

Agnostic learning in (almost) optimal time via Gaussian surface area

Lucas Pesenti*
lucas.pesenti@inf.ethz.ch
ETH Zurich

Lucas Slot†
l.f.h.slot@uva.nl
University of Amsterdam

Manuel Wiedmer*
manuel.wiedmer@inf.ethz.ch
ETH Zurich

March 6, 2026

Abstract

The complexity of learning a concept class under Gaussian marginals in the difficult agnostic model is closely related to its L_1 -approximability by low-degree polynomials. For any concept class with *Gaussian surface area* at most Γ , Klivans et al. (2008) show that degree $d = O(\Gamma^2/\varepsilon^4)$ suffices to achieve an ε -approximation. This leads to the best-known bounds on the complexity of learning a variety of concept classes. In this note, we improve their analysis by showing that degree $d = \tilde{O}(\Gamma^2/\varepsilon^2)$ is enough. In light of lower bounds due to Diakonikolas et al. (2021), this yields (near) optimal bounds on the complexity of agnostically learning polynomial threshold functions in the statistical query model. Our proof relies on a direct analogue of a construction of Feldman et al. (2020), who considered L_1 -approximation on the Boolean hypercube.

1 Introduction

The *agnostic learning* framework of [KSS92] extends the classical PAC model of [Val84] to learning tasks with noisy input data, and has served as one of the standard models for studying the computational complexity of such tasks. It can be described as follows: Let C be a *concept class* (i.e., a set of functions $f : X \rightarrow \{-1, 1\}$) and let $\mathcal{D}$ be a distribution on $X \times \{-1, 1\}$. Given labeled examples (x, y) drawn independently from $\mathcal{D}$, our task is to output a *hypothesis* $\hat{f}$ that predicts the labels of future examples drawn from $\mathcal{D}$.

Contrary to the PAC model, we do not assume any relation between the examples x and their label y . Thus, it is generally impossible to find a hypothesis with nontrivial success rate. Instead, we ask that $\hat{f}$ predicts labels almost as well as the best concept in C . More formally, we say that an algorithm *agnostically learns* C up to excess error $\varepsilon > 0$ if it produces (with high probability) a hypothesis $\hat{f}$ with $\mathbb{P}_{(x,y) \sim \mathcal{D}}(\hat{f}(x) \neq y) \leq \text{opt} + \varepsilon$, where $\text{opt} := \min_{f \in C} \mathbb{P}_{(x,y) \sim \mathcal{D}}(f(x) \neq y)$. The parameter opt measures the “noisiness” of the problem; if $\text{opt} = 0$, we recover the PAC model.

The L_1 -polynomial regression algorithm

Efficient algorithms for agnostic learning under general distributions are unlikely to exist, already for simple concept classes like halfspaces [Dan16, Tie23]. It is therefore natural to restrict to certain

*Supported by the Swiss National Science Foundation (SNSF), grant no. 10004947.

†Supported by the project “Foundations of sum-of-squares algorithms: worst-case and beyond” with file number VI.Veni.242.234 of the research programme Veni ENW which is (partly) financed by the Dutch Research Council (NWO) under the grant <https://doi.org/10.61686/GIEOW65108>.

classes of distributions. Here, we focus on the setting where $\mathcal{X} = \mathbb{R}^n$ and the marginal distribution $\mathcal{D}_X = \mathcal{N}(0, I_n)$ on the examples is the standard Gaussian.

In their seminal work, [KKMS08] give the first efficient algorithm for agnostically learning halfspaces under the Gaussian distribution (among other distributions). Their algorithm computes a best degree- d polynomial approximation of the labeled examples, as measured in the L_1 -norm. This can be done in time $n^{O(d)}$, e.g., via linear programming. Assuming that each concept in C admits a degree- d approximation with L_1 -error at most ε , this procedure (combined with a thresholding scheme) yields an agnostic learner for C with excess error ε .¹ If $d = d(\varepsilon)$ can be chosen independently of the dimension n , this yields a PTAS with runtime $n^{O(d)} \cdot \text{poly}(1/\varepsilon)$.

The L_1 -polynomial regression algorithm remains the standard (and de facto only) choice for efficient agnostic learning under distributional assumptions. It achieves the best-known runtime for a variety of concept classes and distributions [KKMS08, KOS08, DKN10, DHK⁺10, BOW10, Kan11a, KKM13, DKR25, CPS26]. In fact, L_1 -polynomial regression is known to be (essentially) optimal in the Statistical Query (SQ) model for agnostic learning under the Gaussian distribution [DKPZ21].

Degree bounds for L_1 -approximation and Gaussian surface area

The results of [KKMS08] and [DKPZ21] show that the complexity of agnostically learning a concept class under the Gaussian distribution is governed by its approximability by low-degree polynomials in L_1 : if d is the smallest degree such that each $f \in C$ admits an ε -approximation in L_1 of degree d , then the SQ-complexity of agnostically learning C up to error ε is approximately $n^{\Theta(d)}$.

For halfspaces, [KKMS08] show that degree $d = O(1/\varepsilon^4)$ suffices to achieve an ε -approximation in L_1 . Their proof first establishes the bound on the degree required to ε -approximate a halfspace in L_2 -norm; this can then be converted to an L_1 -guarantee using Cauchy-Schwarz. The reason for taking this indirect route is the availability of powerful tools for analyzing approximation in L_2 that are unavailable in the L_1 -setting, most notably Hermite analysis.

To generalize this approach to arbitrary concept classes, [KOS08] consider the *Gaussian surface area* (GSA) of a concept f , which measures the surface area of the set $\{x \in \mathbb{R}^n : f(x) = 1\}$ relative to the Gaussian density function. Using Hermite analysis, they show that any concept class whose elements have GSA at most Γ can be ε -approximated in L_2 -norm by a polynomial of degree $d = O(\Gamma^2/\varepsilon^4)$; thus, $d = O(\Gamma^2/\varepsilon^4)$ also suffices in L_1 . Halfspaces, Euclidean balls and certain cones have constant GSA [KOS08]; intersections of k halfspaces have GSA at most $O(\sqrt{\log k})$ [KOS08]; polynomial threshold functions (PTFs) of degree k have GSA at most $O(k)$ [Kan11a]; convex sets have GSA at most $O(n^{1/4})$ [Bal93].

Optimal degree bounds for L_1 -approximation

A key appeal of the approach of [KOS08] is that it establishes GSA as a universal upper bound on the complexity of agnostic learning. However, their analysis yields a suboptimal guarantee for approximating (and thus for learning) halfspaces: A direct construction of [DKN10] shows that one only needs degree $d = O(1/\varepsilon^2)$ to get an ε -approximation for halfspaces in the L_1 -norm (vs. $d = O(1/\varepsilon^4)$ obtained in [KKMS08, KOS08]).

The bound of [DKN10] is optimal both in the sense that no better L_1 -approximation is possible [Gan08], and that the corresponding runtime of the L_1 -polynomial regression algorithm

¹The choice of L_1 -norm here is crucial: earlier algorithms that compute a best L_2 -approximation yield a hypothesis with (total) error $O(\text{opt}) + \varepsilon$. This is a weaker guarantee, tolerating only small amounts of noise. For algorithmic results in this regime, see for example [KLS09, ABL14, Dan15, DKS18, DKTZ20].

matches the lower bound of [DKPZ21] in the SQ-model (up to polynomial factors in $1/\varepsilon$). On the other hand, their construction does not (obviously) extend beyond halfspaces (see Section 4.2).

This raises the question whether it's possible to prove an approximation guarantee that is general (applies to all concepts of bounded GSA) and simultaneously optimal (in terms of its dependence on ε). In this work, we (almost) answer this question in the positive: We improve the analysis of [KOS08] by showing that degree $d = \tilde{O}(\Gamma^2/\varepsilon^2)$ suffices to achieve an ε -approximation of any concept class with GSA at most Γ (Theorem 1.1). This recovers the (optimal) bound of [DKN10] for halfspaces (up to a $\log(1/\varepsilon)$ -factor). Moreover, it improves the best-known bound for degree- k PTFs from $O(k^2/\varepsilon^4)$ to $\tilde{O}(k^2/\varepsilon^2)$, which (nearly) matches the lower bound of $\Omega(k^2/\varepsilon^2)$ from [DKPZ21]. Finally, it improves the best-known bounds for intersections of halfspaces and convex sets by a factor $\tilde{\Omega}(1/\varepsilon^2)$. See Table 1 for an overview.

Our technical contributions are minor: We show that a construction of [FKV20] involving the noise operator on the Boolean hypercube can be transported to the Gaussian case (mutatis mutandis). Then, we use results of [KOS08] to bound the resulting L_1 -approximation errors in terms of the Gaussian surface area and conclude. Thus, all ingredients of our proof were (more or less) known in the literature, but (to our knowledge) had not been assembled yet in this context.

concept class	LB	UB (previous)	UB (Thm. 1.1)
halfspaces	$\Omega(1/\varepsilon^2)$ [Gan08]	$O(1/\varepsilon^2)$ [DKN10]	$\tilde{O}(1/\varepsilon^2)$
degree- k PTFs	$\Omega(k^2/\varepsilon^2)$ [DKPZ21]	$O(k^2/\varepsilon^4)$ [Kan11a]	$\tilde{O}(k^2/\varepsilon^2)$
intersec. k halfspaces	$\tilde{\Omega}(\sqrt{\log k}/\varepsilon)$ [DKPZ21]	$O(\log k/\varepsilon^4)$ [KOS08]	$\tilde{O}(\log k/\varepsilon^2)$
convex sets	—	$O(\sqrt{n}/\varepsilon^4)$ [KOS08]	$\tilde{O}(\sqrt{n}/\varepsilon^2)$
GSA at most Γ	—	$O(\Gamma^2/\varepsilon^4)$ [KOS08]	$\tilde{O}(\Gamma^2/\varepsilon^2)$

Table 1: Upper bounds (UB) and lower bounds (LB) on the smallest degree needed to achieve an ε -approximation in $L_1(\mathcal{N}_n)$ for concept classes $\mathcal{C}$. These correspond to upper and lower bounds in $n^{O(d)}$ and $n^{\Omega(d)}$ on the (SQ-)complexity of agnostically learning $\mathcal{C}$ up to error ε , respectively.

1.1 Main results

To state our main results, we need two (standard) notions related to Boolean functions. First, the *Gaussian noise sensitivity* $\text{GNS}_\rho(f)$ of $f : \mathbb{R}^n \rightarrow \{\pm 1\}$ at level $\rho \in [0, 1]$ is the probability that $f(X) \neq f(Y)$ for two $(1 - \rho)$ -correlated Gaussian random variables $X, Y \sim \mathcal{N}_n$ (see Definition 2.2). Second, writing $K(f) = \{x \in \mathbb{R}^n : f(x) = 1\}$, the *Gaussian surface area* of f is²

$$\text{GSA}(f) := \lim_{\delta \rightarrow 0} \frac{\text{vol}_{\mathcal{N}_n}(K(f)_\delta \setminus K(f))}{\delta},$$

where $A_\delta := \{x \in \mathbb{R}^n : \text{dist}(x, A) \leq \delta\}$ is the δ -neighborhood of $A \subseteq \mathbb{R}^n$, and $\text{vol}_{\mathcal{N}_n}(\cdot)$ denotes the Gaussian probability mass (see Definition 2.4).

Our main result is the following L_1 -approximation guarantee for f in terms of its GSA.

Theorem 1.1. *Let $f : \mathbb{R}^n \rightarrow \{\pm 1\}$ be a (measurable) function. For every $\varepsilon > 0$, there exists a polynomial p of degree $d \leq O(\log(1/\varepsilon) \cdot \text{GSA}(f)^2/\varepsilon^2)$ with $\mathbb{E}_{x \sim \mathcal{N}_n} [|f(x) - p(x)|] \leq \varepsilon$.*

²In the definition of [KOS08], Gaussian surface area is a property of a subset of $\mathbb{R}^n$. We think of a Boolean function f as the indicator of the subset $K(f) = \{x : f(x) = 1\} \subseteq \mathbb{R}^n$, which allows us to assign a surface area to f itself.

As an intermediate result, we get the following bound in terms of the noise sensitivity of f .

Proposition 1.2. *Let $f : \mathbb{R}^n \rightarrow \{\pm 1\}$ be a (measurable) function. For each $\rho \in [0, 1]$ and $d \in \mathbb{N}$, there is a polynomial p of degree d such that*

$$\mathbb{E}_{x \sim \mathcal{N}_n} [|f(x) - p(x)|] \leq 2\text{GNS}_{1-\rho}(f) + \rho^{d+1}.$$

Together with [KKMS08, Theorem 5] and bounds on the Gaussian surface area from [Kan11a, KOS08], we immediately get the following corollary from Theorem 1.1.

Corollary 1.3. *Let C be a concept class of Boolean functions $f : \mathbb{R}^n \rightarrow \{\pm 1\}$. Let Γ be an upper bound on the GSA of every $f \in C$. Let $\varepsilon > 0$. Then, the L_1 -polynomial regression algorithm agnostically learns C up to error ε in time and sample complexity $n^{\tilde{O}(\Gamma^2/\varepsilon^2)} \cdot \text{poly}(1/\varepsilon)$. In particular, the L_1 polynomial regression algorithm agnostically learns*

- a) degree- k PTFs in time $n^{\tilde{O}(k^2/\varepsilon^2)}$;
- b) intersections of k halfspaces in time $n^{\tilde{O}(\log(k)/\varepsilon^2)}$;
- c) convex sets in time $n^{\tilde{O}(\sqrt{n}/\varepsilon^2)}$.

Lower bounds on the L_1 -approximation degree of concepts

As mentioned, [DKPZ21] show that the SQ-complexity of learning a concept class C up to error ε is approximately $n^{\Omega(d)}$, where $d = d(\varepsilon)$ is the smallest polynomial degree required to ε -approximate any element of C in L_1 . Complementing this result, they establish lower bounds on this degree: For PTFs of degree k , they show that $d = \Omega(k^2/\varepsilon^2)$. The case $k = 1$ corresponds to halfspaces, thus requiring $d = \Omega(1/\varepsilon^2)$.

They also prove the following general bound in terms of noise sensitivity:

Theorem 1.4 ([DKPZ21]). *Let $f : \mathbb{R}^n \rightarrow \{\pm 1\}$ be a (measurable) function. For every polynomial p of degree d we have that*

$$\mathbb{E}_{x \sim \mathcal{N}_n} [|f(x) - p(x)|] \geq \Omega(1/\log d) \cdot \text{GNS}_{(\log d/d)^2}(f)$$

As an application, Theorem 1.4 shows that $d = \tilde{\Omega}(\sqrt{\log k}/\varepsilon)$ is required to approximate an intersection of k halfspaces. Note that Theorem 1.4 is typically not tight, and does not match the upper bound of Proposition 1.2. For example, when f is a halfspace, we have $\text{GNS}_\delta(f) \approx \sqrt{\delta}$. Then, Theorem 1.4 gives a degree bound $d = \tilde{\Omega}(1/\varepsilon)$, whereas Proposition 1.2 shows that $d = \tilde{O}(1/\varepsilon^2)$. The correct degree is $d = \Theta(1/\varepsilon^2)$ [Gan08, DKN10].

Outline

The rest of the paper is organized as follows: In Section 2, we give some preliminaries on Hermite analysis and the Ornstein-Uhlenbeck (noise) operator; in Section 3, we give a proof of our main result Theorem 1.1; in Section 4 we compare our construction to those of [KKMS08, DKN10, FKV20].

2 Hermite analysis and the Ornstein-Uhlenbeck operator

In this section, we recall some basic facts on Hermite analysis and the Ornstein-Uhlenbeck operator. See [LT91, Bak94, Jan97] for general references, and [O'D14, Chapter 11] for a nice overview.

2.1 Hermite analysis

We write $\mathcal{N}_n = \mathcal{N}(0, I_n)$ for the standard Gaussian distribution on $\mathbb{R}^n$. Throughout, we assume that all functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ are (Borel-)measurable. We denote by $L_2(\mathcal{N}_n)$ the space of functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with $\|f\|_{L_2(\mathcal{N}_n)}^2 := \mathbb{E}_{X \sim \mathcal{N}_n}[f(X)^2] < \infty$. It has an inner product given by $\langle f, g \rangle_{\mathcal{N}_n} := \mathbb{E}_{X \sim \mathcal{N}_n}[f(X)g(X)]$. Similarly, $L_1(\mathcal{N}_n)$ denotes the space of functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with $\|f\|_{L_1(\mathcal{N}_n)} := \mathbb{E}_{X \sim \mathcal{N}_n}[|f(X)|] < \infty$. We may use Cauchy-Schwarz to bound

$$\|f\|_{L_1(\mathcal{N}_n)} \leq \|f\|_{L_2(\mathcal{N}_n)}. \quad (1)$$

For $n = 1$, the space $L_2(\mathcal{N}_1)$ has an orthonormal basis of *Hermite polynomials* $(H_k)_{k \geq 0} \subseteq \mathbb{R}[x]$, which are the unique polynomials with positive leading coefficient such that H_k has degree k and

$$\langle H_i, H_j \rangle_{\mathcal{N}_1} = \mathbb{E}_{X \sim \mathcal{N}_1}[H_i(X)H_j(X)] = \delta_{ij} \quad (\forall i, j \in \mathbb{N}).$$

This definition corresponds to the (monic) *probabilist's Hermite polynomials* scaled by a factor $1/\sqrt{k!}$. For $n > 1$, we have an orthonormal basis of *multivariate Hermite polynomials*, given by

$$H_\alpha(x) := H_{\alpha_1}(x_1) \cdot H_{\alpha_2}(x_2) \cdot \dots \cdot H_{\alpha_n}(x_n) \quad (\alpha \in \mathbb{N}^n). \quad (2)$$

Note that H_α is of degree $|\alpha| := \sum_{i=1}^n \alpha_i$. Any $f \in L_2(\mathcal{N}_n)$ has a (unique) *Hermite expansion*

$$f = \sum_{\alpha \in \mathbb{N}^n} \hat{f}(\alpha) \cdot H_\alpha, \quad \text{with} \quad \hat{f}(\alpha) \in \mathbb{R}.$$

We refer to the $\hat{f}(\alpha)$ as the *Hermite coefficients* of f . By orthonormality of the Hermite polynomials, we have that $\|f\|_{L_2(\mathcal{N}_n)}^2 = \sum_{\alpha} \hat{f}(\alpha)^2$ (Parseval's identity). For $d \in \mathbb{N}$, we write $\Pi_d : L_2(\mathcal{N}_n) \rightarrow L_2(\mathcal{N}_n)$ for the orthogonal projection onto $\text{span}\{H_\alpha : |\alpha| \leq d\}$. That is,

$$\Pi_d f = \sum_{|\alpha| \leq d} \hat{f}(\alpha) \cdot H_\alpha.$$

By definition, $\Pi_d f$ is the best L_2 -approximation of f by a degree- d polynomial. It is also referred to as the degree- d *Hermite expansion* of f . Its approximation error can be expressed in terms of the Hermite coefficients of f as follows:

$$\|f - \Pi_d f\|_{L_2(\mathcal{N}_n)}^2 = \left\| \sum_{|\alpha| > d} \hat{f}(\alpha) \cdot H_\alpha \right\|_{L_2(\mathcal{N}_n)}^2 = \sum_{|\alpha| > d} \hat{f}(\alpha)^2. \quad (3)$$

2.2 The Ornstein-Uhlenbeck operator

For $\rho \in [0, 1]$, the *Ornstein-Uhlenbeck operator* $T_\rho : L_2(\mathcal{N}_n) \rightarrow L_2(\mathcal{N}_n)$ is defined by:

$$T_\rho f(x) = \mathbb{E}_{Y \sim \mathcal{N}_n} \left[f \left(\rho x + \sqrt{1 - \rho^2} Y \right) \right].$$

By construction, T_ρ is a linear operator. It is diagonalized by the Hermite polynomials; namely

$$T_\rho H_\alpha = \rho^{|\alpha|} H_\alpha \quad (\forall \alpha \in \mathbb{N}^n). \quad (4)$$

For completeness, we include a proof of this fact in [Lemma A.2](#). As a direct consequence, the Hermite coefficients of $T_\rho f$ decay exponentially fast for any $f \in L_2(\mathcal{N}_n)$. In light of (3), this means that $T_\rho f$ is well-approximated in $L_2(\mathcal{N}_n)$ by low-degree polynomials:

Lemma 2.1. Let $f \in L_2(\mathcal{N}_n)$. For any $\rho \in [0, 1]$ and $d \in \mathbb{N}$, we have

$$\|T_\rho f - \Pi_{\leq d}(T_\rho f)\|_{L_2(\mathcal{N}_n)}^2 \leq \rho^{2d+2} \cdot \|f\|_{L_2(\mathcal{N}_n)}^2.$$

Proof. Writing $f(x) = \sum_\alpha \hat{f}(\alpha) H_\alpha(x)$, and using (4), we find that

$$T_\rho f(x) = \sum_{\alpha \in \mathbb{N}^n} \hat{f}(\alpha) \cdot T_\rho H_\alpha(x) = \sum_{\alpha \in \mathbb{N}^n} \rho^{|\alpha|} \hat{f}(\alpha) \cdot H_\alpha(x).$$

Thus, by (3) we get that

$$\|T_\rho f - \Pi_{\leq d}(T_\rho f)\|_{L_2(\mathcal{N}_n)}^2 = \sum_{|\alpha| > d} \hat{f}(\alpha)^2 \rho^{2|\alpha|} \leq \rho^{2(d+1)} \cdot \sum_{|\alpha| > d} \hat{f}(\alpha)^2 \leq \rho^{2(d+1)} \cdot \|f\|_{L_2(\mathcal{N}_n)}^2. \quad \square$$

2.3 Gaussian noise sensitivity and Gaussian surface area

For $f : \mathbb{R}^n \rightarrow \{\pm 1\}$, the expected difference between f and $T_\rho f$ is equal to twice the *Gaussian noise sensitivity* of f . (In fact, this is sometimes used as the definition of noise sensitivity.) For completeness, we include a proof of this fact in [Section A.2](#).

Definition 2.2 (cf. [\[KOS08, Def. 12\]](#)). The *Gaussian noise sensitivity* of $f : \mathbb{R}^n \rightarrow \{\pm 1\}$ at $\delta \in [0, 1]$ is

$$\text{GNS}_\delta(f) := \mathbb{P}_{(X,Y)}[f(X) \neq f(Y)],$$

where (X, Y) are two $(1 - \delta)$ -correlated standard Gaussians, that is

$$(X, Y) \sim \mathcal{N}\left(\begin{pmatrix} 0_n \\ 0_n \end{pmatrix}, \begin{pmatrix} I_n & (1 - \delta)I_n \\ (1 - \delta)I_n & I_n \end{pmatrix}\right).$$

Lemma 2.3 (see [Section A.2](#)). Let $f : \mathbb{R}^n \rightarrow \{\pm 1\}$. For any $\rho \in [0, 1]$, we have

$$\|f - T_\rho f\|_{L_1(\mathcal{N}_n)} = 2\text{GNS}_{1-\rho}(f).$$

In turn, the noise sensitivity of a function can be bounded in terms of its *Gaussian surface area*. This is due to [\[KOS08\]](#), who prove it using a result of [\[Led94\]](#).

Definition 2.4 (cf. [\[KOS08, Def. 1\]](#)). Let $K \subseteq \mathbb{R}^n$ be a Borel set. Its Gaussian surface area is

$$\text{GSA}(K) := \lim_{\delta \rightarrow 0} \frac{\text{vol}_{\mathcal{N}_n}(K_\delta \setminus K)}{\delta},$$

where $K_\delta := \{x \in \mathbb{R}^n : \text{dist}(x, K) \leq \delta\}$ is the δ -neighborhood of K , and $\text{vol}_{\mathcal{N}_n}(\cdot)$ denotes the Gaussian probability mass.

For $f : \mathbb{R}^n \rightarrow \{\pm 1\}$, we write $K(f) = \{x \in \mathbb{R}^n : f(x) = 1\}$, and set $\text{GSA}(f) := \text{GSA}(K(f))$. We always assume as in [\[KOS08\]](#) that $\text{vol}_{\mathcal{N}_n}(\delta K(f)) = 0$, where $\delta K(f)$ denotes the boundary of $K(f)$.

Lemma 2.5 ([\[KOS08, Corollary 14\]](#)). Let $f : \mathbb{R}^n \rightarrow \{\pm 1\}$. For each $\rho \in [0, 1]$, we have

$$\text{GNS}_{1-\rho}(f) \leq \sqrt{\pi} \sqrt{1 - \rho} \cdot \text{GSA}(f).$$

3 Proof of Theorem 1.1

Our proof of Theorem 1.1 proceeds in two steps. First, we approximate f by $T_\rho f$, for some $\rho \in [0, 1]$ to be chosen later. Then, we approximate $T_\rho f$ by its low-degree Hermite expansion $\Pi_d(T_\rho f)$. Using the triangle inequality, the total approximation error is at most

$$\|f - \Pi_d(T_\rho f)\|_{L_1(\mathcal{N}_n)} \leq \|f - T_\rho f\|_{L_1(\mathcal{N}_n)} + \|T_\rho f - \Pi_d(T_\rho f)\|_{L_1(\mathcal{N}_n)}.$$

The first term on the RHS can be bounded directly in terms of the noise sensitivity of f (Lemma 2.3). For the second term, note that, by (1) and Lemma 2.1,

$$\|T_\rho f - \Pi_d(T_\rho f)\|_{L_1(\mathcal{N}_n)} \leq \|T_\rho f - \Pi_d(T_\rho f)\|_{L_2(\mathcal{N}_n)} \leq \rho^{d+1}.$$

Together, this yields:

Proposition 1.2. *Let $f : \mathbb{R}^n \rightarrow \{\pm 1\}$ be a (measurable) function. For each $\rho \in [0, 1]$ and $d \in \mathbb{N}$, there is a polynomial p of degree d such that*

$$\mathbb{E}_{x \sim \mathcal{N}_n} [|f(x) - p(x)|] \leq 2\text{GNS}_{1-\rho}(f) + \rho^{d+1}.$$

Proof of Theorem 1.1. Let p be the degree- d polynomial of Proposition 1.2. By Lemma 2.5, we can further bound

$$\|f - p\|_{L_1(\mathcal{N}_n)} \leq 2\sqrt{\pi}\sqrt{1-\rho} \cdot \text{GSA}(f) + \rho^{d+1}. \quad (5)$$

Given $\varepsilon > 0$, we want to ensure that both terms on the RHS of (5) are at most $\varepsilon/2$. For the first term, it suffices to set

$$\rho = \max\left(0, 1 - \frac{\varepsilon^2}{16\pi\text{GSA}(f)^2}\right).$$

Now, given this choice of ρ , we pick $d \in \mathbb{N}$ such that $\rho^{d+1} \leq \varepsilon/2$. Without loss of generality, assume that $\rho > 0$ (otherwise, $d = 0$ suffices). We have

$$\begin{aligned} \rho^{d+1} \leq \frac{\varepsilon}{2} &\iff (d+1)\log(\rho) \leq \log(\varepsilon/2) \\ &\iff (d+1)(\rho-1) \leq \log(\varepsilon/2) \quad (\text{as } d \geq 0 \text{ and } \log(t) \leq t-1 \text{ for all } t > 0) \\ &\iff (d+1) \geq \frac{\log(\varepsilon/2)}{\rho-1} \quad (\text{as } \rho-1 < 0) \\ &\iff d \geq \frac{\log(2/\varepsilon)}{1-\rho} - 1. \end{aligned}$$

Plugging in the value of $1-\rho$, it therefore suffices to set

$$d = 16\pi \cdot \text{GSA}(f)^2 \cdot \frac{\log(2/\varepsilon)}{\varepsilon^2} - 1. \quad \square$$

4 Comparison with previous work

In this section, we compare our proof of Theorem 1.1 to three prior works on agnostic learning that established polynomial approximation results in L_1 . These differ by their choice of low-degree approximation for f :

1. the truncated Hermite expansion of f [KOS08];
2. the truncated Hermite expansion of a smoothed version of f , obtained by convolving with a carefully chosen function [DKN10];³
3. the truncated Fourier expansion of a smoothed version of f , obtained by applying the Boolean noise operator [FKV20].

4.1 Comparison to [KOS08]

Building on the argument of [KKMS08] for halfspaces, [KOS08] shows more generally that the degree- d Hermite truncation of f (denoted $\Pi_d f$) satisfies $\|f - \Pi_d f\|_{L_1(\mathcal{N}_n)} \leq \varepsilon$ for $d = O(\text{GSA}(f)^2 / \varepsilon^4)$. Crucially, their proof reduces the desired L_1 -approximation to an L_2 -bound:

$$\begin{aligned} \|f - \Pi_d f\|_{L_1(\mathcal{N}_n)} &\leq \|f - \Pi_d f\|_{L_2(\mathcal{N}_n)} \\ &\leq O(1) \cdot \text{GNS}_{1/d}(f)^{1/2} && \text{(by [KOS02, Prop. 16])} \\ &\leq O(1) \cdot d^{-1/4} \cdot \text{GSA}(f)^{1/2}. && \text{(by Lemma 2.5)} \end{aligned}$$

The gap between the bound obtained in this argument and the guarantees of Theorem 1.1 can be entirely explained by the first step: even for origin-centered halfspaces, the L_2 -error of the degree- d Hermite truncation is $\Omega(d^{-1/4})$ [KKMS08]. As $\Pi_d f$ is the best degree- d approximation to f in L_2 , this shows that the strategy of reducing the whole problem to an L_2 question is inherently lossy.

However, we are not aware of any example showing that $\Pi_d f$ fails to match the guarantees of Theorem 1.1 in L_1 . On the positive side, we show that $\Pi_d f$ does achieve these guarantees for origin-centered halfspaces. After reducing to $n = 1$, it suffices to show the following:

Proposition 4.1. *Let $\text{sign} : \mathbb{R} \rightarrow \{\pm 1\}$ denote the sign function. We have that*

$$\|\text{sign} - \Pi_d \text{sign}\|_{L_1(\mathcal{N}_1)} \leq O\left(\frac{\log d}{\sqrt{d}}\right).$$

Proof. See Section A.4. □

The upshot is that, at least for origin-centered halfspaces, the suboptimal degree bound of [KKMS08, KOS08] with respect to Theorem 1.1 is not due to the choice of approximating polynomial ($\Pi_d f$ vs. $\Pi_d T_\rho f$), but to a loose application of Cauchy-Schwarz.

4.2 Comparison to [DKN10]

[DKN10] introduces a different technique for constructing low-degree approximations, showing in particular that degree $d = O(1/\varepsilon^2)$ suffices to approximate halfspaces in L_1 . In this special case, this improves on the guarantees of Theorem 1.1 by a logarithmic factor.

In light of our proof, their method can be motivated as follows. First, for a halfspace $f(x) = \text{sign}(c - \langle w, x \rangle)$, by projecting onto $\text{span}(w)$ we may assume without loss of generality that $n = 1$. We now identify what properties we would like from a smoothing g of f in order to

³[DKN10] actually consider the Taylor approximation of this smoothed version of f . They are concerned with a stronger notion of approximation than the one considered here (related to proving universality), for which this is important. In our context, we can equivalently consider the truncated Hermite expansion.

implement the strategy from [Section 3](#). By Gaussian integration by parts, the Hermite coefficients bounded in [Lemma 2.1](#) can be interpreted as (cf. [Lemma A.4](#))

$$\langle g, h_k \rangle = \frac{\mathbb{E}_{x \sim \mathcal{N}(0,1)} g^{(k)}(x)}{\sqrt{k!}}. \quad (6)$$

Hence, we would like g to approximate f in L_1 (to replace [Lemma 2.3](#)), while keeping the derivatives of g bounded. In the one-dimensional case, the authors of [\[DKN10\]](#) show that this is possible in a strong sense: for any fixed $\varepsilon > 0$, one can choose g so that (a) $\|g - f\|_{L_1(\mathcal{N})} \leq \varepsilon$, and (b) $\|g^{(k)}\|_\infty \leq O(1/\varepsilon)^k$ for all $k \geq 1$. In the proof, g is defined from f by convolving with a suitable function whose Fourier transform has bounded support. Combined with (6) and following the argument from [Section 3](#), (a) and (b) readily imply that $\|\Pi_d g - f\|_1 \leq \varepsilon$ for $d = O(1/\varepsilon^2)$, as desired.

The main difficulty with this construction lies in extending it to higher dimensions. For agnostic learning over general concept classes with bounded Gaussian surface area, repeating the previous argument appears to incur a dimension-dependent factor. In the special case of polynomial threshold functions, [\[DKN10, Kan11b\]](#) show that one can still reduce to the low-dimensional case, at the cost of a (much) worse dependence on ε .

4.3 Comparison to [\[FKV20\]](#)

The construction in [\[FKV20\]](#) is the Boolean analogue of ours. Indeed, [\[FKV20, Lemma 3.1\]](#) shows that for any Boolean function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ and $\rho \in (0, 1]$, there exists a polynomial p of degree $d = O(\log(1/\varepsilon)/(1 - \rho))$ such that

$$\|f - p\|_{L_1(\mathcal{U}(\{0,1\}^n))} \leq 2\text{BNS}_{1-\rho}(f) + \rho^{d+1} \cdot \|f\|_{L_2(\mathcal{U}(\{0,1\}^n))}, \quad (7)$$

where $\mathcal{U}(\{0, 1\}^n)$ is the uniform distribution on $\{0, 1\}^n$ and BNS is the Boolean noise sensitivity.

Equation (7) is the Boolean analogue of [Proposition 1.2](#). In fact, the proof of (7) in [\[FKV20\]](#) closely parallels our proof of [Proposition 1.2](#). One first approximates f by $T_\rho f$, where T_ρ is the noise operator on the hypercube; the resulting L_1 error is controlled directly by the Boolean noise sensitivity. One then approximates $T_\rho f$ by its degree- d Fourier truncation, and bounds the L_1 error of this truncation by essentially the same argument as in the Gaussian setting.

Acknowledgments

We thank Piotr Buys for a useful discussion at an early stage of this work.

References

- [ABL14] Pranjali Awasthi, Maria Florina Balcan, and Philip M. Long. The power of localization for efficiently learning linear separators with noise. In *Proceedings of the ACM Symposium on Theory of Computing (STOC 2014)*, pages 449–458. ACM, New York, 2014. [2](#)
- [Bak94] Dominique Bakry. *L’hypercontractivité et son utilisation en théorie des semigroupes*, volume 1581 of *Lecture Notes in Math.*, pages 1–114. Springer, Berlin, 1994. [4, 12](#)
- [Bal93] Keith Ball. The reverse isoperimetric problem for Gaussian measure. *Discrete & Computational Geometry*, 10(4):411–420, 1993. [2](#)

- [BOW10] Eric Blais, Ryan O’Donnell, and Karl Wimmer. Polynomial regression under arbitrary product distributions. *Machine Learning*, 80(2-3):273–294, 2010. [2](#)
- [CPS26] Xi Chen, Shyamal Patel, and Rocco A. Servedio. A mysterious connection between tolerant junta testing and agnostically learning conjunctions. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*, 2026. To appear. [2](#)
- [Dan15] Amit Daniely. A PTAS for agnostically learning halfspaces. In *Proceedings of The 28th Conference on Learning Theory (COLT 2015)*, volume 40, pages 484–502, Paris, France, 2015. PMLR. [2](#)
- [Dan16] Amit Daniely. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2016)*, pages 105–117. ACM, New York, 2016. [1](#)
- [DHK⁺10] Ilias Diakonikolas, Prahladh Harsha, Adam Klivans, Raghu Meka, Prasad Raghavendra, Rocco A. Servedio, and Li-Yang Tan. Bounding the average sensitivity and noise sensitivity of polynomial threshold functions. In *Proceedings of the 2010 ACM International Symposium on Theory of Computing (STOC 2010)*, pages 533–542. ACM, New York, 2010. [2](#)
- [DKN10] Ilias Diakonikolas, Daniel M. Kane, and Jelani Nelson. Bounded independence fools degree-2 threshold functions. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS 2010)*, pages 11–20. IEEE Computer Soc., Los Alamitos, CA, 2010. [2](#), [3](#), [4](#), [8](#), [9](#)
- [DKPZ21] Ilias Diakonikolas, Daniel M. Kane, Thanasis Pittas, and Nikos Zarifis. The Optimality of Polynomial Regression for Agnostic Learning under Gaussian Marginals in the SQ Model. In *Proceedings of Thirty Fourth Conference on Learning Theory (COLT 2021)*, volume 134, pages 1552–1584. PMLR, 15–19 Aug 2021. [2](#), [3](#), [4](#)
- [DKR25] Ilias Diakonikolas, Daniel M. Kane, and Lisheng Ren. Faster Algorithms for Agnostically Learning Disjunctions and their Implications. In *Proceedings of Thirty Eighth Conference on Learning Theory (COLT 2025)*, volume 291, pages 1531–1558. PMLR, 2025. [2](#)
- [DKS18] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Learning geometric concepts with nasty noise. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018)*, pages 1061–1073. ACM, New York, 2018. [2](#)
- [DKTZ20] Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, and Nikos Zarifis. Non-convex SGD learns halfspaces with adversarial label noise. In *Advances in Neural Information Processing Systems*, volume 33, pages 18540–18549. Curran Associates, Inc., 2020. [2](#)
- [DLMF] *NIST Digital Library of Mathematical Functions*. <https://dlmf.nist.gov/>, Release 1.2.5 of 2025-12-15. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds. [19](#)
- [FKV20] Vitaly Feldman, Pravesh Kothari, and Jan Vondrák. Tight bounds on L1 approximation and learning of self-bounding functions. *Theoretical Computer Science*, 808:86–98, 2020. [3](#), [4](#), [8](#), [9](#)

- [Gan08] Michael I. Ganzburg. Limit theorems of polynomial approximation with exponential weights. *Memoirs of the American Mathematical Society*, 192(897):viii+161, 2008. [2](#), [3](#), [4](#)
- [Jan97] Svante Janson. *Gaussian Hilbert Spaces*. Cambridge University Press, 1997. [4](#)
- [Kan11a] Daniel M. Kane. The Gaussian surface area and noise sensitivity of degree- D polynomial threshold functions. *Computational Complexity*, 20(2):389–412, 2011. [2](#), [3](#), [4](#)
- [Kan11b] Daniel M. Kane. k -independent Gaussians fool polynomial threshold functions. In *26th Annual IEEE Conference on Computational Complexity (CCC 2011)*, pages 252–261. IEEE Computer Soc., Los Alamitos, CA, 2011. [9](#)
- [KKM13] Daniel M. Kane, Adam Klivans, and Raghu Meka. Learning halfspaces under log-concave densities: Polynomial approximations and moment matching. In *Proceedings of the 26th Annual Conference on Learning Theory (COLT 2013)*, volume 30, pages 522–545, Princeton, NJ, USA, 2013. PMLR. [2](#)
- [KKMS08] Adam Tauman Kalai, Adam R. Klivans, Yishay Mansour, and Rocco A. Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008. [2](#), [4](#), [8](#)
- [KLS09] Adam R. Klivans, Philip M. Long, and Rocco A. Servedio. Learning halfspaces with malicious noise. *Journal of Machine Learning Research (JMLR)*, 10:2715–2740, 2009. [2](#)
- [KOS02] Adam Klivans, Ryan O’Donnell, and Rocco A. Servedio. Learning Intersections and Thresholds of Halfspaces. In *Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS 2002)*, pages 177–186, USA, 2002. IEEE Computer Society. [8](#)
- [KOS08] Adam R. Klivans, Ryan O’Donnell, and Rocco A. Servedio. Learning Geometric Concepts via Gaussian Surface Area. In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008)*, pages 541–550, 2008. [2](#), [3](#), [4](#), [6](#), [8](#)
- [KSS92] Michael J. Kearns, Robert E. Schapire, and Linda M. Sellie. Toward efficient agnostic learning. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, pages 341–352. Association for Computing Machinery, 1992. [1](#)
- [Led94] Michel Ledoux. Semigroup proofs of the isoperimetric inequality in Euclidean and Gauss space. *Bulletin des Sciences Mathématiques*, 118(6):485–510, 1994. [6](#)
- [LT91] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces*, volume 23. Springer-Verlag, Berlin, 1991. [4](#)
- [O’D14] Ryan O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, New York, 2014. [4](#)
- [Sze75] Gábor Szegő. *Orthogonal polynomials*, volume Vol. XXIII of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, fourth edition, 1975. [15](#), [16](#)
- [Tie23] Stefan Tiegel. Hardness of Agnostically Learning Halfspaces from Worst-Case Lattice Problems. In *Proceedings of Thirty Sixth Conference on Learning Theory (COLT 2023)*, volume 195, pages 3029–3064. PMLR, 2023. [1](#)
- [Val84] Leslie G. Valiant. A theory of the learnable. *Communications of the ACM*, 27:1134–1142, 1984. [1](#)

A Technical proofs

In this appendix, we give some technical proofs omitted from the main text.

A.1 Eigenvalues of the Ornstein-Uhlenbeck operator

We show that the Ornstein-Uhlenbeck operator is diagonalized by the Hermite polynomials (equation (4)). This is a standard result which we include for completeness, see, e.g., [Bak94]. We need the following fact about the generating function for the Hermite polynomials.

Fact A.1. *We have that, for every $x, t \in \mathbb{R}$,*

$$\exp\left(xt - \frac{1}{2}t^2\right) = \sum_{n=0}^{\infty} H_n(x) \frac{t^n}{\sqrt{n!}}.$$

By definition of T_ρ , for any f of the form $f(x) = f_1(x_1) \cdot f_2(x_2) \cdot \dots \cdot f_n(x_n)$, we have

$$T_\rho f(x) = T_\rho f_1(x_1) \cdot T_\rho f_2(x_2) \cdot \dots \cdot T_\rho f_n(x_n).$$

Together with (2), it thus suffices to prove the following one-dimensional case to conclude (4).

Lemma A.2. *For any $n \in \mathbb{N}$, we have that*

$$T_\rho H_n = \rho^n H_n.$$

The proof presented below follows [Bak94, Section 0.I].

Proof. We want to use Fact A.1. Define $g(x) = \exp(xt - \frac{1}{2}t^2)$. Then, by definition of T_ρ and g , we get that

$$\begin{aligned} T_\rho g(x) &= \mathbb{E}_{Y \sim \mathcal{N}_n} \left[g\left(\rho x + \sqrt{1 - \rho^2} Y\right) \right] \\ &= \int_{y \in \mathbb{R}} \exp\left(\rho x t + \sqrt{1 - \rho^2} y t - \frac{1}{2}t^2\right) \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}y^2\right) dy. \end{aligned}$$

Rearranging the terms and completing the square, we get that

$$\begin{aligned} T_\rho g(x) &= \exp\left(\rho x t - \frac{1}{2}t^2\right) \int_{y \in \mathbb{R}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(y - \sqrt{1 - \rho^2}t\right)^2\right) \exp\left(\frac{1}{2}(1 - \rho^2)t^2\right) dy \\ &= \exp\left(\rho x t - \frac{1}{2}\rho^2 t^2\right) \int_{y \in \mathbb{R}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(y - \sqrt{1 - \rho^2}t\right)^2\right) dy. \end{aligned}$$

Substituting now $y' = y - \sqrt{1 - \rho^2}t$, we get that

$$\int_{y \in \mathbb{R}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(y - \sqrt{1 - \rho^2}t\right)^2\right) dy = \int_{y' \in \mathbb{R}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(y')^2\right) dy' = 1$$

and thus we get that

$$T_\rho g(x) = \exp\left(\rho x t - \frac{1}{2}\rho^2 t^2\right).$$

We now use [Fact A.1](#) twice. First, we get that (using [Fact A.1](#) for $t' = \rho t$)

$$T_\rho g(x) = \exp\left(\rho x t - \frac{1}{2}\rho^2 t^2\right) = \sum_{n=0}^{\infty} H_n(x) \rho^n \frac{t^n}{\sqrt{n!}}.$$

On the other hand, first using [Fact A.1](#) and then linearity of T_ρ , we get

$$T_\rho g(x) = T_\rho \left(\sum_{n=0}^{\infty} H_n(x) \frac{t^n}{\sqrt{n!}} \right) = \sum_{n=0}^{\infty} T_\rho H_n(x) \frac{t^n}{\sqrt{n!}}.$$

Thus, we can conclude that

$$\sum_{n=0}^{\infty} H_n(x) \rho^n \frac{t^n}{\sqrt{n!}} = \sum_{n=0}^{\infty} T_\rho H_n(x) \frac{t^n}{\sqrt{n!}}$$

and since this holds for all $t \in \mathbb{R}$, we need to have

$$T_\rho H_n = \rho^n H_n. \quad \square$$

A.2 Ornstein-Uhlenbeck operator and noise sensitivity

Here, we give a proof of [Lemma 2.3](#), showing that $\|f - T_\rho f\|_{L_1(\mathcal{N}_n)} = 2\text{GNS}_{1-\rho}(f)$ for $f : \mathbb{R}^n \rightarrow \{\pm 1\}$.

Proof of Lemma 2.3. We have that

$$\begin{aligned} \|f - T_\rho f\|_{L_1(\mathcal{N}_n)} &= \mathbb{E}_{X \sim \mathcal{N}_n} [|f(X) - T_\rho f(X)|] \\ &= \mathbb{E}_{X \sim \mathcal{N}_n} \left[\left| f(X) - \mathbb{E}_{Y \sim \mathcal{N}_n} \left[f\left(\rho X + \sqrt{1-\rho^2}Y\right) \right] \right| \right] \\ &= \mathbb{E}_{X \sim \mathcal{N}_n} \left[\left| \mathbb{E}_{Y \sim \mathcal{N}_n} \left[f(X) - f\left(\rho X + \sqrt{1-\rho^2}Y\right) \right] \right| \right] \\ &= \mathbb{E}_{X, Y \sim \mathcal{N}_n} \left[\left| f(X) - f\left(\rho X + \sqrt{1-\rho^2}Y\right) \right| \right]. \end{aligned}$$

In the last step, we used the triangle inequality. For a function $f : \mathbb{R}^n \rightarrow \{\pm 1\}$ this is in fact an equality because for fixed x , the term $f(x) - f(\rho x + \sqrt{1-\rho^2}y)$ is either always non-negative (if $f(x) = 1$) or always non-positive (if $f(x) = -1$).

Now, if $X, Y \sim \mathcal{N}(0, I_n)$, then X and $\rho X + \sqrt{1-\rho^2}Y$ are ρ -correlated Gaussians. Therefore, as

$$\left| f(x) - f\left(\rho x + \sqrt{1-\rho^2}y\right) \right| = 2 \cdot \mathbb{1} \left\{ f(x) \neq f\left(\rho x + \sqrt{1-\rho^2}y\right) \right\}$$

the above expectation is equal to $2\text{GNS}_{1-\rho}(f)$. $\square$

A.3 Hermite coefficients of a smooth function

We show how to compute the Hermite coefficients of a smooth function in terms of its derivatives. We need the following fact about the Hermite polynomials. Let φ denote the density function of the standard Gaussian on $\mathbb{R}$.

Fact A.3. For $x \in \mathbb{R}$ and $k \in \mathbb{N}_0$ we have that

$$H_k(x)\varphi(x) = \frac{(-1)^k}{\sqrt{k!}} \cdot \frac{d^k \varphi}{dx^k}(x).$$

With this fact, we can now prove the following lemma.

Lemma A.4. Let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be a smooth function with $\left\| \frac{\partial g}{\partial \alpha} \right\|_\infty < \infty$ for all multi-indices α . Let β be a multi-index. Then we have that

$$\hat{g}(\beta) := \langle g, H_\beta \rangle_{\mathcal{N}_n} = \frac{\mathbb{E}_{X \sim \mathcal{N}_n} \left[\frac{\partial g}{\partial \beta}(X) \right]}{\sqrt{\beta!}}.$$

Proof. In order to prove this lemma, we want to use Gaussian integration by parts on every coordinate. For this, consider first the case $n = 1$. In this case we have, by [Fact A.3](#),

$$\langle g, H_k \rangle_{\mathcal{N}_1} = \frac{(-1)^k}{\sqrt{k!}} \int_{x \in \mathbb{R}} g(x) \cdot \frac{d^k \varphi}{dx^k}(x) dx.$$

Using integration by parts k times, we get

$$\langle g, H_k \rangle_{\mathcal{N}_1} = \frac{1}{\sqrt{k!}} \int_{x \in \mathbb{R}} \frac{d^k g}{dx^k}(x) \cdot \varphi(x) dx = \frac{1}{\sqrt{k!}} \mathbb{E}_{X \sim \mathcal{N}_1} \left[\frac{d^k g}{dx^k}(X) \right].$$

Note that the boundary conditions are always 0 because we assume $\left\| \frac{d^\ell g}{dx^\ell} \right\|_\infty < \infty$. This shows the one-dimensional case. For the multivariate case note that we can iteratively consider each coordinate to get

$$\begin{aligned} \hat{g}(\beta) &:= \langle g, H_\beta \rangle_{\mathcal{N}_n} \\ &= \mathbb{E}_{X \sim \mathcal{N}_n} [g(X) \cdot H_{\beta_1}(X_1) \cdot \dots \cdot H_{\beta_n}(X_n)] \\ &= \mathbb{E}_{X_2, \dots, X_n \sim \mathcal{N}_1} [\mathbb{E}_{X_1 \sim \mathcal{N}_1} [g(X) \cdot H_{\beta_1}(X_1) \cdot \dots \cdot H_{\beta_n}(X_n) \mid X_2, \dots, X_n]] \\ &= \mathbb{E}_{X_2, \dots, X_n \sim \mathcal{N}_1} [\mathbb{E}_{X_1 \sim \mathcal{N}_1} [g(X) \cdot H_{\beta_1}(X_1) \mid X_2, \dots, X_n] \cdot H_{\beta_2}(X_2) \cdot \dots \cdot H_{\beta_n}(X_n)]. \end{aligned}$$

Now, we can use the one-dimensional case to get

$$\mathbb{E}_{X_1 \sim \mathcal{N}_1} [g(X) H_{\beta_1}(X_1) \mid X_2, \dots, X_n] = \frac{1}{\sqrt{\beta_1!}} \mathbb{E}_{X_1 \sim \mathcal{N}_1} \left[\frac{\partial^{\beta_1} g}{\partial x_1^{\beta_1}}(X) \mid X_2, \dots, X_n \right]$$

and thus

$$\hat{g}(\beta) = \frac{1}{\sqrt{\beta_1!}} \mathbb{E}_{X \sim \mathcal{N}_n} \left[\frac{\partial^{\beta_1} g}{\partial x_1^{\beta_1}}(X) \cdot H_{\beta_2}(X_2) \cdot \dots \cdot H_{\beta_n}(X_n) \right].$$

We can iteratively use the same argument for the other coordinates to get the result. $\square$

A.4 Hermite truncation for halfspaces

The goal of this subsection is to prove [Proposition 4.1](#). All O notations are with respect to $d \rightarrow \infty$. We denote by φ the density of the standard Gaussian distribution on $\mathbb{R}$.

Note that $\text{sign} \in L^2(\mathcal{N}_1)$, so the Hermite expansion of sign is well-defined. Moreover, since sign is odd, all even-degree Hermite coefficients of sign are 0, so it suffices to consider the case of odd d . By the same symmetry, we can simplify

$$\|\text{sign} - \Pi_d \text{sign}\|_{L_1(\mathcal{N}_1)} = \int_{-\infty}^{\infty} |\text{sign}(x) - \Pi_d \text{sign}(x)| d\varphi(x) = 2 \int_0^{\infty} |1 - \Pi_d \text{sign}(x)| d\varphi(x). \quad (8)$$

To bound the integral on the right-hand side, we first express the low-degree Hermite projection of sign in terms of Hermite polynomials:

Lemma A.5. *If d is odd, then for every $x \geq 0$,*

$$\Pi_d \text{sign}(x) = \sqrt{\frac{2d}{\pi}} H_{d-1}(0) \int_0^x \frac{H_d(t)}{t} dt.$$

The integral on the right-hand side is well-defined for any fixed odd d , because, as $t \rightarrow 0$, we have $H_d(t) \sim C(d)t$ for some constant $C(d) > 0$.

Proof. For k odd, the k -th Hermite coefficient of sign is

$$\langle \text{sign}, H_k \rangle_{\mathcal{N}_1} = \int_{-\infty}^{\infty} \text{sign}(y) H_k(y) d\varphi(y) = 2 \int_0^{\infty} H_k(y) d\varphi(y).$$

Next, we use the identity $H_k \varphi = -\frac{1}{\sqrt{k}}(H_{k-1} \varphi)'$ (which follows from the recurrence relation that the orthonormalized Hermite polynomials satisfy) to obtain

$$\langle \text{sign}, H_k \rangle_{\mathcal{N}_1} = \frac{2}{\sqrt{k}} [-H_{k-1}(y) \varphi(y)]_0^{\infty} = \sqrt{\frac{2}{k\pi}} H_{k-1}(0).$$

Therefore, for any $x \in \mathbb{R}$,

$$\Pi_d \text{sign}(x) = \sum_{k=1}^d \langle \text{sign}, H_k \rangle_{\mathcal{N}_1} H_k(x) = \sqrt{\frac{2}{\pi}} \sum_{k=1}^d \frac{H_k(x) H_{k-1}(0)}{\sqrt{k}}.$$

Differentiating both sides with respect to x and applying the identity $H'_k = \sqrt{k} H_{k-1}$,

$$(\Pi_d \text{sign})'(x) = \sqrt{\frac{2}{\pi}} \sum_{k=1}^d H_{k-1}(x) H_{k-1}(0).$$

Now, we apply the Christoffel-Darboux formula [[Sze75](#), Theorem 3.2.2]:

$$\sum_{k=0}^{d-1} H_k(x) H_k(0) = \sqrt{d} \frac{H_d(x) H_{d-1}(0) - H_{d-1}(x) H_d(0)}{x}.$$

Since d is odd, $H_d(0) = 0$, so we are left with:

$$(\Pi_d \text{sign})'(x) = \sqrt{\frac{2d}{\pi}} H_{d-1}(0) \frac{H_d(x)}{x},$$

and integrating this identity yields the claim. □

A.4.1 Asymptotics of Hermite polynomials

We will use a pointwise asymptotic expansion of Hermite polynomials $H_d(x)$ as $d \rightarrow \infty$.

The expansion holds uniformly in the region $|x| \ll \sqrt{d}$ and is known as the *Plancherel-Rotach asymptotics* [Sze75]. It takes the form

$$H_d(x) = \exp\left(\frac{x^2}{4}\right) \cdot \left(\frac{2}{\pi d}\right)^{\frac{1}{4}} \cdot \left(\sin\left[\frac{1-d}{2}\pi + \sqrt{d}x\right] + r_d(x)\right), \quad (9)$$

which defines the remainder term $r_d(x)$ that we will bound next. Note that the argument of sin can be simplified depending on the parity of d :

$$\sin\left[\frac{1-d}{2}\pi + \sqrt{d}x\right] = \begin{cases} (-1)^{\frac{d-1}{2}} \sin(\sqrt{d}x) & \text{if } d \text{ is odd} \\ (-1)^{\frac{d}{2}} \cos(\sqrt{d}x) & \text{if } d \text{ is even} \end{cases}$$

Lemma A.6. For all $|x| \leq \sqrt{d}$, we have

$$|r_d(x)| \leq O\left(\max\left\{\frac{|x|^3}{\sqrt{d}}, \frac{|x|}{\sqrt{d}}, \frac{x^2}{d}, \frac{1}{d}\right\}\right),$$

where the O is uniform over all $|x| \leq \sqrt{d}$.

Proof. We use the following formula [Sze75, Theorem 8.22.9]:

$$H_d(x) = \exp\left(\frac{x^2}{4}\right) \cdot \left(\frac{2}{\pi d}\right)^{\frac{1}{4}} \cdot \sin\left[\arccos\left(\frac{x}{\sqrt{4d+2}}\right)\right]^{-\frac{1}{2}} \cdot \left(\sin\left[\left(\frac{d}{2} + \frac{1}{4}\right)\left(2\frac{x}{\sqrt{4d+2}}\sqrt{1 - \frac{x^2}{4d+2}} - 2\arccos\left(\frac{x}{\sqrt{4d+2}}\right)\right) + \frac{3\pi}{4}\right] + O\left(\frac{1}{d}\right)\right).$$

This holds whenever $|x| \leq \frac{1}{2}\sqrt{4d+2}$ and the O is uniform over all such x . We use Taylor expansion to get that $\sin[\arccos(x/\sqrt{4d+2})] = 1 + O(x^2/d)$, where the O is again uniform over all $|x| \leq \frac{1}{2}\sqrt{4d+2}$. Defining $g(y) = 2y\sqrt{1-y^2} - 2\arccos(y)$, we can thus write

$$H_d(x) = \exp\left(\frac{x^2}{4}\right) \cdot \left(\frac{2}{\pi d}\right)^{\frac{1}{4}} \cdot \left(\sin\left[\left(\frac{d}{2} + \frac{1}{4}\right) \cdot g\left(\frac{x}{\sqrt{4d+2}}\right) + \frac{3\pi}{4}\right] + O\left(\max\left\{\frac{x^2}{d}, \frac{1}{d}\right\}\right)\right).$$

Using the Taylor expansion of g around 0, we get $g(y) = -\pi + 4y + \frac{1}{2}g''(\xi)y^2$ for some ξ between 0 and y . We can bound the derivative $g''(\xi)$ by $O(|y|)$, for some uniform constant for all $|y| \leq \frac{1}{2}$. Putting this all together and using that sin is 1-Lipschitz, we get

$$\begin{aligned} H_d(x) &= \exp\left(\frac{x^2}{4}\right) \cdot \left(\frac{2}{\pi d}\right)^{\frac{1}{4}} \cdot \left(\sin\left[\left(\frac{d}{2} + \frac{1}{4}\right) \cdot \left(-\pi + 4\frac{x}{\sqrt{4d+2}}\right) + \frac{3\pi}{4}\right] + O\left(\max\left\{\frac{|x|^3}{\sqrt{d}}, \frac{x^2}{d}, \frac{1}{d}\right\}\right)\right) \\ &= \exp\left(\frac{x^2}{4}\right) \cdot \left(\frac{2}{\pi d}\right)^{\frac{1}{4}} \cdot \left(\sin\left[\frac{1-d}{2}\pi + \sqrt{d + \frac{1}{2}}x\right] + O\left(\max\left\{\frac{|x|^3}{\sqrt{d}}, \frac{x^2}{d}, \frac{1}{d}\right\}\right)\right). \end{aligned}$$

Next, we note that $\sqrt{d + \frac{1}{2}} = \sqrt{d} + O(1/\sqrt{d})$ and hence, using Lipschitzness of $\sin$ again, we get

$$H_d(x) = \exp\left(\frac{x^2}{4}\right) \cdot \left(\frac{2}{\pi d}\right)^{\frac{1}{4}} \cdot \left(\sin\left[\frac{1-d}{2}\pi + \sqrt{d}x\right] + O\left(\max\left\{\frac{|x|^3}{\sqrt{d}}, \frac{|x|}{\sqrt{d}}, \frac{x^2}{d}, \frac{1}{d}\right\}\right)\right),$$

as claimed. $\square$

Corollary A.7. *If d is even, then*

$$H_d(0) = (-1)^{\frac{d}{2}} \left(\frac{2}{\pi}\right)^{\frac{1}{4}} d^{-\frac{1}{4}} + O(d^{-\frac{5}{4}}).$$

Corollary A.8. *We have*

$$\sup_{t \in [0, d^{1/6}]} |H_d(t)| e^{-t^2/4} \leq O(d^{-\frac{1}{4}}).$$

Corollary A.9. *If d is odd, then*

$$\sup_{t \in (0, d^{-1/2}]} \frac{|r_d(t)|}{t} \leq O\left(\frac{1}{\sqrt{d}}\right).$$

Proof of Corollary A.9 from Lemma A.6. We apply Taylor's inequality (for d odd, we have $r_d(0) = 0$):

$$\sup_{t \in (0, d^{-1/2}]} \frac{|r_d(t)|}{t} \leq \sup_{t \in (0, d^{-1/2}]} |r'_d(t)|. \quad (10)$$

Next, we compute $r'_d(t)$. Differentiating (9) term by term, we get

$$H'_d(t) = \left(\frac{2}{\pi d}\right)^{1/4} e^{t^2/4} \left(\frac{t}{2}((-1)^{\frac{d-1}{2}} \sin(t\sqrt{d}) + r_d(t)) + \sqrt{d}(-1)^{\frac{d-1}{2}} \cos(t\sqrt{d}) + r'_d(t)\right).$$

Using the identity $H'_d(t) = \sqrt{d}H_{d-1}(t)$, we obtain, after rearranging:

$$r'_d(t) = \sqrt{d} \left[\left(\frac{\pi d}{2}\right)^{1/4} e^{-t^2/4} H_{d-1}(t) - (-1)^{\frac{d-1}{2}} \cos(t\sqrt{d}) \right] - \frac{t}{2} \left[(-1)^{\frac{d-1}{2}} \sin(t\sqrt{d}) + r_d(t) \right].$$

We show separately that the first and the second term are small in absolute value. The second term is $O(d^{-1/2})$ whenever $t \leq d^{-1/2}$. For the first term, we apply Lemma A.6 to $H_{d-1}(t)$ (in the regime $t \leq d^{-1/2}$, the error term $1/d$ dominates):

$$\sup_{t \in [0, d^{-1/2}]} \left| \left(\frac{\pi d}{2}\right)^{1/4} e^{-t^2/4} H_{d-1}(t) - (-1)^{\frac{d-1}{2}} \cos(t\sqrt{d}) \right| \leq O\left(\frac{1}{d}\right).$$

The proof follows from multiplying by $\sqrt{d}$ and substituting the bound in (10). $\square$

A.4.2 Proof of Proposition 4.1

Applying the Hermite asymptotic expansion of $H_{d-1}(0)$ to our target (8),

$$\begin{aligned}\Pi_d \text{sign}(x) &= \sqrt{\frac{2d}{\pi}} H_{d-1}(0) \int_0^x \frac{H_d(t)}{t} dt && \text{(by Lemma A.5)} \\ &= \left((-1)^{\frac{d-1}{2}} \left(\frac{2}{\pi} \right)^{\frac{3}{4}} d^{\frac{1}{4}} + O(d^{-\frac{3}{4}}) \right) \int_0^x \frac{H_d(t)}{t} dt. && \text{(by Corollary A.7)}\end{aligned}$$

The main technical part of the argument lies in bounding the integral of $t \mapsto H_d(t)/t$.

Lemma A.10 (Small t). *If d is odd, then for all $\tau = \tau(d) \leq d^{1/6}$,*

$$\int_0^\tau \frac{|H_d(t)|}{t} dt \leq O(d^{\frac{1}{4}}) \cdot \tau e^{\tau^2/4}.$$

Proof. Since $H_d(0) = 0$ (when d is odd), Taylor's inequality yields:

$$\left| \int_0^\tau \frac{H_d(t)}{t} dt \right| \leq \tau \cdot \sup_{t \in [0, \tau]} |H'_d(t)|.$$

The result follows from the identity $H'_d(t) = \sqrt{d} H_{d-1}(t)$ and Corollary A.8. □

Lemma A.11 (Large t). *For all $\tau \geq 1$,*

$$\int_0^\infty \int_\tau^x \frac{|H_d(t)|}{t} dt d\varphi(x) \leq e^{-\tau^2/4}.$$

Proof. Using the notation $\bar{\Phi}(t) = \Pr_{x \sim \mathcal{N}(0,1)}(x > t)$, we have

$$\begin{aligned}\int_0^\infty \int_\tau^x \frac{|H_d(t)|}{t} dt d\varphi(x) &= \int_\tau^\infty \frac{|H_d(t)|}{t} \bar{\Phi}(t) dt && \text{(by Fubini's theorem)} \\ &\leq \left(\int_\tau^\infty H_d(t)^2 \varphi(t) dt \right)^{\frac{1}{2}} \left(\int_\tau^\infty \frac{\bar{\Phi}(t)^2}{t^2 \varphi(t)} dt \right)^{\frac{1}{2}} && \text{(by Cauchy-Schwarz)} \\ &\leq \left(\int_\tau^\infty \varphi(t) dt \right)^{\frac{1}{2}} && \text{(orthonormality, } \bar{\Phi}(t) \leq \varphi(t), \text{ and } \tau \geq 1) \\ &\leq e^{-\tau^2/4}. && (\bar{\Phi}(\tau) \leq e^{-\tau^2/2})\end{aligned}$$

□

Proof of Proposition 4.1. We start by expanding $H_{d-1}(0)$ with Corollary A.7:

$$\begin{aligned}&\int_0^\infty |1 - \Pi_d \text{sign}(x)| d\varphi(x) \\ &= \int_0^\infty \left| 1 - (-1)^{\frac{d-1}{2}} \left(\frac{2}{\pi} \right)^{\frac{3}{4}} d^{\frac{1}{4}} \int_0^x \frac{H_d(t)}{t} dt \right| d\varphi(x) + O(d^{-\frac{3}{4}}) \cdot \int_0^\infty \int_0^x \frac{|H_d(t)|}{t} dt d\varphi(x).\end{aligned}$$

We bound the second term using [Lemma A.10](#) when $t \leq d^{1/6}$ and [Lemma A.11](#) when $t > d^{1/6}$:

$$d^{-\frac{3}{4}} \int_0^\infty \int_0^x \frac{|H_d(t)|}{t} dt d\varphi(x) \leq d^{-\frac{3}{4}} \left(\int_0^\infty O(d^{1/4}) \cdot x e^{x^2/4} d\varphi(x) + e^{-d^{1/3}/4} \right) \leq O(d^{-\frac{1}{2}}).$$

This concludes the analysis of the second term. Next, we split the inner integral of the first term into $t \in [0, \tau]$ and $t \in [\tau, \infty)$ for $\tau = \sqrt{100 \log d}$. We start with the latter, which is negligible by [Lemma A.11](#):

$$d^{\frac{1}{4}} \int_0^\infty \int_\tau^x \frac{|H_d(t)|}{t} dt d\varphi(x) \leq d^{\frac{1}{4}} e^{-\tau^2/4} = O(d^{-\frac{1}{2}}).$$

We switch to $t \in [0, \tau]$, which gives the main contribution. We decompose as follows:

$$\begin{aligned} \int_0^\infty \left| 1 - (-1)^{\frac{d-1}{2}} \left(\frac{2}{\pi} \right)^{\frac{3}{4}} d^{\frac{1}{4}} \int_0^{\min(x, \tau)} \frac{H_d(t)}{t} dt \right| d\varphi(x) &\leq \int_0^\infty \left| 1 - \frac{2}{\pi} \int_0^{\min(x, \tau)} \frac{\sin(t\sqrt{d})}{t} dt \right| d\varphi(x) \\ &\quad + \frac{2}{\pi} \int_0^\infty \left| \int_0^{\min(x, \tau)} \sin(t\sqrt{d}) \cdot \frac{e^{t^2/4} - 1}{t} dt \right| d\varphi(x) \\ &\quad + \frac{2}{\pi} \int_0^\infty \left| \int_0^{\min(x, \tau)} \frac{e^{t^2/4} r_d(t)}{t} dt \right| d\varphi(x). \end{aligned}$$

where we recall that $r_d(t)$ is defined in (9). We will bound the three terms separately:

$$\int_0^\infty \left| 1 - \frac{2}{\pi} \int_0^{\min(\tau, x)} \frac{\sin(t\sqrt{d})}{t} dt \right| d\varphi(x) \leq O\left(\frac{\log d}{\sqrt{d}}\right) \quad (11)$$

$$\int_0^\infty \left| \int_0^{\min(x, \tau)} \frac{e^{t^2/4} - 1}{t} \sin(t\sqrt{d}) dt \right| d\varphi(x) \leq O\left(\frac{1}{\sqrt{d}}\right) \quad (12)$$

$$\int_0^\infty \int_0^{\min(x, \tau)} \frac{e^{t^2/4} |r_d(t)|}{t} dt d\varphi(x) \leq O\left(\frac{1}{\sqrt{d}}\right). \quad (13)$$

Assuming (11), (12), and (13), the proof of [Proposition 4.1](#) follows from the triangle inequality. $\square$

Proof of (11). First, by a change of variable,

$$\int_0^{\min(\tau, x)} \frac{\sin(t\sqrt{d})}{t} dt = \int_0^{\min(\tau, x) \cdot \sqrt{d}} \frac{\sin u}{u} du.$$

Claim A.12 (See, e.g., [\[DLMF, §6.12\]](#)). *There exists a universal constant $C > 0$ such that for any $z > 0$,*

$$\left| 1 - \frac{2}{\pi} \int_0^z \frac{\sin t}{t} dt \right| \leq C \min\left(1, \frac{1}{z}\right).$$

Applying the bound from [Claim A.12](#), we get

$$\int_0^\infty \left| 1 - \frac{2}{\pi} \int_0^{\min(\tau, x)} \frac{\sin(t\sqrt{d})}{t} dt \right| d\varphi(x) \leq O(1) \cdot \int_0^\infty \min\left(1, \frac{1}{\sqrt{d} \min(\tau, x)}\right) d\varphi(x).$$

We further split this integral into three parts, according to which of the arguments of the minima dominate:

1. between 0 and $\frac{1}{\sqrt{d}}$, the integral is bounded by $O\left(\frac{1}{\sqrt{d}}\right)$ by the first bound.
2. between $\frac{1}{\sqrt{d}}$ and τ , $\min(\tau, x) = x$ so the integral is $O\left(\frac{\log d}{\sqrt{d}}\right)$ by the second bound.
3. between τ and $+\infty$, the integral is $O\left(\frac{1}{\sqrt{d}}\right)$ by the second bound.

Taken together, these observations conclude the proof. $\square$

Proof of (12). Let $u(t) = \frac{e^{t^2/4}-1}{t}$ extended by continuity with $u(0) = 0$. A calculation shows that

$$\max(|u(t)|, |u'(t)|) \leq \frac{e^{t^2/4}}{2} \quad \text{for all } t \geq 0. \quad (14)$$

The inner integral we want to study is a product of u with an oscillatory factor of frequency $\sqrt{d}$, so one expects substantial cancellations as $d \rightarrow \infty$. We make this explicit by integrating by parts:

$$\int_0^{\min(x, \tau)} \sin(t\sqrt{d})u(t) dt = \left[-\frac{\cos(t\sqrt{d})}{\sqrt{d}}u(t) \right]_0^{\min(x, \tau)} + \frac{1}{\sqrt{d}} \int_0^{\min(x, \tau)} \cos(t\sqrt{d})u'(t) dt.$$

Hence, by the triangle inequality and (14),

$$\left| \int_0^{\min(x, \tau)} \sin(t\sqrt{d})u(t) dt \right| \leq \frac{(1+x)e^{x^2/4}}{2\sqrt{d}}.$$

Integrating the right-hand side against φ yields $O(d^{-1/2})$, as required. $\square$

Proof of (13). We decompose the inner integral according to whether $t \leq d^{-1/2}$ or $t > d^{-1/2}$, and use [Corollary A.9](#) on the first part:

$$\begin{aligned} \int_0^\infty \int_0^{\min(x, \tau)} e^{t^2/4} \frac{|r_d(t)|}{t} dt d\varphi(x) &\leq O\left(\frac{1}{\sqrt{d}}\right) \cdot \int_0^\infty x e^{x^2/4} d\varphi(x) \\ &\quad + \int_{d^{-1/2}}^\infty \int_{d^{-1/2}}^{\min(x, \tau)} e^{t^2/4} \frac{|r_d(t)|}{t} dt d\varphi(x). \end{aligned}$$

The first term is $O(d^{-1/2})$, and for the second we use Fubini's theorem:

$$\int_{d^{-1/2}}^\infty \int_{d^{-1/2}}^{\min(x, \tau)} e^{t^2/4} \frac{|r_d(t)|}{t} dt d\varphi(x) = \int_{d^{-1/2}}^\tau \frac{|r_d(t)|}{t} e^{t^2/4} \bar{\Phi}(t) dt \leq \int_{d^{-1/2}}^\tau \frac{|r_d(t)|}{t} e^{-t^2/4} dt,$$

where we used the notation $\bar{\Phi}(t)$ for the Gaussian tails as in the proof of [Lemma A.11](#), and applied the standard bound $\bar{\Phi}(t) \leq e^{-t^2/2}$.

When $t \leq \tau$ and $t \geq \frac{1}{\sqrt{d}}$ (i.e., $\frac{t}{\sqrt{d}} \geq \frac{1}{d}$) the upper bound from [Lemma A.6](#) simplifies to:

$$|r_d(t)| \leq O\left(\frac{t^3}{\sqrt{d}} + \frac{t}{\sqrt{d}} + \frac{t^2}{d}\right).$$

Substituting this into the integral above, we obtain:

$$\int_{d^{-1/2}}^\tau \frac{|r_d(t)|}{t} e^{-t^2/4} dt \leq O\left(\frac{1}{\sqrt{d}}\right),$$

as desired. $\square$