

AdaBoost and robust one-bit compressed sensing

Geoffrey Chinot, Felix Kuchelmeister, Matthias Löffler, and
Sara van de Geer

Abstract. This paper studies binary classification in robust one-bit compressed sensing with adversarial errors. It is assumed that the model is overparametrized and that the parameter of interest is effectively sparse. AdaBoost is considered, and, through its relation to the max ℓ_1 -margin-classifier, prediction error bounds are derived. The developed theory is general and allows for heavy-tailed feature distributions, requiring only a weak moment assumption and an anti-concentration condition. Improved convergence rates are shown when the features satisfy a small deviation lower bound. In particular, the results provide an explanation why interpolating adversarial noise can be harmless for classification problems. Simulations illustrate the presented theory.

1. Introduction

Classification is a fundamental statistical problem in data science, with applications ranging from genomics to character recognition. AdaBoost, proposed by Freund and Schapire [24] and further developed in [53], is a popular and successful algorithm from the machine learning literature to tackle such classifications problems. It is based on building an additive model with coefficients $\tilde{\beta}_T$ composed of simple classifiers such as regression trees and then using the binary classification rule $\text{sgn}(\langle \tilde{\beta}_T, \cdot \rangle)$. At each iteration another simple classifier is added to the model, minimizing a weighted loss-function. Alternatively, AdaBoost can be viewed as a variant of mirror-gradient-descent for the exponential loss [9, 25]. Empirically, it often achieves the best generalization performance when it is overparametrized and runs long after the training error equals zero [21].

However, a theoretical understanding of the generalization properties of AdaBoost, that explains this behaviour, is still missing. Early theoretical results on the generalization error of AdaBoost and other classification algorithms were based on

2020 Mathematics Subject Classification. Primary 62H30; Secondary 94A12.

Keywords. AdaBoost, overparameterization, classification, one-bit compressed sensing, sparsity.

margin-theory [33, 52] and entropy bounds. In high-dimensional situations, where the dimension of the features and number of base classifiers is larger than the number of observations n , these become meaningless. Another approach to explain the success of AdaBoost and other boosting algorithms is based on regularization through early stopping [11, 31, 39, 64]. However, by their nature these bounds can not explain generalization performance when the number of iterations grows large and the empirical training error equals zero. In the population setting [10] showed that the generalization risk of AdaBoost converges to the Bayes risk, but this does also not indicate any performance guarantees for finite data.

A more thorough understanding has developed through the lens of optimisation. Already in [24] it was shown that each iteration of AdaBoost decreases the training error. Moreover, in [9, 25], a close connection to the exponential loss was pointed out and studied. Building on these results, [48, 50, 64] discovered that overparametrized AdaBoost, when run long enough with vanishing learning rate ϵ (see Algorithm 1), has ℓ_1 -margin converging to the maximal ℓ_1 -margin. In particular, this means that given training data $(X_i, y_i)_{i=1}^n$, where the y_i are binary and the X_i are p -dimensional feature vectors, and where $\hat{\beta}_T$ denotes the output of AdaBoost with the canonical basis as simple classifiers, learning rate ϵ and run-time T , we have

$$\min_{1 \leq i \leq n} \frac{\langle y_i X_i, \hat{\beta}_T \rangle}{\|\hat{\beta}_T\|_1} \xrightarrow[\epsilon \rightarrow 0]{T \rightarrow \infty} \max_{\beta \neq 0} \min_{1 \leq i \leq n} \frac{\langle y_i X_i, \beta \rangle}{\|\beta\|_1} =: \gamma, \quad (1)$$

provided that γ is positive. The above holds universally for boosting algorithms that are derived from exponential type loss functions and various possible adaptive step-sizes. For these, general non-asymptotic bounds have been developed in [44, 55].

Any vector $\hat{\beta}$ that maximizes the right-hand side in (1) is proportional to an output of

$$\hat{\beta} \in \arg \min \{ \|\beta\|_1 \text{ subject to } \min_{1 \leq i \leq n} y_i \langle X_i, \beta \rangle \geq 1 \}. \quad (2)$$

From the representation (2), it can be seen that, if $\hat{\beta}$ is well-defined, then $\hat{\beta}$ interpolates the data in the sense that $\langle X_i, \hat{\beta} \rangle$ and y_i have matching signs for all i . Similarly, neural networks and random forests are typically massively overparametrized and trained until they interpolate the data. Empirically, it has been shown that this can lead to smaller test errors compared to algorithms with a smaller number of parameters [7, 59]. Statistical learning theory based on empirical risk minimization techniques and entropy bounds can not explain these empirical findings and a mathematical understanding of this phenomenon has only began to form in recent years. The prevalent explanation so far is that, similar as in (1), these algorithms approximate max margin solutions [30, 54, 55]. As in (2), an algorithm that maximizes a margin is equivalent to a minimum-norm interpolator. It is then argued that this leads to implicit regularization and hence a good fit.

The study of minimum-norm interpolating algorithms has mainly been investigated in three settings so far. The first line of research has focused on a random matrix regime where the number of data points and parameters are proportional. Here precise asymptotic results can be obtained, see for instance [18, 43] for max ℓ_2 -margin interpolation, [38] for max ℓ_1 -margin interpolation and consequently AdaBoost, [40] for 2-layer-neural networks in regression and [28] for minimum- ℓ_2 -norm linear regression. However, these results do not exploit possible low-dimensional structure such as sparsity and they also require a large enough, constant, noise-level, leading to inconsistent estimators.

Another line of work has focused on non-asymptotic results in an Euclidean setting with features that have a covariance matrix with decaying eigenvalues, see [45] for classification with support-vector machines (SVM) and [6, 16] for linear regression. These results rely crucially on Euclidean geometry, which gives explicit formulas for the estimators under consideration, and also do not lead to improved convergence rates in the presence of low-dimensional intrinsic structure.

A third line of work originates in the compressed sensing literature. Here low-dimensional intrinsic structure and often small noise levels, including adversarial noise, are studied. Small noise might be a realistic assumption for many classification data sets from the machine learning literature. On data sets such as CIFAR-10 [22] or MNIST [57] state of the art algorithms achieve test errors smaller than 0.5%, implying that the proportion of flipped labels in the full data set is also small. On the theoretical side, pioneering work by Wojtaszczyk [58] has shown that minimum- ℓ_1 -norm interpolation, introduced by [15] as *basis pursuit*, is robust to small, adversarial errors in sparse linear regression. This has recently been extended to other minimum-norm solutions in linear regression [17], phase-retrieval [34] and heavy-tailed features in sparse linear regression [35]. Moreover, ideas from compressed sensing and interpolation were recently combined in [37] to construct non-parametric prediction bands.

Sparsity enables to model the possibility that only few variables are sufficient to predict well and allows for easier model interpretation. In binary classification, a sparse model with adversarial errors can be described by having access to a data set $\mathcal{D}_n = (X_i, y_i)_{i=1}^n$, where the features (X_i) 's are i.i.d. random vectors in \mathbb{R}^p distributed as X and $X = (x_1, \dots, x_p)$ where $x_j \stackrel{\text{i.i.d.}}{\sim} \mu$ for some distribution μ . For $s > 0$, we are given an effectively s -sparse $\beta^* \in \mathcal{S}^{p-1}$, i.e. a vector β^* such that $\|\beta^*\|_2 = 1$ and $\|\beta^*\|_1 \leq \sqrt{s}$. Finally, for a set $\mathcal{O} \subset [n]$, we have

$$y_i = \begin{cases} \text{sgn}(\langle X_i, \beta^* \rangle) & i \notin \mathcal{O}, \\ -\text{sgn}(\langle X_i, \beta^* \rangle) & i \in \mathcal{O}. \end{cases} \quad (3)$$

The set \mathcal{O} contains the indices of the data that is labelled incorrectly. We do not impose any modelling assumptions on \mathcal{O} . \mathcal{O} may be random, deterministic or adversarially

depend on all features $(X_i)_{i=1}^n$, but we impose that the proportion of flipped labels is small such that $|\mathcal{O}| = o(n)$. In the applied mathematics literature, this model is called *robust one-bit compressed sensing* and in learning theory *agnostic learning of (sparse) half-spaces*.

As far as we know, there are no theoretical results for estimators that necessarily interpolate in the model (3) when $\mathcal{O} \neq \emptyset$. In the noiseless case where $\mathcal{O} = \emptyset$ and for standard Gaussian measurements, the authors of [46] have proposed and investigated an interpolating estimator, similarly defined as (2) with the minimum replaced by an average and an additional matching sign constraint. In particular, they showed that this estimator is able to consistently estimate the direction of β^* .

Subsequent work where the model (3) and variants of it were considered, has focused on regularized estimators in order to adapt to noise or to generalize the required assumptions. First results for the model (3) and a computable algorithm were obtained by [47], where a convex program was proposed and investigated. If β^* is exactly s -sparse, i.e. it has at most s non-zero entries, the attainable convergence rates can be improved and faster performance guarantees were obtained in [3, 29, 63]. Further works investigated non-Gaussian measurements [1], active learning [3, 60, 62], overcomplete dictionaries [5], and random shifts of $\langle X_i, \beta^* \rangle$, called *dithering*, [20, 32].

In this paper, we consider the performance of AdaBoost in the overparametrized regime with small and adversarial noise. We additionally assume that β^* is effectively s -sparse. We leverage the relation in (1) between AdaBoost and the max ℓ_1 -margin estimator (2) to analyze AdaBoost (as described below in Algorithm 1). In particular, we show that when $p \gtrsim n$ and the feature vectors fulfil a weak moment assumption and an anti-concentration assumption, then with high probability AdaBoost has vanishing prediction error, provided $(s + |\mathcal{O}|) \log^c(p) = o(n)$ for some constant $c > 0$ and sufficiently many iterations $T = O(n)$ of AdaBoost are performed. Moreover, when the features are Gaussian or Student- t (with at least $c \log(p)$ degrees of freedom) distributed, we obtain prediction and Euclidean estimation error bounds that scale as

$$\left(\frac{(s + |\mathcal{O}|) \log^c(p)}{n} \right)^{1/3}, \quad (4)$$

which is among the best available convergence guarantees in the one-bit compressed sensing literature so far.

These results are, as far as we know, the first non-asymptotic guarantee for overparametrized and data interpolating AdaBoost in a sparse and noisy setting. We illustrate our theory with Laplace, uniform, Gaussian and Student- t (with at least $c \log(p)$ degrees of freedom) distributed features. Moreover, our main result also explains why interpolating data can perform well in the presence of adversarial noise, providing an

explanation to the question raised in [61]. Numerical experiments complement our theoretical results.

Compared to [38], we consider a completely different regime. In their setting, sparsity can not be assumed and the noise level can neither be adversarial nor small. Hence, in [38] consistent estimation of the direction of β^* is impossible and the resulting generalization error is close to $1/2$ when p is large compared to n .

Notation

The Euclidean norm is denoted by $\|\cdot\|_2$ and induced by the inner product $\langle \cdot, \cdot \rangle$, whilst $\|\cdot\|_1$ denotes the ℓ_1 -norm and $\|\cdot\|_\infty$ the ℓ_∞ -norm for both vectors and matrices. B_1^p and B_2^p denote the unit ℓ_1 -ball and ℓ_2 -ball in \mathbb{R}^p , respectively. In addition, we write \mathcal{S}^{p-1} for the p -dimensional unit sphere. By c , we denote a generic, strictly positive constant, that may change value from line to line. Moreover, for two sequences a_n, b_n we write $a_n \lesssim b_n$ if $a_n \leq cb_n \forall n$. Similarly, $a_n \gtrsim b_n$ if $b_n \lesssim a_n$ and $a_n \asymp b_n$ if $a_n \lesssim b_n$ and $b_n \lesssim a_n$. When an assumption reads ‘Suppose $a_n \lesssim b_n$ ’ this means that we assume that for a small enough constant $c > 0$ we have $a_n \leq cb_n \forall n$. By $[p]$ we denote the enumeration $\{1, \dots, p\}$, by $\{e_j\}_{j \in [p]}$ the set of canonical basis vectors in \mathbb{R}^p , and by X_i the i -th column of the matrix $\mathbb{X} = [X_1, \dots, X_n]$ of feature vectors. We denote the sign function, $\text{sgn}(x) = \mathbf{1}(x > 0) - \mathbf{1}(x < 0)$. Throughout this article, we use bold letters to denote random matrices, upper case letters for random vectors and lower case letters for random variables. For example, we write $\mathbb{X} = (X_i)_{i \in [n]} \in \mathbb{R}^{p \times n}$ and $X_i = (x_{i,1}, \dots, x_{i,p}) \in \mathbb{R}^p$.

2. Main results

2.1. Model and assumptions

We consider a binary classification model, which allows for adversarial flips. In particular, we assume that we have access to a data set $(y_i, X_i)_{i \in [n]}$. The X_i ’s are i.i.d. random vectors in \mathbb{R}^p distributed as X and $X = (x_1, \dots, x_p)$, where $x_j \stackrel{\text{i.i.d.}}{\sim} \mu$ for some distribution μ that is symmetric and has zero mean and finite variance. Without loss of generality, we may assume that the variance is one, as both the considered loss functions and the observed data are scaling invariant. The y_i ’s are generated via

$$y_i = \begin{cases} \text{sgn}(\langle X_i, \beta^* \rangle) & i \notin \mathcal{O}, \\ -\text{sgn}(\langle X_i, \beta^* \rangle) & i \in \mathcal{O}. \end{cases} \quad (5)$$

The set $\mathcal{O} \subset [n]$ is the set of the indices of the mislabelled data. We assume that the fraction of flipped labels is asymptotically vanishing, $|\mathcal{O}| = o(n)$, but that \mathcal{O} may be

picked by an adversary and depend on the data. In particular, this includes parametric noise models such as logistic regression or additive Gaussian noise inside the sign-function above, as long as the variance of the noise decays to zero as n goes to infinity. We assume that $\beta^* \in \mathcal{S}^{p-1}$ is effectively s -sparse, that is $\|\beta^*\|_1 \leq \sqrt{s}$.

For stating our results we will treat all other parameters that do not depend on p, n, s and $|\mathcal{O}|$ as fixed constants. Moreover, we always assume tacitly that $p \geq cn$, for a large enough constant $c > 0$.

We measure the accuracy of recovery by the prediction error

$$d(\tilde{\beta}, \beta^*) := \mathbb{P}(\text{sgn}(\langle X_{n+1}, \tilde{\beta} \rangle) \neq \text{sgn}(\langle X_{n+1}, \beta^* \rangle) \mid (y_i, X_i)_{i \in [n]}), \quad (6)$$

where X_{n+1} is an independent copy of X . This is the quantity which is used empirically to measure the quality of classifiers such as neural networks on standard image benchmark data sets [22, 57].

We now formulate the three main assumptions used throughout this article. They describe the tail-behaviour and behaviour around zero of the features. For the tail-behaviour, the only assumption we make is a weak moment assumption of order $\log(p)$.

Definition 2.1. A centred, scalar random variable x fulfils the weak moment assumption (of order $\log(p)$) with parameter $\zeta \geq 1/2$ if

$$(\mathbb{E}|x|^q)^{1/q} \lesssim q^\zeta \quad \forall 1 \leq q \leq \max(1, \log(p)).$$

For a matrix \mathbb{X} or a vector X , we say that they satisfy the weak moment assumption if their entries satisfy the weak moment assumption.

This assumption is weaker than commonly assumed sub-Gaussian or sub-exponential tail behaviour and allows for feature distributions with heavy-tails such as the Student- t -distribution with $c \log(p)$ degrees of freedom. Under the weak moment assumption we can control the ℓ_∞ norm of $X = (x_1, \dots, x_p)$ composed of i.i.d. random variables satisfying the weak moment assumption with polynomial deviation (see Proposition B.1). Assuming sub-Gaussianity of the x_j 's does not lead to improvement of the convergence rates in our main results except for logarithmic factors. This is different to the theory developed in [20] where the exponent in the obtained convergence results depends on whether the features are sub-Gaussian or heavy-tailed.

The next two assumptions measure the behaviour of the feature distribution around zero.

Definition 2.2. A random vector $X \in \mathbb{R}^p$ fulfils an anti-concentration assumption with parameter $\alpha \in (0, 1]$ if

$$\sup_{\beta \in \mathcal{S}^{p-1}} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) \lesssim \varepsilon^\alpha \quad \forall p^{-1} \leq \varepsilon \leq 1. \quad (7)$$

We say that a matrix $\mathbb{X} \in \mathbb{R}^{p \times n}$ satisfies the anti-concentration assumption with $\alpha \in (0, 1]$ if each column of \mathbb{X} satisfies (7).

Assuming that \mathbb{X} satisfies an anti-concentration assumption will be a necessity for our results, as it ensures that for fixed vector β , the scalar $|\langle X_i, \beta \rangle|$ is not too close to zero for too many indices $i \in [n]$. This in turn would lead to a tiny ℓ_1 -margin and many discontinuities of $\text{sgn}(\langle X_i, \beta \rangle)$ at β , rendering it impossible to prove uniform results.

Similar anti-concentration assumptions were previously introduced in the learning theory literature in the non-sparse setting, see e.g. [4, 19, 23], and were shown to be satisfied by isotropic log-concave distributions [4] via an uniform upper bound for the density of $\langle \beta, X \rangle$.

The next assumption is an optional counterpart to Definition 2.2 and leads to improved convergence rates if it is satisfied.

Definition 2.3. A random vector $X \in \mathbb{R}^p$ fulfils a small deviation assumption with parameter $\theta > 0$ if

$$\inf_{\beta \in \mathcal{S}^{p-1}} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) \gtrsim \varepsilon^\theta \quad \forall p^{-1} \leq \varepsilon \leq 1. \quad (8)$$

We say that a matrix $\mathbb{X} \in \mathbb{R}^{p \times n}$ satisfies the small deviation assumption with parameter $\theta > 0$ if each column of \mathbb{X} satisfies (8).

In [19] a stronger small-deviation assumption was formulated, assuming an uniform lower bound on the density of two-dimensional projections of X . This property is satisfied by isotropic log-concave distributions [4] and implies (8) with $\theta = 1$.

2.2. Main results

2.2.1. AdaBoost, max ℓ_1 -margin and a bound in terms of the margin. AdaBoost, proposed by Freund and Schapire [24], is an algorithm where an additive model for an unnormalized version of β^* is built by iteratively adding weak classifiers to the model. To facilitate our analysis, we assume that the features X_i 's are i.i.d. distributed and that the weak classifiers can be identified with the standard basis vectors in \mathbb{R}^p . We consider AdaBoost as described in Algorithm 1. The main difference to the original proposal by [24] consists of the choice of the step-size α_t , which is obtained by minimizing a quadratic upper bound for the loss-function at each step [55].

Alternative to the interpretation by Freund and Schapire [24], AdaBoost can be viewed as a form of mirror gradient descent on the exponential loss-function [9, 25]. It is thus natural to expect that it converges to the infimum of the loss-function and eventually interpolates the labels if possible. In fact, a stronger statement holds: As described in (1), AdaBoost with infinitesimally small learning rate and a growing

Algorithm 1: AdaBoost for binary classification

Input: Binary data $(y_i)_{i \in [n]}$, features $\mathbb{X} = (X_i)_{i \in [n]}$, run-time T , learning rate ϵ

Output: Vector $\tilde{\beta}_T \in \mathbb{R}^p$

1 Initialize $\tilde{\beta}_{0,i} = 0$ and rescale features $\mathbb{X} = \mathbb{X}/\|\mathbb{X}\|_\infty$

2 For $t = 1, \dots, T$ repeat

- Update weights $w_{t,i} = \frac{\exp(-y_i \langle X_i, \tilde{\beta}_{t-1} \rangle)}{\sum_{j=1}^n \exp(-y_j \langle X_j, \tilde{\beta}_{t-1} \rangle)}$, $i = 1, \dots, n$
- Select coordinate $v_t = \arg \max_{v \in \{e_j\}_{j=1}^p} |\sum_{i=1}^n w_{t,i} y_i \langle X_i, v \rangle|$
- Compute adaptive step size $\alpha_t = \sum_{i=1}^n w_{t,i} y_i \langle X_i, v_t \rangle$
- Update $\tilde{\beta}_t = \tilde{\beta}_{t-1} + \epsilon \alpha_t v_t$

3 Return $\tilde{\beta}_T$

number of iterations T converges to a solution that maximizes the ℓ_1 -margin [48, 55, 64].

This holds even non-asymptotically [55] for many variants of AdaBoost and includes both the exponential and logistic loss-function as well as various choices of adaptive step sizes α_t , for instance logarithmic as originally proposed by [24], line search [53, 64] or quadratic as in Algorithm 1.

To present non-asymptotic results and to ensure that our theory can potentially be applied to other variants of AdaBoost, we introduce the following definition of an approximation of the largest ℓ_1 -margin: We say that $\tilde{\beta} \in \mathbb{R}^p$ provides an approximation of the max ℓ_1 -margin if

$$\min_{1 \leq i \leq n} \frac{y_i \langle X_i, \tilde{\beta} \rangle}{\|\tilde{\beta}\|_1} \geq \frac{1}{2} \max_{\beta \neq 0} \min_{1 \leq i \leq n} \frac{y_i \langle X_i, \beta \rangle}{\|\beta\|_1} =: \frac{\gamma}{2}. \quad (9)$$

The quantity γ is called the max ℓ_1 -margin. Moreover, the factor $1/2$ can be substituted by any other positive constant smaller than one.

The following theorem gives a bound for the prediction error for any $\tilde{\beta}$ that provides an approximation of the max ℓ_1 -margin. The bound itself depends on the max ℓ_1 -margin γ . If the features fulfil an additional small deviation assumption, we obtain improved convergence rates.

Theorem 2.1. *Assume $p \gtrsim n$ and that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$ and the anti-concentration assumption with $\alpha \in (0, 1]$. Let $\tilde{\beta}$ be an approximation of the margin*

and suppose that $\tilde{\beta}$ satisfies with probability at least $1 - t$ that $\gamma \geq \gamma_0$. Define

$$\eta = \left(\frac{\log^{2\zeta+1}(p) \log(n)}{\gamma_0^2 n} \right)^{1/(2+\alpha)},$$

and assume that $\eta \lesssim 1$. Moreover, assume that $|\mathcal{O}| \lesssim \eta^\alpha n$. Then with probability at least $1 - cp^{-1} - t$ we have that

$$d(\tilde{\beta}, \beta^*) \lesssim \eta^\alpha.$$

Moreover, if \mathbb{X} satisfies a small deviation assumption with $\theta > 0$ and $|\mathcal{O}| \lesssim \eta^{\alpha(1+\frac{2}{\theta})} n$, then, with probability at least $1 - cp^{-1} - t$, we have that

$$d(\tilde{\beta}, \beta^*) \lesssim \eta^{\alpha(1+\frac{2}{\theta})}.$$

The proof of Theorem 2.1 involves two main arguments: a bound for the ratio $\|\tilde{\beta}\|_1/\|\tilde{\beta}\|_2$ in terms of the max ℓ_1 -margin and a sparse hyperplane tessellation result that adapts a proof technique introduced by [20]. For the bound on the ratio we argue by contradiction and show that with high probability no β can simultaneously approximate the margin and have small Euclidean norm. If the small deviation assumption is satisfied we obtain an improved bound on the ratio by using a more involved discretization argument via Maurey's empirical method [12, 13, 51]. For the sparse hyperplane tessellation result, we argue similarly as [20], but use again Maurey's empirical method instead of their net argument. Compared to a discretization argument via nets (as in [20, 47]) this has the advantage that we are able to deal with features that only fulfil the weak moment assumption, while still retaining the same rate (up to logarithmic factors) as in the sub-Gaussian case. By contrast, the obtained convergence rates in [20] depend on whether the features are sub-Gaussian or not.

The following lemma shows that AdaBoost, as described in Algorithm 1, provides an approximation of the max ℓ_1 -margin, when it is run long enough. The proof is a simple adaptation of results by [55] to our setting.

Lemma 2.1. *Consider the AdaBoost Algorithm 1 and suppose that $p \gtrsim n$ and that $\mathbb{X} = (X_i)_{i \in [n]}$ satisfies the weak moment assumption with $\zeta \geq 1/2$. Suppose that $\gamma > 0$, that the learning rate ϵ satisfies $\epsilon \leq 1/6$ and that*

$$T \gtrsim \log^{2\zeta+1}(np)/(\epsilon^2 \gamma^2).$$

Then, the output of Algorithm 1 provides an approximation of the max ℓ_1 -margin with probability at least $1 - p^{-1}$.

Hence, both for algorithmic (Lemma 2.1) as well as recovery guarantees (Theorem 2.1), it is necessary to obtain a lower bound on the max ℓ_1 -margin γ .

2.2.2. A bound for the max ℓ_1 -margin. In this section, we obtain a lower on the max ℓ_1 -margin γ , holding with large probability.

Theorem 2.2. *Assume that $p \gtrsim n$ and that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. symmetric, zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$ and the anti-concentration assumption with $\alpha \in (0, 1]$. Then, with probability at least $1 - cn^{-1}$, we have that*

$$\gamma \gtrsim \left[\frac{n}{\log(ep/n)} \left(s + \frac{\log^{1+2\zeta}(n)|\mathcal{O}|}{\log(ep/n)} + \log^{1+2\zeta}(n) \right)^{\alpha/2} \right]^{-1/(2+\alpha)}. \quad (10)$$

Crucial for the proof of Theorem 2.2 is the fact that, defining

$$\hat{\beta} \in \arg \min \{ \|\beta\|_1 \text{ subject to } y_i \langle X_i, \beta \rangle \geq 1, i = 1, \dots, n \}, \quad (11)$$

we have the relation $\gamma = 1/\|\hat{\beta}\|_1$ (see Lemma 4.1). Hence, to obtain a lower bound for γ it suffices to obtain an upper bound for $\|\hat{\beta}\|_1$, which we accomplish by explicitly constructing a β that fulfils the constraints in (11). In particular, we use the ℓ_1 -quotient property [35, 58] to find a perturbation of β^* that has sufficiently small ℓ_1 -norm while still fulfilling the constraint $y_i \langle X_i, \beta \rangle \geq 1$ for all $i \in [n]$.

The dependence in (10) on the $\log(n)$ -factor can be improved in certain situations, e.g. when $\mathcal{O} = \emptyset$ or $s \gtrsim \log(n)$, leading to a slightly sharper bound.

The following proposition shows that even in an idealized setting with no noise and isotropic Gaussian features, where $\alpha = 1$ (see Corollary 2.3), the lower bound on the margin in Theorem 2.2 is, in general, tight (up to logarithmic factors).

Proposition 2.1. *Suppose $p \gtrsim n$, $\mathcal{O} = \emptyset$ and that the entries of \mathbb{X} are i.i.d. $\mathcal{N}(0, 1)$ distributed. Then, for any $\beta^* \in \mathcal{S}^{p-1}$ which satisfies $\|\beta^*\|_\infty \lesssim 1/\sqrt{s}$, we have that*

$$\mathbb{E}\gamma \lesssim \left(\frac{\log(p)}{n} \frac{1}{\sqrt{s}} \right)^{1/3}. \quad (12)$$

2.2.3. Rates for AdaBoost. Combining Theorems 2.1 and 2.2 with Lemma 2.1, we obtain the following corollary that shows convergence rates for AdaBoost.

Corollary 2.1. *Grant the assumptions of Theorem 2.2 and assume that for some large enough constant $\kappa_1 = \kappa_1(\alpha, \zeta)$ the AdaBoost Algorithm 1 is run for*

$$T \gtrsim \left(n(s + |\mathcal{O}|)^{\alpha/2} \right)^{2/(2+\alpha)} \log^{\kappa_1}(p) \epsilon^{-2}$$

iterations with learning rate $\epsilon \leq 1/6$. Then, with probability at least $1 - cn^{-1}$, the output $\tilde{\beta}_T$ of AdaBoost Algorithm 1 satisfies

$$d(\tilde{\beta}_T, \beta^*) \lesssim \left(\frac{(s + |\mathcal{O}|) \log^{\kappa_2}(p)}{n} \right)^{\alpha/(2+\alpha)^2}$$

for some constant $\kappa_2 = \kappa_2(\alpha, \zeta)$. Moreover, if \mathbb{X} satisfies a small deviation assumption with $\theta > 0$, then

$$d(\tilde{\beta}_T, \beta^*) \lesssim \left(\frac{(s + |\mathcal{O}|) \log^{\kappa_2}(p)}{n} \right)^{\alpha(1 + \frac{2}{\theta}) / (2 + \alpha)^2}$$

with probability at least $1 - cn^{-1}$.

As for consistency $(s + |\mathcal{O}|) \log^{\kappa_2}(p) = o(n)$ is required, it is ensured that in this relevant regime AdaBoost is an approximation of the max ℓ_1 -margin if we run AdaBoost for $T \asymp n \log(p)^{\kappa_1} / \epsilon^{-2}$ iterations. Hence, by contrast to other algorithms such as gradient descent (e.g. [8, Section 9.3.1]) where often a logarithmic number of iterations in n suffices, we require in the worst case an approximately linear number of iterations in n to ensure consistency of AdaBoost.

2.2.4. Examples. We now illustrate our developed theory for some specific feature distributions. First, for the density of the x_{ij} 's being continuous, bounded and unimodal, we are able to show that the anti-concentration condition holds with parameter $\alpha = 1/2$.

Corollary 2.2. *Assume that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. symmetric, zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$. Assume that the x_{ij} 's have density f that is continuous, bounded by a constant, and unimodal, i.e. $f(a\varepsilon) \geq f(\varepsilon) \forall a \in (0, 1), \varepsilon \in \mathbb{R}$. Then \mathbb{X} satisfies the anti-concentration condition with parameter $\alpha = 1/2$. In particular, this includes features that are distributed according to the uniform, Gaussian, Student- t with $d \gtrsim \log(p)$, $d \in \mathbb{N}$, degrees of freedom distributions (with $\zeta = 1/2$) and the Laplace distribution (with $\zeta = 1$). Hence, when $p \gtrsim n$ and AdaBoost is for some constant $\kappa_1 = \kappa_1(\zeta)$ run for at least*

$$T \gtrsim (n(s + |\mathcal{O}|)^{1/4})^{4/5} \log(p)^{\kappa_1} \epsilon^{-2}$$

iterations, then with probability at least $1 - cn^{-1}$, we have that

$$d(\tilde{\beta}_T, \beta^*) \lesssim \left(\frac{(s + |\mathcal{O}|) \log^{\kappa_2}(p)}{n} \right)^{2/25}$$

for some constant $\kappa_2 = \kappa_2(\zeta)$.

When the features are Gaussian or Student- t with at least $c \log(p)$ degrees of freedom distributed, we are able to improve upon this and show that the anti-concentration and small deviation conditions are both fulfilled with parameters $\alpha = \theta = 1$. Moreover, for these distributions the prediction and Euclidean estimation error are closely related such that we are also able to obtain error bounds in this distance.

Corollary 2.3. *Assume that the entries of $\mathbb{X} = (X_i)_{i \in [n]}$ are i.i.d. $\mathcal{N}(0, 1)$ or $\sqrt{(d-2)/d} t_d$ distributed for $\log(p) \lesssim d$, $d \in \mathbb{N}$, $p \gtrsim 1$. Then \mathbb{X} satisfies the anti-concentration and small deviation assumptions with $\alpha = \theta = 1$ and the weak moment assumption with $\zeta = 1/2$. In particular, when $p \gtrsim n$, and after at least*

$$T \gtrsim (n(s + |\mathcal{O}|)^{1/2})^{2/3} \log(p)^{\kappa_1} \epsilon^{-2}$$

iterations of AdaBoost, we have with probability at least $1 - cn^{-1}$ that

$$d(\tilde{\beta}_T, \beta^*) \lesssim \left(\frac{(s + |\mathcal{O}|) \log^{\kappa_2}(p)}{n} \right)^{1/3}$$

for some constant κ_2 . Moreover, on the same event, we have that

$$\left\| \frac{\tilde{\beta}_T}{\|\tilde{\beta}_T\|_2} - \beta^* \right\|_2 \lesssim \left(\frac{(s + |\mathcal{O}|) \log^{\kappa_2}(p)}{n} \right)^{1/3}. \quad (13)$$

We now compare the convergence guarantees for AdaBoost with Gaussian or Student- t distributed features with the state of the literature, where mostly Gaussian features and Euclidean estimation error were considered. When $\mathcal{O} = \emptyset$ the performance guarantee in (13) is better than existing bounds for regularized algorithms [47, 63] and match, up to logarithmic factors, the best available bounds that can be obtained by combining the tessellation result in Proposition 4.2 with Plan and Vershynin's [46] linear programming estimator. A straightforward adaptation of the proofs from [20] for the tessellation to our setting, would lead to an exponent of $1/4$ in (13) in case of the Student- t distribution with at least $c \log(p)$ degrees of freedom. We achieve improved rates by replacing the net discretization from [20] with a more involved Maurey argument.

For Gaussian features and in the presence of adversarial errors, the convergence rate obtained in (13) improves over the rate for the regularized estimator by [47] if $(|\mathcal{O}|/n)^4 = o(s/n)$ and otherwise their algorithm achieves faster convergence rates, in both cases up to logarithmic factors. If β^* is exactly s -sparse, i.e. at most s entries of β^* are non-zero, then the rate in (13) is sub-optimal in the dependence on $s \log(p)/n$ and $|\mathcal{O}|/n$ and faster rates were obtained for a (non-interpolating) regularized estimator in [3] for strongly log-concave features.

2.3. Simulations

In this subsection, we provide simulations for various feature distributions to illustrate our theoretical results qualitatively. Alongside Theorem 2.1, we show the empirical prediction error as a function of the sample size n and the number of corrupted

labels $|\mathcal{O}|$. Moreover, to accompany Theorem 2.2, we plot the margin as a function of n .

As illustrated in Corollary 2.2, the developed theory applies to various distributions of the entries of the features X_{ij} , such as continuous, bounded and unimodal distributions. To highlight the universality of our theory, simulations were performed for the standard normal distribution $\mathcal{N}(0, 1)$, the Student- t distribution with $\log(p)$ degrees of freedom, the uniform distribution with unit variance, and the Laplace distribution with zero location and unit scale parameter.

The ground truth β^* was generated randomly, with an s -sparse Rademacher prior. That is, s out of the possible p entries are chosen at random, and set to $\pm 1/\sqrt{s}$ with equal probability. The remaining entries are set to zero, making β^* s -sparse, with $\|\beta^*\|_2 = 1$. The indices for the set of corruptions \mathcal{O} was chosen uniformly at random, such that a predetermined number of labels is corrupted. The sparsity was chosen as $s = 5$. As Theorem 2.1 assumes $p \gtrsim n$, we let $p = 10n$. AdaBoost was executed as described in Algorithm 1, using step size $\epsilon = 0.2$. The number of steps performed was $T = (n\sqrt{s + |\mathcal{O}|})^{2/3} \log(p)/\epsilon^2$ steps, imitating the setting in Corollary 2.3. The simulations are averaged over twenty iterations.

The two plots in Figure 1 show the noisy case, while noise is absent in the two plots in Figure 2. For the max ℓ_1 -margin estimator the prediction error for all simulated features appears to behave identically. By contrast, the ℓ_1 -margin differs widely across the features by a multiplicative constant, but shows the same asymptotic behaviour.

As stated in Lemma 2.1, we see that the margin of AdaBoost is close to the max ℓ_1 -margin and that the performance of AdaBoost is similar to the performance of the max ℓ_1 -margin classifier. The proximity of AdaBoost to its limit appears to depend on the distribution of the features. In particular, the simulations suggest that heavier tails lead to slower convergence. This is reasonable, considering that AdaBoost rescales the features with their ℓ_∞ -norm, see Algorithm 1. This is particularly visible when comparing the uniform distribution, for which the max ℓ_1 -margin estimator and AdaBoost seem to behave almost identically, to the Student- t distribution, for which the margin is close to zero for some n .

3. Conclusion

In this paper, we have shown that AdaBoost, as described in Algorithm 1, achieves consistent recovery in the presence of small, adversarial errors, despite being overparametrized and interpolating the observations. Our results hold under weak assumptions on the tail behaviour and the behaviour around zero of the feature distribution. In addition, for Gaussian features the derived convergence rates in Corollary 2.3 are

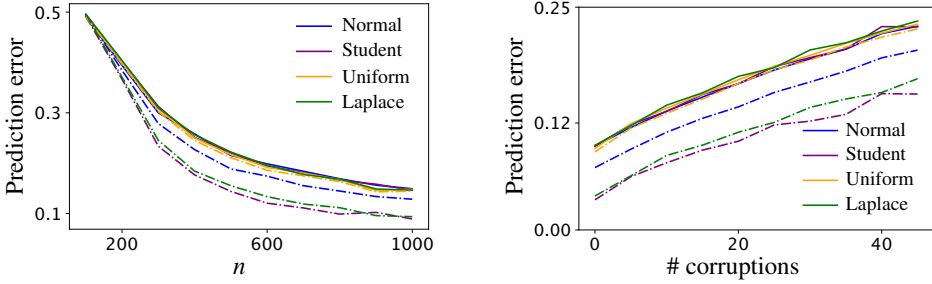


Figure 1. On the left, we plot the prediction error for $|\mathcal{O}| = 40$ corruptions, against the number of samples n , for various features. On the right, we show for $n = 500$ how the prediction error changes as the number of randomly flipped labels $|\mathcal{O}|$ decreases. The solid lines represent the max ℓ_1 -margin estimators $\hat{\beta}$ (2). The dash-dotted lines are instances of AdaBoost $\tilde{\beta}_T$, as defined in Algorithm 1.

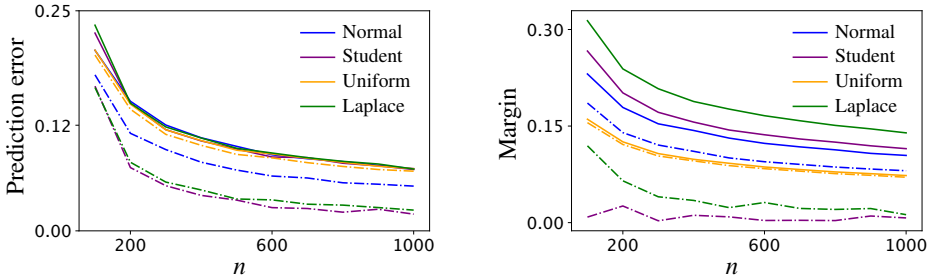


Figure 2. On the left, we consider the same setting as in Figure 1, however in the case of noiseless data $|\mathcal{O}| = 0$. On the right, we plot for noiseless data the margins γ of the max ℓ_1 -margin estimators, as defined in (1), as well as the ℓ_1 -margins of AdaBoost $\tilde{\beta}_T$.

comparable to convergence rates of state-of-the-art regularized estimators [47]. This is a first step for the understanding of overparametrized and interpolating AdaBoost and other interpolating algorithms and shows why such algorithms can generalize well in high-dimensional and noisy situations, despite interpolating the data.

However, in the presence of well-behaved noise, as in logistic regression, our bounds are suboptimal and require that the fraction of mislabelled data points decays to zero. By contrast, regularized estimators [47] are able to achieve faster convergence rates in such settings and do not require that the fraction of mislabelled data is asymptotically vanishing to achieve consistency.

Many open question do remain. The convergence rate for Gaussian features in Corollary 2.3 is among the best available results if β^* is allowed to be genuinely effectively sparse. However, it is not clear whether the exponent in (13) is optimal,

and further research about information theoretic lower bounds is needed. When β^* is exactly sparse, the convergence rate in (13) is sub-optimal and better results for log-concave features have been obtained by [3] for a regularized estimator. Moreover, for noiseless data and exact sparse β^* our simulations suggest that AdaBoost attains a faster rate than in the noisy case. It thus remains as an interesting further research question how to show that AdaBoost attains faster convergence rates for noiseless data and when β^* is exactly s -sparse.

Finally, our results rely heavily on the anti-concentration assumption in Definition 2.2, which is not fulfilled for Rademacher features. Assuming additionally

$$\|\beta^*\|_\infty = o(1),$$

the authors of [1] obtained convergence rates for the regularized estimator proposed in [47]. It is straightforward to adapt our lower bound on the max ℓ_1 -margin in Theorem 2.2 to such a setting. However, by contrast, it is not clear how to modify the uniform tessellation result used in the proof of Theorem 2.1 and consequently show convergence rates for AdaBoost without anti-concentration.

4. Proofs

4.1. Proof of Theorem 2.1

Proof. Let $\tilde{\beta}$ be an approximation of the max ℓ_1 -margin. Defining $\bar{\beta} := 2\tilde{\beta}/(\gamma_0\|\tilde{\beta}\|_1)$, we have on an event of probability at least $1 - t$ that

$$\min_{i \in [n]} y_i \langle X_i, \bar{\beta} \rangle \geq 1,$$

and $\bar{\beta} \in r_n B_1^p$ for $r_n = 2/\gamma_0$. It follows that $\|\tilde{\beta}\|_1/\|\tilde{\beta}\|_2 = \|\bar{\beta}\|_1/\|\bar{\beta}\|_2$, which we bound by applying the following proposition, which we prove in Section 4.3.

Proposition 4.1. *Assume $p \gtrsim n$ and that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$. Let $r_n > 0$ be such that*

$$r_n \lesssim \sqrt{\frac{n}{\log^{2\zeta+1}(p) \log(n)}}.$$

Then with probability at least $1 - cp^{-1}$, for any $\beta \in \mathbb{R}^p$ such that $\|\beta\|_1 \leq r_n$ and $\min_{i \in [n]} y_i \langle X_i, \beta \rangle \geq 1$, we have that

$$\frac{\|\beta\|_1}{\|\beta\|_2} \lesssim r_n.$$

Moreover, assume that \mathbb{X} fulfils a small deviation assumption with parameter $\theta > 0$. Then with probability at least $1 - cp^{-1}$, for any $\beta \in \mathbb{R}^p$ such that $\|\beta\|_1 \leq r_n$ and $\min_{i \in [n]} y_i \langle X_i, \beta \rangle \geq 1$, we have that

$$\frac{\|\beta\|_1}{\|\beta\|_2} \lesssim \frac{r_n}{\tau_n},$$

where

$$\tau_n \asymp \left[\frac{n}{\log^{2\zeta+1}(p) \log(n) r_n^2} \right]^{1/\theta}.$$

Having obtained a bound for the ratio $\|\tilde{\beta}\|_1/\|\tilde{\beta}\|_2$, we next use a sparse hyperplane tessellation result for the pseudo-metric d , arguing by contradiction. The proof of Proposition 4.2 is presented in Section 4.4. Since d is scaling invariant, i.e.

$$d(\beta, \tilde{\beta}) = d(\beta, \tilde{\beta}/\|\tilde{\beta}\|_2),$$

it suffices to consider only elements on the unit sphere.

Proposition 4.2. *Assume $p \gtrsim n$ and that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$ and the anti-concentration assumption with $\alpha \in (0, 1]$. For $a > 0$, define*

$$\eta = c \left(a^2 \frac{\log^{2\zeta+1}(p) \log(n)}{n} \right)^{1/(2+\alpha)}, \tag{14}$$

and assume $\eta \leq 1/2$. Define

$$\mathcal{B}(a, \eta) = \{\beta \in \mathbb{R}^p : d(\beta, \beta^*) \geq c\eta^\alpha\}.$$

Then with probability at least $1 - cp^{-1}$, we have uniformly for $\beta \in aB_1^p \cap \mathcal{S}^{p-1} \cap \mathcal{B}(a, \eta)$ that

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}(\text{sgn}(\langle X_i, \beta \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)) \gtrsim \eta^\alpha.$$

Now we apply Proposition 4.2 with $a \asymp r_n$. Since $|\mathcal{O}| \lesssim \eta^\alpha$ by assumption, we obtain

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}(\text{sgn}(\langle X_i, \tilde{\beta}/\|\tilde{\beta}\|_2 \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)) = \frac{|\mathcal{O}|}{n} \lesssim \eta^\alpha,$$

and hence, adjusting constants, we have on an event of probability at least $1 - t - cp^{-1}$ that $\tilde{\beta}/\|\tilde{\beta}\|_2 \notin \mathcal{B}(a, \eta)$, and hence on the same event $d(\tilde{\beta}, \beta^*) \lesssim \eta^\alpha$.

When \mathbb{X} satisfies the small deviation assumption with parameter $\theta > 0$, we apply Proposition 4.2 with $a \asymp r_n/\tau_n$. Since $|\mathcal{O}| \lesssim \eta^{\alpha(1+\frac{2}{\theta})}$ by assumption, we conclude the proof using the same reasoning. ■

4.2. Upper and lower bounds for the max ℓ_1 -margin

4.2.1. Proof of Theorem 2.2. We start this section with the following lemma. A proof is given in [38].

Lemma 4.1 ([38, Proposition A.2]). *Suppose that*

$$\gamma := \max_{\beta \neq 0} \min_{1 \leq i \leq n} \frac{y_i \langle X_i, \beta \rangle}{\|\beta\|_1} > 0.$$

Then, we have that $\gamma^{-1} = \|\hat{\beta}\|_1$, where

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \{\|\beta\|_1 \text{ subject to } y_i \langle X_i, \beta \rangle \geq 1\}. \quad (15)$$

Hence, in order to lower bound γ it suffices to upper bound $\|\hat{\beta}\|_1$, which is accomplished in the following proposition.

Proposition 4.3. *Assume $p \gtrsim n$ and that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. symmetric, zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$ and the anti-concentration assumption with $\alpha \in (0, 1]$. Then with probability at least $1 - cn^{-2}$, we have that*

$$\|\hat{\beta}\|_1 \lesssim \left[\frac{n}{\log(ep/n)} \left(s + \frac{\log^{1+2\zeta}(n)|\mathcal{O}|}{\log(ep/n)} + \log^{1+2\zeta}(n) \right)^{\alpha/2} \right]^{1/(2+\alpha)}. \quad (16)$$

Proof. We prove Proposition 4.3 by explicitly constructing a β that fulfils the constraints in (15). For $\varepsilon > 0$, we define a lifting function $f_\varepsilon: \mathbb{R} \rightarrow \mathbb{R}$, where

$$f_\varepsilon(x) := \begin{cases} x - \varepsilon & \text{if } 0 \leq x \leq \varepsilon, \\ x + \varepsilon & \text{if } -\varepsilon \leq x < 0, \\ 0 & \text{otherwise.} \end{cases}$$

For $i \in [n]$, we denote

$$z_i = \begin{cases} f_\varepsilon(\langle X_i, \beta^* \rangle) & i \notin \mathcal{O}, \\ 2\langle X_i, \beta^* \rangle - f_\varepsilon(\langle X_i, \beta^* \rangle) & i \in \mathcal{O}, \end{cases}$$

and $Z = (z_1, \dots, z_n)^T$. Finally, we define

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \text{ subject to } \langle X_i, \beta \rangle = z_i, \quad i = 1, \dots, n. \quad (17)$$

By definition of \hat{v} , if $i \in \mathcal{O}$, we have the decomposition

$$\begin{aligned} \langle X_i, \beta^* - \hat{v} \rangle &= -\langle X_i, \beta^* \rangle + f_\varepsilon(\langle X_i, \beta^* \rangle) \\ &= \begin{cases} -\langle X_i, \beta^* \rangle & \text{if } |\langle X_i, \beta^* \rangle| \geq \varepsilon, \\ -\varepsilon & \text{if } 0 \leq \langle X_i, \beta^* \rangle \leq \varepsilon, \\ \varepsilon & \text{if } -\varepsilon \leq \langle X_i, \beta^* \rangle < 0. \end{cases} \end{aligned}$$

A similar decomposition with each equation above multiplied with -1 holds if $i \notin \mathcal{O}$. Hence, we have that $\text{sgn}(\langle X_i, \beta^* - \hat{v} \rangle) = y_i$ and $|\langle X_i, \beta^* - \hat{v} \rangle| \geq \varepsilon$ for $i = 1, \dots, n$. It follows that

$$\|\hat{\beta}\|_1 \leq \frac{\|\beta^* - \hat{v}\|_1}{\varepsilon} \leq \frac{\sqrt{s}}{\varepsilon} + \frac{\|\hat{v}\|_1}{\varepsilon}.$$

We now apply Proposition B.1 and obtain, with probability at least $1 - 2 \exp(-2n)$, that

$$\|\hat{v}\|_1 \lesssim \frac{\|Z\|_2}{\sqrt{\log(ep/n)}} + \|Z\|_\infty.$$

By Lemma B.1, we have with probability at least $1 - n^{-2}$ that

$$\|Z\|_\infty \lesssim \varepsilon + \max_{i \in \mathcal{O}} |\langle X_i, \beta^* \rangle| \leq \varepsilon + \max_{i \in [n]} |\langle X_i, \beta^* \rangle| \lesssim \varepsilon + \log^{1/2+\xi}(n).$$

It is left to bound $\|Z\|_2$. By the triangle inequality, we have that

$$\|Z\|_2 \leq 2 \sqrt{\sum_{i \in \mathcal{O}} |\langle X_i, \beta^* \rangle|^2} + \sqrt{\sum_{i=1}^n f_\varepsilon(\langle X_i, \beta^* \rangle)^2}. \quad (18)$$

By Lemma B.1, we have with probability at least $1 - n^{-2}$ that

$$\sum_{i \in \mathcal{O}} |\langle X_i, \beta^* \rangle|^2 \leq |\mathcal{O}| \max_{i \in [n]} |\langle X_i, \beta^* \rangle|^2 \lesssim |\mathcal{O}| \log^{1+2\xi}(n).$$

We next bound the second term on the right-hand side in (18). Indeed, we have that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n f_\varepsilon(\langle X_i, \beta^* \rangle)^2 &= \frac{1}{n} \sum_{i=1}^n (|\langle X_i, \beta^* \rangle| - \varepsilon)^2 \mathbf{1}(|\langle X_i, \beta^* \rangle| \leq \varepsilon) \\ &\leq \frac{\varepsilon^2}{n} \sum_{i=1}^n \mathbf{1}(|\langle X_i, \beta^* \rangle| \leq \varepsilon). \end{aligned} \quad (19)$$

Let $p_\varepsilon = \mathbb{P}(|\langle X_1, \beta^* \rangle| \leq \varepsilon)$. By Hoeffding's inequality ([26, Theorem 3.1.2]), we have with probability at least $1 - \exp(-2n\varepsilon^{2\alpha})$ that

$$\frac{1}{n} \sum_{i=1}^n f_\varepsilon(\langle X_i, \beta^* \rangle)^2 \leq \frac{\varepsilon^2}{n} \sum_{i=1}^n \mathbf{1}(|\langle X_i, \beta^* \rangle| \leq \varepsilon) \leq \varepsilon^2(p_\varepsilon + \varepsilon^\alpha) \lesssim \varepsilon^{2+\alpha},$$

where the last inequality holds by the anti-concentration assumption and for $p^{-1} \leq \varepsilon \leq 1$. Hence, summarizing, we have with probability at least $1 - e^{-2n\varepsilon^{2\alpha}} - n^{-2} - 2 \exp(-2n)$ that

$$\|\hat{v}\|_1 \lesssim \frac{\log^{1/2+\xi}(n)|\mathcal{O}|^{1/2} + n^{1/2}\varepsilon^{1+\alpha/2}}{\sqrt{\log(ep/n)}} + \varepsilon + \log(n)^{1/2+\xi}.$$

Choosing

$$\varepsilon \asymp \left(\frac{s \log(ep/n)}{n} + \frac{|\mathcal{O}| \log(n)^{1+2\xi}}{n} \right)^{1/(2+\alpha)}$$

concludes the proof. \blacksquare

4.2.2. Proof of Proposition 2.1.

Proof. By the dual formulation of the margin (see Appendix A), we have that

$$\gamma = \inf_{w: w_i \geq 0 \forall i \in [n], \|w\|_1 = 1} \left\| \sum_{i=1}^n w_i y_i X_i \right\|_{\infty}. \quad (20)$$

Hence, for proving an upper bound it suffices to find an appropriate weighting w . For τ_n a sequence to be defined and τ_n^{-1} taking integer values, we define

$$w_i = \begin{cases} \tau_n & i \text{ is among indices of } \tau_n^{-1} \text{ smallest entries of } (\langle X_i, \beta^* \rangle)_{i=1}^n, \\ 0 & \text{otherwise.} \end{cases}$$

We use this choice of w to upper bound γ . We denote the projector onto the space spanned by β^* by P , $P := \beta^*(\beta^*)^T$, and define its orthogonal complement $P^{\perp} := I_p - P$. We have that

$$\left\| \sum_{i=1}^n w_i y_i X_i \right\|_{\infty} \leq \left\| \sum_{i=1}^n w_i y_i P X_i \right\|_{\infty} + \left\| \sum_{i=1}^n w_i y_i P^{\perp} X_i \right\|_{\infty}. \quad (21)$$

We treat the two terms separately. For the first term, we have by [27, Theorems 5 and 7] that

$$\begin{aligned} \mathbb{E} \left\| \sum_{i=1}^n w_i y_i P X_i \right\|_{\infty} &= \mathbb{E} \left\| \sum_{i=1}^n w_i |\langle X_i, \beta^* \rangle| \beta^* \right\|_{\infty} = \|\beta^*\|_{\infty} \mathbb{E} \sum_{i=1}^n w_i |\langle X_i, \beta^* \rangle| \\ &\lesssim \|\beta^*\|_{\infty} \sum_{k=1}^{\tau_n^{-1}} \frac{\tau_n k \log(k+1)}{n} \lesssim \frac{\|\beta^*\|_{\infty} (\tau_n^{-1} + 1) \log(p)}{n}. \end{aligned}$$

We next bound the second term on the right-hand side in (21). Observe that

$$y_i = \operatorname{sgn}(\langle X_i, \beta^* \rangle) = \operatorname{sgn}(\langle P X_i, \beta^* \rangle),$$

and hence y_i is independent of $P^\perp X_i$. Likewise, w is a function of $(PX_i)_i$ and not $(P^\perp X_i)_i$, and hence w and $P^\perp X_i$ are independent for each i . We conclude that

$$\left(\sum_i w_i y_i P^\perp X_i\right)_j \sim \mathcal{N}(0, \|w\|_2^2 \langle e_j, P^\perp e_j \rangle).$$

Hence, using a standard Chernoff-bound, we obtain

$$\mathbb{E} \left\| \sum_i w_i y_i P X_i \right\|_\infty \lesssim \sqrt{\log(p) \|w\|_2^2} = \sqrt{\log(p) \tau_n}.$$

Hence, we obtain

$$\mathbb{E} \gamma \leq \mathbb{E} \left\| \sum_{i=1}^n w_i y_i X_i \right\|_\infty \lesssim \frac{\|\beta^*\|_\infty \log(p)}{n \tau_n} + \sqrt{\log(p) \tau_n}.$$

The final result is obtained by choosing

$$\tau_n^{-1} = \left\lceil \left(\frac{n}{\|\beta^*\|_\infty \log(p)} \right)^{2/3} \right\rceil. \quad \blacksquare$$

4.3. Proof Proposition 4.1

4.3.1. Proof of the first part of Proposition 4.1. In this subsection, we present a result holding only under the weak moment assumption. We will see in the next section how to improve this result when assuming a small deviation assumption.

Proposition 4.4. *Assume that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$. Suppose that $r_n > 0$ satisfies*

$$r_n \lesssim \sqrt{\frac{n}{\log(p)}}.$$

Then, with probability at least

$$1 - np^{-2} - 2 \exp\left(-\frac{cn}{r_n^2 \log^{2\zeta}(p)}\right),$$

for any $\beta \in \mathbb{R}^p$ such that $\|\beta\|_1 \leq r_n$ and $\min_{i \in [n]} y_i \langle X_i, \beta \rangle \geq 1$, we have that

$$\|\beta\|_2 \geq 1/2.$$

Proof. For $r_n > 0$, let $\beta \in \mathbb{R}^p$ such that $\|\beta\|_1 \leq r_n$ and $\min_{i \in [n]} y_i \langle X_i, \beta \rangle \geq 1$. Thus, we have

$$\frac{1}{n} \sum_{i=1}^n |\langle X_i, \beta \rangle| \geq 1. \tag{22}$$

We proceed by contradiction. Assume that $\|\beta\|_2 \leq 1/2$. In this case, we show that equation (22) is not satisfied with large probability, concluding the proof by contradiction.

For $i \in [n]$, using Hölder's inequality, we have that

$$|\langle X_i, \beta \rangle| \leq \|X_i\|_\infty \|\beta\|_1 \leq r_n \|X_i\|_\infty \lesssim r_n \log^\xi(p),$$

where the last inequality follows from Lemma B.1 and holds with probability at least $1 - p^{-2}$. Thus, with probability at least $1 - n/p^2$ we have, for all $i \in [n]$, that

$$|\langle X_i, \beta \rangle| \lesssim r_n \log^\xi(p).$$

Hence, conditioning on this event and using the bounded differences inequality ([26, Theorem 3.3.14]), we obtain with probability at least

$$1 - 2 \exp\left(-\frac{cn}{r_n^2 \log^{2\xi}(p)}\right) - n/p^2$$

that

$$\begin{aligned} \sup_{\beta \in r_n B_1^p \cap (1/2)B_2^p} \frac{1}{n} \sum_{i=1}^n |\langle X_i, \beta \rangle| &\leq \sup_{\beta \in r_n B_1^p \cap (1/2)B_2^p} \mathbb{E} |\langle X_1, \beta \rangle| \\ &+ \mathbb{E} \sup_{\beta \in r_n B_1^p \cap (1/2)B_2^p} \frac{1}{n} \sum_{i=1}^n |\langle X_i, \beta \rangle| - \mathbb{E} |\langle X_i, \beta \rangle| + \frac{1}{4}. \end{aligned}$$

By Jensen's inequality and the fact that X is isotropic with unit variance, we obtain that

$$\sup_{\beta \in r_n B_1^p \cap (1/2)B_2^p} \mathbb{E} |\langle X_1, \beta \rangle| \leq 1/2.$$

Moreover, we have that

$$\begin{aligned} \mathbb{E} \sup_{\beta \in r_n B_1^p \cap (1/2)B_2^p} \frac{1}{n} \sum_{i=1}^n |\langle X_i, \beta \rangle| - \mathbb{E} |\langle X_i, \beta \rangle| &\leq \mathbb{E} \sup_{\beta \in r_n B_1^p \cap (1/2)B_2^p} \frac{4}{n} \sum_{i=1}^n \sigma_i \langle X_i, \beta \rangle \\ &\lesssim r_n \sqrt{\frac{\log(p)}{n}}, \end{aligned}$$

where $(\sigma_i)_{i=1}^n$ are i.i.d. Rademacher random variables independent from $(X_i)_{i=1}^n$. We used in the first line the symmetrization and contraction principles ([26, Theorems 3.1.21 and 3.2.1]), and Proposition B.2 in the second line to bound the Rademacher complexity.

The condition on r_n shows that

$$\sup_{\beta \in r_n B_1^p \cap (1/2)B_2^p} \frac{1}{n} \sum_{i=1}^n |\langle X_i, \beta \rangle| < 1,$$

and the contradiction is established. \blacksquare

4.3.2. Proof of the second part of Proposition 4.1: Small deviation assumption.

In this subsection, we show how to prove the second part of Proposition 4.1 under the small deviation assumption (Definition 2.3).

Proposition 4.5. *Assume $p \gtrsim n$ and that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$. Moreover, assume that \mathbb{X} fulfils a small deviation assumption (Definition 2.3) with constant $\theta > 0$. Let $r_n \geq 1$, define*

$$\tau_n = c \left(\frac{n}{\log^{2\zeta+1}(p) \log(n) r_n^2} \right)^{1/\theta},$$

and suppose that $\tau_n \gtrsim 1$. Then, with probability at least $1 - p^{-1}$ for any $\beta \in \mathbb{R}^p$ such that $\|\beta\|_1 \leq r_n$ and $\min_{i \in [n]} y_i \langle X_i, \beta \rangle \geq 1$, we have that

$$\|\beta\|_1 / \|\beta\|_2 \lesssim r_n / \tau_n.$$

Proof of Proposition 4.5. Let $\{e_j\}_{j=1}^p$ be the set of standard unit vectors in \mathbb{R}^p and $\mathcal{D} := \{\pm e_j\} \cup \{0\} \subset \mathbb{R}^p$ be the set of vectors with all entries equal to zero possibly except just one, where the value is then ± 1 . For $m \in \mathbb{N}$, we define Maurey’s set

$$\mathcal{Z}_m := \left\{ z = \frac{1}{m} \sum_{k=1}^m z_k, z_k \in \mathcal{D} \forall k \right\}.$$

Take

$$m = c \log^{2\zeta}(p) \log(n) r_n^2,$$

and define the event

$$\mathcal{E}_{\max} := \{ \|\mathbb{X}\|_\infty \leq c \log^\zeta(p) \},$$

and observe that by Lemma B.1, the event \mathcal{E}_{\max} occurs with probability at least $1 - p^{-1}$ as $p \gtrsim n$. Then, by Lemma 4.2, for all $\beta \in r_n B_1^p$ such that $\|\beta\|_2 \lesssim \tau_n$, there exists a vector $z_\beta \in r_n \mathcal{Z}_m$ such that on \mathcal{E}_{\max} , we have

$$\max_{1 \leq i \leq n} |\langle X_i, \beta \rangle - \langle X_i, z_\beta \rangle| \lesssim \log^\zeta(p) r_n \sqrt{\frac{\log(2n)}{m}} \leq \frac{1}{2}$$

as well as $\|\beta - z_\beta\|_2 \leq 1/2$ for m defined previously. Thus, by assumption on τ_n , we also have

$$\|z_\beta\|_2 \lesssim \tau_n + 1/2 \lesssim \tau_n.$$

In other words, on \mathcal{E}_{\max} we have that

$$\{z_\beta : \|\beta\|_1 \leq r_n, \|\beta\|_2 \leq c \tau_n\} \subset r_n \mathcal{Z}_m \cap \{\beta : \|\beta\|_2 \leq c \tau_n\} =: \mathcal{Z}_m(r_n, \tau_n).$$

We invoke that

$$\begin{aligned} & \left\{ \sup_{\beta \in r_n B_1^p \cap \{\beta: \|\beta\|_2 \leq \tau_n\}} \min_{1 \leq i \leq n} |\langle X_i, \beta \rangle| \geq 1 \right\} \cap \mathcal{E}_{\max} \\ & \subseteq \left\{ \max_{z \in \mathcal{Z}_m(r_n, \tau_n)} \min_{1 \leq i \leq n} |\langle X_i, z \rangle| \geq \frac{1}{2} \right\}. \end{aligned}$$

For all $z \in \mathcal{Z}_m(r_n, \tau_n)$ and $i \in [n]$, by the small deviation assumption (Definition 2.3), we have

$$\mathbb{P}\left(|\langle X_i, z \rangle| \leq \frac{1}{2}\right) \geq \mathbb{P}\left(\frac{|\langle X_i, z \rangle|}{\|z\|_2} \leq \frac{c}{\tau_n}\right) \gtrsim \tau_n^{-\theta}.$$

Hence,

$$\mathbb{P}\left(|\langle X_i, z \rangle| \geq \frac{1}{2}\right) \leq (1 - c\tau_n^{-\theta}) \leq \exp[-c\tau_n^{-\theta}],$$

and thus we obtain

$$\mathbb{P}\left(\min_{1 \leq i \leq n} |\langle X_i, z \rangle| \geq \frac{1}{2}\right) \leq \exp[-cn\tau_n^{-\theta}].$$

Since

$$|\mathcal{Z}_m(r_n, \tau_n)| \leq |\mathcal{Z}_m| \leq (2p + 1)^m,$$

we obtain by a union bound that

$$\mathbb{P}\left(\max_{z \in \mathcal{Z}_m(r_n, \tau_n)} \min_{1 \leq i \leq n} |\langle X_i, z \rangle| \geq \frac{1}{2}\right) \leq \exp[m \log(2p + 1) - cn\tau_n^{-\theta}].$$

We conclude that

$$\begin{aligned} & \mathbb{P}\left(\sup_{\beta \in r_n B_1^p \cap \{\beta: \|\beta\|_2 \leq c\tau_n\}} \min_{1 \leq i \leq n} |\langle X_i, \beta \rangle| \geq 1\right) \\ & \leq \exp[m \log(2p + 1) - cn\tau_n^{-\theta}] + \mathbb{P}(\mathcal{E}_{\max}^c) \\ & \leq \exp(-cn\tau_n^{-\theta}) + p^{-1} \end{aligned}$$

from our choice of m and applying Lemma B.1. ■

4.4. Tessellation: Proof of Proposition 4.2

Proof. For $a > 0$ and η defined in equation (14), let $\beta \in aB_1^p \cap \mathcal{S}^{p-1} \cap \mathcal{B}(a, \eta)$. By Lemma 4.2, there exists z_β in \mathcal{Z}_m such that

$$\max_{i \in [n]} |\langle X_i, \beta - z_\beta \rangle| \mathbb{1}_{\mathcal{E}_{\max}} \lesssim a \log^\zeta(p) \sqrt{\frac{\log(n)}{m}} \asymp \eta \quad \text{and} \quad \|\beta - z_\beta\|_2 \lesssim \frac{1}{\sqrt{m}},$$

where $\mathcal{E}_{\max} := \{\|\mathbb{X}\|_{\infty} \leq c \log^{\zeta}(p)\}$ and for $m = ca^2 \log^{2\zeta}(p) \log(n)/\eta^2$. We note that by Lemma B.1, \mathcal{E}_{\max} occurs with probability at least $1 - np^{-2}$. In particular, we have

$$1/2 \leq 1 - \eta \leq \|z_{\beta}\|_2 \leq 1 + \eta \leq 3/2$$

for η small enough. Let $z_{\beta} \in \mathcal{Z}_m$. By Bernstein's inequality ([26, Theorem 3.1.7]) and the anti-concentration assumption, we have that

$$\begin{aligned} \sum_{i=1}^n \mathbf{1}\{|\langle X_i, z_{\beta} \rangle| \leq \eta\} &\leq n(\mathbb{P}(|\langle X_1, z_{\beta} \rangle| \leq \eta) + \eta^{\alpha}) \\ &\leq n\left(\mathbb{P}\left(\left|\left\langle X_1, \frac{z_{\beta}}{\|z_{\beta}\|_2}\right\rangle\right| \leq 2\eta\right) + \eta^{\alpha}\right) \\ &\leq n\left(\sup_{\beta \in \mathcal{S}^{p-1}} \mathbb{P}(|\langle X_1, \beta \rangle| \leq 2\eta) + \eta^{\alpha}\right) \\ &\lesssim n\eta^{\alpha} \end{aligned}$$

with probability at least $1 - \exp(-cn\eta^{\alpha})$. Now, define

$$J := \{i \in [n] : \min_{z_{\beta} \in \mathcal{Z}_m} |\langle X_i, z_{\beta} \rangle| \geq \eta\}.$$

Using an bound over \mathcal{Z}_m and that $|\mathcal{Z}_m| \leq (2p + 1)^m$, we obtain that with probability at least

$$1 - 2 \exp[m \log(2p + 1) - cn\eta^{\alpha}] \geq 1 - 2 \exp[-cn\eta^{\alpha}],$$

uniformly for $z_{\beta} \in \mathcal{Z}_m$, we have

$$|J^C| \lesssim \eta^{\alpha} n.$$

For $i \in J$ and working on the event \mathcal{E}_{\max} , we have that

$$|\langle X_i, z_{\beta} \rangle| \geq \eta \quad \text{and} \quad |\langle X_i, \beta - z_{\beta} \rangle| < \eta,$$

and hence $\langle X_i, \beta \rangle$ and $\langle X_i, z_{\beta} \rangle$ have matching signs.

Hence, for $\beta \in aB_1^p \cap \mathcal{S}^{p-1} \cap \mathcal{B}(a, \eta)$ and working on the event \mathcal{E}_{\max} , we have

$$\begin{aligned} &\sum_{i=1}^n \mathbf{1}\{\text{sgn}(\langle X_i, \beta \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)\} \\ &\geq \sum_{i \in J} \mathbf{1}\{\text{sgn}(\langle X_i, \beta \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)\} \\ &= \sum_{i \in J} \mathbf{1}\{\text{sgn}(\langle X_i, z_{\beta} \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)\} \\ &\geq \sum_{i=1}^n (\mathbf{1}\{\text{sgn}(\langle X_i, z_{\beta} \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)\} - c\eta^{\alpha}). \end{aligned}$$

Applying Bernstein's inequality ([26, Theorem 3.1.7]), we have that

$$\sum_{i=1}^n \mathbf{1}(\text{sgn}(\langle X_i, z_\beta \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)) \geq (d(z_\beta, \beta^*) - c\eta^\alpha)n$$

with probability at least $1 - \exp(-cn\eta^\alpha)$. We next lower bound $d(z_\beta, \beta^*)$. Indeed, arguing as above, we have that

$$\begin{aligned} d(z_\beta, \beta^*) &= \mathbb{P}(\text{sgn}(\langle X, z_\beta \rangle) \neq \text{sgn}(\langle X, \beta^* \rangle) | (y_i, X_i)_{i=1}^n) \\ &\geq \mathbb{P}(\text{sgn}(\langle X, z_\beta \rangle) \neq \text{sgn}(\langle X, \beta^* \rangle), |\langle X, z_\beta \rangle| \geq \eta, |\langle X, z_\beta - \beta \rangle| < \eta | (y_i, X_i)_{i=1}^n) \\ &= \mathbb{P}(\text{sgn}(\langle X, \beta \rangle) \neq \text{sgn}(\langle X, \beta^* \rangle), |\langle X, z_\beta \rangle| \geq \eta, |\langle X, z_\beta - \beta \rangle| < \eta | (y_i, X_i)_{i=1}^n) \\ &\geq d(\beta, \beta^*) - \mathbb{P}(|\langle X, z_\beta \rangle| \leq \eta | (y_i, X_i)_{i=1}^n) - \mathbb{P}(|\langle X, z_\beta - \beta \rangle| \geq \eta | (y_i, X_i)_{i=1}^n). \end{aligned}$$

By the anti-concentration assumption (Definition 2.2), we have that

$$\mathbb{P}(|\langle X, z_\beta \rangle| \leq \eta | (y_i, X_i)_{i=1}^n) \lesssim \eta^\alpha.$$

Moreover, by our choice of m and Lemma B.1, we have that

$$\mathbb{P}(|\langle X, z_\beta - \beta \rangle| > \eta | (y_i, X_i)_{i=1}^n) \lesssim n^{-2} \lesssim \eta^\alpha.$$

Hence, and since $d(\beta, \beta^*) \gtrsim \eta^\alpha$, we obtain when the constant in the definition of $\mathcal{B}(a, \eta)$ is large enough that

$$d(z_\beta, \beta^*) \gtrsim \eta^\alpha.$$

Hence, taking another union bound over \mathcal{Z}_m and \mathcal{E}_{\max} and for the constant in the definition of $\mathcal{B}(a, \eta)$ large enough, we obtain with probability at least

$$1 - 2 \exp[m \log(2p + 1) - c\eta^\alpha n] - np^{-2} \geq 1 - 2 \exp[-c\eta^\alpha n] - np^{-2},$$

uniformly for $\beta \in aB_1^p \cap \mathcal{S}^{p-1} \cap \mathcal{B}(a, \eta)$ that

$$\sum_{i=1}^n \mathbf{1}(\text{sgn}(\langle X_i, \beta \rangle) \neq \text{sgn}(\langle X_i, \beta^* \rangle)) \gtrsim \eta^\alpha n,$$

which concludes the proof. ■

4.5. Rest of the proofs

4.5.1. Lemma 4.2. The following lemma applies Maurey's empirical method [12, 13] to construct a set \mathcal{Z}_m that approximates the B_1^p -ball well.

Lemma 4.2 (Maurey’s lemma). *Let $\{e_j\}_{j=1}^p$ be the set of standard unit vectors in \mathbb{R}^p and $\mathcal{D} := \{\pm e_j\} \cup \{0\} \subset \mathbb{R}^p$ be the set of vectors with all entries equal to zero except at most one, where the value is then ± 1 . For $m \in \mathbb{N}$, define Maurey’s set*

$$\mathcal{Z}_m := \left\{ z = \frac{1}{m} \sum_{k=1}^m z_k, z_k \in \mathcal{D} \forall k \right\}.$$

Then, we have that $\mathcal{Z}_m \subset B_1^p$ and that $|\mathcal{Z}_m| \leq (2p + 1)^m$. Moreover, for every $\beta \in B_1^p$, there exists a vector $z_\beta \in \mathcal{Z}_m$ such that for $\mathcal{E}_{\max} := \{\|\mathbb{X}\|_\infty \leq c \log^\zeta(p)\}$, we have that

$$\max_{1 \leq i \leq n} |\langle X_i, \beta \rangle - \langle X_i, z_\beta \rangle| \mathbb{1}_{\mathcal{E}_{\max}} \lesssim \log^\zeta(p) \sqrt{\frac{\log(n)}{m}} \quad \text{and} \quad \|\beta - z_\beta\|_2 \lesssim \frac{1}{\sqrt{m}}.$$

Proof. For $z \in \mathcal{D}$, either $\|z\|_1 = 1$ or $z \equiv 0$. Thus, for $\bar{z} := \sum_{k=1}^m z_k/m \in \mathcal{Z}_m$, we have

$$\|\bar{z}\|_1 \leq \sum_{k=1}^m \|z_k\|_1/m \leq 1.$$

It is moreover clear that $|\mathcal{D}| = (2p + 1)$. Therefore, $|\mathcal{Z}_m| \leq (2p + 1)^m$.

We now turn to the main part of the lemma. Let $\beta \in B_1^p$. Define a random vector $Z \in \mathcal{D}$ by

$$\mathbb{P}(Z = \text{sign}(\beta_j)e_j) = |\beta_j| \quad \text{for } \beta_j \neq 0, j = 1, \dots, p,$$

and

$$\mathbb{P}(Z = 0) = 1 - \|\beta\|_1.$$

Then

$$\mathbb{E}Z = \beta, \quad \mathbb{E}\|\beta - Z\|_2^2 = \|\beta\|_1 - \|\beta\|_2^2 \leq \|\beta\|_1 \leq 1.$$

Let Z_1, \dots, Z_m be independent copies of Z and define $\bar{Z} := \sum_{k=1}^m Z_k/m$. Then we get

$$\mathbb{E}\|\beta - \bar{Z}\|_2^2 \leq \frac{1}{m}.$$

Let $\sigma_1, \dots, \sigma_m$ be a Rademacher sequence independent of $(\mathbb{X}, (Z_1, \dots, Z_m))$. Then we have by the symmetrization inequality ([26, Theorem 3.1.21]) that

$$\mathbb{E} \left[\max_{1 \leq i \leq n} |\langle X_i, \beta \rangle - \langle X_i, \bar{Z} \rangle| \mid \mathbb{X} \right] \leq \frac{2}{m} \mathbb{E} \left[\max_{1 \leq i \leq n} \left| \sum_{k=1}^m \sigma_k \langle X_i, \bar{Z}_k \rangle \right| \mid \mathbb{X} \right].$$

Further, for $i = 1, \dots, n$, we have that

$$\sum_{k=1}^m \langle X_i, Z_k \rangle^2 \leq m \|X_i\|_\infty^2 \leq m \|\mathbb{X}\|_\infty^2.$$

Thus, we obtain

$$\mathbb{E} \left[\max_{1 \leq i \leq n} \left| \sum_{k=1}^m \sigma_k \langle X_i, Z_k \rangle \right| \middle| \mathbb{X}, Z_1, \dots, Z_m \right] \leq \sqrt{2 \log(2n)} \sqrt{m} \|\mathbb{X}\|_\infty.$$

Hence, and since

$$\mathcal{E}_{\max} = \{ \|\mathbb{X}\|_\infty \leq c \log^\zeta(p) \},$$

we obtain that

$$\mathbb{E} \left[\max_{1 \leq i \leq n} |\langle X_i, \beta \rangle - \langle X_i, \bar{Z} \rangle| \middle| \mathcal{E}_{\max} \right] \lesssim \log^\zeta(p) \sqrt{\frac{\log(n)}{m}}.$$

Invoking Jensen's inequality and $\mathbb{E} \|\beta - \bar{Z}\|_2^2 \leq 1/m$, we have that

$$\mathbb{E} \|\beta - \bar{Z}\|_2 \lesssim 1/\sqrt{m}.$$

Hence, we obtain that

$$\mathbb{E} \left[\max_{1 \leq i \leq n} |\langle X_i, \beta \rangle - \langle X_i, \bar{Z} \rangle| \middle| \mathcal{E}_{\max} + \log^\zeta(p) \log^{1/2}(n) \|\beta - \bar{Z}\|_2 \right] \lesssim \log^\zeta(p) \sqrt{\frac{\log(n)}{m}},$$

and hence there exists at least one $z_\beta \in \mathcal{Z}_m$ with the desired properties. \blacksquare

4.5.2. Proof of Lemma 2.1.

Proof. The proof follows closely the arguments in [55]. First, note that rescaling $\mathbb{X} = \mathbb{X}/\|\mathbb{X}\|_\infty$ does not change the approximating properties of $\tilde{\beta}_T$ for the max ℓ_1 -margin. Indeed, if $\tilde{\beta}_T$ fulfils

$$\min_{1 \leq i \leq n} \frac{y_i \langle X_i / \|\mathbb{X}\|_\infty, \tilde{\beta}_T \rangle}{\|\tilde{\beta}_T\|_1} \geq \frac{1}{2} \max_{\beta \neq 0} \min_{1 \leq i \leq n} \frac{y_i \langle X_i / \|\mathbb{X}\|_\infty, \beta \rangle}{\|\beta\|_1} =: \gamma_R = \gamma / \|\mathbb{X}\|_\infty,$$

then, by linearity, $\tilde{\beta}_T$ also fulfils

$$\min_{1 \leq i \leq n} \frac{y_i \langle X_i, \tilde{\beta}_T \rangle}{\|\tilde{\beta}_T\|_1} \geq \frac{1}{2} \max_{\beta \neq 0} \min_{1 \leq i \leq n} \frac{y_i \langle X_i, \beta \rangle}{\|\beta\|_1} = \gamma.$$

Henceforth, we work with the rescaled data $\mathbb{X}/\|\mathbb{X}\|_\infty$, which, in slight abuse of notation, we also denote by \mathbb{X} . Note, that by definition $\|\mathbb{X}\|_\infty \leq 1$. Define the exponential loss

$$\ell(\beta) := \frac{1}{n} \sum_{i=1}^n \exp(-y_i \langle X_i, \beta \rangle).$$

Note that

$$\begin{aligned}\nabla\ell(\beta) &= -\frac{1}{n}\sum_{i=1}^n y_i X_i \exp(-y_i \langle X_i, \beta \rangle), \\ \nabla^2\ell(\beta) &= \frac{1}{n}\sum_{i=1}^n X_i X_i^T \exp(-y_i \langle X_i, \beta \rangle).\end{aligned}$$

Hence, we have that

$$-\langle \nabla\ell(\tilde{\beta}_t), v_t \rangle = \alpha_t \ell(\tilde{\beta}_t).$$

Moreover, note that $|\alpha_t| \leq \sum w_{t,i} |\langle X_i, v_t \rangle| \leq 1$. By second order Taylor expansion, we obtain that

$$\ell(\tilde{\beta}_{t+1}) \leq \ell(\tilde{\beta}_t) + \epsilon \alpha_t \langle \nabla\ell(\tilde{\beta}_t), v_t \rangle + \frac{1}{2} \sup_{r \in [0,1]} \langle \nabla^2\ell(\tilde{\beta}_t + r\epsilon\alpha_t v_t) v_t, v_t \rangle.$$

We next bound the Hessian above. Indeed, we have for any r that

$$\begin{aligned}\langle \nabla^2\ell(\tilde{\beta}_t + r\epsilon\alpha_t v_t) v_t, v_t \rangle &= \frac{1}{n} \sum_{i=1}^n \langle X_i, v_t \rangle^2 \epsilon^2 \alpha_t^2 \exp(-y_i \langle X_i, \tilde{\beta}_t + r\epsilon\alpha_t v_t \rangle) \\ &\leq \epsilon^2 \alpha_t^2 \exp(r|\alpha_t|\epsilon) \ell(\tilde{\beta}_t) \leq \epsilon^2 \alpha_t^2 e^\epsilon \ell(\tilde{\beta}_t).\end{aligned}$$

Hence, we can further bound

$$\begin{aligned}\ell(\tilde{\beta}_{t+1}) &\leq \ell(\tilde{\beta}_t) + \epsilon \alpha_t \langle \nabla\ell(\tilde{\beta}_t), v_t \rangle + \frac{\epsilon^2 \alpha_t^2 e^\epsilon}{2} \ell(\tilde{\beta}_t) \\ &\leq \ell(\tilde{\beta}_t) \left(1 - \epsilon \alpha_t^2 + \frac{3\epsilon^2 \alpha_t^2}{2} \right) \leq \ell(\tilde{\beta}_t) \exp\left(-\epsilon \left(\alpha_t^2 - \frac{3\epsilon \alpha_t^2}{2} \right)\right),\end{aligned}$$

and hence we obtain

$$\ell(\tilde{\beta}_T) \leq \exp\left(-\epsilon \sum_{t=1}^T \left(\alpha_t^2 - \frac{3\epsilon \alpha_t^2}{2} \right)\right).$$

Moreover, we have that

$$\|\tilde{\beta}_T\|_1 = \left\| \sum_{t=1}^T \epsilon \alpha_t v_t \right\|_1 \leq \epsilon \sum_{t=1}^T |\alpha_t|.$$

In addition, we note that by the dual formulation of the margin (see Appendix A) and definition of v_t and α_t , we have that

$$|\alpha_t| = \left\| \sum w_{t,i} y_i X_i \right\|_\infty \geq \inf_{w: w_i \geq 0 \forall i, \|w\|_1 = 1} \left\| \sum w_i y_i X_i \right\|_\infty = \gamma_R.$$

Hence, by Markov's inequality and since $3\epsilon/2 < 1$, we obtain for any positive x that

$$\begin{aligned} \sum_{i=1}^n \mathbf{1}_{\{y_i \langle X_i, \tilde{\beta}_T \rangle \leq \|\tilde{\beta}_T\|_1 x\}} &\leq \sum_{i=1}^n \exp(\|\tilde{\beta}_T\|_1 x - y_i \langle X_i, \tilde{\beta}_T \rangle) = n \ell(\tilde{\beta}_T) \exp(\|\tilde{\beta}_T\|_1 x) \\ &\leq \exp\left(\log(n) - \epsilon \sum_{t=1}^T |\alpha_t| \left(|\alpha_t| - x - \frac{3\epsilon|\alpha_t|}{2}\right)\right) \\ &\leq \exp\left(\log(n) - \epsilon \sum_{t=1}^T |\alpha_t| \left(\gamma_R - x - \frac{3\epsilon\gamma_R}{2}\right)\right). \end{aligned}$$

Hence, choosing $x = \frac{1}{2}\gamma_R$ and using that $\epsilon \leq 1/6$ and that with probability at least $1 - np^{-2}$, we have by Lemma B.1 that

$$T > \frac{2 \log(n)}{3\epsilon^2 \gamma_R^2} = \frac{2 \log(n) \|X\|_\infty^2}{3\epsilon^2 \gamma^2},$$

we obtain

$$\begin{aligned} \sum_{i=1}^n \mathbf{1}_{\{y_i \langle X_i, \tilde{\beta}_T \rangle \leq \|\tilde{\beta}_T\|_1 x\}} &\leq \sum_{i=1}^n \exp(\|\tilde{\beta}_T\|_1 x - y_i \langle X_i, \tilde{\beta}_T \rangle) \\ &\leq \exp(\log(n) - 3T\epsilon^2 \gamma_R^2 / 2) < e^0 = 1. \end{aligned}$$

Since $\sum_{i=1}^n \mathbf{1}_{\{y_i \langle X_i, \tilde{\beta}_T \rangle \leq \|\tilde{\beta}_T\|_1 \frac{1}{2}\gamma_R\}}$ can only take values in $\{0, 1, \dots, n\}$ this implies that

$$\sum_{i=1}^n \mathbf{1}_{\{y_i \langle X_i, \tilde{\beta}_T \rangle \leq \|\tilde{\beta}_T\|_1 \frac{1}{2}\gamma_R\}} = 0,$$

and hence the result follows. \blacksquare

4.5.3. Proof of Corollary 2.3.

Proof. For Gaussian distributed features it is clear that the weak moment assumption with $\zeta = 1/2$ is satisfied. Moreover, since for $\beta \in \mathcal{S}^{p-1}$ we have that $\langle X, \beta \rangle \sim \mathcal{N}(0, 1)$, for any $0 < \epsilon \leq 1$, we then have

$$\sup_{\beta \in \mathcal{S}^{p-1}} \mathbb{P}(|\langle X, \beta \rangle| \leq \epsilon) = \int_{-\epsilon}^{\epsilon} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \asymp \epsilon,$$

and hence both the anti-concentration and small deviation assumptions are fulfilled with $\alpha = \theta = 1$. Finally, to show (13), note that by Grothendieck's identity ([56, Lemma 3.6.6]) and as the geodesic distance on the sphere is lower bounded by the Euclidean distance, we have that

$$d(\beta^*, \tilde{\beta}_T) = \frac{\arccos(\langle \beta^*, \tilde{\beta}_T / \|\tilde{\beta}_T\|_2 \rangle)}{\pi} \geq \left\| \beta^* - \frac{\tilde{\beta}_T}{\|\tilde{\beta}_T\|_2} \right\|_2.$$

For the Student- t -distribution with at least $32 \log(p)$ degrees of freedom, Lemma 4.3 below proves that the weak moment assumption, the anti-concentration, and small deviation assumptions with $\alpha = \theta = 1$ are satisfied. Moreover, Lemma 4.4 below shows that in this case we can also lower bound

$$d(\tilde{\beta}_T, \beta^*) \gtrsim \left\| \beta^* - \frac{\tilde{\beta}_T}{\|\tilde{\beta}_T\|_2} \right\|_2. \quad \blacksquare$$

Lemma 4.3. *Suppose that $X = (x_j)_{j=1}^p$ with $x_j \stackrel{\text{i.i.d.}}{\sim} \sqrt{(d-2)/d} t_d$ for $d \in \mathbb{N}$, $d \geq 32 \log(p)$ and $p \gtrsim 1$. Then for any $0 \leq \varepsilon \leq 1$, we have*

$$\inf_{\beta \in \mathcal{S}^{p-1}} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) \gtrsim \varepsilon.$$

Moreover, under the same assumptions, we have that

$$\sup_{\beta \in \mathcal{S}^{p-1}} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) \lesssim \varepsilon + p^{-1}.$$

Finally, for $2q + 1 \leq d$ and $p \gtrsim 1$, we have that

$$(\mathbb{E}|x_1|^q)^{1/q} \lesssim \sqrt{q}.$$

Proof. Since $x_j \sim \sqrt{(d-2)/d} t_d$, we have that

$$x_j = \frac{\sqrt{((d-2)/d)} z_j}{\sqrt{\chi_{d,j}^2/d}},$$

where z denotes a standard Gaussian random variable and $\chi_{d,j}^2$ a chi-squared random variable with d degrees of freedom that is independent of z . For $\beta \in \mathcal{S}^{p-1}$, we have conditionally on the $\chi_{d,j}^2$ -variables that

$$\langle X, \beta \rangle \mid (\chi_{d,j}^2)_{j=1}^p \sim \mathcal{N}\left(0, \frac{d-2}{d} \sum_{i=1}^n \frac{\beta_i^2 d}{\chi_{d,j}^2}\right).$$

Hence, conditioning on the event where $\chi_{d,j}^2 \geq d/2$ for all $j \in [p]$, and using independence of z and the $\chi_{d,j}^2$ -variables and that $\|\beta\|_2 \leq 1$, we obtain that

$$\begin{aligned} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) &\geq \mathbb{P}(|z| \leq \varepsilon \sqrt{d/(2(d-2))}) \mathbb{P}\left(\min_{j \in [p]} \chi_{d,j}^2 \geq d/2\right) \\ &\gtrsim \varepsilon \mathbb{P}\left(\min_{j \in [p]} \chi_{d,j}^2 \geq d/2\right). \end{aligned}$$

It is left to lower bound the probability involving the minimum. By applying a lower tail bound for chi-square random variables ([36, Lemma 1]) and a union bound, we obtain

$$\mathbb{P}\left(\min_{j \in [p]} \chi_{d,j}^2 < d/2\right) \leq p\mathbb{P}(\chi_{d,1}^2 < d/2) \leq pe^{-d/16} \leq p^{-1} \leq \frac{1}{2},$$

using the conditions on d and p .

For the upper bound we argue similarly. We have

$$\begin{aligned} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) &\leq \mathbb{P}\left(\max_{j \in [p]} \chi_{d,j}^2 \geq 2d\right) + \mathbb{P}(|Z| \leq \varepsilon \sqrt{2d/(d-2)}) \\ &\lesssim \mathbb{P}\left(\max_{j \in [p]} \chi_{d,j}^2 \geq 2d\right) + \varepsilon. \end{aligned}$$

Applying an upper tail bound for chi-square random variables ([36, Lemma 1]), we obtain

$$\mathbb{P}\left(\max_{j \in [p]} \chi_{d,j}^2 > 2d\right) \leq p\mathbb{P}(\chi_{d,1}^2 > 2d) \leq pe^{-d/16} \leq p^{-1},$$

thus proving the claimed result.

Finally, for the claimed moment bound, integration by parts, and for Γ denoting the Gamma function, give

$$\mathbb{E}|x_1|^q = \frac{d^{q/2}\Gamma((q+1)/2)\Gamma((d-q)/2)}{\pi^{1/2}\Gamma(d/2)}.$$

We now only consider the case where q is uneven, as the other case follows along the same lines. Indeed, since $d \gtrsim q$, applying Gautschi's inequality and using that $\Gamma(z) \leq z^{z-1/2}$ for $z \geq 1$, we have that

$$\begin{aligned} &\frac{d^{q/2}\Gamma((q+1)/2)\Gamma((d-q)/2)}{\Gamma(d/2)} \\ &= \frac{d^{q/2}\Gamma((q+1)/2)\Gamma((d-q)/2)}{((d-2)/2 \cdot \dots \cdot (d-q-1)/2)\Gamma((d-q+1)/2)} \\ &\lesssim \frac{c^{(q-1)/2}d^{q/2}\Gamma((q+1)/2)\Gamma((d-q)/2)}{d^{(q-1)/2}\Gamma((d-q+1)/2)} \\ &\lesssim \frac{(cd)^{q/2}\Gamma((q+1)/2)}{d^{q/2}} \lesssim (cq)^{q/2}. \end{aligned}$$

Taking the q -th root concludes the proof. ■

Lemma 4.4. *Suppose that $X = (x_1, \dots, x_p)$ with $x_j \stackrel{\text{i.i.d.}}{\sim} \sqrt{(d-2)/d} t_d$ for $d \gtrsim \log(p)$ and $p \gtrsim 1$. Then for any $\beta, \tilde{\beta} \in \mathcal{S}^{p-1}$, we have*

$$\mathbb{P}(\text{sgn}(\langle X, \beta \rangle) \neq \text{sgn}(\langle X, \tilde{\beta} \rangle)) \gtrsim \|\beta - \tilde{\beta}\|_2.$$

Proof. Since $x_j \sim \sqrt{(d-2)/d} z_j$, we have that

$$x_j = \frac{\sqrt{(d-2)/d} z_j}{\sqrt{\chi_{d,j}^2/d}},$$

where z denotes a standard Gaussian random variable and $\chi_{d,j}^2$ a chi-squared random variable with d degrees of freedom that is independent of z . Denote

$$\beta_\chi = (\beta_j \sqrt{d/\chi_{d,j}^2})_{j \in [p]}$$

and note that on the event $\{d/2 \leq \chi_{d,j}^2 \leq 2d \ \forall 1 \leq j \leq p\}$, we have

$$\sqrt{1/2} \leq \|\beta_\chi\|_2 \leq \sqrt{2}.$$

Then, conditioning on the $\chi_{d,j}^2$ -variables and using Grothendieck's identity ([56, Lemma 3.6.6]), we have

$$\begin{aligned} \mathbb{P}(\text{sgn}(\langle X, \beta \rangle) \neq \text{sgn}(\langle X, \tilde{\beta} \rangle)) &= \frac{\mathbb{E} \arccos(\langle \beta_\chi / \|\beta_\chi\|_2, \tilde{\beta}_\chi / \|\tilde{\beta}_\chi\|_2 \rangle)}{\pi} \\ &\geq \mathbb{E} \left\| \frac{\beta_\chi}{\|\beta_\chi\|_2} - \frac{\tilde{\beta}_\chi}{\|\tilde{\beta}_\chi\|_2} \right\|_2. \end{aligned}$$

Using that β and $\tilde{\beta}$ have unit norm, we further bound

$$\begin{aligned} &\mathbb{E} \left\| \frac{\beta_\chi}{\|\beta_\chi\|_2} - \frac{\tilde{\beta}_\chi}{\|\tilde{\beta}_\chi\|_2} \right\|_2 \\ &\gtrsim \mathbb{E} [(\|\beta_\chi\|_2 \|\tilde{\beta}_\chi\|_2 - \tilde{\beta}_\chi \|\beta_\chi\|_2)_2 \mathbf{1}(d/2 \leq \chi_{d,j}^2 \leq 2d \ \forall 1 \leq j \leq p)] \\ &= \mathbb{E} \left[\sqrt{\sum_{j=1}^p \frac{(\beta_j \|\tilde{\beta}_\chi\|_2 - \tilde{\beta}_j \|\beta_\chi\|_2)^2}{\chi_{d,j}^2/d}} \mathbf{1}(d/2 \leq \chi_{d,j}^2 \leq 2d \ \forall 1 \leq j \leq p) \right] \\ &\gtrsim \mathbb{E} [(\|\beta\|_2 \|\tilde{\beta}_\chi\|_2 - \tilde{\beta} \|\beta_\chi\|_2)_2 \mathbf{1}(d/2 \leq \chi_{d,j}^2 \leq 2d \ \forall 1 \leq j \leq p)] \\ &\geq \min_{a,b \in [1/2,2]} \sqrt{a^2 + b^2 - 2ab \langle \beta, \tilde{\beta} \rangle} \mathbb{P}(d/2 \leq \chi_{d,j}^2 \leq 2d \ \forall 1 \leq j \leq p). \end{aligned}$$

By [36, Lemma 1], a union bound, and by our assumption on d and p , we have that

$$\mathbb{P}(d/2 \leq \chi_{d,j}^2 \leq 2d \ \forall 1 \leq j \leq p) \geq (1 - 2pe^{-d/16}) \geq 1/2.$$

Hence, it is left to lower bound the quadratic equation

$$\min_{a,b \in [1/2,1]} (a^2 + b^2 - 2ab \langle \beta, \tilde{\beta} \rangle).$$

If $\langle \tilde{\beta}, \beta \rangle \leq 0$, it is clear that the minimum is attained at $a = b = 1/2$. Conversely, if $\langle \tilde{\beta}, \beta \rangle > 0$, we have

$$\begin{aligned} \min_{a,b \in [1/2,1]} (a^2 + b^2 - 2ab \langle \beta, \tilde{\beta} \rangle) &\geq \min_{a \in \mathbb{R}, b \in [1/2,1]} (a^2 + b^2 - 2ab \langle \beta, \tilde{\beta} \rangle) \\ &= \min_{b \in [1/2,1]} b^2 (1 - \langle \beta, \tilde{\beta} \rangle^2) \gtrsim (1 - \langle \beta, \tilde{\beta} \rangle). \end{aligned}$$

since $0 < \langle \tilde{\beta}, \beta \rangle \leq 1$. Hence, summarizing, we have that

$$\min_{a,b \in [1/2,2]} \sqrt{a^2 + b^2 - 2ab \langle \beta, \tilde{\beta} \rangle} \gtrsim \sqrt{(2 - 2 \langle \beta, \tilde{\beta} \rangle)} = \|\beta - \tilde{\beta}\|_2,$$

thus concluding the proof. \blacksquare

4.5.4. Proof of Corollary 2.2.

Proof. The proof of Corollary 2.2 follows mainly from Lemma 4.5 below, which shows that the anti-concentration condition is satisfied for unimodal features with bounded density, and by noting that the weak moment assumption is satisfied for Laplace distributed features with $\zeta = 1$, for Student- t with at least $2 \log(p) + 1$ degrees of freedom by Lemma 4.3 with $\zeta = 1/2$, and for uniform and Gaussian features with $\zeta = 1/2$ as they are sub-Gaussian. \blacksquare

Lemma 4.5. *Suppose that $X = (x_1, \dots, x_p)$ consists of i.i.d. symmetric and unit variance scalar random variables with density f . Suppose that $\|f\|_\infty \lesssim 1$ and that f is unimodal, i.e. $f(aw) \geq f(w)$ for any $0 \leq a \leq 1$ and any $w \in \mathbb{R}$. Then, for $0 \leq \varepsilon \leq 1$, we have that*

$$\sup_{\beta \in \mathcal{S}^{p-1}} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) \lesssim \varepsilon^{1/2} \mathbb{E}|x_1|^3.$$

Proof. We consider two cases. If $\|\beta\|_\infty \leq \varepsilon^{1/2}$, then by the Berry–Essen theorem (e.g. [14, Theorem 3.6]), we have

$$\mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) \lesssim \varepsilon + \sum_{j=1}^p |\beta_j|^3 \mathbb{E}|x_1|^3 \lesssim \varepsilon + \varepsilon^{1/2} \mathbb{E}|x_1|^3 \quad (23)$$

$$\lesssim \varepsilon^{1/2} \mathbb{E}|x_1|^3. \quad (24)$$

If $\|\beta\|_\infty \geq \varepsilon^{1/2}$, we argue as follows. Assume, without loss of generality, that $|\beta_1| \geq \varepsilon^{1/2}$. We note that since x_1 is unimodal that $\beta_1 x_1$ is unimodal, too. Then, since $\beta_1 x_1$ is unimodal (see e.g. [2, Theorem 1]), we obtain that

$$\begin{aligned} \mathbb{P}(|\langle X, \beta \rangle| \leq \varepsilon) &= \mathbb{P}\left(|\beta_1 x_1 + \sum_{j=2}^p x_j \beta_j| \leq \varepsilon\right) \leq \mathbb{P}(|\beta_1 x_1| \leq \varepsilon) \leq \mathbb{P}(|x_1| \leq \varepsilon^{1/2}) \\ &\lesssim \varepsilon^{1/2} \leq \varepsilon^{1/2} \mathbb{E}|x_1|^3, \end{aligned}$$

as $\mathbb{E}|x_1|^3 \geq (\mathbb{E}|x_1|^2)^{3/2} = 1$ by Jensen's inequality. \blacksquare

A. Dual formulation of the max ℓ_1 -margin

We use Lagrangian duality to derive the dual version of the max ℓ_1 -margin. Recall that

$$\gamma = \max_{\beta \neq 0} \min_{1 \leq i \leq n} \frac{y_i \langle X_i, \beta \rangle}{\|\beta\|_1} = \frac{1}{\|\widehat{\beta}\|},$$

where we used Lemma 4.1, recalling that

$$\widehat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \text{ subject to } y_i \langle X_i, \beta \rangle \geq 1. \quad (25)$$

For every $\lambda \in \mathbb{R}^n$, define the Lagrangian $\mathcal{L}: \mathbb{R}^p \times \mathbb{R}^n \mapsto \mathbb{R}$ as

$$\mathcal{L}(\beta, \lambda) = \|\beta\|_1 + \sum_{i=1}^n \lambda_i (1 - y_i \langle \beta, X_i \rangle).$$

The dual problem of (25) is defined as

$$\sup_{\lambda \in \mathbb{R}_+^n} \inf_{\beta \in \mathbb{R}^p} \mathcal{L}(\beta, \lambda). \quad (26)$$

We have that

$$\begin{aligned} \inf_{\beta \in \mathbb{R}^p} \mathcal{L}(\beta, \lambda) &= \inf_{\beta \in \mathbb{R}^p} \left\{ \|\beta\|_1 + \sum_{i=1}^n \lambda_i (1 - y_i \langle \beta, X_i \rangle) \right\} \\ &= \sum_{i=1}^n \lambda_i - \sup_{\beta \in \mathbb{R}^p} \left\{ \left\langle \beta, \sum_{i=1}^n \lambda_i y_i X_i \right\rangle - \|\beta\|_1 \right\}. \end{aligned}$$

For any function $f: \mathbb{R}^p \mapsto \mathbb{R}$, the conjugate f^* is defined as

$$f^*(y) = \sup_{x \in \mathbb{R}^p} \{ \langle x, y \rangle - f(x) \}. \quad (27)$$

In particular (see [8, Example 3.26]), when $f(\beta) = \|\beta\|_1$, we have that

$$f^*(y) = \begin{cases} 0 & \text{if } y \in B_\infty, \\ \infty & \text{otherwise,} \end{cases} \quad (28)$$

where B_∞ is the unit ball with respect $\|\cdot\|_\infty$. From (27) and (28), the dual problem (26) can be rewritten as

$$\sup_{\lambda \in \mathbb{R}_+^n} \sum_{i=1}^n \lambda_i \text{ subject to } \left\| \sum_{i=1}^n y_i \lambda_i X_i \right\|_\infty \leq 1.$$

Since the X_i are linearly independent with probability 1 and $p > n$, the Moore–Penrose inverse of $\mathbb{X} = [X_1, \dots, X_n]$ exists, and hence there exists some β in \mathbb{R}^p such that $y_i \langle X_i, \beta \rangle = 1$ for $i = 1, \dots, n$. Hence, Slater’s condition is satisfied and consequently there is no duality gap. It follows that

$$\gamma = \frac{1}{\|\widehat{\beta}\|_1} = \inf_{w: w_i \geq 0 \forall i \in [n], \|w\|_1 = 1} \left\| \sum_{i=1}^n w_i y_i X_i \right\|_{\infty}.$$

B. Extra lemmas

B.1. Lemma B.1

Lemma B.1. *Let $X = (x_1, \dots, x_p)^T$ be a random vector where the x_j ’s are i.i.d. random variables that satisfy the weak moment assumption with $\zeta \geq 1/2$. Then, with probability at least $1 - p^{-2}$, we have that*

$$\|X\|_{\infty} \lesssim \log^{\zeta}(p). \quad (29)$$

Moreover, let X_1, \dots, X_n , $n \leq p$, be n i.i.d. copies of X and $\beta \in \mathcal{S}^{p-1}$. Then, with probability at least $1 - n^{-2}$, we have additionally

$$\max_{i \in [n]} |\langle X_i, \beta^* \rangle| \lesssim \log^{1/2+\zeta}(n).$$

Proof. We have that

$$\left(\mathbb{E} \left(\max_{j \in [p]} |x_j| \right)^q \right)^{1/q} \leq p^{1/q} \left(\mathbb{E} |x_1|^q \right)^{1/q} \lesssim p^{1/q} q^{\zeta}.$$

Hence, by Markov’s inequality, we have

$$\mathbb{P}(\|X\|_{\infty} > t) \leq e^{\log(p) + cq + \zeta q \log(q) - q \log(t)}.$$

Choosing $q = \log(p)$ and $t \asymp \log^{\zeta}(p)$ concludes the proof of the first claim.

For the second claim we argue as follows. By Rio’s version of the Marcinkiewicz–Zygmund inequality ([49, Theorem 2.1]), we have that

$$\begin{aligned} \left(\mathbb{E} \max_{i \in [n]} |\langle X_i, \beta \rangle|^q \right)^{1/q} &\leq n^{1/q} \left(\mathbb{E} |\langle X, \beta \rangle|^q \right)^{1/q} \\ &\leq n^{1/q} q^{1/2} \left(\sum_{j=1}^p |\beta_j|^2 \left(\mathbb{E} |x_j|^q \right)^{2/q} \right)^{1/2} \\ &\lesssim n^{1/q} q^{1/2+\zeta}. \end{aligned}$$

Arguing as before with $q = \log(n)$ concludes the proof. ■

B.2. Proposition B.1

Proposition B.1 ([35, Theorem 5]). *Let X_1, \dots, X_n be i.i.d. random vectors distributed as $X = (x_1, \dots, x_p)^T$, where the x_j 's are i.i.d. symmetric, zero mean and unit variance random variables that satisfy the weak moment assumption. For $Z \in \mathbb{R}^n$ and $\mathbb{X} = [X_1, \dots, X_n]$, define*

$$\hat{v} := \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \quad \text{such that } \mathbb{X}^T \beta = Z.$$

Assume that $p \gtrsim n$. Then, with probability at least $1 - 2 \exp(-2n)$, we have that

$$\|\hat{v}\|_1 \lesssim \frac{\|Z\|_2}{\sqrt{\log(ep/n)}} + \|Z\|_\infty.$$

B.3. Rademacher complexity under weak moment assumption

Proposition B.2. *Assume that $\mathbb{X} = (X_i)_{i \in [n]}$ has i.i.d. zero mean and unit variance entries and satisfies the weak moment assumption with $\zeta \geq 1/2$. For $a \in \mathbb{N}$, we have that*

$$\mathbb{E} \sup_{\beta \in aB_1^p \cap B_2^p} \frac{1}{n} \sum_{i=1}^n \sigma_i \langle X_i, \beta \rangle \lesssim a \sqrt{\frac{\log(p)}{n}}.$$

The proof of Proposition B.2 uses the following bound for sums of order statistics and will be presented below.

Lemma B.2. *Assume that $X = (x_1, \dots, x_p)^T$ has i.i.d. symmetric, zero mean and unit variance entries that satisfy the weak moment assumption with $\zeta \geq 1/2$. Then, for all $1 \leq k \leq p$, we have*

$$\mathbb{E} \left(\sum_{i=1}^k (x_i^*)^2 \right)^{1/2} \lesssim \log^\zeta(p) \sqrt{k},$$

where $(x_i^*)_i^p$ is a monotone non-increasing rearrangement of $(|x_i|)_{i=1}^p$.

Proof. The proof is a small adaptation from [41, Lemma 6.5], where $\zeta = 1/2$ is assumed. Fix $1 \leq j \leq p$, $1 \leq q \leq \log(p)$ and $t > 0$. By the weak moment assumption for some $c_1 > 0$, we have

$$\mathbb{P}(x_j^* \geq t) \leq \binom{p}{j} \mathbb{P}^j(|x_1| \geq t) \leq \binom{p}{j} \left(\frac{\mathbb{E}|x_1|^q}{t^q} \right)^j \leq \binom{p}{j} \left(\frac{c_1 q^\zeta}{t^q} \right)^j.$$

Since $\binom{p}{j} \leq \exp(j \log(p))$, taking $q = \log(p)$, we get

$$\mathbb{P}(x_j^* \geq t) \leq \left(\frac{c \log^\zeta(p)}{t} \right)^{j \log(p)}.$$

Hence, integrating out the tails and using Jensen's inequality it follows that

$$\mathbb{E} \left(\sum_{i=1}^k ((x_i^*)^2) \right)^{1/2} \leq \left(\mathbb{E} \sum_{i=1}^k (x_i^*)^2 \right)^{1/2} \lesssim \log^\zeta(p) \sqrt{k}. \quad \blacksquare$$

Proof of Proposition B.2. From [42, Equation 3.1], we have that

$$\begin{aligned} \mathbb{E} \sup_{\beta \in aB_1^p \cap B_2^p} \frac{1}{n} \sum_{i=1}^n \sigma_i \langle X_i, \beta \rangle &\leq 2 \mathbb{E} \sup_{\beta \in B_0^p(a^2) \cap B_2^p} \frac{1}{n} \sum_{i=1}^n \sigma_i \langle X_i, \beta \rangle \\ &= \frac{2}{\sqrt{n}} \mathbb{E} \sup_{\beta \in B_0^p(a^2) \cap B_2^p} \langle W, \beta \rangle, \end{aligned}$$

where $B_0^p(a^2) = \{\beta \in \mathbb{R}^p : \|\beta\|_0 \leq a^2\}$ and $W = n^{-1/2} \sum_{i=1}^n \sigma_i X_i$, and it follows that

$$\mathbb{E} \sup_{\beta \in aB_1^p \cap B_2^p} \frac{1}{n} \sum_{i=1}^n \sigma_i \langle X_i, \beta \rangle \leq \frac{2}{\sqrt{n}} \mathbb{E} \left(\sum_{i=1}^{a^2} (W_i^*)^2 \right)^{1/2}. \quad (30)$$

The (W_i) 's are centred random variables. For $1 \leq q \leq \log(p)$, using the Khintchine–Kahane inequality ([26, Proposition 3.2.8]) and Jensen's inequality, we have

$$\begin{aligned} (\mathbb{E} |W_j|^q)^{1/q} &= \left(\mathbb{E} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n \sigma_i X_{i,j} \right|^q \right)^{1/q} \lesssim \sqrt{q} \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n X_{i,j}^2 \right)^{1/2} \\ &\leq \sqrt{q} (\mathbb{E} X_{1,1}^2)^{1/2} \leq \sqrt{q}. \end{aligned}$$

Thus, applying Lemma B.2 with $\zeta = 1/2$ to bound (30), we have that

$$\mathbb{E} \sup_{\beta \in aB_1^p \cap B_2^p} \frac{1}{n} \sum_{i=1}^n \sigma_i \langle X_i, \beta \rangle \lesssim a \sqrt{\frac{\log(p)}{n}},$$

concluding the proof. \blacksquare

Acknowledgements. M.L. would like to thank C.S. Lorenz and M.D. Wong for helpful comments and F.K. would like to thank D. Fan for help with the Euler Cluster. The authors would like to thank the editors and an anonymous referee for their helpful suggestions and remarks, which led to an improved version of the paper.

Funding. G.C. and M.L. are funded in part by ETH Foundations of Data Science (ETH-FDS).

References

- [1] A. Ai, A. Lapanowski, Y. Plan, and R. Vershynin, One-bit compressed sensing with non-Gaussian measurements. *Linear Algebra Appl.* **441** (2014), 222–239 Zbl [1332.94041](#) MR [3134344](#)
- [2] T. W. Anderson, The integral of a symmetric unimodal function over a symmetric convex set and some probability inequalities. *Proc. Amer. Math. Soc.* **6** (1955), 170–176 Zbl [0066.37402](#) MR [69229](#)
- [3] P. Awasthi, M. Balcan, N. Haghtalab, and H. Zhang, Learning and 1-bit compressed sensing under asymmetric noise. In *COLT 2016 – 29th Annual Conference on Learning Theory*, pp. 152–192, Proceedings of Machine Learning Research 49, PMLR, 2016
- [4] M. Balcan and H. Zhang, Sample and Computationally Efficient Learning Algorithms under S-Concave Distributions. In *NIPS 2017 – Advances in Neural Information Processing Systems 30*, pp. 4799–4808, Curran Associates, Inc., 2017
- [5] R. Baraniuk, S. Foucart, D. Needell, Y. Plan, and M. Wootters, One-bit compressive sensing of dictionary-sparse signals. *Inf. Inference* **7** (2018), no. 1, 83–104 Zbl [1446.94017](#) MR [3801519](#)
- [6] P. L. Bartlett, P. M. Long, G. Lugosi, and A. Tsigler, Benign overfitting in linear regression. *Proc. Natl. Acad. Sci. USA* **117** (2020), no. 48, 30063–30070 Zbl [1485.62085](#) MR [4263288](#)
- [7] M. Belkin, D. Hsu, S. Ma, and S. Mandal, Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proc. Natl. Acad. Sci. USA* **116** (2019), no. 32, 15849–15854 Zbl [1433.68325](#) MR [3997901](#)
- [8] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge Univ. Press, Cambridge, 2004 Zbl [1058.90049](#) MR [2061575](#)
- [9] L. Breiman, Arcing classifiers. *Ann. Statist.* **26** (1998), no. 3, 801–849 Zbl [0934.62064](#) MR [1635406](#)
- [10] L. Breiman, Population theory for boosting ensembles. *Ann. Statist.* **32** (2004), no. 1, 1–11 Zbl [1105.62308](#) MR [2050998](#)
- [11] P. Bühlmann, Boosting for high-dimensional linear models. *Ann. Statist.* **34** (2006), no. 2, 559–583 Zbl [1095.62077](#) MR [2281878](#)
- [12] B. Carl, Inequalities of Bernstein-Jackson-type and the degree of compactness of operators in Banach spaces. *Ann. Inst. Fourier (Grenoble)* **35** (1985), no. 3, 79–118 Zbl [0564.47009](#) MR [810669](#)
- [13] D. Chafaï, O. Guédon, G. Lecué, and A. Pajor, *Interactions between compressed sensing random matrices and high dimensional geometry*. Panor. Synthèses 37, Soc. Math. France, Paris, 2012 Zbl [1396.94015](#) MR [3113826](#)
- [14] L. H. Y. Chen, L. Goldstein, and Q.-M. Shao, *Normal approximation by Stein’s method*. Probab. Appl. (N. Y.), Springer, Heidelberg, 2011 Zbl [1213.62027](#) MR [2732624](#)
- [15] S. S. Chen, D. L. Donoho, and M. A. Saunders, Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.* **20** (1998), no. 1, 33–61 Zbl [0919.94002](#) MR [1639094](#)
- [16] G. Chinot and M. Lerasle, Benign overfitting in the large deviation regime. To appear in *Bernoulli*

- [17] G. Chinot, M. Löffler, and S. van de Geer, On the robustness of minimum norm interpolators and regularized empirical risk minimizers. To appear in *Ann. Statist.*
- [18] Z. Deng, A. Kammoun, and C. Thrampoulidis, A model of double descent for high-dimensional binary linear classification. To appear in *Inf. Inference*
- [19] I. Diakonikolas, C. Tzamos, V. Kontonis, and N. Zarifis, Non-convex SGD learns half-spaces with adversarial label noise. In *NeurIPS 2020 – Advances in Neural Information Processing Systems 33*, Curran Associates, Inc., 2020
- [20] S. Dirksen and S. Mendelson, Non-Gaussian hyperplane tessellations and robust one-bit compressed sensing. *J. Eur. Math. Soc. (JEMS)* **23** (2021), no. 9, 2913–2947
Zbl [07367681](#) MR [4275468](#)
- [21] H. Drucker and C. Cores, Boosting decision trees. In *NIPS 1995 – Advances in Neural Information Processing Systems 8*, pp. 479–485, MIT Press, 1995
- [22] P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur, Sharpness-aware minimization for efficiently improving generalization. In *ICLR 2021 – International Conference on Learning Representations*, 2021
- [23] S. Frei, Y. Cao, and Q. Gu, Agnostic learning of halfspaces with gradient descent via soft margins. In *ICML 2021 – 38th International Conference on Machine Learning*, pp. 3417–3426, Proceedings of Machine Learning Research 139, PMLR, 2021
- [24] Y. Freund and R. E. Schapire, A decision-theoretic generalization of on-line learning and an application to boosting. *Comput. Syst. Sci.* **55** (1997), no. 1, 119–139
Zbl [0880.68103](#) MR [1473055](#)
- [25] J. Friedman, T. Hastie, and R. Tibshirani, Additive logistic regression: A statistical view of boosting. *Ann. Statist.* **28** (2000), no. 2, 337–407 Zbl [1106.62323](#) MR [1790002](#)
- [26] E. Giné and R. Nickl, *Mathematical foundations of infinite-dimensional statistical models*. Camb. Ser. Stat. Probab. Math. 40, Cambridge Univ. Press, New York, 2016
Zbl [1358.62014](#) MR [3588285](#)
- [27] Y. Gordon, A. E. Litvak, C. Schütt, and E. Werner, On the minimum of several random variables. *Proc. Amer. Math. Soc.* **134** (2006), no. 12, 3665–3675 Zbl [1107.60009](#)
MR [2240681](#)
- [28] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, Surprises in high-dimensional ridgeless least squares interpolation. *Ann. Statist.* **50** (2022), no. 2, 949–986
Zbl [07514075](#) MR [4404925](#)
- [29] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans. Inform. Theory* **59** (2013), no. 4, 2082–2102 Zbl [1364.94130](#) MR [3043783](#)
- [30] Z. Ji and M. Telgarsky, The implicit bias of gradient descent on nonseparable data. In *COLT 2019 – 32nd Annual Conference on Learning Theory*, pp. 1772–1798, Proceedings of Machine Learning Research 99, PMLR, 2019
- [31] W. Jiang, Process consistency for AdaBoost. *Ann. Statist.* **32** (2004), no. 1, 13–29
Zbl [1105.62316](#) MR [2050999](#)
- [32] K. Knudson, R. Saab, and R. Ward, One-bit compressive sensing with norm estimation. *IEEE Trans. Inform. Theory* **62** (2016), no. 5, 2748–2758 Zbl [1359.94117](#)
MR [3493876](#)

- [33] V. Koltchinskii and D. Panchenko, Empirical margin distributions and bounding the generalization error of combined classifiers. *Ann. Statist.* **30** (2002), no. 1, 1–50
Zbl [1012.62004](#) MR [1892654](#)
- [34] F. Krahmer, C. Kümmerle, and O. Melnyik, On the robustness of noise-blind low-rank recovery from rank-one measurements. 2020, arXiv:[2010.12402](#)
- [35] F. Krahmer, C. Kümmerle, and H. Rauhut, A quotient property for matrices with heavy-tailed entries and its application to noise-blind compressed sensing. 2018, arXiv:[1806.04261](#)
- [36] B. Laurent and P. Massart, Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.* **28** (2000), no. 5, 1302–1338 Zbl [1105.62328](#) MR [1805785](#)
- [37] T. Liang, Universal prediction band via semi-definite programming. 2021, arXiv:[2103.17203](#)
- [38] T. Liang and P. Sur, A precise high-dimensional asymptotic theory for boosting and minimum- ℓ_1 -norm interpolated classifiers. *Ann. Statist.* **50** (2022), no. 3, 1669–1695
Zbl [7547946](#) MR [4441136](#)
- [39] G. Lugosi and N. Vayatis, On the Bayes-risk consistency of regularized boosting methods. *Ann. Statist.* **32** (2004), no. 1, 30–55 Zbl [1105.62319](#) MR [2051000](#)
- [40] S. Mei and A. Montanari, The generalization error of random features regression: precise asymptotics and the double descent curve. *Comm. Pure Appl. Math.* **75** (2022), no. 4, 667–766 Zbl [07513415](#) MR [4400901](#)
- [41] S. Mendelson, Learning without concentration. *J. ACM* **62** (2015), no. 3, Art. ID 21
Zbl [1333.68232](#) MR [3367000](#)
- [42] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann, Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geom. Funct. Anal.* **17** (2007), no. 4, 1248–1282 Zbl [1163.46008](#) MR [2373017](#)
- [43] A. Montanari, F. Ruan, Y. Sohn, and J. Yan, The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. 2020, arXiv:[1911.01544](#)
- [44] I. Mukherjee, C. Rudin, and R. E. Schapire, The rate of convergence of AdaBoost. *J. Mach. Learn. Res.* **14** (2013), 2315–2347 Zbl [1317.68176](#) MR [3111364](#)
- [45] V. Muthukumar, A. Narang, V. Subramanian, M. Belkin, D. Hsu, and A. Sahai, Classification vs regression in overparametrized regimes: does the loss function matter? *J. Mach. Learn. Res.* **22** (2021), Art. ID 222 MR [4329801](#)
- [46] Y. Plan and R. Vershynin, One-bit compressed sensing by linear programming. *Comm. Pure Appl. Math.* **66** (2013), no. 8, 1275–1297 Zbl [1335.94018](#) MR [3069959](#)
- [47] Y. Plan and R. Vershynin, Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach. *IEEE Trans. Inform. Theory* **59** (2013), no. 1, 482–494
Zbl [1364.94153](#) MR [3008160](#)
- [48] G. Rätsch, T. Onoda, and K.R. Müller, Soft margins for AdaBoost. *Mach. Learn.* **42** (2001), no. 3, 287–320 Zbl [0969.68128](#)
- [49] E. Rio, Moment inequalities for sums of dependent random variables under projective conditions. *J. Theoret. Probab.* **22** (2009), no. 1, 146–163 Zbl [1160.60312](#)
MR [2472010](#)

- [50] S. Rosset, J. Zhu, and T. Hastie, Boosting as a regularized path to a maximum margin classifier. *J. Mach. Learn. Res.* **5** (2003/04), 941–973 Zbl [1222.68290](#) MR [2248005](#)
- [51] M. Rudelson and R. Vershynin, On sparse reconstruction from Fourier and Gaussian measurements. *Comm. Pure Appl. Math.* **61** (2008), no. 8, 1025–1045 Zbl [1149.94010](#) MR [2417886](#)
- [52] R. E. Schapire, Y. Freund, P. Bartlett, and W. S. Lee, Boosting the margin: a new explanation for the effectiveness of voting methods. *Ann. Statist.* **26** (1998), no. 5, 1651–1686 Zbl [0929.62069](#) MR [1673273](#)
- [53] R. E. Schapire and Y. Singer, Improved boosting algorithms using confidence-rated predictions. In *Mach. Learn.* **37** (1999), no. 3, 297–336 Zbl [0945.68194](#) MR [1811573](#)
- [54] D. Soudry, E. Hoffer, M. S. Nacson, S. Gunasekar, and N. Srebro, The implicit bias of gradient descent on separable data. *J. Mach. Learn. Res.* **19** (2018), Art. ID 70 Zbl [1477.62192](#) MR [3899772](#)
- [55] M. Telgarsky, Margins, shrinkage, and boosting. In *ICML 2013 – 30th International Conference on Machine Learning*, pp. 307–315, Journal of Machine Learning Research 28, Curran Associates, Inc., 2013
- [56] R. Vershynin, *High-dimensional probability*. Camb. Ser. Stat. Probab. Math. 47, Cambridge Univ. Press, Cambridge, 2018 Zbl [1430.60005](#) MR [3837109](#)
- [57] L. Wan, M. Zeiler, S. Zhang, Y. Lecun, and R. Fergus, Regularization of neural networks using DropConnect. In *ICML 2013 – 30th International Conference on Machine Learning*, pp. 1058–1066, Journal of Machine Learning Research 28, JMLR, 2013
- [58] P. Wojtaszczyk, Stability and instance optimality for Gaussian measurements in compressed sensing. *Found. Comput. Math.* **10** (2010), no. 1, 1–13 Zbl [1189.68060](#) MR [2591836](#)
- [59] A. J. Wyner, M. Olson, J. Bleich, and D. Mease, Explaining the success of AdaBoost and random forests as interpolating classifiers. *J. Mach. Learn. Res.* **18** (2017), Art. ID 48 Zbl [1433.68384](#) MR [3670493](#)
- [60] C. Zhang, Efficient active learning of sparse halfspaces. In *COLT 2018 – 31st Conference on Learning Theory*, pp. 1–26, Proceedings of Machine Learning Research 75, PMLR, 2018
- [61] C. Zhang, S. Bengio, M. Hardt, B Recht, and O. Vinyals, Understanding deep learning requires rethinking generalization. *ICLR 2017 – 5th International Conference on Learning Representations*, 2017
- [62] C. Zhang, J. Shen, and P. Awasthi, Efficient active learning of sparse halfspaces with arbitrary bounded noise. In *NeurIPS 2020 – Advances in Neural Information Processing Systems 33*, pp. 7184–7197, Curran Associates, Inc., 2020
- [63] L. Zhang, J. Yi, and R. Jin, Efficient algorithms for robust one-bit compressive sensing. In *ICML 2014 – 31st International Conference on Machine Learning*, pp. 820–828, Journal of Machine Learning Research 32, JMLR, 2014
- [64] T. Zhang and B. Yu, Boosting with early stopping: convergence and consistency. *Ann. Statist.* **33** (2005), no. 4, 1538–1579 Zbl [1078.62038](#) MR [2166555](#)

Received 22 July 2021; revised 10 December 2021.

Geoffrey Chinot

Seminar for Statistics, Department of Mathematics, ETH Zürich, Rämistrasse 101, 8092 Zurich, Switzerland; geoffrey.chinot@stat.math.ethz.ch

Felix Kuchelmeister

Seminar for Statistics, Department of Mathematics, ETH Zürich, Rämistrasse 101, 8092 Zurich, Switzerland; felix.kuchelmeister@stat.math.ethz.ch

Matthias Löffler

Seminar for Statistics, Department of Mathematics, ETH Zürich, Rämistrasse 101, 8092 Zurich, Switzerland; matthias.loeffler@stat.math.ethz.ch

Sara van de Geer

Seminar for Statistics, Department of Mathematics, ETH Zürich, Rämistrasse 101, 8092 Zurich, Switzerland; sara.vandegeer@stat.math.ethz.ch