

On Suboptimality of Least Squares with Application to Estimation of Convex Bodies

Gil Kur

GILKUR@MIT.EDU

MIT

Alexander Rakhlin

RAKHLIN@MIT.EDU

MIT

Adityanand Guntuboyina

ADITYA@STAT.BERKELEY.EDU

UC Berkeley

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We develop a technique for establishing lower bounds on the sample complexity of Least Squares (or, Empirical Risk Minimization) for large classes of functions. As an application, we settle an open problem regarding optimality of Least Squares in estimating a convex set from noisy support function measurements in dimension $d \geq 6$. Specifically, we establish that Least Squares is *mimimax sub-optimal*, and achieves a rate of $\bar{\Theta}_d(n^{-2/(d-1)})$ whereas the minimax rate is $\Theta_d(n^{-4/(d+3)})$.

Keywords: Non-parametric statistics, regression, support function, ERM, Least squares, non-Donsker regime

1. Introduction and main results

Consider the problem of regression where the goal is to estimate a function $f^* : \mathcal{X} \rightarrow \mathbb{R}$ from observations $(X_1, Y_1), \dots, (X_n, Y_n)$ drawn according to the model

$$Y_i = f^*(X_i) + \xi_i \quad \text{for } i = 1, \dots, n$$

where X_1, \dots, X_n are design points (fixed or random) and ξ_1, \dots, ξ_n are i.i.d random variables with mean zero. Assume that the unknown function f^* belongs to a known *convex* class of functions \mathcal{F} . The most natural and basic estimator in this setting is the Least Squares Estimator (LSE)

$$\hat{f}_n \in \operatorname{argmin}_{f \in \mathcal{F}} \sum_{i=1}^n (Y_i - f(X_i))^2,$$

which is also the Maximum Likelihood Estimator (MLE) under Gaussian noise. It is customary to evaluate the performance of \hat{f}_n via the risk

$$\|\hat{f}_n - f^*\|_{\mathbb{P}_n}^2 := \frac{1}{n} \sum_{i=1}^n \left(\hat{f}_n(X_i) - f^*(X_i) \right)^2 \tag{1}$$

in the fixed-design setting where X_1, \dots, X_n are assumed to be non-random points in \mathcal{X} , and via

$$\|\hat{f}_n - f^*\|_{\mathbb{P}}^2 := \int_{\mathcal{X}} \left(\hat{f}_n(x) - f^*(x) \right)^2 d\mathbb{P}(x) \tag{2}$$

in the random design setting where $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} \mathbb{P}$. We shall focus on the optimality of the LSE as measured in the minimax sense (see e.g., [Tsybakov \(2003\)](#)) via the worst-case risk

$$\mathcal{R}_{\mathbb{P}_n}(\hat{f}_n, \mathcal{F}) := \sup_{f^* \in \mathcal{F}} \mathbb{E}_{f^*} \|\hat{f}_n - f^*\|_{\mathbb{P}_n}^2 \quad \text{or} \quad \mathcal{R}_{\mathbb{P}}(\hat{f}_n, \mathcal{F}) := \sup_{f^* \in \mathcal{F}} \mathbb{E}_{f^*} \|\hat{f}_n - f^*\|_{\mathbb{P}}^2,$$

in the fixed and random designs, respectively.

The accuracy of the LSE is one of the most basic questions in statistics ([van de Geer, 2000](#); [Birgé and Massart, 1993](#); [Koltchinskii, 2011](#); [van der Vaart and Wellner, 1996](#)). Typical results impose conditions on the metric entropy on (natural subsets of) the class \mathcal{F} . We shall denote the ϵ -metric entropy of a class \mathcal{G} of functions under the $L^2(\mathbb{Q})$ pseudometric by $\log \mathcal{N}(\epsilon, \mathcal{G}, \mathbb{Q})$. Here $\mathcal{N}(\epsilon, \mathcal{G}, \mathbb{Q})$ is the minimal cardinality of a set \mathcal{N}_ϵ of functions such that for any $g \in \mathcal{G}$ there exists a $g' \in \mathcal{N}_\epsilon$ with $\|g - g'\|_{\mathbb{Q}} \leq \epsilon$.

It is convenient to isolate the existing results on the LSE into two categories depending on whether the class of functions \mathcal{F} is in the *Donsker* regime or in the *non-Donsker* regime. By the Donsker (or, more precisely, \mathbb{P} -Donsker) regime, we mean that the $L^2(\mathbb{P})$ metric entropy of natural totally bounded subsets of \mathcal{F} grows as ϵ^{-p} for some $p < 2$.¹ In contrast, for the non-Donsker regime, the entropy grows as ϵ^{-p} for some $p > 2$.

It should be noted that the two regimes can often occur in different settings of the same general problem. For example, the class of support functions of compact, convex subsets of \mathbb{R}^d is Donsker for $d \leq 4$ and non-Donsker for $d \geq 6$. We shall revisit this example in detail in [Section 2](#).

In the Donsker regime, it is well-known that the LSE achieves the rate $n^{-2/(2+p)}$, which is the minimax rate of convergence, under minimal assumptions ([Barron et al., 1999](#)). Thus, the LSE is minimax optimal when the class \mathcal{F} is Donsker with respect to \mathbb{P} or \mathbb{P}_n . In contrast, the minimax optimality of the LSE in the non-Donsker regime remains unresolved. In a fundamental paper, [Birgé and Massart \(1993\)](#) proved that in the non-Donsker regime, the rate of convergence of the LSE is bounded from above by $n^{-1/p}$. However, it is also well-known (see e.g., [Yang and Barron \(1999\)](#)) that the minimax rate of estimation is still $n^{-2/(2+p)}$.

It was first observed in ([Birgé and Massart, 1993](#), Section 4) and later in [Birgé \(2006\)](#) that it is possible to design “pathological” function classes \mathcal{F} where the LSE provably achieves a risk of $\tilde{\Theta}(n^{-1/p})$. However, for a general non-Donsker class \mathcal{F} , the fundamental question of whether the rate of convergence of the LSE is $n^{-1/p}$ or $n^{-2/(2+p)}$ (or some intermediate rate) is unresolved.

Some important progress on this open problem has been made in the recent papers ([Kur et al., 2019](#); [Han, 2019](#); [Carpenter et al., 2018](#); [Han et al., 2019](#)). Specifically, these papers have shown that there exist “natural” non-Donsker families of functions (including the class of bounded convex functions on a smooth domain in \mathbb{R}^d for $d \geq 3$, the class of multivariate Isotonic functions over \mathbb{R}^d for $d \geq 2$, etc.) where the LSE achieves the minimax rate of $n^{-2/(2+p)}$, in contrast to the pessimistic upper bound of $n^{-1/p}$.

In this paper, we give a recipe for establishing a lower bound for the LSE’s risk for convex families of functions in the *non-Donsker* regime both in fixed and random design. As an application, we complement the aforementioned recent results by proving that there also exist other “natural” non-Donsker families of functions (such as the class of support functions of compact convex subsets of \mathbb{R}^d for $d \geq 6$) where the LSE cannot achieve a rate faster than $n^{-1/p}$, up to logarithmic multiplicative factors. In other words, for these non-Donsker classes, the LSE is provably *suboptimal*.

1. In this section, for brevity we may omit constants that may depend on p .

We start by stating assumptions that imply sub-optimality of LSE in a fixed design. Then, we extend these assumption to the random design setting. These assumptions will be verified for the case of support functions in Section 2.

Throughout, we assume that \mathcal{F} is a *convex* family of functions. Also, in the remainder of the paper, we assume for simplicity that $\xi \sim N(0, 1)$; see Remark 5 for more general assumptions on the noise. We employ the notation $a \lesssim b$ to mean inequality up to a constant that may depend on p . Let

$$B_{\mathbb{P}_n}(f_0, t) := \{f \in \mathcal{F} : \|f - f_0\|_{\mathbb{P}_n} \leq t\}$$

and for a set of functions \mathcal{G} on \mathcal{X} , the Gaussian width of \mathcal{G} is defined, conditionally on X_1, \dots, X_n , as

$$\bar{\mathcal{W}}(\mathcal{G}) := \mathbb{E} \sup_{f \in \mathcal{G}} \frac{1}{n} \sum_{i=1}^n g_i f(X_i), \quad (3)$$

where g_1, \dots, g_n are i.i.d. $N(0, 1)$.

We are now ready to state our key structural assumptions.

Assumptions: Fixed design Let $p > 2$ and $n \gtrsim 1$. Assume there are two functions $f_0, f_{(n)} \in \mathcal{F}$ such that for some $\beta, \gamma \geq 0$,

$$\log \mathcal{N}(\epsilon, B_{\mathbb{P}_n}(f_0, 2\epsilon), \mathbb{P}_n) \gtrsim \epsilon^{-p} \quad \forall \epsilon \gtrsim n^{-1/p} \log(n)^\gamma. \quad (4)$$

Furthermore, suppose

$$\|f_{(n)} - f_0\|_{\mathbb{P}_n} \lesssim n^{-\frac{1}{2p}} \log(n)^{-s_{p,\gamma}}, \quad (5)$$

where $s_{p,\gamma} = \frac{\gamma}{2}(\frac{p}{2} - 1)$, and either

$$\log \mathcal{N}(\epsilon, B_{\mathbb{P}_n}(f_{(n)}, t), \mathbb{P}_n) \lesssim \sqrt{n} \log(n)^{p \cdot s_{p,\gamma}} \left(\log \frac{1}{\epsilon} \right)^\beta \left(\frac{t}{\epsilon} \right)^p \quad (6)$$

for every $n^{-\frac{1}{p}} \log(n)^{-s_{p,\gamma}} \lesssim \epsilon \leq t \lesssim n^{-\frac{1}{2p}} \log(n)^{-\frac{\beta}{p} - \frac{3s_{p,\gamma}}{2}}$ or a weaker condition

$$\bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, t)) \lesssim n^{-\frac{1}{2p}} \log(n)^{\frac{\beta}{p} + s_{p,\gamma}} t, \quad (7)$$

for some $t \lesssim n^{-\frac{1}{2p}} \log(n)^{-\frac{\beta}{p} - \frac{3s_{p,\gamma}}{2}}$ holds.

The following theorem establishes that the rate of convergence of the LSE equals $n^{-1/p}$, up to a poly-logarithmic multiplicative factor, under the aforementioned assumptions.

Theorem 1 *Assume that \mathcal{F} is a convex class of functions satisfying (4), (5), and either (6) or (7) holds for some $p > 2$ and $n \gtrsim 1$. Then the following holds:*

$$\mathcal{R}_{\mathbb{P}_n}(\hat{f}_n, \mathcal{F}) \gtrsim n^{-\frac{1}{p}} \log(n)^{-\frac{2\beta}{p} - 3s_{p,\gamma}}.$$

We shall next extend Theorem 1 to the case of random design, where the design points X_1, \dots, X_n are independent with common distribution \mathbb{P} . It is sufficient for our purposes to establish near-isometry: with high probability,

$$\forall f, g \in \mathcal{F}, \quad \frac{1}{2} \|f - g\|_{\mathbb{P}_n} - d_n \leq \|f - g\|_{\mathbb{P}} \leq 2 \|f - g\|_{\mathbb{P}_n} + d_n \quad (8)$$

for some remainder d_n that decays to zero with increasing n . In the case of uniformly bounded functions, sufficient conditions for the two-sided inequality (8) can be found in the literature on local Rademacher averages (see e.g., (Bousquet, 2002; Bartlett et al., 2005)), while the right-hand side of (8) holds under weaker conditions (Koltchinskii and Mendelson, 2013; Mendelson, 2014, 2017). For the purpose of proving lower bounds for random design, however, we need the more demanding left-hand side of (8). Hence, we shall assume that functions in \mathcal{F} are uniformly bounded.

In addition, we assume the growth of Koltchinskii-Pollard entropy

$$\sup_{n \in \mathbb{N}} \sup_{\mathbb{Q} \in \mathcal{P}_n} \log \mathcal{N}(\epsilon, \mathcal{F}, \mathbb{Q}) \lesssim \epsilon^{-p} \quad (9)$$

where \mathcal{P}_n denotes the set of all probability measures supported on finite subsets of \mathcal{X} of cardinality at most n . Under this assumption, with high probability $d_n \lesssim (\log n)^2 n^{-1/p}$ (Rakhlin et al., 2017). This will allow us to reduce the random design setting to a fixed design and use Theorem 1, since the remainder $(d_n)^2$ is of the lower order than $n^{-1/p}$.²

To establish a lower bound for LSE in random design, we may verify that above assumptions (4), (5), (6) hold with high probability for random measures \mathbb{P}_n and employ near-isometry for the distance between LSE and the regression function. Alternatively, it may be easier to verify the corresponding assumptions in the population. We now state this latter approach.

First, recall that bracketing number $\mathcal{N}_{[]}(\epsilon, \mathcal{G}, \mathbb{P})$ is the minimal cardinality of a set $\mathcal{N}_{[], \epsilon}$ of pairs of functions (g^-, g^+) such that $\|g^+ - g^-\|_{\mathbb{P}} \leq \epsilon$, and for any $g \in \mathcal{G}$ there exists $(g^-, g^+) \in \mathcal{N}_{[], \epsilon}$ such that $g^- \leq g \leq g^+$.

Assumptions: Random design Let $p > 2$. Assume there exists a function $f_0 \in \mathcal{F}$ such that for every integer $m \gtrsim 1$, there exists $f_{(m)} \in \mathcal{F}$ such that

$$\log \mathcal{N}(\epsilon, B_{\mathbb{P}}(f_0, 2\epsilon), \mathbb{P}) \gtrsim \epsilon^{-p} \quad \forall \epsilon \geq 0 \quad (10)$$

and

$$\|f_{(m)} - f_0\|_{\mathbb{P}} \lesssim m^{-\frac{1}{p}}. \quad (11)$$

Furthermore, for some $\beta \geq 0$

$$\log \mathcal{N}_{[]}(\epsilon, B_{\mathbb{P}}(f_{(m)}, t), \mathbb{P}) \lesssim m \left(\log \frac{1}{\epsilon} \right)^{\beta} \left(\frac{t}{\epsilon} \right)^p, \quad (12)$$

for every $0 < \epsilon \leq t \lesssim \log(m)^{-\beta/p} m^{-1/p}$.

We remark that assumption (10) is satisfied for non-Donsker classes of functions: see (Yang and Barron, 1999, Lemma 3) for the proof of existence.

The following is our lower bound for the random-design setting:

Corollary 2 *Assume that \mathcal{F} is a convex class of uniformly bounded functions satisfying (9), (10), (11), and (12) for some $p > 2$ and $n \gtrsim 1$. Then*

$$\mathcal{R}_{\mathbb{P}}(\hat{f}_n, \mathcal{F}) \gtrsim n^{-\frac{1}{p}} \log(n)^{-\frac{2\beta}{p} - 3(\frac{p}{2} - 1)}.$$

2. For some families of functions it might be possible to relax the additional assumptions of uniform boundedness and (9) and achieve a similar result to Corollary 2, see for example Kur et al. (2020)

1.1. The ideas behind our assumptions

Assumption (10) says that the function f_0 is “complex” in that its local neighborhood is as complex as the entire function class \mathcal{F} when $\log \mathcal{N}(\epsilon, \mathcal{F}, \mathbb{P}) \sim \epsilon^{-p}$. The assumption (11) states that f_0 is approximated by $f_{(m)}$ up to the accuracy $m^{-1/p}$ for every m . The assumption (12) captures the “simplicity” of the functions $f_{(m)}$ in relation to the complex function f_0 satisfying (10). Note that when $t = O(\epsilon)$, the right hand side of (12) is logarithmic in ϵ (assuming that m is not too large) while the right hand side of (10) is polynomial in $1/\epsilon$. Thus the local neighborhood of $f_{(m)}$ is smaller than that of f_0 and in this sense $f_{(m)}$ is simpler than f_0 . Note also that there is a factor of m on the right hand side of (12) which means that the complexity of the functions $f_{(m)}$ increases with m .

There exist natural non-Donsker function classes \mathcal{F} which satisfy (10), (11) and (12). In Section 2, we show that the class of support functions of compact convex sets in \mathbb{R}^d satisfies these assumptions with f_0 being the support function of the unit ball and $f_{(m)}$ the support function of a regular polytopal approximation to the unit ball with m vertices.

The proof of Theorem 1 will reveal that \hat{f}_n achieves the rate $n^{-1/p}$ (up to logarithmic factors) when the regression function is in the class $\{f_{(m)}, m \geq 1\}$. In fact, the specific function achieving the rate $n^{-1/p}$ is $f^* = f_{(m)}$ for $m \sim \sqrt{n}$. It is interesting to note that for $m \sim \sqrt{n}$ and $t^2 \sim n^{-1/p}$, the local neighborhood $B_{\mathbb{P}_n}(f_{(m)}, t)$ has the same metric entropy as the right hand side of (10). Our main technical insight is that the sub-optimality occurs at a function $f^* = f_{(m)}$ instead of perhaps a more natural candidate function such as f_0 (we actually believe that the rate at f_0 may be equal to $n^{-2/(2+p)}$).

1.2. Discussion: the non-Donsker regime, revisited

From the recent results on families in the non-Donsker regime (Kur et al., 2019; Han, 2019; Carpenter et al., 2018; Kur et al., 2020; Han et al., 2019) and the present results, the LSE may be optimal or sub-optimal in the non-Donsker regime for natural classes of functions, even if uniformly bounded. These results also indicate that LSE achieves a risk that equals to the Gaussian complexity of the class, up to a constant that depends on dimension. Namely, in contrast to the Donsker regime, there is no localization. In the problem of convex regression with uniformly bounded functions and *Euclidean ball* as domain, as well as in multiple isotopic regression, the minimax rate equals to the Gaussian complexity of the family. In contrast, by Theorem 3, for support function regression and for convex uniformly bounded regression (or Lipschitz-convex regression) with support on the *cube* (Kur et al., 2020), the Gaussian complexity differs from the minimax rate.

2. Application: Sub-optimality of Least Squares for estimating a convex set from noisy support function measurements

In this section, we shall use Corollary 2 to resolve a long-standing open problem on the sub-optimality of the least squares estimator in the problem of estimating a convex set in dimension $d \geq 6$ from noisy support function measurements; see for example (Gardner et al., 2006; Brunel, 2016, 2013; Brunel et al., 2018; Guntuboyina, 2012; Fisher et al., 1997; Balázs et al., 2015; Soh and Chandrasekaran, 2019). Let us first recall that the support function $h_K : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ of a compact

convex set K in \mathbb{R}^d is defined as

$$h_K(u) = \max_{x \in K} \langle x, u \rangle \quad \text{for } u \in \mathbb{S}^{d-1} := \{u : \|u\| = 1\}.$$

The support function uniquely determines the compact convex set K and is a fundamental object in convex geometry (see, for example, (Schneider, 2014, Section 1.7) or (Rockafellar, 1970, Section 13)). Consider now the problem of estimating an unknown compact, convex set K^* from observations $(X_1, Y_1), \dots, (X_n, Y_n)$ drawn according to the model:

$$Y_i = h_{K^*}(X_i) + \xi_i \quad \text{for } i = 1, \dots, n$$

where X_1, \dots, X_n are design points (fixed or random) and ξ_1, \dots, ξ_n are i.i.d. $N(0, 1)$. Recovery of K^* is a fundamental problem in geometric tomography (Prince and Willsky, 1990; Gardner, 2006). The natural estimator in this problem is the least squares estimator defined by

$$\hat{K}_n \in \operatorname{argmin}_{K \in \mathcal{C}} \sum_{i=1}^n (Y_i - h_K(X_i))^2$$

where \mathcal{C} denotes the class of all compact, convex sets in \mathbb{R}^d . Basic properties and algorithms for computing \hat{K}_n can be found in (Prince and Willsky, 1990; Lele et al., 1992) and (Kiderlen et al., 2008). Rigorous accuracy results for \hat{K}_n as an estimator of $K^* \subset B_d$ (here B_d denotes the unit-Euclidean ball) were proved in (Gardner et al., 2006). Specifically, (Gardner et al., 2006, Corollary 5.7) proved that, under the fixed design setting of any X_1, \dots, X_n (fixed points) that form a well-separated set (i.e. a set of n points on the unit-sphere that are at least $c(d)n^{-\frac{1}{d-1}}$ far from each other),

$$\|h_{\hat{K}_n} - h_{K^*}\|_{\mathbb{P}_n}^2 = O_P(\beta_n) \quad \text{where } \beta_n = \begin{cases} n^{-4/(d+3)} & \text{for } d = 2, 3, 4 \\ n^{-4/(d+3)} \cdot \log n & \text{for } d = 5 \\ n^{-2/(d-1)} & \text{for } d \geq 6. \end{cases} \quad (13)$$

Complementarily, Guntuboyina (2012) proved that the minimax rate of estimation in this problem (under the same fixed design version of the problem and the $\|\cdot\|_{\mathbb{P}_n}$ loss function) equals $\Theta_d(n^{-4/(d+3)})$ for all $d \geq 2$. These two results combined imply that the least squares estimator \hat{K}_n is minimax optimal for $d = 2, 3, 4$ and nearly minimax optimal (up to the logarithmic multiplicative factor $\log n$) for $d = 5$. However, there is a gap between the upper bound $n^{-2/(d-1)}$ on the rate of convergence of the least squares estimator (13) and the minimax rate $n^{-4/(d+3)}$ for $d \geq 6$. This gap has remained open since the paper Gardner et al. (2006) where it was suggested that the upper bound is accurate and that the least squares estimator is indeed minimax *suboptimal* for $d \geq 6$. The goal of this section is to confirm the conjecture of Gardner et al. (2006).

Specifically, we use Corollary 2 to prove that the LSE is suboptimal for $d \geq 6$, in the sense that there exist sets for which the rate of convergence for LSE is bounded from below by $n^{-2/(d-1)}$ up to a logarithmic multiplicative factor. Moreover, in this result, f_0 will be the d -dimensional ball, and the sets $\{K_m\}_{m=d+1}^\infty \subset \mathcal{C}$ will be “regular” polytopes that form an optimal approximation to the ball with $m \geq d + 1$ vertices (see Lemma 15 for more details).

Corollary 3 *Let $d \geq 6$. Suppose $X_1, \dots, X_n \sim \mathbb{P}$, where \mathbb{P} is the uniform distribution³ on \mathbb{S}^{d-1} . There exist a positive constants, C_d, γ_d depending only on d such that*

$$\sup_{K^* \in \mathcal{C}(1)} \mathbb{E}_{K^*} \|h_{\hat{K}_n} - h_{K^*}\|_{\mathbb{P}}^2 \geq C_d n^{-2/(d-1)} \log(n)^{-\gamma_d},$$

where $\mathcal{C}(1)$ is the set of all compact convex sets that are contained in the unit Euclidean ball in \mathbb{R}^d .

Observe that here the LSE is not restricted to $\mathcal{C}(1)$, i.e. $h_{\hat{K}_n}$ can return any convex set in \mathbb{R}^d .

Finally, the proof of the above corollary can be modified to hold for any fixed design of n well-separated points. The modification will follow from the fact that n well-separated points are a discrete approximation to the uniform measure on the sphere. Therefore, we settle the question in (Gardner et al., 2006).

Notation Throughout this text, c, C with subscripts are positive absolute constants that do not depend on the dimension d . Additionally, positive constants that only depend on the dimension or on p are explicitly denoted, respectively, by $c(d), C(d), c(p), C(p)$. These constants may change from line to line, and from section to section. B_d denotes the unit Euclidean ball in \mathbb{R}^d . $\tilde{O}(\cdot), \tilde{\Omega}(\cdot)$ denotes behavior up to logarithmic factors in n or ϵ .

3. Proofs

3.1. Proof of Theorem 1

Our main technical tool is the following important result of Chatterjee (2014) which gives sharp upper and lower bounds for the accuracy of an LSE over a convex family of functions in the fixed-design setting. We use the following notation in this result.

Theorem 4 (Theorem 1.1 Chatterjee (2014)) *Let \mathcal{F} be a convex family of functions and consider the LSE for the fixed design setting. Let*

$$t_f := \operatorname{argmax}_{t \geq 0} H_f(t) \quad \text{where } H_f(t) := \bar{\mathcal{W}}(B_{\mathbb{P}_n}(f, t)) - \frac{t^2}{2}. \quad (14)$$

Then, t_f is unique, $H_f(\cdot)$ is a concave function, and we have

$$\mathbb{P} \left\{ 0.5t_f^2 \leq \|\hat{f}_n - f\|_{\mathbb{P}_n}^2 \leq 2t_f^2 \right\} \geq 1 - 3p_n, \quad (15)$$

where $p_n = \exp\left(-cnt_f^2\right)$.

Remark 5 *Based on the proof of (Chatterjee, 2014, Thm 1.1), the assumption on $g = (g_1, \dots, g_n) \sim N(0, I_{n \times n})$ can be relaxed to any isotropic random vector that satisfies a convex Lipschitz concentration inequality (Adamczak, 2015; Boucheron et al., 2013). Also, the same result for an isotropic log-concave noise with $p_n = \exp\left(-c\sqrt{nt_f^2}\right)$ follows from the recent result of Lee and Vempala (2018), which is a concentration inequality for Lipschitz functions for isotropic log-concave distributions. Finally, we remark that for zero-mean ξ with variance σ^2 , all our results scale by σ^2 .*

3. Or for any density $g(x)$ on the sphere, such that $g(x) \geq c_1$ for all $x \in \mathbb{S}^{d-1}$.

The proof of Theorem 1 is reduced to the following ‘‘two points’’ lemma that is mainly based on Theorem 4. The notation $f_{(n)}$ is to emphasize the fact that $f_{(n)}$ is a function that is chosen based on the number of samples.

Lemma 6 [Lower bound for fixed design] *Let $f_0, f_{(n)} \in \mathcal{F}$, and $r_n \geq Cn^{-1/4}$, δ_n, w_n, s_n be positive constants. Also, let $w'_n = \max\{\frac{r_n^2}{8\delta_n}, w_n\}$, and assume the following:*

- $\bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_0, s_n)) \geq r_n^2$
- $\|f_{(n)} - f_0\|_{\mathbb{P}_n} \leq \delta_n$, and $r_n \leq 2\delta_n$, $s_n \leq \delta_n$.
- $\bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, t)) \leq w'_n t$ for some $t \leq \frac{r_n^4}{16\delta_n^2 w'_n}$.

Then, when $Y_i = f_{(n)}(x_i) + \xi_i$, the following holds with high probability

$$\|\hat{f}_n - f_{(n)}\|_{\mathbb{P}_n}^2 \geq c \cdot \min \left\{ \frac{r_n^8}{\delta_n^4 w_n^2}, \frac{r_n^4}{\delta_n^2} \right\}.$$

Proof [Proof of Lemma 6] By Theorem 4, it is enough to show that $t_{f_{(n)}}$ is greater than $\frac{r_n^4}{16\delta_n^2 w_n}$. Namely, the functional $H_{f_{(n)}}(\cdot)$ attains its unique maximum on a $t_{f_{(n)}} \geq \frac{r_n^4}{16\delta_n^2 w_n}$. First, using the second assumption and the upper bound on δ_n , we see that

$$B_{\mathbb{P}_n}(f_0, s_n) \subseteq B_{\mathbb{P}_n}(f_{(n)}, \|f_{(n)} - f_0\|_{\mathbb{P}_n} + s_n) \subseteq B_{\mathbb{P}_n}(f_{(n)}, 2\delta_n).$$

By the convexity of \mathcal{F} , the function $t \mapsto \bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, t))/t$ is nonincreasing. Therefore, by the first assumption and the above inclusion, for all $t \leq 2\delta_n$

$$\bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, t)) \geq \frac{r_n^2}{2\delta_n} t.$$

Now, by the last equation and the second assumption, we know that for $t_1 = \frac{r_n^2}{2\delta_n} \leq 2\delta_n$

$$H_{f_{(n)}}\left(\frac{r_n^2}{2\delta_n}\right) = \bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, \frac{r_n^2}{2\delta_n})) - \frac{\left(\frac{r_n^2}{2\delta_n}\right)^2}{2} \geq \frac{r_n^2}{2\delta_n} \cdot \left(\frac{r_n^2}{2\delta_n}\right) - \frac{\left(\frac{r_n^2}{2\delta_n}\right)^2}{2} \geq \frac{r_n^4}{8\delta_n^2}.$$

Finally, by the last assumption, we know that for $t_2 = \frac{r_n^4}{16\delta_n^2 w'_n} < t_1$ the following holds:

$$H_{f_{(n)}}\left(\frac{r_n^4}{16\delta_n^2 w'_n}\right) \leq \bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, \frac{r_n^4}{16\delta_n^2 w'_n})) \leq w'_n \frac{r_n^4}{16\delta_n^2 w'_n} = \frac{r_n^4}{16\delta_n^2}.$$

Since $H_{f_{(n)}}(\cdot)$ is concave in t , we conclude by the last two equations that $t_{f_{(n)}} \geq \frac{r_n^4}{16\delta_n^2 w'_n}$ and the claim follows from Theorem 4. \blacksquare

In addition to Lemma 6, we need the following two standard facts (which can be found, for example, in Koltchinskii (2011)) to prove Theorem 1.

Lemma 7 (Sudakov Minoration) *There exists a universal positive constant C such that the following holds for any class \mathcal{F} of real-valued functions:*

$$\bar{\mathcal{W}}(\mathcal{F}) \geq \frac{C}{\sqrt{n}} \sup_{\epsilon > 0} \left\{ \epsilon \sqrt{\log \mathcal{N}(\epsilon, \mathcal{F}, \mathbb{P}_n)} \right\}.$$

Lemma 8 (Dudley Integral) *There exists a universal positive constant C such that the following holds for any class \mathcal{F} of real-valued functions:*

$$\bar{\mathcal{W}}(\mathcal{F}) \leq C \inf_{\epsilon > 0} \left(\epsilon + \frac{1}{\sqrt{n}} \int_{\epsilon}^{\text{diam}(\mathcal{F}, \mathbb{P}_n)/2} \sqrt{\log \mathcal{N}(u, \mathcal{F}, \mathbb{P}_n)} du \right),$$

where $\text{diam}(\mathcal{F}, \mathbb{P}_n)$ is the diameter of \mathcal{F} with respect to $L^2(\mathbb{P}_n)$.

Proof [Proof of Theorem 1] We aim to apply Lemma 6. In order to satisfy its first assumption, we need to estimate the Gaussian width of $B_{\mathbb{P}_n}(f_0, C_3 n^{-1/p} \log(n)^\gamma)$. By Lemma 7 and Eq. (10)

$$\begin{aligned} \bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_0, C_3 n^{-1/p} \log(n)^\gamma)) &\geq \frac{C}{\sqrt{n}} \sup_{\epsilon \geq C_1 n^{-1/p} \log(n)^\gamma} \left\{ \epsilon \sqrt{\log \mathcal{N}(\epsilon, B_{\mathbb{P}_n}(f_0, C_3 n^{-1/p} \log(n)^\gamma, \mathbb{P}_n)} \right\} \\ &\geq c_3(p) n^{-\frac{1}{p}} \log(n)^{-\gamma(\frac{p}{2}-1)}, \end{aligned} \tag{16}$$

where we chose $\epsilon = C_1 n^{-1/p} \log(n)^\gamma$. Now, we aim to upper bound $\bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, t))$ when

$$t = c_1(p) n^{-\frac{1}{2p}} \log(n)^{-\frac{\beta}{p} - \frac{3\gamma}{2}(\frac{p}{2}-1)},$$

for some suitable $c_1(p)$. By using Lemma 8, and Eq. (11), we see that

$$\begin{aligned} \bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, t)) &\leq C \inf_{\epsilon > 0} \left(\epsilon + \frac{1}{\sqrt{n}} \int_{\epsilon}^t \sqrt{\log \mathcal{N}(u, B_{\mathbb{P}_n}(f_{(n)}, t), \mathbb{P}_n)} du \right) \\ &\leq C \inf_{\epsilon > 0} \left(\epsilon + \frac{C_2(p)}{\sqrt{n}} \int_{\epsilon}^t \left(\frac{t n^{\frac{1}{2p}} \log(n)^{\frac{\gamma}{2}(\frac{p}{2}-1)} \log(u^{-1})^{\frac{\beta}{p}}}{u} \right)^{\frac{p}{2}} du \right) \\ &\leq C \inf_{\epsilon > 0} \left(\epsilon + \frac{C_3(p)}{\sqrt{n}} \left(\frac{t n^{\frac{1}{2p}} \log(\epsilon^{-1})^{\frac{\beta}{p}} \log(n)^{\frac{\gamma}{2}(\frac{p}{2}-1)}}{\epsilon} \right)^{\frac{p}{2}} \epsilon \right) \\ &\leq C_4(p) n^{-\frac{1}{2p}} \log(n)^{\frac{\beta}{p} + \frac{\gamma}{2}(\frac{p}{2}-1)} t, \end{aligned} \tag{17}$$

where we used the fact that $p > 2$ and set

$$\epsilon = C_1(p) n^{-\frac{1}{2p}} \log(n)^{\frac{\beta}{p} + \frac{\gamma}{2}(\frac{p}{2}-1)} t = C_2(p) n^{-\frac{1}{p}} \log(n)^{-\gamma(\frac{p}{2}-1)}.$$

Hence, it is enough to assume (7) in place of (6). Finally, we can apply Lemma 6 with the following parameters:

$$\|f_{(n)} - f_0\|_{\mathbb{P}_n} \leq c_1(p) n^{-\frac{1}{2p}} \log(n)^{-\frac{\gamma}{2}(\frac{p}{2}-1)} =: \delta_n,$$

where we used Eq. (6). Also, we set $s_n = C_5(p)n^{-1/p} \log(n)^\gamma$,

$$r_n^2 = \min\{c_3(p), c_1(p)/8\} n^{-\frac{1}{p}} \log(n)^{-\gamma(\frac{p}{2}-1)}$$

and

$$w_n = C_4(p)n^{-\frac{1}{2p}} \log(n)^{\frac{\beta}{p} + \frac{\gamma}{2}(\frac{p}{2}-1)}.$$

Therefore, Lemma 6 gives that

$$\mathcal{R}_{\mathbb{P}}(\hat{f}_n, \mathcal{F}) \geq c_5(p) \left(\frac{n^{-\frac{1}{p}} \log(n)^{-\gamma(\frac{p}{2}-1)}}{n^{-\frac{1}{2p}} \log(n)^{\frac{\beta}{p} + \frac{\gamma}{2}(\frac{p}{2}-1)}} \right)^2 \gtrsim n^{-\frac{1}{p}} \log(n)^{-\frac{2\beta}{p} - 3\gamma(\frac{p}{2}-1)},$$

and the claim follows. ■

3.2. Proof of Corollary 2

For completeness, we state a consequence of (van de Geer, 2000, Thm 5.11) (see Appendix for a proof).

Lemma 9 *Let \mathcal{G} be a family of functions uniformly bounded by one, and \mathbb{P} be some probability measure. For any $f \in \mathcal{G}$ and $t \geq 0$, let $\varepsilon(f, t)$ be the stationary point of*

$$\frac{1}{\sqrt{n}} \int_{\frac{\varepsilon}{c}}^t \sqrt{\log \mathcal{N}_{[]} (u, B_{\mathbb{P}}(f, t), \mathbb{P})} du = \delta$$

for some absolute constant $c > 0$. Then, for any $\varepsilon \geq \varepsilon(f, t)$ the following holds with probability of at least $1 - C \exp(-c n \varepsilon^2)$:

$$\bar{\mathcal{W}}(B_{\mathbb{P}}(f, t)) \lesssim \varepsilon + n^{-1/2}.$$

Proof [Proof of Corollary 2] First, we take $m = C(p)\sqrt{n} \log(n)^{\frac{p\gamma}{2}(\frac{p}{2}-1)}$ in (11) and set $f_{(n)} := f_{(m)}$. By using Eq. (8) and the estimate $d_n \lesssim (\log n)^2 n^{-1/p}$ from (Rakhlin et al., 2017), we see that with high probability,

$$\|f_{(n)} - f_0\|_{\mathbb{P}_n} \leq 2\|f_{(n)} - f_0\|_{\mathbb{P}} + C \log^2(n) n^{-\frac{1}{p}} \lesssim n^{-\frac{1}{2p}} \log(n)^{\frac{\gamma}{2}(\frac{p}{2}-1)}.$$

Thus Eq. (5) holds for $f_{(n)}$. Next, by using Eq. (8), it is easy to see that Eq. (4) hold with $\gamma = 2$. It remains to show that (7) holds.

First, by Eq. (8), we know that for $t \geq C d_n$ the following holds:

$$\bar{\mathcal{W}}(B_{\mathbb{P}_n}(f_{(n)}, t)) \leq \bar{\mathcal{W}}(B_{\mathbb{P}}(f_{(n)}, 2t + d_n)) \leq \bar{\mathcal{W}}(B_{\mathbb{P}}(f_{(n)}, 4t)).$$

To estimate the last term, we use our bracketing assumption along with the uniform boundedness of functions in \mathcal{F} . These assumptions let us apply Lemma 9 that yields an upper bound on this term. In order to upper bound $\varepsilon(f_{(n)}, 4t)$ (the fixed point of Lemma 9), we use Eq. (12) and derive that

$$\frac{1}{\sqrt{n}} \int_{\frac{\varepsilon}{c}}^{4t} \left(\frac{t n^{\frac{1}{2p}} \log(n)^{\frac{\gamma}{2}(\frac{p}{2}-1)} \log(u^{-1})^{\frac{\beta}{p}}}{u} \right)^{\frac{p}{2}} du \lesssim \frac{1}{\sqrt{n}} \left(\frac{t n^{\frac{1}{2p}} \log(n)^{\frac{\beta}{p}} \log(n)^{\frac{\gamma}{2}(\frac{p}{2}-1)}}{\varepsilon} \right)^{\frac{p}{2}} \varepsilon$$

where in the last inequality we used the fact that $p > 2$ and $\varepsilon \geq 1/n$. Setting the right-hand side equal to ε , we find an upper bound on the fixed point

$$\varepsilon(f_{(n)}, 4t) \lesssim n^{-\frac{1}{2p}} \log^{\frac{\beta}{p} + \frac{\gamma}{2}(\frac{p}{2}-1)}(n)t.$$

We now set $t \sim n^{-\frac{1}{2p}} \log(n)^{-\frac{\beta}{p} - \frac{3sp\gamma}{2}}$ and apply Lemma 9 with $\varepsilon \sim n^{-1/p} \log(n)^{-\gamma(p/2-1)}$, concluding that with probability of at least $1 - C \exp(-c\tilde{\Omega}(n^{1-2/p}))$

$$\bar{W}(B_{\mathbb{P}}(f_{(n)}, 4t)) \lesssim \varepsilon \sim n^{-\frac{1}{2p}} \log^{\frac{\beta}{p} + \frac{\gamma}{2}(\frac{p}{2}-1)}(n)t.$$

Hence, all the assumptions of Theorem 1 hold with high probability. Therefore, by Eq. (8) and Theorem 1, we conclude that

$$\begin{aligned} \|\hat{f}_n - f^*\|_{\mathbb{P}}^2 &\geq \frac{1}{4} \|\hat{f}_n - f^*\|_{\mathbb{P}_n}^2 - 2d_n^2 \geq \frac{1}{8} \|\hat{f}_n - f^*\|_{\mathbb{P}_n}^2 \\ &\gtrsim n^{-\frac{1}{p}} \log(n)^{-\frac{2\beta}{p} - 3\gamma(\frac{p}{2}-1)}, \end{aligned}$$

and the claim follows. ■

3.3. Proof of Corollary 3

Throughout this subsection \mathbb{P} denotes the uniform measure on the sphere. In this proof we will use the following auxiliary lemmas:

Lemma 10 (Bronshtein (1976); Dudley (1999)) *The following holds for all $0 < \epsilon < 1$,*

$$2^{-d} \epsilon^{-\frac{d-1}{2}} \leq \log \mathcal{N}_2(\epsilon, \mathcal{C}(1), \mathbb{P}) \leq \log \mathcal{N}_{\infty}(\epsilon, \mathcal{C}(1), \mathbb{P}) \leq Cd^{5/2} \epsilon^{-\frac{d-1}{2}},$$

where $\log \mathcal{N}_{\infty}(\epsilon, \mathcal{C}(1), \mathbb{P})$ denotes covering with respect to $\|\cdot\|_{\infty}$.

Lemma 11 (Bronshtein (1976); Dudley (1999)) *Let h_{B_d} be the support of function of the unit ball in \mathbb{R}^d , that is $h_{B_d} \equiv 1$. Then for any $0 \leq \epsilon \leq c$, there exists a set of cardinality $M(\epsilon, d) \geq 2^{10-d} \epsilon^{-(d-1)/2}$, denoted by $h_1, \dots, h_{K_{M(\epsilon, d)}}$, that has the following properties:*

- $\forall 1 \leq i < j \leq M(\epsilon, d): \|h_{K_i} - h_{K_j}\|_{\mathbb{P}} \geq \epsilon.$
- $\forall 1 \leq i \leq M(\epsilon, d): \|h_{K_i} - h_{B_d}\|_{\mathbb{P}} \leq 8\epsilon.$

Definition 12 *We say that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a k -piecewise simplicial linear if f is a convex piecewise linear function and its support can be written as a union of k -simplicials, and f is linear in each of them.*

For the following lemma, let \mathcal{P}_d be the regular simplex with $d + 1$ faces that contains the unit Euclidean ball, with vertices at distance d from the origin.

Lemma 13 (Theorem 4.4 in Kur et al. (2020)) *Let $\mathcal{F}(\Gamma, \mathcal{P}_d)$ be the family of all the convex functions that are uniformly bounded by $\Gamma > 0$, and their domain is a simplex. Let $f_k \in \mathcal{F}(\Gamma, \mathcal{P}_d)$ be k piecewise simplicial linear. Then, the following bound holds for $0 \leq \epsilon \leq t$:*

$$\mathcal{N}_{\square}(\epsilon, B_{\text{Unif}}(f_k, t), \text{Unif}) \leq C(d) (t/\epsilon)^{d/2} \log^{d+1}(\epsilon^{-1}) \log(\Gamma)$$

where $B_{\text{Unif}}(f_k, t) \subset \mathcal{F}(M, \mathcal{P}_d)$ and Unif denotes the Lebesgue volume measure.

In order to prove Corollary 3, we first prove that when we restrict the LSE to $\mathcal{C}(C_1 \cdot d)$ (here C_1 is some absolute constant) we achieve the desired bound:

Lemma 14 *For every $d \geq 6, n \geq C(d)$, the following holds*

$$\operatorname{argmin}_{K \in \mathcal{C}(C_1 \cdot d)} \sum_{i=1}^n (Y_i - h_K(X_i))^2 \geq c(d) \log^{\gamma_d}(n) n^{-2/(d-1)}, \quad (18)$$

whenever $K^* \in \mathcal{C}(1)$.

Then, in Subsection B.2, we prove that with high probability,

$$\operatorname{argmin}_{K \in \mathcal{C}} \sum_{i=1}^n (Y_i - h_K(X_i))^2 = \operatorname{argmin}_{K \in \mathcal{C}(C_1 \cdot d)} \sum_{i=1}^n (Y_i - h_K(X_i))^2$$

whenever $K^* \in \mathcal{C}(1)$, and therefore Corollary 3 follows.

Proof [Proof of Lemma 14] Let $h_K \in \mathcal{C}(1)$ and $t > 0$, denote by

$$B_{\mathbb{P}}(h_K, t) := \{L \in \mathcal{C}(C_1 \cdot d) : \|h_K - h_L\|_{\mathbb{P}} \leq t\}.$$

We will need the following Lemma (its proof appears in Section B.1).

Lemma 15 (The regular polytope lemma) *For every $m \geq 4^d, d \geq 2$ there exists a polytope $P_{m,d} \subset B_d$ with m vertices that satisfies the following:*

1. $\|h_{P_{m,d}}(x) - 1\|_{\infty} \leq C_1 m^{-\frac{2}{d-1}}$
2. $\log \mathcal{N}_{\square}(\epsilon, B_{\mathbb{P}}(h_{P_{m,d}}, t), \mathbb{P}) \leq C(d) m \log(\epsilon^{-1})^d \left(\frac{t}{\epsilon}\right)^{(d-1)/2}$ for all $0 \leq \epsilon \leq t$.

Now, observe that the family is uniformly bounded by $C_1 \cdot d$. Therefore, we can invoke Corollary 2 with $f_{(m)} = h_{P_{m,d}}, \beta = d, f_0 = h_{B_d}$, and $p = \frac{d-1}{2}$. By Lemmas 10, 11, and 15, we satisfy all the corollary's assumptions (observe that an upper bound on ∞ -covering implies that Koltchinski-Pollard entropy satisfies the same upper bound). Thus $\mathcal{R}(\hat{h}_{\hat{K}_n}, \mathcal{C}(1)) \geq \tilde{\Omega}(n^{-\frac{2}{d-1}})$, and the claim follows. \blacksquare

Acknowledgments

The first and second authors acknowledge the support from NSF under award DMS-1953181, and from the Center for Minds, Brains and Machines (CBMM) funded by NSF award CCF-1231216. The third author is supported by NSF CAREER Grant DMS-1654589. We also acknowledge the anonymous referees for their useful comments.

References

- Radoslaw Adamczak. A note on the hanson-wright inequality for random vectors with dependencies. *Electronic Communications in Probability*, 20, 2015.
- Shiri Artstein-Avidan, Apostolos Giannopoulos, and Vitali D Milman. *Asymptotic geometric analysis, Part I*, volume 202. American Mathematical Soc., 2015.
- Gábor Balázs, András György, and Csaba Szepesvári. Near-optimal max-affine estimators for convex regression. In *AISTATS*, 2015.
- Andrew Barron, Lucien Birgé, and Pascal Massart. Risk bounds for model selection via penalization. *Probability theory and related fields*, 113(3):301–413, 1999.
- Peter L Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.
- Lucien Birgé. Model selection via testing: an alternative to (penalized) maximum likelihood estimators. In *Annales de l’IHP Probabilités et statistiques*, volume 42, pages 273–325, 2006.
- Lucien Birgé and Pascal Massart. Rates of convergence for minimum contrast estimators. *Probability Theory and Related Fields*, 97(1-2):113–150, 1993.
- Károly Böröczky and Gergely Wintsche. Covering the sphere by equal spherical balls. In *Discrete and computational geometry*, pages 235–251. Springer, 2003.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Olivier Bousquet. Concentration inequalities and empirical processes theory applied to the analysis of learning algorithms. 2002.
- EM Bronshtein. ε -entropy of convex sets and functions. *Siberian Mathematical Journal*, 17(3): 393–398, 1976.
- Victor-Emmanuel Brunel. Adaptive estimation of convex polytopes and convex sets from noisy data. *Electronic Journal of Statistics*, 7:1301–1327, 2013.
- Victor-Emmanuel Brunel. Adaptive estimation of convex and polytopal density support. *Probability Theory and Related Fields*, 164(1-2):1–16, 2016.
- Victor-Emmanuel Brunel, Jason M Klusowski, and Dana Yang. Estimation of convex supports from noisy measurements. *arXiv preprint arXiv:1804.09879*, 2018.
- Timothy Carpenter, Ilias Diakonikolas, Anastasios Sidiropoulos, and Alistair Stewart. Near-optimal sample complexity bounds for maximum likelihood estimation of multivariate log-concave densities. In *Conference On Learning Theory*, pages 1234–1262, 2018.
- Sourav Chatterjee. A new perspective on least squares under convex constraint. *The Annals of Statistics*, 42(6):2340–2381, 2014.

- Richard M Dudley. *Uniform central limit theorems*. Number 63. Cambridge university press, 1999.
- Nicholas I Fisher, Peter Hall, Berwin A Turlach, and GS Watson. On the estimation of a convex set from noisy data on its support function. *Journal of the American Statistical Association*, 92(437):84–91, 1997.
- Fuchang Gao and Jon A Wellner. Entropy of convex functions on \mathbb{R}^d . *Constructive approximation*, 46(3):565–592, 2017.
- Richard J. Gardner. *Geometric Tomography*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2 edition, 2006.
- Richard J Gardner, Markus Kiderlen, and Peyman Milanfar. Convergence of algorithms for reconstructing convex bodies and directional measures. *The Annals of Statistics*, 34(3):1331–1374, 2006.
- Adityanand Guntuboyina. Optimal rates of convergence for convex set estimation from support functions. *The Annals of Statistics*, 40(1):385–411, 2012.
- Qiyang Han. Global empirical risk minimizers with” shape constraints” are rate optimal in general dimensions. *arXiv preprint arXiv:1905.12823*, 2019.
- Qiyang Han, Tengyao Wang, Sabyasachi Chatterjee, and Richard J Samworth. Isotonic regression in general dimensions. *The Annals of Statistics*, 47(5):2440–2471, 2019.
- Markus Kiderlen et al. A new algorithm for 3d reconstruction from support functions. *IEEE transactions on pattern analysis and machine intelligence*, 31(3):556–562, 2008.
- Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems: Ecole d’Eté de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer Science & Business Media, 2011.
- Vladimir Koltchinskii and Shahar Mendelson. Bounding the smallest singular value of a random matrix without concentration. *arXiv preprint arXiv:1312.3580*, 2013.
- Gil Kur, Yuval Dagan, and Alexander Rakhlin. Optimality of maximum likelihood for log-concave density estimation and bounded convex regression. *arXiv preprint arXiv:1903.05315*, 2019.
- Gil Kur, Fuchang Gao, Adityanand Guntuboyina, and Bodhisattva Sen. Convex regression in multidimensions: Suboptimality of least squares estimators. *arXiv preprint arXiv:2006.02044*, 2020.
- Yin Tat Lee and Santosh S Vempala. Stochastic localization+ stieljtjes barrier= tight bound for log-sobolev. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1122–1129. ACM, 2018.
- Avinash Shreedhar Lele, Sanjeev R Kulkarni, and Alan S Willsky. Convex-polygon estimation from support-line measurements and applications to target reconstruction from laser-radar data. *JOSA A*, 9(10):1693–1714, 1992.
- Shahar Mendelson. Learning without concentration. In *Conference on Learning Theory*, pages 25–39, 2014.

- Shahar Mendelson. On aggregation for heavy-tailed classes. *Probability Theory and Related Fields*, 168(3-4):641–674, 2017.
- Jerry Ladd Prince and Alan S Willsky. Reconstructing convex sets from support line measurements. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(4):377–389, 1990.
- Alexander Rakhlin, Karthik Sridharan, and Alexandre B Tsybakov. Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 23(2):789–824, 2017.
- R Tyrrell Rockafellar. *Convex analysis*. Number 28. Princeton university press, 1970.
- Rolf Schneider. *Convex bodies: the Brunn–Minkowski theory*. Number 151. Cambridge university press, 2014.
- Rolf Schneider and Wolfgang Weil. *Stochastic and integral geometry*. Springer Science & Business Media, 2008.
- Yong Sheng Soh and Venkat Chandrasekaran. Fitting tractable convex sets to support function evaluations. *arXiv preprint arXiv:1903.04194*, 2019.
- Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2003.
- Sara A van de Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- Aad W van der Vaart and Jon A Wellner. Weak convergence. In *Weak convergence and empirical processes*, pages 16–28. Springer, 1996.
- Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, pages 1564–1599, 1999.

Appendix A. Auxiliary results and basic facts

Proof [Proof of Lemma 9] Under the conditions of the Lemma, (van de Geer, 2000, Thm 5.11) implies that for the fixed point ε ,

$$\sup_{g \in B_{\mathbb{P}}(f, t)} |\mathbb{P}_n(g) - \mathbb{P}(g)| \lesssim \varepsilon,$$

with probability of at least $1 - C \exp(-c(n(\varepsilon/t)^2)) \geq 1 - C \exp(-cn\varepsilon^2)$ since $t \leq 1$ due to uniform boundedness. Now, using de-symetrization argument (e.g. (Koltchinskii, 2011)), we know that

$$\mathbb{E} \bar{W}(B_{\mathbb{P}}(f, t)) - \frac{1}{\sqrt{n}} \lesssim \mathbb{E} \sup_{g \in B_{\mathbb{P}}(f, t)} |\mathbb{P}_n(g) - \mathbb{P}(g)| \lesssim \varepsilon.$$

Therefore, in expectation we have the desired bound. Since the class is uniformly bounded by one, we can apply McDiarmid’s inequality which gives

$$\Pr \left(\bar{W}(B_{\mathbb{P}}(f, t)) - C(n^{-1/2} + \varepsilon) \geq u \right) \leq \exp(-cnu^2)$$

and by setting $u = C(\varepsilon + n^{-1/2})$ the claim follows. ■

Lemma 16 (Jacobian of the radial function, see for example [Schneider and Weil \(2008\)](#)) *Let $H = \{x \in \mathbb{R}^d : x^T u = h\}$ be a $d - 1$ hyper-plane and any integrable $f : H \rightarrow \mathbb{R}$, the following holds for the radial function $R(x) = x/\|x\|$:*

$$\int_H f(x) \frac{u^T x}{\|x\|^d} dx = \int_{R(H)} f(R^{-1}(x)) dS(x).$$

Lemma 17 (Basic facts on the support function, see for example [Artstein-Avidan et al., 2015](#))

- For every $K \in \mathcal{C}(1)$, the function h_K can be extended to the whole of \mathbb{R}^d via $h_K(x) = \|x\| h_K(x/\|x\|)$ and this extension makes h_K a convex and 1-Lipschitz function on \mathbb{R}^d .
- For any $\lambda > 0$, $K \in \mathcal{C}$ the following holds for all $x \in \mathbb{R}^d$: $h_{\lambda K}(x) = \lambda h_K(x)$.

Appendix B.

B.1. Proof of Lemma 15

In this proof, we denote by $B_G(x, r)$ a geodesic ball on the sphere with center x and radius r . Let $\mathcal{V} := \{v_i\}_{i=1}^s$ be a set of vertices of size $m \leq s \leq 400d \ln(d)m$ with norm one. This set is a $m^{-\frac{1}{d-1}}$ -net on \mathbb{S}^{d-1} with the following property: In any geodesic ball with radius of $m^{-\frac{1}{d-1}}$ on the sphere there are at most $400d \ln(d)$ points. When $m \geq C^d$ such a set exists by ([Böröczky and Wintsche, 2003](#)).

Now, we define a polytope $P_{m,d} := \text{conv}\{\mathcal{V}\}$. First, observe that for each $x \in \mathbb{R}^d$ the following holds:

$$h_{P_{m,d}}(x) = \max_{v \in \mathcal{V}} \|x\| \|v\| \cos(\angle(x, v)) = \|x\| \cos(\angle(x, v_\pi)), \quad (19)$$

where v_π is the closest point in the net to $x/\|x\|$. Then, using the fact that $\cos(t) = 1 - t^2/2 + O(t^4)$, we know that when $m \geq C_1^d$, for all $u \in \mathbb{S}^{d-1}$

$$0 \leq 1 - h_{P_{m,d}}(u) \leq C m^{-2/(d-1)}$$

and the first part of the Lemma follows.

Now, we prove the last part of Lemma 15, that is an upper bound on the entropy numbers. Let us consider the support function when it is restricted to the facets of regular simplex \mathcal{P}_d , denoted by S_k , $1 \leq k \leq d + 1$. Now we prove the following:

Lemma 18 *For any $k \in \{1, \dots, d + 1\}$, $h_{P_{m,d}} : S_k \rightarrow \mathbb{R}$ is at most $C \ln(d)m$ piecewise linear, and every piece has at most $C(d \ln(d))^2 4^{d-1} (d - 2)$ -facets.*

Proof

First, we use the Voronoi cells of the vertices \mathcal{V} when they are restricted to \mathbb{S}^{d-1} , that is each cell, denoted by $C(v_i)$, $1 \leq i \leq |\mathcal{V}|$, is defined by

$$C(v_i) := \{x \in \mathbb{S}^{d-1} : \|v_i - x\| \leq \|v_j - x\| \ \forall j \neq i\}.$$




Figure 1: Illustration of the proof. The function $h_{P_{m,d}}$ is linear on the set $R^{-1}(C(v_i))$ (in blue) and convex piece-wise linear on S_k .

By Eq. (19) we know that for each $x \in \mathbb{R}^d$ the value of $h_{P_{m,d}}$ is attained on the closest $v_i \in \mathcal{V}$ to $x/\|x\|$ (since all the vertices of $P_{n,d}$ have the same norm). Thus, we can write

$$S_k = \bigcup_{i=1}^{|\mathcal{V}|} R^{-1}(C(v_i)) \cap S_k$$

where R denotes the radial function from \mathcal{P}_d to \mathbb{S}^{d-1} . Moreover, observe that

$$h_{P_{m,d}} : R^{-1}(C(v_i))(x) \equiv v_i^T x.$$

Since $h_{P_{m,d}} : S_k \rightarrow \mathbb{R}$ is a convex function, we conclude that it is also piecewise linear. Due to the regularity of the net \mathcal{V} , we know that number of pieces can be bounded by $2|\mathcal{V}|/(d+1) \leq C \ln(d)m$.

Next, observe that the number of $d-2$ facets of each piece, which corresponds to some $v_i \in \mathcal{V}$, is determined by the number of neighbors of the Voronoi cell $C(v_i)$. By the construction of \mathcal{V} , it can be bounded by $C(d \ln(d))^2 4^{d-1}$. To see this, since \mathcal{V} is a $m^{-1/(d-1)}$ -net, clearly all the neighbors of $v_i \in \mathcal{V}$ lie in $B_G(v_i, 4\epsilon)$. Moreover, we can bound the number of vertices in this ball by

$$\begin{aligned} |\{v \in \mathcal{V} : v \in B_G(v_i, 4\epsilon)\}| &\leq 400d \ln(d) \cdot |\mathcal{N}(\epsilon, B_G(v_i, 4\epsilon), \text{Unif})| \\ &\leq C(d \ln(d))^2 4^{d-1}, \end{aligned}$$

and the claim follows. ■

Now, we are ready to prove Lemma 15. It is easy to see that each $h_{P_{m,d}} : S_k \rightarrow \mathbb{R}$ is uniformly bounded by d (since $K \subset B_d$ implies that h_K restricted to S_k is bounded by d). Moreover, since we assume that all $h_L \in B_{\mathbb{P}}(h_K, t)$ are bounded in $C_1 d$, we also know that our family is uniformly bounded by $C d^2$. Now, recall that all of $h_{P_{m,d}}$ $C\sqrt{m}/(d-1)$ -pieces have at most $(C d \ln(d))^2 4^{d-1}$ facets, hence, it is also $C(d)\sqrt{m}$ piecewise simplicial linear. Thus, we can apply Lemma 13 and find a $(d+1)^{-1/2}\epsilon$ -net, denoted by $\{f_{k,j}\}_{j=1}^{S_\epsilon}$, with respect to $\text{Unif}(S_k)$, where

$$S_\epsilon := \log \mathcal{N}_{\square}((d+1)^{-1/2}\epsilon, B_{\text{Unif}(S_k)}(h_{P_{m,d}}, t), \text{Unif}(S_k)) \leq C(d)m \log(\epsilon^{-1})^d \left(\frac{t}{\epsilon}\right)^{(d-1)/2}.$$

Recall that we aim to bound the entropy numbers of the the support function on the sphere. Therefore, we use the radial function to project each facet S_k onto the sphere, and show that the set of functions

$$\{\|R^{-1}(x)\|^{-1}f_{k,j}(R^{-1}(x))\}_{j=1}^{S_\epsilon}$$

forms an $(d+1)^{-1/2}\epsilon$ -net on $B_{\text{Unif}(R(S_k))}(h_{P_{m,d}}, t)$ with respect to uniform measure.

Let $h_K \in \mathcal{C}(C_1 \cdot d)$ (which is convex and bounded by one) that is at least $(d+1)^{-1/2}\epsilon$ far from $h_{P_{m,d}}$, when the support functions are restricted to $R(S_k)$. Denote by $h_{\pi,K}$ the closest member to the h_K with respect to $\text{Unif}(S_k)$ in the aforementioned set. Then,

$$\begin{aligned} & \sqrt{\int_{R(S_k)} (h_K(x) - \|R^{-1}(x)\|^{-1}h_{\pi,K}(R^{-1}(x)))^2 dS(x)} \\ &= \sqrt{\int_{R(S_k)} \|R^{-1}(x)\|^{-2} (h_K(R^{-1}(x)) - h_{\pi,K}(R^{-1}(x)))^2 dS(x)} \\ &= \sqrt{\int_{S_k} (h_K(x) - h_{\pi,K}(x))^2 \frac{u_i^T x}{\|x\|^{d+2}} dx} \leq \sqrt{\int_{S_k} (h_K(x) - h_{\pi,K}(x))^2 dx} \leq (d+1)^{-1/2}\epsilon, \end{aligned}$$

where we used the homogeneity of the support function, and Lemma 16. Therefore, we found our desired net. Now, we can use (Gao and Wellner, 2017, Lemma 2.8), and conclude that for all $0 \leq \epsilon \leq t$

$$\begin{aligned} \log \mathcal{N}_{[]}(\epsilon, B_{\mathbb{P}}(h_{P_{m,d}}, t), \mathbb{P}) &\leq \sum_{k=1}^{d+1} \log \mathcal{N}_{[]}(\frac{\epsilon}{\sqrt{d+1}}, B_{\text{Unif}(R(S_k))}(h_{P_{m,d}}, t), \text{Unif}(R(S_k))) \\ &\leq \sum_{k=1}^{d+1} \log \mathcal{N}_{[]}(\frac{\epsilon}{\sqrt{d+1}}, B_{S_k}(h_{P_{m,d}}, t), \text{Unif}(S_k)) \\ &\leq C(d)m \log(\epsilon^{-1})^d \left(\frac{t}{\epsilon}\right)^{(d-1)/2}, \end{aligned}$$

and the claim follows.

B.2. Reduction to Lemma 14

We will show that when $K^* \in \mathcal{C}(1)$, the LSE only considers convex sets in $\mathcal{C}(C_1 \cdot d)$. First, observe that the score of h_{K^*} is bounded by 5 when n is large enough. To see this,

$$\begin{aligned} n^{-1} \sum_{i=1}^n (Y_i - h_{K^*}(X_i))^2 &\leq 2n^{-1} \left(\sum_{i=1}^n \xi_i^2 + h_{K^*}(X_i)^2 \right) \\ &\leq 2 \left(\int_{\mathbb{S}^{d-1}} h_{K^*}^2(x) d\mathbb{P}(x) + 1 + O\left(\frac{1}{\sqrt{n}}\right) \right) \leq 5. \end{aligned} \tag{20}$$

where we used the fact that $h_{K^*}(x) \leq 1$ when $x \in \mathbb{S}^{d-1}$.

Now, let $K \notin \mathcal{C}(C_1 \cdot d)$. Then, K has a vertex (denoted by v_1) with $\|v_1\|_2 > C_1 d$. In this case for C_1 that is large enough, there exists a spherical cap of \mathbb{S}^{d-1} , denoted by $C(v_1)$, with center

$v_1/\|v_1\|$ and normalized surface area of $1/3$, such that

$$h_K(u) \geq 100 \quad \forall u \in C(v_1).$$

To see this, observe that if $C(v_1)$ has a normalized surface of $1/3$, then it implies that the geodesic distance between $v_1/\|v_1\|$ and its boundary is at most $\pi/2 - c_2/(d-1)$ (for some fixed $c_2 \geq 0$). Thus, for any $u \in C(v_1)$ the following holds:

$$\begin{aligned} h_K(u) &\geq \|v_1\| \cos(\angle(v_1, u)) \geq C_1 d \cos(\pi/2 - c_2/(d-1)) \\ &= C_1 d \sin(c_2/(d-1)) \geq C_1 c_2/2, \end{aligned}$$

therefore if we choose C_1 to be large enough then $h_K(u) \geq 100$ for all $u \in C(v_1)$.

Now, using the fact that X_i are i.i.d uniform on the sphere, we know that with probability of at least $1 - e^{-cn}$, there are at least $n/4$ points in this cap (using concentration for Bernoulli random variables). Moreover, since $\xi_i \sim N(0, 1)$, we know that with high probability, at least 0.49 of the points that lie in this cap, the ξ_i are negative. Thus, we conclude that with high probability

$$n^{-1} \sum_{i=1}^n (Y_i - h_K(X_i))^2 \geq n^{-1} \sum_{X_i \in C(v_1), \xi_i \leq 0} (h_K(X_i) - 1 - \xi_i)^2 \geq 11.$$

Using the last equation and Eq. (20), only sets in $\mathcal{C}(C_1 \cdot d)$ will be considered, and the claim follows.