolean functions which can have an ϵ-approximating polynomial supported on $\mathcal{S}$ is at most $\exp\left(-\frac{\epsilon^2}{72} 2^n \right)$.*

Proof: Suppose g is an ϵ-approximating polynomial for f. Then for every x, $f(x)g(x) > 1 - \epsilon$. Hence:

$$
\begin{aligned}
(1-\epsilon)^2 \;&\leq\; \mathbf{E}[fg]^2 \\[4pt]
&= \left(\sum_{S \subseteq [n]} \hat{f}(S)\hat{g}(S) \right)^2 && \text{(Plancherel)} \\[4pt]
&= \left(\sum_{S \in \mathcal{S}} \hat{f}(S)\hat{g}(S) \right)^2 && \text{(since } g \text{ is supported on } \mathcal{S}) \\[4pt]
&\leq \left(\sum_{S \in \mathcal{S}} \hat{f}(S)^2 \right)\left(\sum_{S \in \mathcal{S}} \hat{g}(S)^2 \right) && \text{(Cauchy-Schwarz)} \\[4pt]
&= \left(\sum_{S \in \mathcal{S}} \hat{f}(S)^2 \right) \mathbf{E}[g^2] && \text{(Parseval)} \\[4pt]
&\leq \left(\sum_{S \in \mathcal{S}} \hat{f}(S)^2 \right)(1+\epsilon)^2,
\end{aligned}
$$

where the last step uses the fact that $|g(x)| \leq 1+\epsilon$ for every x. Hence we have that $\sum_{S \in \mathcal{S}} \hat{f}(S)^2 \geq \frac{(1-\epsilon)^2}{(1+\epsilon)^2} = 1 - \frac{4\epsilon}{(1+\epsilon)^2} > 1-4\epsilon$. But since the expected value of $\sum_{S \in \mathcal{S}} \hat{f}(S)^2$ is $1-5\epsilon$ for a random f, by Proposition 6 at most an $\exp\left(-\frac{\epsilon^2}{72}2^n\right)$ fraction of all Boolean functions have $\sum_{S \in \mathcal{S}} \hat{f}(S)^2 > 1-4\epsilon$.
$\square$

This theorem immediately yields the following interesting corollary by taking $\mathcal{S}$ to be the $(1-5\epsilon)2^n$ smallest monomials in $\mathcal{M}$:

Corollary 12 *For any $2^{-n/2} \ll \epsilon \leq \frac{1}{5}$ it holds that almost all Boolean functions have ϵ-approximating polynomial degree at least $n/2 + \Omega(\sqrt{n\log\frac{1}{\epsilon}})$.*

Note that in the case $\epsilon = \frac{1}{n}$ the lower bound of Corollary 12 essentially matches the upper bound given by Corollary 9.

6 Weak PTF density

In this section we give an upper bound on weak PTF density which holds for all Boolean functions and a stronger upper bound which holds for almost all Boolean functions. These bounds give a negative answer to a question of M. Saks. We also give a lower bound on weak PTF density which holds for almost all Boolean functions and a stronger lower bound which holds for a particular Boolean function. To the best of our knowledge these are the only lower bounds known for weak PTF density.

6.1 Upper bounds for weak PTF density

Since any strong representation of a Boolean function f is also a weak representation, Theorem 3 implies that for any function f and any set $\mathcal{S} \subseteq \mathcal{M}$ of monomials either f has a weak representation with support contained in $\mathcal{S}$ or f has a weak representation with support contained in $\mathcal{M} - \mathcal{S}$ (or

both). Taking $\mathcal{S}$ to be any set of $\frac{1}{2}2^n$ monomials, it follows that every Boolean function has weak density at most $\frac{1}{2}2^n$.

M. Saks has asked the following question (Question 2.28.2 of [28]): is it true that for all $\epsilon > 0$ almost all Boolean functions have weak density at least $(\frac{1}{2} - \epsilon)2^n$? Our next two theorems show that the answer is "no" in a rather strong sense:

Theorem 13 *Almost all Boolean functions have weak density at most $\frac{2}{n}2^n$.*

Theorem 14 *All Boolean functions have weak density $o(1)2^n$.*

The intuition behind the proof of Theorem 13 is straightforward: with high probability a random Boolean function f has some small subcube on which f is "simple." We take advantage of this simplicity to construct a low-density polynomial p which weakly represents f on this subcube. Multiplying p by another polynomial which is 0 off of the subcube, we obtain a weak representative for f. More precisely, we use the following lemma:

Lemma 15 *Let τ be a restriction which fixes $n - k$ variables from $x_1, \ldots, x_n$ and keeps k variables free. Let D denote the weak density of $f|_\tau$. Then the weak density of f is at most $2^{n-k}D$.*

Proof: Without loss of generality we can suppose that τ is the restriction which maps variables $x_1, \ldots, x_{n-k}$ to 1 and leaves the remaining k variables free. Let p be a polynomial over $x_{n-k+1}, \ldots, x_n$ which weakly represents $f|_\tau$ and has D nonzero monomials. Then the polynomial

$$P(x_1, \ldots, x_n) \quad = \quad (x_1 + 1)(x_2 + 1) \cdots (x_{n-k} + 1) \cdot p(x_{n-k+1}, \ldots, x_n)$$

has density $2^{n-k}D$. To see that P weakly represents f, note that on any input $x = 1^{n-k}y$ we have $P(x) = 2^{n-k}p(x)$, while on any other input we have $P(x) = 0$. Since p is a weak representative of $f|_\tau$ it must be somewhere nonzero, so the same is true for P. $\qquad\square$

Proof of Theorem 13: Let f be a random Boolean function. Consider the 2^{n-k} disjoint k-dimensional subcubes of $\{+1, -1\}$ corresponding to restrictions τ which fix variables $x_1, \ldots, x_{n-k}$. For any such restriction τ we have

$$\Pr[f|_\tau \text{ is not identically } 1] = 1 - \frac{1}{2^{2^k}}$$

and hence

$$\Pr[f|_\tau \not\equiv 1 \text{ for all such } \tau] = \left(1 - \frac{1}{2^{2^k}}\right)^{2^{n-k}}.$$

Taking $k = \log n - 1$, the above probability is $(1 - 2^{-n/2})^{2^{n-\log n + 1}} < e^{-2^{n/2+1}/n}$. Thus with overwhelmingly high probability there is some restriction τ fixing $n - \log n + 1$ variables such that $f|_\tau$ is identically 1, and hence the weak density of $f|_\tau$ is 1. Now use Lemma 15. $\qquad\square$

Using Lemma 15 it is easy to prove an upper bound of $\frac{1}{2}2^n$ on the weak density of all Boolean functions without using Theorem 3. For any Boolean function f on n variables, the polynomial

$$(x_1 + 1)(x_2 + 1) \cdots (x_{n-1} + 1)y$$

is easily seen to be a weak representative of f which has density $\frac{1}{2}2^n$, where $y \in \{-1, 1, -x_n, x_n\}$ is suitably chosen depending on the two values of $f(1^{n-1}, 1)$ and $f(1^{n-1}, -1)$.

By looking at subcubes of dimension greater than 1 it is possible to improve this bound. A straightforward case analysis shows the following:

Fact 16 *Every Boolean function on 3 variables has weak density at most* 3.

Together with Lemma 15, this yields

Corollary 17 *Every Boolean function has weak density at most* $\frac{3}{8}2^n$.

While Corollary 17 already gives a strong negative answer to the question of Saks, we can obtain the stronger upper bound of Theorem 14 by using more powerful tools from Ramsey theory. A *k-dimensional affine subspace* of a vector space V is a translate of a k-dimensional vector subspace of V. The following is a special case of the Affine Ramsey Theorem of Graham *et al.* [15, 16]:

Theorem 18 *Let A be a finite field. For all $r, k \geq 1$ there exists n such that if the points of A^n are r-colored, then some k-dimensional affine subspace of A^n has all of its points the same color.*

Taking $r = 2$ and $A = GF_2$, we can rephrase this as:

Corollary 19 *There is a function $g(n) = \omega(1)$ such that for any Boolean function $f : (GF_2)^n \to \{-1, 1\}$, there is some $g(n)$-dimensional affine subspace of $(GF_2)^n$ on which f is constant.*

Proof of Theorem 14: Let f be any Boolean function on n variables and let W' be the affine subspace whose existence is asserted by Corollary 19. Any $g(n)$-dimensional vector subspace W of $(GF_2)^n$ is the set of solutions to some system of $n - g(n)$ homogeneous linear equations, i.e.,

$$W = \{x \in (GF_2)^n : Ax = (0)^{n-g(n)}\}$$

where A is an $(n - g(n)) \times n$ matrix over GF_2. Thus the $g(n)$-dimensional affine subspace W' is the set of solutions to some system of $n - g(n)$ not necessarily homogeneous linear equations, i.e.,

$$W' = \{x \in (GF_2)^n : Ax = b\}$$

for some $b \in (GF_2)^{n-g(n)}$. If we identify GF_2 with the set $\{+1, -1\}$, then this system of equations becomes:

$$\prod_{j:A_{1,j}=1} x_j = b_1,$$

$$\prod_{j:A_{2,j}=1} x_j = b_2,$$

$$\vdots$$

$$\prod_{j:A_{n-g(n),j}=1} x_j = b_{n-g(n)}.$$

Without loss of generality we may suppose that $f(x) = 1$ for all $x \in W'$. It is easy to see that the points of $\{+1, -1\}^n$ on which the polynomial

$$\prod_{i=1}^{n-g(n)} \left(b_i \left(\prod_{j:A_{i,j}=1} x_i \right) + 1 \right)$$

is nonzero are exactly the points in W', and that moreover this polynomial always takes value exactly $2^{n-g(n)}$ on W'. Thus this polynomial is a weak representative for f of density $2^{n-g(n)} = o(1)2^n$, and Theorem 14 is proved. $\qquad \square$

6.2 Lower bounds for weak PTF density

Here we give our lower bounds for weak PTF density. The first lower bound holds for almost every Boolean function:

Theorem 20 *Almost all Boolean functions have weak PTF density at least $\frac{1}{2\sqrt{n}}2^{n/2}$.*

Proof: Recall the proof of Theorem 8; in particular, equation (7). If we take $\epsilon = \frac{9}{10}$ (this choice is somewhat arbitrary since any value less than 1 will do) and consider a fixed set $\mathcal{S}$ of size $(1-\tau)2^n$, then Equation (7) tells us that the probability that f has no PTF over $\mathcal{S}$ is bounded by

$$2\exp\left(-\frac{\epsilon^2|\mathcal{S}|}{4(2^n-|\mathcal{S}|)}\right)\cdot 2^n = 2^{n+1}\exp\left(-\frac{81(1-\tau)}{400\tau}\right) < 2^{n+1}\exp\left(-\frac{1}{5\tau}\right)$$

where the extra 2^n factor comes from a union bound over all $x \in \{+1,-1\}^n$ and the last inequality holds for small enough τ (any $\tau = o(1)$ suffices). There are exactly $\binom{2^n}{\tau 2^n}$ sets $\mathcal{S}$ of size $(1-\tau)2^n$. Hence if we select τ such that $\binom{2^n}{\tau 2^n}2^{n+1}\exp\left(-\frac{1}{5\tau}\right)$ is at most $1/2^n$, then a union bound tells us that all but a $1/2^n$ fraction of Boolean functions can be sign-represented using *any* set of $(1-\tau)2^n$ monomials. In this case Theorem 3 implies that for almost every Boolean function, no set of $\tau 2^n$ monomials can serve as the support of a weak sign-representation. Taking $\tau = \frac{1}{2\sqrt{n}}2^{-n/2}$, it is easily shown that $\binom{2^n}{\tau 2^n}\cdot 2^{n+1}\exp\left(-\frac{1}{5\tau}\right) \ll 1/2^n$, and the theorem is proved. $\qquad\square$

We can give a slightly better bound for an explicit Boolean function. For $n = 2k$ let IP denote the "inner product mod 2" function, i.e. $IP(x_1,\ldots,x_k,y_1,\ldots,y_k) = \oplus_{i=1}^k(x_i \wedge y_i)$ where $\oplus$ denotes exclusive-OR (parity) and $\wedge$ denotes AND.

Theorem 21 *IP has weak density at least $2^{n/2}$.*

Proof: It is known [10, 22] that IP is a *bent* function, i.e. a function for which $|\hat{f}(S)| = \frac{1}{2^{n/2}}$ for all $S \subseteq [n]$. Consequently, for any set $\mathcal{S}$ of $2^n - 2^{n/2} + 1$ monomials, the function $\text{sgn}(f_{\mathcal{S}}(x))$ is a strong representative of f by Fact 4. By Theorem 3 this means that for any set $\mathcal{T}$ of $2^{n/2} - 1$ monomials, it is not the case that f has a weak representative whose support is contained in $\mathcal{T}$. Hence the weak density of f is at least $2^{n/2}$. $\qquad\square$

7 PTF degree of sparse functions

The following question was posed by R. Beigel [7]: are sparse sets easy for low-degree polynomial threshold functions? More concretely, let $f : \{+1,-1\}^n \to \{+1,-1\}$ be a Boolean function such that $|f^{-1}(1)| = m \ll 2^n$, so f is the characteristic function of a sparse subset of the Boolean cube. What is the maximum polynomial threshold function degree for such an f? The following theorem gives a complete answer for all values of m:

Theorem 22 *For $1 \le m \le \frac{1}{2}2^n$, let $\mathcal{F}_m$ be the set of all Boolean functions $f : \{+1,-1\}^n \to \{+1,-1\}$ such that $m = \min\{|f^{-1}(1)|,|f^{-1}(-1)|\}$. Then the maximum PTF degree over all $f \in \mathcal{F}_m$ is exactly $\lfloor \log m \rfloor + 1$.*

Proof: We assume without loss of generality that $1 \le |f^{-1}(1)| = m \le \frac{1}{2}2^n$. For the lower bound, let f be any function which is such that if the last $n - (\lfloor \log m \rfloor + 1)$ inputs are fixed to 1 then f computes parity on the first $\lfloor \log m \rfloor + 1$ inputs. (Note that this uses up $2^{\lfloor \log m \rfloor} \le m$ of the ones in

f's output; any remaining ones can be located arbitrarily). Since any polynomial threshold function which computes parity on k variables must have degree at least k, it follows that any polynomial threshold function for f must have degree at least $\lfloor \log m \rfloor + 1$.

For the upper bound, we begin by constructing an m-leaf decision tree over variables $x_1, \ldots, x_n$ such that each string in $f^{-1}(1)$ arrives at a different leaf. Such a tree can be constructed by a greedy algorithm: initially all strings in $f^{-1}(1)$ are at the root of the tree. Let x_i be any variable such that there are two strings in $f^{-1}(1)$ which disagree on x_i (such a variable must exist as long as $|f^{-1}(1)| \geq 2$). Label the root with x_i. The strings $\{x : x \in f^{-1}(1), x_i = -1\}$ go to the left child and the strings $\{x : x \in f^{-1}(1), x_i = 1\}$ go to the right child. Now recurse on each child. At the end of this process we have an m-leaf tree in which each (unlabeled) leaf has a unique string in $f^{-1}(1)$ which reaches that leaf.

Let ℓ be a leaf in this tree and let z be the element of $f^{-1}(1)$ which reaches that leaf. We label ℓ with the degree-1 polynomial threshold function $\mathrm{sgn}(p(x))$ where $p(x) = x_1 z_1 + \cdots + x_n z_n - n + \frac{1}{2}$. Note that $p(z) = \frac{1}{2}$, and $p(x) \leq -\frac{1}{2}$ for all binary inputs $x \neq z$. Thus we now have an m-leaf decision tree T in which internal nodes are labeled with variables and leaves are labeled with degree-1 polynomial threshold functions, such that T computes exactly f.

The rest of our proof follows the proof of Theorem 2 in [20]. Recall that the rank of a decision tree T is defined inductively as follows:

- If T is a single leaf then $rank(T) = 0$.

- If T has subtrees T_0 and T_1 then $rank(T)$ equals $\max(rank(T_0), rank(T_1))$ if $rank(T_0) \neq rank(T_1)$ and equals $rank(T_0) + 1$ if $rank(T_0) = rank(T_1)$.

It follows from this definition that the rank of an m-leaf tree is at most $\lfloor \log m \rfloor$. Now we use the fact (see [9]) that a rank-r decision tree with functions $f_1, f_2, \ldots, f_m$ at the leaves is equivalent to some r-decision list, i.e., to a function "if $C_1(x)$ then output $f_1(x)$ else if $C_2(x)$ then output $f_2(x)$ else $\ldots$ else output $f_m(x)$" where each C_i is a conjunction on at most r variables. Thus, our decision tree T is equivalent to such a decision list, where $r = \lfloor \log m \rfloor$ and each f_i is a degree-1 polynomial threshold function $\mathrm{sgn}(p_i)$ as described above.

We now show that the degree-$(\lfloor \log m \rfloor + 1)$ polynomial threshold function $\mathrm{sgn}(P(x))$ computes T, where $P(x)$ equals

$$A_1 \tilde{C}_1(x) p_1(x) + A_2 \tilde{C}_2(x) p_2(x) + \cdots A_m \tilde{C}_m(x) p_m(x).$$

Here $\tilde{C}_i$ is the polynomial of degree at most $\lfloor \log m \rfloor$ which outputs 1 if C_i is true and 0 if C_i is false, and $A_1 \gg A_2 \gg A_3 \gg \cdots \gg A_m > 0$ are appropriately chosen positive values. To see that this works, note that if C_i is the first conjunction in the decision list which is satisfied by x, then we have

$$P(x) = A_i p_i(x) + \sum_{j > i, C_j(x) = 1} A_j p_j(x).$$

Since $|p_i(x)| \geq \frac{1}{2}$ and $A_i \gg A_j > 0$ for $j > i$, the sign of $P(x)$ is the same as the sign of $p_i(x)$, and we are done. $\qquad \square$

8 Conclusion

While we have made significant progress on extremal bounds for PTF degree and PTF density, there is still room for improvement. One goal is to improve the lower order term in our $n/2 + O(\sqrt{n \log n})$

upper bound for the PTF degree of almost every Boolean function. Another goal is to give tighter bounds on the maximum PTF density of Boolean functions. Saks [28] has asked whether almost all Boolean functions have PTF density at least $(1 - \epsilon)2^n$ for some $\epsilon > 0$. We conjecture that the answer is "no" in a strong sense:

Conjecture 23 *For n sufficiently large, every Boolean function $f : \{+1, -1\}^n \to \{+1, -1\}$ has PTF density at most $\frac{1}{2}2^n$.*

Finally, a large gap remains between our upper and lower bounds for weak PTF density; it would be interesting to tighten these bounds.

9   Acknowledgements

We would like to thank Johan Håstad for telling us the proof of Theorem 11 and letting us reproduce it here.

References

[1] N. Alon. Personal communication to M. Saks, reported in M. Saks, "Slicing the Hypercube", 1993.

[2] A. Ambainis. A note on quantum black-box complexity of almost all Boolean functions. *Information Processing Letters*, 71, 1999.

[3] M. Anthony. Classification by polynomial surfaces. *Discrete Applied Mathematics*, 61:91–103, 1995.

[4] J. Aspnes, R. Beigel, M. Furst, and S. Rudich. The expressive power of voting polynomials. *Combinatorica*, 14(2):1–14, 1994.

[5] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001.

[6] R. Beigel. Perceptrons, PP, and the Polynomial Hierarchy. *Computational Complexity*, 4:339–349, 1994.

[7] R. Beigel. . Personal communication, 2000.

[8] R. Beigel, N. Reingold, and D. Spielman. PP is closed under intersection. *Journal of Computer and System Sciences*, 50(2):191–202, 1995.

[9] A. Blum. Rank-r decision trees are a subclass of r-decision lists. *Information Processing Letters*, 42(4):183–185, 1992.

[10] J. Bruck. Harmonic analysis of polynomial threshold functions. *SIAM Journal on Discrete Mathematics*, 3(2):168–177, 1990.

[11] J. Bruck and R. Smolensky. Polynomial threshold functions, AC^0 functions and spectral norms. *SIAM Journal on Discrete Mathematics*, 21:33–42, 1992.

[12] T. Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Transactions on Electronic Computers*, EC-14(3):326–334, 1965.

[13] B. Fu. Separating ph from pp by relativization. *Acta Math. Sinica*, 8(3):329–336, 1992.

[14] C. Gotsman. On boolean functions, polynomials and algebraic threshold functions. Technical Report TR-89-18, Department of Computer Science, Hebrew University, 1989.

[15] R. L. Graham, K. Leeb, and B. L. Rothschild. Ramsey's theorem for a class of categories. *Adv. Math.*, 8:417–433, 1972.

[16] R. L. Graham, B. L. Rothschild, and J. L. Spencer. *Ramsey Theory*. Wiley-Interscience, New York, 1990.

[17] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

[18] J. Jackson. An efficient membership-query algorithm for learning DNF with respect to the uniform distribution. *Journal of Computer and System Sciences*, 55:414–440, 1997.

[19] A. Klivans, R. O'Donnell, and R. Servedio. Learning intersections and thresholds of halfspaces. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science*, pages 177–186, 2002.

[20] A. Klivans and R. Servedio. Learning DNF in time $2^{\tilde{O}(n^{1/3})}$. In *Proceedings of the Thirty-Third Annual Symposium on Theory of Computing*, pages 258–265, 2001.

[21] M. Ledoux. *Concentration of measure and logarithmic Sobolev inequalities. In Séminaire de Probabilités XXXIII*, pages 120–216. Springer, 1999.

[22] F.J. MacWilliams and N.J.A. Sloane. *The Theory of Error-Correcting Codes*. North-Holland, New York, 1977.

[23] M. Minsky and S. Papert. *Perceptrons: an introduction to computational geometry*. MIT Press, Cambridge, MA, 1968.

[24] N. Nisan and M. Szegedy. On the degree of Boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1994.

[25] R. O'Donnell and R. Servedio. New degree bounds for polynomial threshold functions. In *Proceedings of the 35th ACM Symposium on Theory of Computing*, pages 325–334, 2003.

[26] R. Paturi. On the degree of polynomials that approximate symmetric Boolean functions. In *Proceedings of the 24th Symposium on Theory of Computing*, pages 468–474, 1992.

[27] A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55(1):24–35, 1997.

[28] M. Saks. *Slicing the hypercube*, pages 211–257. London Mathematical Society Lecture Note Series 187, 1993.

[29] A. Samorodnitsky. Unpublished manuscript, 2000. Personal communication, 2003.

[30] M. Talagrand. An isoperimetric theorem on the cube and the Khintchine-Kahane inequalities. *Proceedings of the American Mathematical Society*, 104:905–909, 1988.

[31] N. Vereshchagin. Lower bounds for perceptrons solving some separation problems and oracle separation of AM from PP. In *Proceedings of the Third Annual Israel Symposium on Theory of Computing and Systems*, 1995.

[32] C. Wang and A.C. Williams. The threshold order of a boolean function. *Discrete Applied Mathematics*, 31:51–69, 1991.