
Fast Rates for Noisy Interpolation Require Rethinking the Effects of Inductive Bias

Konstantin Donhauser^{1,2} Nicolò Ruggeri^{1,3} Stefan Stojanovic⁴ Fanny Yang¹

Abstract

Good generalization performance on high-dimensional data crucially hinges on a simple structure of the ground truth and a corresponding strong inductive bias of the estimator. Even though this intuition is valid for regularized models, in this paper we caution against a strong inductive bias for interpolation in the presence of noise: While a stronger inductive bias encourages a simpler structure that is more aligned with the ground truth, it also increases the detrimental effect of noise. Specifically, for both linear regression and classification with a sparse ground truth, we prove that minimum ℓ_p -norm and maximum ℓ_p -margin interpolators achieve fast polynomial rates close to order $1/n$ for $p > 1$ compared to a logarithmic rate for $p = 1$. Finally, we provide preliminary experimental evidence that this trade-off may also play a crucial role in understanding non-linear interpolating models used in practice.

1. Introduction

Despite being extremely overparameterized, large complex models such as deep convolutional neural networks generalize surprisingly well, even when interpolating noisy training data. If the noise in the training data is low, a natural explanation could be that these models are biased towards having a certain structural simplicity. For example, in deep learning theory, a long line of work studies the *implicit bias* of standard optimization algorithms towards solutions with a small structured norm, see e.g. (Ji & Telgarsky, 2019; Lyu & Li, 2020; Soudry et al., 2018; Chizat & Bach, 2020; Arora et al., 2019a; Jacot et al., 2018). If the optimal prediction model also has a small corresponding norm, then, intuitively, the

implicit bias effectively reduces the search space to a "good" subset that includes a good model.

Even though this intuition is valid in the low-noise regime, it is unclear why the generalization error might stay low when the (structured) models are forced to interpolate non-negligible noise in the data. Further, it is hard to mathematically characterize structural simplicity for complex prediction models, and theoretical analysis is difficult.

Interestingly, some of the fundamental phenomena revolving around the generalization behavior of interpolating complex overparameterized models even occur for high-dimensional linear models (see e.g., (Hastie et al., 2019; Muthukumar et al., 2020; Bartlett et al., 2020; Deng et al., 2021) and references therein). Although the latter are significantly simpler to analyze, the literature has yet to provide a comprehensive theoretical understanding for interpolating models as it exists for example for regularized estimators. This paper aims to take an important step in that direction.

As mentioned above, high-dimensional parametric models with structural simplicity may correspond to parameter vectors with a small particular norm. For linear models, the common structural simplicity assumption is sparsity, induced by the ℓ_0/ℓ_1 -norms – the concrete example we focus on in this paper. We say that estimators with small ℓ_0/ℓ_1 -norm have a *strong inductive bias* towards "simple", in this case sparse, solutions. In contrast, the more frequently studied rotationally invariant ℓ_2 -norm uniformly shrinks the estimator in all directions, thereby inducing only weak to no inductive bias towards sparse solutions.

For noiseless interpolation, it is well-known that the $\min\text{-}\ell_1$ -norm interpolator (aka basis pursuit) yields exact recovery for sparse ground truths (Chen et al., 1998). Moreover, when the measurements are noisy, its regularized variant, the LASSO (Tibshirani, 1996), achieves minimax optimal rates of order $\frac{s \log(d)}{n}$ (Van de Geer, 2008) (where s is the ℓ_0 -norm of the ground truth). However, when forcing structured models to interpolate the noisy samples, it is unclear why the generalization error might stay low. In fact, the $\min\text{-}\ell_1$ -norm interpolator achieves rates of order $\frac{1}{\log(d/n)}$ (Wang et al., 2022; Muthukumar et al., 2020), suggesting that models with strong inductive biases suffer heavily from noise.

¹ETH Zurich, Department of Computer Science ²ETH AI Center ³Max-Planck-Institute for Intelligent Systems, Tübingen, Germany ⁴ETH Zurich. Correspondence to: Konstantin Donhauser <konstantin.donhauser@ai.ethz.ch>.

On the other hand, a long line of work establishes how min- ℓ_2 -norm interpolators with a weak inductive bias benefit from noise resilience in high dimensions (see e.g., (Hastie et al., 2019; Bartlett et al., 2020; Muthukumar et al., 2020; 2021) and references therein). However, these uniform shrinkage estimators do not encode structural assumptions on the ground truth — thus fail to learn the signal for inherently high-dimensional covariates such as isotropic Gaussians (see e.g., (Muthukumar et al., 2020)).

The apparent trade-off between structural simplicity and noise resilience in high dimensions raises a natural question:

Can min-norm interpolators in noisy high dimensional settings achieve fast or even close to minimax optimal rates using a moderate inductive bias?

To the best of our knowledge, in this paper, we first provide a positive answer for sparse linear models. Specifically, we bound the (directional) estimation error of min- ℓ_p -norm interpolators (for regression) and max- ℓ_p -margin interpolators (for classification). For isotropic Gaussian features with dimension $d \asymp n^\beta$ for $\beta > 1$ and in the presence of observation noise in the data,

- we provide upper and lower bounds for min- ℓ_p -norm interpolators with $p \in (1, 2)$. For large enough d, n , the estimation error decays at polynomial rates close to order $1/n$ compared to logarithmic or constant rates for $p = 1$ or 2 (Section 2).
- we further provide upper bounds for the max- ℓ_p -margin interpolators (or equivalently, hard ℓ_p -margin SVM) with $p \in (1, 2)$. Surprisingly, for large enough d, n , they even match minimax optimal rates of order $1/n$ up to logarithmic factors in the regime $\beta \geq 2$ (Section 3).

We confirm the better generalization properties for the choice $p \in (1, 2)$ compared to $p = 1$ and $p = 2$ on synthetic and real-world data in Section 4. Our results on linear models suggest that when interpolating noisy data, a moderate inductive bias yields the optimal performance. Additional experiments with convolutional neural tangent kernels in Section 7 provide preliminary evidence that this intuition may also extend to non-linear models, prompting an exciting line of future work. In particular, we hypothesize that this trade-off between structural simplicity and noise resilience may be an important ingredient for understanding the good generalization capabilities of overparameterized interpolating models used in practice.

2. Minimum-norm Interpolators for Regression

In this section we derive non-asymptotic bounds for sparse linear regression for the estimation error of min- ℓ_p -norm interpolators with $p \in (1, 2)$. We describe the setting in Section 2.1 and present the main theorem followed by a discussion in Section 2.2.

2.1. Setting

We study a standard linear regression model where we observe n pairs of standard normal distributed features $x_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$ and observations $y_i = \langle x_i, w^* \rangle + \xi_i$ with Gaussian noise $\xi_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$. For simplicity, we consider the 1-sparse ground truth $w^* = (1, 0, \dots, 0)$ and discuss in Section 5.2 how to generalize this assumption. Given the data set $\{(x_i, y_i)\}_{i=1}^n$, the goal is to find an estimator \hat{w} that has small estimation error

$$\mathcal{R}_{\mathcal{R}}(\hat{w}) := \|\hat{w} - w^*\|_2^2 = \mathbb{E}_{x \sim \mathcal{N}(0, I_d)} \langle x, \hat{w} - w^* \rangle^2,$$

which also corresponds to the irreducible prediction error. We specifically study *min- ℓ_p -norm interpolators* with $p \in (1, 2)$, given by

$$\hat{w} = \arg \min_w \|w\|_p \text{ s.t. } \forall i : \langle x_i, w \rangle = y_i, \quad (1)$$

where $\|w\|_p := (\sum_i w_i^p)^{1/p}$.

2.2. Main Result

We now state our main result for regression that provides a non-asymptotic upper bound for the estimation error of min- ℓ_p -norm interpolators.¹

Theorem 2.1. *Let the data distribution be as described in Section 2.1 and assume that $\sigma \asymp 1$. Further, let q be such that $\frac{1}{p} + \frac{1}{q} = 1$. Then there exist universal constants $\kappa_1, \dots, \kappa_7 > 0$ such that for any $n \geq \kappa_1$ and any $p \in \left(1 + \frac{\kappa_2}{\log \log d}, 2\right)$ and $n \log^{\kappa_3} n \lesssim d \lesssim n^{q/2} \log^{-\kappa_4 q} n$, the estimation error of the min- ℓ_p -norm interpolator (1) satisfies*

$$\begin{aligned} \mathcal{R}_{\mathcal{R}}(\hat{w}) &\gtrsim \frac{\sigma^{4-2p} q^p d^{2p-2}}{n^p} \sqrt{\frac{\sigma^2 n}{d}} \text{ and} \\ \mathcal{R}_{\mathcal{R}}(\hat{w}) &\lesssim \frac{\sigma^{4-2p} q^p d^{2p-2}}{n^p} \sqrt{\frac{\sigma^2 n \exp(\kappa_5 q)}{qd}}, \end{aligned} \quad (2)$$

with probability at least $1 - \kappa_6 d^{-\kappa_7}$ over the draws of the data set.

¹We use \lesssim, \gtrsim and \asymp to hide universal constants, without any hidden dependence on d, n or p . Further, $\tilde{O}(\cdot), \tilde{\Theta}(\cdot)$ hide logarithmic factors in d, n or p and $a \vee b = \max(a, b)$.

The proof of the theorem can be found in Appendix C. We now discuss the implications of the theorem and refer to Section 5.2 for a discussion on the assumptions in the theorem. Throughout the discussion we consider the regime $d \asymp n^\beta$ with $\beta > 1$.

Close to minimax optimal rates. Theorem 2.1 implies that for any fixed $\beta > 1$, the estimation error $R(\hat{w}) = \tilde{O}(\alpha)$ of the min- ℓ_p -norm interpolator with $p \in \left(1 + \frac{\kappa_2}{\log \log d}, \frac{2\beta}{2\beta-1}\right)$ vanishes at a polynomial rate ($\alpha < 0$). To illustrate these rates, Figure 1 plots the exponent of the rate α as a function of β for different values of p , assuming that d, n are sufficiently large. We compare the rates against the minimax optimal rate for sparse regression of order $\frac{\sigma^2 \log d}{n}$ (dotted horizontal line at $\alpha = -1$) (Raskutti et al., 2011). We can clearly see in Figure 1 that for $\beta > 2$ and small values of p , the rates of the error are close to the minimax optimal rate. In fact, when choosing $\beta = 2$ and $p = 1 + \frac{\kappa_4}{\log \log d}$, Theorem 2.1 shows that, for n, d sufficiently large, the rate of the error equals the minimax optimal rate up to logarithmic factors.

Faster rates than for $p = 1$ For comparison, we also indicate the rates of the min- ℓ_2 - and min- ℓ_1 -norm interpolators by the dashed horizontal line at $\alpha = 0$, which are of constant and $\frac{\sigma^2}{(\beta-1)\log n}$ order, respectively, (Wang et al., 2022; Muthukumar et al., 2020). Clearly, we can see that the minimum- ℓ_p -norm interpolator with $p \in (1, 2)$ achieves faster rates than with $p = 1$ and 2. We emphasize that Figure 1 depicts the exponent of the rate at which the error vanishes up to logarithmic factors for fixed values of p as $d, n \rightarrow \infty$. For fixed n, d , our non-asymptotic bounds only hold for $p \in (1 + \frac{\kappa_4}{\log \log d}, 2)$. We refer to Section 5.2 for a discussion of this limitation and motivate future work on tight bounds for the full interval $p \in [1, 2]$ in Section 7.1.

Tightness of the upper and lower bounds in Theorem 2.1

We note that for any fixed p and n, d sufficiently large, the upper and lower bounds in Theorem 2.1 are of the same order. In fact, the second term in the lower bound of order $\frac{n}{d}$ is a universal lower bound for all interpolators (Muthukumar et al., 2020). This bound is tight when the term $\frac{\sigma^2 n \exp(\kappa_5 q)}{qd}$ dominates the upper bound in Theorem 2.1, which is the case when $\beta \leq 2$ and p is small constant.

Comparison with existing bounds for min- ℓ_p -norm interpolators. We now discuss existing results for min- ℓ_p -norm interpolators from the literature. To the best of our knowledge, previous works do not study the rates of the min- ℓ_p -norm interpolator. However, we may obtain an upper bound as a consequence of Theorem 4 in (Koehler et al., 2021) which follows straightforwardly from applying Lemma G.2 in Appendix C.2 and Lemma G.1 in Ap-




Figure 1. Depiction of Theorem 2.1 for regression when $d \asymp n^\beta$. We plot the exponent α of the resulting estimation error rate $\tilde{O}(n^\alpha)$ for different strengths of inductive bias p at different high-dimensional regimes β .

pendix G.

Theorem 2.2 (Corollary of Theorem 4 from (Koehler et al., 2021) (informal)). *Let the data distribution be as described in Section 2.1 with general ground truth $w^* \in \mathbb{R}^d$. Under the same conditions on d, n as in Theorem 2.1, there exist universal constants $\kappa_1, \kappa_2, \kappa_3 > 0$ such that for any $p \in (1, 2]$ and q such that $\frac{1}{p} + \frac{1}{q} = 1$, we have that*

$$R_{\mathcal{R}}(\hat{w}) \lesssim \sigma^2 \left(\frac{\sqrt{\log d}}{\sqrt{q}d^{1/q}} + \sqrt{\frac{\log d}{n}} + \frac{n \exp(\kappa_1 q)}{qd} \right) + \sigma \|w^*\|_p d^{1/q} \sqrt{\frac{q}{n}} + \|w^*\|_p^2 \frac{qd^{2/q}}{n}$$

with probability at least $1 - \kappa_2 d^{-\kappa_3}$ over the draws of the data set.

We remark that it already follows from Theorem 2.2, in combination with the results in (Wang et al., 2022), that $p > 1$ achieves faster rates compared to $p = 1$. However, for any p and $\beta > 1$ with $d = n^\beta$ the rates in the upper bound in Theorem 2.2 are at most of order $n^{-1/4}$, and thus much slower than the rates in Theorem 2.1 that can reach orders even close to n^{-1} . For a better comparison, we illustrate both the rates in Theorem 2.1 and Theorem 2.2 in Figure 5 in Appendix B, which also provides a detailed comparison of the proof techniques used to derive Theorems 2.1 and 2.2.

Comparison with ℓ_p -norm regularized estimators We are only aware of existing bounds in the literature for ℓ_p -norm regularized estimators that characterize the ground truth by the ℓ_p -norm (e.g. see (Lecué & Mendelson, 2017)). These bounds are at least of order $\|w^*\|_p \frac{d^{2-2/p}}{n}$ (see Theorem 5.4 in (Lecué & Mendelson, 2017)), and thus slower

than the bounds for the corresponding interpolating estimator in Theorem 2.1, for some choices of β . Furthermore, we hypothesize when the term $\frac{\sigma^{4-2p} q^p d^{2p-2}}{n^p}$ on the RHS in Equation (2) dominates, the optimally ℓ_p -norm regularized estimator achieves the same rates as the corresponding interpolator. As can be seen in the proof of Theorem 2.1 in Appendix C, this term captures the error that arises from the orthogonal projection of \hat{w} onto the direction of the signal w^* . This error is expected to increase when adding explicitly regularization, and thus shrinking the estimator.

3. Maximum-margin Interpolators for Classification

We now establish upper bounds for max- ℓ_p -margin interpolators, also called hard-margin ℓ_p -SVMs or sparse-SVMs (Blanco et al., 2020; Bennett & Brendensteiner, 2000). We show that these interpolators achieve fast rates for $p \in (1, 2)$ and even match minimax optimal rates up to logarithmic factors.

3.1. Setting

We study a discriminative linear classification setting with n pairs of random input features $x_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$ and associated labels $y_i = \text{sgn}(\langle x_i, w^* \rangle) \xi_i$ where the label noise $\xi_i \in \{+1, -1\}$ follows the conditional distribution

$$\xi_i | x_i \sim \mathbb{P}_\sigma(\cdot; \langle x_i, w^* \rangle)$$

for some parameter σ . We again choose $w^* = (1, 0, \dots, 0)$ for the same reason as in Section 2.1. Notice that the noise only depends on the input features in the direction of the ground truth. More specifically, we make the following assumption on the noise distribution \mathbb{P}_σ :

The function $z \rightarrow \mathbb{P}_\sigma(\xi = 1; z)$ is a piece-wise continuous function such that the minimum $\nu_f := \arg \min \mathbb{E}_{Z \sim \mathcal{N}(0,1)} \mathbb{E}_{\xi \sim \mathbb{P}_\sigma(\cdot; Z)} (1 - \xi \nu | Z|)_+^2$ exists and is positive $\nu_f > 0$.

Assumption A is rather weak and satisfied by most noise models in the literature, such as

- *Logistic regression* with $\mathbb{P}_\sigma(\xi_i = 1; z) = h(z\sigma)$ and $h(z) = \frac{e^{|z|}}{1+e^{|z|}}$ and $\sigma > 0$.
- *Random label flips* with $\mathbb{P}_\sigma(\xi = 1; \langle x_i, w^* \rangle) = 1 - \sigma$ and $\sigma \in (0, \frac{1}{2})$.
- *Random noise before quantization* where $y_i = \text{sgn}(\langle w^*, x_i \rangle + \xi_i)$ with $\xi_i | x_i \sim \mathcal{N}(0, \sigma^2)$ and $\sigma^2 > 0$.

Given the data set $\{(x_i, y_i)\}_{i=1}^n$, the goal is to obtain an estimate \hat{w} that directionally aligns with the normalized

ground truth w^* and thus has a small directional estimation error

$$R_C(\hat{w}) := \left\| \frac{\hat{w}}{\|\hat{w}\|_2} - w^* \right\|_2^2. \quad (3)$$

This classification error is also studied for example in the 1-bit compressed sensing literature (see e.g., (Boufounos & Baraniuk, 2008; Plan & Vershynin, 2012) and references therein). Note that a small value $R_C(\hat{w})$ corresponds to a small expected (noiseless) 0-1 error via the relation

$$\begin{aligned} & \mathbb{E}_{x \sim \mathcal{N}(0, I_d)} \mathbb{1}[\text{sgn}(\langle x, \hat{w} \rangle) \neq \text{sgn}(\langle x, w^* \rangle)] \\ &= \frac{1}{\pi} \arccos \left(1 - \frac{R_C(\hat{w})}{2} \right) \approx \frac{1}{\pi} \sqrt{R_C(\hat{w})}. \end{aligned}$$

Throughout this section, we study the *max- ℓ_p -margin interpolators*, or equivalently, the hard-margin ℓ_p -SVM solutions for $p \in [1, 2]$ defined by

$$\hat{w} = \arg \min_w \|w\|_p \quad \text{s.t.} \quad \forall i : y_i \langle x_i, w \rangle \geq 1. \quad (4)$$

3.2. Maximum- ℓ_p -margin Interpolation

We are the first to present non-asymptotic upper bounds for the directional estimation error of max- ℓ_p -margin interpolators.

Theorem 3.1. *Let the data distribution be as described in Section 3.1 and assume that the noise model \mathbb{P}_σ is independent of n, d and p . Let q be such that $\frac{1}{p} + \frac{1}{q} = 1$. There exist universal constants $\kappa_1, \dots, \kappa_7 > 0$ such that for any $n \geq \kappa_1$, any $p \in \left(1 + \frac{\kappa_2}{\log \log d}, 2\right)$ and any $n \log^{\kappa_3} n \lesssim d \lesssim \frac{n^{q/2}}{\log^{\kappa_4 q} n}$, the directional estimation error of the max- ℓ_p -norm interpolator (4) satisfies*

$$R_C(\hat{w}) \lesssim \frac{q^{\frac{3}{2}p} d^{3p-3} \log^{3/2} d}{n^{\frac{3}{2}p}} \vee \frac{n \exp(\kappa_4 q)}{qd} \vee \frac{\log^{\kappa_5} d}{n}, \quad (5)$$

with probability at least $1 - \kappa_6 d^{-\kappa_7}$ over the draws of the data set.

The proof of the result is presented in Appendix D. The dependence on the noise \mathbb{P}_σ is hidden in the universal constants but can be made explicit when carefully following the steps in the proof. We refer to Section 5.2 for a discussion on the assumptions in the theorem. Throughout the discussion we consider the regime $d \asymp n^\beta$ with $\beta > 1$.

Close to minimax optimal rates As for regression, Theorem 3.1 implies that the estimation error $R(\hat{w}) = \tilde{O}(\alpha)$ of the max- ℓ_p -margin interpolator with $p \in \left(1 + \frac{\kappa_2}{\log \log d}, \frac{3\beta}{3\beta-1.5}\right)$ vanishes at a polynomial rate ($\alpha < 0$). We again illustrate the rates in Figure 2 where




Figure 2. Depiction of Theorem 3.1 for classification when $d \asymp n^\beta$. We plot the exponent α of the resulting upper bound of the (directional) estimation error rate $\tilde{O}(n^\alpha)$ for different strengths of inductive bias p and different high-dimensional regimes β .

we plot the exponent of the rate α as a function of β for different values of p , assuming that d, n is sufficiently large. Furthermore, we plot the minimax optimal lower bounds (dotted line at $\alpha = -1$) for the directional estimation error $R_{\mathcal{C}}$, which are known to be of order $\frac{\log d}{n}$ (Wainwright, 2009; Abramovich & Grinshtein, 2018). When $\beta > 2$ and p is small, the rates of the max- ℓ_p -norm interpolator are of order $\tilde{O}(\frac{1}{n})$ and hence minimax optimal up to logarithmic factors. More specifically, for any $\beta > 2$ and $p \in \left(1 + \frac{\kappa_2}{\log \log d}, \frac{6\beta-2}{6\beta-3}\right)$, the term $\frac{\log^{\kappa_5} d}{n}$ dominates the RHS in Equation (5) and we obtain a minimax optimal rate up to logarithmic factors.

Faster rates than $p = 1$ So far, previous non-asymptotic upper bounds for the error of the max- ℓ_1 -margin interpolator (Chinot et al., 2021a) are non-vanishing, while to the best of our knowledge we are not aware of any lower bounds. However, using the same tools as used for bounding the error of the min- ℓ_1 -norm interpolator in (Wang et al., 2022) and the tools introduced in the proof of Theorem 3.1, we can upper and lower bound² the directional estimation error (3) by $R_{\mathcal{C}}(\hat{w}) \asymp \frac{1}{\log(d/n)}$.

Analogous to regression, Figure 2 shows how the max- ℓ_p -margin interpolators with $p > 1$ achieve faster rates compared to $p = 1$ and 2 (dashed line at $\alpha = 0$ in Figure 2). More specifically, for any fixed $\beta > 1$, we obtain faster rates with any $p \in \left(1 + \frac{\kappa_2}{\log \log d}, \frac{3\beta}{3\beta-1.5}\right)$ than with $p = 1$

²The explicit theorem statement is moved from the camera ready version of this paper as a response to the reviewers concerns on the length of the paper. The statement will instead appear in an extended version of (Wang et al., 2022).

and 2.

Lower bounds Unlike Theorem 2.1, Theorem 3.1 does not provide matching lower bounds. However, when $\beta \geq 2$ and p is small, the rates in Theorem 3.1 are close to minimax lower bounds and therefore cannot be improved. Furthermore, similar to regression, in Proposition 3.2 we provide a uniform lower bound of order $\frac{n}{d}$ for all interpolating classifiers that matches the upper bound in Theorem 3.1 for the regime $\beta \leq 2$ and when p is small. The proof can be found in Appendix E.

Proposition 3.2 (Universal lower bound for all interpolating classifiers). *Let the data distribution be as described in Section 3.1 and assume that the noise model \mathbb{P}_σ is independent of n, d and p . There exist universal constants $\kappa_1, \kappa_2 > 0$ such that for all non-zero ground truths $w^* \in \mathbb{R}^d$, with probability $\geq 1 - \exp(-\kappa_1 n) - \exp(-\kappa_2 d)$ for some $\kappa > 0$, the directional estimation error of any interpolator \hat{w} satisfying $\forall i : y_i \langle x_i, \hat{w} \rangle \geq 0$ satisfies*

$$R_{\mathcal{C}}(\hat{w}) \geq \frac{cn}{d}.$$

Comparison with regression. Comparing the rates in Figures 1 and 2, we can see that the max- ℓ_p -margin interpolators in classification achieve faster rates than the corresponding min- ℓ_p -norm interpolators in regression, in the sense that they are of order of minimax lower bounds for a wider range of β and $p \in (1, 2)$. We remark that the proofs of Theorems 2.1 and 3.1 follow a similar scheme (see detailed discussion in Section 5). Intuitively, the difference in performance mostly originates from the fact that the directional estimation error $R_{\mathcal{C}}$ merely depends on the direction of the interpolator \hat{w} and not its magnitude, as it is the case for the estimation error $R_{\mathcal{R}}$ in regression. We note that the authors in (Muthukumar et al., 2021) observe a similar difference in performance between max-margin and min-norm interpolators for $p = 2$ and Gaussians features with spiked covariance matrices.

4. Numerical Simulations and Interpretation of the Results

We now present empirical evidence for our theory³, which predicts the superiority of min- ℓ_p -norm/max- ℓ_p -margin interpolators with moderate inductive bias $p \in (1, 2)$ compared to $p = 1, 2$. Further, we discuss how varying the strength of the inductive bias induces a trade-off between structural simplicity aligning with the ground truth and noise resilience in high dimensions. Finally, we show experimentally that this trade-off can also be intuitively understood

³We release the code for the replication of the experiments at <https://github.com/nickkruggeri/fast-rates-for-noisy-interpolation>

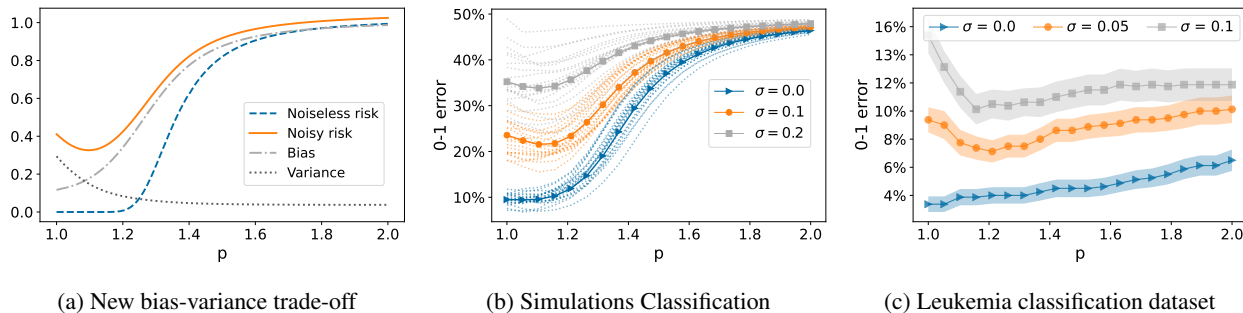


Figure 3. (a) The averaged estimation error (risk) as well as the estimated bias and variance of the min- ℓ_p -norm interpolator (1) as a function of p . We run 50 independent simulations with data drawn from the distribution in Section 2.1 and choose $n = 100$, $d = 5000$, $\sigma = 1$ (noisy) and $\sigma = 0$ (noiseless). (b) Average and individual classification error for the max- ℓ_p -margin interpolator (4) with $n = 100$, $d = 5000$ over 50 independent runs. (c) Mean and variance of the classification error for the max- ℓ_p -margin interpolator (4) on the Leukemia data set with $d = 7070$ and 72 data points in total. See Section 4 for further details.

as a “new” bias-variance trade-off induced by varying the inductive bias.

Simulations for regression and classification While Theorems 2.1 and 3.1 hold for “large enough” d, n , we now illustrate experimentally in Figure 3a and 3b that the superiority of the choice $p \in (1, 2)$ over $p = 1$ and $p = 2$ also holds for finite d, n of practically relevant orders: For $d = 5000$, $n = 100$, the min- ℓ_p -norm/max- ℓ_p -margin interpolators with $p \in (1, 2)$ achieve lowest (directional) estimation error in the noisy case ($\sigma > 0$), while $p = 1$ is optimal in the noiseless case ($\sigma = 0$). Both experiments are run on a synthetic dataset with a 1-sparse ground truth w^* , as described in Section 2.1 (regression) and Section 3.1 (classification). For the regression experiment, we plot the average risks over 50 runs for $\sigma = 1$ (noisy) and $\sigma = 0$ (noiseless) and the bias and variance in the noisy case. For the classification experiment, we randomly flip a fixed amount of $\sigma\%$ of the labels. We plot the error of individual random draws and their mean over 50 independent runs.

Furthermore, we examine the performance of the max- ℓ_p -margin classifier on the Leukemia classification dataset (Golub et al., 1999) with $d = 7070$. In line with our theory, Figure 3 shows that the value of p which minimizes the (directional) estimation error shifts from $p = 1$, in the noiseless case, to $p \in (1, 2)$ in the presence random label flips in the data. Noiseless, i.e. $\sigma = 0$, here means that we are not adding artificial label noise. We plot the averaged error and its variance over 100 random train-test splits with training set size $n = 65$. Finally, note that in practical applications with noisy data, we do not advocate the use of interpolators but recommend using regularized estimators instead.

The strength of the inductive bias induces a new bias-variance trade-off. Why the optimal choice of p does not correspond to the strongest inductive bias ($p = 1$) can be explained by a trade-off between two competing factors: the

regularizing effect of high-dimensionality and the effective sparsity of the solutions, matching the simple structure of the ground truth. In particular, with increasing inductive bias (decreasing $p \rightarrow 1$), noise resilience due to high dimensions decreases while the interpolator becomes effectively sparser and hence achieves better performance on noiseless data. This trade-off is also reflected in Theorems 2.1 and 3.1: The terms of order $\frac{\sigma^2 n \exp(\kappa_4 q)}{q^d}$ capture the noise resilience due to high-dimensionality, or rather vulnerability, and monotonically decrease with d but increase as $p \rightarrow 1$. On the other hand, the terms of order $\frac{q^p \sigma^{4-2p} d^{2p-2}}{n^p}$ and $\frac{q^{\frac{3}{2}p} d^{3p-3} \log^{3/2} d}{n^{\frac{3}{2}p}}$ capture the benefits of the structural alignment with the sparse ground truth as they decrease with $p \rightarrow 1$.

We now give a more intuitive reasoning for this trade-off, which may also translate to more general models. Clearly, to recover noiseless signals in high dimensions, the space of possible solutions must be restricted by using an inductive bias that encourages the structure of the interpolator to match that of the ground truth. For instance, the solution of the min- ℓ_1 -norm interpolator is always n -sparse despite having d parameters. However, exactly this restriction towards a certain structure (such as sparsity) becomes harmful when fitting noisy labels: Instead of low ℓ_2 -norm solutions that can distribute noise across all dimensions, the interpolator is forced to find n -sparse solutions with a higher ℓ_2 -norm.

In Figure 3a, we demonstrate how the trade-off can also be viewed as a novel kind of bias-variance trade-off for interpolating models. In the classical bias-variance trade-off, increasing model complexity (e.g., by decreasing the regularization penalty) leads to larger variance but smaller (statistical) bias. For interpolating models, we observe a similar trade-off when varying the strength of the inductive. Figure 3a depicts this trade-off and shows that the optimal (directional) estimation error is attained at the p where both terms are approximately the same.

5. Proof Idea and Discussion of the Assumptions

We now provide the proof sketch followed by a discussion of the assumptions in Theorem 2.1 and 3.1.

5.1. Proof idea

The proofs of Theorem 2.1 and 3.1 follow a localized uniform convergence argument which is standard in the literature on empirical risk minimization. More specifically, we first upper bound the ℓ_p -norm of the interpolators by

$$\min_{\forall i: \langle x_i, w \rangle = y_i} \|w\|_p^p \leq M_{\mathcal{R}, \xi} \text{ and } \min_{\forall i: y_i \langle x_i, w \rangle \geq 1} \|w\|_p^p \leq M_{\mathcal{C}, \xi}.$$

In a second step, we uniformly bound the risk over all interpolating models with $M_{\mathcal{R}, \xi}/M_{\mathcal{C}, \xi}$ -bounded ℓ_p -norm

$$\begin{aligned} R_{\mathcal{R}}(\hat{w}) &\leq \max_{\substack{\|w\|_p^p \leq M_{\mathcal{R}, \xi} \\ \forall i: \langle x_i, w \rangle = y_i}} \|w - w^*\|_2^2 =: \Phi_{\mathcal{R}} \\ R_{\mathcal{C}}(\hat{w}) &\leq 2 - 2 \underbrace{\min_{\substack{\|w\|_p^p \leq M_{\mathcal{C}, \xi} \text{ and } \\ \forall i: y_i \langle x_i, w \rangle \geq 1}} \frac{\langle w, w^* \rangle}{\|w\|_2}}_{=: \Phi_{\mathcal{C}}}. \end{aligned}$$

We now only discuss the uniform convergence argument as the localization step to upper bound $M_{\mathcal{R}, \xi}$ and $M_{\mathcal{C}, \xi}$ follows from a similar argument.

The proof exploits the assumption that the ground truth is 1-sparse (the extension to s -sparse ground truths is discussed in Appendix 5.2). We decompose $w = (w', w'')$, with w' the first entry of w (since $w^* = (1, 0, \dots, 0)$), and abbreviate $\eta = \|w''\|_2$. Furthermore, we define $\nu = (w - w^*)_{[1]} = w' - 1$ for regression and $\nu = w'$ for classification. Define

$$\begin{aligned} \phi_{\mathcal{R}} &:= \max_{(\nu, \eta) \in \Gamma_{\mathcal{R}}} \nu^2 + \eta^2 \leq \max_{\nu \in \Gamma_{\mathcal{R}}} \nu^2 + \max_{\eta \in \Gamma_{\mathcal{R}}} \eta^2 \\ \phi_{\mathcal{C}} &:= \min_{(\nu, \eta) \in \Gamma_{\mathcal{C}}} \frac{\nu}{\sqrt{\nu^2 + \eta^2}} \geq \left(1 + \frac{\max_{\eta \in \Gamma_{\mathcal{C}}} \eta^2}{\min_{\nu \in \Gamma_{\mathcal{C}}} \nu^2}\right)^{-1/2}, \end{aligned}$$

with constraint sets

$$\Gamma_{\mathcal{R}} = \left\{ (\nu, \eta) \mid \exists b > 0 \text{ s.t. } (1 + \nu)^p + b^p \leq M_{\mathcal{R}, \xi} \text{ and } \frac{1}{n} \|H\|_q^2 b^2 \geq \frac{1}{n} \sum_{i=1}^n (\xi_i - Z_i \nu - \tilde{Z}_i \eta)^2 \right\},$$

and

$$\Gamma_{\mathcal{C}} = \left\{ (\nu, \eta) \mid \exists b > 0 \text{ s.t. } \nu^p + b^p \leq M_{\mathcal{C}, \xi} \text{ and } \frac{1}{n} \|H\|_q^2 b^2 \geq \frac{1}{n} \sum_{i=1}^n \left(1 - \xi_i |Z_i| \nu + \tilde{Z}_i \eta\right)_+^2 \right\},$$

where $(\cdot)_+ = \max(0, \cdot)$ and H, Z, \tilde{Z} are i.i.d. Gaussian random vectors.

The key ingredient for the proofs is now that via the (Convex) Gaussian Minimax Theorem (C)GMT (Thrapoulidis et al., 2015; Gordon, 1988), we can show that $\mathbb{P}(\Phi_{\mathcal{R}} > t | \xi, X') \leq 2\mathbb{P}(\phi_{\mathcal{R}} \geq t | \xi, X')$ and $\mathbb{P}(\Phi_{\mathcal{C}} < t | \xi, X') \leq 2\mathbb{P}(\phi_{\mathcal{C}} \leq t | \xi, X')$ with X' being the subset of features X in the direction of the ground truth. In short, high probability upper and lower bounds for $\phi_{\mathcal{R}}$ and $\phi_{\mathcal{C}}$, yield corresponding high probability bounds for $\Phi_{\mathcal{R}}$ and $\Phi_{\mathcal{C}}$. The (C)GMT has been used previously to obtain similar bounds for regression in (Koehler et al., 2021; Wang et al., 2022).

The result can then be obtained by carefully bounding the constraint sets $\Gamma_{\mathcal{C}}$ and $\Gamma_{\mathcal{R}}$ using tight concentration inequalities for the ℓ_q -norm of i.i.d. Gaussian random vectors in (Paouris et al., 2017).

5.2. Limitations and discussion of the assumptions

In the following subsections, we now discuss the assumptions of our main results and how they can be generalized.

Assumption on p & d While for very large d , Theorems 2.1 and 3.1 hold for most $p \in (1, 2)$, even close to 1, for fixed n, d our theorems apply to the range $p \in \left(1 + \frac{\kappa_2}{\log \log d}, 2\right)$. This assumption is used to obtain a high probability concentration of the dual norm $\|H\|_q$ in Lemma C.4. It is possible and straightforward to relax the assumption on p to allow for smaller p , such as $p \in \left(1 + \frac{\kappa \log \log d}{\log d}, 2\right)$, by choosing $\epsilon(n, d)$ in Lemma C.4 larger. However, this choice comes at the price of non-matching upper and lower bounds. In fact, when choosing $p = 1 + \frac{\kappa}{\log d}$, the ℓ_q -norm $\|H\|_q$ behaves similarly as the ℓ_∞ -norm $\|H\|_\infty$ (see (Paouris et al., 2017)).

Assumption on the sparsity of w^* In this paper we study the special case where the ground truth is 1-sparse and thus aligns "maximally" with the sparse inductive bias of the ℓ_1 -norm. Precisely, choosing a 1-sparse ground truth significantly simplifies the analysis (see Section 5) and the presentation of the bounds in Theorem 2.1 and 3.1. However, the application of the (C)GMT in Propositions C.1 and D.2 holds more generally for s -sparse ground truths and even non-sparse ground truths as exploited in (Koehler et al., 2021). Therefore, the proof methodology presented in this paper can also be employed to bound the risk of general s -sparse ground truths. However, this would come at the cost of more involved theorems statements, and non-tight upper and lower bounds in s .

Assumption on the noise Theorems 2.1 and 3.1 assume that the amount of noise σ is fixed and non-vanishing as $n \rightarrow \infty$. This setting is of particular interest as the ground truth is still consistently learnable (unlike in settings where the noise dominates), while the noise prevents the min- ℓ_1 -norm/max- ℓ_1 -margin interpolators from generalizing well. In fact, we are the first to prove, for a constant noise setting, that we can consistently learn with min-norm/max-margin interpolators at fast rates.

On the other hand, several works including (Chinot et al., 2021b;a; Wojtaszczyk, 2010) have also studied the low noise regime where $\sigma \rightarrow 0$ as n, d tend to infinity. We remark that our proof of Theorem 2.1 can be directly extended to cover vanishing noise for regression. On the other hand, the methodology used for the classification results, i.e. Theorem 3.1, strongly relies on the noise model and we leave the extension of our results to low noise settings (i.e. vanishing fraction of flipped labels) as an interesting future work.

Assumption on the distribution of the features Our proofs strongly rely on the (C)GMT (Propositions C.1 and D.2), which crucially hinges on the assumption that the input features are Gaussian. We believe that generalizing the input distribution \mathbb{P}_X requires the development of novel tools – an important task for future work. For instance, the small-ball method that is known to yield tight bounds for general input distributions for many estimators (Mendelson, 2014; Koltchinskii & Mendelson, 2015) results in loose bounds when applied directly ((Chinot et al., 2021b), see also the discussion on this topic in (Wang et al., 2022)). It remains to be seen whether a modified technique based on the small-ball method can be powerful enough to yield tight bounds for min-norm/max-margin interpolators studied in this paper.

6. Related Work

The majority of works that attempt to rigorously understand minimum-norm/max-margin interpolation from a non-asymptotic viewpoint have so far focused on ℓ_2 -norm interpolators, (Bartlett et al., 2020; Tsigler & Bartlett, 2020; Muthukumar et al., 2020). However, to be asymptotically consistent as $d, n \rightarrow \infty$, the covariates need to be effectively low-dimensional and aligned with the direction of the ground truth, e.g. via a spiked covariance structure.

On the other hand, the existing literature on structured interpolators primarily focuses on the min- ℓ_1 -norm interpolator (Wang et al., 2022; Muthukumar et al., 2020; Chinot et al., 2021b; Wojtaszczyk, 2010; Chatterji & Long, 2022), showing exact rates of order $\frac{1}{\log(d/n)}$. Even though (Koehler et al., 2021; Zhou et al., 2021; Chinot et al., 2021b) present frameworks to obtain non-asymptotic bounds for general

min-norm interpolators, the discussion of the implications of their results center around the ℓ_1/ℓ_2 -norms. For classification, so far the only known non-asymptotic upper bound holds for the max- ℓ_1 -margin interpolator — however, assuming that a constant fraction of data points are mislabeled, the upper bound diverges with growing sample size (Chinot et al., 2021a).

Beyond the mentioned non-asymptotic results, many more papers study linear regression and classification in the limit as $d, n \rightarrow \infty$ (Muthukumar et al., 2021; Hastie et al., 2019; Dobriban & Wager, 2018; Deng et al., 2021; Li & Wei, 2021). In contrast to this paper, these works study the linear regime where $d/n \rightarrow \gamma > 0$ and thus where the errors of interpolators do not vanish.

7. Generality and Future Work

Our results naturally suggests a variety of impactful follow-up work that may advance the understanding of the effect of inductive bias and interpolation both from a theoretical and empirical perspective.

7.1. Future Work on Theory for Linear Interpolators

As a specific question for future work, a natural quantity of interest would be the optimal choice among all $p \in (1, 2)$ as a function of n, d . Unfortunately, our presented bounds are not sufficient to provide an explicit answer. Note that for $\beta < 2$ and $p = 1 + \frac{\kappa_2}{\log \log d}$, our rates match the uniform lower bounds for all interpolators when n, d are sufficiently large. On the other hand, for larger β , the optimum of the upper bound is achieved at $p < 1 + \frac{\kappa_2}{\log \log d}$ which is beyond the range of our analysis. We refer to Section 5.2 for a discussion on smaller choices of p . Giving a precise expression and tight lower bounds supporting the choice remains a challenging task.

7.2. Non-linear Overparameterized Interpolators

The discussion in Section 4 suggests that good generalization of interpolating models in the noisy case hinges on a careful choice of the inductive bias and that stronger is not automatically better, as is perhaps widely assumed. This stands in contrast to regularized models and noiseless interpolation, where the best performance is usually attained at the strongest inductive bias with the right amount of regularization. One analogous conclusion is that the optimal inductive bias for *noiseless* interpolation may not be optimal for *noisy* interpolation, in particular for small sample sizes.

As a step towards more complex non-linear models, we provide experimental evidence to support our claim on convolutional neural tangent kernels (Arora et al., 2019b). Kernel regression with kernels with convolutional filter structure




Figure 4. Classification with convolutional neural tangent kernels on binarized MNIST. We plot the 0 – 1 classification error on the test set as a function of the depth of the model, i.e. the number of stacked convolutional layers with ReLU activation. Means and confidence bands are computed across 50 independent draws of the training set with $n = 500.m$

has shown good performance on real world image datasets (Lee et al., 2020) and is thus a good candidate to move towards state-of-the-art models. In Figure 4, we plot the test error as a function of the depth in a binary classification task on the MNIST data for a small sample size regime (see Appendix A for the experimental details). When we do not add artificial label noise in the training data, we observe how increasing the depth leads to a monotonic increase in error, suggesting that the inductive bias decreases with depth. However, when adding label noise, the optimal performance is attained at a medium depth and hence a moderate inductive bias. These findings suggest a promising avenue for empirical and theoretical investigations evolving around the new bias variance trade-off.

8. Conclusion

In this paper, we showed that $\min\text{-}\ell_p\text{-norm}/\max\text{-}\ell_p\text{-margin}$ interpolators can achieve much faster rates with a moderate inductive bias, i.e., $p \in (1, 2)$, compared to a strong inductive bias, i.e., $p = 1$. This arises from a novel bias-variance type trade-off induced by the inductive bias of the interpolating model, balancing the regularizing effect of high-dimensionality and the structural alignment with the ground truth. Based on preliminary experiments on image data with the CNTK in Figure 4, we further hypothesize that this trade-off carries over to more complex interpolating models and datasets used in practice and leave a thorough investigation as future work.

Acknowledgements

K.D. is supported by the ETH AI Center and the ETH Foundations of Data Science. N.R. is supported by the Max

Planck ETH Center for Learning Systems. We would like to thank Nikita Zhivotovskiy for helpful discussions and Kai Lion for his preliminary experiments. Finally we would like to thank the anonymous reviewers for their feedback.

References

- Abramovich, F. and Grinshtein, V. High-dimensional classification by sparse logistic regression. *IEEE Transactions on Information Theory*, 65(5):3068–3079, 2018.
- Adamczak, R. A tail inequality for suprema of unbounded empirical processes with applications to markov chains. *Electronic Journal of Probability*, 13:1000–1034, 2008.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R., and Wang, R. On exact computation with an infinitely wide neural net. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019a.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R., and Wang, R. On exact computation with an infinitely wide neural net. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019b.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- Bennett, K. P. and Bredensteiner, E. J. Duality and geometry in svm classifiers. *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 57–64, 2000.
- Blanco, V., Puerto, J., and Rodriguez-Chia, A. M. On lp-support vector machines and multidimensional kernels. *Journal of Machine Learning Research (JMLR)*, 21:14–1, 2020.
- Boucheron, S. and Thomas, M. Concentration inequalities for order statistics. *Electronic Communications in Probability*, 17:1–12, 2012.
- Boufounos, P. T. and Baraniuk, R. G. 1-bit compressive sensing. *Proceedings of the Conference on Information Sciences and Systems*, pp. 16–21, 2008.
- Chatterji, N. S. and Long, P. M. Foolish crowds support benign overfitting. *Journal of Machine Learning Research*, 23(125):1–12, 2022.
- Chen, S. S., Donoho, D. L., and Saunders, M. A. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.
- Chinot, G., Kuchelmeister, F., Löffler, M., and van de Geer, S. Adaboost and robust one-bit compressed sensing. *arXiv preprint arXiv:2105.02083*, 2021a.

- Chinot, G., Löffler, M., and van de Geer, S. On the robustness of minimum-norm interpolators. *arXiv preprint arXiv:2012.00807*, 2021b.
- Chizat, L. and Bach, F. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. *Proceedings of the Conference on Learning Theory (COLT)*, pp. 1305–1338, 2020.
- Deng, L. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- Deng, Z., Kammoun, A., and Thrampoulidis, C. A model of double descent for high-dimensional binary linear classification. *Information and Inference: A Journal of the IMA*, 2021.
- Dobriban, E. and Wager, S. High-dimensional asymptotics of prediction: Ridge regression and classification. *Annals of Statistics*, 46(1):247–279, 2018.
- Golub, T. R., Slonim, D. K., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J. P., Coller, H., Loh, M. L., Downing, J. R., Caligiuri, M. A., et al. Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science*, 286(5439):531–537, 1999.
- Gordon, Y. On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n . *Geometric aspects of functional analysis*, pp. 84–106, 1988.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018.
- Ji, Z. and Telgarsky, M. The implicit bias of gradient descent on nonseparable data. *Proceedings of the Conference on Learning Theory (COLT)*, pp. 1772–1798, 2019.
- Ju, P., Lin, X., and Liu, J. Overfitting Can Be Harmless for Basis Pursuit, But Only to a Degree. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 2020.
- Koehler, F., Zhou, L., Sutherland, D. J., and Srebro, N. Uniform convergence of interpolators: Gaussian width, norm bounds and benign overfitting. *Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- Koltchinskii, V. and Mendelson, S. Bounding the smallest singular value of a random matrix without concentration. *International Mathematics Research Notices*, 2015(23): 12991–13008, 2015.
- Lecué, G. and Mendelson, S. Regularization and the small-ball method II: complexity dependent error rates. *Journal of Machine Learning Research (JMLR)*, 18(146):1–48, 2017.
- Ledoux, M. A heat semigroup approach to concentration on the sphere and on a compact riemannian manifold. *Geometric & Functional Analysis GAFA*, 2:221–224, 1992.
- Lee, J., Schoenholz, S., Pennington, J., Adlam, B., Xiao, L., Novak, R., and Sohl-Dickstein, J. Finite versus infinite neural networks: an empirical study. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:15156–15172, 2020.
- Li, Y. and Wei, Y. Minimum ℓ_1 -norm interpolators: Precise asymptotics and multiple descent. *arXiv preprint arXiv:2110.09502*, 2021.
- Lyu, K. and Li, J. Gradient descent maximizes the margin of homogeneous neural networks. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.
- Massart, P. The Tight Constant in the Dvoretzky-Kiefer-Wolfowitz Inequality. *The Annals of Probability*, 18(3): 1269 – 1283, 1990.
- Mendelson, S. Learning without concentration. *Proceedings of the Conference on Learning Theory (COLT)*, pp. 25–39, 2014.
- Muthukumar, V., Vodrahalli, K., Subramanian, V., and Sahai, A. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 1(1):67–83, 2020.
- Muthukumar, V., Narang, A., Subramanian, V., Belkin, M., Hsu, D., and Sahai, A. Classification vs regression in overparameterized regimes: Does the loss function matter? *Journal of Machine Learning Research*, 22(222): 1–69, 2021.
- Novak, R., Xiao, L., Hron, J., Lee, J., Alemi, A. A., Sohl-Dickstein, J., and Schoenholz, S. S. Neural tangents: Fast and easy infinite neural networks in python. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.
- Paouris, G., Valettas, P., and Zinn, J. Random version of dvoretzky’s theorem in ℓ_p^n . *Stochastic Processes and their Applications*, 127(10):3187–3227, 2017.
- Plan, Y. and Vershynin, R. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1):482–494, 2012.

- Raskutti, G., Wainwright, M. J., and Yu, B. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE Transactions on Information Theory*, 57: 6976–6994, 2011.
- Schechtman, G. and Zinn, J. On the volume of the intersection of two l_p^n balls. *Proceedings of the American Mathematical Society*, 110, 1989.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research (JMLR)*, 19(1):2822–2878, 2018.
- Thrapoulidis, C., Oymak, S., and Hassibi, B. Regularized linear regression: A precise analysis of the estimation error. *Proceedings of the Conference on Learning Theory (COLT)*, pp. 1683–1709, 2015.
- Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58(1): 267–288, 1996.
- Tsigler, A. and Bartlett, P. L. Benign overfitting in ridge regression. *arXiv preprint arXiv:2009.14286*, 2020.
- Van de Geer, S. A. High-dimensional generalized linear models and the lasso. *Annals of Statistics*, 36(2):614–645, 2008.
- Wainwright, M. J. Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting. *IEEE Transactions on Information Theory*, 55:5728–5741, 2009.
- Wang, G., Donhauser, K., and Yang, F. Tight bounds for minimum l_1 -norm interpolation of noisy data. *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.
- Wojtaszczyk, P. Stability and instance optimality for Gaussian measurements in compressed sensing. *Foundations of Computational Mathematics*, 10(1):1–13, 2010.
- Zhou, L., Koehler, F., Sutherland, D. J., and Srebro, N. Optimistic rates: A unifying theory for interpolation learning and regularization in linear regression. *arXiv preprint arXiv:2112.04470*, 2021.

A. Experimental details for the convolutional NTK

We include here details on the experiments with the convolutional neural tangent kernel (CNTK) presented in Section 7.

Binarized MNIST We perform a binary classification task on a reduced version of the MNIST data set (Deng, 2012). In particular, we only utilize the digits from 0 to 5 (included), and define the binary labels as $y = -1$ if the digit is even, $y = +1$ if it is odd. To stay closer to the high dimensional regime, we randomly subsample only $n = 500$ data points for training, the remaining ones are used for testing. The label noise is added by randomly flipping exactly $\sigma\%$ of the labels only on the training set.

The CNTK The CNTK is a kernel method based on the analytical solutions of infinite width neural networks trained via gradient flow. In particular, a convolutional neural network with a given architecture has a corresponding CNTK, defined as its limit for an infinite number of channels and trained to convergence via MSE loss. We include here a short description of the model. For additional details we refer to (Arora et al., 2019b; Novak et al., 2020) and references therein.

Consider a given neural network architecture f_θ , parameterized by θ . (Jacot et al., 2018; Arora et al., 2019b) showed that, under initialization and optimization conditions, the CNTK can be formalized as a kernel method with

$$k(x, x') = \mathbb{E}_\theta \left\langle \frac{\partial f_\theta(x)}{\partial \theta}, \frac{\partial f_\theta(x')}{\partial \theta} \right\rangle.$$

For Gaussian weights, the expected value above can be computed in closed form via recursion through the layers, which allows this model to be used in practice and without sampling.

Architecture All the implementations of the CNTK are done via the Neural Tangents Python library (Novak et al., 2020). In our experiments, we utilize CNTK architectures with the following structure:

- we stack convolutional layers followed by a ReLU activation l times, where l is what we call the depth of the architecture. We utilize a kernel size 3, stride 1 and padding at every layer.
- a final flatten layer followed by a linear layer with one output neuron, containing the logit of the classification probability.

The results for this experiment are presented in Figure 4.

B. Comparison of Theorem 2.1 and 2.2

We now discuss how our proof of Theorem 2.1 differs from the proof of Theorem 2.2 (i.e., Theorem 4 (Koehler et al., 2021)). Theorem 4 in (Koehler et al., 2021) (summarized in Theorem 2.2) also applies to non-sparse ground truths w^* and is therefore expected to be less tight. Nevertheless, we discuss the major differences with respect to their proof, which allow us to obtain tighter bounds and illustrate the different rates in Figure 5. Balancing the terms $\sqrt{\frac{d^{2/q}}{n}}$ and $\frac{1}{d^{1/q}}$, we obtain an optimal rate of order $n^{-1/4}$ in Theorem 2.2. We now discuss how each of these terms can be further tightened in the bounds in Theorem 2.1.

Use of triangular inequality in $\Phi_{\mathcal{R},N}$, $\Phi_{\mathcal{R},+}$ and $\Phi_{\mathcal{R},-}$. The definitions of the terms $\Phi_{\mathcal{R},N}$, $\Phi_{\mathcal{R},+}$ and $\Phi_{\mathcal{R},-}$ in the proof of Theorem 2.1 in Appendix C differ from the ones used in (Koehler et al., 2021) for the localization and uniform convergence steps. In (Koehler et al., 2021) the authors rely on the triangular inequality in the localization and uniform convergence steps by bounding

$$\begin{aligned} \Phi_{\mathcal{R},N} &:= \|\hat{w}\|_p^p = \min_{Xw=\xi} \|w + w^*\|_p^p \leq \|w^*\|_p + \min_{Xw=\xi} \|w\|_p \leq \widetilde{M} \quad \text{and} \\ \Phi_{\mathcal{R},+} &:= \max_{\substack{\|w+w^*\|_p \leq \widetilde{M} \\ Xw=\xi}} \|w\|_2^2 \leq \max_{\|w\|_p \leq \widetilde{M} + \|w^*\|_p} \|w\|_2^2. \end{aligned}$$

While such a procedure also covers non-sparse ground truth, the resulting bounds are not tight when w^* is sparse. Instead, we directly bound the terms $\Phi_{\mathcal{R},N}$, $\Phi_{\mathcal{R},+}$ (for which we make use of the sparsity of w^*), which effectively allows us to improve the terms of order $\sqrt{\frac{d^{2/q}}{n}}$ in Theorem 2.2 to the term of order $\frac{\sigma^{4-2p} q^p d^{2p-2}}{n^p}$ in Theorem 2.1.




Figure 5. Depiction of the rates in Theorem 2.1 (solid line) and Theorem 2.2 (dashed lines): classification. Order of the theoretical rates $\tilde{O}(n^\alpha)$ for different strengths of inductive bias p at different asymptotic regimes. On the x-axis we represent the value β defining the regime $d = n^\beta$, on the y-axis the risk decay α .

Bounds for $f_n^{\mathcal{R}}$. Another important improvement is in the bounds used for $f_n^{\mathcal{R}}$. While the proofs in (Koehler et al., 2021) rely on a bound of the form of Equation (13), in order to obtain faster rates, our analysis requires on the tighter bound in Equation (8). More specifically, in Theorem 2.1 this allows us to avoid the term $\sqrt{\frac{\log d}{n}}$.

Concentration inequalities for $\|H\|_q$. The concentration inequalities for the dual norm $\|H\|_q$ are a crucial ingredient for the proof of the theorem. The ones used in the proofs in (Koehler et al., 2021) rely on the concentration of Lipschitz continuous functions (see, (Ledoux, 1992)) and also apply to more general norms. However, for the specific case of ℓ_p -norms we can make use of much tighter concentration inequalities from (Paouris et al., 2017) (see Appendix G) to avoid the terms involving $\sqrt{\frac{\log(d)}{q} \frac{1}{d^{1/q}}}$ in the bound in Theorem 2.2.

C. Proof of Theorem 2.1

Throughout the remainder of the appendix we use $\kappa_1, \kappa_2, \dots$ and c_1, c_2, \dots for generic universal positive constants independent of d, n or p . The value c_1, c_2, \dots may change from display to display throughout the derivations. The standard notations $O(\cdot), o(\cdot), \Omega(\cdot), w(\cdot)$ and $\Theta(\cdot)$, as well as \lesssim, \gtrsim and \asymp , are utilized to hide universal constants, without any hidden dependence on d, n or p . Throughout the proof, whenever we say with high-probability, we mean with probability $\geq 1 - c_1 d^{-c_2}$ with universal constants $c_1, c_2 > 0$.

The lower bound $R_{\mathcal{R}}(\hat{w}) \gtrsim \frac{n}{d}$ follows simply from the uniform lower bound for all interpolators in Theorem 1 and Corollary 1 in (Muthukumar et al., 2020).

For the rest of the proof of the statement, we use a localization/maximization approach, similar to the papers (Chinot et al., 2021b; Koehler et al., 2021; Ju et al., 2020; Muthukumar et al., 2020; Wang et al., 2022), and common in the literature. More specifically, the proof consists of two major parts:

1. *Localization.* We derive a high-probability upper bound $M_{\mathcal{R}, \xi}$ in Proposition C.3 that only depends on the noise ξ but not on the draws of the features X . More specifically, for the norm of the min- ℓ_p -norm interpolator \hat{w}

$$\Phi_{\mathcal{R}, N} := \|\hat{w}\|_p^p = \min_{X(w-w^*)=\xi} \|w\|_p^p = \min_{Xw=\xi} \|w + w^*\|_p^p \leq M_{\mathcal{R}, \xi}$$

and consequently $\|\hat{w}\|_p^p \leq M_{\mathcal{R}, \xi}$, with high probability.

2. *Uniform convergence.* Conditioning on the draw of the noise ξ , we derive high-probability uniform upper and lower

bounds (Proposition C.5) on the prediction error for all interpolators w with $\|w\|_p^p \leq M_{\mathcal{R},\xi}$,

$$\begin{aligned}\Phi_{\mathcal{R},+} &:= \max_{\substack{\|w\|_p^p \leq M_{\mathcal{R},\xi} \\ X(w-w^*)=\xi}} \|w - w^*\|_2^2 = \max_{\substack{\|w+w^*\|_p^p \leq M_{\mathcal{R},\xi} \\ Xw=\xi}} \|w\|_2^2 \quad \text{and} \\ \Phi_{\mathcal{R},-} &:= \min_{\substack{\|w\|_p^p \leq M_{\mathcal{R},\xi} \\ X(w-w^*)=\xi}} \|w - w^*\|_2^2 = \min_{\substack{\|w+w^*\|_p^p \leq M_{\mathcal{R},\xi} \\ Xw=\xi}} \|w\|_2^2.\end{aligned}$$

C.1. Application of the (C)GMT

Similar to (Koehler et al., 2021), we now discuss how the (Convex) Gaussian Minimax Theorem ((C)GMT) (Gordon, 1988; Thrampoulidis et al., 2015) can be used to simplify the study of the primal optimization problems defining $\Phi_{\mathcal{R},N}$, $\Phi_{\mathcal{R},+}$ and $\Phi_{\mathcal{R},-}$ via the introduction of auxiliary optimization problems $\phi_{\mathcal{R},N}$, $\phi_{\mathcal{R},+}$ and $\phi_{\mathcal{R},-}$ defined below. This approach also applies to general sparse ground truths with $\|w^*\|_0 = s$.

We begin with the introduction of some additional notation. Assume that w^* is s -sparse which motivates us to separate w into the vector $w' \in \mathbb{R}^s$ consisting of the s -sparse entries containing the support of w^* and $w'' \in \mathbb{R}^{d-s}$ consisting of the remaining entries. Similarly, decompose w^* into w'_* and $w''_* = 0$. Using this notation, we can express $y_i \langle x_i, w \rangle = y_i (\langle x'_i, w' \rangle + \langle x''_i, w'' \rangle)$ with x'_i and x''_i the corresponding subvectors of x_i . We denote by $X' = (x'_1, \dots, x'_n)$ and $X'' = (x''_1, \dots, x''_n)$ the sub-matrices of $X = (x_1, \dots, x_n)$.

We now make use of the fact that x''_i is independent of y_i and x'_i and define $f_n^{\mathcal{R}} : \mathbb{R}^s \times \mathbb{R} \rightarrow \mathbb{R}_+$ with

$$f_n^{\mathcal{R}}(w', \|w''\|_2) = \frac{1}{n} \|\xi - X'w' - G\|_2^2 \|w''\|_2^2.$$

with $G \sim \mathcal{N}(0, I_n)$, By reformulating the optimization problems via Lagrange multipliers, we can then apply the CGMT to $\Phi_{\mathcal{R},N}$ and the GMT to $\Phi_{\mathcal{R},+}$, $\Phi_{\mathcal{R},-}$. As a result, we obtain the following result:

Proposition C.1. *Let $p \geq 1$ be some constant and assume that w^* is s -sparse. For $H \sim \mathcal{N}(0, I_{d-s})$ define the stochastic auxiliary optimization problems:⁴*

$$\begin{aligned}\phi_{\mathcal{R},N} &= \min_{(w', w'')} \|w' + w'_*\|_p^p + \|w''\|_p^p \quad \text{s.t.} \quad \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^{\mathcal{R}}(w', \|w''\|_2) \\ \phi_{\mathcal{R},+} &= \max_{(w', w'')} \|w'\|_2^2 + \|w''\|_2^2 \quad \text{s.t.} \quad \begin{cases} \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^{\mathcal{R}}(w', \|w''\|_2) \\ \|w' + w'_*\|_p^p + \|w''\|_p^p \leq M_{\mathcal{R},\xi} \end{cases} \\ \phi_{\mathcal{R},-} &= \min_{(w', w'')} \|w'\|_2^2 + \|w''\|_2^2 \quad \text{s.t.} \quad \begin{cases} \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^{\mathcal{R}}(w', \|w''\|_2) \\ \|w' + w'_*\|_p^p + \|w''\|_p^p \leq M_{\mathcal{R},\xi} \end{cases}\end{aligned}$$

where $M_{\mathcal{R},\xi} > 0$ is a constant possibly depending on ξ and X' . Then for any $t \in \mathbb{R}$, we have:

$$\begin{aligned}\mathbb{P}(\Phi_{\mathcal{R},N} > t|\xi, X') &\leq 2\mathbb{P}(\phi_{\mathcal{R},N} \geq t|\xi, X') \\ \mathbb{P}(\Phi_{\mathcal{R},+} > t|\xi, X') &\leq 2\mathbb{P}(\phi_{\mathcal{R},+} \geq t|\xi, X') \\ \mathbb{P}(\Phi_{\mathcal{R},-} < t|\xi, X') &\leq 2\mathbb{P}(\phi_{\mathcal{R},-} \leq t|\xi, X'),\end{aligned}$$

where the probabilities on the LHS and RHS are over the draws of X'' and of G, H , respectively.

The proof is analogous to Lemma 4 and Lemma 7 in (Koehler et al., 2021). Proposition C.1 allows us to reduce the optimization problems depending on the random matrix $X \in \mathbb{R}^{n \times d}$ to optimization problems depending on the much smaller random matrix $X' \in \mathbb{R}^{n \times s}$ and two additional random vectors H, G .

C.2. Localization step - bounding $\phi_{\mathcal{R},N}$

The goal of this first step is to give a high probability upper bound on $\phi_{\mathcal{R},N}$ from Proposition C.1. We now return to the 1-sparse ground truth $w^* = (1, 0, \dots, 0)$ assumed in the theorem statement, which allows us to write $\phi_{\mathcal{R},N}$ as

$$\phi_{\mathcal{R},N} = \min_{(w', w'')} |w' + 1|^p + \|w''\|_p^p \quad \text{s.t.} \quad \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^{\mathcal{R}}(w', \|w''\|_2) \quad (6)$$

⁴We define $\Phi_{\mathcal{R},N}$, $\Phi_{\mathcal{R},-}$, $\phi_{\mathcal{R},N}$, $\phi_{\mathcal{R},-} = \infty$ and $\Phi_{\mathcal{R},+}$, $\phi_{\mathcal{R},+} = -\infty$ if the corresponding optimization problems have no feasible solution.

Remark C.2 (A simplifying modification for the proof.) By definition, $H \in \mathbb{R}^{d-1}$. Instead, for the ease of presentation, we prove the result assuming that $w^*, x_i \in \mathbb{R}^{d+1}$ and hence $H \in \mathbb{R}^d$, which does not affect the expression for the bounds in Equation (2) in Theorem 2.1.

We now derive a bound for $\phi_{\mathcal{R},N}$ that holds with high probability over the draws of H, G, X', ξ and define $\mu_{d,q} = \mathbb{E}\|H\|_q$ and $\sigma_\xi^2 = \frac{1}{n}\|\xi\|_2^2$ with $H \sim N(0, I_d)$ and $\tilde{\mu}_d := \frac{\mu_{d,2q/p}^{2q/p}}{\mu_{d,q}^{2q/p}}$.

Proposition C.3. *Let the data distribution be as described in Section 2.1 and assume that $\sigma = \Theta(1)$. Under the same conditions as in Theorem 2.1 for n, d, p , there exist universal constants $c_1, c_2, \dots, c_8 > 0$ such that with probability at least $1 - c_1 d^{-c_2}$ over the draws of G, H, X', ξ , it holds that*

$$\phi_{\mathcal{R},N} \leq \left(\frac{n\sigma_\xi^2}{\mu_{d,q}^2} \right)^{p/2} \left(1 + \frac{p\nu_0^2}{2\sigma_\xi^2} (1 + c_6\rho) + \mathcal{O}_c + c_7 \left(\rho^2 + \frac{\rho|\nu_0|}{\sigma_\xi} \right) \right) + 1 + p\nu_0 (1 + c_8\nu_0) =: M_{\mathcal{R},\xi}, \quad (7)$$

with $\mathcal{O}_c = c_5 \frac{n \exp(c_3 q)}{dq}$, $\rho = \frac{\log^{c_4} d}{\sqrt{n}}$ and $\nu_0 = -\sigma_\xi^2 \left(\frac{\mu_{d,q}^2}{n\sigma_\xi^2} \right)^{p/2} = \Theta \left(\frac{q^p d^{2p-2}}{n^p} \right)^{1/2}$.

Note that $M_{\mathcal{R},\xi}$ only depends on ξ but not on H and G , and thus, $M_{\mathcal{R},\xi}$ is a valid choice for the upper bound of the ℓ_p -norm of w in the constraints of $\Phi_{\mathcal{R},+}, \Phi_{\mathcal{R},-}$ and Proposition C.1. Furthermore, we note that the bound in the proposition holds with high-probability over the draws of ξ despite explicitly depending on σ_ξ . However, replacing σ_ξ with σ in the Equation (7) by applying the high-probability bound $|\sigma_\xi^2 - \sigma^2| \lesssim \sqrt{\frac{\log(d)}{n}}$ would lead to a loose bound in the subsequent uniform convergence step (Section C.3).

Proof. We condition on the high probability event $|\sigma_\xi - \sigma| \lesssim \sqrt{\frac{\log d}{n}}$ and $\sigma \asymp 1$ by assumption, which we will use in multiple occasions implicitly throughout the remainder of the proof. By definition, any feasible $w = (w', w'')$ in Equation (6) must satisfy $\frac{1}{n} \langle w'', H \rangle^2 \geq f_n^{\mathcal{R}}(w', \|w''\|_2)$. Since the goal is to upper bound the solution of the minimization problem in Equation (6), it suffices to find one feasible point. In a first step, following standard concentration arguments (as used for instance in the proof of Lemma 5 in (Koehler et al., 2021)), we can show that with probability at least $1 - c_1 d^{-c_2}$ the following equation holds uniformly over all w', w''

$$|f_n^{\mathcal{R}}(w', \|w''\|_2) - \sigma_\xi^2 - (w')^2 - \|w''\|_2^2| \lesssim \tilde{\rho} \left[\sigma_\xi (\|w''\|_2 + |w'|) + (w')^2 + \|w''\|_2^2 \right] \quad (8)$$

with $\tilde{\rho} = \sqrt{\frac{\log(d)}{n}}$. We now show that we can choose $w' = \nu_0$ and w'' to be a rescaled version of the subgradient of H , i.e. $w'' = b w''_{\text{grad}} = b \partial \|H\|_q$ with some $b \in \mathbb{R}_+$. We have $\langle w'', H \rangle = \|w\|_p \|H\|_q = b \left\| w''_{\text{grad}} \right\|_p \|H\|_q = b \|H\|_q$ and thus suffices to show that

$$\begin{aligned} \frac{b^2 \|H\|_q^2}{n} &\geq \sigma_\xi^2 + (\nu_0^2 + b^2 \|w''_{\text{grad}}\|_2^2) (1 + c_1 \tilde{\rho}) + c_2 \tilde{\rho} \sigma_\xi (|\nu_0| + b \|w''_{\text{grad}}\|_2) \\ \iff b^2 &\geq \frac{\sigma_\xi^2 + \nu_0^2 (1 + c_1 \tilde{\rho}) + c_2 \tilde{\rho} \sigma_\xi (|\nu_0| + b \|w''_{\text{grad}}\|_2)}{\frac{\|H\|_q^2}{n} - \left\| w''_{\text{grad}} \right\|_2^2 (1 + c_1 \tilde{\rho})}. \end{aligned} \quad (9)$$

Finding feasible b We begin by studying the denominator, for which we can use concentration results for the ℓ_q -norm of Gaussian random vectors in (Paouris et al., 2017), summarized in the following lemma

Lemma C.4. *There exist universal constants $c_1, c_2, \dots, c_6 > 0$ such that for any $n \geq c_1$, $n \lesssim d \lesssim \exp(n^{c_2})$, $q \lesssim \log \log(d)$ and for $\epsilon = \max \left(\frac{n}{d}, \frac{\log^{c_3}(d)}{n} \right)$ with probability at least $1 - c_4 d^{-c_5}$,*

$$\langle w''_{\text{grad}}, H \rangle = \mu_{d,q} (1 + O(\epsilon)) \quad \text{and} \quad (10)$$

$$\left\| w''_{\text{grad}} \right\|_2^2 = \tilde{\mu}_d (1 + O(\epsilon q)) \quad (11)$$

where p is such that $\frac{1}{p} + \frac{1}{q} = 1$. Further it holds that

$$\mu_{d,q} = \Theta(d^{\frac{1}{q} + \frac{\log q}{2 \log d}}) \text{ and } \tilde{\mu}_d = O\left(d^{\frac{1}{q} - \frac{1}{p} + \frac{c_6 q}{\log(d)}}\right) \vee \Omega\left(d^{\frac{1}{q} - \frac{1}{p} - \frac{c_6 q}{\log(d)}}\right).$$

In particular, as a consequence of Lemma C.4 and Proposition G.2, we have with high probability that

$$\begin{aligned} \frac{\|H\|_q^2}{n} - \|w''_{\text{grad}}\|_2^2 (1 + c_1 \tilde{\rho}) &= \frac{\mu_{d,q}^2}{n} (1 + O(\epsilon)) - \tilde{\mu}_d (1 + O(q\epsilon)) (1 + c_1 \tilde{\rho}) \\ &= \frac{\mu_{d,q}^2}{n} \left(1 + O(\epsilon) - \frac{\tilde{\mu}_d n}{\mu_{d,q}^2} (1 + O(q\epsilon + \tilde{\rho})) \right), \end{aligned} \quad (12)$$

with $\epsilon = \max\left(\frac{n}{d}, \frac{\log^2 d}{n}\right) = O\left(\frac{n \exp(c_3 q)}{dq} + \rho^2\right)$ and $\rho = \frac{\log^c d}{\sqrt{n}} \geq \tilde{\rho}$ with universal constant $c > 1$. Note that this step requires the assumption that $q \lesssim \log \log d$ or equivalently $p \geq 1 + \frac{\kappa_2}{\log \log d}$ when choosing ϵ . In fact, the small choice of ϵ is necessary such that the error term $1 + O(\epsilon)$ arising from the concentration in Equation (12) does not affect the final error bound in Theorem 2.1.

Next, it is straight forward to check that both terms $\nu_0, \frac{\tilde{\mu}_d n}{\mu_{d,q}^2} \rightarrow 0$ using Proposition G.2 and the assumptions in the theorem statement on n, d and p . More precisely, we choose $\kappa_1, \dots, \kappa_4 > 0$ in the theorem statement such that

$$\frac{\tilde{\mu}_d n}{\mu_{d,q}^2} = O\left(\frac{n \exp(c_3 q)}{dq}\right) = O\left(\frac{1}{\log d}\right)$$

As a result, Equation (12) becomes

$$\frac{\|H\|_q^2}{n} - \|w''_{\text{grad}}\|_2^2 (1 + c_1 \tilde{\rho}) = \frac{\mu_{d,q}^2}{n} \left(1 + O\left(\frac{n \exp(c_3 q)}{dq} + \rho^2\right) \right),$$

where we further used that $O(q\epsilon) = O\left(\frac{\log \log d}{\log d}\right)$.

We now study the the nominator. Recall that we conditioned on the event where σ_ξ is lower bounded by a constant. We again use Lemma C.4 for the concentration of $\|w''_{\text{grad}}\|_2$ around $\sqrt{\tilde{\mu}_d}$ and $\sqrt{\tilde{\mu}_d} = \frac{\mathcal{O}_c^{1/2} \mu_{d,q}}{n^{1/2}}$, with $\mathcal{O}_c = c_5 \frac{n \exp(c_3 q)}{dq}$, which yields that it suffices to satisfy

$$b^2 \geq \frac{\sigma_\xi^2 \left(1 + \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_1 \rho) + c_2 \frac{\rho}{\sigma_\xi} \left(|\nu_0| + b \frac{\mathcal{O}_c^{1/2} \mu_{d,q}}{n^{1/2}} \right) \right)}{\frac{\mu_{d,q}^2}{n} \left(1 + c_4 \left(\frac{n \exp(c_3 q)}{dq} + \rho^2 \right) \right)}.$$

We can now combine our bounds on the denominator and nominator using Equation (9). First recall that $\tilde{\rho} \leq \rho$. We can now take the Taylor series approximation of $\frac{1}{1+x} = 1 - x + O(x^2)$ to further tighten the inequality in Equation (9) so that it suffices to find b satisfying

$$\begin{aligned} b^2 &\geq \frac{n \sigma_\xi^2}{\mu_{d,q}^2} \left(1 + \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_1 \rho) + c_2 \frac{\rho}{\sigma_\xi} \left(|\nu_0| + b \frac{\mathcal{O}_c^{1/2} \mu_{d,q}}{n^{1/2}} \right) \right) (1 + c_3 \rho^2 + \mathcal{O}_c), \\ \iff b^2 &\geq \frac{n \sigma_\xi^2}{\mu_{d,q}^2} \left(1 + \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_1 \rho) + c_2 \frac{\rho |\nu_0|}{\sigma_\xi} + c_3 \rho^2 + \mathcal{O}_c \right) + \frac{c_4 b n^{1/2} \rho \mathcal{O}_c^{1/2}}{\mu_{d,q}}. \end{aligned}$$

Furthermore, we choose $\kappa_1, \dots, \kappa_4 > 0$ in the theorem statement such that $\nu_0^2 = \frac{\mu_{d,q}^2 \sigma_\xi^{4-2p}}{n^p} = \Theta\left(\frac{q^p d^{2p-2}}{n^p}\right) = O\left(\frac{1}{\log d}\right)$ and therefore, we can choose b satisfying

$$b^2 = \frac{n \sigma_\xi^2}{\mu_{d,q}^2} \left(1 + \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_1 \rho) + c_2 \left(\frac{\rho |\nu_0|}{\sigma_\xi} + \rho^2 \right) + \mathcal{O}_c \right).$$

to satisfy Equation (9) and hence the feasibility constraint for $\phi_{\mathcal{R},N}$.

We can now obtain the upper bound in Proposition C.3 noting that $\phi_{\mathcal{R},N} \leq b^p + (1 + \nu_0)^p$ and taking the Taylor series expansion $(1+x)^p \approx 1 + px + O(x^2)$, using that $\nu_0^2 = O\left(\frac{1}{\log d}\right)$ \square

C.3. Uniform convergence step

We use the same notation as in Appendix C.2. We now prove upper and lower bounds for $\phi_{\mathcal{R},+}$, $\phi_{\mathcal{R},-}$, respectively, for the choice $M_{\mathcal{R},\xi}$ from Proposition C.3. Note that both problems only depend on w'' via $\|w''\|_2$ and $\|w''\|_p$. We are now going to decouple these two quantities. Define

$$\Gamma = \{(\nu, \eta) \in \mathbb{R} \times \mathbb{R}_+ \mid \exists b > 0 \text{ s.t. } \frac{1}{n} \|H\|_q^2 b^2 \geq f_n^{\mathcal{R}}(\nu, \eta) \text{ and } (\nu + 1)^p + b^p \leq M_{\mathcal{R},\xi}\}$$

Using Hölder's inequality (i.e., $\langle w, H \rangle \leq \|H\|_q \|w\|_p$) to relax the optimization problems defining $\phi_{\mathcal{R},+}$, $\phi_{\mathcal{R},-}$ in Proposition C.1, we can show that

$$\phi_{\mathcal{R},+} \leq \left[\max_{(\nu, \eta) \in \Gamma} \nu^2 + \eta^2 \right] \text{ and } \phi_{\mathcal{R},-} \geq \left[\min_{(\nu, \eta) \in \Gamma} \nu^2 + \eta^2 \right].$$

Hence, Equation (2) in Theorem 2.1 follows directly from the following proposition.

Proposition C.5. *Let the data distribution be as described in Section 2.1 and assume that $\sigma = \Theta(1)$. Under the same conditions as in Theorem 2.1 for n, d, p and for the choice $M_{\mathcal{R},\xi}$ in Proposition C.3, there exist universal constants $c_1, c_2, \dots, c_5 > 0$ such that with probability at least $1 - c_1 d^{-c_2}$ over the draws of G, H, X', ξ , it holds that*

$$\Gamma \subseteq \left\{ (\nu, \eta) \mid (\nu - \nu_0)^2 \vee \eta^2 \lesssim \nu_0^{3/2} + \nu_0^2 \rho^{1/2} + \sigma^2 (\mathcal{O}_c + c_4 \rho^2) + \sigma \rho \nu_0 \right\},$$

with $\mathcal{O}_c = c_5 \frac{n \exp(c_3 q)}{dq}$, $\rho = \frac{\log^{c_4} d}{\sqrt{n}}$ and $\nu_0 = -\sigma_\xi^2 \left(\frac{\mu_{d,q}^2}{n \sigma_\xi^2} \right)^{p/2} = \Theta \left(\frac{q^p d^{2p-2}}{n^p} \right)^{1/2}$.

Proof. First we show that any $\nu, \eta \in \Gamma$ vanishes, i.e. $\nu, \eta \rightarrow 0$ as $d, n \rightarrow \infty$, which we then use for the second step where we derive a tight bound. We condition on the high probability event $|\sigma_\xi - \sigma| \lesssim \sqrt{\frac{\log d}{n}}$ and $\sigma \asymp 1$ by assumption, which we will use in multiple occasions implicitly throughout the remainder of the proof.

Step 1: $\Gamma \subset \left\{ (\nu, \eta) \mid \nu^2 + \eta^2 = O \left(\left(\frac{\mu_{d,q}^2}{n \sigma_\xi^2} \right)^{p/2} + \rho \right) + \mathcal{O}_c \right\}$ First note that we can relax the constraints in Γ to $\frac{1}{n} \|H\|_q^2 M_{\mathcal{R},\xi}^{2/p} \geq f_n^{\mathcal{R}}(\nu, \eta)$. Conditioning on Equation (8) to control $f_n^{\mathcal{R}}$ we then obtain

$$\frac{M_{\mathcal{R},\xi}^{2/p} \|H\|_q^2}{n} \geq f_n^{\mathcal{R}}(\nu, \eta) \geq (\sigma^2 + \nu^2 + \eta^2) (1 + c_3 \rho), \quad (13)$$

which holds with probability at least $1 - c_1 d^{-c_2}$, where we again choose $\rho = \frac{\log^{c_4} d}{\sqrt{n}}$ with universal constant $c_4 > 1$.

Recalling the definition of $M_{\mathcal{R},\xi}$ from Proposition C.3, and using the Taylor series approximation $\frac{1}{1+x} = 1 - x + O(x^2)$, it holds that

$$M_{\mathcal{R},\xi}^{2/p} \leq \frac{n \sigma_\xi^2}{\mu_{d,q}^2} \left(1 + \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_6 \rho) + \mathcal{O}_c + c_7 (\rho^2 + \rho \nu_0) + \frac{2 + 2p\nu_0 + c_8 \nu_0^2}{p \left(\frac{n \sigma_\xi^2}{\mu_{d,q}^2} \right)^{p/2}} \right), \quad (14)$$

and hence by Lemma C.4 we have

$$\frac{M_{\mathcal{R},\xi}^{2/p} \|H\|_q^2}{n} \leq \sigma_\xi^2 \left(1 + c_3 \left(\frac{\mu_{d,q}^2}{n \sigma_\xi^2} \right)^{p/2} + \mathcal{O}_c \right),$$

where we only keep track of the dominating terms. In particular, the following upper bound holds from Equation (13),

$$\nu^2 + \eta^2 = O\left(\left(\frac{\mu_{d,q}^2}{n\sigma_\xi^2}\right)^{p/2} + \rho\right) + \mathcal{O}_c = O(\nu_0 + \rho) + \mathcal{O}_c =: \mathcal{O}_b = O\left(\frac{1}{\log d}\right). \quad (15)$$

Step 2: Bound in Proposition C.5 Conditioning on the event where Equation (15) holds, we now show how we can obtain a tighter bound on ν^2 and η^2 using a more refined analysis. As in the proof of Proposition C.3, we condition on the event where Equation (8) holds, which allows us to relax the first constraint in Γ to

$$\frac{1}{n} \|H\|_q^2 b^2 \geq f_n^{\mathcal{R}}(\nu, \eta) \geq \sigma_\xi^2 + (\nu^2 + \eta^2)(1 + c_1\rho) + c_2\sigma_\xi\rho(\nu + \eta). \quad (16)$$

Unlike the previous step, we do not further simplify the second constraint in Γ but instead use that $\nu^2 = \mathcal{O}_b \rightarrow 0$ which allows us to take the Taylor series $(1 + \nu)^p \approx 1 + p\nu + O(\nu^2) = 1 + p\nu(1 + \mathcal{O}_b^{1/2})$, and hence

$$\begin{aligned} b^p &\leq M_{\mathcal{R},\xi} - (\nu + 1)^p \leq M_{\mathcal{R},\xi} - 1 - p\nu(1 + \mathcal{O}_b^{1/2}) \\ &= \left(\frac{n\sigma_\xi^2}{\mu_{d,q}^2}\right)^{p/2} \left(1 + \frac{p}{2} \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_6\rho) + \mathcal{O}_c + c_7(\rho^2 + \frac{\rho\nu_0}{\sigma_\xi})\right) + p(\nu_0 - \nu) \left(1 + \mathcal{O}_b^{1/2}\right). \end{aligned}$$

Further, using the Taylor series approximation $(1 + x)^{2/p} = 1 + \frac{2}{p}x + O(x^2)$ and following the same reasoning as for Equation (14) (by applying Lemma C.4), we get

$$\frac{1}{n} \|H\|_q^2 b^2 \leq \sigma_\xi^2 \left(1 + \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_6\rho) + \mathcal{O}_c + c_7(\rho^2 + \frac{\rho\nu_0}{\sigma_\xi}) + \frac{2(\nu_0 - \nu)}{\left(\frac{n\sigma_\xi^2}{\mu_{d,q}^2}\right)^{p/2}} \left(1 + \mathcal{O}_b^{1/2}\right)\right). \quad (17)$$

Plugging the inequality from Equation (17) into Equation (16) and using that by definition $\nu_0 = -\frac{\sigma_\xi^2 \mu_{d,q}^p}{n^{p/2} \sigma_\xi^p}$, we obtain from completing the square:

$$\begin{aligned} &\sigma_\xi^2 \left(1 + \frac{\nu_0^2}{\sigma_\xi^2} (1 + c_6\rho) + \mathcal{O}_c + c_7(\rho^2 + \frac{\rho\nu_0}{\sigma_\xi}) + \frac{2(\nu_0 - \nu)}{\left(\frac{n\sigma_\xi^2}{\mu_{d,q}^2}\right)^{p/2}} \left(1 + \mathcal{O}_b^{1/2}\right)\right) \\ &\geq \sigma_\xi^2 + (\nu^2 + \eta^2)(1 + c_1\rho) + c_2\sigma_\xi\rho(\nu + \eta) \\ &\implies c_1\nu_0^2(\rho^{1/2} + |\nu_0|^{1/2}) + \sigma_\xi^2\mathcal{O}_c + c_7(\sigma_\xi^2\rho^2 + \sigma_\xi\rho\nu_0) \geq (\nu - \nu_0(1 + \mathcal{O}_b^{1/2}) + c_2\sigma_\xi\rho)^2 + (\eta + c_3\sigma_\xi\rho)^2, \end{aligned}$$

which implies

$$c_1\rho^2 + \mathcal{O}_c + \nu_0^2 \gtrsim \nu^2 + \eta^2 \quad \text{and} \quad \nu^2 \gtrsim \nu_0^2 + c_2\rho^2 + \mathcal{O}_c$$

Hence, we can conclude the proof when plugging in the definition of ν_0 and using that $|\sigma^2 - \sigma_\xi^2| \lesssim \rho$, and additionally note that we can choose κ_2 in Theorem 2.1 such that for any $p \in \left(1 + \frac{\kappa_2}{\log \log(d)}, 2\right)$ and any $d \geq n$, $\frac{\log^{\kappa_2} d}{n} = O(\nu_0^2)$. \square

D. Proof of Theorem 3.1

Since $R_C(\hat{w}) = 2 - 2\frac{\langle \hat{w}, w^* \rangle}{\|\hat{w}\|_2}$, in order to obtain a valid upper bound for $R_C(\hat{w})$, it is sufficient to lower bound $\frac{\langle \hat{w}, w^* \rangle}{\|\hat{w}\|_2}$. The proof of the theorem is again divided into several steps and has essentially the same structure as the proof of Theorem 2.1. Let again $X' = \{x'_i\}_{i=1}^n$ be the features containing the support of w^* (i.e., the first entries of x_i).

Define the set $\mathcal{S}_w(\delta) := \{(w', w'') : |\langle w', w'' \rangle| \geq \delta\}$ with small $\delta > 0$. The proof consists of two major parts:

1. *Localization.* We derive a high-probability upper bound $M_{C,\xi}$ (Proposition D.4), which only depends on ξ and X' , on the norm of the max- ℓ_p -margin interpolator \hat{w} , by finding $M_{C,\xi} > 0$ such that

$$\Phi_{C,N} := \|\hat{w}\|_p^p = \min_{\substack{\forall i: y_i \langle x_i, w \rangle \geq 1 \\ w \in \mathcal{S}_w(\delta)}} \|w\|_p^p \leq M_{C,\xi}.$$

2. *Uniform convergence.* Conditioning on ξ and X' , we derive high-probability uniform upper bound on the classification error for all interpolators w with $\|w\|_p^p \leq M_{C,\xi}$. More precisely, we find a high-probability lower bound (Proposition D.5) for

$$\Phi_{C,-} := \min_{\substack{\|w\|_p^p \leq M_{C,\xi} \\ \forall i: y_i \langle x_i, w \rangle \geq 1 \\ w \in \mathcal{S}_w(\delta)}} \frac{\langle w, w^* \rangle}{\|w\|_2},$$

which in turn then gives us a high probability upper bound for the classification error using that $R_C(\hat{w}) = 2 - 2 \frac{\langle \hat{w}, w^* \rangle}{\|\hat{w}\|_2}$.

Remark D.1. The constraint $(w', w'') \in \mathcal{S}_w(\delta)$ in the definitions of $\Phi_{C,N}, \Phi_{C,-}$ is needed to ensure that the optimization objective $\frac{\langle w', w'_* \rangle}{\sqrt{\|w'\|_2^2 + \|w''\|_2^2}}$ is continuous. Since we can choose δ arbitrarily small, we can neglect the constraint $w \in \mathcal{S}_w(\delta)$ as long as, with high probability, the set of feasible points satisfying the other constraints in $\Phi_{C,N}, \Phi_{C,-}$ does not contain an open ball around $(0, 0)$.

D.1. Application of the (C)GMT

While we only need bounds for $\Phi_{C,N}$ and $\Phi_{C,-}$ for the proof of Theorem 3.1, for completeness of the following Proposition D.2, we also define

$$\max_{\substack{\|w\|_p^p \leq M_{C,\xi} \\ \forall i: y_i \langle x_i, w \rangle \geq 1 \\ w \in \mathcal{S}_w(\delta)}} \frac{\langle w, w^* \rangle}{\|w\|_2} =: \Phi_{C,+}$$

and the function $f_n^C : \mathbb{R}^s \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$,

$$f_n^C(w', \|w''\|_2) = \frac{1}{n} \sum_{i=1}^n (1 - \xi_i \text{sgn}(\langle x'_i, w'_* \rangle) \langle x'_i, w' \rangle - G_i \|w''\|_2)_+^2, \quad (18)$$

with $G \sim \mathcal{N}(0, I_n)$. Using the same notation as for the regression setting (Proposition C.1), and by following a standard argument relying on the Lagrange multiplier, we can bring $\Phi_{C,N}, \Phi_{C,+}$ and $\Phi_{C,-}$ in a suitable form to apply the (C)GMT (see e.g. (Deng et al., 2021) for a similar application), we have

Proposition D.2. *Let $p \geq 1$ be some constant and assume that w^* is s -sparse. Further, let $H \sim \mathcal{N}(0, I_{d-s})$ and let $\delta > 0$ be an arbitrary constant. Define the stochastic auxiliary optimization problems:⁵*

$$\begin{aligned} \phi_{C,N} &= \min_{(w', w'')} \|w'\|_p^p + \|w''\|_p^p \text{ s.t. } \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^C(w', \|w''\|_2) \\ \phi_{C,+} &= \max_{(w', w'') \in \mathcal{S}_w(\delta)} \frac{\langle w', w'_* \rangle}{\sqrt{\|w'\|_2^2 + \|w''\|_2^2}} \text{ s.t. } \begin{cases} \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^C(w', \|w''\|_2) \\ \|w'\|_p^p + \|w''\|_p^p \leq M_{C,\xi} \end{cases} \\ \phi_{C,-} &= \min_{(w', w'') \in \mathcal{S}_w(\delta)} \frac{\langle w', w'_* \rangle}{\sqrt{\|w'\|_2^2 + \|w''\|_2^2}} \text{ s.t. } \begin{cases} \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^C(w', \|w''\|_2) \\ \|w'\|_p^p + \|w''\|_p^p \leq M_{C,\xi} \end{cases} \end{aligned}$$

⁵We define $\Phi_{C,N}, \Phi_{C,-}, \phi_{C,N}, \phi_{C,-} = \infty$ and $\Phi_{C,+}, \phi_{C,+} = -\infty$ if the corresponding optimization problems have no feasible solution.

where $M_{\mathcal{C},\xi} > 0$ is a constant possibly depending on ξ and X' . Then for any $t \in \mathbb{R}$, we have:

$$\begin{aligned}\mathbb{P}(\Phi_{\mathcal{C},N} > t | \xi, X') &\leq 2\mathbb{P}(\phi_{\mathcal{C},N} \geq t | \xi, X') \\ \mathbb{P}(\Phi_{\mathcal{C},+} > t | \xi, X') &\leq 2\mathbb{P}(\phi_{\mathcal{C},+} \geq t | \xi, X') \\ \mathbb{P}(\Phi_{\mathcal{C},-} < t | \xi, X') &\leq 2\mathbb{P}(\phi_{\mathcal{C},-} \leq t | \xi, X'),\end{aligned}$$

where the probabilities on the LHS and RHS are over the draws of X'' and of G, H , respectively.

The proof of the proposition is presented in Appendix D.4.

D.2. Localization step - bounding $\phi_{\mathcal{C},N}$

Similar to proof of Theorem 2.1 in Appendix C.2, we first derive an upper bound for $\phi_{\mathcal{C},N}$, where we again assume for simplicity that $w^*, x_i \in \mathbb{R}^{d+1}$ and thus $H \in \mathbb{R}^d$ (see Remark C.2). In particular, using that $w^* = (1, 0, \dots, 0)$ we can rewrite the optimization problem defining $\phi_{\mathcal{C},N}$ in Proposition D.2 as:

$$\phi_{\mathcal{C},N} = \min_{(w', w'')} |w'|^p + \|w''\|_p^p \text{ s.t. } \frac{1}{n} \langle w'', H \rangle^2 \geq f_n^{\mathcal{C}}(w', \|w''\|_2),$$

$$\text{with } f_n^{\mathcal{C}}(w', \|w''\|_2) = \frac{1}{n} \sum_{i=1}^n (1 - w' \xi_i |x'_i| - \|w''\|_2 G_i)_+^2.$$

Before stating the proof, we first introduce some additional notation. We define $\nu := w'$ and $\eta = \|w''\|_2$, and define

$$\begin{aligned}\zeta_{n,f} &:= \min_{\nu} f_n^{\mathcal{C}}(\nu, 0) \text{ and } \nu_{n,f} := \arg \min_{\nu} f_n^{\mathcal{C}}(\nu, 0) \text{ and} \\ \zeta_{n,\nu\nu} &:= \frac{\partial^2}{\partial \nu^2} f_n^{\mathcal{C}}(\nu_{n,f}, 0) = \frac{2}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] (z'_i)^2 \text{ and } \zeta_{n,\eta\eta} := \frac{2}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|].\end{aligned}$$

Furthermore, we define $f^{\mathcal{C}}(\nu, \eta) := \mathbb{E}_{Z_1, Z_2 \sim \mathcal{N}(0,1)} \mathbb{E}_{\xi \sim \mathbb{P}(\cdot | Z_1)} (1 - \xi \nu |Z_1| - \eta Z_2)_+^2$ and $\nu_f := \arg \min_{\nu} f^{\mathcal{C}}(\nu, 0)$ and define the quantities $\zeta_f := f^{\mathcal{C}}(\nu_f, 0)$, $\zeta_{\eta\eta} := \frac{d^2}{d^2 \eta} |_{(\nu,\eta)=(\nu_f,0)} f^{\mathcal{C}}(\nu, \eta)$ and $\zeta_{\nu\nu} := \frac{d^2}{d^2 \nu} |_{(\nu,\eta)=(\nu_f,0)} f^{\mathcal{C}}(\nu, \eta)$, which are all non-zero positive constants independent of n, d and p as they only depend on \mathbb{P}_{σ} . Throughout the proof, we implicitly make use of the fact that, with high probability over the draws of X' and ξ , the following lemma holds.

Lemma D.3. *There exists universal constants $c_1, c_2 > 0$ such that with probability $\geq 1 - c_1 d^{-c_2}$ over the draws of z' and ξ , we have that*

$$\max(|\nu_f - \nu_{n,f}|, |\zeta_f - \zeta_{n,f}|, |\zeta_{n,\nu\nu} - \zeta_{\nu\nu}|, |\zeta_{n,\eta\eta} - \zeta_{\eta\eta}|) \lesssim \left(\frac{\log d}{n}\right)^{1/4}$$

Proof. Both $f_n^{\mathcal{C}}$ and $f^{\mathcal{C}}$ are convex functions with unique global minima with high probability. In particular, because of Lemma F.1 we know that $\nu_{n,f}$ is attained in a bounded domain around zero with probability $\geq 1 - c_1 d^{-c_2}$. Further, by Lemma F.2 we see that, with high probability, $|\zeta_{n,f} - \zeta_f| \lesssim \frac{\log d}{\sqrt{n}}$ and $|\nu_f - \nu_{n,f}| \lesssim \frac{\log^{1/2} d}{n^{1/4}}$.

To control the quantities $\zeta_{n,\nu\nu}, \zeta_{n,\eta\eta}$, note that we can use exactly the same uniform convergence argument as for Proposition F.4 to show that the functions $\nu \mapsto \frac{2}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu |z'_i|] (z'_i)^2$ and $\nu \mapsto \frac{2}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu |z'_i|]$ also converge uniformly for any bounded domain $\nu \in [\nu_f - \delta, \nu_f + \delta]$ with $\delta > 0$. Thus, the proof then follows from $\nu_{n,f} \rightarrow \nu_f$ and the fact that $h_1(\nu) = \mathbb{E}_{z', \xi} 2\mathbb{1}[1 - \xi \nu |z'|] (z')^2$ and $h_2(\nu) = \mathbb{E}_{z', \xi} 2\mathbb{1}[1 - \xi \nu |z'|]$ are both Lipschitz continuous functions in the domain $[\nu_f - \delta, \nu_f + \delta]$ with $h_1(\nu_f) = \zeta_{\nu\nu}, h_2(\nu_f) = \zeta_{\eta\eta}$. \square

The goal of this section is to show the following upper bound on $\phi_{\mathcal{C},N}$:

Proposition D.4. *Let the data distribution be as described in Section 3.1 and assume that the noise model \mathbb{P}_{σ} is independent of n, d and p . Under the same assumptions as in Theorem 3.1 for n, d, p , there exists universal constants $c_1, c_2, \dots, c_8 > 0$*

such that with probability at least $1 - c_1 d^{-c_2}$ over the draws of G, H, X', ξ , it holds that

$$\begin{aligned} \phi_{\mathcal{C}, N} &\leq \left(\frac{n \zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2} \left(1 + \frac{p}{2} \frac{\zeta_{n,\nu\nu} \Delta \nu_0^2}{2 \zeta_{n,f}} \left(1 + \mathcal{O}_c^{1/2} \log^{3/2} d + c_6 \left(|\Delta \nu_0| \log^{3/2} d + \log d \rho \right) \right) + \mathcal{O}_c + c_7 \rho^2 \right) \\ &\quad + \nu_{n,f}^p + p \nu_{n,f}^{p-1} \Delta \nu_0 (1 + c_8 \Delta \nu_0) =: M_{\mathcal{C}, \xi}, \end{aligned}$$

with $\mathcal{O}_c = c_5 \frac{n \exp(c_3 q)}{dq}$, $\rho = \frac{\log^c d}{\sqrt{n}}$ and $\Delta \nu_0 := -\frac{2 \nu_{n,f}^{p-1} \zeta_{n,f} \mu_{d,q}^p}{\zeta_{n,\nu\nu} (n \zeta_{n,f})^{p/2}}$.

Proof. The proof essentially follows again from exactly the same argument as used in the proof of Proposition C.3, where the goal is again to find one feasible point. We choose $w'' = b w''_{\text{grad}} = b \partial \|H\|_q$ and $w' = \nu_{n,f} + \Delta \nu_0$. Conditioning on the event where the lower bound for $f_n^{\mathcal{C}}$ in Equation (28) from Lemma F.3 holds, we obtain the following lower bound on b (we choose again $w'' = b w''_{\text{grad}} = b \partial \|H\|_q$):

$$\begin{aligned} \frac{1}{n} \langle w'', H \rangle^2 &\geq f_n^{\mathcal{C}}(w', \|w''\|_2) = f_n^{\mathcal{C}}(\nu_{n,f} + \Delta \nu_0, \|w''\|_2) \\ &\Leftarrow b^2 \frac{\|H\|_q^2}{n} \geq \zeta_{n,f} + \left(\frac{\zeta_{n,\nu\nu} \Delta \nu_0^2}{2} + b^2 \frac{\zeta_{n,\eta\eta}}{2} \|w''_{\text{grad}}\|_2^2 \right) (1 + c_1 \rho) + c_2 \rho b \|w''_{\text{grad}}\|_2 \\ &\quad + c_3 (\rho + \sqrt{\log d} (|\Delta \nu_0| + b \|w''_{\text{grad}}\|_2))^3 \\ &\Leftarrow b^2 \geq \frac{\zeta_{n,f} + \frac{\zeta_{n,\nu\nu} \Delta \nu_0^2}{2} (1 + c_1 \rho) + c_2 \rho b \|w''_{\text{grad}}\|_2 + c_3 (\rho + \sqrt{\log d} (|\Delta \nu_0| + b \|w''_{\text{grad}}\|_2))^3}{\frac{\|H\|_q^2}{n} - \frac{\zeta_{n,\eta\eta}}{2} (1 + c_1 \rho) \|w''_{\text{grad}}\|_2^2}, \end{aligned} \tag{19}$$

where we can again choose $\rho = \frac{\log^c d}{\sqrt{n}}$ with $c > 1$ some universal constant. We can now apply the concentration inequalities for $\|H\|_q$ and $\|w''_{\text{grad}}\|_2$ from Lemma C.4 to show that

$$b^2 = \frac{n \zeta_{n,f}}{\mu_{d,q}^2} \left(1 + \frac{\zeta_{n,\nu\nu} \Delta \nu_0^2}{2 \zeta_{n,f}} \left(1 + \mathcal{O}_c^{1/2} \log^{3/2} d + c_6 \left(|\Delta \nu_0| \log^{3/2} d + \log d \rho \right) \right) + \mathcal{O}_c + c_7 \rho^2 \right)$$

satisfies Equation (19). Unlike in the proof of Proposition C.3, we used that

$$\mathcal{O}_c = O(\log^{-5} d) \text{ and } \Delta \nu_0^2 = \Theta \left(\frac{q^p d^{2p-2}}{n^p} \right) = O(\log^{-5} d), \tag{20}$$

which is again satisfied when applying Proposition G.2 and choosing $\kappa_1, \dots, \kappa_4 > 0$ characterizing n, d and p in Theorem 3.1 appropriately. Hence, we can conclude the proof. \square

D.3. Uniform convergence step

We use the same notation as in Appendix D.2 and Appendix C.2. Similarly to Appendix C.3, we can again relax the constraints in the optimization problem defining $\phi_{\mathcal{C}, -}$ in Proposition D.2 by using Hölders inequality $\langle w, H \rangle \leq \|w\|_p \|H\|_q$. Note that both problems again only depend on w'' via $\|w''\|_2$ and $\|w''\|_p$. Define

$$\Gamma = \{(\nu, \eta) | \exists b > 0 \text{ s.t. } \frac{1}{n} \|H\|_q^2 b^2 \geq f_n^{\mathcal{C}}(\nu, \eta) \text{ and } \nu^p + b^p \leq M\}$$

. It is again straight forward to verify that

$$\phi_{\mathcal{C}, -} \geq \left[\min_{(\nu, \eta) \in \Gamma} \frac{\nu}{\sqrt{\nu^2 + \eta^2}} \right] \geq \left(1 + \frac{\max_{(\nu, \eta) \in \Gamma} \eta^2}{\min_{(\nu, \eta) \in \Gamma} \nu^2} \right)^{-1/2},$$

where the last inequality holds when $\min_{\nu \in \Gamma} \nu > 0$. The goal of this section is to prove the following proposition, from which the theorem then straightforwardly follows.

Proposition D.5. *Let the data distribution be as described in Section 3.1 and assume that the noise model \mathbb{P}_σ is independent of n, d and p . Under the same conditions as in Theorem 3.1 for n, d and p and for the choice of $M_{\mathcal{C}, \xi}$ as in Proposition D.4, there exists universal constants $c_1, c_2, c_3, c_4, c_5 > 0$ such that with probability at least $1 - c_1 d^{-c_2}$ over the draws of G, H, X', ξ , it holds that*

$$\Gamma \subseteq \left\{ (\nu, \eta) \in \mathbb{R} \times \mathbb{R}_+ \mid (\nu - \nu_{n,f})^2 \lesssim \mathcal{O}_c + \Delta \nu_0^2 + \rho^2 \text{ and } \eta^2 \lesssim \mathcal{O}_c + \Delta \nu_0^3 \log^{3/2} d + \rho^2 \right\}$$

with $\mathcal{O}_c = c_5 \frac{n \exp(c_3 q)}{dq}$, $\rho = \frac{\log^4 d}{\sqrt{n}}$ and $\Delta \nu_0 := -\frac{2\nu_{n,f}^{p-1} \zeta_{n,f} \mu_{d,q}^p}{\zeta_{n,\nu\nu} (n\zeta_{n,f})^{p/2}}$.

Proof. The proof is very similar to the proof of the uniform convergence bound in Appendix C.3, and consists of iteratively bounding $\eta + \nu$. The primary difference lies in the control of $f_n^{\mathcal{C}}$, which requires more care.

Step 1: $\Gamma \subset \{(\nu, \eta) \mid \nu^2 + \eta^2 = O(1)\}$ A first bound is obtained from the relaxation $b^p \leq M_{\mathcal{C}, \xi}$, which implies that $\frac{1}{n} \|H\|_q^2 M_{\mathcal{C}, \xi}^{2/p} \geq f_n^{\mathcal{C}}(\nu, \eta)$.

Conditioning on the event where Equation (26) in Lemma F.1 holds, we have that $f_n^{\mathcal{C}}(\nu, \eta) \geq c_1 \nu^2 + c_2 \eta^2$, and thus we can relax

$$\frac{M_{\mathcal{C}, \xi}^{2/p} \|H\|_q^2}{n} \geq c_1 \nu^2 + c_2 \eta^2$$

We obtain the desired upper bound from the concentration of $\|H\|_q^2$ from Lemma C.4 and when noting that the dominating term in $M_{\mathcal{C}, \xi}$ is $\left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2}\right)^{p/2}$:

$$\nu^2 + \eta^2 = O(\zeta_{n,f}) = O(1). \quad (21)$$

Step 2: $\Gamma \subset \{(\nu, \eta) \mid (\nu - \nu_{n,f})^2 + \eta^2 = O(\log^{-5} d)\}$ Define $\Delta \nu := \nu - \nu_{n,f}$. Conditioning on Equation (21), we can now repeat the same argument but using the tighter lower bound for $f_n^{\mathcal{C}}$ from Equation (27) in Lemma F.2, where we let ϵ in Lemma F.2 be $\frac{\log^c d}{\sqrt{n}} =: \rho$ with universal constant $c > 1$. Thus, we obtain the high probability upper bound:

$$\begin{aligned} \frac{M_{\mathcal{C}, \xi}^{2/p} \|H\|_q^2}{n} &\geq f_n^{\mathcal{C}}(\nu, \eta) \geq \zeta_f + \tilde{c}_\nu (\nu - \nu_f)^2 + \tilde{c}_\eta \eta^2 + c_1 \rho \\ &= \zeta_f + \tilde{c}_\nu (\Delta \nu^2 + 2\Delta \nu (\nu_{n,f} - \nu_f) + (\nu_{n,f} - \nu_f)^2) + \tilde{c}_\eta \eta^2 + c_1 \rho, \\ &\geq \zeta_{n,f} + \tilde{c}_\nu \Delta \nu^2 + \tilde{c}_\eta \eta^2 + c_2 (\rho^{1/2} + \Delta \nu \rho^{1/2}). \end{aligned} \quad (22)$$

In particular, we can take the Taylor series approximation of $M_{\mathcal{C}, \xi}^{2/p}$ which gives us

$$M_{\mathcal{C}, \xi}^{2/p} = \frac{n\zeta_{n,f}}{\mu_{d,q}^2} \left(1 + \mathcal{O}_c + O \left(\frac{\zeta_{n,\nu\nu} \Delta \nu_0^2}{2\zeta_{n,f}} + \rho^2 + \frac{\nu_{n,f}^p}{\left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2}\right)^{p/2}} \right) \right).$$

And hence, applying again Lemma C.4 to control $\|H\|_q^2$ as in Appendix C.2, the bound in Equation (22) implies

$$\tilde{c}_\nu \Delta \nu^2 + \tilde{c}_\eta \eta^2 = \mathcal{O}_c + O \left(\left(\frac{\mu_{d,q}^2}{n\zeta_{n,f}} \right)^{p/2} + \rho^{1/2} \right) = O(\log^{-5} d), \quad (23)$$

where in the last line we used Equation (20).

Step 3: $\Gamma \subset \{(\nu, \eta) | (\nu - \nu_{n,f})^2 \vee \eta^2 = \mathcal{O}_c + O(\Delta\nu_0^2 + \rho^2)$ We are now able to derive a tighter bound conditioning on the event where Equation (23) holds and thus $\Gamma \in \{(\nu, \eta) | |\nu - \nu_{n,f}|^2 \lesssim \log^{-5} d$ and $\eta^2 \lesssim \log^{-5} d\}$. We can apply Lemma F.3, which allows us, with high probability, to uniformly bound

$$\left| f_n^c(\nu, \eta) - \zeta_{n,f} - \Delta\nu^2 \frac{\zeta_{n,\nu\nu}}{2} - \eta^2 \frac{\zeta_{n,\eta\eta}}{2} \right| \lesssim (\Delta\nu^2 + \eta^2) \log^{-1} d + \rho\eta + \rho^3$$

where we choose again $\rho = \frac{\log^c d}{\sqrt{n}}$. Furthermore, we can relax the second constraint in Γ to

$$\begin{aligned} b^p &\leq M_{C,\xi} - (\Delta\nu + \nu_{n,f})^p = \left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2} \left(1 + \mathcal{O}_c + c_1(\Delta\nu_0^2 + \rho^2) + \frac{p\nu_{n,f}^{p-1}}{\left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2}} c_2 (\Delta\nu_0 - \Delta\nu) \right) \\ &= \left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2} (1 + \mathcal{O}_c + c_1(\Delta\nu_0^2 + \rho^2) + c_3 (\Delta\nu_0^2 - \Delta\nu_0\Delta\nu)) \\ &= \left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2} (1 + \mathcal{O}_c + c_4(\Delta\nu_0^2 + \rho^2 + \Delta\nu_0\Delta\nu)) \end{aligned} \quad (24)$$

with $M_{C,\xi}$ from Proposition D.4 and where we used that by definition $\Delta\nu_0 = -\frac{2\nu_{n,f}^{p-1}\zeta_{n,f}}{\zeta_{n,\nu\nu} \left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2}}$. In summary, we can

again obtain an upper bound for η^2 and $\Delta\nu^2$ from $\frac{1}{n} \|H\|_q^2 b^2 \geq f_n^c(\nu, \eta)$ and when following the same argument as in Appendix C.2. We have

$$\begin{aligned} &\frac{1}{n} \|H\|_q^2 b^2 \geq f_n^c(\nu, \eta) \\ \implies &\frac{\zeta_{n,f}}{\frac{n\zeta_{n,f}}{\mu_{d,q}^2}} (1 + \epsilon) \frac{n\zeta_{n,f}}{\mu_{d,q}^2} (1 + \mathcal{O}_c + c_4(\Delta\nu_0^2 + \rho^2 + \Delta\nu_0\Delta\nu)) \\ &\geq \zeta_{n,f} + \Delta\nu^2 \frac{\zeta_{n,\nu\nu}}{2} + \eta^2 \frac{\zeta_{n,\eta\eta}}{2} + c_5((\Delta\nu^2 + \eta^2) \log^{-1} d + \rho\eta + \rho^3) \\ \implies &\mathcal{O}_c + c_4(\Delta\nu_0^2 + \rho^2 + \Delta\nu_0\Delta\nu) \geq \Delta\nu^2 \frac{\zeta_{n,\nu\nu}}{2} + \eta^2 \frac{\zeta_{n,\eta\eta}}{2} + c_5((\Delta\nu^2 + \eta^2) \log^{-1} d + \rho\eta + \rho^3) \\ \implies &\Delta\nu^2 \vee \eta^2 = \mathcal{O}_c + O(\Delta\nu_0^2 + \rho^2), \end{aligned}$$

where we used in the second line that, with high probability, by Lemma C.4, $\frac{\|H\|_q}{n} \leq \frac{\zeta_{n,f}}{\frac{n\zeta_{n,f}}{\mu_{d,q}^2}} (1 + \epsilon)$ with $\epsilon \lesssim \mathcal{O}_c + \rho^2$.

Step 4: Bound in Proposition D.5 We are now ready to prove the bounds in Proposition D.5. We already know from the previous steps that $\Delta\nu \rightarrow 0$ and hence ν concentrates around $\nu_{n,f}$, which itself concentrates around ν_f by Lemma D.3. However, the classification error in Theorem 3.1 depends on the term $\frac{\eta^2}{\nu^2}$, which allows us to obtain a tighter bound when further bounding η^2 .

Lemma F.3 together with the previous bound on η^2 and $\Delta\nu^2$ implies that uniformly over all $\Delta\nu^2 \vee \eta^2 = \mathcal{O}_c + O(\Delta\nu_0^2 + \rho^2)$, we have:

$$\left| f_n^c(\nu, \eta) - \zeta_{n,f} - \Delta\nu^2 \frac{\zeta_{n,\nu\nu}}{2} - \eta^2 \frac{\zeta_{n,\eta\eta}}{2} \right| \lesssim (\Delta\nu^2 + \eta^2) \mathcal{O}_b + \rho^2 + \rho\eta,$$

where we define $\mathcal{O}_b = c_1 \log^{3/2} d (\sqrt{\mathcal{O}_c} + O(\rho + \Delta\nu_0))$.

Instead of Equation (24), we can analogously obtain the tighter expression (using that $\Delta\nu = \mathcal{O}_b$)

$$\begin{aligned} b^p &\leq M_{C,\xi} - (\Delta\nu + \nu_{n,f})^p \\ &\leq \left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2} \left(1 + \frac{p}{2} \frac{\zeta_{n,\nu\nu} \Delta\nu_0^2}{2\zeta_{n,f}} (1 + \mathcal{O}_b) + \mathcal{O}_c + c_1 \rho^2 + \frac{p\nu_{n,f}^{p-1}}{\left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2} \right)^{p/2}} (\Delta\nu_0 - \Delta\nu) (1 + \mathcal{O}_b) \right). \end{aligned}$$

Therefore, applying the Taylor series approximation and Lemma C.4 as in Appendix C.2, we can upper bound

$$\begin{aligned}
 \frac{1}{n} \|H\|_q^2 b^2 &\leq \frac{1}{n} \|H\|_q^2 \frac{n\zeta_{n,f}}{\mu_{d,q}^2} \left(1 + \frac{\zeta_{n,\nu\nu}\Delta\nu_0^2}{2\zeta_{n,f}} (1 + \mathcal{O}_b) + \mathcal{O}_c + c_1\rho^2 + \frac{2\nu_{n,f}^{p-1}}{\left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2}\right)^{p/2}} (\Delta\nu_0 - \Delta\nu) (1 + \mathcal{O}_b) \right) \\
 &= (1 + \epsilon)\zeta_{n,f} \left(1 + \frac{\zeta_{n,\nu\nu}\Delta\nu_0^2}{2\zeta_{n,f}} (1 + \mathcal{O}_b) + \mathcal{O}_c + c_1\rho^2 + \frac{2\nu_{n,f}^{p-1}}{\left(\frac{n\zeta_{n,f}}{\mu_{d,q}^2}\right)^{p/2}} (\Delta\nu_0 - \Delta\nu) (1 + \mathcal{O}_b) \right) \\
 &= \zeta_{n,f} + \frac{\zeta_{n,\nu\nu}}{2} (2\Delta\nu_0\Delta\nu - \Delta\nu_0^2)(1 + \mathcal{O}_b) + \mathcal{O}_c + c_1\rho^2,
 \end{aligned}$$

where we used the same argument as in the previous step. As a result, we obtain the upper bound

$$\begin{aligned}
 \frac{1}{n} b^2 \|H\|_q^2 &\geq f_n^c(\nu, \eta) \\
 \implies \zeta_{n,f} + \frac{\zeta_{n,\nu\nu}}{2} (2\Delta\nu_0\Delta\nu - \Delta\nu_0^2)(1 + \mathcal{O}_b) + \mathcal{O}_c + c_1\rho^2 \\
 &\geq \zeta_{n,f} + \Delta\nu^2 \frac{\zeta_{n,\nu\nu}}{2} + \eta^2 \frac{\zeta_{n,\eta\eta}}{2} + (\Delta\nu^2 + \eta^2)\mathcal{O}_b + c_2\rho^2 + c_3\rho\eta \\
 \implies \Delta\nu_0^2\mathcal{O}_b + \mathcal{O}_c + c_1\rho^2 &\geq \frac{\zeta_{n,\eta\eta}}{2}(\eta - c_2\rho)^2 + \frac{\zeta_{n,\nu\nu}}{2}(\Delta\nu - \Delta\nu_0)^2.
 \end{aligned}$$

Finally, we get the desired result from the fact that $\zeta_{n,\nu\nu}$ and $\zeta_{n,\eta\eta}$ concentrate around $\zeta_{\nu\nu}$ and $\zeta_{\eta\eta}$ by Lemma D.3. Hence we can conclude the proof. \square

D.4. Proof of Proposition D.2: Application of the (C)GMT for classification

The proof essentially follows exactly the same steps as in (Koehler et al., 2021) except for a few simple modifications which we describe next. First we introduce a more general form of the (C)GMT:

Lemma D.6. *Let $X'' \in \mathbb{R}^{n \times d-s}$ be a matrix with i.i.d. $\mathcal{N}(0, 1)$ entries and let $G \sim \mathcal{N}(0, I_n)$ and $H \sim \mathcal{N}(0, I_{d-s})$ be independent random vectors. Let $S_w \subset \mathbb{R}^s \times \mathbb{R}^{d-s}$ and $S_v \subset \mathbb{R}^n$ be compact sets, and let $\psi : S_w \times S_v \rightarrow \mathbb{R}$ be a continuous function. Then for the following two optimization problems:*

$$\Phi = \min_{(w', w'') \in S_w} \max_{v \in S_v} \langle v, X'' w'' \rangle + \psi((w', w''), v)$$

$$\phi = \min_{(w', w'') \in S_w} \max_{v \in S_v} \|w''\|_2 \langle v, G \rangle + \|v\|_2 \langle w'', H \rangle + \psi((w', w''), v)$$

and any $t \in \mathbb{R}$ holds that:

$$\mathbb{P}(\Phi < t) \leq 2\mathbb{P}(\phi \leq t)$$

If in addition ψ is convex-concave function we also have for any $t \in \mathbb{R}$:

$$\mathbb{P}(\Phi > t) \leq 2\mathbb{P}(\phi \geq t)$$

In both inequalities the probabilities on the LHS and RHS are over the draws of X'' , and of G, H , respectively.

Proof. The first part of the lemma is equivalent to Theorem 10 in (Koehler et al., 2021). The proof of the second part follows from Theorem 9 in (Koehler et al., 2021) and proof of Theorem 10 in (Koehler et al., 2021). \square

We first rewrite $\Phi_{C,N}$ using the Lagrange multipliers $v \in \mathbb{R}^n$ to be able to apply Lemma D.6:

$$\Phi_{C,N} = \min_w \max_{v \geq 0} \|w\|_p^p + \langle v, 1 - D_y X w \rangle = \min_{(w', w'')} \max_{v \geq 0} \|w'\|_p^p + \|w''\|_p^p + \langle v, 1 - D_y X' w' \rangle - \langle D_y v, X'' w'' \rangle$$

where $D_y = \text{diag}(y_1, y_2, \dots, y_n)$ and X'' is the sub-matrix containing the $d - s$ columns of X from the complement of the support of w^* . We note that $D_y X \in \mathbb{R}^{n \times d}$ has i.i.d. entries distributed according to the standard normal distribution and that the random matrices $-X''$ and X'' have the same distribution. Also, the function $\psi((w', w''), v) := \|w'\|_p^p + \|w''\|_p^p + \langle v, 1 - D_y X' w' \rangle$ is a continuous convex-concave function on the whole domain. We further define

$$\begin{aligned} \phi_{C,N} &= \min_{(w', w'')} \max_{v \geq 0} \|w'\|_p^p + \|w''\|_p^p + \langle v, 1 - D_y X' w' \rangle - \|w''\|_2 \langle D_y v, G \rangle - \|D_y v\|_2 \langle w'', H \rangle \\ &= \min_{(w', w'')} \max_{\lambda \geq 0} \|w'\|_p^p + \|w''\|_p^p - \lambda \left(\langle w'', H \rangle - \left\| \left(1 - D_y X' w' - D_y G \|w''\|_2 \right)_+ \right\|_2 \right) \\ &= \min_{(w', w'')} \|w'\|_p^p + \|w''\|_p^p \text{ s.t. } \langle w'', H \rangle \geq \left\| \left(1 - D_y X' w' - D_y G \|w''\|_2 \right)_+ \right\|_2 \end{aligned}$$

where in the second equality we set $\lambda := \|v\|_2$. Since $\|w'\|_p^p$ and $\|w''\|_2$ do not depend on the signs of the entries of w'' , any minimizer $w = (w', w'')$ of $\phi_{C,N}$ satisfies $\langle w'', H \rangle \geq 0$. Hence squaring the last inequality and scaling with $\frac{1}{n}$, we note that the RHS is given by

$$\frac{1}{n} \left\| \left(1 - D_y X' w' - D_y G \|w''\|_2 \right)_+ \right\|_2^2 = \frac{1}{n} \sum_{i=1}^n (1 - \xi_i \text{sgn}(\langle x'_i, w'_* \rangle) \langle x'_i, w' \rangle - G_i \|w''\|_2)_+^2,$$

which is exactly the function $f_n^C(w', \|w''\|_2)$, as defined in Equation (18). Hence, we obtain the desired expression for $\phi_{C,N}$ in Proposition D.2.

In order to complete the proof of the proposition, we need to discuss compactness of the feasible sets in the optimization problem so that we can apply Lemma D.6. For this purpose, we define the following truncated optimization problems $\Phi_{C,N}^r(t)$ and $\phi_{C,N}^r(t)$ for some $r, t \geq 0$:

$$\begin{aligned} \Phi_{C,N}^r(t) &:= \min_{\|w\|_p^p \leq t} \max_{\substack{\|v\| \leq r \\ v \geq 0}} \|w\|_p^p + \langle v, 1 - D_y X w \rangle \\ \phi_{C,N}^r(t) &:= \min_{\|w'\|_p^p + \|w''\|_p^p \leq t} \max_{0 \leq \lambda \leq nr} \|w'\|_p^p + \|w''\|_p^p - \lambda \left(\frac{1}{n} \langle w'', H \rangle^2 - f_n^C(w', \|w''\|_2) \right). \end{aligned}$$

By definition it then follow that

$$\mathbb{P}(\Phi_{C,N} > t | \xi, X') \leq \inf_{r \geq 0} \mathbb{P}(\Phi_{C,N}^r(t) > t | \xi, X').$$

Furthermore, by making use of the simple (linear) dependency on λ in the optimization objective in the definition of $\phi_{C,N}$, a standard limit argument as in the proof of Lemma 7 in (Koehler et al., 2021) shows that:

$$\mathbb{P}(\phi_{C,N} \geq t | \xi, X') \geq \inf_{r \geq 0} \mathbb{P}(\phi_{C,N}^r(t) \geq t | \xi, X')$$

Finally, the proof follows when noting that we can apply Lemma D.6 directly to $\Phi_{C,N}^r(t)$ and $\phi_{C,N}^r(t)$ for any $r, t \geq 0$, which gives us $\mathbb{P}(\Phi_{C,N}^r > t | \xi, X') \leq 2\mathbb{P}(\phi_{C,N}^r \geq t | \xi, X')$. Combining the previous two statements completes the proof for $\Phi_{C,N}$.

The proof for $\Phi_{C,+}$ and $\Phi_{C,-}$ uses the same steps as discussed above. We only detail the proof for $\Phi_{C,-}$ here, as the the proof for $\Phi_{C,+}$ follows from the exact same reasoning. Let $\mathcal{B}_p(M) = \{w \in \mathcal{R}^d : \|w\|_p^p \leq M\}$ be an l_p -ball of radius M and note that we optimize over $(w', w'') \in S_w$ where $S_w = \mathcal{S}_w(\delta) \cap \mathcal{B}_p(M_C, \xi)$ is a compact set for $\delta > 0$ sufficiently small. Furthermore, define the function ψ by $\psi((w', w''), v) := \frac{\langle w', w'_* \rangle}{\sqrt{\|w'\|_2^2 + \|w''\|_2^2}} + \langle v, 1 - D_y X' w' \rangle$, which is a continuous

function on S_w . Similarly as above, we can overcome the issue with the compactness of the set S_v by using a truncation argument as proposed in Lemma 4 in (Koehler et al., 2021). In particular, we define

$$\begin{aligned}\Phi_{\mathcal{C},-}^r &:= \min_{\substack{w \in S_w(\delta) \\ \|w\|_p^p \leq M_{\mathcal{C},\xi}}} \max_{\substack{\|v\| \leq r \\ v \geq 0}} \frac{\langle w, w^* \rangle}{\|w\|_2} + \langle v, 1 - D_y X w \rangle, \\ \phi_{\mathcal{C},-}^r &:= \min_{\substack{(w', w'') \in S_w(\delta) \\ \|w'\|_p^p + \|w''\|_p^p \leq M_{\mathcal{C},\xi}}} \max_{0 \leq \lambda \leq nr} \frac{\langle w', w'' \rangle}{\sqrt{\|w'\|_2^2 + \|w''\|_2^2}} - \lambda \left(\frac{1}{n} \langle w'', H \rangle^2 - f_n^{\mathcal{C}}(w', \|w''\|_2) \right)\end{aligned}$$

for which we have

$$\begin{aligned}\mathbb{P}(\Phi_{\mathcal{C},-} < t | \xi, X') &\leq \inf_{r \geq 0} \mathbb{P}(\Phi_{\mathcal{C},-}^r < t | \xi, X') \\ \text{and } \mathbb{P}(\phi_{\mathcal{C},-} \leq t | \xi, X') &\geq \inf_{r \geq 0} \mathbb{P}(\phi_{\mathcal{C},-}^r \leq t | \xi, X').\end{aligned}$$

We note that the first statement again follows from the definition of $\Phi_{\mathcal{C},-}$, while the second statement follows from a limit argument as in Lemma 4 in (Koehler et al., 2021). Finally, we conclude the proof by applying the first part of Lemma D.6 to $\Phi_{\mathcal{C},-}^r$ and $\phi_{\mathcal{C},-}^r$.

E. Proof of Proposition 3.2: Uniform lower bound for interpolating classifiers

First, note that we can assume without loss of generality that w^* is normalized, i.e. $\|w^*\|_2^2 = 1$. Furthermore, because the classification error is invariant under rotations of the input features, we can assume without loss of generality that $w^* = (1, 0, \dots, 0)$. This trick only works for rotational invariant distributions, including the Gaussian distribution, where the marginal distribution for every sample rotated by a orthogonal matrix is again i.i.d. Gaussian. More precisely, let $\tilde{X} = OX$ with O an orthogonal matrix and $X \sim \mathcal{N}(0, I_d)$, then $\tilde{X} \sim \mathcal{N}(0, I_d)$ as well.

We now distinguish between the two cases where $\langle \hat{w}, w^* \rangle \leq 0$ and $\langle \hat{w}, w^* \rangle > 0$. Clearly, in the first case we have that $R_{\mathcal{C}}(\hat{w}) \geq 2$, and hence, we only need to bound the second case. Furthermore, because the risk is invariant under rescalings of the vector \hat{w} , we can assume without loss of generality that $\langle \hat{w}, w^* \rangle = 1$.

For any $B > 0$ define

$$\Phi_+^{(B)} = \max_w \frac{\langle w, w^* \rangle}{\|w\|_2} \text{ s.t. } \begin{cases} \min_i y_i \langle x_i, w \rangle \geq 0 \\ \langle w^*, w \rangle = 1 \\ \|w\|_2^2 \leq B, \end{cases}$$

where we remark that we could also choose for this proof any other norm or compact set to bound w instead of the ℓ_2 -norm.

We can now apply again the GMT (with slight trivial modifications) from Proposition D.2 to show that for any $t \in \mathbb{R}$, we have

$$\mathbb{P}(\Phi_+^{(B)} > t) \leq 2\mathbb{P}(\phi_+^{(B)} \geq t),$$

with

$$\phi_+^{(B)} = \max_{(\nu, w'')} \frac{\nu}{\sqrt{\nu^2 + \|w''\|_2^2}} \text{ s.t. } \begin{cases} \frac{1}{n} \langle w'', G \rangle^2 \geq \tilde{f}_n(\nu, \|w''\|_2) \\ \nu^2 + \|w''\|_2^2 \leq B \\ \nu = 1 \end{cases}$$

and $\tilde{f}_n(\nu, \eta) = \frac{1}{n} \sum_{i=1}^n (-\xi_i \text{sgn}(\langle w^*, x'_i \rangle) \langle x'_i, \nu \rangle - G_i y)_+^2$. Note that since we only assume that $\min_i y_i \langle x_i, \hat{w} \rangle \geq 0$ and not ≥ 1 , the constant 1 factor in $f_n^{\mathcal{C}}$ from Proposition D.2 in the term $(\cdot)_+$ disappears in \tilde{f}_n .

We can now again lower bound \tilde{f}_n , where a straight forward modification of Lemma F.1 shows that, with probability at least $1 - \exp(-cn)$, uniformly for all ν, η , we have

$$\tilde{f}_n(\nu, \eta) \geq c_\nu \nu^2 + c_\eta \eta^2,$$

with $c_\nu, c_\eta > 0$ some universal constants. Further, using Cauchy-Schwarz, we can upper bound $\frac{1}{n} \langle w'', H \rangle^2 \leq \frac{1}{n} \|H\|_2^2 \|w''\|_2^2 \leq \frac{2d \|w''\|_2^2}{n}$ where the last inequality holds with probability at least $1 - \exp(-cd)$ using standard concentration arguments. Hence, in summary we have the upper bound

$$\frac{2d \|w''\|_2^2}{n} \geq c_\nu \nu^2 \implies \frac{\|w''\|_2^2}{\nu^2} \geq \Omega\left(\frac{n}{d}\right)$$

Therefore, we have that for all $B > 0$, $\phi_+^{(B)} \leq 1 - \Omega\left(\frac{n}{d}\right)$ with universal constants independent of B and thus the proof is complete.

F. Technical lemmas: Bounds for f_n^C

Throughout this section, for simplicity of the notation, we abbreviate $\{x'_i\}_{i=1}^n$ as z' and G from Proposition D.2 as z'' . Thus, we have $f_n^C(\nu, \eta) := \frac{1}{n} \sum_{i=1}^n (1 - \xi_i |z'_i| \nu - \eta z''_i)_+^2$ and

$$f_n^C(\nu, \eta) := \mathbb{E} f_n^C(\nu, \eta) = \mathbb{E}_{Z_1, Z_2 \sim \mathcal{N}(0,1)} \mathbb{E}_{\xi \sim \mathbb{P}_\sigma(\cdot; Z_1)} (1 - \xi \nu |Z_1| - \eta Z_2)_+^2. \quad (25)$$

F.1. Lower bounding f_n^C by a quadratic form

We show the following lemma.

Lemma F.1. *There exist universal positive constants c_ν, c_η only depending on \mathbb{P}_σ and c such that for any ν, η we have that:*

$$f_n^C(\nu, \eta) \geq c_\nu \nu^2 + c_\eta \eta^2 \quad (26)$$

with probability at least $1 - \exp(-cn)$ over the draws of z', z'', ξ .

Proof. We can assume that $\nu, \eta \geq 0$ as the other cases follow from exactly the same argument. First we show an auxiliary statement which we use later in the proof. Namely, we claim that there exists some positive constant c_1 such that for all $z \in [z_1, z_2]$, $\mathbb{P}_\sigma(\xi = -1; z) > c_1$ for some $z_1, z_2 \in \mathbb{R}$ and $z_1 \neq z_2$. Let's prove this statement by contradiction and assume that there exists no $z \in [z_1, z_2]$ which satisfy the previous equation. Then, for almost any $z \sim N(0, 1)$, we have $\mathbb{P}_\sigma(\xi; z) = +1$ and hence the minimum of the function $f_n^C(\nu, \eta) = \mathbb{E} f_n^C(\nu, \eta)$ is obtained for $\nu = \infty$. However, this is in contradiction with the assumption on \mathbb{P}_σ in Section 3.1. Hence there exists some z for which $\mathbb{P}(\xi = -1; z) > c_1$. By the assumption on \mathbb{P}_σ in Section 3.1 we assume piece-wise continuity of $z \rightarrow \mathbb{P}_\sigma(\xi = -1; z)$ and hence there exists some interval $[z - \delta, z + \delta] =: [z_1, z_2]$ in which the given probability is bounded away from zero.

We can assume without loss of generality that this interval does not contain zero, since in that case we can always define a new interval of the form $[\epsilon, z_2]$ or $[z_1, -\epsilon]$ for $\epsilon > 0$ small enough, which does not contain zero. Let's define $\tilde{z} = \max\{|z_1|, |z_2|\}$.

We can now bound $f_n^C(\nu, \eta)$ as follows:

$$\begin{aligned} f_n^C(\nu, \eta) &= \frac{1}{n} \sum_{i=1}^n (1 - \xi_i |z'_i| \nu - \eta z''_i)_+^2 \geq \frac{1}{n} \sum_{i=1}^n \mathbb{1}[\xi_i = -1, z' \in [z_1, z_2], z'' < -c_2] (1 - \xi_i |z'_i| \nu - \eta z''_i)_+^2 \\ &\geq (1 + \tilde{z} |\nu| + c_2 \eta)^2 \frac{1}{n} \sum_{i=1}^n \mathbb{1}[\xi_i = -1, z'_i \in [z_1, z_2], z''_i < -c_2] \end{aligned}$$

for arbitrary positive constant c_2 . From section D.1 we have that z'' is independent of ξ and z' and hence:

$$\begin{aligned} \mathbb{P}(\xi = \text{sign}(\nu), z' \in [z_1, z_2], z'' < -c_2) &= \mathbb{P}(\xi = \text{sign}(\nu) | z' \in [z_1, z_2]) \mathbb{P}(z' \in [z_1, z_2]) \mathbb{P}(z'' < -c_2) \\ &\geq c_1 \left(\Phi^G(z_1) - \Phi^G(z_2) \right) \Phi^G(c_2) \geq c \end{aligned}$$

for some positive constant c . Now using concentration of i.i.d. Bernoulli random variables we obtain:

$$f_n^C(\nu, \eta) \geq (1 + \tilde{z} |\nu| + c_2 \eta)^2 \frac{c}{2} \geq \frac{\tilde{z}^2 c_3}{2} \nu^2 + \frac{c_2^2 c_3}{2} \eta^2$$

with universal constant $c_3 > 0$ and with probability at least $1 - \exp(-nD(c|\frac{c}{2})) \geq 1 - \exp(-nc)$. \square

F.2. Lower bounding f_n^C by a quadratic form with constant

We use the notation as described in Appendix D.2 and the beginning of Appendix F. We show the following lemma.

Lemma F.2. *Let $B_\nu, B_\eta > 0$ be two positive constants. Then, there exist positive constants $\tilde{c}_\nu, \tilde{c}_\eta > 0$ and $c_1, c_2, c_3 > 0$ only depending on \mathbb{P}_σ , such that for any $\epsilon \geq c_1 \sqrt{\frac{\log(n)}{n}}$ and any $\nu \leq B_\nu, \eta \leq B_\eta$ we have that:*

$$f_n^C(\nu, \eta) \geq \zeta_f + \tilde{c}_\nu (\Delta\nu)^2 + \tilde{c}_\eta \eta^2 - \epsilon \quad (27)$$

with probability at least $1 - c_2 \exp\left(-\frac{c_3 n \epsilon^2}{\log(n)}\right)$ over the draws of z', z'', ξ .

Proof. First note that from the uniform convergence result in Proposition F.4 we have that $f^C(\nu, \eta) \geq f_n^C(\nu, \eta) - \epsilon$, with f from Equation (25), with probability at least $1 - c_2 \exp\left(-\frac{c_3 n \epsilon^2}{\log(n)}\right)$. Thus, it is sufficient to study f_n^C . Clearly, by the convexity of f^C we have that $f_n^C \geq \zeta_f$ with $\zeta_f = f^C(\nu_f, 0)$ where we use the simple fact that $(\nu_f, 0)$ is the global minimizer of f^C , which follows from the assumption on \mathbb{P}_σ in Section 3.1. Furthermore, it is not difficult to check that for any ν, η , $\nabla^2 f^C(\nu, \eta) \succ 0$ and therefore, f^C is strictly convex on every compact set. Hence, the proof follows. \square

F.3. Tighter bound for f_n^C around $(\nu_{n,f}, 0)$

The bound in Lemma F.2 contains a term ϵ which is at least of order $\frac{1}{\sqrt{n}}$. This term arises from the uniform convergence of f_n^C to f^C in Proposition F.4 and cannot be avoided. Instead, we show that f_n^C uniformly converges much faster to a different expression.

Lemma F.3. *With probability at least $1 - c_1 d^{-c_2}$ over the draws of z', z'', ξ , we have that*

$$\sup_{\eta \geq 0, \Delta\nu} \left\| f_n^C(\nu, \eta) - \zeta_{n,f} - \Delta\nu^2 \frac{\zeta_{n,\nu\nu}}{2} - \eta^2 \frac{\zeta_{n,\eta\eta}}{2} \right\| \lesssim (\Delta\nu^2 + \eta^2) \rho + \rho \eta + (\rho + \sqrt{\log(d)}(|\Delta\nu| + \eta))^3, \quad (28)$$

with $\rho = \frac{\sqrt{\log(d)}}{\sqrt{n}}$.

Proof. We write

$$\begin{aligned} f_n^C(\nu, \eta) &= \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu |z'_i| - \eta z''_i] \left(\underbrace{(1 - \xi_i \nu_{n,f} |z'_i|)^2}_{=: T_{1,i}} + \underbrace{\xi_i^2 \Delta\nu^2 |z'_i|^2}_{=: T_{2,i}} + \underbrace{\eta^2 (z''_i)^2}_{=: T_{3,i}} \right) \\ &\quad - \underbrace{2\xi_i \Delta\nu |z'_i| (1 - \xi_i \nu_{n,f} |z'_i|)}_{T_{4,i}} - \underbrace{2\eta z''_i (1 - \xi_i \nu_{n,f} |z'_i|)}_{=: T_{5,i}} + \underbrace{2\xi_i \Delta\nu |z'_i| \eta z''_i}_{T_{6,i}} \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu |z'_i| - \eta z''_i] (T_{1,i} + T_{2,i} + T_{3,i} + T_{4,i} + T_{5,i} + T_{6,i}) \end{aligned}$$

We now separately apply concentration inequalities to the different terms. We first need to control the difference of the term $\mathbb{1}[1 - \xi_i \nu |z'_i| - \eta z''_i]$ to the term $\mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|]$. Denote by $F_n(x) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[\xi_i |z'_i| \leq x]$ and by $F(x) = \mathbb{P}(\xi |z'| \leq x)$. By the Dvoretzky-Kiefer-Wolfowitz inequality (see (Massart, 1990)), we have that, with probability at least $1 - \frac{2}{d}$,

$$\sup_{x \in \mathbb{R}} |F_n(x) - F(x)| \leq \sqrt{\frac{\log(d)}{2n}} =: \rho$$

for any $n \geq 3$. Furthermore, note that with probability $\geq 1 - \frac{1}{d}$, we have that $\max_i \{|z''_i|, |z'_i|\} = \gamma_m \lesssim \sqrt{\log(d)}$ (see e.g.,

(Boucheron & Thomas, 2012)). Therefore, we can upper bound

$$\begin{aligned}
 & \frac{1}{n} \sum_{i=1}^n (\mathbb{1}[1 - \xi_i \nu |z'_i| - \eta z''_i] - \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|]) \\
 & \leq F_n \left(\frac{1}{\nu_{n,f}} (1 + \gamma_m(|\Delta\nu| + \eta)) \right) - F_n \left(\frac{1}{\nu_{n,f}} (1 - \gamma_m(|\Delta\nu| + \eta)) \right) \\
 & \leq 2\rho + F \left(\frac{1}{\nu_{n,f}} (1 + \gamma_m(|\Delta\nu| + \eta)) \right) - F \left(\frac{1}{\nu_{n,f}} (1 - \gamma_m(|\Delta\nu| + \eta)) \right) \lesssim \rho + \gamma_m(|\Delta\nu| + \eta)
 \end{aligned}$$

where we condition on the event where $\nu_{n,f}$ concentrates around the constant ν_f in Lemma D.3. Therefore, we get

$$\begin{aligned}
 & \left| f_n^{\mathcal{C}}(\nu, \eta) - \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] (T_{1,i} + T_{2,i} + T_{3,i} + T_{4,i} + T_{5,i} + T_{6,i}) \right| \\
 & = \left| \frac{1}{n} \sum_{i=1}^n (\mathbb{1}[1 - \xi_i \nu |z'_i| - \eta z''_i] - \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|]) (T_{1,i} + T_{2,i} + T_{3,i} + T_{4,i} + T_{5,i} + T_{6,i}) \right| \\
 & \lesssim (\rho + \gamma_m(|\Delta\nu| + \eta))^3
 \end{aligned}$$

Thus, we only need to study $\frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] (T_{1,i} + T_{2,i} + T_{3,i} + T_{4,i} + T_{5,i} + T_{6,i})$. We now separately bound the different terms involving $T_{j,i}$ for every j which all hold with probability at least $1 - c_1 d^{-c_2}$ with universal constants $c_1, c_2 > 0$.

$T_{1,i}$ First, note that by definition, $\frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] T_{1,i} = \zeta_{n,f}$.

$T_{2,i}$ Recall that $\frac{\zeta_{n,\nu\nu}}{2} = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] (z'_i)^2$.

$T_{3,i}$ Recall that $\frac{\zeta_{n,\eta\eta}}{2} = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|]$.

$$\left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] \eta^2 (z''_i)^2 - \eta^2 \frac{\zeta_{n,\eta\eta}}{2} \right| \leq \eta^2 \left| \frac{1}{n} \sum_{i=1}^k ((z''_i)^2 - 1) \right| \lesssim \eta^2 \rho$$

with probability at least $1 - \frac{1}{d}$ over the draws of z''_i . We used that $\sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] = k \leq n$. We apply the same reasoning in the following concentration inequalities.

$T_{4,i}$ Note that by definition of $\nu_{n,f}$ we have that

$$\frac{d}{d\nu} |_{\nu_{n,f}} f_n(\nu, 0) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] 2(1 - \xi_i \nu_{n,f} |z'_i|) \xi_i |z'_i| = 0$$

$T_{5,i}$ We can bound the term involving $T_{5,i}$ by

$$\left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] 2\eta z''_i (1 - \xi_i \nu_{n,f} |z'_i|) \right| \leq 2\eta \left| \frac{1}{n} \sum_{i=1}^n (1 + \nu_{n,f} |z'_i|) z''_i \right| \lesssim \eta \rho$$

$T_{6,i}$ We can bound the term involving $T_{6,i}$ by

$$\left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}[1 - \xi_i \nu_{n,f} |z'_i|] 2\xi_i \Delta\nu |z'_i| \eta z''_i \right| \leq 2\eta |\Delta\nu| \left| \frac{1}{n} \sum_{i=1}^n z'_i z''_i \right| \lesssim \eta |\Delta\nu| \rho$$

□

F.4. Uniform convergence of f_n^C to f^C

We use the notation as described in the beginning of Appendix F and denote with $Z', Z'' \sim \mathcal{N}(0, 1)$ independent Gaussian random variables and let $\xi_{RV} | Z' \sim \mathbb{P}_\sigma(\cdot; Z')$.

Let P be the distribution of (Z', Z'', ξ_{RV}) and let $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{z'_i, z''_i, \xi_i}$. Define $g_{\nu, \eta}(Z, \xi) = (1 - \xi_{RV} |Z'| \nu - \eta Z'')_+^2$ with $Z = (Z', Z'')$. Note that using this notation, we have $P g_{\nu, \eta} = \mathbb{E} g_{\nu, \eta}(Z, \xi) = f^C(\nu, \eta)$, with f from Equation (25), and $P_n g_{\nu, \eta} = f_n^C(\nu, \eta)$. Let $\mathcal{G}_B = \{g_{\nu, \eta} \mid \nu \leq B_\nu, \eta \leq B_\eta\}$ with positive constants B_ν and B_η . Moreover define:

$$\|P_n - P\|_{\mathcal{G}_B} := \sup_{g_{\nu, \eta} \in \mathcal{G}_B} |(P_n - P)g_{\nu, \eta}|$$

We then obtain from an application of Theorem 4 in (Adamczak, 2008) the following uniform convergence result

Proposition F.4. *There exist positive universal constants $c_1, c_2, c_3 > 0$ such that*

$$\mathbb{P} \left(\|P_n - P\|_{\mathcal{G}_B} \leq \frac{c_1}{\sqrt{n}} + \epsilon \right) \geq 1 - c_2 \exp \left(-c_3 \frac{n\epsilon^2}{\log n} \right)$$

Proof. The proof follows from an application of the tail-bound result from Theorem 4 in (Adamczak, 2008). More precisely,

Theorem F.5 (Corollary of Theorem 4 in (Adamczak, 2008)). *For any $0 < t < 1$, $\delta > 0$, $\alpha \in (0, 1]$ there exists a constant $C = C(\alpha, t, \delta)$ such that*

$$\mathbb{P} \left(\|P_n - P\|_{\mathcal{G}_B} \geq (1+t)\mathbb{E} \|P_n - P\|_{\mathcal{G}_B} + \epsilon \right) \leq \exp \left(-\frac{\epsilon^2}{2(1+\delta)\sigma_{\mathcal{G}_B}^2} \right) + 3 \exp \left(-\left(\frac{\epsilon}{C\psi_{\mathcal{G}_B}} \right)^\alpha \right)$$

with variance term

$$\sigma_{\mathcal{G}_B}^2 := \sup_{g_{\nu, \eta} \in \mathcal{G}_B} \frac{1}{n} \left(\mathbb{E} [g_{\nu, \eta}^2] - (\mathbb{E} [g_{\nu, \eta}])^2 \right),$$

and $\psi_{\mathcal{G}_B}$ defined as:

$$\psi_{\mathcal{G}_B} = \left\| \max_{1 \leq i \leq n} \sup_{g_{\nu, \eta} \in \mathcal{G}_B} \frac{1}{n} |g_{\nu, \eta}(z_i, \xi_i) - \mathbb{E}[g_{\nu, \eta}]| \right\|_{\Psi_\alpha}$$

where $\|\cdot\|_{\Psi_\alpha}$ denotes Orlicz norm.

We choose $\alpha = 1$ and explicitly show that condition from Theorem 4 in (Adamczak, 2008) requiring finite Orlicz norms is indeed satisfied for this choice of α . We separate the proof into two steps, where in a first step we bound the $\psi_{\mathcal{G}_B}$ and then apply Theorem F.5.

Step 1: Bounding $\psi_{\mathcal{G}_B}$ By the definition of Orlicz norms, $\psi_{\mathcal{G}_B}$ is given by:

$$\begin{aligned} \psi_{\mathcal{G}_B} &= \left\| \max_{1 \leq i \leq n} \sup_{\nu, \eta} \frac{1}{n} \left| (1 - \xi_i |z'_i| \nu - z''_i \eta)_+^2 - \mathbb{E} \left[(1 - \xi |z'| \nu - z'' \eta)_+^2 \right] \right| \right\|_{\Psi_\alpha} \\ &= \inf \left\{ \lambda > 0 : \mathbb{E} \left[\exp \left(\frac{1}{\lambda} \max_i \sup_{\nu, \eta} \frac{1}{n} \left| (1 - \xi_i |z'_i| \nu - z''_i \eta)_+^2 - \mathbb{E} \left[(1 - \xi |z'| \nu - z'' \eta)_+^2 \right] \right| \right) \right] - 1 \right\} \end{aligned} \quad (29)$$

Applying the triangle inequality and using that ν, η are bounded by constants, we can bound the expectation from the expression above as:

$$\begin{aligned} &\mathbb{E} \left[\exp \left(\frac{1}{\lambda} \max_i \sup_{\nu, \eta} \frac{1}{n} \left| (1 - \xi_i |z'_i| \nu - z''_i \eta)_+^2 - \mathbb{E} \left[(1 - \xi |z'| \nu - z'' \eta)_+^2 \right] \right| \right) \right] \\ &\leq \mathbb{E} \left[\exp \left(\frac{1}{n\lambda} \max_i \sup_{\nu, \eta} (1 - \xi_i |z'_i| \nu - z''_i \eta)_+^2 \right) \right] \exp \left(\frac{1}{n\lambda} \sup_{\nu, \eta} \mathbb{E} \left[(1 - \xi |z'| \nu - z'' \eta)_+^2 \right] \right) \\ &\leq \mathbb{E} \left[\exp \left(\frac{B}{n\lambda} z_{(1)}^2 \right) \right] \exp \left(\frac{B}{n\lambda} \right) \end{aligned}$$

with $B := 3(1 + B_\nu^2 + B_\eta^2)$ and $z_{(1)} = \max_{1 \leq i \leq 2n} z_i = \max\{|z'_1|, |z''_1|, \dots, |z'_n|, |z''_n|\}$. Now we split the expectation from the above inequality into two terms:

$$\tilde{T}_1 = \mathbb{E} \left[1_{[z_{(1)} < \sqrt{2 \log n}] \exp \left(\frac{B}{n\lambda} z_{(1)}^2 \right) \right] \leq \exp \left(\frac{2B \log n}{n\lambda} \right)$$

and

$$\begin{aligned} \tilde{T}_2 &= \mathbb{E} \left[1_{[z_{(1)} \geq \sqrt{2 \log n}] \exp \left(\frac{B}{n\lambda} z_{(1)}^2 \right) \right] = 4n \int_{\sqrt{2 \log n}}^{\infty} \int_{-z_1}^{z_1} \cdots \int_{-z_1}^{z_1} \exp \left(\frac{B}{n\lambda} z_1^2 \right) \prod_{i=1}^{2n} \frac{\exp(-\frac{1}{2} z_i^2)}{\sqrt{2\pi}} dz_i \\ &\leq 4n \int_{\sqrt{2 \log n}}^{\infty} \frac{1}{\sqrt{2\pi}} \exp \left(-\frac{1}{2} z_1^2 \right) \exp \left(\frac{B}{n\lambda} z_1^2 \right) dz_1 \lesssim \frac{n^{\frac{2B}{n\lambda}}}{\sqrt{\log n} (1 - \frac{2B}{n\lambda})} \end{aligned}$$

where we assumed that $\lambda > \frac{2B}{n}$. Now choosing $\lambda = c_\lambda \frac{\log n}{n}$ with a positive constant c_λ sufficiently large, we find that the condition in Equality (29) is satisfied for this λ , which implies that $\psi_{\mathcal{G}_B} \leq c_\lambda \frac{\log n}{n}$.

Step 2: Proof of the statement To apply Theorem F.5, we also need to bound the variance $\sigma_{\mathcal{G}_B}^2$ and use that $\mathbb{E} \|P_n - P\|_{\mathcal{G}_B} \leq 2\mathcal{R}_n(\mathcal{G}_B)$ where $\mathcal{R}_n(\mathcal{G}_B)$ is the Rademacher complexity given by $\mathcal{R}_n(\mathcal{G}_B) = \mathbb{E} \left[\sup_{g_{\nu, \eta} \in \mathcal{G}_B} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i g_{\nu, \eta}(z_i, \xi_i) \right| \right]$. The bound on the variance follows from a straightforward calculation:

$$\sigma_{\mathcal{G}_B}^2 \leq \frac{1}{n} \sup_{g_{\nu, \eta} \in \mathcal{G}_B} \mathbb{E} [g_{\nu, \eta}^2] \leq \frac{1}{n} c_{\sigma_{\mathcal{G}_B}} (1 + B_\nu^4 + B_\eta^4)$$

which holds for some positive universal constant $c_{\sigma_{\mathcal{G}_B}} > 0$.

Next, note that we can upper bound the Rademacher complexity using the triangle inequality and the fact that $(\cdot)_+$ is 1-Lipschitz:

$$\begin{aligned} \mathcal{R}_n(\mathcal{G}_B) &\leq \mathbb{E} \left[\sup_{\nu \leq B_\nu, \eta \leq B_\eta} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i (1 - \xi_i |z'_i| \nu - \eta z''_i)^2 \right| \right] \\ &\lesssim \mathbb{E} \left[\sup_{\nu \leq B_\nu, \eta \leq B_\eta} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i \right| \right] + \mathbb{E} \left[\sup_{\nu \leq B_\nu, \eta \leq B_\eta} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i (\xi_i |z'_i| \nu)^2 \right| \right] \\ &\quad + \mathbb{E} \left[\sup_{\nu \leq B_\nu, \eta \leq B_\eta} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i (\eta z''_i)^2 \right| \right] \lesssim \frac{1}{\sqrt{n}} \end{aligned}$$

where the last line follows from applying the standard concentration results. Thus, we obtain:

$$\mathbb{P} (\|P_n - P\|_{\mathcal{G}_B} \geq 2(1+t)\mathcal{R}_{\mathcal{G}_B} + \epsilon) \leq \exp(-c_2 n \epsilon^2) + 3 \exp\left(-c_3 \frac{n\epsilon}{\log n}\right)$$

with $c_2^{-1} = 2(1+\delta)c_{\sigma_{\mathcal{G}_B}} B^2$ and $c_3^{-1} = Cc_\lambda$, which concludes the proof. \square

G. Proof of Lemma C.4

First note that by taking subgradients we directly obtain $\langle w''_{\text{grad}}, H \rangle = \|H\|_q$ and $\|w''_{\text{grad}}\|_2^2 = \left(\frac{\|H\|_{2q/p}}{\|H\|_q} \right)^{2q/p}$.

Recall the definitions $\mu_{d,q} = \mathbb{E} \|H\|_q$ and $\tilde{\mu}_d := \left(\frac{\mu_{d,2q/p}}{\mu_{d,q}} \right)^{2q/p}$. We use the following well-known concentration result on ℓ_q -norms to control these quantities.

Theorem G.1 (Corollary of Theorem 1 (Paouris et al., 2017)). *For all sufficiently large d and any $2 < q < c_1 \log(d)$ and $\epsilon \in (0, 1)$, we have:*

$$\mathbb{P} (\| \|H\|_q - \mu_{d,q} \geq \epsilon \|H\|_q) \leq c_2 \exp(-c_3 \min(q^2 2^{-q} \epsilon^2 d, (\epsilon d)^{2/q})) \quad (30)$$

where $c_1 \in (0, 1)$, $c_2, c_3 > 0$ are universal constants.

The proof of the lemma follows by showing that the choice of ϵ leads to sufficiently small terms in the exponent on the RHS of Equation (30). More precisely, we show that there exists a universal constant $c > 0$, such that for all n, d sufficiently large with $d \geq n$ and $2 \leq q \leq c \log \log(d)$, we have that

$$\min(q^2 2^{-q} \epsilon^2 d, (\epsilon d)^{2/q}) \geq \log(d), \quad (31)$$

and similarly for $2q/p$ instead of q .

Equation (10) then follows directly from the claim and Theorem G.1. To see Equation (11), we have that with probability $\geq 1 - c_1 d^{-c_2}$,

$$\|H\|_{2q/p}^{2q/p} = \mu_{d,2q/p}^{2q/p} (1 + O(\epsilon))^{2q/p} = \mu_{d,2q/p}^{2q/p} (1 + O(q\epsilon)),$$

where we have used in the last equality that $q \lesssim \log \log(d)$ and that $\epsilon = O(\log^{-1}(d))$ (by assumption on d, n and ϵ). Furthermore, the proof follows from the fact that a similar argument also applies to $\|H\|_q^{2q/p}$.

Thus, the only thing left to show is Equation (31). We separately check the different conditions in Theorem G.1. First, note that it is sufficient to only consider the case where $q = c \log \log(d)$. We separate the proof into two cases where $n^2 \geq d \log^{c_2}(d)$ and $n^2 \leq d \log^{c_2}(d)$ with constant $c_2 > 0$.

Case $n^2 \geq d \log^{c_2}(d) \implies \epsilon = \frac{n}{d}$ Looking only at the first term on the RHS in Equation (31), we have

$$\log(q^2 2^{-q} \epsilon^2 d) = 2 \log(q) - q \log(2) + 2 \log(n) - \log(d) \geq \log \log(d),$$

Furthermore, looking only at the second term on the RHS in Equation (31), we have

$$\frac{2}{q} \log(\epsilon d) = \frac{2}{q} \log(n) \geq \log \log(d),$$

which holds also under weaker assumptions on n, d . Hence the proof for the first case is complete.

Case $n^2 \leq d \log^{c_2}(d) \implies \epsilon = \frac{\log^{c_1}(d)}{n}$. From straight forward calculation, we note that there exists a universal constant $c_1 > 0$ only depending on c_2 such that

$$\log(q^2 2^{-q} \epsilon^2 d) = 2 \log(q) - q \log(2) + \log(d) - 2 \log(n) + 2c_1 \log \log(d) \geq \log \log(d),$$

for any n such that $n^2 \leq d \log^{c_2}(d)$. Furthermore, under the same assumption on d, n , we have

$$\frac{2}{q} \log(\epsilon d) = \frac{2}{q} (\log(d) - \log(n) + c_1 \log \log(d)) \geq \log \log(d).$$

Finally the expectations $\tilde{\mu}_d, \mu_{d,q}$ are directly obtained using the following well-known result on the expectations of the ℓ_q norms of Gaussian vectors

Proposition G.2 (Proposition 2.4 (Paouris et al., 2017), (Schechtman & Zinn, 1989)). *For all $q \leq \log(d)$, we have that $\mu_{d,q} \asymp \sqrt{q} d^{1/q}$.*

This concludes the proof.