

Average sensitivity and noise sensitivity of polynomial threshold functions

Ilias Diakonikolas* Prasad Raghavendra† Rocco A. Servedio‡
Li-Yang Tan§

October 19, 2009

Abstract

We give the first non-trivial upper bounds on the average sensitivity and noise sensitivity of degree- d polynomial threshold functions (PTFs). These bounds hold both for PTFs over the Boolean hypercube $\{-1, 1\}^n$ and for PTFs over \mathbb{R}^n under the standard n -dimensional Gaussian distribution $\mathcal{N}(0, I_n)$. Our bound on the Boolean average sensitivity of PTFs represents progress towards the resolution of a conjecture of Gotsman and Linial [GL94], which states that the symmetric function slicing the middle d layers of the Boolean hypercube has the highest average sensitivity of all degree- d PTFs. Via the L_1 polynomial regression algorithm of Kalai et al. [KKMS08], our bounds on Gaussian and Boolean noise sensitivity yield polynomial-time agnostic learning algorithms for the broad class of constant-degree PTFs under these input distributions.

The main ingredients used to obtain our bounds on both average and noise sensitivity of PTFs in the Gaussian setting are tail bounds and anti-concentration bounds on low-degree polynomials in Gaussian random variables [Jan97, CW01]. To obtain our bound on the Boolean average sensitivity of PTFs, we generalize the “critical-index” machinery of [Ser07] (which in that work applies to halfspaces, i.e. degree-1 PTFs) to general PTFs. Together with the “invariance principle” of [MOO05], this lets us extend our techniques from the Gaussian setting to the Boolean setting. Our bound on Boolean noise sensitivity is achieved via a simple reduction from upper bounds on average sensitivity of Boolean PTFs to corresponding bounds on noise sensitivity.

*Department of Computer Science, Columbia University. Email: ilias@cs.columbia.edu. Research supported by NSF grant CCF-0728736, and by an Alexander S. Onassis Foundation Fellowship. Part of this work was done while visiting IBM Almaden.

†Microsoft Research, New England. Email: prasad@cs.washington.edu. Part of the research done while at the University of Washington and visiting Carnegie Mellon University supported by NSF CCF-0343672.

‡Department of Computer Science, Columbia University. Email: rocco@cs.columbia.edu. Supported by NSF grants CCF-0347282, CCF-0523664 and CNS-0716245, and by DARPA award HR0011-08-1-0069.

§Department of Computer Science, Columbia University. Email: liyong@cs.columbia.edu. Supported by DARPA award no. HR0011-08-1-0069 and NSF Cybertrust grant no. CNS-0716245.

1 Introduction

A degree- d polynomial threshold function (PTF) over a domain $X \subseteq \mathbb{R}^n$ is a Boolean-valued function $f : X \rightarrow \{-1, +1\}$,

$$f(x) = \text{sign}(p(x_1, \dots, x_n))$$

where $p : X \rightarrow \mathbb{R}$ is a degree- d polynomial with real coefficients. When $d = 1$ polynomial threshold functions are simply linear threshold functions (also known as halfspaces or LTFs), which play an important role in complexity theory, learning theory, and other fields such as voting theory. Low-degree PTFs (where d is greater than 1 but is not too large) are a natural generalization of LTFs which are also of significant interest in these fields.

Over more than twenty years much research effort in the study of Boolean functions has been devoted to different notions of the “sensitivity” of a Boolean function to small perturbations of its input, see e.g. [KKL88, BT96, BK97, Fri98, BKS99, Shi00, MO03, MOO05, OSSS05, OS07] and many other works. In this work we focus on two natural and well-studied measures of this sensitivity, the “average sensitivity” and the “noise sensitivity.” As our main results, we give the first non-trivial upper bounds on average sensitivity and noise sensitivity of low-degree PTFs. These bounds have several applications in learning theory and complexity theory as we describe later in this introduction.

We now define the notions of average and noise sensitivity in the setting of Boolean functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$. (Our paper also deals with average sensitivity and noise sensitivity of functions $f : \mathbb{R}^n \rightarrow \{-1, 1\}$ under the Gaussian distribution, but the precise definitions are more involved than in the Boolean case so we defer them until later.)

1.1 Average Sensitivity and Noise Sensitivity

The *sensitivity* of a Boolean function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ on an input $x \in \{-1, 1\}^n$, denoted $s_f(x)$, is the number of Hamming neighbors $y \in \{-1, 1\}^n$ of x (i.e. strings which differ from x in precisely one coordinate) for which $f(x) \neq f(y)$. The *average sensitivity* of f , denoted $\text{AS}(f)$, is simply $\mathbf{E}[s_f(x)]$ (where the expectation is with respect to the uniform distribution over $\{-1, 1\}^n$). An alternate definition of average sensitivity can be given in terms of the influence of individual coordinates on f . For a Boolean function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and a coordinate index $i \in [n]$, the *influence of coordinate i on f* is the probability that flipping the i -th bit of a uniform random input $x \in \{-1, 1\}^n$ causes the value of f to change, i.e. $\text{Inf}_i(f) = \mathbf{Pr}[f(x) \neq f(x^{\oplus i})]$ (where the probability is with respect to the uniform distribution over $\{-1, 1\}^n$). The sum of all n coordinate influences, $\sum_{i=1}^n \text{Inf}_i(f)$, is called the *total influence* of f ; it is easily seen to equal $\text{AS}(f)$. Bounds on average sensitivity have been of use in the structural analysis of Boolean functions (see e.g. [KKL88, Fri98, Shi00]) and in developing computationally efficient learning algorithms (see e.g. [BT96, OS07]).

The average sensitivity is a measure of how f changes when a single coordinate is perturbed. In contrast, the noise sensitivity of f measures how f changes when a random collection of coordinates are all perturbed simultaneously. More precisely, given a noise parameter

$0 \leq \epsilon \leq 1$ and a Boolean function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, the *noise sensitivity of f at noise rate ϵ* is defined to be

$$\text{NS}_\epsilon(f) = \Pr_{x,y}[f(x) \neq f(y)]$$

where x is uniform from $\{-1, 1\}^n$ and y is obtained from x by flipping each bit independently with probability ϵ . Noise sensitivity has been studied in a range of contexts including Boolean function analysis, percolation theory, and computational learning theory [BKS99, KOS04, MO03, SS, KOS08].

1.2 Main Results: Upper Bounds on Average Sensitivity and Noise Sensitivity

1.2.1 Boolean PTFs

In 1994 Gotsman and Linial [GL94] conjectured that the symmetric function slicing the middle d layers of the Boolean hypercube has the highest average sensitivity among all degree- d PTFs. Since this function has average sensitivity $\Theta(d\sqrt{n})$ for every $1 \leq d \leq \sqrt{n}$, this conjecture implies (and is nearly equivalent to) the conjecture that every degree- d PTF f over $\{-1, 1\}^n$ has $\text{AS}(f) \leq d\sqrt{n}$.

Our first main result is an upper bound on average sensitivity which makes progress toward this conjecture:

Theorem 1.1 *For any degree- d PTF f over $\{-1, 1\}^n$, we have*

$$\text{AS}(f) \leq 2^{O(d)} \cdot \log n \cdot n^{1-1/(4d+2)}.$$

Using a completely different set of techniques, we also prove a different bound which improves on Theorem 1.1 for $d \leq 4$:

Theorem 1.2 *For any degree- d PTF f over $\{-1, 1\}^n$, we have*

$$\text{AS}(f) \leq 2n^{1-1/2^d}.$$

We give a simple reduction which translates any upper bound on average sensitivity for degree- d PTFs over Boolean variables into a corresponding upper bound on noise sensitivity. Combining this reduction with Theorems 1.1 and 1.2, we establish:

Theorem 1.3 *For any degree- d PTF f over $\{-1, 1\}^n$ and any $0 \leq \epsilon \leq 1$, we have*

$$\begin{aligned} \text{NS}_\epsilon(f) &\leq 2^{O(d)} \cdot \epsilon^{1/(4d+2)} \log(1/\epsilon) \\ \text{NS}_\epsilon(f) &\leq O(\epsilon^{1/2^d}). \end{aligned}$$

1.2.2 Gaussian PTFs

Looking beyond the Boolean hypercube, there are well-studied notions of average sensitivity and noise sensitivity for Boolean-valued functions over \mathbb{R}^n , where we view \mathbb{R}^n as endowed with the standard multivariate Gaussian distribution $\mathcal{N}(0, I_n)$ [Bog98, MOO05]. Let $\text{GAS}(f)$ denote the Gaussian average sensitivity of a function $f : \mathbb{R}^n \rightarrow \{-1, 1\}$, and let $\text{GNS}_\epsilon(f)$ denote the Gaussian noise sensitivity at noise rate ϵ . (See Section 2 for precise definitions of these quantities; here we just note that these are natural analogues of their uniform-distribution Boolean hypercube counterparts defined above.) We prove an upper bound on Gaussian average sensitivity of low-degree PTFs:

Theorem 1.4 *For any degree- d PTF f over \mathbb{R}^n , we have*

$$\text{GAS}(f) \leq O(d^2 \cdot \log n \cdot n^{1-1/2d}).$$

We remark that in the case of degree- d *multilinear* PTFs it is possible to obtain a slightly stronger bound of $\text{GAS}(f) \leq O(d \cdot \log n \cdot n^{1-1/2d})$ using our approach. We also prove an upper bound on the Gaussian noise sensitivity of degree- d PTFs:

Theorem 1.5 *For any degree- d PTF f over \mathbb{R}^n and any $0 \leq \epsilon \leq 1$, we have*

$$\text{GNS}_\epsilon(f) \leq O(d \cdot \log^{1/2}(1/\epsilon) \cdot \epsilon^{1/2d}).$$

1.3 Application: agnostically learning constant-degree PTFs in polynomial time

Our bounds on noise sensitivity, together with machinery developed in [KOS04, KKMS08, KOS08], let us obtain the first efficient agnostic learning algorithms for low-degree polynomial threshold functions. In this section we state our new learning results; details are given in Section 8.

We begin by briefly reviewing the fixed-distribution agnostic learning framework that has been studied in several recent works, see e.g. [KKMS08, KOS08, BOW08, GKK08, KMV08, SSS09]. Let \mathcal{D}_X be a (fixed, known) distribution over an example space X such as the uniform distribution over $\{-1, 1\}^n$ or the standard multivariate Gaussian distribution $\mathcal{N}(0, I_n)$ over \mathbb{R}^n . Let \mathcal{C} denote a class of Boolean functions, such as the class of all degree- d PTFs. An algorithm A is said to be an *agnostic learning algorithm for \mathcal{C} under distribution \mathcal{D}_X* if it has the following property: Let \mathcal{D} be any distribution over $X \times \{-1, 1\}$ such that the marginal of \mathcal{D} over X is \mathcal{D}_X . Then if A is run on a sample of labeled examples drawn independently from \mathcal{D} , with high probability A outputs a hypothesis $h : X \rightarrow \{-1, 1\}$ such that $\Pr_{(x,y) \sim \mathcal{D}}[h(x) \neq y] \leq \text{opt} + \epsilon$, where $\text{opt} = \min_{f \in \mathcal{C}} \Pr_{(x,y) \sim \mathcal{D}}[f(x) \neq y]$. In words, A 's hypothesis is nearly as accurate as the best hypothesis in \mathcal{C} .

Kalai et al. [KKMS08] gave an L_1 polynomial regression algorithm and showed that it can be used for agnostic learning. More precisely, they showed that for a class \mathcal{C} of functions and a distribution \mathcal{D} , if every function in \mathcal{C} has a low-degree polynomial approximator (in the L_2 norm) under the marginal distribution \mathcal{D}_X , then the L_1 polynomial regression algorithm is an efficient agnostic learning algorithm for \mathcal{C} under \mathcal{D}_X . They used this L_1 polynomial regression algorithm together with the existence of low-degree polynomial approximators for halfspaces (under the uniform distribution on $\{-1, 1\}^n$ and the standard Gaussian distribution $\mathcal{N}(0, I_n)$ on \mathbb{R}^n) to obtain $n^{O(1/\epsilon^4)}$ -time agnostic learning algorithms for halfspaces under these distributions.

Using ingredients from [KOS04], one can easily convert upper bounds on Boolean noise sensitivity (such as [Theorem 1.3](#)) into results asserting the existence of low-degree L_2 -norm polynomial approximators under the uniform distribution on $\{-1, 1\}^n$. We thus obtain the following agnostic learning result (a more detailed proof is given in [Section 8](#)):

Theorem 1.6 *The class of degree- d PTFs is agnostically learnable under the uniform distribution on $\{-1, 1\}^n$ in time*

$$n^{2^{O(d^2)}(\log 1/\epsilon)^{4d+2}/\epsilon^{8d+4}}.$$

For $d \leq 4$, this bound can be improved to $n^{O(1/\epsilon^{2^{d+1}})}$.

Similarly, using ingredients from [KOS08], one can easily convert upper bounds on Gaussian noise sensitivity (such as [Theorem 1.5](#)) into results asserting the existence of low-degree L_2 -norm polynomial approximators under $\mathcal{N}(0, I_n)$. This lets us obtain

Theorem 1.7 *The class of degree- d PTFs is agnostically learnable under any n -dimensional Gaussian distribution in time $n^{(d/\epsilon)^{O(d)}}$.*

For ϵ constant, these results are the first polynomial-time agnostic learning algorithms for constant-degree PTFs.

1.4 Other applications

The results and approaches of this paper have found other recent applications beyond the agnostic learning results presented above; we describe two of these below.

Gopalan and Servedio [GS09] have combined the average sensitivity bound given by [Theorem 1.1](#) with techniques from [LMN93] to give the first sub-exponential time algorithms for learning AC^0 circuits augmented with a small (but super-constant) number of arbitrary threshold gates, i.e. gates that compute arbitrary LTFs which may have weights of any magnitude. (Previous work using different techniques [JKS02] could only handle AC^0 circuits augmented with majority gates.)

In other recent work Diakonikolas et al. [DSTW09] have refined the approach used to prove [Theorem 1.1](#) to establish a “regularity lemma” for low-degree polynomial threshold functions. Roughly speaking, this lemma says that any degree- d PTF can be decomposed into a constant number of subfunctions, almost all of which are “regular” degree- d PTFs. [DSTW09] apply this regularity lemma to extend the positive results on the existence of low-weight approximators for LTFs, proved in [Ser07], to low-degree PTFs.

Related work. Simultaneously and independently of this work, Harsha et al. [HKM09] have obtained very similar results on average sensitivity, noise sensitivity, and agnostic learning of low-degree PTFs using techniques very similar to ours.

1.5 Techniques

In this section we give a high-level overview of how [Theorems 1.1](#), [1.4](#) and [1.5](#) are proved. (As mentioned earlier, [Theorem 1.2](#) is proved using completely different techniques; see [Section 6](#).) The arguments are simpler for the Gaussian setting so we begin with these.

1.5.1 The Gaussian case

We sketch the argument for the Gaussian noise sensitivity bound [Theorem 1.5](#); the Gaussian average sensitivity bound [Theorem 1.4](#), follows along similar lines

Let $f = \text{sign}(p)$ where $p : \mathbb{R}^n \rightarrow \mathbb{R}$ is a degree- d polynomial. The Gaussian noise sensitivity $\text{GNS}_\epsilon(f)$ of f is equal to $\Pr_{x,y}[f(x) \neq f(y)]$ where x is distributed according to $\mathcal{N}(0, I_n)$ and y is an “ ϵ -perturbed” version of x (see [Section 2](#) for the precise definition). Intuitively, the event $f(x) \neq f(y)$ can only take place if either

- x lies close to the boundary of p , i.e. $|p(x)|$ is “small”, or
- $|p(x) - p(y)|$ is “large”.

We use an anti-concentration result for polynomials in Gaussian random variables, due to Carbery and Wright [CW01], to show that $|p(x)|$ is “small” only with low probability. For the second bullet, it turns out that $p(x) - p(y)$ can be expressed as a low-degree polynomial in independent Gaussian random variables, and thus we can apply tail bounds for this setting [Jan97] to show that $|p(x) - p(y)|$ is “large” only with low probability. We can thus argue that $\Pr_{x,y}[f(x) \neq f(y)]$ is low, and bound the Gaussian noise sensitivity of f . (We note that this high-level explanation glosses over some significant technical issues. In particular, since we are dealing with general degree- d PTFs which may not be multilinear, it is nontrivial to establish the conditions that allow us to apply the tail bound; see the proof of [Claim 4.1](#) in [Section 4.1](#).)

1.5.2 The Boolean case

One advantage of working over the Boolean domain $\{-1, 1\}^n$ is that without loss of generality we may consider only *multilinear* PTFs, where $f = \text{sign}(p(x))$ for p a multilinear polynomial. However, this advantage is offset by the fact that the uniform distribution on $\{-1, 1\}^n$ is less symmetric than the Gaussian distribution; for example, every degree-1 PTF under the Gaussian distribution \mathcal{N}^n is equivalent simply to $\text{sign}(x_1 - \theta)$, but this is of course not true for degree-1 PTFs over $\{-1, 1\}^n$. Our upper bound on Boolean average sensitivity uses ideas from the Gaussian setting but also requires significant additional ingredients.

An important notion in the Boolean case is that of a “regular” PTF; this is a PTF $f = \text{sign}(p)$ where every variable *in the polynomial* p has low influence. (See [Section 2](#) for a definition of the influence of a variable on a real-valued function; note that the definition from [Section 1.1](#) applies only for Boolean-valued functions.) If f is a regular PTF, then the “invariance principle” of [\[MOO05\]](#) tells us that $p(x)$ (where x is uniform from $\{-1, 1\}^n$) behaves much like $p(\mathcal{G})$ (where \mathcal{G} is drawn from $\mathcal{N}(0, I_n)$), and essentially the arguments from the Gaussian case can be used.

It remains to handle the case where f is not a regular PTF, i.e. some variable has high influence in p . To accomplish this, we generalize the notion of the “critical-index” of a halfspace (see [\[Ser07, DGJ⁺09\]](#)) to apply to PTFs. We show that a carefully chosen random restriction (one which fixes only the variables up to the critical index – very roughly speaking, only the highest-influence variables – and leaves the other ones free) has non-negligible probability of causing f to collapse down to a regular PTF. This lets us give a recursive bound on average sensitivity which ends up being not much worse than the bound that can be obtained for the regular case; see [Section 5.1](#) for a detailed explanation of the recursive argument.

1.6 Organization

Formal definitions of average sensitivity and noise sensitivity (especially in the Gaussian case), and tail bounds and anticoncentration results for low degree polynomials are presented in [Section 2](#). In [Section 3](#), we show an upper bound on the Gaussian average sensitivity of PTFs ([Theorem 1.4](#)). Upper bounds on Gaussian noise sensitivity ([Theorem 1.5](#)) are obtained in the section that follows ([Section 4](#)).

The main result of the paper – a bound on the Boolean average sensitivity ([Theorem 1.1](#)) – is proved in [Section 5](#). In [Section 6](#), an alternate bound for Boolean average sensitivity that is better for degrees $d \leq 4$ ([Theorem 1.2](#)) is shown. This is followed by a reduction from Boolean average sensitivity bounds to corresponding noise sensitivity bounds ([Theorem 7.1](#)) in [Section 7](#). We present the applications of these upper bounds to agnostic learning of PTFs in [Section 8](#). [Section 9](#) concludes by proposing a direction for future work towards the resolution of the Gotsman–Linial conjecture.

2 Definitions and Background

2.1 Basic Definitions

In this subsection we record the basic notation and definitions used throughout the paper. For $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, 2, \dots, n\}$. We write \mathcal{N} to denote the standard univariate Gaussian distribution $\mathcal{N}(0, 1)$.

For a degree- d polynomial $p : X \rightarrow \mathbb{R}$ we denote by $\|p\|_2$ its l_2 norm, $\|p\|_2 = \mathbf{E}_x[p(x)^2]^{1/2}$, where the intended distribution over $x \in \mathbb{R}^n$ (which will always be either uniform over $\{-1, 1\}^n$, or the \mathcal{N}^n distribution) will always be clear from context. We note that for multilinear p the two notions are always equal (see e.g. Proposition 3.5 of [MOO05]).

We now proceed to define the notion of influence for real-valued functions in a product probability space. Throughout this paper we consider either the uniform distribution on the hypercube $\{\pm 1\}^n$ or the standard n -dimensional Gaussian distribution in \mathbb{R}^n . However, for the sake of generality, we adopt this more general setting.

Let $(\Omega_1, \mu_1), \dots, (\Omega_n, \mu_n)$ be probability spaces and let $(\Omega = \otimes_{i=1}^n \Omega_i, \mu = \otimes_{i=1}^n \mu_i)$ denote the corresponding product space. Let $f : \Omega \rightarrow \mathbb{R}$ be any square integrable function on (Ω, μ) , i.e. $f \in L^2(\Omega, \mu)$. The influence of the i th coordinate on f [MOO05] is

$$\text{Inf}_i^\mu(f) \stackrel{\text{def}}{=} \mathbf{E}_\mu[\text{Var}_{\mu_i}[f]]$$

and the total influence of f is $\text{Inf}^\mu(f) \stackrel{\text{def}}{=} \sum_{i=1}^n \text{Inf}_i^\mu(f)$.

For a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ over the Boolean hypercube endowed with the uniform distribution, the influence of variable i on f can be expressed in terms of the Fourier coefficients of f as,

$$\text{Inf}_i(f) = \sum_{S \ni i} \widehat{f}(S)^2,$$

and as mentioned in the introduction it is easily seen that $\text{AS}(f) = \text{Inf}(f)$ for Boolean-valued functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$.

In this paper we are concerned with variable influences for functions defined over $\{-1, 1\}^n$ under the uniform distribution, and over \mathbb{R}^n under $\mathcal{N}(0, I_n)$; we shall adopt the convention that $\text{Inf}_i(f)$ denotes the former and $\text{GI}_i(f)$ the latter. We also denote by $\text{GAS}(f) = \sum_{i \in [n]} \text{GI}_i(f)$ the Gaussian average sensitivity.

Note that for a function $f : \mathbb{R}^n \rightarrow \{-1, 1\}$, the Gaussian influence $\text{GI}_i(f)$ can be equivalently written as: $\text{GI}_i(f) = 2 \mathbf{Pr}_{x, x^i}[f(x) \neq f(x^i)]$, where $x \sim \mathcal{N}^n$ and x^i is obtained by replacing the i^{th} coordinate of x by an independent random sample from \mathcal{N} .

We proceed to define the notion of noise sensitivity for Boolean-valued functions in $(\mathbb{R}^n, \mathcal{N}^n)$. For the domain $\{-1, 1\}^n$, the notion has been defined already in the introduction. (We remark that “noise sensitivity” can be defined in a much more general setting and also for real-valued functions; but such generalizations are not needed here.)

Definition 1 (Gaussian Noise Sensitivity) Given $f : \mathbb{R}^n \rightarrow \{-1, 1\}$, the “Gaussian noise sensitivity of f at noise rate $\epsilon \in [0, 1]$ ” is

$$\text{GNS}_\epsilon(f) \stackrel{\text{def}}{=} \Pr_{x,z}[f(x) \neq f(y)];$$

where $x \sim \mathcal{N}^n$ and $y \stackrel{\text{def}}{=} (1 - \epsilon)x + \sqrt{2\epsilon - \epsilon^2}z$ for an independent Gaussian noise vector $z \sim \mathcal{N}^n$.

Fourier and Hermite Analysis. We assume familiarity with the basics of Fourier analysis over the Boolean hypercube $\{-1, 1\}^n$. We will also require similar basics of Hermite analysis over the space \mathbb{R}^n equipped with the standard n -dimensional Gaussian distribution \mathcal{N}^n ; a brief review is provided in Appendix A.

2.2 Probabilistic Facts

In this subsection, we record the basic probabilistic tools we use in our proofs.

We first recall the following well-known consequence of hypercontractivity (see e.g. Lecture 16 of [O’D07] for the boolean setting and [Bog98] for the Gaussian setting):

Theorem 2.1 Let $p : X \rightarrow \mathbb{R}$ be a degree- d polynomial, where X is either $\{-1, 1\}^n$ under the uniform distribution or \mathbb{R}^n under \mathcal{N}^n , and fix $q > 2$. Then

$$\|p\|_q^2 \leq (q - 1)^d \|p\|_2^2.$$

We will need a concentration bound for low-degree polynomials over independent random signs or standard Gaussians. It can be proved (in both cases) using Markov’s inequality and hypercontractivity, see e.g. [Jan97, O’D07, AH09].

Theorem 2.2 (“degree- d Chernoff bound”) Let $p(x)$ be a degree- d polynomial. Let x be drawn either from the uniform distribution in $\{-1, 1\}^n$ or from \mathcal{N}^n . For any $t > e^d$, we have

$$\Pr_x[|p(x)| \geq t\|p\|_2] \leq \exp(-\Omega(t^{2/d})).$$

The second fact is a powerful anti-concentration bound for low-degree polynomials over Gaussian random variables. (We note that this result does not hold in the Boolean setting.)

Theorem 2.3 ([CW01]) Let $p : \mathbb{R}^n \rightarrow \mathbb{R}$ be a degree- d polynomial. Then for all $\epsilon > 0$, we have

$$\Pr_{x \sim \mathcal{N}^n}[|p(x)| \leq \epsilon\|p\|_2] \leq O(d\epsilon^{1/d}).$$

We also make essential use of a (weak) anti-concentration property of low-degree polynomials over the hypercube $\{-1, 1\}^n$:

Theorem 2.4 ([DFKO06, AH09]) *Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a degree- d polynomial with $\text{Var}[p] \equiv \sum_{0 < |S| \leq d} \widehat{p}(S)^2 = 1$ and $\mathbf{E}[p] = \widehat{p}(\emptyset) = 0$. Then we have*

$$\Pr[p(x) > 1/2^{O(d)}] > 1/2^{O(d)} \quad \text{and hence} \quad \Pr[|p(x)| \geq 1/2^{O(d)}] > 1/2^{O(d)}.$$

The following is a restatement of the invariance principle, specifically Theorem 3.19 under hypothesis **H4** in [MOO05].

Theorem 2.5 ([MOO05]) *Let $p(x) = \sum_{|S| \leq d} \widehat{p}(S) x_S$ be a degree- d multilinear polynomial with $\sum_{0 < |S| \leq d} \widehat{p}(S)^2 = 1$. Suppose each variable $i \in [n]$ has low influence $\text{Inf}_i(p) \leq \tau$, i.e. $\sum_{S \ni i} \widehat{p}(S)^2 \leq \tau$. Let x be drawn uniformly from $\{-1, 1\}^n$ and $\mathcal{G} \sim \mathcal{N}^n$. Then,*

$$\sup_{t \in \mathbb{R}} |\Pr[p(x) \leq t] - \Pr[p(\mathcal{G}) \leq t]| \leq O(d\tau^{1/(4d+1)}).$$

3 Gaussian Average Sensitivity

In this section we prove an upper bound on the Gaussian average sensitivity of degree- d PTFs (Theorem 1.4).

The following lemma, which relates the influence of a variable on f to its influence on the polynomial p , is central to the argument.

Lemma 3.1 *Let $p : \mathbb{R}^n \rightarrow \mathbb{R}$ be a degree- d polynomial over Gaussian inputs with $\text{Var}[p] = 1$ and let $f = \text{sign}(p)$. Then for each $i \in [n]$,*

$$\text{GI}_i(f) \leq O(d^2 \cdot \text{GI}_i(p)^{1/(2d)} \cdot \log(1/\text{GI}_i(p))).$$

Proof: [of Lemma 3.1] Let $p(x)$ be a degree- d polynomial where $\|p\|_2 = 1$. For notational convenience let us fix $i = 1$ and let $\tau = \text{GI}_1(p)$. We may assume that $\tau < 1/4$ since otherwise the claimed bound holds trivially. We express $p(x)$ as a univariate polynomial in x_1 as follows,

$$p(x) = p(x_1, \dots, x_n) = \sum_{i=0}^d p_i(x_2, \dots, x_n) \cdot h_i(x_1)$$

where $h_i(x_1)$ is the univariate degree- i Hermite polynomial. Note that for any multi-index $S = (S_2, \dots, S_n) \in \mathbb{N}^{n-1}$ and $0 \leq i \leq d$, we have $\widehat{p}_i(S) = \widehat{p}(S')$ where $S' = (i, S_2, \dots, S_n) \in \mathbb{N}^n$. As a result, using Parseval's identity for the Hermite basis, we have that

$$\|p\|^2 = \sum_{i=0}^d \|p_i\|^2.$$

We further have

$$\|p_i\|^2 = \sum_{S \in \mathbb{N}^{n-1}} \widehat{p}_i(S)^2 \quad \text{and} \quad \text{GI}_i(p) = \sum_{S: S_i > 0} \widehat{p}(S)^2.$$

Consequently the 2-norms of p_1, \dots, p_d are “small” and the 2-norm of p_0 is “large”:

$$\sum_{i=1}^d \|p_i\|^2 = \sum_{S: S_1 > 0} \widehat{p}(S)^2 = \text{GI}_1(p) = \tau \quad \text{and} \quad \|p_0\|^2 = 1 - \tau \geq 1/2.$$

Let $t = C^{d/2} \tau^{1/2} \log^{d/2}(1/\tau)$ and $\gamma = d^2 \cdot \tau^{1/2d} \log(1/\tau)$ where C is an absolute constant that will be defined later in [Claim 3.3](#). We can assume that $\gamma < 1/10$ since otherwise the bound of [Lemma 3.1](#) holds trivially. For these values of t and γ , the proof strategy is as follows:

- We use the “small ball probability” bound ([Theorem 2.3](#)) to argue that with high probability $p_0(g_2, \dots, g_n)$ is not too small: more precisely, $\Pr_{(g_2, \dots, g_n) \sim \mathcal{N}^{n-1}} [|p_0(g_2, \dots, g_n)| \leq td(2ed \log(1/\gamma))^{d/2}] \leq O(\gamma)$ (see [Claim 3.2](#)).
- We use the concentration bound ([Theorem 2.2](#)) to argue that with high probability each $p_i(g_2, \dots, g_n)$, $i \in [d]$, is not too large: more precisely, $\Pr_{(g_2, \dots, g_n) \sim \mathcal{N}^{n-1}} [|p_i(g_2, \dots, g_n)| \geq t] \leq O(\gamma)$ (see [Claim 3.3](#)).
- We use elementary properties of the $\mathcal{N}(0, 1)$ distribution to argue that if $|a| \geq td(2ed \log(1/\gamma))^{d/2}$ and $|b_i| \leq t$, then the function $\text{sign}(a + \sum_{i=1}^d b_i h_i(g_1))$ (a function of one $\mathcal{N}(0, 1)$ random variable g_1) is $O(\gamma)$ -close to the constant function $\text{sign}(a)$ (see [Claim 3.5](#)).
- Thus we know that with probability at least $1 - O(\gamma)$ over the choice of g_2, \dots, g_n , we have $\text{Var}_{g_1}[\text{sign}(p(g_1, \dots, g_n))] \leq O(\gamma(1 - \gamma)) \leq O(\gamma)$. For the remaining (at most) $O(\gamma)$ fraction of outcomes for g_2, \dots, g_n we always have $\text{Var}_{g_1}[\text{sign}(p(g_1, \dots, g_n))] \leq 1$, so overall we get $\text{GI}_1(\text{sign}(p)) \leq O(\gamma)$.

Thus, to complete the proof of the lemma, it suffices to prove the three aforementioned claims.

Claim 3.2 *With probability at least $1 - O(\gamma)$ over draws $(g_2, \dots, g_n) \sim \mathcal{N}^{n-1}$, the polynomial $p_0(g_2, \dots, g_n)$ has magnitude at least $td(2ed \log(1/\gamma))^{d/2}$.*

Proof: Applying [Theorem 2.3](#) to the polynomial $p_0(x_2, \dots, x_n)$ we get:

$$\Pr_{g_2, \dots, g_n} [|p_0(g_2, \dots, g_n)| \leq td(2ed \log(1/\gamma))^{d/2}] \leq O(d) \cdot \left(\frac{td(2ed \log(1/\gamma))^{d/2}}{\|p_0\|} \right)^{1/d}.$$

Recall that $\|p_0\| \geq \frac{1}{2}$, and so by our choice of t and γ it follows that the right hand side is:

$$O(d^{3/2}) \cdot O(\tau^{1/2d} \log^{1/2}(1/\tau) \cdot \log^{1/2}(1/\gamma)) = O(\gamma).$$

■

Claim 3.3 For each $i \in [d]$, the polynomial $p_i(g_2, \dots, g_n)$ has magnitude larger than t with probability at most γ/d . Therefore, the probability that any $p_i(g_2, \dots, g_n)$ has magnitude larger than t is at most γ .

Proof: First note that since $\sum_{i=1}^d \|p_i\|^2 = \tau$, certainly for each $i \in [d]$ we have $\|p_i\| \leq \sqrt{\tau}$. Therefore,

$$|\mathbf{E}[p_i]| \leq \mathbf{E}[p_i^2]^{1/2} = \|p_i\| \leq \sqrt{\tau}.$$

Let $p'_i = p_i - \mathbf{E}[p_i]$, so $\mathbf{E}[p'_i] = 0$. Applying [Theorem 2.2](#), we get:

$$\Pr_{g_2, \dots, g_n} \left[|p'_i(g_2, \dots, g_n)| > \frac{t - \sqrt{\tau}}{\|p'_i\|} \cdot \|p'_i\| \right] \leq 2 \exp \left(-\Omega \left(\left(\frac{t - \sqrt{\tau}}{\|p'_i\|} \right)^{2/d} \right) \right).$$

Given our bound on $\|p'_i\| \leq \|p_i\| \leq \sqrt{\tau}$ and choice of t , we know that the probability bound is at most $2 \exp(-\Omega(C \log(1/\tau)))$. For a sufficiently large absolute constant C this is at most $\exp(-4 \log(1/\tau)) = \tau^4 \leq \gamma/d$. To complete the proof note that if $|p'_i| \leq t - \sqrt{\tau}$ then certainly $|p_i| \leq t$. \blacksquare

We will need the following lemma in the proof of [Claim 3.5](#):

Lemma 3.4 The degree- d Hermite polynomial $h_d(x)$, $d \geq 1$, satisfies the following bound for all x :

$$|h_d(x)| \leq (ed)^{d/2} \cdot \max\{1, |x|^d\}.$$

Proof: The lemma is immediate for $d = 1$. For $d \geq 2$, we note that the polynomial $h_d(x)$ has at most d terms, each of which has coefficients of magnitude at most $\sqrt{d!} \leq d^{d-1}/\sqrt{d!}$. This directly gives $|h_d(x)| \leq (d^d/\sqrt{d!}) \cdot \max\{1, |x|^d\}$. The claimed equality follows easily from this using Stirling's approximation. \blacksquare

Claim 3.5 Suppose $|a| \geq td(2ed \log(1/\gamma))^{d/2}$, $|b_i| \leq t$ for all $i \in [d]$, and $\gamma < 1/10$. Then,

$$\Pr_{g_1 \sim \mathcal{N}(0,1)} \left[\text{sign}(a + \sum_{i=1}^d b_i h_i(g_1)) \neq \text{sign}(a) \right] \leq O(\gamma).$$

Proof: If $\text{sign}(a + \sum_{i=1}^d b_i h_i(x)) \neq \text{sign}(a)$ then it has to be the case that:

$$\left| \sum_{i=1}^d b_i h_i(x) \right| \geq |a|.$$

By [Lemma 3.4](#) we know that for all x , we have

$$\left| \sum_{i=1}^d b_i h_i(x) \right| \leq td \cdot \max_{1 \leq i \leq d} |h_i(x)| < td(ed)^{d/2} \cdot \max\{1, |x|^d\}.$$

Now if $|x|$ is at most $\sqrt{2\log(1/\gamma)}$, since $\gamma < 1/10$ we have $\sqrt{2\log(1/\gamma)} > 1$ and hence it follows that

$$\left| \sum_{i=1}^d b_i h_i(x) \right| \leq td(2ed\log(1/\gamma))^{d/2} \leq |a|.$$

In other words, if $\text{sign}(a + \sum_{i=1}^d b_i h_i(x))$ differs from $\text{sign}(a)$, it must necessarily be the case that $|x| \geq \sqrt{2\log(1/\gamma)}$. The standard tail bound on Gaussians,

$$\Pr_{g_1 \sim \mathcal{N}(0,1)}[g_1 < c] \leq \frac{1}{\sqrt{2\pi}|c|} \exp(-c^2/2) \quad \text{for } c < 0,$$

completes the proof. ■

The proof of [Lemma 3.1](#) is now complete. ■

We can now complete the proof of [Theorem 1.4](#).

Proof:[Proof of [Theorem 1.4](#)] Let us denote $\text{GI}_i(p)$ by τ_i for $i \in [n]$. Note that since p is of degree d , we have

$$\sum_{i \in [n]} \tau_i = \sum_{i \in [n]} \sum_{S \ni i} \widehat{p}(S)^2 = \sum_{|S| \leq d} |S| \cdot \widehat{p}(S)^2 \leq d. \quad (1)$$

Let $a_d(x) = d^2 x^{1/2d} \log(1/x)$. By [Lemma 3.1](#) the average sensitivity of f can be bounded as

$$\text{GAS}(f) = \sum_{i \in [n]} \text{GI}_i(f) \leq O\left(\sum_{i \in [n]} a_d(\tau_i)\right).$$

The function $a_d(x)$ is monotone increasing and concave in $[0, e^{-2d}]$. In this light, we split the summation into terms greater than e^{-2d} and the rest. Let $S = \{i | \tau_i \geq e^{-2d}\}$ and $T = [n] \setminus S$. From (1), we have $|S| \leq de^{2d}$. Observe that for $n < (27d^2)^{2d}$, [Theorem 1.4](#) holds trivially since $\text{GAS}(f) \leq n \leq 27d^2 n^{1-1/2d} \leq 27d^2 n^{1-1/2d} \log n$. Hence we may assume $n \geq (27d^2)^{2d}$, and consequently $|T|$ is at least $n/2$. Using concavity and monotonicity of a_d , we can write

$$\sum_{i \in T} a_d(\tau_i) \leq |T| \cdot a_d\left(\frac{\sum_{i \in T} \tau_i}{|T|}\right) \leq na_d\left(\frac{2d}{n}\right) \leq O(d^2 n^{1-1/2d} \log n).$$

Therefore, the average sensitivity of f is bounded by

$$\begin{aligned} \text{GAS}(f) &= \sum_{i \in S} \text{GI}_i(f) + \sum_{i \in T} \text{GI}_i(f) \\ &\leq |S| + O\left(\sum_{i \in T} a_d(\tau_i)\right) \leq de^{2d} + O(d^2 n^{1-1/2d} \log n). \end{aligned}$$

For all $d \geq 1$ we have

$$de^{2d} < e^{3d} < (3d^{1/3})^{3d} \leq (27d^2)^d \leq n^{1/2}, \quad \text{since } n \geq (27d^2)^{2d}.$$

Consequently we have $\text{GAS}(f) \leq n^{1/2} + O(d^2 n^{1-1/2d} \log n) = O(d^2 n^{1-1/2d} \log n)$, and the proof is complete. ■

4 Gaussian Noise Sensitivity

In this section we prove an upper bound on the noise sensitivity of degree- d PTFs.

Proof:[of [Theorem 1.5](#)] Let $f = \text{sign}(p)$, where $p = p(x_1, \dots, x_n)$ is a degree- d polynomial with $\mathbf{E}_{x \sim \mathcal{N}^n}[p(x)^2]^{1/2} = \|p\|_2 = 1$. Recall that $\text{GNS}_\epsilon(f)$ equals $\Pr_{x,z}[f(x) \neq f(y)]$ where $x \sim \mathcal{N}^n, z \sim \mathcal{N}^n$; x and z are independent; and $y = \alpha x + \beta z$, with $\alpha \stackrel{\text{def}}{=} 1 - \epsilon$ and $\beta = \sqrt{2\epsilon - \epsilon^2}$.

We can assume wlog that $\epsilon \leq 2^{-2d-1}$, since otherwise the theorem trivially holds.

Let us define the function

$$q(x, z) = p(x) - p(y).$$

Note that q is a degree- d polynomial over $2n$ variables.

Fix a real number $t^* > 0$. It is easy to see that $f(x) \neq f(y)$ only if at least one of the following two events hold:

$$(\text{Event } \mathcal{E}_1) \quad |p(x)| \leq t^* \quad \text{OR} \quad (\text{Event } \mathcal{E}_2) \quad |q(x, z)| \geq t^*.$$

We will upper bound the probability of these two events for a carefully chosen t^* . We will bound the probability of the event \mathcal{E}_1 using Carbery-Wright ([Theorem 2.3](#)), the probability of event \mathcal{E}_2 using the tail bound for degree- d polynomials ([Theorem 2.2](#)) and then apply a union bound.

The choice of t^* will be dictated by [Theorem 2.2](#). More precisely, to apply [Theorem 2.2](#), a bound on $\|q\|_2$ is needed. To this end, we show the following claim:

Claim 4.1 *We have $\|q\|_2 = O(d \cdot \sqrt{\epsilon})$.*

The proof of this claim is somewhat involved and is deferred to [Section 4.1](#).

Fix $t^* = \Theta(d\sqrt{\epsilon} \log^{d/2}(1/\epsilon))$. By [Theorem 2.3](#), we have:

$$\Pr_{x \sim \mathcal{N}^n}[|p(x)| \leq t^*] = O(d \cdot (t^*)^{1/d}) = O(d \cdot \epsilon^{1/(2d)} \cdot \log^{1/2}(1/\epsilon)).$$

Since both x and y are individually distributed according to \mathcal{N}^n , we have $\mathbf{E}[q(x, z)] = \mathbf{E}[p(x) - p(y)] = 0$. By [Theorem 2.2](#) and [Claim 4.1](#), we get

$$\Pr_{x,z \sim \mathcal{N}^{2n}} \left[|q(x, z)| \geq \frac{t^*}{\|q\|_2} \cdot \|q\|_2 \right] \leq 2 \exp \left(-\Omega \left(\left(\frac{t^*}{\|q\|_2} \right)^{2/d} \right) \right) \leq \epsilon.$$

Hence, by a union bound the noise sensitivity is $O(d \cdot \epsilon^{1/(2d)} \cdot \log^{1/2}(1/\epsilon))$. This completes the proof of [Theorem 1.5](#). ■

4.1 Proof of Claim 4.1

Let $p : \mathbb{R}^n \rightarrow \mathbb{R}$ be a degree- d polynomial over independent standard Gaussian random variables. Let us assume that $\|p\|_2 = 1$ and that $\epsilon \leq 2^{-2d-1}$. We will show that

$$\|q\|_2 = O(d\sqrt{\epsilon}).$$

It will be convenient for the proof to express p in an appropriate orthonormal basis. Let $p(x) = \sum_{S \in \mathcal{S}} \widehat{p}(S) H_S(x)$ be its Hermite expansion; \mathcal{S} is a family of multi-indices where each $\{H_S\}_{S \in \mathcal{S}}$ has degree at most d . By orthonormality of the basis we have that

$$\|p\|_2^2 = \sum_{S \in \mathcal{S}} \widehat{p}(S)^2.$$

Note that $q(x, z) = \sum_{S \in \mathcal{S}} \widehat{p}(S) (H_S(x) - H_S(y))$ and

$$q^2(x, z) = \sum_{S \in \mathcal{S}} \widehat{p}^2(S) (H_S(x) - H_S(y))^2 + \sum_{S, T \in \mathcal{S}, S \neq T} \widehat{p}(S) \widehat{p}(T) (H_S(x) - H_S(y))(H_T(x) - H_T(y)).$$

Let us denote the second summand in the above expression by $q'(x, z)$. We will first show that

$$\mathbf{E}_{x, z} [q'(x, z)] = 0.$$

By linearity of expectation we can write

$$\mathbf{E}_{x, z} [q'(x, z)] = \sum_{S, T \in \mathcal{S}, S \neq T} \widehat{p}(S) \widehat{p}(T) \mathbf{E}_{x, z} \left[(H_S(x) - H_S(y))(H_T(x) - H_T(y)) \right] = 0.$$

Hence, it suffices to show that for all $S \neq T$ we have

$$\mathbf{E}_{x, z} \left[(H_S(x) - H_S(y))(H_T(x) - H_T(y)) \right] = 0.$$

By orthogonality of the Hermite basis, and the fact that y is distributed according to \mathcal{N}^n , the above expression equals

$$-\mathbf{E}_{x, z} \left[H_S(x) H_T(y) \right] - \mathbf{E}_{x, z} \left[H_S(y) H_T(x) \right].$$

Thus, the desired result follows from the following lemma:

Lemma 4.2 *For all $S \neq T$ it holds*

$$\mathbf{E}_{x, z} \left[H_S(x) H_T(y) \right] = 0.$$

Proof: Since $S \neq T$, it suffices to prove the result for univariate Hermite polynomials. The result for the multivariate case then follows by independence. That is, for $x_1, z_1 \in \mathcal{N}(0, 1)$ and $s \neq t \in [d]$, we need to show that

$$\mathbf{E}_{x_1, z_1} [h_s(x_1)h_t(\alpha x_1 + \beta z_1)] = 0.$$

Since $\alpha^2 + \beta^2 = 1$, we have that the joint distribution of $(x_1, \alpha x_1 + \beta z_1)$ is identical to the joint distribution of $(\alpha x_1 + \beta z_1, x_1)$, and thus we can assume wlog that $s > t$. Since $h_t(\alpha x_1 + \beta z_1)$ is a degree- t polynomial in x_1, z_1 it can be written in the form

$$\sum_{i, j=0}^t c_{ij} h_i(x_1) h_j(z_1)$$

for some real coefficients c_{ij} . Hence, by linearity of expectation and independence, the desired expectation is

$$\sum_{i, j=0}^t c_{ij} \mathbf{E}[h_i(x_1)h_s(x_1)] \cdot \mathbf{E}[h_j(z_1)]$$

which equals 0 by orthogonality of the Hermite basis. ■

At this point, we need the following claim whose proof is deferred to the following subsection:

Claim 4.3 *Let $H_d(x)$ be a degree- d multivariate Hermite polynomial. Then*

$$\|H_d(x) - H_d(y)\|_2 = O(d \cdot \sqrt{\epsilon}).$$

Repeated applications of Claim 4.3 now yield

$$\begin{aligned} \mathbf{E}_{x, z}[q^2] &= \sum_{S \in \mathcal{S}} \hat{p}^2(S) \mathbf{E}_{x, z} [(H_S(x) - H_S(y))^2] \\ &\leq \sum_{S \in \mathcal{S}} \hat{p}^2(S) \cdot O(d^2 \cdot \epsilon) = O(d^2 \cdot \epsilon) \end{aligned}$$

concluding the proof.

4.1.1 Proof of Claim 4.3

We can assume wlog that

$$H_d(x) = \prod_{i=1}^j h_{k_i}(x_i)$$

where $j \in [d]$, $k_i \geq 1$, and $\sum_{i=1}^j k_i = d$.

For $i \in [j]$, we denote by $\Delta h_{k_i}(x_i, y_i) = h_{k_i}(y_i) - h_{k_i}(x_i)$. Then we can write

$$\begin{aligned} H_d(y) &= \prod_{i=1}^j h_{k_i}(y_i) = \prod_{i \in [j]} (h_{k_i}(x_i) + \Delta h_{k_i}(x_i, y_i)) \\ &= H_d(x) + \sum_{\emptyset \neq I \subseteq [j]} \prod_{i \in I} \Delta h_{k_i}(x_i, y_i) \cdot \prod_{i \in [j] \setminus I} h_{k_i}(x_i). \end{aligned}$$

We will need the following claim whose proof lies in the next subsection:

Claim 4.4 *Let $h_d(x)$ be a degree- d univariate Hermite polynomial. Then*

$$\|\Delta h_d(x, y)\|_2 = \|h_d(x) - h_d(y)\|_2 \leq 8\sqrt{d} \cdot \sqrt{\epsilon}.$$

The triangle inequality for norms combined with independence now yields

$$\|H_d(x) - H_d(y)\|_2 \leq \sum_{\emptyset \neq I \subseteq [j]} \prod_{i \in I} \|\Delta h_{k_i}(x_i, y_i)\|_2 \cdot \prod_{i \in [j] \setminus I} \|h_{k_i}(x_i)\|_2$$

Noting that $\|h_{k_i}(x_i)\|_2 = 1$ for all i , and $\|\Delta h_{k_i}(x_i, y_i)\|_2 \leq 8\sqrt{k_i} \cdot \sqrt{\epsilon}$ by Claim 4.4 above, we obtain

$$\begin{aligned} \|H_d(x) - H_d(y)\|_2 &\leq \sum_{\emptyset \neq I \subseteq [j]} \prod_{i \in I} \|\Delta h_{k_i}(x_i, y_i)\|_2 \\ &= \sum_{i=1}^j \|\Delta h_{k_i}(x_i, y_i)\|_2 + \sum_{I \subseteq [j], |I| \geq 2} \prod_{i \in I} \|\Delta h_{k_i}(x_i, y_i)\|_2 \\ &\leq 8 \left(\sum_{i=1}^j \sqrt{k_i} \right) \cdot \sqrt{\epsilon} + \sum_{\substack{d \\ |I|=2}} \binom{d}{|I|} (8\sqrt{d}\sqrt{\epsilon})^{|I|} \\ &\leq 8d \cdot \sqrt{\epsilon} + (1 + 8\sqrt{d}\sqrt{\epsilon})^d - (1 + 8d^{3/2}\sqrt{\epsilon}) \\ &\leq O(d \cdot \sqrt{\epsilon}) \end{aligned}$$

where the last inequality follows from the elementary bound $(1 + 8\sqrt{d}\sqrt{\epsilon})^d \leq 1 + 8d^{3/2}\sqrt{\epsilon} + O(d^3\epsilon)$ and the fact that $\epsilon \leq 2^{-2d-1}$. This completes the proof of Claim 4.3.

4.1.2 Proof of Claim 4.4

We will need a crucial lemma:

Lemma 4.5 *For all $k \in [d]$ we have*

$$\|h_k(x) - h_k(x - \epsilon x)\|_2 \leq 3k\epsilon.$$

Proof: Note that $h_k(x - \epsilon x)$ is a degree- k polynomial in x . Hence, by Taylor's theorem we deduce

$$h_k(x) - h_k(x - \epsilon x) = - \sum_{i=1}^k h_k^{(i)}(x) (-\epsilon x)^i / i!.$$

The triangle inequality for norms now yields

$$\|h_k(x) - h_k(x - \epsilon x)\|_2 \leq \sum_{i=1}^k (\epsilon^i / i!) \cdot \|h_k^{(i)}(x) x^i\|_2.$$

It thus suffices to bound the term $\|h_k^{(i)}(x) x^i\|_2$. Recalling that $(h_k^{(i)}(x))^2 = i! \binom{k}{i} (h_{k-i}(x))^2$ we have

$$\mathbf{E}_x[(h_k^{(i)}(x))^2 x^{2i}] = i! \binom{k}{i} \cdot \mathbf{E}_x[h_{k-i}^2(x) x^{2i}].$$

For $i = 1$, using the well-known relation

$$\sqrt{k} h_k(x) + \sqrt{k-1} h_{k-2}(x) = x h_{k-1}(x)$$

and the orthonormality of the h_i 's, an easy calculation gives $\mathbf{E}_x[h_{k-1}^2(x) x^2] = 2k - 1$; hence,

$$\|h_k'(x) x\|_2 \leq \sqrt{2k}.$$

For $i > 1$, by Cauchy-Schwartz we get

$$\mathbf{E}_x[h_{k-i}^2(x) x^{2i}] \leq \sqrt{\mathbf{E}_x[h_{k-i}^4(x)] \cdot \mathbf{E}_x[x^{4i}]}.$$

We now proceed to bound the RHS. By hypercontractivity, the first term can be bounded as follows

$$\|h_{k-i}\|_4^2 \leq 3^{k-i} \|h_{k-i}\|_2^2 = 3^{k-i}.$$

For the second term we recall that, for $x \sim \mathcal{N}$, we have $\mathbf{E}_x[x^{4i}] = \frac{(4i)!}{2^{2i} (2i)!}$. Using the elementary inequality $(2j)!/j! < 2^{2j} j!$ we conclude

$$\mathbf{E}_x[h_{k-i}^2(x) x^{2i}] \leq 3^{k-i} \cdot 2^i \sqrt{(2i)!} \leq 3^k \cdot (4/3)^i \cdot i! \leq 4^k i!$$

hence,

$$\|h_k^{(i)}(x) x^i\|_2 \leq \sqrt{\binom{k}{i}} 2^k i! \leq 2^{3k/2} \cdot i!.$$

Therefore,

$$\|h_k(x) - h_k(x - \epsilon x)\|_2 \leq \sqrt{2k} \cdot \epsilon + \epsilon \cdot 2^{3k/2} \cdot \sum_{j=1}^{k-1} \epsilon^j \leq 3k \cdot \epsilon$$

where we used the fact $\epsilon \leq 2^{-2d} \leq 2^{-2k}$ which yields $\sum_{j=1}^{k-1} \epsilon^j \leq \sum_{j=1}^{\infty} 2^{-2kj} \leq 2^{-2k+1}$. The proof of the lemma is now complete. \blacksquare

We now proceed to complete the proof of our claim. Let us write

$$\Delta h_d(x, y) = h_d(x) - h_d(y) = q_1(x) + q_2(x, z)$$

where $q_1(x) = h_d(x) - h_d(x - \epsilon x)$ and $q_2(x, z) = h_d(x - \epsilon x) - h_d(x - \epsilon x + \beta z)$.

By the triangle inequality for norms it holds that

$$\|\Delta h_d(x, y)\|_2 \leq \|q_1\|_2 + \|q_2\|_2$$

hence it suffices to bound each of the terms in the RHS.

By Lemma 4.5 it follows that

$$\|q_1\|_2 \leq 3d\epsilon.$$

For the second term, we will show that

$$\|q_2\|_2 \leq 5\sqrt{d} \cdot \sqrt{\epsilon}.$$

Note that this suffices to complete the proof, since by our assumption on ϵ , we have $d \cdot \epsilon < 1$, which implies that

$$\|\Delta h_d(x, y)\|_2 \leq 8\sqrt{d}\sqrt{\epsilon}$$

as desired.

Now observe that $h_d(x - \epsilon x + \beta z)$ is a degree- d polynomial in x, z . Let us denote $x' = (1 - \epsilon)x$. By Taylor's theorem we can write

$$h_d(x' + \beta z) = h_d(x') + \sum_{i=1}^d (\beta^i / i!) h_d^{(i)}(x') z^i$$

or

$$q_2(x, z) = - \sum_{i=1}^d (\beta^i / i!) h_d^{(i)}(x') z^i.$$

By triangle inequality

$$\|q_2\|_2 \leq \sum_{i=1}^d (\beta^i / i!) \|h_d^{(i)}(x') z^i\|_2$$

For the terms in the RHS by independence we get

$$\|h_d^{(i)}(x') z^i\|_2 = \|h_d^{(i)}(x')\|_2 \cdot \|z^i\|_2$$

For the second term above we have that $\|z^i\|_2 \leq 2^{i/2} \cdot \sqrt{i!}$.

Recalling that $h_d^{(i)}(x')^2 = i! \binom{d}{i} h_{d-i}^2(x')$ for the first term we have

$$\|h_d^{(i)}(x')\|_2 = \sqrt{i!} \sqrt{\binom{d}{i}} \cdot \|h_{d-i}(x')\|_2.$$

Since $x' = x - \epsilon x$ we apply Lemma 4.5 for $k = d - i$ and get

$$\|h_{d-i}(x')\|_2 \leq \|h_{d-i}(x)\|_2 + 3(d-i)\epsilon \leq 2$$

where the second inequality uses the assumption on the range of ϵ .

Therefore,

$$\begin{aligned} \|q_2\|_2 &\leq \sum_{i=1}^d 2^{i/2+1} \sqrt{\binom{d}{i}} \beta^i \\ &\leq 4\sqrt{d\epsilon} + \beta \cdot \sum_{i=2}^d 2^{i/2+1} \sqrt{\binom{d}{i}} \beta^{i-1} \\ &\leq 4\sqrt{d\epsilon} + \sqrt{2\epsilon} \cdot \sum_{i=2}^d 2^{i/2+1} \sqrt{\binom{d}{i}} 2^{-d(i-1)} \\ &\leq 5\sqrt{d\epsilon} \end{aligned}$$

This completes the proof of Claim 4.4.

5 Boolean Average Sensitivity

Let $\text{AS}(n, d)$ denote the maximum possible average sensitivity of any degree- d PTF over n Boolean variables. In this section we prove the claimed bound in Theorem 1.1:

$$\text{AS}(n, d) \leq 2^{O(d)} \cdot \log n \cdot n^{1-1/(4d+2)} \quad (2)$$

For $d = 1$ (linear threshold functions) it is well known that $\text{AS}(n, 1) = 2^{-n} \binom{n}{n/2} = \Theta(\sqrt{n})$. Also, notice that the RHS of (2) is larger than n for $d = \omega(\sqrt{\log n})$, yielding a trivial bound of $\text{AS}(n, d) \leq n$. Therefore throughout this section we shall assume d satisfies $2 \leq d \leq O(\sqrt{\log n})$.

5.1 Overview of proof

The high-level approach to proving Theorem 1.1 is a combination of a case analysis and a recursive bound.

For certain types of PTFs (“ τ -regular” PTFs; see Section 5.2 for a precise definition) we argue directly that the average sensitivity is small, using arguments similar to the Gaussian case together with the invariance principle. In particular, we show:

Claim 5.1 *Suppose $f = \text{sign}(p)$ is a τ -regular degree- d PTF where $\tau \stackrel{\text{def}}{=} n^{-(4d+1)/(4d+2)}$. Then,*

$$\text{AS}(f) \leq O(d \cdot n^{1-1/(4d+2)})$$

Claim 5.1 follows directly from Lemma 5.8, which we prove in Section 5.4.

For PTFs that are not τ -regular, we show that there is a not-too-large value of k (at most $K \stackrel{\text{def}}{=} 2d \log n / \tau$), and a collection of k variables (the variables whose influence in p are largest), such that the following holds: if we consider all 2^k subfunctions of f obtained by fixing the variables in all possible ways, a “large” (at least $1/2^{O(d)}$) fraction of the restricted functions have low average sensitivity. More precisely, we show:

Claim 5.2 *Let $K \stackrel{\text{def}}{=} 2d \log n / \tau$ where $\tau \stackrel{\text{def}}{=} n^{-(4d+1)/(4d+2)}$. Suppose $f = \text{sign}(p)$ is a degree- d PTF that is not τ -regular. Then for some $1 \leq k \leq K$, there is a set of k variables with the following property: for at least a $1/2^{O(d)}$ fraction of all 2^k assignments ρ to those k variables, we have*

$$\text{AS}(f_\rho) \leq O(d \cdot (\log n)^{1/4} \cdot n^{1-1/(4d+2)})$$

The proof of Claim 5.2 is given in Section 5.7. We do this by generalizing the “critical index” case analysis from [Ser07]. We define a notion of the τ -critical index of a degree- d polynomial; a τ -regular polynomial p is one for which the τ -critical index is 0. If the τ -critical index of p is some value $k \leq 2d \log n / \tau$, we restrict the k largest-influence variables (see Section 5.5). If the τ -critical index is larger than $2d \log n / \tau$, we restrict the $k = 2d \log n / \tau$ largest-influence variables in p (see Section 5.6).

5.1.1 Proof of main result (Theorem 1.1) assuming Claim 5.1 and Claim 5.2

Given these two claims it is not difficult to obtain the final result. In Claim 5.2, we note that the k restricted variables may each contribute at most 1 to the average sensitivity of f (recall that average sensitivity is equal to the sum of influences of each variable), and that the total influence of the remaining variables on f is equal to the expected average sensitivity of f_ρ , where the expectation is taken over all 2^k restrictions ρ . Since each function f_ρ is itself a degree- d PTF over at most n variables, we have the following recursive constraint on $\text{AS}(n, d)$:

$$\text{AS}(n, d) \leq \max\left\{ O(d \cdot n^{1-1/(4d+2)}), \max_{1 \leq k \leq K, 1/2^{O(d)} \leq \alpha \leq 1} \left\{ k + \alpha \cdot O(d \cdot (\log n)^{1/4} \cdot n^{1-1/(4d+2)}) + (1 - \alpha) \text{AS}(n, d) \right\} \right\}.$$

It is easy to see that the maximum possible value of $\text{AS}(n, d)$ subject to the above constraint is at most the maximum possible value of $\text{AS}'(n, d)$ that satisfies the following weaker constraint:

$$\text{AS}'(n, d) \leq K + \left(1 - \frac{1}{2^{O(d)}}\right) \text{AS}'(n, d)$$

which is satisfied by $\text{AS}'(n, d) \leq 2^{O(d)} \cdot \log n \cdot n^{1-1/(4d+2)}$.

5.2 Regularity and the critical index of polynomials

In [Ser07] a notion of the “critical index” of a linear form was defined and subsequently used in [OS08, DS09, DGJ⁺09]. We now give a generalization of the critical index notion for polynomials.

Definition 2 *Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ and $\tau > 0$. Assume the variables are ordered such that $\text{Inf}_i(f) \geq \text{Inf}_{i+1}(f)$ for all $i \in [n - 1]$. The τ -critical index of f is the least i such that:*

$$\frac{\text{Inf}_{i+1}(p)}{\sum_{j=i+1}^n \text{Inf}_j(p)} \leq \tau. \quad (3)$$

If (3) does not hold for any i we say that the τ -critical index of p is $+\infty$. If p has τ -critical index 0, we say that p is τ -regular.

The following simple lemma will be useful for us. It says that the total influence $\sum_{i=j+1}^n \text{Inf}_i(p)$ goes down exponentially as a function of j prior to the critical index:

Lemma 5.3 *Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ and $\tau > 0$. Let k be the τ -critical index of p . For $0 \leq j \leq k$ we have*

$$\sum_{i=j+1}^n \text{Inf}_i(p) \leq (1 - \tau)^j \cdot \text{Inf}(p).$$

Proof: The lemma trivially holds for $j = 0$. In general, since j is at most k , we have that

$$\text{Inf}_j(p) \geq \tau \cdot \sum_{i=j}^n \text{Inf}_i(p),$$

or equivalently

$$\sum_{i=j+1}^n \text{Inf}_i(p) \leq (1 - \tau) \cdot \sum_{i=j}^n \text{Inf}_i(p)$$

which yields the claimed bound. ■

Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a degree- d polynomial. We note here that the total influence of p is within a factor of d of the sum of squares of the non-constant coefficients of p :

$$\sum_{S \neq \emptyset} \widehat{p}(S)^2 \leq \sum_{i=1}^n \sum_{S \ni i} \widehat{p}(S)^2 = \sum_{i=1}^n \text{Inf}_i(p) = \sum_{S \subseteq [n]} |S| \cdot \widehat{p}(S)^2 \leq d \sum_{S \neq \emptyset} \widehat{p}(S)^2,$$

where the final inequality holds since $\widehat{p}(S) \neq 0$ only for sets $|S| \leq d$.

5.3 Restrictions and the influences of variables in polynomials

Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a degree- d polynomial. The goal of this section is to understand what happens to the influences of a variable x_ℓ , $\ell > k$, when we do a random restriction to variables x_1, \dots, x_k .

We start with the following elementary claim:

Claim 5.4 *Let ρ be a randomly chosen assignment to the variables x_1, \dots, x_k . Fix any $S \subseteq \{k+1, \dots, n\}$. Then for any polynomial $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ we have*

$$\widehat{p}_\rho(S) = \sum_{T \subseteq [k]} \widehat{p}(S \cup T) \rho_T,$$

and so we have

$$\mathbf{E}_\rho[\widehat{p}_\rho(S)^2] = \sum_{T \subseteq [k]} \widehat{p}(S \cup T)^2. \quad (4)$$

In words, all the Fourier weight on sets of the form $S \cup \{\text{some restricted variables}\}$ “collapses” down onto S in expectation. A corollary of this is that in expectation, the influence of an unrestricted variable x_ℓ does not change when we do a restriction:

Corollary 5.5 *Let ρ be a randomly chosen assignment to the variables x_1, \dots, x_k . Fix any $\ell \in \{k+1, \dots, n\}$. Then for any polynomial $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ we have*

$$\mathbf{E}_\rho[\text{Inf}_\ell(p_\rho)] = \text{Inf}_\ell(p).$$

Proof:

$$\begin{aligned} \mathbf{E}_\rho[\text{Inf}_\ell(p_\rho)] &= \mathbf{E}_\rho \left[\sum_{\ell \in S \subseteq \{k+1, \dots, n\}} \widehat{p}_\rho(S)^2 \right] \\ &= \sum_{T \subseteq [k]} \sum_{\ell \in S \subseteq \{k+1, \dots, n\}} \widehat{p}(S \cup T)^2 \\ &= \sum_{U \ni \ell} \widehat{p}(U)^2 = \text{Inf}_\ell(p). \end{aligned}$$

■

5.3.1 Influences of low-degree polynomials behave nicely under restrictions

In this subsection we prove the following lemma: For a low-degree polynomial, a random restriction with very high probability does not cause any variable’s influence to increase by more than a polylog(n) factor.

Lemma 5.6 *Let $p(x_1, \dots, x_n)$ be a degree- d polynomial. Let ρ be a randomly chosen assignment to the variables x_1, \dots, x_k . Fix any $t > e^{2d}$ and any $\ell \in [k+1, n]$. With probability at least $1 - \exp(-\Omega(t^{1/d}))$ over the choice of ρ , we have*

$$\text{Inf}_\ell(p_\rho) \leq t \cdot 3^d \text{Inf}_\ell(p).$$

In particular, for $t = \log^d n$, we have that with probability at least $1 - n^{-\omega(1)}$, every variable $\ell \in [k+1, n]$ has $\text{Inf}_\ell(p_\rho) \leq (3 \log n)^d \cdot \text{Inf}_\ell(p)$.

Proof: Since $\text{Inf}_\ell(p_\rho)$ is a degree- $2d$ polynomial in ρ , [Lemma 5.6](#) follows as an immediate consequence of [Theorem 2.2](#) if we can upper bound $\|\text{Inf}_\ell(p_\rho)\|_2$. We use the bound in [Lemma 5.7](#), stated and proven below. ■

Lemma 5.7 *Let $p(x_1, \dots, x_n)$ be a degree- d polynomial. Let ρ be a randomly chosen assignment to the variables x_1, \dots, x_k , and let $\ell \in [k+1, n]$. Then $\text{Inf}_\ell(p_\rho)$ is a degree- $2d$ polynomial in variables ρ_1, \dots, ρ_k , and*

$$\|\text{Inf}_\ell(p_\rho)\|_2 \leq 3^d \cdot \text{Inf}_\ell(p).$$

Proof: The triangle inequality tells us that we may bound the 2-norm of each squared-coefficient separately:

$$\|\text{Inf}_\ell(p_\rho)\|_2 \leq \sum_{\ell \in S \subseteq [k+1, n]} \|\widehat{p}_\rho(S)^2\|_2.$$

Since $\widehat{p}_\rho(S)$ is a degree- d polynomial, Bonami-Beckner (i.e., $(4, 2)$ -hypercontractivity) tells us that

$$\|\widehat{p}_\rho(S)^2\|_2 = \|\widehat{p}_\rho(S)\|_4^2 \leq 3^d \|\widehat{p}_\rho(S)\|_2^2,$$

hence

$$\|\text{Inf}_\ell(p_\rho)\|_2 \leq 3^d \sum_{\ell \in S \subseteq [k+1, n]} \|\widehat{p}_\rho(S)\|_2^2 = 3^d \cdot \text{Inf}_\ell(p)$$

where the last equality is by [Corollary 5.5](#). ■

5.4 The regular case

In this section we prove that regular degree- d PTF's have low average sensitivity. In particular, we show:

Lemma 5.8 *Fix $\tau = n^{-\Theta(1)}$. Let f be a τ -regular degree- d PTF. Then,*

$$\text{AS}(f) \leq O(d \cdot n \cdot \tau^{1/(4d+1)})$$

Claim 5.1 follows directly from the above lemma, recalling we choose $\tau \stackrel{\text{def}}{=} n^{-(4d+1)/(4d+2)}$. However, the lemma will also be useful in the “small critical index” case for a slightly larger regularity parameter τ .

Proof: Let $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a degree- d PTF, i.e. $f = \text{sign}(p)$ where p is τ -regular. We may assume that p is normalized such that $\sum_{0 < |S| \leq d} \widehat{p}(S)^2 = 1$.

First we note that flipping the i -th bit of an input $x \in \{-1, 1\}^n$ changes the value of p by the magnitude of its partial derivative with respect to i :

$$2D_i p(x) = 2 \sum_{S \ni i} \widehat{p}(S) x_{S-\{i\}}$$

It follows that:

$$\text{Inf}_i(f) \leq \mathbf{Pr}_{x \in \{-1, 1\}^n} [|p(x)| \leq |2D_i p(x)|]$$

Therefore, bounding from above the influence of variable i in f can be done by showing the following:

1. $p(x)$ has small magnitude, $|p(x)| \leq t$ for some threshold t , with small probability.
2. $2D_i p(x)$ has large magnitude, $|2D_i p(x)| \geq t$, with small probability.

We bound the probability of the first event using the anti-concentration property of regular low-degree polynomials, as implied by the invariance principle along with [Theorem 2.3](#). For the second event we use the tail bound for degree- d polynomials ([Theorem 2.2](#)).

We will take our threshold t to be $t \stackrel{\text{def}}{=} \tau^{1/4}$, where τ is the regularity parameter of p .

5.4.1 Bounding the probability of the first event

By the τ -regularity of p , for all $i \in [n]$ we have $\text{Inf}_i(p) \leq \tau \cdot \text{Inf}(p) \leq d \cdot \tau$ where the last inequality follows by the assumed normalization. With this bound, the invariance principle ([Theorem 2.5](#)) tells us that $\mathbf{Pr}_{x \in \{-1, 1\}^n} [|p(x)| \leq \tau^{1/4}]$ differs from $\mathbf{Pr}_{\mathcal{G}_1, \dots, \mathcal{G}_n} [|p(\mathcal{G})| \leq \tau^{1/4}]$ by at most $O(d \cdot (d\tau)^{1/(4d+1)}) = O(d \cdot \tau^{1/(4d+1)})$. Applying the anti-concentration bound of Carbery and Wright for polynomials in Gaussian random variables ([Theorem 2.3](#)), we get:

$$\begin{aligned} \mathbf{Pr}_x [|p(x)| \leq \tau^{1/4}] &\leq \mathbf{Pr}_{\mathcal{G}_1, \dots, \mathcal{G}_n} [|p(\mathcal{G})| \leq \tau^{1/4}] + O(d\tau^{1/(4d+1)}) \\ &\leq O(d \cdot \tau^{1/4d}) + O(d \cdot \tau^{1/(4d+1)}) \\ &= O(d \cdot \tau^{1/(4d+1)}). \end{aligned}$$

5.4.2 Bounding the probability of the second event

Next we consider $\Pr_x[|2D_i p(x)| \geq \tau^{1/4}]$. Note that $2D_i p$ is a degree- $(d-1)$ polynomial whose l_2 norm is small:

$$\|2D_i p\| = 2 \sqrt{\sum_{S \ni i} \widehat{p}(S)^2} = 2 \sqrt{\text{Inf}_i(p)} \leq 2\sqrt{d \cdot \tau}.$$

By (Theorem 2.2), we get that

$$\begin{aligned} \Pr_x[|2D_i p(x)| \geq \tau^{1/4}] &\leq \Pr_x[|2D_i p(x)| \geq \tau^{-1/4}/(2\sqrt{d}) \cdot \|2D_i p\|] \\ &\leq \exp(-\tau^{-1/(2d)}/(2\sqrt{d})^{2/d}) = \exp(-\Theta(1) \cdot \tau^{-1/(2d)}) \ll O(d \cdot \tau^{1/(4d+1)}). \end{aligned}$$

(In the second inequality, we were able to apply the concentration bound since, by our assumptions on d and τ , we indeed have that $\tau^{-1/4}/(2\sqrt{d}) > e^d$.)

Hence, we have shown that:

$$\begin{aligned} \text{Inf}_i(f) &\leq \Pr_{x \in \{-1,1\}^n}[|p(x)| \leq |2D_i p(x)|] \\ &\leq \Pr_x[|p(x)| \leq \tau^{1/4}] + \Pr_x[|2D_i p(x)| \geq \tau^{1/4}] \\ &= O(d \cdot \tau^{1/(4d+1)}). \end{aligned}$$

Since this holds for all indices $i \in [n]$, we have the following bound on the average sensitivity of $f = \text{sign}(p)$:

$$\text{AS}(f) \leq O(d \cdot n \cdot \tau^{1/(4d+1)}).$$

■

5.5 The small critical index case

Let $f = \text{sign}(p)$ be such that the τ -critical index of p is some value k between 1 and $K = 2d \log n / \tau$. By definition, the sequence of influences $\text{Inf}_{k+1}(p), \dots, \text{Inf}_n(p)$ is τ -regular. We essentially reduce this case to the regular case for a regularity parameter τ' somewhat larger than τ .

Consider a random restriction ρ of all the variables up to the critical index. We will show the following:

Lemma 5.9 *For a $1/2^{O(d)}$ fraction of restrictions ρ , the sequence of influences $\text{Inf}_{k+1}(p_\rho), \dots, \text{Inf}_n(p_\rho)$ is τ' -regular, where $\tau' \stackrel{\text{def}}{=} (3 \log n)^d \cdot \tau$.*

By our choice of $\tau = n^{-(4d+1)/(4d+2)}$, we have that $\tau' = n^{-\Theta(1)}$, and so we may apply Lemma 5.8 to these restrictions to conclude that the associated PTFs have average sensitivity at most $O(d \cdot n \cdot (\tau')^{1/(4d+1)})$.

Proof:

Since the sequence of influences $\text{Inf}_{k+1}(p), \dots, \text{Inf}_n(p)$ is τ -regular, we have

$$\frac{\text{Inf}_i(p)}{\sum_{j=k+1}^n \text{Inf}_j(p)} \leq \tau$$

for all $i \in [k+1, n]$.

We want to prove that for a $1/2^{O(d)}$ fraction of all 2^k restrictions ρ to x_1, \dots, x_k we have

$$\frac{\text{Inf}_i(p_\rho)}{\sum_{j=k+1}^n \text{Inf}_j(p_\rho)} \leq \tau'$$

for all $i \in [k+1, n]$.

To do this we proceed as follows: [Lemma 5.6](#) implies that, with very high probability over the random restrictions, we have $\text{Inf}_i(p_\rho) \leq (3 \log n)^d \cdot \text{Inf}_i(p)$, for all $i \in [k+1, n]$. We need to show that for a $1/2^{O(d)}$ fraction of all restrictions the denominator of the fraction above is at least $\sum_{j=k+1}^n \text{Inf}_j(p)$ (its expected value). The lemma then follows by a union bound.

We consider the degree- $2d$ polynomial $A(\rho_1, \dots, \rho_k) \stackrel{\text{def}}{=} \sum_{j=k+1}^n \text{Inf}_j(p_\rho)$ in variables ρ_1, \dots, ρ_k . The expected value of A is $\mathbf{E}_\rho[A] = \sum_{j=k+1}^n \text{Inf}_j(p) = \widehat{A}(\emptyset)$. We apply the [Theorem 2.4](#) for $B = A - \widehat{A}(\emptyset)$. We thus get $\Pr_\rho[B > 0] > 1/2^{O(d)}$. We thus get $\Pr_\rho[A > \mathbf{E}_\rho[A]] > 1/2^{O(d)}$ and we are done. \blacksquare

5.6 The large critical index case

Finally we consider PTFs $f = \text{sign}(p)$ with τ -critical index greater than $K = 2d \log n / \tau$. Let ρ be a restriction of the first K variables $\mathcal{H} = \{1, \dots, K\}$; we call these the ‘‘head’’ variables. We will show the following:

Lemma 5.10 *For a $1/2^{O(d)}$ fraction of restrictions ρ , the function $\text{sign}(p_\rho(x))$ is a constant function.*

Proof: By [Lemma 5.3](#), the surviving variables x_{K+1}, \dots, x_n have very small total influence in p :

$$\sum_{i=K+1}^n \text{Inf}_i(p) = \sum_{i=K+1}^n \sum_{S \ni i} \widehat{p}(S)^2 \leq (1 - \tau)^K \cdot \text{Inf}(p) \leq d/n^{2d}. \quad (5)$$

Therefore, if we let p' be the truncation of p comprising only the monomials with all variables in \mathcal{H} ,

$$p'(x_1, \dots, x_k) = \sum_{S \subset \mathcal{H}} \widehat{p}(S) x_S$$

we know that almost all of the original Fourier weight of p is on the coefficients of p' :

$$1 \geq \sum_{\substack{S \subseteq \mathcal{H} \\ |S| > 0}} \widehat{p}(S)^2 \geq 1 - \sum_{i=K+1}^n \text{Inf}_i(p) \geq 1 - d/n^{2d}$$

We now apply [Theorem 2.4](#) to p' ¹ and get:

$$\Pr_{x \in \{-1,1\}^K} [|p'(x)| \geq 1/2^{O(d)}] \geq 1/2^{O(d)}.$$

In words, for a $1/2^{O(d)}$ fraction of all restrictions ρ to x_1, \dots, x_K , the value $p'(\rho)$ has magnitude at least $1/2^{O(d)}$.

For any such restriction, if the function $f_\rho(x) = \text{sign}(p_\rho(x_{K+1}, \dots, x_n))$ is not a constant function it must necessarily be the case that:

$$\sum_{0 < |S| \subseteq \{x_{K+1}, \dots, x_n\}} |\widehat{p}_\rho(S)| \geq 1/2^{O(d)}$$

As noted in [\(5\)](#), each tail variable $\ell > K$ has very small influence in p :

$$\text{Inf}_\ell(p) \leq \sum_{i=K+1}^n \text{Inf}_i(p) = d/n^{2d}$$

Applying [Lemma 5.6](#), we get that for the overwhelming majority of the $1/2^{O(d)}$ fraction of restrictions mentioned above, the influence of ℓ in p_ρ is not much larger than the influence of ℓ in p :

$$\text{Inf}_\ell(p_\rho) \leq (3 \log n)^d \cdot \text{Inf}_\ell(p) \leq d \cdot (3 \log n)^d / n^{2d} \tag{6}$$

Using Cauchy-Schwarz, we have

$$\begin{aligned} \sum_{S \ni \ell, S \subseteq \{x_{K+1}, \dots, x_n\}} |\widehat{p}_\rho(S)| &\leq n^{d/2} \cdot \sqrt{\sum_{S \ni \ell, S \subseteq \{x_{K+1}, \dots, x_n\}} \widehat{p}_\rho(S)^2} \\ &= n^{d/2} \sqrt{\text{Inf}_\ell(p_\rho)} \\ &\leq n^{-\Omega(1)} \end{aligned}$$

where we have used [\(6\)](#) (and our upper bound on d). From this we easily get that

¹after a very slight rescaling so the non-constant Fourier coefficients of p' have sum of squares equal to 1; this does not affect the bound we get because of the big-O.

$$\sum_{0 < |S| \subseteq \{x_{K+1}, \dots, x_n\}} |\widehat{p}_\rho(S)| \leq n^{-\Omega(1)} \ll 1/2^{O(d)}$$

We have established that for a $1/2^{O(d)}$ fraction of all restrictions to x_1, \dots, x_K , the function $f_\rho = \text{sign}(p_\rho)$ is a constant function, and the lemma is proved. \blacksquare

5.7 Proof of Claim 5.2

If f is a degree- d PTF that is not τ -regular, then its τ -critical index is either in the range $\{1, \dots, K\}$ or it is greater than K .

In the first case (small critical index case), as shown in Section 5.5, we have that for a $1/2^{O(d)}$ fraction of restrictions ρ to variables x_1, \dots, x_k , the total influence of $f_\rho = \text{sign}(p_\rho)$ is at most

$$O(d \cdot n \cdot (\tau')^{1/(4d+1)}) = O(d \cdot (\log n)^{1/4} \cdot n^{1-1/(4d+2)}),$$

so the conclusion of Claim 5.2 holds in this case.

In the second case (large critical index case), as shown in Section 5.6, for a $1/2^{O(d)}$ fraction of restrictions ρ to x_1, \dots, x_K the function f_ρ is constant and hence has zero influence, so the conclusion of Claim 5.2 certainly holds in this case as well. \blacksquare

6 A Fourier-Analytic Bound on Boolean Average Sensitivity

In this section, we present a simple proof of the following upper bound on the average sensitivity of a degree- d PTF (Theorem 1.2):

$$\text{AS}(n, d) \leq 2n^{1-1/2^d}.$$

We recall here the definition of the formal derivative of a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$.

$$D_i p(x) = \sum_{S \ni i} \widehat{p}_S x_{S-\{i\}}.$$

It is easy to see that,

$$D_i p(x) = \frac{1}{2} x_i [p(x) - p(x^{\oplus i})] = \frac{1}{2} \left(\frac{p(x) - p(x^{\oplus i})}{x_i} \right) \quad (7)$$

where “ $x^{\oplus i}$ ” means “ x with the i -th bit flipped.”

For a Boolean function f , we have $D_i f(x) = \pm 1$ iff flipping the i th bit flips f ; otherwise $D_i f(x) = 0$. So we have

$$\text{Inf}_i(f) = \mathbf{E}[|D_i f(x)|].$$

Lemma 6.1 Fix $i \neq j \in [n]$. Let $f, g : \{-1, 1\}^n \rightarrow \mathbb{R}$ be functions such that f is independent of the i^{th} bit x_i and g is independent of the j^{th} bit x_j . Then

$$\mathbf{E}_x[x_i x_j f(x) g(x)] \leq \frac{\text{Inf}_i(g) + \text{Inf}_j(f)}{2}.$$

Proof: First, note that the influence of i^{th} coordinate on a function f can be written as:

$$\text{Inf}_i(f) = \mathbf{E}_{x_{-i}}[\text{Var}_{x_i}[f(x)]] = \mathbf{E}_x \left[\left(\frac{|f(x^{\oplus i}) - f(x)|}{2} \right)^2 \right] = \mathbf{E}_{x_{-i}} [|\mathbf{E}_{x_i}[x_i f(x)]|^2] \quad (8)$$

As f is independent of x_i and g is independent of x_j , we can write,

$$\begin{aligned} \mathbf{E}_x[x_i x_j f(x) g(x)] &= \mathbf{E}_{x_{-\{i,j\}}} \mathbf{E}_{x_i, x_j} [x_i x_j f(x) g(x)] \\ &= \mathbf{E}_{x_{-\{i,j\}}} [\mathbf{E}_{x_i}[x_i g(x)] \mathbf{E}_{x_j}[x_j f(x)]] \\ &\leq \mathbf{E}_{x_{-\{i,j\}}} \left[\frac{1}{2} |\mathbf{E}_{x_i}[x_i g(x)]|^2 + \frac{1}{2} |\mathbf{E}_{x_j}[x_j f(x)]|^2 \right] \quad (\text{using } ab \leq \frac{1}{2}(a^2 + b^2)) \\ &\leq \frac{\text{Inf}_j(f) + \text{Inf}_i(g)}{2} \quad (\text{using Equation 8}) \end{aligned}$$

■

[Theorem 1.2](#) is shown using an inductive argument over the degree d . Central to this inductive argument is the following lemma relating the influences of a degree- d PTF $\text{sign}(p(x))$ to the degree- $(d-1)$ PTFs obtained by taking formal derivatives of p .

Lemma 6.2 For a PTF $f = \text{sign}(p(x))$ on n variables and $i \in [n]$, $\text{Inf}_i(f) = \mathbf{E}[f(x) x_i \text{sign}(D_i p(x))]$.

The following simple claim will be useful in the proof of the above lemma.

Claim 6.3 For two real numbers a, b , if $\text{sign}(a) \neq \text{sign}(b)$ then

$$\text{sign}(\text{sign}(a) - \text{sign}(b)) = \text{sign}(a - b)$$

Proof: If $\text{sign}(a) = 1$ and $\text{sign}(b) = -1$ ($a \geq 0, b < 0$) then $a - b \geq 0$. Hence in this case, $\text{sign}(a - b) = 1 = \text{sign}(1 - (-1)) = \text{sign}(\text{sign}(a) - \text{sign}(b))$. On the other hand, if $\text{sign}(a) = -1$ and $\text{sign}(b) = 1$, then $\text{sign}(a - b) = -1 = \text{sign}((-1) - 1) = \text{sign}(\text{sign}(a) - \text{sign}(b))$. ■

Proof:[of [Lemma 6.2](#)] The influence of the i^{th} coordinate is given by,

$$\begin{aligned} \text{Inf}_i(f) &= \mathbf{E} \left[\frac{1}{2} |f(x) - f(x^{\oplus i})| \right] \\ &= \mathbf{E} \left[\frac{1}{2} (f(x) - f(x^{\oplus i})) \text{sign}(f(x) - f(x^{\oplus i})) \right] \end{aligned} \quad (9)$$

Consider an x for which $f(x) \neq f(x^{\oplus i})$. In this case, we can use [Claim 6.3](#) to conclude:

$$\begin{aligned} \text{sign}(f(x) - f(x^{\oplus i})) &= \text{sign}(p(x) - p(x^{\oplus i})) , \\ &= \text{sign}(2x_i D_i p(x)) = x_i \text{sign}(D_i p(x)) . \quad (\text{using (7)}) \end{aligned}$$

Hence for an x with $f(x) \neq f(x^{\oplus i})$,

$$(f(x) - f(x^{\oplus i})) \text{sign}(f(x) - f(x^{\oplus i})) = (f(x) - f(x^{\oplus i})) x_i \text{sign}(D_i p(x)) .$$

On the other hand, if $f(x) = f(x^{\oplus i})$ then the above equation continues holds since both the sides evaluate to 0. Substituting this equality into [Equation 9](#) yields,

$$\text{Inf}_i(f) = \frac{1}{2} \mathbf{E} [f(x) x_i \text{sign}(D_i p(x))] - \frac{1}{2} \mathbf{E} [f(x^{\oplus i}) x_i \text{sign}(D_i p(x))] .$$

Notice that the i^{th} coordinate $(x^{\oplus i})_i$ of $x^{\oplus i}$ is given by $-x_i$. Since $D_i p$ is independent of the i^{th} coordinate x_i , we have $D_i p(x) = D_i p(x^{\oplus i})$. Rewriting the above equation, we get

$$\begin{aligned} \text{Inf}_i(f) &= \frac{1}{2} \mathbf{E} [f(x) x_i \text{sign}(D_i p(x))] + \frac{1}{2} \mathbf{E} [f(x^{\oplus i}) (x^{\oplus i})_i \text{sign}(D_i p(x^{\oplus i}))] , \\ &= \mathbf{E} [f(x) x_i \text{sign}(D_i p(x))] \quad ((x^{\oplus i}) \text{ is also uniformly distributed}) \end{aligned}$$

■

Theorem 6.4 *Let $\text{AS}(n, d)$ denote the max possible average sensitivity of any degree- d PTF on n variables. Then we have*

$$\text{AS}(n, d) \leq \sqrt{n + n \cdot \text{AS}(n, d - 1)} .$$

Proof:

$$\begin{aligned} \text{Inf}(f) &= \sum_i \text{Inf}_i(f) \\ &= \sum_i \mathbf{E} [f(x) x_i \text{sign}(D_i p(x))] \quad (\text{by Lemma 6.2}) \\ &= \mathbf{E} [f(x) \sum_i x_i \text{sign}(D_i p(x))] \\ &\leq \sqrt{\mathbf{E} [f(x)^2]} \cdot \sqrt{\mathbf{E} [(\sum_i x_i \text{sign}(D_i p(x)))^2]} \end{aligned} \tag{10}$$

$$= 1 \cdot \sqrt{\mathbf{E} [\sum_{i,j} x_i x_j \text{sign}(D_i p(x)) \text{sign}(D_j p(x))]} \tag{11}$$

$$\leq \sqrt{\mathbf{E} [\sum_i x_i^2 \text{sign}(D_i p(x))^2] + \sum_{i \neq j} \text{Inf}_i(\text{sign}(D_j p(x)))} \tag{12}$$

$$= \sqrt{n + \sum_{i \neq j} \text{Inf}_i(\text{sign}(D_j p(x)))} . \tag{13}$$

Here (10) is the Cauchy-Schwarz inequality, (11) is expanding the square. Step (12) uses Lemma 6.1 which we may apply since $D_i p(x)$ does not depend on x_i .

Observe that for any fixed j' , we have $D_{j'} p(x)$ is a degree- $(d-1)$ polynomial and $\text{sign}(D_{j'} p(x))$ is a degree- $(d-1)$ PTF. Hence, by definition we have,

$$\sum_{i \neq j'} \text{Inf}(\text{sign}(D_i p(x))) \leq \text{AS}(n, d-1),$$

for all $j' \in [n]$. Therefore the quantity $\sum_{i \neq j} \text{Inf}(\text{sign}(D_i p(x))) \leq n \cdot \text{AS}(n, d-1)$, finishing the proof. \blacksquare

The bound on average sensitivity (Theorem 1.2) follows immediately from the above recursive relation.

Proof:[of Theorem 1.2] Clearly, we have $\text{AS}(n, 0) = 0$. For $d = 1$, Theorem 6.4 yields $\text{AS}(n, 1) \leq \sqrt{n}$. Now suppose $\text{AS}(n, d) = 2n^{1-1/2^d}$ for $d \geq 1$, then by Theorem 6.4,

$$\text{AS}(n, d+1) \leq \sqrt{n + n \cdot \text{AS}(n, d)} \leq \sqrt{4n^{2-1/2^d}} = 2n^{1-1/2^{d+1}},$$

finishing the proof. \blacksquare

7 Boolean average sensitivity vs noise sensitivity

Our results on Boolean noise sensitivity are obtained via the following simple reduction which translates any upper bound on average sensitivity for degree- d PTFs over Boolean variables into a corresponding upper bound on noise sensitivity. This theorem is inspired by the proof of noise sensitivity of halfspaces by Peres [Per04].

Theorem 7.1 *Let $\text{NS}(\epsilon, d)$ denote the maximum noise sensitivity of a degree d -PTF at a noise rate of ϵ . For all $0 \leq \epsilon \leq 1$ if $m = \lfloor \frac{1}{\epsilon} \rfloor$ then,*

$$\text{NS}(\epsilon, d) \leq \frac{1}{m} \text{AS}(m, d).$$

Theorem 1.3 follows immediately from this reduction along with our bounds on Boolean average sensitivity (Theorems 1.1 and 1.2), so it remains for us to prove Theorem 7.1.

7.1 Proof of Theorem 7.1

Let $f(x) = \text{sign}(p(x))$ be a degree d -PTF. Let us denote $\delta = \frac{1}{m}$. As $\delta \geq \epsilon$, by the monotonicity of noise sensitivity we have $\text{NS}_\epsilon(f) \leq \text{NS}_\delta(f)$. In the following, we will show that $\text{NS}_\delta(f) \leq \frac{1}{m} \text{AS}(m, d)$ which implies the intended result. Recall that $\text{NS}_\delta(f)$ is defined as

$$\text{NS}_\delta(f) = \Pr_{x \sim_\delta y} [f(x) \neq f(y)],$$

where $x \sim_\delta y$ denotes that y is generated by flipping each bit of x independently with probability δ . An alternate way to generate y from x is as follows:

- Sample $r \in \{1, \dots, m\}$ uniformly at random.
- Partition the bits of x into $m = \frac{1}{\delta}$ sets S_1, S_2, \dots, S_m by independently assigning each bit to a uniformly random set. Formally, a partition α is specified by a function $\alpha : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$ mapping bit locations to their partition numbers, i.e., $i \in S_{\alpha(i)}$. A uniformly random partition is picked by sampling $\alpha(i)$ for each $i \in \{1, \dots, n\}$ uniformly at random from $\{1, \dots, m\}$.
- Flip the bits of x contained in the set S_r to obtain y .

Each bit of x belongs to the set S_r independently with probability $\frac{1}{m} = \delta$. Therefore, the vector y generated by the above procedure can equivalently be generated by flipping each bit of x with probability δ .

Inspired by the above procedure, we now define an alternate equivalent procedure to generate the pair $x \sim_\delta y$.

- Sample $a \in \{-1, 1\}^n$ uniformly at random.
- Sample a uniformly random partition $\alpha : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$ of the bits of a .
- Sample $z \in \{-1, 1\}^m$ uniformly at random.
- Sample $r \in \{1, \dots, m\}$ uniformly at random. Let $\tilde{z} = z^{\oplus r}$ and

$$x_i = a_i z_{\alpha(i)} \qquad y_i = a_i \tilde{z}_{\alpha(i)}$$

Notice that x is uniformly distributed in $\{-1, 1\}^n$, since both a and z are uniformly distributed in $\{-1, 1\}^n$ and $\{-1, 1\}^m$ respectively. Furthermore, $\tilde{z}_i = z_i$ for all $i \neq r$ and $\tilde{z}_r = -z_r$. Therefore, y is obtained by flipping the bits of x in the coordinates belonging to the r^{th} partition. As the partition α is generated uniformly at random, this amounts to flipping each bit of x with probability exactly $\frac{1}{m} = \delta$.

The noise sensitivity of f can be rewritten as,

$$\text{NS}_\delta(f) = \Pr_{a,\alpha,z,r} [f(x) \neq f(y)]$$

For a fixed choice of a and α , $f(x)$ is a function of z . In this light, let us define the function $f_{a,\alpha} : \{-1, 1\}^m \rightarrow \{-1, 1\}$ for each a, α as $f_{a,\alpha}(z) = f(x)$. Returning to the expression for noise sensitivity we get:

$$\begin{aligned} \text{NS}_\delta(f) &= \Pr_{a,\alpha,z,r} [f_{a,\alpha}(z) \neq f_{a,\alpha}(\tilde{z})] \\ &= \mathbf{E}_{a,\alpha,z,r} [\mathbf{1}[f_{a,\alpha}(z) \neq f_{a,\alpha}(z^{\oplus r})]] \\ &= \mathbf{E}_{a,\alpha,z} \left[\frac{1}{m} \sum_{r=1}^m \mathbf{1}[f_{a,\alpha}(z) \neq f_{a,\alpha}(z^{\oplus r})] \right] \\ &= \mathbf{E}_{a,\alpha} \left[\frac{1}{m} \sum_{r=1}^m \mathbf{E}_z [\mathbf{1}[f_{a,\alpha}(z) \neq f_{a,\alpha}(z^{\oplus r})]] \right]. \end{aligned}$$

In the above calculation, the notation $\mathbf{1}[E]$ refers to the indicator function of the event E . Recall that, by definition of influences,

$$\text{Inf}_r(f_{a,\alpha}) = \mathbf{E}_z [\mathbf{1}[f_{a,\alpha}(z) \neq f_{a,\alpha}(z^{\oplus r})]] ,$$

for all r . Thus, we can rewrite the noise sensitivity of f as

$$\text{NS}_\delta(f) = \mathbf{E}_{a,\alpha} \left[\frac{1}{m} \sum_{r=1}^m \text{Inf}_r(f_{a,\alpha}) \right] = \frac{1}{m} \mathbf{E}_{a,\alpha} [\text{Inf}(f_{a,\alpha})] . \quad (14)$$

We claim that $f_{a,\alpha}$ is a degree d -PTF in m variables. To see this observe that

$$f_{a,\alpha}(z) = \text{sign}(p(x_1, \dots, x_n)) = \text{sign}(p(a_1 z_{\alpha(1)}, \dots, a_n z_{\alpha(n)})) ,$$

which for a fixed choice of a, α is a degree d -PTF in z . Consequently, by definition of $\text{AS}(m, d)$ we have $\text{Inf}(f_{a,\alpha}) \leq \text{AS}(m, d)$ for all a and α . Using this in (14), the result follows.

8 Application to Agnostic Learning

In this section, we outline the applications of the noise sensitivity bounds presented in this work to agnostic learning of PTFs. Specifically, we will present the proofs of [Theorem 1.6](#) and [Theorem 1.7](#). To begin with, we recall the main theorem of [\[KKMS08\]](#) about the L_1 polynomial regression algorithm:

Theorem 8.1 *Let \mathcal{D} be a distribution over $X \times \{-1, 1\}$ (where $X \subseteq \mathbb{R}^n$) which has marginal \mathcal{D}_X over X . Let \mathcal{C} be a class of Boolean-valued functions over X such that for every $f \in \mathcal{C}$, there is a degree- d polynomial $p(x_1, \dots, x_n)$ such that $\mathbf{E}_{x \sim \mathcal{D}_X} [(p(x) - f(x))^2] \leq \epsilon^2$. Then given independent draws from \mathcal{D} , the L_1 polynomial regression algorithm runs in time $\text{poly}(n^d, 1/\epsilon, \log(1/\delta))$ and with probability $1 - \delta$ outputs a hypothesis $h : X \times \{-1, 1\}$ such that $\mathbf{Pr}_{(x,y) \sim \mathcal{D}} [h(x) \neq y] \leq \text{opt} + \epsilon$, where $\text{opt} = \min_{f \in \mathcal{C}} \mathbf{Pr}_{(x,y) \sim \mathcal{D}} [f(x) \neq y]$.*

We first consider the case where \mathcal{D}_X is the uniform distribution over the n -dimensional Boolean hypercube $\{-1, 1\}^n$. Klivans et al. [\[KOS04\]](#) observed that Boolean noise sensitivity bounds are easily shown to imply the existence of low-degree polynomial approximators in the L_2 norm under the uniform distribution on $\{-1, 1\}^n$:

Fact 8.2 *For any Boolean function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and any value $0 \leq \gamma < 1/2$, there is a polynomial $p(x)$ of degree at most $d = 1/\gamma$ such that $\mathbf{E}[(p(x) - f(x))^2] \leq \frac{2}{1-e^{-2}} \text{NS}_\gamma(f)$.*

[Theorem 1.6](#) follows directly from [Theorem 8.1](#), [Fact 8.2](#) and [Theorem 1.3](#).

Next we turn to the case where \mathcal{D}_X is the $\mathcal{N}(0, I_n)$ distribution over \mathbb{R}^n . In [\[KOS08\]](#) observed that using entirely similar arguments to the Boolean case, Gaussian noise sensitivity bounds imply the existence of low-degree polynomial approximators in the L_2 norm:

Fact 8.3 For any Boolean function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and any value $0 \leq \gamma < 1/2$, there is a polynomial $p(x)$ of degree at most $d = 1/\gamma$ such that $\mathbf{E}_{\mathcal{G} \sim \mathcal{N}(0, I_n)}[(p(\mathcal{G}) - f(\mathcal{G}))^2] \leq \frac{2}{1-e^{-1}} \text{GNS}_\gamma(f)$.

For the special case of learning under the standard multivariate Gaussian \mathcal{N}^n , [Theorem 1.7](#) follows directly from [Theorem 8.1](#), [Fact 8.3](#) and [Theorem 1.5](#). Since our results hold for *all* degree- d PTFs, the extension to arbitrary Gaussian distributions follows exactly as described in Appendix C of [\[KOS08\]](#).

9 Discussion

An obvious question left open by this work is to actually resolve the Gotsman-Linial conjecture and show that every degree- d PTF over $\{-1, 1\}^n$ has average sensitivity at most $O(d\sqrt{n})$. [\[GS09\]](#) show that this would have interesting implications in computational learning theory beyond the obvious strengthenings of the agnostic learning results presented in this paper.

In this section we observe ([Proposition 9.1](#)) that this conjecture is in fact equivalent to a strong upper bound on the Boolean noise sensitivity of degree- d PTFs. We further point out ([Proposition 9.2](#)) that Gaussian noise sensitivity of degree- d PTFs is upper bounded by Boolean noise sensitivity. Thus, we propose working on improved upper bounds for the Gaussian noise sensitivity of degree- d PTFs as a preliminary – in fact, necessary – step to settling the Gotsman-Linial conjecture.

Proposition 9.1 *The following two statements are equivalent:*

1. Every degree- d PTF over $\{-1, 1\}^n$ has $\text{AS}(f) \leq O(d\sqrt{n})$.
2. Every degree- d PTF over $\{-1, 1\}^n$ has $\text{NS}_\epsilon(f) \leq O(d\sqrt{\epsilon})$ for all ϵ .

Proof:

1) \Rightarrow 2): This follows immediately from [Theorem 7.1](#)

2) \Rightarrow 1): Let $f = \text{sign}(p)$ be a degree- d PTF. We have

$$\begin{aligned}
 \text{NS}_{1/n}(f) &= \Pr_{x,y}[f(x) \neq f(y)] \\
 &= \sum_{k=0}^n \Pr_{x,y}[f(x) \neq f(y) \mid y \text{ flips } k \text{ of } x\text{'s bits}] \cdot \Pr_{x,y}[y \text{ flips } k \text{ of } x\text{'s bits}] \\
 &\geq \Pr_{x,y}[f(x) \neq f(y) \mid y \text{ flips 1 of } x\text{'s bits}] \cdot \Pr_{x,y}[y \text{ flips 1 of } x\text{'s bits}] \\
 &\geq (1/n)\text{AS}(f) \cdot \Theta(1),
 \end{aligned}$$

where the last inequality holds because at noise rate $1/n$, there is constant probability that y flips exactly 1 of x 's bits, and conditioned on this taking place, the probability that $f(x) \neq f(y)$ is exactly $\text{AS}(f)/n$. Taking $\epsilon = 1/n$ in 2) and rearranging, we get 1). \blacksquare

Proposition 9.2 *Let $\text{NS}(\epsilon, d)$ and $\text{GNS}_{\epsilon, d}$ denote the maximum noise sensitivity of a degree d PTF in the Boolean and Gaussian domains respectively. For all ϵ and d , we have*

$$\text{NS}(\epsilon, d) \geq \text{GNS}(\epsilon, d).$$

Proof: Consider a degree- d PTF $f = \text{sign}(p(x))$ in the Gaussian setting. We will define a sequence of degree- d PTFs $\{h_k\}_{k=1}^{\infty}$ over the Boolean domain. The function $h_k : \{-1, 1\}^{nk} \rightarrow \{-1, 1\}$ is on nk input bits $\{y_i^{(j)} | i \in [n], j \in [k]\}$ and is given by,

$$h_k(y_1^{(1)}, y_1^{(2)}, \dots, y_n^{(k)}) \stackrel{\text{def}}{=} \text{sign} \left(p \left(\frac{\sum_{j \in [k]} y_1^{(j)}}{\sqrt{k}}, \frac{\sum_{j \in [k]} y_2^{(j)}}{\sqrt{k}}, \dots, \frac{\sum_{j \in [k]} y_n^{(j)}}{\sqrt{k}} \right) \right).$$

By the Central Limit Theorem, the normalized sum $\frac{\sum_{j \in [k]} y_i^{(j)}}{\sqrt{k}}$ of k independent random values from $\{-1, 1\}$, tends to in distribution to the normal distribution $\mathcal{N}(0, 1)$ as $k \rightarrow \infty$. Intuitively, this implies that as $k \rightarrow \infty$, among other things the Boolean noise sensitivity of h_k approaches the noise sensitivity of f . However, since h_k is a Boolean PTF its noise sensitivity is bounded by $\text{NS}(\epsilon, d)$.

We now present the details of the above argument. Consider the random variables $y = (y_1, \dots, y_n), \tilde{y} = (\tilde{y}_1, \dots, \tilde{y}_n) \in \{-1, 1\}^n$ generated by setting each y_i to an uniform random value in $\{-1, 1\}$ and \tilde{y}_i as

$$\tilde{y}_i = \begin{cases} y_i & \text{with probability } 1 - \epsilon \\ \text{uniform value in } \{-1, 1\} & \text{with probability } \epsilon. \end{cases}$$

It is clear that $\mathbf{E}[y_i \tilde{y}_i] = 1 - \epsilon$ for all $i \in [n]$ and all other pairwise correlations are 0. Let $\{(y^{(1)}, \tilde{y}^{(1)}), \dots, (y^{(k)}, \tilde{y}^{(k)})\}$ be k independent samples of (y, \tilde{y}) . By definition of Boolean noise sensitivity,

$$\begin{aligned} \text{NS}_{\epsilon}(h_k) &= \Pr[h_k(y) \neq h_k(\tilde{y})] \\ &= \Pr \left[p \left(\frac{\sum_{j \in [k]} y^{(j)}}{\sqrt{k}} \right) \cdot p \left(\frac{\sum_{j \in [k]} \tilde{y}^{(j)}}{\sqrt{k}} \right) \leq 0 \right]. \end{aligned}$$

Let $x \sim \mathcal{N}^n, z \sim \mathcal{N}^n$ be independent and let $\tilde{x} = \alpha x + \beta z$, with $\alpha = 1 - \epsilon$ and $\beta = \sqrt{2\epsilon - \epsilon^2}$. By the Multidimensional Central Limit Theorem [Fel68], as $k \rightarrow \infty$ we have the following convergence in distribution,

$$\left(\frac{\sum_{j \in [k]} y^{(j)}}{\sqrt{k}}, \frac{\sum_{j \in [k]} \tilde{y}^{(j)}}{\sqrt{k}} \right) \xrightarrow{\mathcal{D}} (x, \tilde{x}).$$

Since the function $a(x, \tilde{x}) = p(x) \cdot p(\tilde{x})$ is a continuous function we get

$$\begin{aligned} \lim_{k \rightarrow \infty} \text{NS}_\epsilon(h_k) &= \lim_{k \rightarrow \infty} \Pr \left[p \left(\frac{\sum_{j \in [k]} y^{(j)}}{\sqrt{k}} \right) \cdot p \left(\frac{\sum_{j \in [k]} \tilde{y}^{(j)}}{\sqrt{k}} \right) \leq 0 \right] \\ &= \Pr_{x, \tilde{x}} [p(x)p(\tilde{x}) \leq 0] \\ &= \text{GNS}_\epsilon(f) \end{aligned}$$

and the result is proved. ■

References

- [AH09] Per Austrin and Johan Håstad. Randomly supported independence and resistance. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 483–492. ACM, 2009. [8](#), [9](#)
- [BK97] J. Bourgain and G. Kalai. Influences of variables and threshold intervals under group symmetries. *GAFSA*, 7:438–461, 1997. [1](#)
- [BKS99] I. Benjamini, G. Kalai, and O. Schramm. Noise sensitivity of Boolean functions and applications to percolation. *Inst. Hautes Études Sci. Publ. Math.*, 90:5–43, 1999. [1](#), [2](#)
- [Bog98] V. Bogachev. *Gaussian measures*. Mathematical surveys and monographs, vol. 62, 1998. [3](#), [8](#)
- [BOW08] E. Blais, R. O’Donnell, and K. Wimmer. Polynomial regression under arbitrary product distributions. In *Proc. 21st Annual Conference on Learning Theory (COLT)*, pages 193–204, 2008. [3](#)
- [BT96] N. Bshouty and C. Tamon. On the Fourier spectrum of monotone functions. *Journal of the ACM*, 43(4):747–770, 1996. [1](#)
- [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. [0](#), [5](#), [8](#)
- [DFKO06] Irit Dinur, Ehud Friedgut, Guy Kindler, and Ryan O’Donnell. On the fourier tails of bounded functions over the discrete cube. In *Proc. 38th ACM Symp. on Theory of Computing*, pages 437–446, 2006. [9](#)
- [DGJ+09] I. Diakonikolas, P. Gopalan, R. Jaiswal, R. Servedio, and E. Viola. Bounded independence fools halfspaces. To appear in *Proc. 50th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2009. [6](#), [21](#)

- [DS09] I. Diakonikolas and R. Servedio. Improved approximation of linear threshold functions. In *Proc. 24th Annual IEEE Conference on Computational Complexity (CCC)*, pages 161–172, 2009. [21](#)
- [DSTW09] I. Diakonikolas, R. Servedio, L.-Y. Tan, and A. Wan. A regularity lemma, and low-weight approximators, for low-degree polynomial threshold functions. manuscript, 2009. [5](#)
- [Fel68] W. Feller. *An introduction to probability theory and its applications*. John Wiley & Sons, 1968. [35](#)
- [Fri98] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):474–483, 1998. [1](#)
- [GKK08] P. Gopalan, A. Kalai, and A. Klivans. Agnostically learning decision trees. In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 527–536, 2008. [3](#)
- [GL94] C. Gotsman and N. Linial. Spectral properties of threshold functions. *Combinatorica*, 14(1):35–50, 1994. [0](#), [2](#)
- [GS09] P. Gopalan and R. Servedio. Learning threshold-of- ac^0 circuits. manuscript, 2009. [4](#), [34](#)
- [HKM09] P. Harsha, A. Klivans, and R. Meka. Bounding the sensitivity of polynomial threshold functions. manuscript, 2009. [5](#)
- [Jan97] S. Janson. *Gaussian Hilbert Spaces*. Cambridge University Press, Cambridge, UK, 1997. [0](#), [5](#), [8](#)
- [JKS02] J. Jackson, A. Klivans, and R. Servedio. Learnability beyond AC^0 . In *Proc. 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 776–784, 2002. [4](#)
- [KKL88] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proc. 29th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 68–80, 1988. [1](#)
- [KKMS08] A. Kalai, A. Klivans, Y. Mansour, and R. Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008. [0](#), [3](#), [4](#), [33](#)
- [KMV08] A. Kalai, Y. Mansour, and E. Verbin. On agnostic boosting and parity learning. In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 629–638, 2008. [3](#)
- [KOS04] A. Klivans, R. O’Donnell, and R. Servedio. Learning intersections and thresholds of halfspaces. *Journal of Computer & System Sciences*, 68(4):808–840, 2004. [2](#), [3](#), [4](#), [33](#)

- [KOS08] A. Klivans, R. O’Donnell, and R. Servedio. Learning geometric concepts via Gaussian surface area. In *Proc. 49th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 541–550, 2008. [2](#), [3](#), [4](#), [33](#), [34](#)
- [LMN93] N. Linial, Y. Mansour, and N. Nisan. Constant depth circuits, Fourier transform and learnability. *Journal of the ACM*, 40(3):607–620, 1993. [4](#)
- [MO03] E. Mossel and R. O’Donnell. On the noise sensitivity of monotone functions. *Random Structures and Algorithms*, 23(3):333–350, 2003. [1](#), [2](#)
- [MOO05] E. Mossel, R. O’Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Proc. 46th Symposium on Foundations of Computer Science (FOCS)*, pages 21–30, 2005. [0](#), [1](#), [3](#), [6](#), [7](#), [9](#)
- [Mos05] E. Mossel. Lecture 4. <http://www.stat.berkeley.edu/~mossel/teach/206af05/scribes/sep8.pdf>, 2005. [40](#)
- [O’D07] R. O’Donnell. Lecture 16: The hypercontractivity theorem. <http://www.cs.cmu.edu/~odonnell/boolean-analysis/lecture16.pdf>, 2007. [8](#)
- [OS07] R. O’Donnell and R. Servedio. Learning monotone decision trees in polynomial time. *SIAM J. Comput.*, 37(3):827–844, 2007. [1](#)
- [OS08] R. O’Donnell and R. Servedio. The Chow Parameters Problem. In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 517–526, 2008. [21](#)
- [OSSS05] R. O’Donnell, M. Saks, O. Schramm, and R. Servedio. Every decision tree has an influential variable. In *Proc. 46th Symposium on Foundations of Computer Science (FOCS)*, pages 31–39, 2005. [1](#)
- [Per04] Y. Peres. Noise stability of weighted majority, 2004. [31](#)
- [Ser07] R. Servedio. Every linear threshold function has a low-weight approximator. *Computational Complexity*, 16(2):180–209, 2007. [0](#), [5](#), [6](#), [20](#), [21](#)
- [Shi00] Y. Shi. Lower bounds of quantum black-box complexity and degree of approximating polynomials by influence of boolean variables. *Inform. Process. Lett.*, 75(1-2):79–83, 2000. [1](#)
- [SS] O. Schramm and J. Steif. Quantitative noise sensitivity and exceptional times for percolation. *Ann. Math.*, to appear. [2](#)
- [SSS09] S. Shalev Shwartz, O. Shamir, and K. Sridharan. Agnostically learning halfspaces with margin errors. TTI Technical Report, 2009. [3](#)

A Basics of Hermite Analysis

Here we briefly review the basics of Hermite analysis over \mathbb{R}^n under the distribution \mathcal{N}^n . The reader who is unfamiliar with Hermite analysis should note the many similarities to Fourier analysis over $\{-1, 1\}^n$.

We work within $L^2(\mathbb{R}^n, \mathcal{N}^n)$, the vector space of all functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $\mathbf{E}_{x \sim \mathcal{N}^n}[f(x)^2] < \infty$. This is an inner product space under the inner product

$$\langle f, g \rangle = \mathbf{E}_{x \sim \mathcal{N}^n}[f(x)g(x)].$$

This inner product space has a complete orthonormal basis given by the *Hermite polynomials*. In the case $n = 1$, this basis is the sequence of polynomials

$$h_0(x) = 1, \quad h_1(x) = x, \quad h_2(x) = \frac{x^2 - 1}{\sqrt{2}}, \quad h_3(x) = \frac{x^3 - 3x}{\sqrt{6}}, \quad \dots,$$

$$h_j(x) = \frac{\sqrt{j!}}{(j-0)!0!2^0}x^j - \frac{\sqrt{j!}}{(j-2)!1!2^1}x^{j-2} + \frac{\sqrt{j!}}{(j-4)!2!2^2}x^{j-4} - \frac{\sqrt{j!}}{(j-6)!3!2^3}x^{j-6} + \dots$$

which may equivalently be defined by

$$h_j(x) = \frac{(-1)^d}{\sqrt{d!} \exp(-x^2/2)} \cdot \frac{d^j}{dx^j} \exp(-x^2/2).$$

We note that $h_d(x)$ is a polynomial of degree d . For general n , the basis for $L^2(\mathbb{R}^n, \mathcal{N}^n)$ is formed by all products of these polynomials, one for each coordinate. In other words, for each n -tuple $S \in \mathbb{N}^n$ we define the n -variate Hermite polynomial $H_S : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$H_S(x) = \prod_{i=1}^n h_{S_i}(x_i);$$

then the collection $(H_S)_{S \in \mathbb{N}^n}$ is a complete orthonormal basis for the inner product space. By orthonormal we mean that

$$\langle H_S, H_T \rangle = \begin{cases} 1 & \text{if } S = T, \\ 0 & \text{if } S \neq T. \end{cases}$$

By complete, we mean that every function $f \in L^2$ can be uniquely expressed as

$$f(x) = \sum_{S \in \mathbb{N}^n} \hat{f}(S) H_S(x),$$

where the coefficients $\hat{f}(S)$ are real numbers and the infinite sum converges in the sense that

$$\lim_{d \rightarrow \infty} \mathbf{E} \left[\left(f(x) - \sum_{|S| \leq d} c_S H_S(x) \right)^2 \right] = 0;$$

here we have used the notation

$$|S| = \sum_{i=1}^n S_i,$$

which is also the total degree of $H_S(x)$ as a polynomial.

We call $\widehat{f}(S)$ the S *Hermite coefficient* of f . By orthonormality of the basis $(H_S)_{S \in \mathcal{N}^n}$, we have the following:

$$\begin{aligned} \widehat{f}(S) &= \langle f, H_S \rangle = \mathbf{E}[f(x)H_S(x)]; \\ \|f\|_2^2 &\stackrel{\text{def}}{=} \langle f, f \rangle = \sum_{S \in \mathcal{N}^n} \widehat{f}(S)^2 \quad (\text{“Parseval’s identity”}); \\ \langle f, g \rangle &= \sum_{S \in \mathcal{N}^n} \widehat{f}(S)\widehat{g}(S) \quad (\text{“Plancherel’s identity”}). \end{aligned}$$

In particular, if $f : \mathbb{R}^n \rightarrow \{-1, 1\}$, then $\sum_S \widehat{f}(S)^2 = 1$.

Using the definition of influence from Section 2.1, it is not difficult to show that for any $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and any $i \in [n]$, we have $\text{GI}_i(f) = \sum_{S: S_i > 0} \widehat{f}(S)^2$ (see e.g. Lecture 4 of [Mos05]).