THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Average Sensitivity and Noise Sensitivity of Polynomial Threshold Functions

Citation for published version:

Diakonikolas, I, Raghavendra, P, Servedio, RA & Tan, L 2014, 'Average Sensitivity and Noise Sensitivity of Polynomial Threshold Functions', *SIAM Journal on Computing*, vol. 43, no. 1, pp. 231-253.
<https://doi.org/10.1137/110855223>

Digital Object Identifier (DOI):

[10.1137/110855223](https://doi.org/10.1137/110855223)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

SIAM Journal on Computing

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

AVERAGE SENSITIVITY AND NOISE SENSITIVITY OF POLYNOMIAL THRESHOLD FUNCTIONS*

ILIAS DIAKONIKOLAS[†], PRASAD RAGHAVENDRA[‡], ROCCO A. SERVEDIO[§], AND
LI-YANG TAN[¶]

Abstract. We give the first non-trivial upper bounds on the Boolean average sensitivity and noise sensitivity of degree- d polynomial threshold functions (PTFs). Our bound on the Boolean average sensitivity of PTFs represents the first 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 bound on Boolean noise sensitivity yields the first polynomial-time agnostic learning algorithm for the broad class of constant-degree PTFs under the uniform distribution.

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 [MOO10], this allows us to essentially reduce the Boolean setting to the Gaussian setting. The main ingredients used to obtain our bound in the Gaussian setting are tail bounds and anti-concentration bounds on low-degree polynomials in Gaussian random variables [Jan97, CW01]. 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.

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. (The function $\text{sign}(z)$ takes value 1 for $z \geq 0$ and -1 for $z < 0$.) 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, MOO10, 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

*An extended abstract merging results from an earlier version of this paper [DRST09] with results from [HKM09] appeared as “Bounding the average sensitivity and noise sensitivity of polynomial threshold functions” in the *Proceedings of the 42nd ACM Symposium on Theory of Computing* [DHK⁺10].

[†]School of Informatics, University of Edinburgh. Email: ilias.d@ed.ac.uk. This research was done while at Columbia University and visiting IBM Almaden. Supported by NSF grant CCF-0728736, and by an Alexander S. Onassis Foundation Fellowship.

[‡]Computer Science Division, UC Berkeley. Email: prasad@cs.berkeley.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.

“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 $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$.

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. $\text{Inf}_i(f) = \Pr[f(x) \neq f(x^{\oplus i})]$ (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 . We note that the noise sensitivity can be equivalently expressed as a function of the Fourier coefficients of f and the noise rate ϵ as follows:

$$\text{NS}_\epsilon(f) = \frac{1}{2} - \frac{1}{2} \sum_{S \subseteq [n]} (1 - 2\epsilon)^{|S|} \hat{f}(S)^2. \quad (1.1)$$

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. 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 for $d = O(\sqrt{n})$ it is equivalent to) the conjecture that every degree- d PTF f over $\{-1, 1\}^n$ has $\text{AS}(f) = O(d\sqrt{n})$.

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*

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

In [Section 7](#) we point out ([Proposition 7.1](#)) that, assuming the Gotsman-Linial conjecture, any degree- d PTF $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ has noise sensitivity $\text{NS}_\epsilon(f) = O(d\sqrt{\epsilon})$.

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 6](#).

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](#), [KMOV08](#), [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 6](#)):

THEOREM 1.4. *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 $n^{O(1/\epsilon^{2^{d+1}})}$.

For ϵ constant, this result is 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 [[GS10](#)] 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. [[DSTW10](#)] 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. [[DSTW10](#)] 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.

A preliminary version of this paper [[DRST09](#)] gave quantitatively similar upper bounds on the *Gaussian* average sensitivity and noise sensitivity of degree- d PTFs. (See [Section 2.1](#) for a definition.) A few months after [[DRST09](#)] appeared, Daniel Kane gave an elegant proof of optimal upper bounds on the Gaussian noise sensitivity of degree- d PTFs by showing that for any degree- d PTF f it holds $\text{GNS}_\epsilon(f) \leq d\sqrt{2\epsilon}/(2\pi)$ [[Kan10](#)]. As we point out in [Section 7](#), this is a necessary step towards the resolution of the Boolean version of the Gotsman-Linial conjecture.

1.5. Techniques. In this section we give a high-level overview of how [Theorem 1.1](#) is proved. (As mentioned earlier, [Theorem 1.2](#) is proved using completely different techniques; see [Section 4](#).)

An important notion in our proof 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 [[MOO10](#)] 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 $\mathcal{N}(0, I_n)$).

We start by sketching the argument for the regular case. Let $f = \text{sign}(p)$ be a regular PTF, where $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ is a degree- d polynomial. Recall that the average sensitivity of f is equal to the sum of the individual influences $\sum_{i=1}^n \text{Inf}_i(f)$. We proceed by showing that each individual influence is small. It follows by the

definition that the influence of the i th coordinate is bounded by the probability that p is smaller than its partial derivative in the i th direction, $D_i p$. Intuitively, this event can only take place if either

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

We use an anti-concentration result for polynomials in Gaussian random variables, due to Carbery and Wright [CW01], combined with the invariance principle, to show that $|p(x)|$ is “small” only with low probability. For the second bullet, note that $D_i p$ is a low-degree polynomial in independent random variables. As a consequence of the regularity of f , $D_i p$ has small 2-norm and thus a tail bound for this setting [AH09] implies that $|D_i p(x)|$ is “large” only with low probability. We can thus argue that $\text{Inf}_i(f)$ is low, and bound the Boolean average sensitivity of f .

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 3.1 for a detailed explanation of the recursive argument.

1.6. Organization. Formal definitions of average sensitivity and noise sensitivity, and tail bounds and anti-concentration results for low degree polynomials are presented in Section 2. The main result of the paper – a bound on the Boolean average sensitivity (Theorem 1.1) – is proved in Section 3. In Section 4, 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 5.1) in Section 5. We present the applications of these upper bounds to agnostic learning of PTFs in Section 6. Section 7 concludes by proposing directions 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 $n \in \mathbb{Z}^+$, we denote by $[n]$ the set $\{1, 2, \dots, n\}$. For $i \leq j \in \mathbb{Z}^+$, we denote by $[i, j]$ the set $\{i, i+1, \dots, j\}$. We write $\mathcal{N}$ to denote the standard univariate Gaussian distribution $\mathcal{N}(0, 1)$.

For a degree- d polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ and $q \geq 1$, we denote by $\|p\|_q$ its l_q norm, $\|p\|_q = \mathbf{E}_x[|p(x)|^q]^{1/q}$, 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 [MOO10]).

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

[MOO10] 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,

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

and as mentioned in the introduction it is easily seen that $\text{AS}(f) = \text{Inf}(f)$ for Boolean-valued functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$. For $i \in [n]$ the i -th partial derivative of $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ is $(D_i f)(x) = \sum_{S \ni i} \hat{f}(S) \chi_{S - \{i\}}$ and so $\text{Inf}_i(f) = \|D_i f\|_2^2$ by Parseval's identity.

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 2.1 (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*

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

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 Analysis. We assume familiarity with the basics of Fourier analysis over the Boolean hypercube $\{-1, 1\}^n$. Here we consider functions $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, and we think of the inputs x to f as being distributed according to the uniform probability distribution. The set of such functions forms a 2^n -dimensional inner product space with inner product given by $\langle f, g \rangle = \mathbf{E}_x[f(x)g(x)]$. The set of functions $(\chi_S)_{S \subseteq [n]}$ defined by $\chi_S(x) = \prod_{i \in S} x_i$ forms a complete orthonormal basis for this space. We will also often write simply x_S for $\prod_{i \in S} x_i$. Given a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ we define its *Fourier coefficients* by $\hat{f}(S) = \mathbf{E}_x[f(x)x_S]$, and we have that $f(x) = \sum_S \hat{f}(S)x_S$. Finally, we have *Plancherel's identity* $\langle f, g \rangle = \sum_S \hat{f}(S)\hat{g}(S)$, which has as a special case *Parseval's identity*, $\mathbf{E}_x[f(x)^2] = \sum_S \hat{f}(S)^2$. From this it follows that for every $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ we have $\sum_S \hat{f}(S)^2 = 1$.

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 [Bon70, Bec75]:

THEOREM 2.2. *Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a degree- d polynomial, where $\{-1, 1\}^n$ is endowed with the uniform distribution, 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. It can be proved (in both cases) using Markov's inequality and hypercontractivity, see e.g. [AH09].

THEOREM 2.3 (“degree- d Chernoff bound”). *Let $p(x)$ be a degree- d polynomial. Let x be drawn either from the uniform distribution over $\{-1, 1\}^n$ 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 2.4 ([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 2.5 ([DFKO06, AH09]). *Let $p : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a degree- d polynomial with $\text{Var}[p] \equiv \sum_{0 < |S| \leq d} \hat{p}(S)^2 = 1$ and $\mathbf{E}[p] = \hat{p}(\emptyset) = 0$. 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 [MOO10].

THEOREM 2.6 ([MOO10]). *Let $p(x) = \sum_{|S| \leq d} \hat{p}(S)x_S$ be a degree- d multilinear polynomial with $\sum_{0 < |S| \leq d} \hat{p}(S)^2 = 1$. Suppose each variable $i \in [n]$ has low influence $\text{Inf}_i(p) \leq \tau$, i.e. $\sum_{S \ni i} \hat{p}(S)^2 \leq \tau$. 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. 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:

$$\text{AS}(n, d) \leq 2^{O(d)} \cdot \log n \cdot n^{1-1/(4d+2)}. \quad (3.1)$$

For $d = 1$ (linear threshold functions) it is well known (see e.g. [GL94]) that $\text{AS}(n, 1) \leq n2^{-n+1} \binom{n}{n/2} = \Theta(\sqrt{n})$. Also, notice that the RHS of (3.1) 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 $2 \leq d \leq O(\sqrt{\log n})$.

3.1. 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 3.2 for a precise definition) we argue directly that the average sensitivity is small. In particular, we show:

CLAIM 3.1. *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 3.1 follows directly from Lemma 3.10, which we prove in Section 3.4.

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. Let ρ be an assignment to a subset of the variables. In the following, we will denote by f_ρ the function obtained from f after fixing these variables. More precisely, we show:

CLAIM 3.2. *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 3.2** is given in Section 3.7. 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 3.5). 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 3.6).

3.1.1. Proof of main result (Theorem 1.1) assuming Claim 3.1 and Claim 3.2. Given these two claims it is not difficult to obtain the final result. In **Claim 3.2**, 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 $\text{AS}(n, d)$:

$$\begin{aligned} \text{AS}(n, d) &\leq \max\{O(d \cdot n^{1-1/(4d+2)}), \\ &\quad \max_{1 \leq k \leq K, \ 1/2^{O(d)} \leq \alpha \leq 1} \{k + \alpha \cdot O(d \cdot (\log n)^{1/4} \cdot n^{1-1/(4d+2)}) + (1 - \alpha)\text{AS}(n, d)\}\}, \end{aligned}$$

where K is a function of n and d as defined in the statement of **Claim 3.2**. 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 $\text{AS}'(n, d) \leq 2^{O(d)} \cdot \log n \cdot n^{1-1/(4d+2)}$.

3.2. 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 3.3. 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:

$$\frac{\text{Inf}_{i+1}(p)}{\sum_{j=i+1}^n \text{Inf}_j(p)} \leq \tau. \quad (3.2)$$

If (3.2) does not hold for any i we say that the τ -critical index of p is $+\infty$. If p is has τ -critical index 0, we say that p is τ -regular.

DEFINITION 3.4. Let $\tau > 0$. A sequence $a_1 \geq \dots \geq a_n \geq 0$ of non-negative numbers is τ -regular if $a_i / (\sum_{j=i+1}^n a_j) \leq \tau$ for all $i \in [n]$.

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 3.5. 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. $\square$

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 $|S| \leq d$.

3.3. 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 apply a random restriction to variables $x_1, \dots, x_k$.

We start with the following elementary claim:

CLAIM 3.6. 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

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

where ρ_T denotes the product of the values ρ assigns to the variables in T . Therefore, we have

$$\mathbf{E}_\rho[\widehat{p}_\rho(S)^2] = \sum_{T \subseteq [k]} \widehat{p}(S \cup T)^2. \quad (3.3)$$

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 3.7. *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.

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

□

3.3.1. 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 $\text{polylog}(n)$ factor.

LEMMA 3.8. *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 $\rho_1, \dots, \rho_k$, 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, the Bonami-Beckner inequality (Theorem 2.2) applied for $q = 4$ 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 3.7](#). $\square$

LEMMA 3.9. *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 3.9](#) follows as an immediate consequence of [Theorem 2.3](#) and the upper bound on $\|\text{Inf}_\ell(p_\rho)\|_2$ given by [Lemma 3.8](#). $\square$

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

LEMMA 3.10. *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 3.1](#) follows directly from the above lemma, recalling that we choose $\tau \stackrel{\text{def}}{=} n^{-(4d+1)/(4d+2)}$. However, the lemma will also be useful in the “small critical index” case of [Section 3.5](#) 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:

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

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 2.4](#). For the second event we use the tail bound for degree- d polynomials ([Theorem 2.3](#)).

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

3.4.1. 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 2.6](#)) tells us that $\Pr_{x \in \{-1,1\}^n}[|p(x)| \leq \tau^{1/4}]$ differs from $\Pr_{\mathcal{G} \sim \mathcal{N}^n}[|p(\mathcal{G})| \leq \tau^{1/4}]$ 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 2.4](#)), we get:

$$\begin{aligned} \Pr_x[|p(x)| \leq \tau^{1/4}] &\leq \Pr_{\mathcal{G} \sim \mathcal{N}^n}[|p(\mathcal{G})| \leq \tau^{1/4}] + O(d\tau^{1/(4d+1)}) \\ &\leq O(d \cdot \tau^{1/4d}) + O(d \cdot \tau^{1/(4d+1)}) \\ &= O(d \cdot \tau^{1/(4d+1)}). \end{aligned}$$

3.4.2. 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:

$$\|2D_i p\|_2 = 2 \sqrt{\sum_{S \ni i} \widehat{p}(S)^2} = 2\sqrt{\text{Inf}_i(p)} \leq 2\sqrt{d \cdot \tau}.$$

By ([Theorem 2.3](#)), 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 $f = \text{sign}(p)$:

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

□

3.5. 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 3.11. *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 3.10](#) 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 3.9](#) 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 $\rho_1, \dots, \rho_k$. The expected value of A is $\mathbf{E}_\rho[A] = \sum_{j=k+1}^n \text{Inf}_j(p) = \hat{A}(\emptyset)$. We apply the [Theorem 2.5](#) for $B = A - \hat{A}(\emptyset)$. We thus get $\mathbf{Pr}_\rho[B > 0] > 1/2^{O(d)}$. We thus get $\mathbf{Pr}_\rho[A > \mathbf{E}_\rho[A]] > 1/2^{O(d)}$ and we are done. $\square$

3.6. 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 3.12. *For a $1/2^{O(d)}$ fraction of restrictions ρ , the function $\text{sign}(p_\rho(x))$ is a constant function.*

Proof. By [Lemma 3.5](#), the surviving variables $x_{K+1}, \dots, x_n$ have very small total influence in p :

$$\sum_{i=K+1}^n \text{Inf}_i(p) = \sum_{i=K+1}^n \sum_{S \ni i} \hat{p}(S)^2 \leq (1 - \tau)^K \cdot \text{Inf}(p) \leq d/n^{2d}, \quad (3.4)$$

where we use the fact that $\text{Inf}(f) \leq d$ and our choice of K for the final inequality. Therefore, if we let p' be the truncation of p comprising only the monomials with all variables in $\mathcal{H}$,

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

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

$$1 \geq \sum_{\emptyset \neq S \subseteq \mathcal{H}} \hat{p}(S)^2 \geq 1 - \sum_{i=K+1}^n \text{Inf}_i(p) \geq 1 - d/n^{2d}.$$

We now apply [Theorem 2.5](#) to p' and get:

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

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:

$$\sum_{\emptyset \neq S \subseteq [K+1, \dots, n]} |\widehat{p}_\rho(S)| \geq 1/2^{O(d)}.$$

As noted in (3.4), 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 3.9, 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 :

$$\text{Inf}_\ell(p_\rho) \leq (3 \log n)^d \cdot \text{Inf}_\ell(p) \leq d \cdot (3 \log n)^d / n^{2d}. \quad (3.5)$$

Using Cauchy-Schwarz, we have

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

where we have used (3.5) (and our upper bound on d). From this we easily get that

$$\sum_{\emptyset \neq S \subseteq [K+1, n]} |\widehat{p}_\rho(S)| \leq n^{-\Omega(1)} \ll 1/2^{O(d)}.$$

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. $\square$

3.7. Proof of Claim 3.2. 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 3.5, 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 3.2 holds in this case.

In the second case (large critical index case), as shown in Section 3.6, 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 3.2 certainly holds in this case as well. $\blacksquare$

4. 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 $f : \{-1, 1\}^n \rightarrow \mathbb{R}$.

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

It is easy to see that

$$D_i f(x) = \frac{1}{2} x_i [f(x) - f(x^{\oplus i})] = \frac{1}{2} \left(\frac{f(x) - f(x^{\oplus i})}{x_i} \right) \quad (4.1)$$

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 4.1. *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:

$$\text{Inf}_i(f) = \mathbf{E}_{x_{-i}}[\text{Var}_{x_i}[f(x)]] = \mathbf{E}_x \left[\left(\frac{|f(x^{\oplus i}) - f(x)|}{2} \right)^2 \right] = \mathbf{E}_{x_{-i}} \left[|\mathbf{E}_{x_i}[x_i f(x)]|^2 \right] \quad (4.2)$$

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

$$\begin{aligned} \mathbf{E}_x[x_i x_j f(x) g(x)] &= \mathbf{E}_{x_{-\{i,j\}}} \mathbf{E}_{x_i, x_j} [x_i x_j f(x) g(x)] \\ &= \mathbf{E}_{x_{-\{i,j\}}} [\mathbf{E}_{x_i}[x_i g(x)] \mathbf{E}_{x_j}[x_j f(x)]] \\ &\leq \mathbf{E}_{x_{-\{i,j\}}} \left[\frac{1}{2} |\mathbf{E}_{x_i}[x_i g(x)]|^2 + \frac{1}{2} |\mathbf{E}_{x_j}[x_j f(x)]|^2 \right] \quad (\text{using } ab \leq \frac{1}{2}(a^2 + b^2)) \\ &\leq \frac{\text{Inf}_j(f) + \text{Inf}_i(g)}{2} \quad (\text{using Equation 4.2}). \end{aligned}$$

□

[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 4.2. For a PTF $f = \text{sign}(p(x))$ on n variables and $i \in [n]$, $\text{Inf}_i(f) = \mathbf{E}[f(x)x_i\text{sign}(D_ip(x))]$. The following simple claim will be useful in the proof of the above lemma.

CLAIM 4.3. 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 $\text{sign}(a - b) = -1 = \text{sign}((-1) - 1) = \text{sign}(\text{sign}(a) - \text{sign}(b))$. $\square$

Proof. [of Lemma 4.2] The influence of the i^{th} coordinate is given by,

$$\begin{aligned} \text{Inf}_i(f) &= \mathbf{E} \left[\frac{1}{2} |f(x) - f(x^{\oplus i})| \right] \\ &= \mathbf{E} \left[\frac{1}{2} (f(x) - f(x^{\oplus i})) \text{sign}(f(x) - f(x^{\oplus i})) \right]. \end{aligned} \quad (4.3)$$

Consider an x for which $f(x) \neq f(x^{\oplus i})$. In this case, we can use Claim 4.3 to conclude:

$$\begin{aligned} \text{sign}(f(x) - f(x^{\oplus i})) &= \text{sign}(p(x) - p(x^{\oplus i})) \\ &= \text{sign}(2x_i D_i p(x)) = x_i \text{sign}(D_i p(x)). \quad (\text{using (4.1)}) \end{aligned}$$

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 to hold since both sides evaluate to 0. Substituting this equality into Equation 4.3 yields,

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

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

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

$\square$

THEOREM 4.4. 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.

$$\begin{aligned}
\text{Inf}(f) &= \sum_i \text{Inf}_i(f) \\
&= \sum_i \mathbf{E}[f(x)x_i \text{sign}(D_i p(x))] \quad (\text{by Lemma 4.2}) \\
&= \mathbf{E}\left[f(x) \sum_i x_i \text{sign}(D_i p(x))\right] \\
&\leq \sqrt{\mathbf{E}[f(x)^2]} \cdot \sqrt{\mathbf{E}\left[\left(\sum_i x_i \text{sign}(D_i p(x))\right)^2\right]} \quad (4.4)
\end{aligned}$$

$$= 1 \cdot \sqrt{\mathbf{E}\left[\sum_{i,j} x_i x_j \text{sign}(D_i p(x)) \text{sign}(D_j p(x))\right]} \quad (4.5)$$

$$\leq \sqrt{\mathbf{E}\left[\sum_i x_i^2 \text{sign}(D_i p(x))^2\right] + \sum_{i \neq j} \text{Inf}_i(\text{sign}(D_j p(x)))} \quad (4.6)$$

$$= \sqrt{n + \sum_{i \neq j} \text{Inf}_i(\text{sign}(D_j p(x)))}. \quad (4.7)$$

Here (4.4) is the Cauchy-Schwarz inequality, (4.5) is expanding the square. Step (4.6) uses Lemma 4.1 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,

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

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

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 4.4 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 4.4,

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

finishing the proof. $\square$

5. 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 5.1. *Let $\text{NS}(\epsilon, d)$ denote the maximum noise sensitivity of a degree d -PTF at a noise rate of ϵ . For all $0 < \epsilon < 1/2$, 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 5.1](#).

5.1. Proof of [Theorem 5.1](#). Let $f(x) = \text{sign}(p(x))$ be a degree d -PTF. Let us denote $\delta = \frac{1}{m} \leq \frac{1}{2}$. We remind the reader that, for any $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, the noise sensitivity $\text{NS}_\gamma(f)$ is a non-decreasing function of γ in the range $[0, 1/2]$. (This fact follows immediately from the Fourier expression [\(1.1\)](#).) As $1/2 \geq \delta \geq \epsilon$, it follows that $\text{NS}_\epsilon(f) \leq \text{NS}_\delta(f)$. 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

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

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:

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

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

$$\text{NS}_\delta(f) = \mathbf{E}_{a,\alpha} \left[\frac{1}{m} \sum_{r=1}^m \text{Inf}_r(f_{a,\alpha}) \right] = \frac{1}{m} \mathbf{E}_{a,\alpha} [\text{Inf}(f_{a,\alpha})] . \quad (5.1)$$

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 (5.1), the result follows.

6. Application to Agnostic Learning. In this section, we outline the application of the noise sensitivity bound presented in this work to agnostic learning of PTFs. Specifically, we will present the proof of [Theorem 1.4](#). To begin with, we recall the main theorem of [\[KKMS08\]](#) about the L_1 polynomial regression algorithm:

THEOREM 6.1. *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 with parameters δ and ϵ 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 6.2. *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 $\mathbf{E}[(p(x) - f(x))^2] \leq \frac{2}{1-\epsilon^{-2}} \text{NS}_\gamma(f)$.*

[Theorem 1.4](#) follows directly from [Theorem 6.1](#), [Fact 6.2](#) and [Theorem 1.3](#).

7. 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 $O(d\sqrt{n})$. [\[GS10\]](#) show that this would have interesting

implications in computational learning theory beyond the obvious strengthenings of the agnostic learning results presented in this paper. Currently, it seems that our techniques cannot avoid the $n^{1-\Omega(1/d)}$ dependency due to the inherent loss in the invariance principle (Theorem 2.6).

In this section we point out (Proposition 7.1) that the Gotsman-Linial conjecture is in fact equivalent to a strong upper bound on the Boolean noise sensitivity of degree- d PTFs. We further point out (Proposition 7.2) that Gaussian noise sensitivity of degree- d PTFs is upper bounded by Boolean noise sensitivity. Thus, improved upper bounds for the Gaussian noise sensitivity of degree- d PTFs is a necessary step to settling the Gotsman-Linial conjecture.

PROPOSITION 7.1. *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 5.1.

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

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

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). $\square$

PROPOSITION 7.2. *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 $\{y_i^{(j)} \mid i \in [n], j \in [k]\}$ 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,

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

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

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

and the result is proved. $\square$

REFERENCES

- [AH09] P. Austrin and J. Håstad. Randomly supported independence and resistance. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 483–492. ACM, 2009. [5](#), [7](#)
- [Bec75] W. Beckner. Inequalities in Fourier analysis. *Annals of Mathematics*, 102:159–182, 1975. [6](#)
- [BK97] J. Bourgain and G. Kalai. Influences of variables and threshold intervals under group symmetries. *Geometric and Functional Analysis (GAFA)*, 7:438–461, 1997. [1](#)
- [BKS99] I. Benjamini, G. Kalai, and O. Schramm. Noise sensitivity of Boolean functions and applications to percolation. *Publications Mathématiques de l’Institut des Hautes Études Scientifiques*, 90:5–43, 1999. [1](#), [2](#)
- [Bon70] A. Bonami. Etude des coefficients Fourier des fonctions de $L^p(G)$. *Annales de l’institut Fourier (Grenoble)*, 20(2):335–402, 1970. [6](#)
- [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](#), [2](#)
- [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. [1](#), [5](#), [7](#)

- [DFKO06] I. Dinur, E. Friedgut, G. Kindler, and R. O’Donnell. On the Fourier tails of bounded functions over the discrete cube. In *Proc. 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 437–446, 2006. 7
- [DGJ⁺09] I. Diakonikolas, P. Gopalan, R. Jaiswal, R. Servedio, and E. Viola. Bounded independence fools halfspaces. In *Proc. 50th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 171–180, 2009. 5, 9
- [DHK⁺10] I. Diakonikolas, P. Harsha, A. Klivans, R. Meka, P. Raghavendra, R. Servedio, and L.-Y. Tan. Bounding the average sensitivity and noise sensitivity of polynomial threshold functions. In *Proc. 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 533–542, 2010. 1
- [DRST09] I. Diakonikolas, P. Raghavendra, R. Servedio, and L.-Y. Tan. Average sensitivity and noise sensitivity of polynomial threshold functions, 2009. Available at <http://arxiv.org/abs/0909.5011>. 1, 4
- [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. 9
- [DSTW10] I. Diakonikolas, R. Servedio, L.-Y. Tan, and A. Wan. A regularity lemma, and low-weight approximators, for low-degree polynomial threshold functions. In *Proc. 25th Annual IEEE Conference on Computational Complexity (CCC)*, pages 211–222, 2010. 4
- [Fel68] W. Feller. *An introduction to probability theory and its applications*. John Wiley & Sons, 1968. 21
- [Fri98] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):474–483, 1998. 1, 2
- [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. 1, 2, 7
- [GS10] P. Gopalan and R. Servedio. Learning and lower bounds for AC^0 with threshold gates. In *Proc. 14th Intl. Workshop on Randomization and Computation (RANDOM)*, pages 588–601, 2010. 4, 19
- [HKM09] P. Harsha, A. Klivans, and R. Meka. Bounding the sensitivity of polynomial threshold functions. Available at <http://arxiv.org/abs/0909.5175>, 2009. 1, 4
- [Jan97] S. Janson. *Gaussian Hilbert Spaces*. Cambridge University Press, Cambridge, UK, 1997. 1
- [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
- [Kan10] D.M. Kane. The Gaussian surface area and noise sensitivity of degree-d polynomial threshold functions. In *Proc. 25th Annual IEEE Conference on Computational Complexity (CCC)*, pages 205–210, 2010. 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, 2
- [KKMS08] A. Kalai, A. Klivans, Y. Mansour, and R. Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008. 1, 3, 19
- [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, 19
- [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
- [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
- [MOO10] E. Mossel, R. O’Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, 171:295–341, 2010.

- 1, 4, 5, 6, 7
- [OS07] R. O'Donnell and R. Servedio. Learning monotone decision trees in polynomial time. *SIAM Journal on Computing*, 37(3):827–844, 2007. 1, 2
 - [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. 9
 - [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. Available at <http://arxiv.org/abs/math/0412377>. 17
 - [Ser07] R. Servedio. Every linear threshold function has a low-weight approximator. *Computational Complexity*, 16(2):180–209, 2007. 1, 4, 5, 8, 9
 - [Shi00] Y. Shi. Lower bounds of quantum black-box complexity and degree of approximating polynomials by influence of boolean variables. *Information Processing Letters*, 75(1-2):79–83, 2000. 1, 2
 - [SS] O. Schramm and J. Steif. Quantitative noise sensitivity and exceptional times for percolation. *Annals of Mathematics*, to appear. 2
 - [SSS09] S. Shalev Shwartz, O. Shamir, and K. Sridharan. Agnostically learning halfspaces with margin errors. TTI Technical Report, 2009. 3