

Finite Sample Rates for Logistic Regression with Small Noise or Few Samples

Journal Article**Author(s):**

Kuchelmeister, Felix; van de Geer, Sara

Publication date:

2024

Permanent link:

<https://doi.org/10.3929/ethz-b-000676418>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

Sankhya A, <https://doi.org/10.1007/s13171-024-00358-6>



Finite Sample Rates for Logistic Regression with Small Noise or Few Samples

Felix Kuchelmeister and Sara van de Geer

*ETH Zürich D-MATH: Eidgenössische Technische Hochschule Zurich
Departement Mathematik, Zürich, Switzerland*

Abstract

The logistic regression estimator is known to inflate the magnitude of its coefficients if the sample size n is small, the dimension p is (moderately) large or the signal-to-noise ratio $1/\sigma$ is large (probabilities of observing a label are close to 0 or 1). With this in mind, we study the logistic regression estimator with $p \ll n/\log n$, assuming Gaussian covariates and labels generated by the Gaussian link function, with a mild optimization constraint on the estimator's length to ensure existence. We provide finite sample guarantees for its direction, which serves as a classifier, and its Euclidean norm, which is an estimator for the signal-to-noise ratio. We distinguish between two regimes. In the low-noise/small-sample regime ($\sigma \lesssim (p \log n)/n$), we show that the estimator's direction (and consequentially the classification error) achieve the rate $(p \log n)/n$ - up to the log term as if the problem was noiseless. In this case, the norm of the estimator is at least of order $n/(p \log n)$. If instead $(p \log n)/n \lesssim \sigma \lesssim 1$, the estimator's direction achieves the rate $\sqrt{\sigma p \log n/n}$, whereas its norm converges to the true norm at the rate $\sqrt{p \log n/(n\sigma^3)}$. As a corollary, the data are not linearly separable with high probability in this regime. In either regime, logistic regression provides a competitive classifier.

AMS (2000) subject classification. Primary Statistics; Secondary Linear inference regression 62J12, Generalized linear models.

Keywords and phrases. Binary classification, probit, ridge regularization, halfspace learning, small sample, nonasymptotic.

1 Introduction

We study a binary classification problem, where the covariates x_i are standard Gaussian and the labels y_i stem from a probit model. More precisely, let x_1, \dots, x_n be independent copies of a standard Gaussian random variable taking values in \mathbb{R}^p . The results readily generalize to arbitrary covariance matrices Σ . Fix an unknown $\beta^* \in S^{p-1} := \{\beta \in \mathbb{R}^p : \|\beta\|_2 = 1\}$, and for $i \in \{1, \dots, n\}$ define $y_i := \text{sign}(x_i^T \beta^* + \sigma \epsilon_i)$, where ϵ_i are independent standard Gaussians, and $\sigma > 0$ is a fixed, unknown parameter - the *noise-to-*

signal ratio. In the statistics literature, the domain of the labels y_i is often $\{0, 1\}$. Moreover, it is common to fix $\sigma = 1$ and allow the length of β^* to be unknown. This leads to the same model, taking $\|\beta^*\|_2 := 1/\sigma$. If both $1/\sigma$ and $\|\beta^*\|_2$ were unknown, the model would not be identifiable. Those two quantities should be understood as a single parameter, the signal-to-noise ratio, in our case $1/\sigma$.

To make the parameter σ more tangible, we mention that it determines the probability of observing a ‘wrong’ label, i.e. $\text{sign}(x^T \beta^*) \neq y$. For example, if $\sigma = 1$, then $\mathbb{P}[y \neq \text{sign}(x^T \beta^*)] = 1/4$. Moreover, for small σ , the probability of observing a wrong label behaves as σ up to multiplicative constants.

The classical choice to estimate β^*/σ would be maximum-likelihood estimation, the *probit* model, which is a *generalized linear model* (McCullagh and Nelder, 1989). The logistic regression estimator has very nice mathematical properties, which pair well with the probit model. For this reason, we will use it to estimate β^* and σ . The logistic regression estimator $\hat{\gamma}$ is a solution to the following objective:

$$\arg \min_{\gamma \in \mathbb{R}^p} \sum_{i=1}^n \log(1 + \exp(-y_i x_i^T \gamma)). \quad (1)$$

The estimator $\hat{\gamma}$ can be decomposed into its length $\hat{\tau} := \|\hat{\gamma}\|_2$, and its direction $\hat{\beta} := \hat{\gamma}/\|\hat{\gamma}\|_2$. The label predicted by $\hat{\gamma}$ at a point $x \in \mathbb{R}^p$ is $\text{sign}(x^T \hat{\gamma})$, which coincides with $\text{sign}(x^T \hat{\beta})$. In other words, the classification error is only dependent on the orientation $\hat{\beta}$. On the other hand, $\hat{\tau}$ serves as a natural estimate of the signal-to-noise ratio $1/\sigma$.

Logistic regression is perhaps *the* canonical classification method in statistics, it has been studied for several decades. Nonetheless, it has not ceased to challenge and surprise theoreticians and practitioners alike. In the following, we look into a few problems that this setup poses.

1.1 Problems of Logistic Regression

1.1.1 Linear Separation. The first problem which is encountered in theory and practice is linear separation. We say that the data are *linearly separable*, if there exists a $\gamma \in \mathbb{R}^p$, such that for all $i \in \{1, \dots, n\}$, $y_i x_i^T \gamma > 0$. In this case, Eq. 1 has no solution in \mathbb{R}^p . To see this, take any γ' which separates the data. As we multiply γ' with a positive scalar tending to infinity, the loss function in Eq. 1 vanishes.

To develop a rigorous non-asymptotic analysis of logistic regression, one must specify how linear separation is dealt with. If one is solely interested in estimating the direction β^* and not the signal-to-noise ratio $1/\sigma$, one can

allow for infinite values $\|\hat{\gamma}\|_2 \rightarrow \infty$. In this case, given linear separation, the set of solutions to Eq. 1 are all estimators, which separate the data. Alternatively, one can modify the objective (1). For example, Agresti (2015) suggests Bayesian methods or penalized likelihood functions. We constrain the estimator to a Euclidean ball of radius $M > 0$, $B_M := \{\gamma \in \mathbb{R}^p : \|\gamma\|_2 \leq M\}$:

$$\arg \min_{\gamma \in B_M} \sum_{i=1}^n \log(1 + \exp(-y_i x_i^T \gamma)). \quad (2)$$

In our analysis, we allow the radius M to be arbitrarily large - but finite. Consequentially, we could choose M so large, that it coincides with the unconstrained objective (1), but for the case where the data are linearly separable. This is the least intrusive modification of the logistic regression, which still enforces that it takes finite values with probability 1. By Lagrangian duality, there exists a $\lambda \geq 0$, such that (1) with ridge penalty $+\lambda\|\gamma\|_2^2$ coincides with Eq. 2 (this is also known as ℓ^2 - or Tikhonov-regularization). Whenever $\lambda > 0$, the coefficients of the ridge estimator $\hat{\gamma}$ are biased towards zero, which is not the case for Eq. 2. On the other hand, if $\lambda = 0$, we have Eq. 1, and the problem of linear separation appears again. Hence, we prefer (2), the constrained form.

In some settings, linear separation occurs with vanishingly small probability and hence can be neglected. This is the case in the classical asymptotic literature (van der Vaart, 2000). In the high-dimensional asymptotic setting, where both p and n tend to infinity, Candès and Sur (2020) have shown that the data are linearly separable, if and only if the ratio p/n in the limit is lower than a constant, depending on the signal-to-noise ratio $1/\sigma$. They study logistic noise. This result has been generalized in Montanari et al. (2019), which for example includes the Gaussian noise model (probit). For finite n and p , Cover (1965) showed that if there is no signal (y and x are independent, i.e. $\sigma \rightarrow \infty$), that the probability of linear separation is exactly $2^{n-1} \sum_{k=0}^{p-1} \binom{n-1}{k}$. However, for finite n the literature on linear separability in generalized linear models is still lacking.

Although we are primarily concerned with providing guarantees for the estimated signal strength $\|\hat{\gamma}\|_2$ and the estimated direction $\hat{\gamma}/\|\hat{\gamma}\|_2$, our results allow conclusions for the linear separation problem. From Theorem 2.1.1, we derive an upper bound for the probability of linear separation. Namely, if $p \log(n)/n \lesssim \sigma \lesssim 1$, the data are not linearly separable with large probability (see Proposition 2.3.1).

1.1.2 Overestimating Coefficient Magnitude. Logistic regression appears to overestimate the signal-to-noise ratio $1/\sigma$ in some settings. In other words, the magnitude of the coefficients $|\hat{\gamma}_j|$ is too large, there is a bias away from zero. For example, this was observed in Nemes et al. (2009) in a simulation where this upward bias was especially present for smaller sample sizes. A similar observation was made in a simulation in Sur and Candès (2019). They performed a simulation with $n = 4000$, $p = 800$ and $\sigma = 1/5$ (using logistic noise). Moreover, they proved that under some regularity conditions in a high-dimensional asymptotic setting (letting $n, p \rightarrow \infty$ at a constant ratio), the estimator $\hat{\gamma}$ does not converge to β^*/σ , but instead fluctuates around $\alpha_*\beta^*/\sigma$, where $\alpha_* > 1$. In other words, the magnitude of the coefficients is overestimated.

In Theorem 2.2.1, we see that if $\sigma \lesssim (p \log n)/n$, the estimate of the signal-to-noise ratio $\hat{\tau}$ is at least of order $n/(p \log n)$, although it may be larger. On the other hand, we see in Theorem 2.1.1 that if $\sigma \gtrsim (p \log n)/n$, the deviation of $\hat{\tau}$ from its target is of order no larger than $\sqrt{(p \log n)/(n\sigma^3)}$. Our findings suggest that in the regime $\sigma \gtrsim (p \log n)/n$ the magnitude of the coefficients can be reliably estimated.

1.1.3 Inadequacy of Classical Asymptotic Approximations. It is known that as $n \rightarrow \infty$, for fixed p and σ , the logistic regression estimator converges in distribution to a normally distributed random variable. This is exploited in statistical software, such as the function *glm* in the R-package ‘stats’, to provide approximate p-values.

However, classical asymptotic approximations can be very poor in logistic regression. A famous early reference for this is Hauck Jr and Donner (1977), who showed that the Wald statistic is unreliable if the signal-to-noise ratio is large. A more recent example is Candès et al. (2016), who observed in a simulation with $n = 500$, $p = 200$ and no signal (y independent of x , in our language $\sigma \rightarrow \infty$), that the p-values from asymptotic maximum likelihood theory for the global null-hypothesis are left-skewed. The same observation was made in Sur and Candès (2019) with $n = 4000$, $p = 800$.

An attempt to overcome these shortcomings is to use a different regime, such as high dimensional asymptotics (Montanari et al., 2019, Sur and Candès, 2019). While in this regime, the bias of the coefficients becomes visible, it only applies when the maximum likelihood estimator exists asymptotically almost surely, i.e. when the data are not linearly separable. For a finite number of samples, however, there is no way around dealing with possible linear separability.

1.2 Related Results The literature for finite sample guarantees for unregularized logistic regression in the regime $p < n$ is still lacking, in particular for the case of small noise σ . We consider the case $p \ll n/\log n$. There are finite sample results for logistic regression in high dimensions $p > n$ (with Lasso penalty, e.g. van de Geer (2008), and ℓ^0 -penalty, e.g. Abramovich and Grinshtein (2018)).

In the noiseless case, if the radius M is large enough, the logistic regression estimator always separates the data. Hence, it is also a minimizer of the 0/1-loss. It is known that in the noiseless case, the optimal rate for estimating β^* is p/n (Long, 1995), which is achieved by any such minimizer of the 0/1-loss (Long, 2003). We show that logistic regression obtains the same rate - up to the log term as if the problem were noiseless - even in the presence of a little noise, namely $\sigma \lesssim (p \log n)/n$. If instead $(p \log n)/n \lesssim \sigma \lesssim 1$, the classification error achieves the parametric rate $\sqrt{(\sigma p \log n)/n}$.

Some state-of-the-art finite-sample rates for logistic regression can be obtained in Ostrovskii and Bach (2021).¹ They show the rate $\sqrt{(1 + 1/\sigma^3)p/n}$ for $\|\hat{\gamma} - \gamma^*\|_2$. In the regime $1 \lesssim 1/\sigma \lesssim n/(p \log n)$, we obtain the same rate with a logarithmic factor (see Theorem 2.1.1 and Lemma 3.3.1). In the regime $\sigma \lesssim p \log(n)/n$, such a control without further constraints is generally not possible. To see why, note that if σ is small enough, the probability of linear separability is close to 1. Consequentially, for small enough σ , $\|\hat{\gamma} - \gamma^*\|_2 = +\infty$ with large probability.

Our analysis exposes that this slow dependence $\sqrt{(1 + 1/\sigma^3)p/n}$ on the signal-to-noise ratio $1/\sigma$ is only present when estimating the signal strength $\tau^* := \|\gamma^*\|_2 \sim 1/\sigma$, and not when performing classification. In fact, Theorems 2.1.1 and 2.2.1 show that the classifier of logistic regression $\hat{\beta} := \hat{\gamma}/\|\hat{\gamma}\|_2$ converges faster to the target classifier $\beta^* = \gamma^*/\|\gamma^*\|_2$, if the signal-to-noise ratio is larger, with rates $\|\hat{\beta} - \beta^*\|_2 \lesssim \sqrt{\sigma p \log(n)/n}$ if $p \log(n)/n \lesssim \sigma \lesssim 1$ and $\|\hat{\beta} - \beta^*\|_2 \lesssim p \log(n)/n$ if $\sigma \lesssim p \log(n)/n$. To obtain these results, it was crucial in our analysis to treat the estimated signal-to-noise ratio $\hat{\tau} := \|\hat{\gamma}\|_2$ and the estimated classifier $\hat{\beta}$ separately.

Shortly after the appearance of the first draft of this work, Hsu and Mazumdar (2024) appeared. They study finite sample classification with

¹ The work Ostrovskii and Bach (2021) does not deal with linear separation, while assuming that the parameter space is a subset of \mathbb{R}^p . So, for their results to apply, one technically needs to assume that τ^* and $\hat{\tau}$ are bounded. Otherwise, the estimator is not well-defined with positive probability. However, the statistician then needs to know an upper bound on the length of the target $\|\gamma^*\|_2 \sim 1/\sigma$. We do not need to make such an assumption on τ^* . However, we impose that $\hat{\tau} \leq M$, where $M \geq n/(p \log(n) + t)$. Although this is a technical contribution of our work, we suppose that such an argument could be added to works avoiding the discussion of linear separability, such as Ostrovskii and Bach (2021).

logistic noise. Their parameter space is S^{p-1} , so they estimate only the direction β^* , not the signal-to-noise ratio $\tau^* \sim 1/\sigma$. Although they study logistic noise, while we study Gaussian noise, the results should be comparable. They also consider the case $\sigma \gtrsim 1$, which is out of the scope of this work. Moreover, they provide upper bounds that match their lower bounds up to constant factors, but with estimators other than the logistic regression estimator.

They show that in order to achieve a rate ε for $\|\beta^* - \hat{\beta}\|_2$, at least $n \gtrsim \sigma p/\varepsilon^2$ samples are needed, and if $\sigma \lesssim \varepsilon$ then $n \gtrsim p/\varepsilon$ samples are needed (ignoring a logarithmic factor in ε). Up to logarithmic terms, this matches our upper bounds in Theorems 2.1.1 and 2.2.1.

1.3 Outline In this paper, we prove guarantees for the estimated signal-to-noise ratio $\hat{\tau} := \|\hat{\gamma}\|_2$ and the estimated direction $\hat{\beta} := \hat{\gamma}/\|\hat{\gamma}\|_2$, for the logistic regression estimator constrained to the Euclidean ball of radius $M > 0$, see Eq. 2. We distinguish between two regimes. The first regime covers the case $\sigma \gtrsim (p \log n)/n$, whereas the second regime covers $\sigma \lesssim (p \log n)/n$, the “large noise case” and the “small noise case”. We prove two theorems, one for each case. They are presented in Section 2 (Theorems 2.1.1 and 2.2.1). There, we provide sketches of the two proofs, along with a consequence for the probability of linear separability (Proposition 2.3.1). The proofs of the main theorems are given in Section 6, where we prove two slightly more general statements (Theorems 6.1.1 and 6.2.1). Sections 3, 4, 5, as well as Appendix A all lead up to the proof of the main theorems, see Fig. 1.

Section 3 establishes the relationships between classification and Euclidean geometry, the noise-to-signal ratio σ and the length of the target $\tau^* := \|\gamma^*\|_2$, and introduces the $*$ -norm, which is used to quantify the distance $d_*(\hat{\gamma})$ between the estimator $\hat{\gamma}$ and its target γ^* in the large noise case.

Throughout our analysis, we decompose the logistic loss as follows:

$$\log(1 + \exp(-yx^T \gamma)) = \log(1 + \exp(-|x^T \gamma|)) + |x^T \gamma| \mathbb{1}\{yx^T \gamma < 0\}.$$

Introducing some notation, we re-state this equality as:

$$l(\gamma, x, y) = b(\gamma, x) + u(\gamma, x, y).$$

We refer to b as the ‘bounded term’ and call u the ‘unbounded term’. Thanks to this decomposition, we can prove strong upper bounds on their




Figure 1: This figure shows the structure of the paper. The main results are given in Section 2 and are proved in 6. These proofs rely on Sections 3, 4 and 5. Appendices A.1 and A.2 contain supporting inequalities for Sections 4 and 5

empirical processes in Sections 4 and 5, which in turn gives us a strong control of the empirical process of l .

1.4 Notation

- n is the number of observations and $p < n$ is the dimension of the problem.
- $x \sim \mathcal{N}(0, I_p)$ are the observed covariates.
- $\epsilon \sim \mathcal{N}(0, 1)$ is unobserved noise independent of x .
- $S^{p-1} := \{\beta \in \mathbb{R}^p : \|\beta\|_2 = 1\}$ is the unit sphere. For $\delta > 0$ and $\beta' \in S^{p-1}$, we use the notation $S(\delta, \beta') := \{\beta \in S^{p-1} : \|\beta - \beta'\|_2 \leq \delta\}$ for a spherical cap.
- For $r > 0$, $B_r := \{\gamma \in \mathbb{R}^p : \|\gamma\|_2 \leq r\}$ is the closed Euclidean ball of radius r . We use the notation B_r° for the open ball of radius $r > 0$.
- $\beta^* \in S^{p-1}$ is the unknown orientation we want to learn.
- $\sigma > 0$ is the noise-to-signal ratio.
- $y = \text{sign}(x^T \beta^* + \sigma \epsilon)$ are the observed labels.
- $M > 0$ is a known hyperparameter, chosen by the statistician. It determines the upper bound on the Euclidean norm of the estimator, defined below.

- We use P to denote the expectation and P_n to denote the average. More formally, P_n is the expectation with respect to the empirical measure, that is $n^{-1} \sum_{i=1}^n \delta_{(x_i, y_i)}$.
- The logistic loss function is the following mapping:

$$l : \mathbb{R}^p \times \mathbb{R}^p \times \{-1, 1\} \rightarrow (0, \infty), \quad (\gamma, x, y) \mapsto \log(1 + \exp(-yx^T \gamma)).$$

- Our estimator is:

$$\hat{\gamma} := \arg \min_{\gamma \in B_M} \sum_{i=1}^n \log(1 + \exp(-y_i x_i^T \gamma)) = \arg \min_{\gamma \in B_M} P_n l(\gamma, x, y).$$

- The target of the estimator is:

$$\gamma^* := \arg \min_{\gamma \in \mathbb{R}^p} P l(\gamma, x, y).$$

- We will frequently decouple a vector γ into its orientation $\beta := \gamma / \|\gamma\|_2$ and its length $\tau := \|\gamma\|_2$. This may seem like an abuse of notation, as β^* would both denote the unknown orientation generating the data, and $\gamma^* / \|\gamma^*\|_2$. However, we will see in Lemma 3.2.3 that those two coincide. As a side remark, we will also see that $1/\sigma$ and $\tau^* := \|\gamma^*\|_2$ generally do not coincide (Lemma 3.2.1).
- We write the logistic loss as the function $l : \mathbb{R}^p \times \mathbb{R}^p \times \{-1, 1\} \rightarrow \mathbb{R}$, mapping $(\gamma, x, y) \mapsto \log(1 + \exp(-yx^T \gamma))$. Occasionally, we omit the arguments x, y , writing simply $l(\gamma)$.
- The loss function l can be decomposed into two parts. One is bounded b , the other unbounded u :

$$\begin{aligned} \log(1 + \exp(-yx^T \gamma)) &= \log(1 + \exp(-|x^T \gamma|)) + |x^T \gamma| \mathbf{1}\{yx^T \gamma < 0\} \\ &=: b(\gamma, x) + u(\gamma, x, y). \end{aligned}$$

If clear from context, we omit the arguments x and y , writing simply $b(\gamma)$ and $u(\gamma)$. We will frequently use the (un)bounded term, subtracting the term evaluated at γ^* :

$$\tilde{b}(\gamma) := \frac{b(\gamma) - b(\gamma^*)}{\log 2}.$$

When studying the large noise case, we will use the notation, $\tilde{u}(\gamma) := u(\gamma) - u(\gamma^*)$ whereas in the small noise case, we use $\tilde{u}(\beta) := u(\beta) - u(\beta^*)$. Moreover, we define the following set of functions:

$$\mathcal{B} := \left\{ \tilde{b}(\gamma) : \gamma \in \mathbb{R}^p \right\}.$$

- We use $e_j \in \mathbb{R}^p$ to denote the j -th coordinate unit vector. For example, $e_1 = (1, 0, \dots, 0)$.
- We define a weighted Euclidean norm:

$$\|\cdot\|_* : \mathbb{R}^{p+1} \rightarrow [0, \infty), \quad (\tau, \beta) \mapsto \sqrt{\frac{|\tau|^2}{\tau^{*3}} + \tau^* \|\beta\|_2^2}.$$

- We define a metric, to quantify the distance between a vector γ and γ^* in terms of the norm $\|\cdot\|_*$:

$$d_* : \mathbb{R}^p \setminus \{0\} \rightarrow [0, \infty), \quad \gamma \mapsto \left\| \left(\|\gamma\|_2, \frac{\gamma}{\|\gamma\|_2} \right) - (\tau^*, \beta^*) \right\|_*.$$

- We define the ball of radius δ in the $\|\cdot\|_*$ -norm, centered at (τ^*, β^*) as:

$$B_\delta^* := \{ \gamma \in \mathbb{R}^p \setminus \{0\} : d_*(\gamma) \leq \delta \}.$$

- For the ‘peeling’ argument, we will define the ‘peels’ $B_{r,R}^* := B_R^* \setminus B_r^*$ for $r, R > 0$.
- In some occasions, to compare two quantities f, g , we use the notation $f \lesssim g$. By this, we mean that there exists an absolute constant $C > 0$, not depending on any model parameters such as n, p, σ, β^* , etc, such that $f \leq Cg$. If $f \lesssim g$ and $g \lesssim f$, we write $f \sim g$.

2 Main Results

Here, we state the main results of the paper: Theorems 2.1.1 and 2.2.1. We start with the regime where $(p \log n)/n \lesssim \sigma \lesssim 1$, and continue with the regime $\sigma \lesssim (p \log n)/n$. The proofs of both theorems are given in Section 6, although we provide a sketch for each in this section. Finally, we show that the data are not linearly separable with large probability in the first regime, which follows from the behavior of the logistic regression estimator.

2.1 Regime 1 (“Large Noise”)

2.1.1 Statement of the Result. We state the main result for regime 1, where the noise is ‘large’. The proof is given in Section 6, where we prove the slightly stronger Theorem 6.1.1. The exact universal constants can be found in the proof. Recall that $\hat{\beta} := \hat{\gamma}/\|\hat{\gamma}\|_2$ and $\hat{\tau} := \|\hat{\gamma}\|_2$, where $\hat{\gamma}$ is the logistic regression estimator constrained to a Euclidean ball of radius M , see Eq. 2.

Theorem 2.1.1 *For any $t > 0$, if*

$$\frac{p \log n + t}{n} \lesssim \sigma \leq \frac{1}{\sqrt{6}}, \quad M \geq \frac{n}{p \log n + t},$$

then, with probability at least $1 - 4 \exp(-t)$,

$$\|\hat{\beta} - \beta^*\|_2 \lesssim \sqrt{\sigma \frac{p \log n + t}{n}}, \quad |\hat{\tau} - \tau^*| \lesssim \sqrt{\frac{1}{\sigma^3} \frac{p \log n + t}{n}}.$$

In other words, as long as σ is not too small (or n is very large) and M is large enough, up to a factor involving σ , we get the rate $\sqrt{(p \log n)/n}$ both for the estimates of the orientation as well as the signal-to-noise ratio (for the relationship between τ^* and $1/\sigma$, see Lemmas 3.2.1 and 3.2.2).

The result translates into a bound for $\|\hat{\gamma} - \gamma^*\|_2$ of order $\sqrt{\tau^{*3}(p \log n)/n}$, using Lemma 3.3.1. This matches the rate in Ostrovskii and Bach (2021). Theorem 2.1.1 reveals a more subtle insight into how the signal-to-noise ratio (SNR) $1/\sigma$ affects the estimated orientation $\hat{\beta}$ and the estimated SNR $\hat{\tau}$. As the noise level is smaller, it is easier to classify, but harder to make statements about the SNR, see also our discussion in Section 3.3.

The assumption $\sigma \leq 1/\sqrt{6}$ is made due to technical reasons. In other words, the probability of observing a wrong label ($y \neq \text{sign}(x^T \beta^*)$) is no more than $\arccos(1/\sqrt{1+1/6})/\pi \approx 0.123$ (see Eq. 4). We rely on this assumption in particular when using lower bounds on Gaussian integrals, see Appendix A.1. Observe that in the case where $\sigma \geq 1/\sqrt{6}$ does not grow with n, p, t , we are back in the classical regime, where Ostrovskii and Bach (2021) provides the optimal rate of $\sqrt{p/n}$.

2.1.2 Idea of the Proof. The proof of Theorem 2.1.1 is given in Section 6. Here, we provide a sketch of the main ideas. To obtain a fast rate, we use localization. In other words, we exploit that the empirical process of the excess risk is smaller if $\hat{\gamma}$ is closer to γ^* . To exploit this, we first create a convex combination $\tilde{\gamma} := \alpha \hat{\gamma} + (1 - \alpha) \gamma^*$. Here, α is chosen such that on the one hand $\|\tilde{\gamma} - \gamma^*\|_2 \leq \tau^*/6$, but on the other hand once $\|\tilde{\gamma} - \gamma^*\|_2$ is small enough, $\|\hat{\gamma} - \gamma^*\|_2 \leq 2\|\tilde{\gamma} - \gamma^*\|_2$. Moreover, as we assume $\tau^* \leq M$,

it holds that $P_n l(\tilde{\gamma}) \leq P_n l(\gamma^*)$. By convexity, also $P_n l(\tilde{\gamma}) \leq P_n l(\gamma^*)$. The separate Lemma 6.1.5 shows that for such $\tilde{\gamma}$, the distance of $\tilde{\gamma}$ to γ^* in the $*$ -norm is small with high probability, i.e. $d_*(\tilde{\gamma}) \lesssim \sqrt{(p \log(n) + t)/n}$. On this event of large probability, it follows that $\tilde{\gamma}$ does not deviate from γ^* by more than $\tau^*/6$ in Euclidean distance. It follows that $\|\hat{\gamma} - \gamma^*\|_2 \leq \tau^*/6$ too. So, we can apply Lemma 6.1.5 again, with $\hat{\gamma}$ taking the role of $\tilde{\gamma}$, which allows us to conclude the proof.

So, the heart of the proof is Lemma 6.1.5. It states that the localization $\|\tilde{\gamma} - \gamma^*\|_2 \leq \tau^*/6$, together with the condition that the empirical loss of $\tilde{\gamma}$ is smaller than the empirical loss of γ^* allows to conclude that $d_*(\tilde{\gamma}) \lesssim \sqrt{(p \log(n) + t)/n}$. This expression $d_*(\tilde{\gamma})$ occurs naturally from the Taylor expansion of the excess risk. It provides a lower bound for the latter, which in turn can be bounded by the empirical process:

$$d_*(\tilde{\gamma})^2 := \frac{|\tilde{\tau} - \tau^*|^2}{\tau^{*3}} + \tau^* \|\tilde{\beta} - \beta^*\|_2^2 \lesssim P(l(\tilde{\gamma}) - l(\gamma^*)) \leq (P - P_n)(l(\tilde{\gamma}) - l(\gamma^*)).$$

Here, as throughout this paper, we use the notation $\tilde{\tau} := \|\tilde{\gamma}\|_2$ and $\tilde{\beta} := \tilde{\gamma}/\|\tilde{\gamma}\|_2$. This step already relies on the localization argument through the condition $|\tilde{\tau} - \tau^*|_2 \leq \tau^*/6$. Moreover, it requires that $\tilde{\tau}, \tau^* \geq \sqrt{6 + \sqrt{51}}$. These conditions appear throughout the proof of the large noise case. They stem from lower bounds on Gaussian integrals, see Appendix A.1.

To upper bound the empirical process, we split the loss into the bounded and unbounded parts. These two are controlled with Bousquet's and Bernstein's inequality, together with covering and peeling arguments. In both parts, we rely on the localization and that $\tau^* \gtrsim 1$ in several parts of the proof, for the same reasons as when proving the margin condition. These calculations are given in Sections 4 and 5. We look at two cases, depending on whether the bounded or unbounded term dominates. By doing so, we arrive at a bound that allows us to conclude that $d_*(\tilde{\gamma}) \lesssim \sqrt{(p \log(n) + t)/n}$, which completes the proof.

Remark 2.1.1 In the part of the proof where we assume that the bounded term dominates, the logarithmic factor $\log(en/p)$ would be sufficient. In the unbounded part, however, we incur the logarithmic factor $\log(p)$. Together, this is equivalent to the factor $\log(n) = \log(n/p) + \log(p)$. The same holds for the small noise case.

2.2 Regime 2 (“Small Noise”)

2.2.1 Statement of the Result. We state the main result for regime 2, where the noise is ‘small’. The proof is given in Section 6. There, we provide a slightly more general version of the theorem.

Theorem 2.2.1 *For any $t > 0$, if*

$$\sigma \leq \frac{p \log n + t}{n} \lesssim 1, \quad M \geq \frac{n}{p \log n + t},$$

then, with probability at least $1 - 6 \exp(-t)$,

$$\|\hat{\beta} - \beta^*\|_2 \lesssim \frac{p \log n + t}{n}, \quad \hat{\tau} \gtrsim \frac{n}{p \log n + t}.$$

We remark that the constants can be chosen, such that there is no gap between the two regimes. The more general Theorem 6.2.1 allows arbitrary values $\sigma \lesssim 1$. Hence, instead of $\sigma \leq (p \log n + t)/n$, one could impose the condition $\sigma \leq c(p \log n + t)/n$, with a constant c matching the universal constant in the large noise case Theorem 2.1.1. We chose $c = 1$ here for aesthetic reasons.

The first interpretation of this result is, that if σ is close to zero, the classifier $\hat{\beta}$ achieves the fast rate $(p \log n)/n$. This is the same rate as a minimizer of the 0-1 loss obtains in the noiseless case, up to the factor $\log n$ (see Long (2003)). Thus, in some sense, Theorem 2.2.1 tells us that if $\sigma \leq (p \log n)/n$, the logistic regression classifier performs as if there was no noise.

There is another viewpoint, that may be interesting from a practical standpoint. The condition could be rewritten as $n/\log(n) \leq p/\sigma$. Informally, this suggests collecting data until regime 1 is reached, as improving the classifier is ‘cheaper’ in regime 2. Moreover, the first regime also brings the benefit that the signal-to-noise ratio can be estimated with $\hat{\tau}$, whereas here we only get a lower bound. We note that this small sample conclusion which Theorem 6.2.1 allows is of a more philosophical than practical nature since the constants in the theorem are too large to be interesting for small sample sizes.

2.2.2 Idea of the Proof. The proof of Theorem 2.2.1 is given in Section 6. Here, we provide a sketch of the main ideas. In the small noise regime, we cannot hope to estimate $1/\sigma \sim \tau^*$ (see e.g. the discussion in Section 3.3). However, we should be able to classify well and thus estimate $\hat{\beta}$. Furthermore, we can show that $\hat{\tau} := \|\hat{\gamma}\|_2$ is somewhat large.

FINITE SAMPLE RATES FOR LOGISTIC REGRESSION...

The big technical obstacle is, that we have no lower bound on the noise-to-signal ratio σ . The first problem this brings is that no matter how large we choose M , since we do not know σ it may be that $\|\gamma^*\|_2 > M$. If so, then potentially $P_n l(\hat{\tau}\hat{\beta}) > P_n l(\tau^*\beta^*)$. However, $P_n l(\hat{\tau}\hat{\beta}) \leq P_n l(\hat{\tau}\beta^*)$ still holds. While comparing $\hat{\tau}\hat{\beta}$ to $\hat{\tau}\beta^*$ is not enough to estimate $\|\hat{\gamma} - \gamma^*\|_2$, this is not our ambition here in the first place. It is enough to allow us to make statements about $\|\hat{\beta} - \beta^*\|_2$ and $\hat{\tau}$, which is all that we want.

Therefore, we compare $\hat{\tau}\hat{\beta}$ with $\hat{\tau}\beta^*$. Here, the excess risk is directly linked to the Euclidean distance as follows:

$$\hat{\tau}\|\hat{\beta} - \beta^*\|_2^2 \sim P(u(\hat{\tau}\hat{\beta}) - u(\hat{\tau}\beta^*)) = P(l(\hat{\tau}\hat{\beta}) - l(\hat{\tau}\beta^*)).$$

If we had access to a good lower bound on $\hat{\tau}$, we would be close to being finished. However, at this point no lower bound for $\hat{\tau}$ is available (although our proof eventually provides such a lower bound). A strategy to overcome this obstacle is given by the inequality of arithmetic and geometric means:

$$\|\hat{\beta} - \beta^*\|_2 = \sqrt{\frac{1}{\hat{\tau}} \sqrt{\hat{\tau}\|\hat{\beta} - \beta^*\|_2^2}} \leq \frac{1}{2\hat{\tau}} + \frac{\hat{\tau}\|\hat{\beta} - \beta^*\|_2^2}{2}.$$

This shows that if we could get the extra summand $1/\hat{\tau}$, then we could upper bound the distance between $\hat{\beta}$ and β^* with the empirical process of the loss function. We obtain this extra summand using that:

$$\frac{1}{\hat{\tau}} \sim Pb(\hat{\tau}\hat{\beta}) - Pb(2\hat{\tau}\beta^*). \tag{3}$$

Simultaneously, the unbounded term is homogeneous in $\hat{\tau}$, and easy to control for small σ . This brings us to the comparison of $\hat{\tau}\hat{\beta}$ with $2\hat{\tau}\beta^*$. Unfortunately, it is too early to celebrate. Now we face the issue that possibly $2\hat{\tau} > M$, which we wanted to avoid in the first place. However, this can be resolved with a case distinction: We either consider the case $M/2 \leq \hat{\tau} \leq M$ or $\hat{\tau} \leq M/2$.

In case 1 we assume $M/2 \leq \hat{\tau} \leq M$, so we have excellent control on $\hat{\tau} \sim M$. We then use that:

$$\frac{1}{\hat{\tau}} \lesssim Pb(\hat{\tau}\hat{\beta}) - Pb(2\hat{\tau}\beta^*) \lesssim Pb(\hat{\tau}\hat{\beta}) - Pb(\hat{\tau}\beta^*) + \frac{1}{M}.$$

Since we require that the radius M is large, this term is fine. We are back to comparing $\hat{\tau}\hat{\beta}$ with $\hat{\tau}\beta^*$, so we can use that $P_n l(\hat{\tau}\hat{\beta}) \leq P_n l(\hat{\tau}\beta^*)$. We proceed to upper bound the bounded and unbounded terms separately, using

Bousquet’s inequality and Bernstein’s inequality. In these upper bounds, we benefit from the assumption that σ is small, to get fast rates.

In case 2, we can directly compare $\hat{\tau}\hat{\beta}$ with $2\hat{\tau}\beta^*$ since $2\hat{\tau} \leq M$ and so $P_n l(\hat{\tau}\hat{\beta}) \leq P_n l(2\hat{\tau}\beta^*)$. The upper bounds on the bounded and unbounded terms are proved similarly as in case 1. Although here, we additionally need to prove that $\hat{\tau}$ is lower bounded by $n/(p \log(n) + t)$. This is possible since $1/\hat{\tau}$ is a lower-bound for the expectation of the difference $Pl(\hat{\tau}\hat{\beta}) - Pl(2\hat{\tau}\beta^*)$, which we show is small.

There is a final complication, which we have not mentioned yet. For the lower bound in the trick (3) to work, we need to assume that $\hat{\tau}$ is somewhat large. Since if $\hat{\tau}$ is upper-bounded by a constant (say $\sqrt{6}$), then in fact:

$$\hat{\tau} \sim Pb(\hat{\tau}\hat{\beta}) - Pb(2\hat{\tau}\beta^*),$$

see Lemma A.1.6. This is a problem, as we intended to use (3) to show that $\hat{\tau}$ cannot be too small. What saves us is that the event $\hat{\tau} \lesssim 1$ occurs with small probability, if n is large and τ^* is large. Intuitively, as $\hat{\gamma}$ ‘converges to’ γ^* , it cannot stay in a ball with a small constant radius with large probability, if the target γ^* lies outside this ball. We thus show that the event $\hat{\tau} \leq \sqrt{6}$ does not occur with large probability if n is somewhat large. To do so, we show that on the event $\hat{\tau} \leq \sqrt{6}$, with large probability,

$$\sqrt{\frac{p \log(n) + t}{n}} \gtrsim Pl(\hat{\tau}\hat{\beta}) - Pl(2\sqrt{6}\beta^*) \geq Pl(\sqrt{6}\beta^*) - Pl(2\sqrt{6}\beta^*) \sim 1.$$

Here, we use that the signal-to-noise ratio is large, in particular $\tau^* > \sqrt{6}$. This leads to a contradiction as soon as n is large enough, so this case cannot occur. The proof is thereby completed.

2.3 A Consequence About Linear Separation Theorem 2.1.1 also allows a consequence for linear separation. In particular, Proposition 2.3.1 below shows that in regime 1 the data are not linearly separable with large probability. This means, for example, that logistic regression in its original formulation can be applied, as the classical optimization problem (1) is well-posed with large probability. This is a step towards a non-asymptotic analog of the results in Candès and Sur (2020).

Proposition 2.3.1 *For any $t > 0$, if:*

$$\frac{p \log n + t}{n} \lesssim \sigma \leq \frac{1}{\sqrt{6}},$$

then, the data are linearly separable with probability at most $4 \exp(-t)$.

The proof is given in Section 6.3.

3 Geometry & Fisher Consistency

In this section, we introduce some results which are fundamental in several parts of the proofs. First, we state Grothendieck’s identity, which provides us with a duality between the classification error and the Euclidean distance. Second, we show how to recover the unknown parameters β^* and σ , given that we have a misspecified model. Third, we introduce a new metric, which is tailored to our problem, allowing us to lower-bound the excess risk with a distance between our estimator and its target.

3.1 The Classification Error and the Euclidean Distance Here, we recall the duality between the classification error for linear classifiers and the angle between two vectors, which is given by Grothendieck’s identity. We show how to translate this to Euclidean distances. Finally, we exploit this equivalence and the Gaussianity of the vector $(x^T, \epsilon)^T$ to control the probability of observing a “wrong” label.

The label predicted by a linear classifier $\gamma \in \mathbb{R}^p \setminus \{0\}$ at location $x \in \mathbb{R}^p$ is $\text{sign}(x^T \gamma)$ ². Note that scaling γ by a positive factor does not change this prediction. Consequentially, it suffices to study vectors on the sphere S^{p-1} . Grothendieck’s identity states that the probability that two vectors disagree about the label at a random position $x \sim \mathcal{N}(0, I_p)$ is proportional to their angle. It is originally due to (Grothendieck, 1956, p. 50), who used it to prove Grothendieck’s inequality (see e.g. Vershynin, 2018).

Theorem 3.1.1 *Let $a, b \in \mathbb{R}^p \setminus \{0\}$ and $x \sim \mathcal{N}(0, I_p)$. It holds that:*

$$\mathbb{P}[a^T x x^T b \leq 0] = \frac{1}{\pi} \arccos \left(\frac{a^T b}{\|a\|_2 \|b\|_2} \right).$$

We now know how the disagreement between two vectors $a, b \in S^{p-1}$ about the label at a point x relates to the angle between, the two vectors, i.e. their geodesic distance on S^{p-1} . Part of our analysis, later on, is of Euclidean nature, hence we would prefer to measure with the Euclidean distance. Fortunately, the geodesic and Euclidean distances are bilipschitz equivalent.

² We note that $\text{sign}(z)$ is not defined if $z = 0$. As the distribution of our covariates x are absolutely continuous with respect to the Lebesgue measure, these events occur with probability zero. Hence, we ignore this issue.

Proposition 3.1.1 For any $\beta, \beta' \in S^{p-1}$,

$$\|\beta - \beta'\|_2 \leq \arccos(\beta^T \beta') \leq \frac{\pi}{2} \|\beta - \beta'\|_2.$$

Using that $\sigma\epsilon$ is a centred Gaussian with variance σ^2 , the probability of observing a wrong label is:

$$\mathbb{P}[yx^T \beta^* < 0] = \frac{1}{\pi} \arccos\left(\frac{1}{\sqrt{1 + \sigma^2}}\right), \quad (4)$$

by Grothendieck's identity. One easily verifies that:

$$\frac{\sigma}{\pi(1 + \sigma^2)} \leq \mathbb{P}[yx^T \beta^* < 0] \leq \frac{\sigma}{\pi}. \quad (5)$$

3.2 Fisher Consistency and Misspecification Our loss function (2) minimizes an empirical risk. We will exploit techniques from M-estimation (see e.g. van de Geer, 2000) to prove that this empirical risk minimizer converges to the minimizer of the 'true risk', i.e.

$$\gamma^* := \arg \min_{\gamma \in \mathbb{R}^p} P \log(1 + \exp(-yx^T \gamma)).$$

Here and throughout, we use P to denote the expectation. If γ^* equals the parameter of interest β^*/σ , we say that our estimator is *Fisher consistent*.

Recall that our data comes from a probit model, a generalized linear model with the Gaussian link function. To estimate β^*/σ , we use the logistic link function, so we cannot exploit maximum likelihood theory to conclude that $\gamma^* = \beta^*/\sigma$. In fact, this turns out to be false. However, there is a one-to-one correspondence between τ^* and σ , so that finding τ^* is equivalent to finding σ . We start by establishing this relationship. Then, we show that the estimator is Fisher consistent in direction, i.e. β^* . Finally, we show that up to constant factors, $\|\gamma^*\|_2$ behaves like $1/\sigma$, aiding us in our calculations.

3.2.1 Fisher Consistency for the Signal-to-noise Ratio. Recall the notation $\tau^* := \|\gamma^*\|_2$. As hinted before, the signal-to-noise ratio $1/\sigma$ and the length of the target vector $\tau^* := \|\gamma^*\|_2$ are generally not equal. Yet, they are in a one-to-one correspondence. This follows from the next lemma.

Lemma 3.2.1 Let $\sigma > 0$ and $z \sim \mathcal{N}(0, 1)$. Then:

$$P \frac{|z|}{1 + \exp(\tau^* |z|)} = \frac{1}{\sqrt{2\pi}} \left(1 - \frac{1}{\sqrt{1 + \sigma^2}}\right).$$

Proof We recall:

$$\log(1 + \exp(-yx^T \gamma)) = \log(1 + \exp(-|x^T \gamma|)) + |x^T \gamma| \mathbb{1}\{yx^T \gamma < 0\}.$$

Let $R : (0, \infty) \rightarrow \mathbb{R}$, $\tau \mapsto P \log(1 + \exp(-yx^T(\tau\beta^*)))$. Differentiating R twice, we see that R is strictly convex. So, it reaches its minimum (if at all) at a root of its derivative. This gives us the following condition:

$$P \frac{|x^T \beta^*|}{1 + \exp(\tau^* |x^T \beta^*|)} = P |x^T \beta^*| \mathbb{1}\{yx^T \beta^* < 0\}.$$

Since $x^T \beta^* \sim \mathcal{N}(0, 1)$, the left-hand side is equal to the left-hand side in the conclusion. The result now follows from Corollary A.2.1. \square

Using Lemma 3.2.1, we can show that τ^* and $1/\sigma$ behave similarly, up to multiplicative constants.

Lemma 3.2.2 For $\sigma \in (0, 1/\sqrt{2}]$,

$$1 \leq \sigma \tau^* \leq \sqrt{2\pi}.$$

Proof We first prove the lower bound. By Lemma 3.2.1, if $\sigma \leq 1/\sqrt{2}$ then $\tau^* > \sqrt{6}$. Therefore, by Lemma A.1.2, Lemma 3.2.1 and since $1 - \frac{1}{\sqrt{1+\sigma^2}} \leq \frac{\sigma^2}{2}$,

$$\begin{aligned} \frac{1}{\sqrt{2\pi}} \frac{1}{2\tau^{*2}} &\leq \sqrt{\frac{1}{2\pi}} \frac{1}{\tau^{*2}} \left(1 - \frac{3}{\tau^{*2}}\right) \leq P \frac{|z|}{1 + \exp(\tau^* |z|)} \\ &= \frac{1}{\sqrt{2\pi}} \left(1 - \frac{1}{\sqrt{1+\sigma^2}}\right) \leq \frac{1}{\sqrt{2\pi}} \frac{\sigma^2}{2}. \end{aligned}$$

For the upper bound, we use Lemma A.1.2, Lemma 3.2.1 and that for $\sigma \leq 1/\sqrt{2}$, $\sigma^2/\pi \leq 1 - \frac{1}{\sqrt{1+\sigma^2}}$, which gives:

$$\sqrt{\frac{2}{\pi}} \frac{1}{\tau^{*2}} \geq P \frac{|z|}{1 + \exp(\tau^* |z|)} = \frac{1}{\sqrt{2\pi}} \left(1 - \frac{1}{\sqrt{1+\sigma^2}}\right) \geq \frac{1}{\sqrt{2\pi}} \frac{\sigma^2}{\pi}.$$

\square

3.2.2 Fisher Consistency of Orientation. Here, we show that the logistic regression estimator achieves Fisher consistency in our model for the direction, for any fixed length. The proof is readily generalized to centrally symmetric distributions, non-identity covariance matrices, and other monotone loss functions.

Lemma 3.2.3 Fix an element $\beta^* \in S^{p-1}$. Then,

$$\beta^* = \arg \min_{\beta \in S^{p-1}} Pl(yx^T \beta).$$

Proof Fix any $\beta \in S^{p-1}$ and define the corresponding label $\hat{y} := \text{sign}(x^T \beta)$, as well as $y' := \text{sign}(x^T \beta^* - \text{sign}(x^T \beta^*)\epsilon)$. Since ϵ is independent of x and ϵ has the same distribution as $-\epsilon$, y and y' have the same distribution. Moreover, $yx^T \beta$ and $y'x^T \beta$ have the same distribution. To simplify notation, define the auxiliary function $L(z) := l(-|z|) - l(|z|)$. We find:

$$\begin{aligned} Pl(yx^T \beta) &= Pl(y'x^T \beta) \\ &= Pl(|x^T \beta|) (1\{y' = y^*\}1\{y^* = \hat{y}\} + 1\{y' \neq y^*\}1\{y^* \neq \hat{y}\}) \\ &\quad + Pl(-|x^T \beta|) (1\{y' \neq y^*\}1\{y^* = \hat{y}\} + 1\{y' = y^*\}1\{y^* \neq \hat{y}\}) \\ &= Pl(|x^T \beta|) ((1 - 1\{y' \neq y^*\})(1 - 1\{y^* \neq \hat{y}\}) + 1\{y' \neq y^*\}1\{y^* \neq \hat{y}\}) \\ &\quad + Pl(-|x^T \beta|) (1\{y' \neq y^*\}(1 - 1\{y^* \neq \hat{y}\}) + (1 - 1\{y' \neq y^*\})1\{y^* \neq \hat{y}\}) \\ &= P1\{\hat{y} \neq y^*\}(1 - 2 \cdot 1\{y' \neq y^*\})L(x^T \beta) + Pl(|x^T \beta|) + PL(x^T \beta)1\{y' \neq y^*\} \\ &\stackrel{(i)}{=} P1\{\hat{y} \neq y^*\}(1 - 2\Phi(-|x^T \beta^*|))L(x^T \beta) + Pl(|x^T \beta|) + PL(x^T \beta)1\{|x^T \beta^*| \leq \epsilon\}. \end{aligned}$$

Equality (i) follows after conditioning on x , and recalling that $\{y' \neq y^*\} = \{|x^T \beta^*| \leq \epsilon\}$. Note that if $\beta = \beta^*$, this reads,

$$Pl(yx^T \beta^*) = Pl(|x^T \beta^*|) + PL(x^T \beta^*)1\{|x^T \beta^*| \leq \epsilon\}. \quad (6)$$

Since l is strictly decreasing, $L(x^T \beta) > 0$ almost surely. So, if instead $\beta^* \neq \beta$, using that by rotational invariance, $Pl(|x^T \beta|) = Pl(|x^T \beta^*|)$,

$$\begin{aligned} Pl(yx^T \beta) - Pl(|x^T \beta^*|) &> PL(x^T \beta)1\{|x^T \beta^*| \leq \epsilon\} \\ &= PL(x^T \beta)1\{|x^T \beta^*| \leq \epsilon\}1\{|x^T \beta| \leq \epsilon\} + L(x^T \beta)1\{|x^T \beta^*| \leq \epsilon\}1\{|x^T \beta| > \epsilon\} \\ &\stackrel{(ii)}{>} PL(x^T \beta^*)1\{|x^T \beta^*| \leq \epsilon\}1\{|x^T \beta| \leq \epsilon\} + L(x^T \beta^*)1\{|x^T \beta^*| \leq \epsilon\}1\{|x^T \beta| > \epsilon\} \\ &= PL(x^T \beta^*)1\{|x^T \beta^*| \leq \epsilon\}. \end{aligned}$$

Inequality (ii) follows since in the first summand, by rotational invariance we may exchange $x^T \beta$ and $x^T \beta^*$, while in the second one, $|x^T \beta| > |x^T \beta^*|$. By Eq. 6, we conclude that for any $\beta \neq \beta^*$, $Pl(yx^T \beta) > Pl(yx^T \beta^*)$. \square

*3.3 Introducing the *-Norm* In the proof for the large noise case, we establish an upper bound on the excess risk $P(l(\hat{\gamma}) - l(\gamma^*))$. In the end, however, we want to show that $\hat{\gamma}$ is close to γ^* . Ideally, we would want to lower-bound the excess risk with the (squared) Euclidean distance (this is sometimes referred to as the *margin condition*). This is not generally possible in our setting, as we argue below. Yet, we introduce a new metric, for which we can prove such a margin condition.

Before we introduce the metric we will use, let us consider an example that shows why the Euclidean distance fails. In a data set of realistic sample size, we see the correct labels with high probability, whether γ^* is equal to $10^{10}e_1$, or $10^{20}e_1$. So, both these vectors would give us the same observations with high probability. However, the Euclidean distance between these two vectors is of order 10^{20} . Therefore, attempting to explain the behavior of logistic regression with the Euclidean distance is hopeless, at least in the absence of bounds on the signal-to-noise ratio (SNR).

We will overcome this issue by separately comparing the vector's lengths and orientations. We will invoke a distance over $\mathbb{R}^p \setminus \{0\} \times \mathbb{R}^p \setminus \{0\}$, or alternatively a norm on \mathbb{R}^{p+1} :

$$\|\cdot\|_* : \mathbb{R}^{p+1} \rightarrow [0, \infty), \quad (\tau, \beta) \mapsto \sqrt{\frac{|\tau|^2}{\tau^{*3}} + \tau^* \|\beta\|_2^2}.$$

The reader will already anticipate from our notation that the first dimension is used to compare the length of two vectors γ, γ' , and the remaining p dimensions are used to compare orientations. In particular, we apply this norm as follows:

$$\|(\tau, \beta) - (\tau^*, \beta^*)\|_* = \sqrt{\frac{|\tau - \tau^*|^2}{\tau^{*3}} + \tau^* \|\beta - \beta^*\|_2^2}.$$

As we will primarily employ this norm to quantify the distance between γ and γ^* , we use the following, more compact notation:

$$d_* : \mathbb{R}^p \setminus \{0\} \rightarrow [0, \infty), \quad \gamma \mapsto \left\| \left(\|\gamma\|_2, \frac{\gamma}{\|\gamma\|_2} \right) - (\tau^*, \beta^*) \right\|_*.$$

If the SNR τ^* is large, more weight is given to the orientation. The latter can still be estimated in this case. On the other hand, if the SNR is small, it can be estimated, and so the $\|\cdot\|_*$ -norm gives more weight to the difference between the two lengths of the vectors.

Although this will not play a role in our analysis, we briefly comment in what sense $\|\cdot\|_*$ is a norm. As it is just a weighted Euclidean norm, it is itself a norm on \mathbb{R}^{p+1} . However, if understood as the function $\mathbb{R}^p \setminus \{0\} \rightarrow [0, \infty)$, mapping $\gamma \mapsto \|(\|\gamma\|_2, \gamma/\|\gamma\|_2)\|_*$, it is not a norm, as it is not absolutely homogeneous, and not defined at $\gamma = 0$. From this viewpoint, it is quickly verified that this function still is a metric for the space $\mathbb{R}^p \setminus \{0\}$.

In the following, we show that the $\|\cdot\|_*$ -norm and the Euclidean distance behave similarly, up to a constant depending on τ^* . We again use the notation $\tau := \|\gamma\|_2$ and $\beta := \gamma/\tau$ for an arbitrary vector $\gamma \in \mathbb{R}^p$.

Lemma 3.3.1 *If $\tau^* := \|\gamma^*\|_2 \geq 1$, then for all $\gamma \in \mathbb{R}^p \setminus \{0, \gamma^*\}$,*

$$\frac{\sqrt{\tau^*}}{3} \leq \frac{\|\gamma - \gamma^*\|_2}{d_*(\gamma)} \leq \sqrt{2\tau^*3}.$$

Proof We prove the upper bound. Using a triangular inequality, Jensen's inequality, and that $\tau^* \geq 1$, we find:

$$\begin{aligned} \|\gamma - \gamma^*\|_2 &\leq |\tau - \tau^*| + \tau^* \|\beta - \beta^*\|_2 \\ &\stackrel{(i)}{\leq} \sqrt{2} \sqrt{|\tau - \tau^*|^2 + \tau^* \|\beta - \beta^*\|_2^2} \leq \sqrt{2\tau^*3} \sqrt{\frac{|\tau - \tau^*|^2}{\tau^*3} + \tau^* \|\beta - \beta^*\|_2^2}. \end{aligned}$$

This proves the upper bound. For the lower bound, we combine two triangular inequalities, namely $\|\gamma - \gamma^*\|_2 \geq |\tau - \tau^*|$ and $\|\gamma - \gamma^*\|_2 \geq \tau^* \|\beta - \beta^*\|_2 - |\tau - \tau^*|$. Using that $\tau^* \geq 1$, we find:

$$3\|\gamma - \gamma^*\|_2 \geq \tau^* \|\beta - \beta^*\|_2 + |\tau^* - \tau| \geq \sqrt{\tau^*} \left(\sqrt{\tau^*} \|\beta - \beta^*\|_2 + \frac{|\tau^* - \tau|}{\tau^*3/2} \right).$$

The inequality $a + b \geq \sqrt{a^2 + b^2}$, which holds for any positive scalar $a, b > 0$, completes the proof. \square

4 Concentration for the Bounded Term

In this section, we control the empirical process of the bounded term. We define the bounded term as:

$$b : \mathbb{R}^p \times \mathbb{R}^p \rightarrow (0, \log 2], \quad (\gamma, x) \mapsto \log(1 + \exp(-|x^T \gamma|)).$$

We often omit the argument x , and write $b(\gamma)$, which is to be understood as the random variable $b(\gamma, x)$. Moreover, we are often interested in the

difference of the bounded term evaluated at γ with the same term evaluated at γ^* , which we give the notation $\tilde{b}(\gamma) := (b(\gamma) - b(\gamma^*)) / \log 2$.

4.1 Bousquet's Inequality We will use Bousquet's inequality to control the bounded term's empirical process. It was introduced in (Bousquet, 2002, Theorem 2.3).

Theorem 4.1.1 (Bousquet (2002)) *Let X_1, \dots, X_n be identically distributed random variables with values in \mathcal{X} . Let \mathcal{F} be a countable class of measurable mappings $\mathcal{X} \rightarrow \mathbb{R}$. Suppose that there exists a $\varsigma > 0$, such that for all $f \in \mathcal{F}$, $\text{Var}[f(X_1)] \leq \varsigma^2$ and $\|f\|_\infty \leq 1$. Define:*

$$Z := \sup_{f \in \mathcal{F}} |(P_n - P)f|.$$

For $t > 0$, with probability at most $\exp(-t)$,

$$Z \geq PZ + \sqrt{\frac{2t(\varsigma^2 + 2PZ)}{n}} + \frac{t}{3n}.$$

We will use two simplifications for this lower bound. Using the sub-additivity of the square root and the inequality between arithmetic and geometric means (AM-GM),

$$PZ + \sqrt{\frac{2t(\varsigma^2 + 2PZ)}{n}} + \frac{t}{3n} \leq 2PZ + \varsigma \sqrt{\frac{2t}{n}} + \frac{4t}{3n}.$$

On the other hand, for any $\lambda > 0$, by AM-GM,

$$PZ + \sqrt{\frac{2t(\varsigma^2 + 2PZ)}{n}} + \frac{t}{3n} \leq PZ(1 + \lambda^{-1}) + \frac{\varsigma^2}{2\lambda} + \frac{t(\lambda + 1/3)}{n}.$$

The reader will already anticipate that we will choose \mathcal{F} as a subset of all functions $\tilde{b}(\gamma) := (b(\gamma) - b(\gamma^*)) / \log 2$, which we denote by³:

$$\mathcal{B} := \left\{ \tilde{b}(\gamma) : \gamma \in \mathbb{R}^p \right\}.$$

In particular, due to our localization argument, we will only need to control the terms \tilde{b} indexed by γ , which are somewhat close to γ^* . Moreover, recall that in the end, our goal is to prove that some γ are close to γ^* .

³ Bousquet's inequality requires that the set \mathcal{F} is countable to avoid measurability issues. We will omit this technical discussion.

Therefore, we do not need to control terms indexed by γ which are very close to γ^* , as for such γ there is nothing left to prove from the beginning. This will allow us to apply the peeling technique, see e.g. van de Geer (2000).

4.1.1 Expectation of the Supremum of the Empirical Process. Here, we bound the term PZ , as introduced in Bousquet's inequality. In the proof, we will take a covering with respect to the $L^1(P_n)$ -norm. Since P_n is a random measure, we will bound the covering number with the largest $L^1(Q)$ -covering number, where Q is any probability measure on \mathbb{R}^p . This is possible with the help of Vapnik-Chervonenkis (VC) theory. We denote the set of all probability measures over \mathbb{R}^p by Π , and a minimal covering of a set $S \subset \mathcal{B}$ by balls of radius ϵ in the $L^1(Q)$ -norm by $\mathcal{N}(\epsilon, S, L^1(Q))$.

Lemma 4.1.1 *For any $\epsilon > 0$, and any probability measure Q on \mathbb{R}^p ,*

$$|\mathcal{N}(\epsilon, \mathcal{B}, L^1(Q))| \leq 2e(p + 2.5) \left(\frac{4e}{\epsilon}\right)^{2(p+2)}.$$

This result follows from Vapnik-Chervonenkis theory. If ν is the VC-index of the hypographs of \mathcal{B} , then by (Haussler, 1995, Corollary 3), for any $\epsilon > 0$,

$$|\mathcal{N}(\epsilon, \mathcal{B}, L^1(Q))| \leq e\nu \left(\frac{4e}{\epsilon}\right)^{\nu-1}.$$

Using elementary properties of the VC-index (van der Vaart and Wellner, 1996, Lemma 2.6.17-18), we can bound ν from above with $2(p + 2) + 1$.

We now turn to the bound for the expectation of the supremum of the empirical process of the bounded term $\tilde{b}(\gamma) := (b(\gamma) - b(\gamma^*)) / \log(2)$. We recall that this is a stepping stone for Bousquet's inequality, which gives us a tail-bound for the supremum of the empirical process. In the following, for a set $S \subset \mathcal{B}$, by $\tilde{b}^{-1}(S) := \{\gamma \in \mathbb{R}^p : \tilde{b}(\gamma) \in S\}$ we mean the pre-image of S through the mapping $\mathbb{R}^p \rightarrow \mathcal{B}$, $\gamma \mapsto \tilde{b}(\gamma)$.

For any subset $S \subset \mathcal{B}$, and $\epsilon > 0$, we define:

$$N_r(S) := \sup_{Q \in \Pi} |\mathcal{N}(r, S, L^1(Q))|, \quad \varsigma_S := \sup_{\gamma \in \tilde{b}^{-1}(S)} \sqrt{P\tilde{b}(\gamma)^2}.$$

Proposition 4.1.1 *Let $S \subset \mathcal{B}$. Then,*

$$\mathbb{E} \sup_{\gamma \in \tilde{b}^{-1}(S)} |(P_n - P)\tilde{b}(\gamma)| \leq 2r + \frac{16 \log(2N_r(S))}{n} + 8\varsigma_S \sqrt{\frac{\log 2N_r(S)}{n}}.$$

This proposition presupposes that we have good control over the covering number $N_r(S)$ and the wimpy variance ς_S . The covering number can be controlled with Lemma 4.1.1, and the wimpy variance will later be bounded using Lemma 4.1.2.

Proof The proof uses standard steps from empirical process theory. The first step is symmetrization (see e.g. van der Vaart and Wellner, 1996, Lemma 2.3.1). Let $\varepsilon_1, \dots, \varepsilon_n$ be Rademacher, independent of each other and independent of $(x_1, y_1), \dots, (x_n, y_n)$. Then,

$$\mathbb{E} \left(\sup_{\gamma \in \tilde{b}^{-1}(S)} |(P_n - P)\tilde{b}(\gamma)| \right) \leq 2\mathbb{E} \left(\sup_{\gamma \in \tilde{b}^{-1}(S)} |P_n \varepsilon \tilde{b}(\gamma)| \right),$$

where $P_n \varepsilon \tilde{b}(\gamma) := n^{-1} \sum_{i=1}^n \varepsilon_i \tilde{b}(\gamma, x_i)$. Next, we discretize. There exists a covering $\mathcal{N} := \mathcal{N}(r, S, L^1(P_n))$ with cardinality $N \leq N_r(S)$. So,

$$\begin{aligned} 2\mathbb{E} \left(\sup_{\gamma \in \tilde{b}^{-1}(S)} |P_n \varepsilon \tilde{b}(\gamma)| \right) &= 2\mathbb{E} \left(\sup_{\gamma \in \tilde{b}^{-1}(S)} \min_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon \tilde{b}(\gamma) - P_n \varepsilon \tilde{b}' + P_n \varepsilon \tilde{b}'| \right) \\ &\leq 2\mathbb{E} \left(\sup_{\gamma \in \tilde{b}^{-1}(S)} \min_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon (\tilde{b}(\gamma) - \tilde{b}')| + \max_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon \tilde{b}'| \right) \leq 2r + 2\mathbb{E} \max_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon \tilde{b}'|. \end{aligned}$$

The next step is concentration. Let \mathbb{P}_ε be the conditional expectation with respect to $\varepsilon_1, \dots, \varepsilon_n$, i.e. conditional on x_1, \dots, x_n . By (Dümbgen et al. 2010, Lemma 3.4),

$$\mathbb{P}_\varepsilon \max_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon \tilde{b}'| \leq \sqrt{\frac{2 \log 2N}{n}} \sqrt{\max_{\tilde{b}' \in \mathcal{N}} P_n \tilde{b}'^2}.$$

Adding and subtracting the expectation, and by sub-additivity of the square root,

$$\begin{aligned} \sqrt{\frac{2 \log 2N}{n}} \sqrt{\max_{\tilde{b}' \in \mathcal{N}} P_n \tilde{b}'^2} &\leq \sqrt{\frac{2 \log 2N}{n}} \left(\sqrt{\max_{\tilde{b}' \in \mathcal{N}} |(P_n - P)\tilde{b}'^2|} + \sqrt{\max_{\tilde{b}' \in \mathcal{N}} P \tilde{b}'^2} \right) \\ &\leq \frac{4 \log 2N}{n} + \frac{1}{8} \max_{\tilde{b}' \in \mathcal{N}} |(P_n - P)\tilde{b}'^2| + \varsigma_S \sqrt{\frac{4 \log 2N}{n}}. \end{aligned}$$

In the last inequality, we used that the second summand is bounded by ς_S and that by the inequality of arithmetic and geometric means, for all

$x, x' \geq 0, \sqrt{xx'} \leq x/8 + 2x'$.

Finally, we take the expectation and rearrange the expressions. Using that N is bounded by the deterministic $N_r(S)$, we conclude that:

$$\begin{aligned} \mathbb{E} \max_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon \tilde{b}'| &\leq \frac{4 \log 2N_r(S)}{n} + \frac{1}{8} \mathbb{E} \max_{\tilde{b}' \in \mathcal{N}} |(P_n - P) \tilde{b}'^2| + \varsigma_S \sqrt{\frac{4 \log 2N_r(S)}{n}} \\ &\stackrel{(i)}{\leq} \frac{4 \log 2N_r(S)}{n} + \frac{1}{8} 4 \mathbb{E} \max_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon \tilde{b}'| + \varsigma_S \sqrt{\frac{4 \log 2N_r(S)}{n}}. \end{aligned}$$

Inequality (i) follows from symmetrization (see e.g. van der Vaart and Wellner, 1996, Lemma 2.3.1), as well as the contraction principle (see e.g. Boucheron et al., 2013, Theorem 11.6), noting that the mapping $z \mapsto z^2/2$ is a contraction on $[-1, 1]$. We subtract $\mathbb{E} \max_{\tilde{b}' \in \mathcal{N}} |P_n \varepsilon \tilde{b}'|/2$ from both sides and conclude the proof. \square

4.1.2 Bounding the Wimpy Variance. Here, we bound the supremum of the variances of the bounded terms. We call this the *wimpy variance*, in line with the terminology used in Boucheron et al. (2013). It turns out, that this expression is smaller, the closer the γ are to γ^* , which pairs very well with our localization. The first result will be for the large-noise case and the second for the small-noise case.

Large noise In the large noise case, we measure the proximity of γ to γ^* using the $*$ -norm, respectively $d_*(\gamma)$, as introduced in Section 3.3. We use that $\tau^* \gtrsim 1$ and $\tau \gtrsim \tau^*$. This is possible due to the localization step in the proof of Theorem 6.1.1, which gives $\|\gamma - \gamma^*\|_2 \lesssim \tau^*$. If it were that $\tau, \tau^* \lesssim 1$, then the term $P(b(\tau\beta) - b(\tau^*\beta))^2$ would behave differently, see Lemma A.1.7.

Lemma 4.1.2 Fix $\gamma \in \mathbb{R}^p \setminus \{0\}$ and define $\tau := \|\gamma\|_2$ and $\beta := \gamma/\tau$. If $\tau^* \geq 6/5$ and $\tau \geq \tau^* 5/6$ then:

$$\sqrt{P \tilde{b}(\gamma)^2} \leq 3.57 d_*(\gamma).$$

Proof We start with a triangular inequality:

$$\sqrt{P(b(\gamma) - b(\gamma^*))^2} \leq \sqrt{P(b(\tau\beta) - b(\tau^*\beta))^2} + \sqrt{P(b(\tau^*\beta) - b(\tau^*\beta^*))^2}.$$

FINITE SAMPLE RATES FOR LOGISTIC REGRESSION...

We bound the first term with Lemma A.1.7. If $\tau^* \geq \tau$, using that $\tau \geq \tau^*5/6$,

$$\begin{aligned} \sqrt{P(b(\tau\beta) - b(\tau^*\beta))^2} &\leq \left(\frac{1}{2^5\pi}\right)^{1/4} \frac{|\tau - \tau^*|}{\tau} \sqrt{\frac{2}{\tau} + \frac{3}{4\tau^3}} \\ &\leq \left(\frac{1}{2^5\pi}\right)^{1/4} \frac{|\tau - \tau^*|}{\tau^{3/2}} \sqrt{2 + \frac{3}{4}} \leq \sqrt{\frac{11}{4}} \sqrt{\frac{1}{2^5\pi}} \left(\frac{6}{5}\right)^{3/2} \frac{|\tau - \tau^*|}{\tau^{*3/2}}. \end{aligned}$$

If instead $\tau \geq \tau^*$, we directly get:

$$\sqrt{P(b(\tau\beta) - b(\tau^*\beta))^2} \leq \sqrt{\frac{11}{4}} \sqrt{\frac{1}{2^5\pi}} \frac{|\tau - \tau^*|}{\tau^{*3/2}}.$$

The second term is bounded with Lemma A.1.8, again using $\tau^* \geq 1$. Then,

$$\sqrt{P(b(\tau^*\beta) - b(\tau^*\beta^*))^2} \leq \sqrt{2} \left(\frac{1}{(2\pi)^{1/4}} + \sqrt{\frac{11}{8}} \sqrt{\frac{2}{\pi}} \right) \sqrt{\tau^*} \|\beta - \beta^*\|_2.$$

Since:

$$\sqrt{\frac{11}{4} \sqrt{\frac{1}{2^5\pi}} \left(\frac{6}{5}\right)^3 + 2 \left(\frac{1}{(2\pi)^{1/4}} + \sqrt{\frac{11}{8}} \sqrt{\frac{2}{\pi}} \right)^2} \leq 3.57 \log(2),$$

it follows from Cauchy-Schwarz that:

$$\sqrt{P(b(\tau\beta) - b(\tau^*\beta))^2} + \sqrt{P(b(\tau^*\beta) - b(\tau^*\beta^*))^2} \leq 3.57 \log(2) d_*(\gamma).$$

This completes the proof. \square

Small noise Controlling the wimpy variance in the case where the noise is small is analogous to the large noise case. Thanks to the case distinction in the proof of Theorem 6.2.1, we will only need to compare $\tau\beta$ with $k\tau\beta^*$, where $k \geq 1$ is some constant factor and $\tau \geq 1$ is arbitrary. For $\delta > 0$, we define the spherical cap of radius δ centered at β^* as $S(\delta, \beta^*) := \{\beta \in S^{p-1} : \|\beta - \beta^*\|_2 \leq \delta\}$.

Lemma 4.1.3 *Let $\tau \geq 1$, $k \geq 1$ and $\delta > 0$. Then:*

$$\sup_{\beta \in S^{p-1} \setminus S(\delta, \beta^*)} \sqrt{P\left(\frac{b(\tau\beta) - b(k\tau\beta^*)}{\log(2)}\right)^2} \leq 5.37\delta\sqrt{\tau} + 0.76\frac{(k-1)}{\sqrt{\tau}}.$$

Proof To bound the wimpy variance, we use that by the triangular inequality, for any $\beta \in S^{p-1}$,

$$\sqrt{P(b(\tau\beta) - b(k\tau\beta^*))^2} \leq \sqrt{P(b(\tau\beta) - b(\tau\beta^*))^2} + \sqrt{P(b(\tau\beta^*) - b(k\tau\beta^*))^2}.$$

By Lemma A.1.8, since $\tau \geq 1$ and $\sqrt{2}((2\pi)^{1/4} + \sqrt{11/8\sqrt{2/\pi}}) \leq \log(2)5.37$, we have:

$$\sqrt{P(b(\tau\beta) - b(\tau\beta^*))^2} \leq 3.72\sqrt{\tau}\|\beta - \beta^*\|_2.$$

Next, by Lemma A.1.7, since $\tau \geq 1$, the second term is bounded by:

$$P(b(\tau\beta^*) - b(k\tau\beta^*))^2 \leq \frac{11}{4} \sqrt{\frac{1}{2^5\pi}} \frac{(k-1)^2}{\tau} \leq 0.76 \log(2) \frac{(k-1)^2}{\tau}.$$

The proof is complete. □

4.2 The Empirical Process of the Bounded Part

4.2.1 The Large Noise Case. We consolidate our results so far to give local control of the empirical process of the bounded part b with Bousquet's inequality (Theorem 4.1.1). We prove a local result, in a $\|\cdot\|_*$ -ball of radius $R > 0$ around the ground truth (τ^*, β^*) . This allows us to get a faster rate: if $d_*(\gamma)$ is small, so is the wimpy variance (see Lemma 4.1.2), and consequentially also the expectation of the supremum of the empirical process (see Proposition 4.1.1). For the peeling argument to work, we remove a small ball of radius $r > 0$ from the supremum. This does not affect our results in the end, since if $d_*(\gamma) \lesssim r$ for r a small enough, the desired result already holds. Moreover, as we controlled the wimpy variance for $\tau \geq \tau^*5/6$, we exclude all γ with $\|\gamma\|_2 < \tau^*5/6$ from the supremum.

In the following, B_δ denotes the Euclidean ball of radius δ , and for the $*$ -norm, we define the ball $B_\delta^* := \{\gamma \in \mathbb{R}^p \setminus \{0\} : d_*(\gamma) \leq \delta\}$. Moreover, for $r, R > 0$ we define the 'peels' $B_{r,R}^* := B_R^* \setminus B_r^*$.

Proposition 4.2.1 *Let $R > 0$ and $r \in (0, 1)$, and suppose that $\tau^* \geq 6/5$. Let $t' := t + \log\lceil\log R/r\rceil$. With probability at least $1 - \exp(-t)$,*

$$\sup_{\gamma \in B_{r,R}^* \setminus B_{\tau^*5/6}} \frac{|(P_n - P)\tilde{b}(\gamma)| - 4r - \frac{32(3p+4)\log(4e/r) + \frac{4}{3}t'}{n}}{d_*(\gamma)e16 \cdot 3.57\sqrt{\frac{2(3p+4)\log(4e/r) + t'}{n}}} \leq 1.$$

Proof We prepare the use of Bousquet's inequality (Theorem 4.1.1). Since $r \leq 1$, by Lemma 4.1.1:

$$N := |\mathcal{N}(r, \mathcal{B}, L^1(Q))| \leq 2e(p + 2.5)(4e/r)^{2(p+2)} \leq \frac{1}{2}(4e/r)^{3p+4}.$$

Moreover, by Lemma 4.1.2,

$$\varsigma := \sup_{\gamma \in B_\delta^* \setminus B_{\tau^*5/6}} \sqrt{P\tilde{b}(\gamma)^2} \leq 3.57\delta.$$

By Proposition 4.1.1,

$$\begin{aligned} \mathbb{E} \sup_{\gamma \in B_\delta^* \setminus B_{\tau^*5/6}} |(P_n - P)\tilde{b}(\gamma)| &\leq 2r + \frac{16\log(2N)}{n} + 8\varsigma\sqrt{\frac{\log(2N)}{n}} \\ &\leq 2r + \frac{16(3p+4)\log(4e/r)}{n} + 8 \cdot 3.57\delta\sqrt{\frac{(3p+4)\log(4e/r)}{n}}. \end{aligned}$$

We now use Bousquet's inequality (Theorem 4.1.1). On an event with probability at least $1 - \exp(-t)$,

$$\begin{aligned} &\sup_{\gamma \in B_\delta^* \setminus B_{\tau^*5/6}} |(P_n - P)\tilde{b}(\gamma)| \\ &\leq 4r + \frac{32(3p+4)\log(4e/r) + \frac{4}{3}t}{n} + 16 \cdot 3.57\delta\sqrt{\frac{2(3p+4)\log(4e/r) + t}{n}} \\ &=: \bar{\psi}_b(t) + \underline{\psi}_b(\delta, t). \end{aligned}$$

Equivalently,

$$\sup_{\gamma \in B_\delta^* \setminus B_{\tau^*5/6}} \frac{|(P_n - P)\tilde{b}(\gamma)| - \bar{\psi}_b(t)}{\underline{\psi}_b(\delta, t)} \leq 1. \quad (7)$$

We now apply the peeling technique (see e.g. van de Geer (2000)), to replace δ by $d_*(\gamma)$. We partition $B_{r,R}^*$ into the disjoint union $B_{r,R}^* = \cup_{j=0}^{J-1} B_{re^j, re^{j+1}}^*$, with $J = \lceil \log R/r \rceil$. Then,

$$\begin{aligned} & \mathbb{P} \left[\bigcup_{\gamma \in B_{r,R}^* \setminus B_{\tau^{*5/6}}} \frac{|(P_n - P)\tilde{b}(\gamma)| - \bar{\psi}_b(t + \log J)}{\underline{\psi}_b(ed_*(\gamma), t + \log J)} > 1 \right] \\ & \leq \sum_{j=0}^{J-1} \mathbb{P} \left[\bigcup_{\gamma \in B_{re^j, re^{j+1}}^* \setminus B_{\tau^{*5/6}}} \frac{|(P_n - P)\tilde{b}(\gamma)| - \bar{\psi}_b(t + \log J)}{\underline{\psi}_b(ed_*(\gamma), t + \log J)} > 1 \right] \\ & \leq \sum_{j=0}^{J-1} \mathbb{P} \left[\bigcup_{\gamma \in B_{re^j, re^{j+1}}^* \setminus B_{\tau^{*5/6}}} \frac{|(P_n - P)\tilde{b}(\gamma)| - \bar{\psi}_b(t + \log J)}{\underline{\psi}_b(re^{j+1}, t + \log J)} > 1 \right] \\ & \stackrel{(ii)}{\leq} J \exp(-t - \log J) = \exp(-t). \end{aligned}$$

The inequality (ii) follows from Eq. 7. The proof is complete. \square

4.2.2 Small Noise Case. Here, we prove two bounds for the empirical process of the bounded term in the case of small noise, based on Bousquet's inequality. We will prove a tail bound on:

$$\sup_{\beta \in S^{p-1} \setminus S(\delta, \beta^*)} |(P - P_n)(b(\tau\beta) - b(k\tau\beta^*))|.$$

We will use different k in the proof of the main theorem. For part 1 of the proof ($M/2 \leq \hat{\tau} \leq M$), we will use $k = 1$, for part 2 ($1 \leq \hat{\tau} \leq M/2$), we will use $k = 2$. In the following, we use the notation $\tilde{b}_k(\tau\beta) := (b(\tau\beta) - b(k\tau\beta^*)) / \log 2$.

Proposition 4.2.2 *Fix $k, \tau, \lambda \geq 1$ as well as $r \in (0, 1)$ and $t > 0$ and define $t' := t + \log \lceil \log 2/r \rceil$. With probability at least $1 - \exp(-t)$,*

$$\begin{aligned} & \sup_{\beta \in S^{p-1} \setminus S(r, \beta^*)} |(P - P_n)\tilde{b}_k(\tau\beta)| \\ & \leq 13851 \frac{\|\beta - \beta^*\|_2^2 \tau}{\lambda} + 4r + 30 \frac{(k-1)^2}{\lambda \tau} + \frac{\lambda 34(3p+4) \log(4e/r) + t'(\lambda + \frac{1}{3})}{n}. \end{aligned}$$

FINITE SAMPLE RATES FOR LOGISTIC REGRESSION...

Proof From Lemma 4.1.1, it follows that to cover the set of functions $\{(b(\gamma) - b(k\tau\beta^*)) / \log(2) : \gamma \in \mathbb{R}^p\}$ with $L^1(Q)$ -balls of radius r for any probability measure Q , one needs at most N balls, where $\log 2N \leq (3p + 4) \log(4e/r)$, arguing as in the proof of Proposition 4.2.1. With Lemma 4.1.3, the wimpy variance can be bounded as follows:

$$\varsigma := \sup_{\beta \in S^{p-1} \setminus S(\delta, \beta^*)} \sqrt{P\tilde{b}_k(\tau\beta)^2} \leq 5.37\delta\sqrt{\tau} + 0.76\frac{(k-1)}{\sqrt{\tau}}.$$

By Proposition 4.1.1,

$$\begin{aligned} & \mathbb{E} \sup_{\beta \in S(\delta, \beta^*) \setminus S(\delta, \beta^*)} \left| (P - P_n)\tilde{b}_k(\tau\beta) \right| \\ & \leq 2r + \frac{16 \log(2N)}{n} + 8\frac{\varsigma}{\sqrt{\lambda}} \sqrt{\lambda \frac{\log 2N}{n}} \\ & \leq 2r + 16\frac{\varsigma^2}{\lambda} + \frac{\lambda 17 \log(2N)}{n} \leq 2r + 16\frac{\varsigma^2}{\lambda} + \frac{\lambda 17(3p+4) \log(4e/r)}{n}. \end{aligned}$$

By Bousquet's inequality (Theorem 4.1.1), with probability at least $1 - \exp(-t)$,

$$\begin{aligned} & \sup_{\beta \in S^{p-1} \setminus S(\delta, \beta^*)} \left| (P - P_n)\tilde{b}_k(\tau\beta) \right| \leq 2PZ + \frac{\varsigma^2}{2\lambda} + \frac{t(\lambda + 1/3)}{n} \\ & \leq 4r + 32.5\frac{\varsigma^2}{\lambda} + \frac{\lambda 34(3p+4) \log(4e/r) + t(\lambda + 1/3)}{n} \\ & \leq \frac{65 \cdot 5.37^2}{\lambda} \delta^2 \tau + 4r + \frac{65 \cdot 0.76^2 (k-1)^2}{\lambda \tau} + \frac{\lambda 34(3p+4) \log(4e/r) + t(\lambda + 1/3)}{n} \\ & =: \delta^2 \underline{\psi}'_b + \overline{\psi}'_b(t). \end{aligned}$$

In other words,

$$\mathbb{P} \left[\sup_{\beta \in S^{p-1} \setminus S(\delta, \beta^*)} \frac{|(P - P_n)\tilde{b}_k(\tau\beta)| - \overline{\psi}'_b(t)}{\delta^2 \underline{\psi}'_b} > 1 \right] \leq \exp(-t). \quad (8)$$

Now, we use the peeling device. For $a, b > 0$, define: $S_{a,b}^* := \{\beta \in S^{p-1} : a \leq \|\beta - \beta^*\|_2 \leq b\}$. Note that for any $\beta \in S^{p-1}$, $\|\beta - \beta^*\|_2 \leq 2$. So, letting

$J := \lceil \log 2/r \rceil$, we find: Consequentially, for $r > 0$,

$$S^{p-1} \setminus S(r, \beta^*) = \bigcup_{j=0}^{J-1} S_{r e^j, r e^{j+1}}^*.$$

We want to upper bound:

$$\begin{aligned} & \mathbb{P} \left[\sup_{\beta \in S^{p-1} \setminus S(r, \beta^*)} \frac{|(P - P_n) \tilde{b}_k(\tau\beta)| - \bar{\psi}'_b(t + \log J)}{e^2 \|\beta - \beta^*\|_2^2 \underline{\psi}'_b} > 1 \right] \\ & \leq \sum_{j=0}^{J-1} \mathbb{P} \left[\sup_{\beta \in S_{r e^j, r e^{j+1}}^*} \frac{|(P - P_n) \tilde{b}_k(\tau\beta)| - \bar{\psi}'_b(t + \log J)}{e^2 \|\beta - \beta^*\|_2^2 \underline{\psi}'_b} > 1 \right] \\ & \leq \sum_{j=0}^{J-1} \mathbb{P} \left[\sup_{\beta \in S_{0, r e^{j+1}}^*} \frac{|(P - P_n) \tilde{b}_k(\tau\beta)| - \bar{\psi}'_b(t + \log J)}{e^2 r^2 e^{2j} \underline{\psi}'_b} > 1 \right]. \\ & \leq J \exp(-t - \log J) = \exp(-t). \end{aligned}$$

The last inequality follows from Eq. 8. In the final result, we use that $65 \cdot 5.37^2 e^2 \leq 13851$, as well as the approximation $65 \cdot 0.76^2 \leq 30$. The proof is complete. \square

5 Concentration for the Unbounded Term

Here, we give two concentration inequalities for the empirical process of the unbounded term:

$$u : \mathbb{R}^p \setminus \{0\} \times \mathbb{R}^p \times \{-1, 1\} \rightarrow \mathbb{R}, \quad (\gamma, x, y) \mapsto |x^T \gamma| \mathbb{1}\{yx^T \gamma < 0\}.$$

We will omit the dependence of u on x, y , if it is clear from context, and write $u(\gamma)$ instead of $u(\gamma, x, y)$. Moreover, we are frequently interested in the difference between the unbounded term evaluated at some point γ and the parameter γ^* , for which we use the notation:

$$\tilde{u}(\gamma) := u(\gamma) - u(\gamma^*).$$

The first is for the “large noise”-regime and the second is for the “small noise”-regime. In the former, we can exploit that due to localization, $\|\gamma - \gamma^*\|_2 \lesssim \tau^* \sim 1/\sigma$, which is not too large. In the small noise case, this is not helpful, as $1/\sigma$ may be arbitrarily large. However, since we do not have the

ambition to estimate τ^* in that case, we only need to compare β and β^* . Moreover, we can exploit that σ is vanishingly small.

Both results are based on Bernstein's inequality. In both cases, we proceed as follows: First, we show that Bernstein's condition is satisfied, in a small region around the true parameter. Second, we use a covering argument to control the empirical process. Third, we apply peeling, to obtain a fast global rate, exploiting the good control locally around the true parameter.

5.1 Bernstein's Inequality & Covering Although we need two separate results for the large and the small-noise cases, there are some results we can use in both cases. One is Bernstein's inequality itself, and the second is the covering argument.

5.1.1 Bernstein's Inequality. Bernstein's inequality is a classical concentration inequality for the tail of sub-exponential random variables. It states that under the "Bernstein condition", the average of centered random variables concentrates around zero. Bernstein's condition dictates that the moments of the random variable may not grow too fast. The parameters in Bernstein's condition may reflect, that the moments of the random variables are small. For example, if the second moment is small, κ can be chosen as small (simply take $m = 2$ to see this). We will exploit this, to obtain a fast convergence rate.

Definition 5.1.1 Let $\{z_i\}_{i \in [n]}$ be a sequence of random variables. We say that $\{z_i\}_{i \in [n]}$ satisfies Bernstein's condition with constants $\{\kappa_i\}_{i=1}^n$ and K , if for all $i \in [n]$, $Pz_i = 0$, and for all $m \in \{2, 3, \dots\}$,

$$P|z_i|^m \leq \frac{m!}{2} K^{m-2} \kappa_i^2.$$

Bernstein's inequality is formulated for centered random variables. However, it is sometimes easier to verify Bernstein's condition for a sequence of random variables $\{z_i\}_{i \in [n]}$, instead of their centered version $\{z_i - Pz_i\}_{i \in [n]}$. In the following remark, we show that this suffices, at the cost of a constant factor.

Remark 5.1.1 Note that:

$$\begin{aligned} P(2|z_i|)^m &= 2^{m-1} P(|z_i|^m + P|z_i|^m) \\ &\stackrel{(i)}{\geq} 2^{m-1} P(|z_i|^m + |Pz_i|^m) \stackrel{(ii)}{\geq} P(|z_i| + |-Pz_i|)^m \stackrel{(iii)}{\geq} P|z_i - Pz_i|^m. \end{aligned}$$

Here, (i) and (ii) follow from Jensen's inequality, while (iii) follows from the triangular inequality. So, Bernstein's condition is met for $z_i - Pz_i$ with

constants $\{\kappa_i\}_{i=1}^n$ and K , if for all $m \in \{2, 3, \dots\}$:

$$P|z_i|^m \leq \frac{m!}{2} \left(\frac{K}{2}\right)^{m-2} \left(\frac{\kappa_i}{2}\right)^2.$$

We will need a Bernstein-type inequality for a supremum of an empirical process. To this end, we use a covering argument and a finite union bound.

Lemma 5.1.1 *Let $L \in \mathbb{N}$. For $l \in [L]$, let $\{z_{i,l}\}_{i \in [n]}$ be a sequence of random variables with $Pz_{i,l} = 0$, satisfying Bernstein's condition with $K_l \leq K$ and $\kappa_{i,l} \leq \kappa$. Let $t > 0$. With probability at least $1 - 2 \exp(-t)$:*

$$\max_{l \in [L]} |P_n z_l| \leq \kappa \sqrt{\frac{2(t + \log L)}{n}} + K \frac{t + \log L}{n}.$$

Proof By Bernstein's inequality (Boucheron et al., 2013, Theorem 2.10), for all $l \in [L]$, with probability at most $2 \exp(-t)$,

$$|P_n z_l| \geq \kappa \sqrt{\frac{2t}{n}} + K \frac{t}{n}.$$

In particular, with probability at most $2 \exp(-t - \log L)$,

$$|P_n z_l| \geq \kappa \sqrt{\frac{2(t + \log L)}{n}} + K \frac{t + \log L}{n} =: \psi.$$

Using a union bound,

$$P \left[\max_{l \in [L]} |P_n z_l| \geq \psi \right] \leq 2L \exp(-t - \log L) = 2 \exp(-t).$$

The proof is complete. □

5.1.2 Covering. To apply Bernstein's inequality (Lemma 5.1.1), which is for a maximum over a finite set, we use a covering argument. For the unbounded terms u , we introduce an integrable envelope U , which dominates them pointwise. Such an envelope can be found with the Cauchy-Schwarz inequality. However, an upper bound on $\|\gamma\|_2$ will be needed, so we will assume that for some $r > 0$, it holds that $\|\gamma\|_2 \leq r$. This is not a problem, as we can assume $\|\gamma - \gamma^*\|_2 \lesssim \tau^*$ in the large noise case, and use $r = 1$ in the small noise case. We define:

$$U : \mathbb{R}^p \times (0, \infty) \rightarrow \mathbb{R}, \quad (x, r) \mapsto r \|x\|_2.$$

For any subset $S \subset \mathbb{R}^p \setminus \{0\}$, we define:

$$\mathcal{U}(S) := \{u(\gamma) : \gamma \in S\}, \quad \mathcal{U}(S)/U_r := \left\{ \frac{u(\gamma)}{U(\cdot, r)} : \gamma \in S \right\}.$$

Note that for any $\gamma, \gamma' \in \mathbb{R}^p$, it holds that $u(\gamma) - u(\gamma') = \tilde{u}(\gamma) - \tilde{u}(\gamma')$. So, a covering for u is equivalent to a covering for \tilde{u} .

With Lemma 5.1.2, we show that a covering of $\mathcal{U}(S)/U_r$ with the supremum norm does not impair the rate of our empirical process significantly. It then remains to calculate the covering number. After a technical Lemma 5.1.3, we see that u/U can be upper bounded by a quantity resembling the $\|\cdot\|_*$ -norm in Lemma 5.1.4. From there, we can derive a covering number using an upper bound for the covering number of the Euclidean unit ball in \mathbb{R}^{p+1} .

Lemma 5.1.2 *Fix $r > 0$. Let $S \subset \mathbb{R}^p \setminus \{0\}$ and $\mathcal{N} \subset S$ be sets such that $\mathcal{U}(\mathcal{N})/U_r$ is an ε -covering of $\mathcal{U}(S)/U_r$ in the supremum-norm. With probability at least $1 - \exp(-t)$,*

$$\sup_{\gamma \in S} |(P_n - P)\tilde{u}(\gamma)| \leq \sup_{\gamma \in \mathcal{N}} |(P_n - P)\tilde{u}(\gamma)| + \varepsilon r \left(\sqrt{\frac{2t}{n}} + 2\sqrt{p} \right).$$

Proof Fix any $\gamma \in S$ and an ε -neighbour $\gamma_\varepsilon \in \mathcal{N}$. By a triangular inequality,

$$|(P_n - P)\tilde{u}(\gamma)| \leq |(P_n - P)(\tilde{u}(\gamma) - \tilde{u}(\gamma_\varepsilon))| + |(P_n - P)\tilde{u}(\gamma_\varepsilon)|.$$

All that is left is to upper bound the first summand on the right-hand side. Note that $\tilde{u}(\gamma) - \tilde{u}(\gamma_\varepsilon) = u(\gamma) - u(\gamma_\varepsilon)$. Using another triangular inequality,

$$\begin{aligned} |(P_n - P)(u(\gamma) - u(\gamma_\varepsilon))| &\leq (P_n + P)U \left| \frac{u(\gamma) - u(\gamma_\varepsilon)}{U} \right| \\ &\leq \varepsilon(P_n + P)U = \varepsilon((P_n - P)U + 2PU). \end{aligned}$$

Now, we use a concentration result. The random variable $U(\cdot, r)/r$ is a 1-Lipschitz function of a Gaussian random variable. So, using concentration for 1-Lipschitz functions of normally distributed random variables (see e.g. Boucheron et al., 2013, Theorem 5.5]), we find that on an event with probability at least $1 - \exp(-t)$, it holds that $(P_n - P)U(\cdot, r)/r \leq \sqrt{2t/n}$. On this

event,

$$\varepsilon((P_n - P)U + 2PU) \leq \varepsilon \left(r\sqrt{\frac{2t}{n}} + 2PU \right) \leq \varepsilon r \left(\sqrt{\frac{2t}{n}} + 2\sqrt{p} \right).$$

This completes the proof. \square

Before we turn our attention to controlling the supremum distance over $\mathcal{U}(S)/U_r$, we state an auxiliary result.

Lemma 5.1.3 *Let $v_1, v_2 \in S^{p-1}$. Then:*

$$\sup_{\alpha \in S^{p-1}: \alpha^T v_2 \leq 0} \alpha^T v_1 \leq \|v_1 - v_2\|_2.$$

Proof Choose any $\alpha \in S^{p-1}$. First, suppose that $v_1^T v_2 \leq 0$. Then, it holds that $\|v_1 - v_2\|_2 \geq \sqrt{2} > 1 \geq \alpha^T v_1$. Now suppose that $v_1^T v_2 \geq 0$. Expand:

$$\begin{aligned} \alpha^T v_1 &= \alpha^T v_2 v_2^T v_1 + \alpha^T (I - v_2 v_2^T) v_1 \stackrel{(i)}{\leq} \|(I - v_2 v_2^T) v_1\|_2 \\ &= \sqrt{v_1^T (I - v_2 v_2^T) v_1} = \sqrt{1 - (v_1^T v_2)^2} \leq \sqrt{2 - 2v_1^T v_2} = \|v_1 - v_2\|_2. \end{aligned}$$

In (i), we used that $v_1^T v_2 \geq 0$ for the first term and Cauchy-Schwarz for the second. The proof is complete. \square

Here, we prove an upper bound for the supremum norm in terms of the orientation β and the length τ . As a consequence, we obtain a bound for the covering number.

Lemma 5.1.4 *Let $r > 0$ and define $\tau_1, \tau_2 \in (0, r]$ as well as $\beta_1, \beta_2 \in S^{p-1}$. Then:*

$$\sup_{(x,y) \in \mathbb{R}^p \times \{-1,1\}} \left| \frac{u(\tau_1 \beta_1, x, y) - u(\tau_2 \beta_2, x, y)}{U(x, r)} \right| \leq \|\beta_1 - \beta_2\|_2 + \frac{|\tau_1 - \tau_2|}{r}.$$

Consequently, it takes at most $(12/\varepsilon)^{p+1}$ many elements to cover $\mathcal{U}(B_r)/U_r$ with balls of radius ε in the supremum distance.

FINITE SAMPLE RATES FOR LOGISTIC REGRESSION...

Proof Fix any $(x, y) \in \mathbb{R}^p \times \{-1, +1\}$. For $j \in \{1, 2\}$, we define $1\{A_j\} := 1\{\text{sign}(x^T \beta_j) \neq y\}$. We write out the numerator:

$$|\tau_1 |x^T \beta_1| 1\{A_1\} - \tau_2 |x^T \beta_2| 1\{A_2\}| = \begin{cases} |\tau_1 |x^T \beta_1| - \tau_2 |x^T \beta_2|| & \text{if } A_1 \cap A_2 \\ \tau_1 |x^T \beta_1| & \text{if } A_1 \cap A_2^c \\ \tau_2 |x^T \beta_2| & \text{if } A_1^c \cap A_2 \\ 0 & \text{else.} \end{cases}$$

If $A_1 \cap A_2$, then:

$$\frac{|\tau_1 |x^T \beta_1| - \tau_2 |x^T \beta_2||}{\|x\|_2 r} \leq \frac{\|\tau_1 \beta_1 - \tau_2 \beta_2\|_2}{r} \leq \frac{\tau_1 \|\beta_1 - \beta_2\|_2 + |\tau_1 - \tau_2|}{r}.$$

Now suppose $A_1 \cap A_2^c$ holds. By Lemma 5.1.3:

$$\frac{\tau_1 |x^T \beta_1|}{\|x\|_2 r} \leq \frac{\tau_1 \|\beta_1 - \beta_2\|_2}{r}.$$

The case $A_1^c \cap A_2$ is analogous. To upper bound the covering number, take any $\tau_1 \beta_1, \tau_2 \beta_2 \in U(B_r)/U_r$. We just established that:

$$\left\| \frac{u(\tau_1 \beta_1) - u(\tau_2 \beta_2)}{U(\cdot, r)} \right\|_\infty \leq 2 \sqrt{\left(\frac{\|\beta_1 - \beta_2\|_2^2}{2} + \frac{|\tau_1 - \tau_2|^2}{2r^2} \right)}.$$

Note that on the right-hand side, we have the Euclidean distance over the set $\{(\beta/2, t/2) : (\beta, t) \in S^{p-1} \times [0, 1]\}$, which is a subset of the Euclidean unit ball in \mathbb{R}^{p+1} . To cover the unit ball with elements in itself takes at most $(3/\varepsilon)^{p+1}$ balls, so to cover a subset of it with elements in itself takes at most $(6/\varepsilon)^{p+1}$ elements. Taking care of the extra factor 2 completes the proof. \square

5.2 The Large Noise Case Here we provide an inequality for the empirical process of the unbounded term, in the case where the noise is large. Due to a localization argument, we can assume that $\|\gamma - \gamma^*\|_2 \lesssim \tau^*$ throughout, i.e. $\gamma \in B_r(\gamma^*)$, where $r \lesssim \tau^*$. After verifying Bernstein's condition locally, in terms of the $\|\cdot\|_*$ -norm, we use Lemma 5.1.1, based on Bernstein's inequality, to control the empirical process.

5.2.1 Local Verification of Bernstein's Condition. The following lemma is an intermediate step to verify Bernstein's condition. We do not yet exploit any localization.

Lemma 5.2.1 *Let $\gamma \in \mathbb{R}^p \setminus \{0\}$, $\tau := \|\gamma\|_2$ and $\beta := \gamma/\tau$. Then, for all $m \in \{2, 3, \dots\}$,*

$$\begin{aligned} & \frac{P|\tilde{u}(\gamma)|^m}{\Gamma\left(\frac{m+1}{2}\right) 2^{m-1}} \\ & \leq (\sqrt{32}\tau\|\beta - \beta^*\|_2)^m \left(\|\beta - \beta^*\|_2 + \frac{\sigma}{4\pi^{3/2}} \right) + \frac{\sigma}{4\pi} (\sqrt{8}\pi\sigma|\tau\beta^T\beta^* - \tau^*|)^m. \end{aligned}$$

Proof Define $y^* := \text{sign}(x^T\gamma^*)$. To start, note that:

$$\begin{aligned} |\tilde{u}(\gamma, x, y)|^m &= \left| \frac{|x^T\gamma| - yx^T\gamma - (|x^T\gamma^*| - yx^T\gamma^*)}{2} \right|^m \\ &= \left| \frac{|x^T\gamma| - y^*x^T\gamma}{2} - \frac{(y - y^*)x^T(\gamma - \gamma^*)}{2} \right|^m \\ &\stackrel{(i)}{\leq} 2^{m-1} \left(|x^T\gamma|^m \mathbf{1}\{\beta^T x x^T \beta^* < 0\} + \frac{|(y - y^*)x^T(\gamma - \gamma^*)|^m}{2} \right). \end{aligned}$$

Inequality (i) follows from Jensen's inequality, and the observation that the term $|x^T\gamma| - y^*x^T\gamma/2$ is equal to $|x^T\gamma| \mathbf{1}\{\beta^T x x^T \beta^* < 0\}$. We proceed to upper bound the expectations of those two summands separately. The left term can be bounded using Corollary A.2.1:

$$\begin{aligned} P|x^T\gamma|^m \mathbf{1}\{\beta^T x x^T \beta^* < 0\} &\leq \tau^m \frac{1}{\sqrt{2\pi}} \frac{\Gamma(m/2 + 1)}{m + 1} \left(\frac{\pi}{\sqrt{2}} \|\beta - \beta^*\|_2 \right)^{m+1} \\ &\stackrel{(ii)}{\leq} \tau^m \frac{1}{\sqrt{2\pi}} \Gamma\left(\frac{m+1}{2}\right) \left(\frac{\pi}{\sqrt{2}} \|\beta - \beta^*\|_2 \right)^{m+1}. \end{aligned}$$

In (ii), we used that $\Gamma(m/2 + 1)/(m + 1) \leq \Gamma((m + 1)/2)$. The right term is bounded with Lemma A.2.1:

$$\begin{aligned} & P \frac{|(y - y^*)x^T(\gamma - \gamma^*)|^m}{2} \\ & \leq \frac{\sigma}{4\pi} \Gamma\left(\frac{m+1}{2}\right) \left(\frac{1}{\sqrt{\pi}} (\sqrt{32}\tau\|\beta - \beta^*\|_2)^m + (\sqrt{8}\pi|\tau\beta^T\beta^* - \tau^*|\sigma)^m \right). \end{aligned}$$

Combining these bounds we find:

$$\frac{P|\tilde{u}(\gamma)|^m}{\Gamma\left(\frac{m+1}{2}\right) 2^{m-1}} \leq (\tau\|\beta - \beta^*\|_2)^m \left(\left(\frac{\pi}{\sqrt{2}} \right)^{m+1} \frac{\|\beta - \beta^*\|_2}{\sqrt{2\pi}} + \sqrt{32}^m \frac{\sigma}{4\pi^{3/2}} \right)$$

$$\begin{aligned}
 & + \frac{\sigma}{4\pi} (\sqrt{8\pi}\sigma |\tau\beta^T\beta^* - \tau^*|)^m \\
 \leq & (\sqrt{32}\tau\|\beta - \beta^*\|_2)^m \left(\|\beta - \beta^*\|_2 + \frac{\sigma}{4\pi^{3/2}} \right) + \frac{\sigma}{4\pi} (\sqrt{8\pi}\sigma |\tau\beta^T\beta^* - \tau^*|)^m.
 \end{aligned}$$

The proof is complete. \square

In the following, we assume that $\tau \gtrsim \tau^*$, which is the case if $\|\gamma - \gamma^*\|_2 \lesssim \tau^*$. This assumption allows us to express Lemma 5.2.1 in terms of the $\|\cdot\|_*$ -norm, giving us nice constants for Bernstein's condition. Moreover, we use that $\tau^*\sigma \lesssim 1$, which holds true for $\sigma \lesssim 1$, see Lemma 3.2.2.

Proposition 5.2.1 *Suppose that $\sigma \leq 1/\sqrt{2}$. Let $\beta \in S^{p-1}$ and $\tau > 0$ such that $|\tau - \tau^*| \leq \tau^*/6$. Then, Bernstein's condition is met for $\tilde{u}(\tau\beta) - P\tilde{u}(\tau\beta)$ with constants:*

$$K := 2^8 \sqrt{\tau^*} d_*(\tau\beta), \quad \kappa := 2K \sqrt{\frac{d_*(\tau\beta)}{\sqrt{\tau^*}} + 0.13\sigma}.$$

Proof The result follows from Lemma 5.2.1. Since $\tau \leq \tau^*7/6$,

$$\begin{aligned}
 |\tau\beta^T\beta^* - \tau^*| &= \tau|\beta^T\beta^* - 1| + |\tau - \tau^*| \leq \frac{7}{6}\tau^* \frac{\|\beta - \beta^*\|_2^2}{2} + \tau^{*3/2} \frac{|\tau - \tau^*|}{\tau^{*3/2}} \\
 &\leq \frac{7}{6} \sqrt{\tau^*} \sqrt{\tau^*} \|\beta - \beta^*\|_2 + \tau^{*3/2} \frac{|\tau - \tau^*|}{\tau^{*3/2}} \stackrel{(i)}{\leq} \sqrt{\tau^*} \sqrt{\frac{49}{36} + \tau^{*2} d_*(\tau\beta)} \leq \sqrt{2\tau^{*3}} d_*(\tau\beta).
 \end{aligned}$$

Inequality (i) follows from the Cauchy-Schwarz inequality. In the last inequality we recall from Lemma 3.2.1 that $\sigma \leq 1/\sqrt{2}$ implies $\tau^* > \sqrt{6}$. Now, using the upper bound $\tau\|\beta - \beta^*\|_2 \leq 7/6\sqrt{\tau^*}d_*(\tau\beta)$ we can write the bound in Lemma 5.2.1 as:

$$\begin{aligned}
 & \frac{P|\tilde{u}(\tau\beta)|^m}{\Gamma\left(\frac{m+1}{2}\right) 2^{m-1}} \\
 \leq & (\sqrt{\tau^*}d_*(\tau\beta))^m \left(\left(\sqrt{32}\frac{7}{6}\right)^m \left(\|\beta - \beta^*\|_2 + \frac{\sigma}{4\sqrt{\pi^3}} \right) + \frac{\sigma}{4\pi} (4\pi\sigma\tau^*)^m \right) \\
 \stackrel{(ii)}{\leq} & (\sqrt{\tau^*}d_*(\tau\beta))^m \left(\left(\sqrt{32}\frac{7}{6}\right)^m \left(\|\beta - \beta^*\|_2 + \frac{\sigma}{4\sqrt{\pi^3}} \right) + \frac{\sigma}{4\pi} \sqrt{32\pi^3}^m \right) \\
 \leq & (\sqrt{32}\pi^{3/2}\sqrt{\tau^*}d_*(\tau\beta))^m \left(\|\beta - \beta^*\|_2 + \sigma \left(\frac{1}{4\sqrt{\pi^3}} + \frac{1}{4\pi} \right) \right).
 \end{aligned}$$

Finally, we use that $\|\beta - \beta^*\|_2 \leq d_*(\tau\beta)/\sqrt{\tau^*}$ and simplify the constants $2\sqrt{32}\pi^{3/2} < 2^6$ and $1/(4\sqrt{\pi^3}) + 1/(4\pi) < 0.13$. To find the constants in Bernstein's condition, note that $\Gamma((m+1)/2) \leq m!/2$ for $m \geq 2$. Recall that an extra factor of 2 is necessary, as \tilde{u} is not centered. The proof is complete. \square

5.2.2 Controlling the Empirical Process. All ingredients are prepared to prove an upper bound for the empirical process of the unbounded term, in the large noise case. We will exploit Bernstein's inequality, or actually, the derived Lemma 5.1.1, using the constants from Proposition 5.2.1, and making use of the discretization result (Lemma 5.1.4). To exploit that the constants in Proposition 5.2.1 are better if $d_*(\gamma)$ is small, we start the proof with a local bound, and then move on to a global bound with peeling. Recall the definition of the δ -ball for the Euclidean distance $B_\delta := \{\gamma \in \mathbb{R}^p : \|\gamma\|_2 \leq r\}$, as well as for the $\|\cdot\|_*$ -norm: $B_\delta^* := \{\gamma \in \mathbb{R}^p \setminus \{0\} : d_*(\gamma) \leq \delta\}$.

Proposition 5.2.2 *Suppose that $\sigma \leq 1/\sqrt{2}$. Let $R > 0$, $r \in (0, 1)$, fix $t > 0$ and define $t' := t + \log\lceil \log R/r \rceil$. Then, with probability at least $1 - 3\exp(-t)$,*

$$\sup_{\gamma \in S_{r,R}} \frac{|(P_n - P)\tilde{u}(\gamma)| - \Delta(t')}{\psi_u(ed_*(\gamma), t')} \leq 1.$$

Here,

$$S_{r,R} := \left\{ \gamma \in \mathbb{R}^p \setminus \{0\} : r \leq d_*(\gamma) \leq R, \quad \left| \|\gamma\|_2 - \tau^* \right| < \frac{\tau^*}{6} \right\},$$

$$\rho(t) := \frac{2(t + (p+1)\log(12\tau^*/r))}{n}, \quad \Delta(t) := \frac{7r}{6} \left(\sqrt{\frac{2t}{n}} + 2\sqrt{p} \right),$$

$$\psi_u(\delta, t) := 2^8 \sqrt{\tau^*} \delta \left(2\sqrt{\rho(t) \left(\frac{\delta}{\sqrt{\tau^*}} + 0.13\sigma \right)} + \rho(t) \right).$$

Proof Let $\delta > 0$ be arbitrary. We start with discretization. Let $\mathcal{N} \subset S_{r,\delta}$ be such that $\mathcal{U}(\mathcal{N})/U_{\tau^*7/6}$ is a covering of $\mathcal{U}(S_{r,\delta})/U_{\tau^*7/6}$ in the supremum norm of radius r/τ^* . By Lemma 5.1.4, \mathcal{N} can be chosen so that it has at most $(12\tau^*/r)^{p+1}$ elements. We apply Lemma 5.1.2 with radius $\tau^*7/6$ and approximation error r/τ^* . On an event with probability at least $1 - \exp(-t)$,

$$\sup_{\gamma \in S_{r,\delta}} |(P_n - P)\tilde{u}(\gamma)| \leq \sup_{\gamma \in \mathcal{N}} |(P_n - P)\tilde{u}(\gamma)| + \Delta(t).$$

FINITE SAMPLE RATES FOR LOGISTIC REGRESSION...

By Proposition 5.2.1, $\{\tilde{u}(\tau\beta) - P\tilde{u}(\tau\beta) : \beta \in S_{r,\delta}\}$ satisfies Bernstein's condition with constants:

$$K := 2^8 \sqrt{\tau^*} \delta, \quad \kappa := 2K \sqrt{\frac{\delta}{\sqrt{\tau^*}} + 0.13\sigma}.$$

Therefore, by Lemma 5.1.1, on an event with probability at least $1 - 2 \exp(-t)$,

$$\sup_{\gamma \in \mathcal{N}} |(P_n - P)\tilde{u}(\gamma)| \leq 2^8 \sqrt{\tau^*} \delta \left(2 \sqrt{\rho(t) \left(\frac{\delta}{\sqrt{\tau^*}} + 0.13\sigma \right)} + \rho(t) \right) = \psi_u(\delta, t),$$

where ρ and ψ_u are defined above. We conclude that for any $\delta > 0$,

$$\mathbb{P} \left[\bigcup_{\gamma \in S_{r,\delta}} \frac{|(P_n - P)\tilde{u}(\gamma)| - \Delta(t)}{\psi_u(\delta, t)} > 1 \right] \leq 3 \exp(-t).$$

Now we use peeling, to replace the dependency on δ with a dependency on $d_*(\gamma)$. Define $J := \lceil \log R/r \rceil$ and $t' := t + \log J$. Then:

$$\begin{aligned} & \mathbb{P} \left[\bigcup_{\gamma \in S_{r,R}} \frac{|(P_n - P)\tilde{u}(\gamma)| - \Delta(t')}{\psi_u(ed_*(\gamma), t')} > 1 \right] \\ & \leq \sum_{j=0}^{J-1} \mathbb{P} \left[\bigcup_{\gamma \in S_{re^j, re^{j+1}}} \frac{|(P_n - P)\tilde{u}(\gamma)| - \Delta(t')}{\psi_u(ed_*(\gamma), t')} > 1 \right] \\ & \leq \sum_{j=0}^{J-1} \mathbb{P} \left[\bigcup_{\gamma \in S_{re^j, re^{j+1}}} \frac{|(P_n - P)\tilde{u}(\gamma)| - \Delta(t')}{\psi_u(re^{j+1}, t')} > 1 \right] \\ & \leq J3 \exp(-t - \log J) = 3 \exp(-t). \end{aligned}$$

The proof is complete. □

5.3 The Small Noise Case We provide an upper bound for the unbounded term, in the case where the noise is small. Compared to the large noise case, we do not have any upper bound for τ^* . On the other hand, we can exploit that σ is very small. Moreover, the proof for the small noise

case only requires that we control $\|\beta - \beta^*\|_2$, which simplifies the expressions considerably. As before, we first verify Bernstein's condition. In this subsection, we use the notation:

$$\tilde{u}(\beta) := u(\beta) - u(\beta^*).$$

5.3.1 Local Verification of Bernstein's Condition. Similarly to the large noise case, we verify Bernstein's condition. The verification is somewhat more straightforward, as we can compare β on the sphere with β^* directly.

Lemma 5.3.1 *Let $\beta \in S^{p-1}$. Bernstein's condition is met for $\tilde{u}(\beta) - P\tilde{u}(\beta)$ with constants:*

$$K := \left(\frac{\pi}{\sqrt{2}} \right) (\|\beta - \beta^*\|_2 + 2\sigma), \quad \kappa := K \sqrt{\|\beta - \beta^*\|_2 + 2\sigma}.$$

Proof It holds that:

$$\begin{aligned} & P|x^T \beta^*|^m \mathbf{1}\{yx^T \beta^* < 0\} - |x^T \beta|^m \mathbf{1}\{yx^T \beta < 0\} \\ & \leq P|x^T \beta^*|^m \mathbf{1}\{yx^T \beta^* < 0\} + |x^T \beta|^m \mathbf{1}\{yx^T \beta < 0\}. \end{aligned}$$

We first consider the summand with β . To control it, we use Corollary A.2.1. If $\tilde{x} \sim \mathcal{N}(0, I_{p+1})$:

$$\begin{aligned} P|x^T \beta|^m \mathbf{1}\{yx^T \beta < 0\} &= P|x^T \beta|^m \mathbf{1}\{(\beta, 0)^T \tilde{x} \tilde{x}^T (\beta^*, \sigma) < 0\} \\ &\stackrel{(i)}{\leq} \frac{1}{\sqrt{2\pi}} \frac{\Gamma(m/2 + 1)}{m + 1} \left(\frac{\pi}{\sqrt{2}} \left\| (\beta, 0) - \frac{1}{\sqrt{1 + \sigma^2}} (\beta^*, \sigma) \right\|_2 \right)^{m+1} \\ &\stackrel{(ii)}{\leq} \frac{1}{2} \frac{\Gamma(m/2 + 1)}{m + 1} \left(\frac{\pi}{\sqrt{2}} \right)^m (\|\beta - \beta^*\|_2 + \sigma)^{m+1} \\ &\stackrel{(iii)}{\leq} \frac{m!}{2} \left(\frac{\pi}{\sqrt{8}} \right)^m (\|\beta - \beta^*\|_2 + \sigma)^{m+1}. \end{aligned}$$

In (i), we used Corollary A.2.1. In (ii), we used the triangular inequality and that $1 - \frac{1}{\sqrt{1 + \sigma^2}} \leq \frac{\sigma^2}{2}$. In (iii), we used that for $m \geq 2$ it holds that $\Gamma(m/2 + 1)/(m + 1) \leq m!/2^m$. We can apply the same reasoning to the term $P|x^T \beta^*|^m \mathbf{1}\{yx^T \beta^* < 0\}$ by choosing β as β^* . Exploiting the fact that integer monomials are superadditive over positive scalars, we find:

$$P|\tilde{u}(\beta)|^m \leq \frac{m!}{2} \left(\frac{\pi}{\sqrt{8}} \right)^m (\|\beta - \beta^*\|_2 + 2\sigma)^{m+1}.$$

This gives constants for Bernstein's condition. Recall the extra factor 2, since \tilde{u} are not necessarily centered. The proof is complete. \square

5.3.2 Controlling the Empirical Process. Here, we prove a high probability upper bound for the empirical process of the unbounded term, in the small noise case. The procedure is analogous to the large noise case. We recall the notation $S(\delta, \beta') := \{\beta \in S^{p-1} : \|\beta - \beta'\|_2 \leq \delta\}$ for $\delta > 0$ and $\beta' \in S^{p-1}$.

Proposition 5.3.1 *Let $r, R > 0$. Fix any $t > 0$ and let $t' := t + \log\lceil \log R/r \rceil$. With probability at least $1 - \exp(-t)$,*

$$\sup_{\beta \in S(R, \beta^*) \setminus S(r, \beta^*)} \frac{|(P_n - P)\tilde{u}(\beta)| - \Delta(t')}{\frac{\pi}{\sqrt{2}}\xi(\beta)(\sqrt{\xi(\beta)\rho(t')} + \rho(t'))} \leq 1.$$

Here,

$$\Delta(t) := r \left(\sqrt{\frac{2t}{n}} + 2\sqrt{p} \right), \quad \rho(t) := \frac{t + (p+1) \log\left(\frac{12}{r}\right)}{n},$$

$$\xi(\beta) := e\|\beta - \beta^*\|_2 + 2\sigma.$$

Proof Let $\delta > 0$ be arbitrary. We start with discretization. By Lemma 5.1.4, it takes at most $(12/r)^{p+1}$ balls of radius r to cover $\{\tilde{u}(\beta)/U(x, 1) : \beta \in S^{p-1}\}$ with elements in itself in the supremum norm. Let $\mathcal{N} \subset S(\delta, \beta^*) \setminus S(r, \beta^*)$ be such a covering of $S(\delta, \beta^*) \setminus S(r, \beta^*)$. We use Lemma 5.1.2 with bound 1 and covering radius $\varepsilon := r$. On an event with probability at least $1 - \exp(-t)$,

$$\sup_{\beta \in S(\delta, \beta^*) \setminus S(r, \beta^*)} |(P_n - P)\tilde{u}(\beta)| \leq \sup_{\beta \in \mathcal{N}} |(P_n - P)\tilde{u}(\beta)| + \Delta(t).$$

By Lemma 5.3.1, the centered \tilde{u} with parameters in $S(\delta, \beta^*) \setminus S(r, \beta^*)$ satisfy Bernstein's condition, with parameters $K := \frac{\pi}{\sqrt{2}}(\delta + 2\sigma)$ and $\kappa := K\sqrt{\delta + 2\sigma}$. By Lemma 5.1.1, with probability at least $1 - 2\exp(-t)$,

$$\sup_{\beta \in \mathcal{N}} \frac{|(P_n - P)\tilde{u}(\beta)|}{\frac{\pi}{\sqrt{2}}(\delta + 2\sigma)} \leq \sqrt{(\delta + 2\sigma)\rho(t)} + \rho(t).$$

We use peeling to replace the dependency on δ with dependency on $\|\beta - \beta^*\|_2$. Define $J := \lceil \log R/r \rceil$ and $t' := t + \log J$.

$$\mathbb{P} \left[\bigcup_{\beta \in S(R, \beta^*) \setminus S(r, \beta^*)} \frac{|(P_n - P)\tilde{u}(\beta)| - \Delta(t')}{\frac{\pi}{\sqrt{2}} \xi(\beta) (\sqrt{\xi(\beta)\rho(t')} + \rho(t'))} > 1 \right] \\ \stackrel{(i)}{\leq} J 3 \exp(-t - \log J) = 3 \exp(-t).$$

As in the large noise proof, (i) follows from peeling, i.e. a union bound over a partition of the form $S(e^{j+1}r, \beta^*) \setminus S(e^j r, \beta^*)$. The proof is complete. \square

6 Proofs of the Main Theorems

6.1 Large Noise Case In this subsection, we prove the main theorem for the case where σ is large. As a final auxiliary step, we prove the margin condition. It provides a lower bound for the excess risk $Pl(\gamma) - Pl(\gamma^*)$ in terms of $d_*(\gamma)$. Once this result is established, we are ready to prove the main result.

6.1.1 The Margin Condition. Recall our notation $\gamma \in \mathbb{R}^p \setminus \{0\}$ and $\tau := \|\gamma\|_2$, as well as $\beta := \gamma/\tau$. To derive bounds on $\|\hat{\beta} - \beta^*\|_2$ and $|\hat{\tau} - \tau^*|$ from our bounds on the excess risk, we would like to have an inequality of the following type, for some $\alpha_1, \alpha_2 > 0$:

$$\alpha_1 |\hat{\tau} - \tau^*| + \alpha_2 \|\hat{\beta} - \beta^*\|_2 \leq Pl(\gamma) - Pl(\gamma^*).$$

This is what we call the *margin condition*. It was originally introduced in Mammen and Tsybakov (1999). The statement we use is reminiscent of the version in (Bühlmann and van de Geer, 2011, Chapter 6.4). In the following, we show that indeed the margin condition is satisfied, with:

$$d_*(\gamma)^2 := \|(\tau, \beta) - (\tau^*, \beta^*)\|_*^2 \lesssim Pl(\gamma) - Pl(\gamma^*).$$

This result is the motivation for our introduction of the $*$ -norm (see Section 3.3).

Re-parametrization As the signal strength $1/\sigma \sim \tau^*$ grows, it becomes more difficult to estimate τ^* but easier to estimate β^* . A lower bound on the excess risk must reflect this. Hence, in our attempt to derive such a bound, we re-parametrize γ into its length $\tau := \|\gamma\|_2$ and its orientation $\beta := \gamma/\|\gamma\|_2$. We will not work with β directly but rather exploit that β

lives in a $p - 1$ -dimensional manifold. To do so, we re-parametrize β in the following.

We find an orthonormal basis, which includes β^* as one of its elements. For any $\beta^* \in S^{p-1}$, there exists a set $\{b_1, \dots, b_{p-1}\} \subset S^{p-1}$, such that the collection $\{\beta^*, b_1, \dots, b_{p-1}\}$ forms an orthonormal basis. So, we can find the Fourier expansion of any vector $v \in \mathbb{R}^p$ as:

$$v = v^T \beta^* \beta^* + \sum_{i=1}^{p-1} v^T b_i b_i.$$

We write this more compactly for vectors on S^{p-1} . Let V be the $p \times (p-1)$ matrix, whose columns are the vectors $\{b_1, \dots, b_{p-1}\}$. For any $v \in \mathbb{R}^p$, we define the coefficient vector $h(v) := (v^T b_1, \dots, v^T b_{p-1})$. So, for any $\beta \in S^{p-1}$,

$$\beta = \text{sign}(\beta^T \beta^*) \sqrt{1 - \|h(\beta)\|_2^2} \beta^* + Vh(\beta). \quad (9)$$

The following result shows that the Euclidean norm of $h(\beta)$ behaves as the distance between β and β^* , up to a constant factor.

Lemma 6.1.1 *For any $\beta^*, \beta \in S^{p-1}$ with $\beta^T \beta^* \geq 0$, it holds that:*

$$\|h(\beta)\|_2 \leq \|\beta - \beta^*\|_2 \leq \sqrt{2} \|h(\beta)\|_2.$$

Proof Fix $\beta \in S^{p-1}$. To reduce notation, we write $h := h(\beta)$. By definition,

$$\|\beta - \beta^*\|_2^2 = \left(1 - \sqrt{1 - \|h\|_2^2}\right)^2 + \|h\|_2^2 = 2 \left(1 - \sqrt{1 - \|h\|_2^2}\right).$$

From this, the lower bound already follows. For the upper bound, multiply out the first term, simplify, and then use that $1 - \sqrt{1 - \|h\|_2^2} = \|h\|_2^2 / (1 + \sqrt{1 - \|h\|_2^2}) \leq \|h\|_2^2$. This completes the proof. \square

We now re-write the true risk (the expectation of the loss function) at a point $\gamma \in \mathbb{R}^p \setminus \{0\}$. For $\gamma \in \mathbb{R}^p$ such that $\gamma^T \gamma^* \geq 0$, we write the ‘true risk’ as:

$$\begin{aligned} & P \log(1 + \exp(-yx^T \gamma)) \\ &= P \log \left(1 + \exp \left(-\tau yx^T \left(\sqrt{1 - \|h(\beta)\|_2^2} \beta^* + Vh(\beta) \right) \right) \right). \end{aligned}$$

The expansion (9) shows that on the hemisphere $H := \{b \in S^{p-1} : b^T \beta^* > 0\}$, the map $\beta \mapsto h(\beta)$ is injective. So, we can reparametrize $(\tau, \beta) \in$

$(0, \infty) \times H$ with $(\tau, h) \in (0, \infty) \times B_1^\circ$, where B_1° is the open Euclidean unit ball, and write the true risk as:

$$P \log \left(1 + \exp \left(-\tau y x^T \left(\sqrt{1 - \|h\|_2^2} \beta^* + Vh \right) \right) \right) =: R(\tau, h). \quad (10)$$

The mapping $R : (0, \infty) \times B_1^\circ \rightarrow \mathbb{R}$ is twice continuously differentiable. In the following, we will study its Hessian.

The Hessian of the re-parametrized risk

Lemma 6.1.2 *Let $z \sim \mathcal{N}(0, 1)$. For any $\tau > 0$ and any $h \in B_1^\circ$, we have:*

$$\ddot{R}(\tau, h) := \begin{pmatrix} P \frac{|z|^2}{(\exp(\tau|z|/2) + \exp(-\tau|z|/2))^2}, & -\frac{1}{\sqrt{2\pi(1+\sigma^2)}\sqrt{1-\|h\|_2^2}} h^T \\ -\frac{1}{\sqrt{2\pi(1+\sigma^2)}\sqrt{1-\|h\|_2^2}} h, & \frac{\tau}{\sqrt{2\pi(1+\sigma^2)}\sqrt{1-\|h\|_2^2}} \left(I + \frac{hh^T}{1-\|h\|_2^2} \right) \end{pmatrix}.$$

Proof Fix $\tau > 0$ and $h \in B_1^\circ$. Let $\beta \in S^{p-1}$ be the unique vector such that $\beta^T \beta^* > 0$ and the identity (9) holds. Let $\gamma := \tau \beta$. We use the decomposition $\log(1 + \exp(-yx^T \gamma)) = \log(1 + \exp(-|\gamma^T x|)) + |x^T \gamma| \mathbb{1}\{yx^T \gamma < 0\}$. Using Corollary A.2.2 for the right summand, we find:

$$R(\tau, h) = P \log(1 + \exp(-\tau|z|)) + \tau \frac{1}{\sqrt{2\pi}} \left(1 - \frac{\beta^T \beta^*}{\sqrt{1 + \sigma^2}} \right).$$

By construction of β , it holds that $\beta^T \beta^* = \sqrt{1 - \|h\|_2^2}$. Differentiating twice, we find the expression stated in the claim. The proof is complete. \square

We now give a lower bound on the second-order term of the Taylor expansion of R near γ^* , i.e. $(\tau^*, 0)$. At this moment, the $*$ -norm appears in our calculation.

Lemma 6.1.3 *Let $\tau, \tau^* \geq \sqrt{6 + \sqrt{51}}$ and $\kappa_\tau > 0$ such that $|\tau - \tau^*| \leq \kappa_\tau \tau^*$. Let $(\bar{\tau}, \bar{h})$ be a convex combination of (τ, h) and $(\tau^*, 0)$. Then,*

$$\begin{pmatrix} \tau - \tau^* \\ h \end{pmatrix}^T \ddot{R}(\bar{\tau}, \bar{h}) \begin{pmatrix} \tau - \tau^* \\ h \end{pmatrix} \geq \frac{1}{\sqrt{8\pi}} \left(\frac{1}{(1 + \kappa_\tau)^3} \wedge \frac{1 - 3\kappa_\tau}{\sqrt{1 + \sigma^2}} \right) d_*(\tau\beta(h))^2.$$

Proof We make use of the expression given in Lemma 6.1.2. We bound each term separately. First, by Lemma A.1.5,

$$P \frac{|z|^2}{(e^{\bar{\tau}|z|/2} + e^{-\bar{\tau}|z|/2})^2} (\tau - \tau^*)^2 \geq \sqrt{\frac{1}{8\pi}} \frac{(\tau - \tau^*)^2}{\bar{\tau}^3} \geq \sqrt{\frac{1}{8\pi}} \frac{(\tau - \tau^*)^2}{(\tau^*(1 + \kappa_\tau))^3}.$$

For the cross terms,

$$\begin{aligned}
 h^T \ddot{R}((\bar{\tau}, \bar{h}))_{21}(\tau - \tau^*) &= -\frac{\tau - \tau^*}{\sqrt{2\pi(1 + \sigma^2)}} \frac{\bar{h}^T h}{\sqrt{1 - \|\bar{h}\|_2^2}} \\
 &\stackrel{(i)}{\geq} -\frac{|\tau - \tau^*|}{\sqrt{2\pi(1 + \sigma^2)}} \frac{\|h\|_2^2}{\sqrt{1 - \|\bar{h}\|_2^2}} \geq -\frac{\kappa_\tau \tau^*}{\sqrt{2\pi(1 + \sigma^2)}} \frac{\|h\|_2^2}{\sqrt{1 - \|\bar{h}\|_2^2}}.
 \end{aligned}$$

In (i), we used that \bar{h} is an intermediate point between h and 0, $\|\bar{h}\|_2 \leq \|h\|_2$. The inequality then follows from Cauchy-Schwarz. The last product gives:

$$\begin{aligned}
 h^T \ddot{R}(\bar{\theta})_{22}h &\geq \frac{\bar{\tau}}{\sqrt{2\pi(1 + \sigma^2)}} \frac{\|h\|_2^2}{\sqrt{1 - \|\bar{h}\|_2^2}} \\
 &\geq \frac{\tau^*(1 - \kappa_\tau)}{\sqrt{2\pi(1 + \sigma^2)}} \frac{\|h\|_2^2}{\sqrt{1 - \|\bar{h}\|_2^2}}.
 \end{aligned}$$

Together with the two cross terms, this is:

$$\frac{\tau^*(1 - 3\kappa_\tau)}{\sqrt{2\pi(1 + \sigma^2)}} \frac{\|h\|_2^2}{\sqrt{1 - \|\bar{h}\|_2^2}} \geq \frac{\tau^*(1 - 3\kappa_\tau)}{\sqrt{2\pi(1 + \sigma^2)}} \|h\|_2^2.$$

By Lemma 6.1.1, $\|h\|_2^2 \geq \|\beta - \beta^*\|_2^2/2$. Combining these lower bounds, we find:

$$\begin{aligned}
 &\frac{1}{\sqrt{8\pi}} \left(\frac{1}{(1 + \kappa_\tau)^3} \frac{(\tau - \tau^*)^3}{\tau^{*3}} + \frac{1 - 3\kappa_\tau}{\sqrt{1 + \sigma^2}} \tau^* \|\beta - \beta^*\|_2^2 \right) \\
 &\geq \frac{1}{\sqrt{8\pi}} \left(\frac{1}{(1 + \kappa_\tau)^3} \wedge \frac{1 - 3\kappa_\tau}{\sqrt{1 + \sigma^2}} \right) d_*(\tau\beta)^2.
 \end{aligned}$$

The proof is complete. \square

Proof of the margin condition

Lemma 6.1.4 *Let $\tau, \tau^* \geq \sqrt{6 + \sqrt{51}}$ and $\kappa_\tau > 0$ such that $|\tau - \tau^*| \leq \kappa_\tau \tau^*$. Suppose that $\beta, \beta^* \in S^{p-1}$ such that $\beta^T \beta^* > 0$. Then,*

$$Pl(\tau\beta) - l(\tau^*\beta^*) \geq \frac{1}{\sqrt{32\pi}} \left(\frac{1}{(1 + \kappa_\tau)^3} \wedge \frac{1 - 3\kappa_\tau}{\sqrt{1 + \sigma^2}} \right) d_*(\tau\beta)^2.$$

Proof Using the re-parametrization, we have $Pl(\tau\beta) - l(\tau^*\beta^*) = R(\tau, h(\beta)) - R(\tau^*, 0)$. By Taylor's Theorem, there exists a convex combination of $(\tau, h(\beta))$ and $(\tau, h(\beta))$, say $(\bar{\tau}, \bar{h})$, such that:

$$R(\tau, h(\beta)) - R(\tau^*, 0) \geq \dot{R}(\tau^*, 0)^T \begin{pmatrix} \tau - \tau^* \\ h \end{pmatrix} + \frac{1}{2} \begin{pmatrix} \tau - \tau^* \\ h \end{pmatrix}^T \ddot{R}(\bar{\tau}, \bar{h}) \begin{pmatrix} \tau - \tau^* \\ h \end{pmatrix}.$$

By convexity and the definition of $\tau^*\beta^*$, $\dot{R}(\tau^*, 0) = 0$. The result now follows from Lemma 6.1.3. \square

6.1.2 Proof of Large Noise Theorem. Here, we prove Theorem 2.1.1. In fact, we prove Theorem 6.1.1, which is a slightly more general version, as it allows for a weaker restriction on the choice of M . Moreover, it is marginally less restrictive on σ , as with Lemma 3.2.1 one can verify that $\sigma \leq 1/\sqrt{6}$ implies $\tau^* > \frac{6}{5}\sqrt{6 + \sqrt{51}}$.

Theorem 6.1.1 *Let $\hat{\gamma}$ be a solution to Eq. 2. There exists universal constants $C > 0$ and $0 < c < 1$, such that for any $t > 0$, if:*

$$\frac{6}{5}\sqrt{6 + \sqrt{51}} < \tau^* \leq c \frac{n}{p \log n + t} \leq M,$$

then with probability at least $1 - 4 \exp(-t)$,

$$d_*(\hat{\gamma}) \leq C \sqrt{\frac{p \log n + t}{n}}.$$

We refer the reader to Section 2 for a sketch of the proof.

Proof We define:

$$\alpha := \frac{1}{1 + \frac{6\|\hat{\gamma} - \gamma^*\|_2}{\tau^*}}, \quad \tilde{\gamma} := \alpha\hat{\gamma} + (1 - \alpha)\gamma^*$$

We show that we can apply Lemma 6.1.5 to $\tilde{\gamma}$. By construction of $\tilde{\gamma}$,

$$\|\tilde{\gamma} - \gamma^*\|_2 = \alpha\|\hat{\gamma} - \gamma^*\|_2 \leq \frac{\tau^*}{6}.$$

By convexity, $P_n l(\tilde{\gamma}) \leq P_n l(\gamma^*)$ holds on an event with probability 1. So, by Lemma 6.1.5, there exists a constant $C' > 0$, such that on an event B_t

with probability at least $1 - 4 \exp(-t)$,

$$d_*(\tilde{\gamma}) \leq C' \sqrt{\frac{p \log(n) + t}{n}}.$$

By Lemma 3.3.1, it follows that:

$$\|\tilde{\gamma} - \gamma^*\|_2 \leq \sqrt{2\tau^*} C' \sqrt{\frac{p \log(n) + t}{n}} \leq c \cdot C' \sqrt{2} \cdot \tau^* \stackrel{(i)}{\leq} \frac{\tau^*}{12}.$$

The inequality (i) holds if $c \leq 1/(C'12\sqrt{2})$, which we require. If so, after rearranging we find:

$$\|\hat{\gamma} - \gamma^*\|_2 = \frac{\|\tilde{\gamma} - \gamma^*\|_2}{1 - \frac{6\|\tilde{\gamma} - \gamma^*\|_2}{\tau^*}} \leq 2\|\tilde{\gamma} - \gamma^*\|_2.$$

Consequentially, $\|\hat{\gamma} - \gamma^*\|_2 \leq \tau^*/6$. So, by Lemma 6.1.5 on the event B_t ,

$$d_*(\hat{\gamma}) \leq C' \sqrt{\frac{p \log(n) + t}{n}}.$$

The proof is complete. □

Lemma 6.1.5 Let $\mu := \frac{p \log(n) + t}{n}$ and suppose that:

$$\frac{6}{5} \sqrt{6 + \sqrt{51}} \leq \tau^* \leq \frac{1}{\mu}.$$

There exists a universal constant $C > 0$ and for any $t > 0$ an event B_t with probability at least $1 - 4 \exp(-t)$, such that the following is true. Let $\tilde{\gamma}$ be any random vector in \mathbb{R}^p which satisfies $\|\tilde{\gamma} - \gamma^*\|_2 \leq \tau^*/6$ and $P_n l(\tilde{\gamma}) \leq P_n l(\gamma^*)$ on an event A . Then, on $A \cap B_t$, it holds that:

$$d_*(\tilde{\gamma}) \leq C \sqrt{\mu}.$$

Proof We will apply the margin condition (Lemma 6.1.4) with $\kappa_\tau = 1/6$, which gives us:

$$\begin{aligned} \frac{1}{\sqrt{27\pi}} \frac{d_*(\tilde{\gamma})^2}{\sqrt{1 + \sigma^2}} &\leq Pl(\tilde{\gamma}) - l(\gamma^*) \leq (P - P_n)l(\tilde{\gamma}) - l(\gamma^*) \\ &= (P - P_n) \log(2) \tilde{b}(\tilde{\gamma}) + \tilde{u}(\tilde{\gamma}). \end{aligned}$$

In the following, we distinguish whether the bounded or unbounded term dominates. Suppose that the bounded term dominates. If so, we apply Proposition 4.2.1. Henceforth, we borrow its notation. Since $\|\tilde{\gamma} - \gamma^*\|_2 \leq \tau^*/6$, we have $\tilde{\tau} := \|\tilde{\gamma}\|_2 \geq \tau^*5/6$ and by Lemma 3.3.1 $d_*(\tilde{\gamma}) \leq \sqrt{\tau^*}/2$. So, we may choose the bigger radius in Proposition 4.2.1 as $R := \sqrt{\tau^*}/2$. Moreover, we choose $r := \mu$. By Proposition 4.2.1, there exists a universal constant C_b such that on an event B_b , which has probability at least $1 - \exp(-t)$,

$$(P - P_n)\tilde{b}(\tilde{\gamma}) \lesssim \mu + \sqrt{\mu}d_*(\tilde{\gamma}).$$

Distinguishing whether the first or second term dominates, we obtain $d_*(\tilde{\gamma}) \lesssim \mu$. Here, we used that $\sigma \lesssim 1$ by Lemma 3.2.2 and since $\tau^* \gtrsim 1$.

Now suppose instead that the unbounded term dominates. We apply Proposition 5.2.2, borrowing its notation. The radii are chosen as $r := \mu/\sqrt{p}$ and $R := \sqrt{\tau^*}/2$. We see that $\Delta(t') \lesssim \mu$ and $\rho(t') \lesssim \mu$, as well as:

$$\psi_u(d_*(\tilde{\gamma}), t') \lesssim \sqrt{\tau^*}d_*(\tilde{\gamma}) \left(\sqrt{\left(\frac{d_*(\tilde{\gamma})}{\sqrt{\tau^*}} + \sigma\right)} \mu + \mu \right).$$

So, by Proposition 5.2.2, on an event B_u with probability at least $1 - 3\exp(-t)$,

$$\begin{aligned} (P - P_n)\tilde{u}(\tilde{\gamma}) &\lesssim d_*(\tilde{\gamma})^{3/2}\sqrt{\mu\tau^{*1/4}} + d_*(\tilde{\gamma})\sqrt{\mu\tau^*\sigma} + d_*(\tilde{\gamma})\mu\sqrt{\tau^*} + \mu \\ &\lesssim d_*(\tilde{\gamma})^{3/2}\mu^{1/4} + d_*(\tilde{\gamma})\sqrt{\mu} + \mu. \end{aligned}$$

Here, we used Lemma 3.2.2 to simplify $\sigma\tau^* \lesssim 1$. Whichever term dominates, we obtain $d_*(\tilde{\gamma}) \lesssim \sqrt{\mu}$.

If we denote the event $B_b \cap B_u$ by B_t , the event B_t has probability at least $1 - 4\exp(-t)$, and on $A \cap B_t$, the event $d_*(\tilde{\gamma}) \lesssim \mu$ holds. The proof is complete. \square

6.2 Proof for Small Noise In this subsection, we prove Theorem 2.2.1, the main result for the case where the noise is small. In fact, we prove the slightly stronger Theorem 6.2.1, which neither requires that $M \geq n/(p \log n + t)$ nor that $\sigma \leq (p \log n + t)/n$. Let $\hat{\gamma}$ be a solution to Eq. 2. Define $\hat{\tau} := \|\hat{\gamma}\|_2$ and $\hat{\beta} := \hat{\gamma}/\hat{\tau}$. It can be verified that almost surely $\hat{\tau} \neq 0$. Consequentially, $\hat{\beta}$ is well-defined.

Theorem 6.2.1 Fix $t > 0$ and let $\mu := (p \log(n) + t)/n$. Assume $\sigma \leq 1/\sqrt{7}$ and require $M > 2\sqrt{6}$ and:

$$0.051 > \sqrt{\frac{9t + 3p(2 + \log(3\sqrt{6n}))}{n}}. \quad (11)$$

Then, with probability at least $1 - 6 \exp(-t)$,

$$\|\hat{\beta} - \beta^*\|_2 \lesssim \frac{1}{M} \vee \sigma \vee \mu, \quad (12)$$

and:

$$M \wedge \frac{1}{\sigma} \wedge \frac{1}{\mu} \lesssim \hat{\tau}. \quad (13)$$

As a direct consequence of Theorem 6.2.1, we get Theorem 2.2.1, which was stated in Section 2. To recover its statement, we require $\sigma \leq \mu$ and $M \geq 1/\mu$. We point the reader to Section 2 for a sketch of the proof.

Proof of Theorem 6.2.1 The proof is divided into 3 cases. Case 3 describes $\hat{\tau} \leq \sqrt{6}$, which we show happens only with small probability. Cases 1-2 capture $\hat{\tau} > \sqrt{6}$. In case 1, we additionally assume $M/2 < \hat{\tau} \leq M$, whereas in case 2 we assume $\sqrt{6} < \hat{\tau} \leq M/2$. On these events, we show that with large probability, Eqs. 12 and 13 hold.

Case 1: $M/2 \vee \sqrt{6} \leq \hat{\tau} \leq M$. In case 1, Eq. 13 follows directly by assumption, so we only need to prove (12). By the lower bound in Lemma A.1.6, using that $\hat{\tau} \geq \sqrt{6}$,

$$\begin{aligned} \frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} &\leq P(b(\hat{\tau}\beta^*) - b(2\hat{\tau}\beta^*)) \stackrel{(i)}{=} P(b(\hat{\tau}\hat{\beta}) - b(2\hat{\tau}\beta^*)) \\ &\stackrel{(ii)}{\leq} P(b(\hat{\tau}\hat{\beta}) - b(\hat{\tau}\beta^*)) + \frac{1}{\hat{\tau}} \sqrt{\frac{2}{\pi}} \stackrel{(iii)}{\leq} P(b(\hat{\tau}\hat{\beta}) - b(\hat{\tau}\beta^*)) + \frac{1}{M} \sqrt{\frac{8}{\pi}}. \end{aligned}$$

The equality in (i) follows from rotational invariance, whereas the inequality (ii) follows from the upper bound in Lemma A.1.6. Inequality (iii) follows from the assumption $M/2 \leq \hat{\tau}$. It turns out that this is the only step where the assumption $\sqrt{6} \leq \hat{\tau} \leq M$ would not be enough in the proof of part 1. On the other hand, Corollary A.2.2 gives:

$$P(u(\hat{\tau}\hat{\beta}) - u(\hat{\tau}\beta^*)) = \frac{\hat{\tau}}{\sqrt{8\pi(1 + \sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2.$$

Together, this gives:

$$\begin{aligned} \frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \leq P \left(b(\hat{\tau}\hat{\beta}) - b(\hat{\tau}\beta^*) + u(\hat{\tau}\hat{\beta}) - u(\hat{\tau}\beta^*) \right) \\ + \frac{1}{M} \sqrt{\frac{8}{\pi}}. \end{aligned}$$

If the term $1/M$ dominates the right-hand side, using the inequality of arithmetic and geometric means we get:

$$\frac{1}{\sqrt{8\pi}} \frac{\|\hat{\beta} - \beta^*\|_2}{(1+\sigma^2)^{1/4}} \leq \frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \leq \frac{2}{M} \sqrt{\frac{8}{\pi}}.$$

Additionally, since $\tau^* \gtrsim 1$, we have $\sigma \lesssim 1$ by Lemma 3.2.2. We conclude that the inequality (12) holds. If instead, the other terms dominate, we have the inequality:

$$\begin{aligned} \frac{1}{16\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{2\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \leq P \left(b(\hat{\tau}\hat{\beta}) - b(\hat{\tau}\beta^*) + u(\hat{\tau}\hat{\beta}) - u(\hat{\tau}\beta^*) \right) \\ \stackrel{(iv)}{\leq} (P - P_n) \left(b(\hat{\tau}\hat{\beta}) - b(\hat{\tau}\beta^*) + u(\hat{\tau}\hat{\beta}) - u(\hat{\tau}\beta^*) \right). \end{aligned}$$

In (iv) we used that the estimator $\hat{\tau}\hat{\beta}$ minimizes the empirical risk over a Euclidean ball of radius M , and $\|\hat{\tau}\beta^*\|_2 = \hat{\tau} \leq M$. We distinguish whether the bounded or the unbounded term dominates. Suppose first that the unbounded term dominates. Then, we use Proposition 5.3.1. We borrow its notation henceforth. Let $R := 2$, since $\|\hat{\beta} - \beta^*\|_2 \leq 2$, and $r := \mu^2/\sqrt{p}$. This choice of r is possible since if $\|\hat{\beta} - \beta^*\|_2 \leq \mu^2/\sqrt{p}$ holds then (12) is true. We note that $\rho(t') \lesssim \mu$ and $\Delta(t') \lesssim \mu^2$. On an event $A_{1,u}$ with probability at least $1 - \exp(-t)$,

$$\begin{aligned} \frac{1}{4\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \leq (P - P_n) \left(u(\hat{\beta}) - u(\beta^*) \right) = (P - P_n) \tilde{u}(\hat{\beta}) \\ \lesssim (\|\hat{\beta} - \beta^*\|_2 + \sigma) \left(\sqrt{(\|\hat{\beta} - \beta^*\|_2 + \sigma)\mu} + \mu \right) + \mu^2 \\ \stackrel{(v)}{\lesssim} \|\hat{\beta} - \beta^*\|_2 \sqrt{\mu\|\hat{\beta} - \beta^*\|_2} + \mu^2. \end{aligned}$$

FINITE SAMPLE RATES FOR LOGISTIC REGRESSION...

In (v) we assume that $\|\hat{\beta} - \beta^*\|_2 \geq \sigma \vee \mu$, otherwise (12) is true. Whichever of the two terms dominates in this expression (12) is true.

Next, suppose that the bounded term dominates. We will use Proposition 4.2.2 with $k = 1$. Taking $r = \mu$, if $\|\hat{\beta} - \beta^*\|_2 \leq \mu$, Eq. 12 holds. Else, there exist constants $c_1, c_2 > 0$, such that on an event $A_{1,b}$ with probability at least $1 - \exp(-t)$,

$$\begin{aligned} \frac{1}{4 \log 2} \left(\frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \right) &\leq (P - P_n) \frac{b(\hat{\tau}\hat{\beta}) - b(\hat{\tau}\beta^*)}{\log 2} \\ &\leq c_1 \frac{\|\hat{\beta} - \beta^*\|_2^2 \hat{\tau}}{\lambda} + c_2 \lambda \mu. \end{aligned}$$

Choosing $\lambda := 2c_1 \log(2) 4\sqrt{8\pi(1+\sigma^2)}$, we get:

$$\frac{1}{4 \log 2} \left(\frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{2\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \right) \lesssim \mu.$$

Using the inequality of arithmetic and geometric means, we get $\|\hat{\beta} - \beta^*\|_2 \lesssim \mu$, so Eq. 12 holds. This concludes case 1.

Case 2: $\sqrt{6} \leq \hat{\tau} \leq M/2$. As in case 1, we get:

$$\frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} \leq P \left(b(\hat{\tau}\hat{\beta}) - b(2\hat{\tau}\beta^*) \right),$$

as well as:

$$\begin{aligned} \frac{\hat{\tau}}{\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 &= P \left(u(\hat{\tau}\hat{\beta}) - u(\hat{\tau}\beta^*) \right) \\ &= P \left(u(\hat{\tau}\hat{\beta}) - u(2\hat{\tau}\beta^*) + u(\hat{\tau}\beta^*) \right) \stackrel{(i)}{\leq} P \left(u(\hat{\tau}\hat{\beta}) - u(2\hat{\tau}\beta^*) \right) + \frac{\hat{\tau}\sigma^2}{\sqrt{8\pi}}. \end{aligned}$$

Inequality (i) follows from Corollary A.2.2 and that $1 - \frac{1}{\sqrt{1+\sigma^2}} \leq \frac{\sigma^2}{2}$. Together, this gives:

$$\begin{aligned} \frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \\ \leq P \left(b(\hat{\tau}\hat{\beta}) - b(2\hat{\tau}\beta^*) + u(\hat{\tau}\hat{\beta}) - u(2\hat{\tau}\beta^*) \right) + \frac{\hat{\tau}\sigma^2}{\sqrt{8\pi}}. \end{aligned}$$

If the last summand on the right-hand side dominates, we conclude that Eqs. 12 and 13 hold. If not, we can use that $\hat{\tau}\hat{\beta}$ minimizes the empirical risk over a Euclidean ball of radius M and that $\|2\hat{\tau}\beta^*\|_2 \leq M$. We remark that this was not true in case 1. It follows that:

$$\begin{aligned} & \frac{1}{8\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \\ & \leq 2(P - P_n) \left(b(\hat{\tau}\hat{\beta}) - b(2\hat{\tau}\beta^*) + u(\hat{\tau}\hat{\beta}) - u(2\hat{\tau}\beta^*) \right). \end{aligned}$$

Again, we make a case distinction. If the unbounded term dominates, we try to control:

$$\begin{aligned} & \frac{1}{4 \cdot 8\sqrt{2\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{4\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \leq (P - P_n) \left(u(\hat{\tau}\hat{\beta}) - u(2\hat{\tau}\beta^*) \right) \\ & = \hat{\tau}(P - P_n)\tilde{u}(\hat{\beta}) - \hat{\tau}(P - P_n)u(\beta^*). \end{aligned}$$

Here, $\tilde{u}(\hat{\beta}) := u(\hat{\tau}\hat{\beta}) - u(\hat{\tau}\beta^*)$. If the right term dominates, we can directly use Bernstein's inequality (Boucheron et al., 2013, Theorem 2.10). With Corollary A.2.1, we see that Bernstein's condition is met with constants $K \lesssim \sigma$ and $\kappa \lesssim \sigma^{3/2}$. Therefore, by Bernstein's inequality, on an event $A_{2,u}$ with probability at least $1 - \exp(-t)$, it holds that:

$$-\hat{\tau}(P - P_n)u(\beta^*) \lesssim \sigma^{3/2}\sqrt{\mu} + \sigma\mu \leq \sigma^2 \vee \mu^2.$$

So, both Eqs. 12 and 13 hold. If instead the term in $\tilde{u}(\hat{\beta})$ dominates, we use Proposition 5.3.1 to bound it. So, as in case 1, we find that on the event $A_{1,u}$ which has probability at least $1 - \exp(-t)$,

$$\frac{\hat{\tau}}{4\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \lesssim \hat{\tau} \left(\|\hat{\beta} - \beta^*\|_2 \sqrt{\mu \|\hat{\beta} - \beta^*\|_2} + \mu^2 \right).$$

So, Eq. 12 holds. Consequentially, on the event $A_{1,u}$, we also have:

$$\frac{1}{4 \cdot 8\sqrt{2\pi}} \frac{1}{\hat{\tau}} \lesssim \hat{\tau} \left(\|\hat{\beta} - \beta^*\|_2 \sqrt{\mu \|\hat{\beta} - \beta^*\|_2} + \mu^2 \right) \lesssim \hat{\tau}\mu^2.$$

So, Eq. 13 is true.

Now suppose that the bounded term dominates. Then:

$$\frac{1}{4\sqrt{8\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{4\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \leq (P - P_n) \left(b(\hat{\tau}\hat{\beta}) - b(2\hat{\tau}\beta^*) \right).$$

As in case 1, we use Proposition 4.2.2 with $r := \mu$, although now with $k = 2$. There are absolute constants $c_1, c_2, c_3 > 0$, such that on the event $A_{2,b}$ which has probability at least $1 - \exp(-t)$,

$$(P - P_n) \frac{b(\hat{\tau}\hat{\beta}) - b(2\hat{\tau}\beta^*)}{\log 2} \leq c_1 \frac{\|\hat{\beta} - \beta^*\|_2^2 \hat{\tau}}{\lambda} + c_2 \frac{1}{\lambda \hat{\tau}} + c_3 \lambda \mu.$$

We choose $\lambda := 2 \max\{c_1 4\sqrt{8\pi(1+\sigma^2)}, c_2 4 \cdot 8\sqrt{2\pi}\}$, so that we get:

$$\frac{1}{8\sqrt{8\pi}} \frac{1}{\hat{\tau}} + \frac{\hat{\tau}}{8\sqrt{8\pi(1+\sigma^2)}} \|\hat{\beta} - \beta^*\|_2^2 \lesssim \mu.$$

Ignoring the second summand on the left-hand side, we see that Eq. 13 is met. By the inequality of arithmetic and geometric means, we see that Eq. 12 is met. This concludes case 2.

Case 3: $\hat{\tau} \leq \sqrt{6}$. Before we start with case 3, we show that assumption (11) implies:

$$Pl(\sqrt{6}\beta^*) - Pl(\min\{\tau^*, M\}\beta^*) > \sqrt{\frac{9t + 3p(2 + \log(3\sqrt{6n}))}{n}} \quad (14)$$

Let $\tau^\# := \min\{\tau^*, M\}$. From Lemma 3.2.2 and $\sigma \leq 1/\sqrt{7}$ we get $\tau^* \geq 2\sqrt{6}$. It follows that $\tau^\# \geq 2\sqrt{6}$, as we assumed $M \geq 2\sqrt{6}$. By Corollary A.2.2,

$$\begin{aligned} P \left(l(\sqrt{6}\beta^*) - l(\tau^\#\beta^*) \right) &= P \left(b(\sqrt{6}\beta^*) - b(\tau^\#\beta^*) \right) - \frac{\tau^\# - \sqrt{6}}{\sqrt{2\pi}} \left(1 - \frac{1}{\sqrt{1+\sigma^2}} \right). \\ &\geq P \left(b(\sqrt{6}\beta^*) - b(2\sqrt{6}\beta^*) \right) - \frac{\tau^* - \sqrt{6}}{\sqrt{2\pi}} \left(1 - \frac{1}{\sqrt{1+\sigma^2}} \right) \\ &\geq P \left(b(\sqrt{6}\beta^*) - b(2\sqrt{6}\beta^*) \right) - \frac{\tau^* - \sqrt{6}}{\sqrt{2\pi}} \frac{\sigma^2}{2} \\ &\stackrel{(i)}{\geq} P \left(b(\sqrt{6}\beta^*) - b(2\sqrt{6}\beta^*) \right) - \frac{\sqrt{2\pi}\sigma - 6\sigma^2}{\sqrt{8\pi}} \stackrel{(ii)}{\geq} 0.051. \end{aligned}$$

In (i) we used Lemma 3.2.2. Inequality (ii) follows from our bound on σ , the constant is found by simulation of the integrals. Consequentially, assumption (11) implies (14).

Now we start with case 3. First, by Lemma 3.2.3, $Pl(\hat{\gamma}) \geq Pl(\hat{\tau}\beta^*)$. Moreover, since $\hat{\tau} \leq 6$, by convexity $Pl(\hat{\tau}\beta^*) \geq Pl(\sqrt{6}\beta^*)$. Therefore,

$$P\left(l(\hat{\gamma}) - l(\tau^\sharp\beta^*)\right) \geq P\left(l(\sqrt{6}\beta^*) - l(\tau^\sharp\beta^*)\right) \stackrel{(iii)}{\geq} \sqrt{\frac{9t + 3p(2 + \log(3\sqrt{6n}))}{n}}.$$

Inequality (iii) follows from Eq. 14. By Lemma 6.2.1, this event together with $\hat{\tau} \leq 6$ occurs on an event A_3 , which has probability at most $2 \exp(-t)$. This concludes case 3.

We conclude that Eqs. 12 and 13 are true on the event $A_{1,u} \cap A_{1,b} \cap A_{2,u} \cap A_{2,b} \cap A_3$, which has probability at least $1 - 6 \exp(-t)$. The proof is complete. \square

The following result was used in the proof of case 3 in Theorem 6.2.1. There, we use it with γ^\sharp as $M\beta^*$ or $\tau^*\beta^*$. The idea is, that two events are unlikely to occur together. The first one states that $\hat{\tau} \leq r$, and the second one states that the excess risk is larger than $\sqrt{(t + p \log(n))/n}$ up to constant factors. In Theorem 6.2.1, we assume that this lower bound is small, implying that the second statement holds. Consequentially, $\hat{\tau} > r$ must be true with high probability.

Lemma 6.2.1 *Fix $r > 0$, and $\gamma^\sharp \in \mathbb{R}^p \setminus \{0\}$ such that $\|\gamma^\sharp\|_2 \leq M$. Let A be the event that $\|\hat{\gamma}\|_2 \leq r$. Moreover, let B be the event that:*

$$P\left(l(\hat{\gamma}) - l(\gamma^\sharp)\right) > \sqrt{\frac{9t + 3p(2 + \log(3r\sqrt{n}))}{n}}.$$

Then, $\mathbb{P}[A \cap B] \leq 2 \exp(-t)$.

Proof Let B_r be the Euclidean ball of radius r centered at zero. We start by bounding the following expression with high probability:

$$\sup_{\gamma \in B_r} (P - P_n)\left(l(\gamma) - l(\gamma^\sharp)\right).$$

The ball B_r can be covered with at most $(3r/\epsilon)^p$ balls of radius $\epsilon \in (0, 1)$. Let $\mathcal{N} \subset B_r$ be such a covering.

$$\sup_{\gamma \in B_r} (P - P_n)\left(l(\gamma) - l(\gamma^\sharp)\right)$$

$$\begin{aligned}
&= \inf_{\gamma' \in \mathcal{N}} \sup_{\gamma \in B_r} (P - P_n)(l(\gamma) - l(\gamma')) + (P - P_n)(l(\gamma') - l(\gamma^\sharp)) \\
&\leq \inf_{\gamma' \in \mathcal{N}} \sup_{\gamma \in B_r} (P - P_n)(l(\gamma) - l(\gamma')) + \sup_{\gamma' \in \mathcal{N}} (P - P_n)(l(\gamma') - l(\gamma^\sharp)). \quad (15)
\end{aligned}$$

To bound the first summand, we proceed as in the proof of Lemma 5.1.2.

$$\begin{aligned}
\inf_{\gamma' \in \mathcal{N}} \sup_{\gamma \in B_r} \left| (P - P_n)\|x\|_2 \frac{l(\gamma) - l(\gamma')}{\|x\|_2} \right| &\leq \inf_{\gamma' \in \mathcal{N}} \sup_{\gamma \in B_r} (P + P_n)\|x\|_2 \left| \frac{l(\gamma) - l(\gamma')}{\|x\|_2} \right| \\
&\stackrel{(i)}{\leq} \epsilon (P + P_n)\|x\|_2 \leq \epsilon ((P_n - P)\|x\|_2 + 2P\|x\|_2).
\end{aligned}$$

In (i), we used that the function $l(\cdot, yx) : \mathbb{R}^p \rightarrow \mathbb{R}$, mapping $\gamma \mapsto l(\gamma, yx)$ is $\|x\|_2$ -Lipschitz. By concentration for 1-Lipschitz functions of standard normally distributed random variables (see e.g. Boucheron et al. 2013, Theorem 5.6), with probability at most $\exp(-t)$, it holds that $(P_n - P)\|x\|_2 > \sqrt{2t/n}$. By Jensen's inequality, $P\|x\|_2 \leq \sqrt{p}$. We conclude that with probability at least $1 - \exp(-t)$,

$$\inf_{\gamma' \in \mathcal{N}} \sup_{\gamma \in B_r} \left| (P - P_n)\|x\|_2 \frac{l(\gamma) - l(\gamma')}{\|x\|_2} \right| \leq \epsilon \left(\sqrt{\frac{2t}{n}} + 2\sqrt{p} \right).$$

For the second term in Eq. 15, we use that the function $\mathbb{R}^p \rightarrow \mathbb{R}$, mapping $\gamma \mapsto l(\gamma, yx) - l(\gamma^\sharp)$ takes values in $(-\log(2), \log(2))$. So, by Hoeffding's inequality (see e.g. Boucheron et al. 2013, Theorem 2.8), for any $\gamma' \in \mathcal{N}$,

$$\mathbb{P} \left[(P - P_n)(l(\gamma') - l(\gamma^\sharp)) > \sqrt{\frac{t}{n}} \right] \leq \exp \left(-\frac{t}{2 \log(2)^2} \right) \leq \exp(-t).$$

Moreover, by a union bound,

$$\mathbb{P} \left[\max_{\gamma' \in \mathcal{N}} (P - P_n)(l(\gamma') - l(\gamma^\sharp)) > \sqrt{\frac{t + \log |\mathcal{N}|}{n}} \right] \leq \exp(-t).$$

Moreover, by choice of \mathcal{N} , it holds that $\log |\mathcal{N}| \leq p \log(3r/\epsilon)$. Altogether, we find that with probability at least $1 - 2 \exp(-t)$,

$$\sup_{\gamma \in B_r} (P - P_n) \left(l(\gamma) - l(\gamma^\sharp) \right) \leq \epsilon \left(\sqrt{\frac{2t}{n}} + 2\sqrt{p} \right) + \sqrt{\frac{t + p \log(3r/\epsilon)}{n}}.$$

Choosing $\epsilon := 1/\sqrt{n}$ and Jensen's inequality gives:

$$\begin{aligned} \sup_{\gamma \in B_r} (P - P_n) \left(l(\gamma) - l(\gamma^\sharp) \right) &\leq \frac{\sqrt{2t}}{n} + 2\sqrt{\frac{p}{n}} + \sqrt{\frac{t + p \log(3r\sqrt{n})}{n}} \\ &\leq \sqrt{\frac{9t + 3p(2 + \log(3r\sqrt{n}))}{n}}. \end{aligned}$$

Denote this event by C . So, we concluded that $\mathbb{P}[C] \geq 1 - 2 \exp(-t)$. Now suppose that $A \cap B$ occurs. Since B occurs,

$$\sqrt{\frac{9t + 3p(2 + \log(3r\sqrt{n}))}{n}} < P \left(l(\hat{\gamma}) - l(\gamma^\sharp) \right).$$

Since $\|\gamma^\sharp\|_2 \leq M$, $P_n l(\hat{\gamma}) \leq P_n l(\gamma^\sharp)$. Consequentially,

$$P \left(l(\hat{\gamma}) - l(\gamma^\sharp) \right) \leq (P - P_n) \left(l(\hat{\gamma}) - l(\gamma^\sharp) \right).$$

Since A occurs, $\hat{\tau} \leq r$. So:

$$\leq \sup_{\gamma \in B_r} (P - P_n) \left(l(\gamma) - l(\gamma^\sharp) \right).$$

We conclude that $A \cap B$ implies C^c . Consequentially, $\mathbb{P}[A \cap B] \leq 1 - P[C] \leq 2 \exp(-t)$. The proof is complete. \square

6.3 Proof of the Consequence for Linear Separation Here, we provide a proof of Proposition 2.3.1.

Proof of Proposition 2.3.1 Let A_M be the event with probability at least $1 - 4 \exp(-t)$ on which the conclusion in Theorem 2.1.1 holds, for a valid choice of M . Let B be the event that the data are linearly separable. We show that if M is chosen large enough, the event $A_M \cap B$ is empty. Hence, $\mathbb{P}[B] \leq \mathbb{P}[A_M^c] \leq 4 \exp(-t)$. We start with the observation that on B , it holds that $\hat{\tau} = M$. So, on B , by a triangular inequality,

$$d_*(\hat{\gamma}) \geq \frac{M - \tau^*}{\tau^{*3/2}}.$$

By Theorem 2.1.1, on A_M the quantity $d_*(\hat{\gamma})$ is upper bounded by a deterministic expression $c > 0$ not depending on M . So, if $M > c\tau^{*3/2} + \tau^*$, we reach a contradiction, meaning that the event $A_M \cap B$ is empty. The proof is complete. \square

A Approximations for the (un)bounded Term

A.1 Approximations for Bounded Term In this subsection, we provide approximations for several quantities related to the bounded term.

A.1.1 Gaussian Integrals of Exponential Functions

Lemma A.1.1 *Let $\tau > 0$ and $z \sim \mathcal{N}(0, 1)$. We have:*

$$\sqrt{\frac{2}{\pi}} \left(\frac{1}{\tau} - \frac{1}{\tau^3} \right) \leq P \exp(-\tau|z|) \leq \sqrt{\frac{2}{\pi}} \frac{1}{\tau}.$$

The result follows from completing the square and repeated integration by parts.

Lemma A.1.2 *Let $z \sim \mathcal{N}(0, 1)$. Then, for $\tau > 0$,*

$$\sqrt{\frac{1}{2\pi}} \frac{1}{\tau^2} \left(1 - \frac{3}{\tau^2} \right) \leq P \frac{|z|}{1 + \exp(\tau|z|)} \leq \sqrt{\frac{2}{\pi}} \frac{1}{\tau^2}.$$

Proof Since $\tau > 0$,

$$P \frac{|z|}{2 \exp(\tau|z|)} \leq P \frac{|z|}{1 + \exp(\tau|z|)} \leq P \frac{|z|}{\exp(\tau|z|)}.$$

We continue working with the expression for the lower bound. The derivation for the upper bound is the same, with a factor 2. It holds that:

$$\begin{aligned} P \frac{|z|}{2 \exp(\tau|z|)} &= 2 \int_0^\infty \frac{z}{2\sqrt{2\pi}} \exp\left(-\tau z - \frac{z^2}{2}\right) dz \\ &= \sqrt{\frac{1}{2\pi}} e^{\tau^2/2} \left(\int_0^\infty (z + \tau) \exp\left(-\frac{(z + \tau)^2}{2}\right) dz - \int_0^\infty \tau \exp\left(-\frac{(z + \tau)^2}{2}\right) dz \right). \\ &= \sqrt{\frac{1}{2\pi}} \left(1 - e^{\tau^2/2} \tau \int_\tau^\infty \exp\left(-\frac{z^2}{2}\right) dz \right) \end{aligned} \quad (16)$$

Repeated integration by parts shows:

$$\frac{1}{\tau} - \frac{1}{\tau^3} \leq e^{\tau^2/2} \int_\tau^\infty \exp\left(-\frac{z^2}{2}\right) dz \leq \frac{1}{\tau} - \frac{1}{\tau^3} + \frac{3}{\tau^5}. \quad (17)$$

Inserting these bounds in Eq. 16 completes the proof.

Lemma A.1.3 Let $z \sim \mathcal{N}(0, 1)$. Then, for $\tau > 0$,

$$\sqrt{\frac{2}{\pi}} \left(\frac{2}{\tau^3} - \frac{12}{\tau^5} - \frac{15}{\tau^7} \right) \leq P \frac{|z|^2}{\exp(\tau|z|)} \leq \sqrt{\frac{2}{\pi}} \left(\frac{2}{\tau^3} + \frac{3}{\tau^5} \right).$$

The lower bound is non-zero if $\tau > \sqrt{3 + \sqrt{33/2}} \approx 2.657$.

Proof We decompose the integral into two parts, with I_1 and I_2 defined below.

$$P \frac{|z|^2}{\exp(\tau|z|)} = \sqrt{\frac{2}{\pi}} \int_0^\infty z^2 \exp\left(-\tau z - \frac{z^2}{2}\right) dz =: \sqrt{\frac{2}{\pi}} (I_1 - I_2).$$

To evaluate the first integral I_1 , we use integration by parts.

$$\begin{aligned} I_1 &:= e^{\tau^2/2} \int_0^\infty (z + \tau)^2 \exp\left(-\frac{(z + \tau)^2}{2}\right) dz \\ &= e^{\tau^2/2} \left(\left[(z + \tau) \left(-\exp\left(-\frac{(z + \tau)^2}{2}\right) \right) \right]_0^\infty + \int_0^\infty \exp\left(-\frac{(z + \tau)^2}{2}\right) dz \right) \\ &= \tau + e^{\tau^2/2} \int_\tau^\infty \exp\left(-\frac{z^2}{2}\right) dz. \end{aligned}$$

The second integral is:

$$\begin{aligned} I_2 &:= e^{\tau^2/2} \int_0^\infty (2z\tau + \tau^2) \exp\left(-\frac{(z + \tau)^2}{2}\right) dz \\ &= 2\tau \left(1 - \tau e^{\tau^2/2} \int_\tau^\infty \exp\left(-\frac{z^2}{2}\right) dz \right) + \tau^2 e^{\tau^2/2} \int_\tau^\infty \exp\left(-\frac{z^2}{2}\right) dz \\ &= 2\tau - \tau^2 e^{\tau^2/2} \int_\tau^\infty \exp\left(-\frac{z^2}{2}\right) dz. \end{aligned}$$

We conclude that:

$$P \frac{|z|^2}{\exp(\tau|z|)} = \sqrt{\frac{2}{\pi}} (I_1 - I_2) = \sqrt{\frac{2}{\pi}} \left(e^{\tau^2/2} \int_\tau^\infty \exp\left(-\frac{z^2}{2}\right) dz (1 + \tau^2) - \tau \right).$$

Repeated integration by parts shows:

$$\frac{1}{t} - \frac{1}{t^3} + \frac{3}{t^5} - \frac{15}{t^7} \leq \exp\left(\frac{t^2}{2}\right) \int_t^\infty \exp\left(-\frac{x^2}{2}\right) dx \leq \frac{1}{t} - \frac{1}{t^3} + \frac{3}{t^5}.$$

Inserting this concludes the proof. \square

A.1.2 A First-Order Bound

Lemma A.1.4 *Suppose that $x \sim \mathcal{N}(0, I)$ and $\gamma, \gamma' \in \mathbb{R}^p \setminus \{0\}$. Moreover, define $\rho := \gamma^T \gamma' / (\|\gamma\|_2 \|\gamma'\|_2)$. Then:*

$$\begin{aligned} & \sqrt{P \frac{|x^T(\gamma - \gamma')|^2}{\exp(2|\gamma'^T x|)}} \\ & \leq \|\gamma\|_2 \sqrt{\frac{1}{\sqrt{2\pi}} \frac{(1 - \rho^2)}{\|\gamma'\|_2}} + |\rho\|\gamma\|_2 - \|\gamma'\|_2| \sqrt{\frac{2}{\pi} \left(\frac{1}{4} \frac{1}{\|\gamma'\|_2^3} + \frac{3}{32} \frac{1}{\|\gamma'\|_2^5} \right)}. \end{aligned}$$

If $\|\gamma'\|_2 \geq 1$, we can bound this term as:

$$\sqrt{P \frac{|x^T(\gamma - \gamma')|^2}{\exp(2|\gamma'^T x|)}} \leq \|\gamma\|_2 \sqrt{\frac{1}{\sqrt{2\pi}} \frac{(1 - \rho^2)}{\|\gamma'\|_2}} + \sqrt{\frac{11}{32}} \sqrt{\frac{2}{\pi}} \frac{|\rho\|\gamma\|_2 - \|\gamma'\|_2|}{\|\gamma'\|_2^{3/2}}.$$

In addition, suppose that $\|\gamma\|_2 = \|\gamma'\|_2 = \tau$, and define $\beta := \gamma/\tau$ as well as $\beta' := \gamma'/\tau$. Since $1 - \rho^2 \leq 2(1 - \rho) = \|\beta - \beta'\|_2^2$, we have:

$$\sqrt{P \frac{|x^T(\gamma - \gamma')|^2}{\exp(2|\gamma'^T x|)}} \leq \sqrt{\frac{\tau}{\sqrt{2\pi}}} \|\beta - \beta'\|_2 + \sqrt{\frac{11}{32}} \sqrt{\frac{2}{\pi}} \frac{\|\beta - \beta'\|_2^2}{\sqrt{\tau}}.$$

Proof Let $\Pi_{\gamma'} := \gamma' \gamma'^T / \|\gamma'\|_2^2$ be the projection onto the span of γ' . Then,

$$x^T(\gamma - \gamma') = x^T((I - \Pi_{\gamma'})\gamma + \Pi_{\gamma'}(\gamma - \gamma')).$$

By the triangular inequality,

$$\sqrt{P \frac{|x^T(\gamma - \gamma')|^2}{\exp(2|\gamma'^T x|)}} \leq \sqrt{P \frac{|x^T(I - \Pi_{\gamma'})\gamma|^2}{\exp(2|\gamma'^T x|)}} + \sqrt{P \frac{|x^T \Pi_{\gamma'}(\gamma - \gamma')|^2}{\exp(2|\gamma'^T x|)}}.$$

We start with the first term on the right-hand side. The random variables $\gamma'^T x$ and $x^T(I - \Pi_{\gamma'})\gamma$ are uncorrelated. Since x is Gaussian, they are independent. So:

$$\sqrt{P \frac{|x^T(I - \Pi_{\gamma'})\gamma|^2}{\exp(2|\gamma'^T x|)}} = \sqrt{P \exp(-2|\gamma'^T x|) P |x^T(I - \Pi_{\gamma'})\gamma|^2}.$$

By Lemma A.1.1, $P \exp(-2|\gamma'^T x|) \leq 1/(\sqrt{2\pi}\|\gamma'\|_2)$. Moreover, by definition of ρ , we have: $P|x^T(I - \Pi_{\gamma'})\gamma|^2 = \|\gamma\|_2^2(1 - \rho^2)$. So, the first term is bounded by:

$$\sqrt{P \exp(-2|\gamma'^T x|)P|x^T(I - \Pi_{\gamma'})\gamma|^2} \leq \|\gamma\|_2 \sqrt{\frac{1}{\sqrt{2\pi}\|\gamma'\|_2}(1 - \rho^2)}.$$

The second term can be bounded as follows:

$$\begin{aligned} \sqrt{P \frac{|x^T \Pi_{\gamma'}(\gamma - \gamma')|^2}{\exp(2|\gamma'^T x|)}} &= |\rho\|\gamma\|_2 - \|\gamma'\|_2| \sqrt{P \frac{(|x^T \gamma'|/\|\gamma'\|_2)^2}{\exp(2|\gamma'^T x|)}} \\ &\stackrel{(i)}{\leq} |\rho\|\gamma\|_2 - \|\gamma'\|_2| \sqrt{\sqrt{\frac{2}{\pi}} \left(\frac{2}{(2\|\gamma'\|_2)^3} + \frac{3}{(2\|\gamma'\|_2)^5} \right)} \\ &= |\rho\|\gamma\|_2 - \|\gamma'\|_2| \sqrt{\sqrt{\frac{2}{\pi}} \left(\frac{1}{4} \frac{1}{\|\gamma'\|_2^3} + \frac{3}{32} \frac{1}{\|\gamma'\|_2^5} \right)}. \end{aligned}$$

In (i), we used Lemma A.1.3. The proof is complete. \square

A.1.3 A Second-Order Bound

Lemma A.1.5 *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be the mapping $\tau \mapsto P \log(1 + \exp(-\tau|z|))$, where $z \sim \mathcal{N}(0, 1)$. Suppose that for some $\kappa > \sqrt{3 + \sqrt{33/2}}$, we have $\min\{\tau, \tau^*\} \geq \kappa$. Then, for all $\bar{\tau} \in [\tau, \tau^*] \cup [\tau^*, \tau]$,*

$$\frac{1}{\sqrt{8\pi}} \left(2 - \frac{12}{\kappa^2} - \frac{15}{\kappa^4} \right) \frac{1}{\bar{\tau}^3} \leq \ddot{f}(\bar{\tau}).$$

In particular, if $\min\{\tau, \tau^*\} \geq \sqrt{6 + \sqrt{51}}$,

$$\sqrt{\frac{1}{8\pi}} \frac{1}{\bar{\tau}^3} \leq \ddot{f}(\bar{\tau}).$$

Proof Using Lebesgue's dominated convergence theorem, we find for any $t > 0$,

$$\ddot{f}(t) = P \frac{z^2 \exp(t|z|)}{(1 + \exp(t|z|))^2} \geq P \frac{z^2}{4 \exp(t|z|)}.$$

Since $\bar{\tau}$ is an intermediate point of τ and τ^* , it holds that $\bar{\tau} \geq \kappa$. Consequently, by Lemma A.1.3,

$$\sqrt{\frac{2}{\pi}} \left(2 - \frac{12}{\kappa^2} - \frac{15}{\kappa^4} \right) \frac{1}{\bar{\tau}^3} \leq P \frac{|z|^2}{\exp(\bar{\tau}|z|)}.$$

The proof is complete. □

A.1.4 Differences Between Bounded Terms

Lemma A.1.6 *Let $k > 1$, and $z \sim \mathcal{N}(0, 1)$. Then,*

$$\frac{1}{\sqrt{2\pi}} \frac{k-1}{k^2\tau} \left(1 - \frac{3}{\tau^2} \right) \leq P \log \left(\frac{1 + e^{-\tau|z|}}{1 + e^{-k\tau|z|}} \right) \leq \frac{k-1}{\tau} \sqrt{\frac{2}{\pi}}.$$

Additionally, if $\tau \leq l$,

$$\frac{1}{\sqrt{2\pi}} \frac{\tau(k-1)}{k^2l^2} \left(1 - \frac{3}{k^2l^2} \right) \leq P \log \left(\frac{1 + e^{-\tau|z|}}{1 + e^{-k\tau|z|}} \right) \leq \frac{\tau(k-1)}{\sqrt{2\pi}}.$$

For $\tau \leq \sqrt{3}$, the first lower bound is void. In this case, we can use the second lower bound, with a large enough bound l , to make the bound non-trivial. In the first lower bound, one could assume $\tau > \sqrt{6}$, to get a nicer lower bound.

Proof Let $h_z : (0, \infty) \rightarrow (0, \log 2)$ be the mapping $\tau \mapsto \log(1 + \exp(-\tau|z|))$. By the mean value theorem, for some (random) $t_z \in [\tau, k\tau]$,

$$Ph_z(\tau) - h_z(k\tau) = \tau(1-k)Ph'_z(t_z) = \tau(k-1)P \frac{|z|}{1 + \exp(t_z|z|)}.$$

For the upper bound, by Lemma A.1.2:

$$\tau(k-1)P \frac{|z|}{1 + \exp(t_z|z|)} \leq \tau(k-1)P \frac{|z|}{1 + \exp(\tau|z|)} \leq \sqrt{\frac{2}{\pi}} \frac{k-1}{\tau}.$$

On the other hand, since $\tau|z| > 0$,

$$\tau(k-1)P \frac{|z|}{1 + \exp(\tau|z|)} \leq \tau(k-1)P \frac{|z|}{2} \leq \frac{\tau(k-1)}{\sqrt{2\pi}}.$$

For the first lower bound, by Lemma A.1.2:

$$\tau(k-1)P \frac{|z|}{1 + \exp(t_z|z|)} \geq \tau(k-1)P \frac{|z|}{1 + \exp(k\tau|z|)} \geq \frac{1}{\sqrt{2\pi}} \frac{k-1}{k^2\tau} \left(1 - \frac{3}{k^2\tau^2} \right).$$

For the second lower bound, use that since $\tau \leq l$, $t_z \leq kl$. Using Lemma A.1.2,

$$\begin{aligned} \tau(k-1)P \frac{|z|}{1 + \exp(t_z|z|)} &\geq \tau(k-1)P \frac{|z|}{1 + \exp(kl|z|)} \\ &\geq \frac{1}{\sqrt{2\pi}} \frac{\tau(k-1)}{k^2 l^2} \left(1 - \frac{3}{k^2 l^2}\right). \end{aligned}$$

□

Lemma A.1.7 *Let $k > 1$, and $z \sim \mathcal{N}(0, 1)$. Then,*

$$P \left(\log \frac{1 + e^{-\tau|z|}}{1 + e^{-k\tau|z|}} \right)^2 \leq (k-1)^2 \min \left\{ \sqrt{\frac{1}{2^5\pi}} \left(\frac{2}{\tau} + \frac{3}{4\tau^3} \right), \frac{\tau^2}{4} \right\}.$$

Additionally, if $\tau \leq l$,

$$\tau^2(k-1)^2 \frac{1}{\sqrt{2^7\pi}} \left(\frac{1}{l^3} - \frac{3}{2l^5} - \frac{15}{2^5 l^7} \right) \leq P \log \left(\frac{1 + e^{-\tau|z|}}{1 + e^{-k\tau|z|}} \right)^2.$$

Note that in the lower bound, we may simply take $l = \tau$ to match the upper bound if τ is large.

Proof Let $h_z : (0, \infty) \rightarrow (0, \log 2)$ be the mapping $\tau \mapsto \log(1 + \exp(-\tau|z|))$. By the mean value theorem, for some $t_z \in [\tau, k\tau]$,

$$P(h_z(\tau) - h_z(k\tau))^2 = \tau^2(1-k)^2 P h'_z(t_z)^2 = \tau^2(k-1)^2 P \frac{|z|^2}{(1 + \exp(\tau|z|))^2}.$$

For the upper bound, by Lemma A.1.3:

$$\begin{aligned} \tau^2(k-1)^2 P \frac{|z|^2}{(1 + \exp(\tau|z|))^2} &\leq \tau^2(k-1)^2 P \frac{|z|^2}{\exp(2\tau|z|)} \\ &\leq \tau^2(k-1)^2 \sqrt{\frac{2}{\pi}} \left(\frac{2}{2^3\tau^3} + \frac{3}{2^5\tau^5} \right) \leq \sqrt{\frac{1}{2^5\pi}} (k-1)^2 \left(\frac{2}{\tau} + \frac{3}{4\tau^3} \right). \end{aligned}$$

On the other hand, since $\tau|z| > 0$,

$$\tau^2(k-1)^2 P \frac{|z|^2}{(1 + \exp(\tau|z|))^2} \leq \tau^2(k-1)^2 P \frac{|z|^2}{2^2} \leq \frac{\tau^2(k-1)^2}{4}.$$

For the lower bound, since $\tau \leq l$,

$$\begin{aligned} \tau^2(k-1)^2 P \frac{|z|^2}{(1 + \exp(\tau|z|))^2} &\geq \tau^2(k-1)^2 P \frac{|z|^2}{2^2 \exp(2l|z|)} \\ &\stackrel{(i)}{\geq} \tau^2(k-1)^2 \frac{1}{\sqrt{2^7 \pi}} \left(\frac{1}{l^3} - \frac{3}{2l^5} - \frac{15}{2^5 l^7} \right). \end{aligned}$$

Inequality (i) follows from Lemma A.1.3. The proof is complete. \square

Lemma A.1.8 Fix $\beta, \beta' \in S^{p-1}$ and $\tau \geq 1$. Then,

$$P \left(\log \frac{1 + \exp(-\tau|x^T \beta|)}{1 + \exp(-\tau|x^T \beta'|)} \right)^2 \leq 2 \|\beta - \beta'\|_2^2 \left(\sqrt{\frac{\tau}{\sqrt{2\pi}}} + \sqrt{\frac{11}{8}} \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\tau}} \right)^2.$$

Proof Let $h : (0, \infty) \rightarrow (0, \log 2)$ be the mapping $z \mapsto \log(1 + \exp(-\tau|z|))$. There exists a $|z| \in [\min\{|x^T \beta|, |x^T \beta'|\}, \max\{|x^T \beta|, |x^T \beta'|\}]$, such that:

$$\begin{aligned} P (h(|x^T \beta|) - h(|x^T \beta'|))^2 &= P \left(|x^T(\beta - \beta')| \frac{\tau}{1 + \exp(\tau|z|)} \right)^2 \\ &\stackrel{(i)}{\leq} \tau^2 \left(P \frac{|x^T(\beta - \beta')|^2}{\exp(2\tau|x^T \beta|)} + P \frac{|x^T(\beta - \beta')|^2}{\exp(2\tau|x^T \beta'|)} \right) = 2P \frac{|x^T(\tau\beta - \tau\beta')|^2}{\exp(2\tau|x^T \beta|)}. \end{aligned}$$

In (i), we used that the derivative of a convex function is increasing, so it reaches its maximum on the boundary of its domain. By Lemma A.1.4, the right-hand-side is bounded from above by:

$$\begin{aligned} &2 \left(\sqrt{\tau} \sqrt{\frac{1}{\sqrt{2\pi}}} \|\beta - \beta'\|_2 + \sqrt{\frac{11}{32}} \sqrt{\frac{2}{\pi}} \frac{\|\beta - \beta'\|_2^2}{\sqrt{\tau}} \right)^2 \\ &= 2 \|\beta - \beta'\|_2^2 \left(\sqrt{\frac{\tau}{\sqrt{2\pi}}} + 2 \sqrt{\frac{11}{32}} \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\tau}} \right)^2. \end{aligned}$$

\square

A.2 Approximations and Identities for Unbounded Term Recall the definition of the unbounded term $u(\gamma, x, y) = |x^T \gamma| \mathbf{1}\{y x^T \gamma < 0\}$. Here we derive several results to control u . The Gaussian distribution proves very handy here, as we can obtain sharp upper and lower bounds, in some cases even identities. As u is positivey homogeneous, for all $\beta \in S^{p-1}$ and $\tau > 0$, it holds that $u(\tau\beta) = \tau u(\beta)$. Consequentially, we restrict our attention to the domain S^{p-1} .

A.2.1 A Trigonometric Argument and its Consequences The following identity is the core of our analysis of the unbounded term. An approximation of its right-hand side will serve us to derive the parameters for Bernstein's inequality. The case $m = 1$ will form the basis of lower and upper bounds for the excess risk. The integral of the sine function raised to the m -th power can be calculated using some trigonometry and basic integration techniques. However, we will only calculate the case $m = 1$, while for $m > 1$, we will use the bound $\sin(x) \leq x$ for $x \geq 0$.

Proposition A.2.1 *Let $x \sim \mathcal{N}(0, I)$ and $\beta, \beta' \in S^{p-1}$. Let $m \in \mathbb{N}$. Then,*

$$P|x^T \beta|^m \mathbf{1}\{\beta^T x x^T \beta' < 0\} = 2^{1+m/2} \Gamma(1 + m/2) \int_0^{\arccos(\beta^T \beta')} \sin(\alpha)^m \frac{d\alpha}{2\pi}.$$

Remark A.2.1 (Non-identity covariance matrix) We prove Proposition A.2.1 and related results for $x \sim \mathcal{N}(0, I)$. Yet, the conversion to $x \sim \mathcal{N}(0, \Sigma)$ with $\Sigma \neq I$ is straightforward: Since Σ is a covariance matrix, it is positive definite. Consequentially, there exists a unique positive definite matrix $\sqrt{\Sigma}$, such that $\Sigma = \sqrt{\Sigma} \sqrt{\Sigma}$. So, $x = \sqrt{\Sigma} z$, where $z \sim \mathcal{N}(0, I)$. Therefore, $x^T \beta = z^T (\sqrt{\Sigma} \beta)$. So, Proposition A.2.1 gives us:

$$\begin{aligned} & P|x^T \beta|^m \mathbf{1}\{\beta^T x x^T \beta' < 0\} \\ &= (\beta^T \Sigma \beta)^{m/2} 2^{1+m/2} \Gamma(1 + m/2) \int_0^{\arccos(\beta^T \Sigma \beta')} \sin(\alpha)^m \frac{d\alpha}{2\pi}. \end{aligned}$$

Note that the expression is no longer symmetric in β, β' , but for specific choices of Σ .

Proof There are two special cases: $\beta \in \{\beta', -\beta'\}$. They follow from continuity or separate verification. Excluding those two cases, β, β' are in general position. We show that we may take $p = 2$. Let $\Pi : \mathbb{R}^p \rightarrow \mathbb{R}^p$ be the orthogonal projection of \mathbb{R}^p onto the span of β and β' . Then:

$$P|x^T \beta|^m \mathbf{1}\{\beta^T x x^T \beta' < 0\} = P|x^T (\Pi\beta)|^m \mathbf{1}\{(\Pi\beta)^T x x^T (\Pi\beta') < 0\}$$

$$= P|(\Pi x)^T \beta|^m \mathbf{1}\{\beta^T (\Pi x) (\Pi x)^T \beta' < 0\}.$$

Note that Πx , restricted to the span of β and β' , has a standard Gaussian distribution on a 2-dimensional linear subspace. Consequentially, we may take $p = 2$, as claimed.

Now we assume $p = 2$. Let $\tilde{x} := x/\|x\|_2$. Then:

$$P|x^T \beta|^m \mathbf{1}\{\beta^T x x^T \beta' < 0\} = P\|x\|_2^m |\tilde{x}^T \beta|^m \mathbf{1}\{\beta^T \tilde{x} \tilde{x}^T \beta' < 0\}.$$

Since x is Gaussian, $\|x\|_2$ is independent of \tilde{x} . Moreover, since x is Gaussian, $\|x\|_2$ follows a χ -distribution, with $p = 2$ degrees of freedom. Therefore, $P\|x\|_2^m = 2^{m/2} \Gamma(1 + m/2)$. Consequentially:

$$= 2^{m/2} \Gamma(1 + m/2) P|\tilde{x}^T \beta|^m \mathbf{1}\{\beta^T \tilde{x} \tilde{x}^T \beta' < 0\}.$$

Note that \tilde{x} is uniformly distributed on S^1 . By symmetry, $\mathbb{P}[\beta^T \tilde{x} \tilde{x}^T \beta' < 0] = 2\mathbb{P}[\beta^T \tilde{x} > 0 \cap \tilde{x}^T \beta' < 0]$. So,

$$P|\tilde{x}^T \beta|^m \mathbf{1}\{\beta^T \tilde{x} \tilde{x}^T \beta' < 0\} = 2P|\tilde{x}^T \beta|^m \mathbf{1}\{\beta^T \tilde{x} > 0 \cap \tilde{x}^T \beta' < 0\}.$$

The random variable \tilde{x} follows a uniform distribution on S^1 . Consequentially, by rotational invariance, we may assume that $\beta = e_1$.

$$P|\tilde{x}^T \beta|^m \mathbf{1}\{\beta^T \tilde{x} > 0 \cap \tilde{x}^T \beta' < 0\} = P|\tilde{x}_1|^m \{\tilde{x}_1 > 0 \cap \tilde{x}^T \beta' < 0\}.$$

Let $\mathcal{C} := \{z \in S^1 : z^T \beta > 0 \cap z^T \beta' < 0\}$. Note that the kernel of a non-zero vector in \mathbb{R}^2 is spanned by a 90-degree rotation R of the vector. So, the boundaries of the cap $\mathcal{C} := \{z \in S^1 : z^T \beta > 0 \cap z^T \beta' < 0\}$ are given by $R\beta$ and $R\beta'$. Since angles are preserved under rotation, the arc length of \mathcal{C} is $\arccos((R\beta)^T (R\beta')) = \arccos(\beta^T \beta')$. Hence:

$$P|\tilde{x}_1|^m \{\tilde{x}_1 > 0 \cap \tilde{x}^T \beta' < 0\} = \int_0^{\arccos(\beta^T \beta')} \sin(\alpha)^m \frac{d\alpha}{2\pi}.$$

This proves the result for β, β' in general position. The proof is complete. \square

As a consequence of Proposition A.2.1, we obtain the following identity.

Corollary A.2.1 *Let $x \sim \mathcal{N}(0, I)$ and $\beta, \beta' \in S^{p-1}$. Then,*

$$P|x^T \beta| \mathbf{1}\{\beta^T x x^T \beta' < 0\} = \frac{\|\beta - \beta'\|_2^2}{\sqrt{8\pi}}.$$

Moreover,

$$P|x^T \beta|^m \mathbf{1}\{\beta^T x x^T \beta' < 0\} \leq \frac{1}{\sqrt{2\pi}} \frac{\Gamma(m/2 + 1)}{m + 1} \left(\frac{\pi}{\sqrt{2}} \|\beta - \beta'\|_2 \right)^{m+1}.$$

Proof The first identity follows directly from Proposition A.2.1, applied with $m = 1$. To see why, note that:

$$\begin{aligned} \int_0^{\arccos(\beta^T \beta')} \sin(\alpha) \frac{d\alpha}{2\pi} &= \frac{1}{2\pi} (-\cos(\arccos(\beta^T \beta')) + 1) = \frac{1}{2\pi} (-\beta^T \beta' + 1) \\ &= \frac{1}{4\pi} \|\beta - \beta'\|_2^2. \end{aligned}$$

Using that $\Gamma(3/2) = \sqrt{\pi}/2$ establishes the identity. For $m > 1$, we use the bound $\sin(x) \leq x$ for $x \geq 0$. Therefore:

$$\begin{aligned} \int_0^{\arccos(\beta^T \beta')} \sin(\alpha)^m \frac{d\alpha}{2\pi} &\leq \frac{1}{m + 1} \arccos(\beta^T \beta')^{m+1} \frac{1}{2\pi} \\ &\stackrel{(i)}{\leq} \frac{1}{m + 1} \left(\frac{\pi}{2} \|\beta - \beta'\|_2 \right)^{m+1} \frac{1}{2\pi}. \end{aligned}$$

In (i), we upper bounded the geodesic distance on S^{p-1} with the Euclidean distance as in Proposition 3.1.1.

The bound now follows from Proposition A.2.1. \square

The following result gives a closed-form solution for the expectation of the unbounded term. This identity relies on the assumption that (x^T, ϵ) is standard Gaussian. This motivates the assumption of the probit model in conjunction with the logistic loss.

Corollary A.2.2 *Let $x \sim \mathcal{N}(0, I)$, $\epsilon \sim \mathcal{N}(0, 1)$, $\beta, \beta^* \in S^{p-1}$ and $\sigma \geq 0$. Then,*

$$P|x^T \beta| \mathbf{1}\{y x^T \beta < 0\} = \frac{1}{\sqrt{2\pi}} \left(1 - \frac{\beta^T \beta^*}{\sqrt{1 + \sigma^2}} \right).$$

In particular,

$$P|x^T \beta^*| \mathbf{1}\{y x^T \beta^* < 0\} = \frac{1}{\sqrt{2\pi}} \left(1 - \frac{1}{\sqrt{1 + \sigma^2}} \right).$$

Consequentially:

$$P|x^T\beta|1\{yx^T\beta < 0\} - |x^T\beta^*|1\{yx^T\beta^* < 0\} = \frac{1}{\sqrt{8\pi(1+\sigma^2)}}\|\beta - \beta^*\|_2^2.$$

Proof Apply Proposition A.2.1 with $\tilde{x} := (x, \epsilon) \sim \mathcal{N}(0, I_{p+1})$ and $\tilde{\beta} := (\beta, 0)$, as well as $\tilde{\beta}' := (\beta, \sigma)/\sqrt{1+\sigma^2}$. Note that y is the sign of $\tilde{x}^T\tilde{\beta}'$. \square

A.2.2 A Projection-Based Argument In the following, we see another bound for the unbounded term. We use it in the verification of Bernstein's condition. Even in the proof of this result, we use Proposition A.2.1. Recall that $\sigma, \tau^* > 0$, $\beta^* \in S^{p-1}$, $(x, \epsilon) \sim \mathcal{N}(0, I_{p+1})$, $\gamma^* := \tau^*\beta^*$ and $y := \text{sign}(x^T\beta^* + \sigma\epsilon)$. Moreover, we define $y^* := \text{sign}(x^T\beta^*)$, as well as for a fixed $\gamma \in \mathbb{R}^p \setminus \{0\}$, $\beta := \gamma/\|\gamma\|_2$ and $\tau := \|\gamma\|_2$.

Lemma A.2.1 *For any $\gamma \in \mathbb{R}^p \setminus \{0\}$ and $m \in \mathbb{N}$, it holds that:*

$$\begin{aligned} & P|(y - y^*)x^T(\gamma - \gamma^*)|^m \\ & \leq \frac{\sigma}{2\pi}\Gamma\left(\frac{m+1}{2}\right)\left(\frac{1}{\sqrt{\pi}}(\sqrt{32}\tau\|\beta - \beta^*\|_2)^m + (\sqrt{8}\pi|\tau\beta^T\beta^* - \tau^*|\sigma)^m\right). \end{aligned}$$

Proof Let $\Pi := \beta^*\beta^{*T}$ be the orthogonal projection onto the span on β^* , and let $\Pi^\perp := I - \Pi$ be the orthogonal projection to its kernel. We write:

$$\gamma - \gamma^* = \tau\beta - \tau^*\beta^* = \tau\Pi^\perp\beta + \tau\Pi\beta - \tau^*\beta^* = \tau\Pi^\perp\beta + (\tau\beta^T\beta^* - \tau^*)\beta^*.$$

Let $\Phi : \mathbb{R} \rightarrow [0, 1]$ be the cumulative distribution function of a standard Gaussian random variable. Then, after conditioning on x , we have:

$$\begin{aligned} & P|(y - y^*)x^T(\gamma - \gamma^*)|^m = P2^m\Phi\left(-\frac{|x^T\beta^*|}{\sigma}\right)|x^T(\gamma - \gamma^*)|^m \\ & \stackrel{(i)}{\leq} 2^{2m-1}P\Phi\left(-\frac{|x^T\beta^*|}{\sigma}\right)\left(\left|\tau x^T\Pi^\perp\beta\right|^m + |(\tau\beta^T\beta^* - \tau^*)x^T\beta^*|^m\right). \quad (18) \end{aligned}$$

The inequality (i) follows from Jensen's inequality, which implies that for any positive scalars $a, b > 0$ we have $(a/2 + b/2)^m \leq a^m/2 + b^m/2$.

We bound the two terms in the expression (18) separately. We start with the left term. Since x is Gaussian and β^* is orthogonal to $\Pi^\perp\beta$, it follows that $x^T\beta^*$ and $x^T\Pi^\perp\beta$ are independent. Consequentially,

$$P\Phi\left(-\frac{|x^T\beta^*|}{\sigma}\right)\left|\tau x^T\Pi^\perp\beta\right|^m = P\left[\Phi\left(-\frac{|x^T\beta^*|}{\sigma}\right)\right]P\left|\tau x^T\Pi^\perp\beta\right|^m.$$

The first term can be bounded using (5):

$$P \left[\Phi \left(-\frac{|x^T \beta^*|}{\sigma} \right) \right] = \mathbb{P}[y \neq y^*] \leq \frac{\sigma}{\pi}.$$

To bound the second term, note that the random variable $\beta^T \Pi^\perp x$ is Gaussian with mean zero and variance $1 - (\beta^T \beta')^2 \leq 2(1 - \beta^T \beta') = \|\beta - \beta'\|_2^2$. Using that the m -th absolute moment of the a standard Gaussian random variable is $\sqrt{2}^m \Gamma((m+1)/2) / \sqrt{\pi}$, we find:

$$P \left| \tau x^T \Pi^\perp \beta \right|^m \leq \tau^m \left(\sqrt{2} \|\beta - \beta^*\|_2 \right)^m \Gamma \left(\frac{m+1}{2} \right) \frac{1}{\sqrt{\pi}}.$$

Now for the right term in expression (18). It carries the deterministic factor $|\tau \beta^{*T} \beta^* - \tau^*|^m$, which we omit in the following calculation. We define the unit vectors $\tilde{\beta}^* := (\beta^{*T}, 0)^T \in S^p$ and $\tilde{\beta}' := (\beta^{*T}, \sigma)^T / \sqrt{1 + \sigma^2} \in S^p$ and a standard Gaussian $\tilde{x} \sim \mathcal{N}(0, I_{p+1})$. Then:

$$P \Phi \left(-\frac{|x^T \beta^*|}{\sigma} \right) |x^T \beta^*|^m = P1\{\tilde{\beta}^T \tilde{x} \tilde{x}^T \tilde{\beta}^* < 0\} \left| \tilde{x}^T \tilde{\beta} \right|^m.$$

We bound this expression with Corollary A.2.1 and use that $1 - \frac{1}{\sqrt{1+\sigma^2}} \leq \frac{\sigma^2}{2}$, which gives $\|\tilde{\beta} - \tilde{\beta}^*\|_2 \leq \sigma$. To make the expression comparable with the first term, we use that $\Gamma(m/2 + 1)/(m+1) \leq \Gamma((m+1)/2)$. We conclude:

$$P1\{\tilde{\beta}^T \tilde{x} \tilde{x}^T \tilde{\beta}^* < 0\} \left| \tilde{x}^T \tilde{\beta} \right|^m \leq \frac{1}{\sqrt{2}\pi} \Gamma \left(\frac{m+1}{2} \right) \left(\frac{\pi}{\sqrt{2}} \sigma \right)^{m+1}.$$

Combining these bounds completes the proof. \square

Acknowledgements. FK would like to thank Christoph Schultheiss, Malte Londschien, Fadoua Balabdaoui, and Nikita Zhivotovskiy for helpful comments and discussions. The authors would like to thank the anonymous reviewers for helpful comments.

Funding Information Open access funding provided by Swiss Federal Institute of Technology Zurich.

Declarations

Conflict of interest and funding statement The authors have no conflicts of interest to declare. F.K. is funded by the Swiss National Science Foundation (SNF).

Open Access. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aad W van der Vaart and Jon A Wellner. *Weak Convergence and Empirical Processes With Applications to Statistics*. Springer, 1996.
- Aad W van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000.
- Felix Abramovich and Vadim Grinshtein. High-dimensional classification by sparse logistic regression. *IEEE Transactions on Information Theory*, 65(5): 3068–3079, 2018.
- Alan Agresti. *Foundations of Linear and Generalized Linear Models*. John Wiley & Sons, 2015.
- Alexandre Grothendieck. *Résumé de la théorie métrique des produits tensoriels topologiques*. Soc. de Matemática de São Paulo, 1956.
- Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. [arXiv:1911.01544](https://arxiv.org/abs/1911.01544), 2019.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Olivier Bousquet. A Bennett concentration inequality and its application to suprema of empirical processes. *Comptes Rendus Mathématique*, 334 (6): 495–500, 2002.
- Daniel Hsu and Arya Mazumdar. On the sample complexity of parameter estimation in logistic regression with normal design. [arXiv:2307.04191v2](https://arxiv.org/abs/2307.04191v2), 2024.
- Dmitrii M Ostrovskii and Francis Bach. Finite-sample analysis of M-estimators using self-concordance. *Electronic Journal of Statistics*, 15: 326–391, 2021.
- Emmanuel J Candès and Pragya Sur. The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression. *The Annals of Statistics*, 48 (1): 27–42, 2020.
- Emmanuel Jean Candès, Yingying Fan, Lucas Janson, and Jinchi Lv. *Panning for gold: Model-free knockoffs for high-dimensional controlled variable selection*, volume 1610. Department of Statistics, Stanford University Stanford, CA, USA, 2016.

FINITE SAMPLE RATES FOR LOGISTIC REGRESSION...

- Enno Mammen and Alexandre B Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27 (6): 1808–1829, 1999.
- David Haussler. Sphere packing numbers for subsets of the boolean n-cube with bounded vapnik-chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69 (2): 217–232, 1995.
- Lutz Dümbgen, Sara van de Geer, Mark C Veraar, and Jon A Wellner. Nemirovski’s inequalities revisited. *The American Mathematical Monthly*, 117 (2): 138–160, 2010.
- Peter Bühlmann and Sara van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Science & Business Media, 2011.
- Peter McCullagh and John A Nelder. *Generalized Linear Models, Vol. 37 of Monographs on Statistics and Applied Probability London*. 1989.
- Philip M Long. An upper bound on the sample complexity of pac-learning halfspaces with respect to the uniform distribution. *Information Processing Letters*, 87 (5): 229–234, 2003.
- Philip M Long. On the sample complexity of pac learning half-spaces against the uniform distribution. *IEEE Transactions on Neural Networks*, 6 (6): 1556–1559, 1995.
- Pragya Sur and Emmanuel J Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116 (29): 14516–14525, 2019.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47. Cambridge University Press, 2018.
- Sara A van de Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge University Press, 2000.
- Sara A van de Geer. High-dimensional generalized linear models and the lasso. *The Annals of Statistics*, pages 614–645, 2008.
- Szilard Nemes, Junmei Miao Jonasson, Anna Genell, and Gunnar Steineck. Bias in odds ratios by logistic regression modelling and sample size. *BMC Medical Research Methodology*, 9: 1–5, 2009.
- Thomas M Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Transactions on Electronic Computers*, (3): 326–334, 1965.
- Walter W Hauck Jr and Allan Donner. Wald’s test as applied to hypotheses in logit analysis. *Journal of the American Statistical Association*, 72 (360a): 851–853, 1977.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

FELIX KUCHELMEISTER AND SARA VAN
DE GEER
ETH ZÜRICH D-MATH:
EIDGENOSSISCHE TECHNISCHE

HOCHSCHULE ZÜRICH DEPARTEMENT
MATHEMATIK, ZÜRICH, SWITZERLAND
E-mail: felixku@ethz.ch

Paper received: 30 May 2023