

A MODEL OF DOUBLE DESCENT FOR HIGH-DIMENSIONAL BINARY LINEAR CLASSIFICATION

Zeyu Deng Abba Kammoun Christos Thrampoulidis

ABSTRACT

We consider a model for logistic regression where only a subset of features of size p is used for training a linear classifier over n training samples. The classifier is obtained by running gradient-descent (GD) on the logistic-loss. For this model, we investigate the dependence of the generalization error on the overparameterization ratio $\kappa = p/n$. First, building on known deterministic results on convergence properties of the GD, we uncover a phase-transition phenomenon for the case of Gaussian regressors: the generalization error of GD is the same as that of the maximum-likelihood (ML) solution when $\kappa < \kappa_*$, and that of the max-margin (SVM) solution when $\kappa > \kappa_*$. Next, using the convex Gaussian min-max theorem (CGMT), we sharply characterize the performance of both the ML and SVM solutions. Combining these results, we obtain curves that explicitly characterize the generalization error of GD for varying values of κ . The numerical results validate the theoretical predictions and unveil “double-descent” phenomena that complement similar recent observations in linear regression settings.

1. INTRODUCTION

Motivation. Modern learning architectures are chosen such that the number of training parameters overly exceeds the size of the training set. Despite their increased complexity, such *over-parametrized* architectures are known to generalize well in practice. In principle, this contradicts conventional statical wisdom of the so-called approximation-generalization tradeoff. The latter suggests a U-shaped curve for the generalization error as a function of the number of parameters, over which the error initially decreases, but increases after some point due to overfitting. In contrast, several authors uncover a peculiar W-shaped curve for the generalization error of neural networks as a function of the number of parameters. After the “classical” U-shaped curve,

it is seen that a second approximation-generalization tradeoff (hence, a second U-shaped curve) appears for large enough number of parameters. The authors of [BMM18, BHMM18, BHM18] coin this the “double-descent” risk curve.

Recent efforts towards theoretically understanding the phenomenon of double-descent focus on linear regression with Gaussian features [HMRT19, MVS19, BHX19, BLLT19]; also [BRT18, XH19, MM19] for related efforts. These works investigate how the generalization error of gradient descent (GD) on square-loss depends on the overparameterization ratio $\kappa = p/n$, where p number of features are divided by the size n of the training set. On one hand, when $\kappa < 1$, GD iterations converge to the least-squares solution for which the generalization performance is well-known [HMRT19]. On the other hand, in the overparameterized regime $\kappa > 1$, GD iterations converge to the min-norm solution for which the generalization performance is sharply characterized in [HMRT19, BHX19] using random matrix-theory (RMT). Using these sharp asymptotics, these papers identify regimes (in terms of model parameters such as SNR) for which a double-descent curve appears.

Contributions. This paper investigates the dependence of generalization error on the overparameterization ratio in binary linear classification with Gaussian regressors. In short, we obtain results that parallel previous studies for linear regression [HMRT19, BHX19]. In more detail, we propose studying gradient descent on logistic loss for two simple, yet popular, models: logistic model and gaussian mixtures (GM) model. Known results establish that GD iterations converge to either the support-vector machines (SVM) solution or the logistic maximum-likelihood (ML) solution, depending on whether the training data is separable or not. For the proposed learning model, we compute a phase-transition threshold κ_* and show that when problem dimensions are large, then data are separable if and only if $\kappa > \kappa_*$. Connecting the two, we redefine our task to that of studying the performance of ML and SVM. In contrast to linear regression, where the corresponding task can be accomplished using RMT, here we employ machinery based on Gaussian process inequalities. In particular, we obtain sharp asymptotics using the framework of the convex Gaussian

Zeyu Deng and Christos Thrampoulidis are with the Electrical and Computer Engineering Department at the University of California, Santa Barbara, USA. Abba Kammoun is with the Electrical Engineering Department at King Abdullah University of Science and Technology, Saudi Arabia.

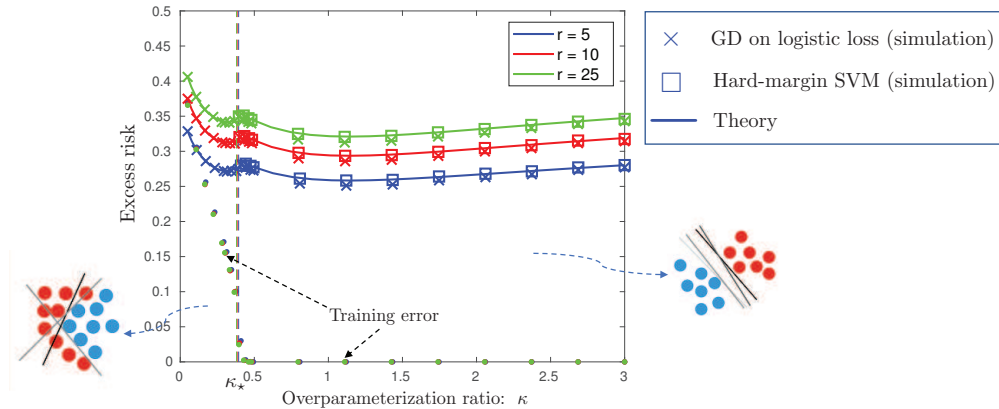


Fig. 1: Excess risk as a function of the overparameterization ratio κ for binary logistic regression under a polynomial feature selection rule. Our asymptotic predictions (theory) match the simulation results for gradient descent (GD) on logistic loss and for hard-margin support-vector machines (SVM). The theory also predicts a sharp phase-transition phenomenon: data is linearly separable (eqv. training error is zero) with high-probability iff $\kappa > \kappa_*$. The threshold κ_* is depicted by dashed lines (different for different values of the SNR parameter r , but almost indistinguishable due to scaling of the figure). Observe that the risk curve experiences a “double-descent”. As such, it has two local minima: one in the underparameterized and one in the overparameterized regime, respectively. Furthermore, the global minimizer of the excess risk is achieved in the latter regime in which data are linearly separable. The minimizer corresponds to the optimal number of features that should be chosen during training for the model under consideration. Please refer to Section 4 for further details on the chosen simulation parameters.

min-max theorem (CGMT) [Sto13, TOH15, TAH18]. Finally, we corroborate our theoretical findings with numerical simulations and build on the former to identify regimes where double-descent occurs. Figure 1 contains a pictorial preview of our results ¹.

Other related works. Our results on the performance of logistic ML and SVM fit in the rapidly growing recent literature on sharp asymptotics of (possibly non-smooth) convex optimization based estimators; [DMM11, BM12, Sto13, TOH15, DM16, TAH18, EK18, WWM19] and many references therein. Most of these works study linear regression models, for which the CGMT framework has been shown to be powerful and applicable under several variations; see for example [TAH18, CM19] and references therein. In contrast, our results hold for binary classification. For the derivations we utilize the machinery recently put forth in [KA19, TPT19, SAH19], which demonstrate that the CGMT framework can be also applied to classification problems. Closely related ideas were previously introduced in [TAH15, DTL18]. Here, we introduce necessary adjustments to accommodate the needs of the specific data generation model and focus on classification error which has not been studied previously. There are several other works on sharp asymptotics

¹The y-axis represents excess risk. Specifically, we show curves for the *excess* risk $\mathcal{R}_{\text{excess}} := \mathcal{R} - \mathcal{R}_{\text{best}}$ defined as the difference of the *absolute* expected risk \mathcal{R} minus the risk of the best linear classifier $\mathcal{R}_{\text{best}}$ (see Eqn. (5)). Naturally both \mathcal{R} and $\mathcal{R}_{\text{best}}$ are decreasing functions of the SNR parameter r . However, $\mathcal{R}_{\text{best}}$ is decreasing faster than \mathcal{R} . This explains why the value of the excess risk $\mathcal{R}_{\text{excess}}$ is smaller for larger values of the signal strength r in Figure 1. In particular, it holds $\mathcal{R}_{\text{best}} = 0.133, 0.098$ and 0.084 for $r = 5, 10$ and 25 , respectively. Contrast this to the values of the absolute risk $\mathcal{R} = 0.434, 0.429$ and 0.428 at (say) $\kappa = 0.05$.

of binary linear classification both for the logistic [CS18, SC19, SAH19] and the GM model [MLC19b, MLC19a, Hua17]. While closely related, these works differ in terms of motivation, model specification, end-results, and proof techniques.

2. LEARNING MODEL

We study supervised binary classification under two popular data models (e.g., [WR06, Sec. 3.1])

- a *discriminative* model: mixtures of Gaussian.
- a *generative* model: logistic regression.

Let $\mathbf{x} \in \mathbb{R}^d$ denote the feature vector and $y \in \pm 1$ denote the class label.

Logistic model. First, we consider a discriminative approach which models the marginal probability $p(y|\mathbf{x})$ as follows:

$$y = \begin{cases} +1, & \text{w.p. } \ell'(\mathbf{x}^T \boldsymbol{\eta}_0), \\ -1, & \text{w.p. } 1 - \ell'(\mathbf{x}^T \boldsymbol{\eta}_0), \end{cases}$$

where $\boldsymbol{\eta}_0 \in \mathbb{R}^d$ is the unknown weight (or, regressor) vector and

$$\ell'(t) := (1 + e^{-t})^{-1}.$$

Throughout, we assume IID Gaussian feature vectors $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)$. For compactness let $\text{Rad}(p)$ denote a symmetric Bernoulli distribution with probability p for the value $+1$ and probability $1 - p$ for the value -1 . We summarize the logistic model with Gaussian features:

$$y \sim \text{Rad}(\ell'(\mathbf{x}^T \boldsymbol{\eta}_0)), \quad \mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d). \quad (1)$$

Gaussian mixtures (GM) model. A common choice for the generative case is to model the class-conditional densities $p(\mathbf{x}|y)$ with Gaussians (e.g., [WR06, Sec. 3.1]). Specifically, each data point belongs to one of two classes $\mathcal{C}_{\pm 1}$ with probabilities $\pi_{\pm 1}$ such that $\pi_{+1} + \pi_{-1} = 1$. If it comes from class $\mathcal{C}_{\pm 1}$, then the feature vector $\mathbf{x} \in \mathbb{R}^d$ is an iid Gaussian vector with mean $\pm \boldsymbol{\eta}_0 \in \mathbb{R}^d$ and the response variable y takes the value of the class label ± 1 :

$$y = \pm 1 \Leftrightarrow \mathbf{x} \sim \mathcal{N}(\pm \boldsymbol{\eta}_0, \mathbf{I}_d). \quad (2)$$

2.1. Training data: Feature selection

During training, we assume access to n data points generated according to either (1) or (2). We allow for the possibility that only a subset $\mathcal{S} \subset [d]$ of the total number of features is known during training. Concretely, for each feature vector \mathbf{x}_i , $i \in [n]$, the learner only knows a sub-vector $\mathbf{w}_i := \mathbf{x}_i(\mathcal{S}) := \{\mathbf{x}_i(j), j \in \mathcal{S}\}$, for a chosen set $\mathcal{S} \subset [d]$. We denote the size of the *known* feature sub-vector as $p := |\mathcal{S}|$. Onwards, we choose $\mathcal{S} = \{1, 2, \dots, p\}$ ², i.e., select the features sequentially in the order in which they appear.

Clearly, $1 \leq p \leq d$. Overall, the training set $\mathcal{T}(\mathcal{S})$ consists of n data pairs:

$$\begin{aligned} \mathcal{T}(\mathcal{S}) &:= \{(\mathbf{w}_i, y_i), i = 1, \dots, n\}, \\ \mathbf{w}_i &\in \mathbb{R}^p \text{ where } (y_i, \mathbf{x}_i) \sim (1) \text{ or } (2). \end{aligned} \quad (3)$$

When clear from context, we simply write

$$\mathcal{T} \sim (1) \text{ or } \mathcal{T} \sim (2),$$

if the training set \mathcal{T} is generated according to (1) or (2), respectively.

2.2. Classification rule

Having access to the training set $\mathcal{T}(\mathcal{S})$, the learner obtains an estimate $\boldsymbol{\beta}_\star \in \mathbb{R}^p$ of the weight vector $\boldsymbol{\eta}_0 \in \mathbb{R}^d$. Then, for a newly generated sample (\mathbf{x}, y) (and $\mathbf{w} := \mathbf{x}(\mathcal{S})$), she forms a *linear classifier* and decides a label \hat{y} for the new sample as:

$$\hat{y} = \text{sign}(\mathbf{w}^T \boldsymbol{\beta}_\star).$$

The estimate $\boldsymbol{\beta}_\star$ is minimizing the empirical risk

$$\widehat{\mathcal{R}}_{\text{emp}}(\boldsymbol{\beta}) := \frac{1}{n} \sum_{i=1}^n \ell(y_i \mathbf{w}_i^T \boldsymbol{\beta}),$$

for certain loss function $\ell: \mathbb{R} \rightarrow \mathbb{R}$. A traditional way to minimize $\widehat{\mathcal{R}}_{\text{emp}}$ is by constructing gradient descent (GD) iterates $\boldsymbol{\beta}_k$, $k \geq 0$ via

$$\boldsymbol{\beta}_{(k+1)} = \boldsymbol{\beta}_{(k)} - \eta_k \nabla \widehat{\mathcal{R}}_{\text{emp}}(\boldsymbol{\beta}_{(k)}),$$

²This assumption is without loss of generality for our asymptotic model specified in Section 3.1.

for step-size $\eta_k > 0$ and arbitrary $\boldsymbol{\beta}_{(0)}$. We run GD until convergence and set

$$\boldsymbol{\beta}_\star = \boldsymbol{\beta}^\infty := \lim_{k \rightarrow \infty} \boldsymbol{\beta}^k.$$

In this paper, we focus on empirical logistic risk

$$\ell(t) := \log(1 + e^{-t}).$$

Generalization error. For a new sample (\mathbf{x}, y) we measure generalization error of \hat{y} by the expected risk

$$\mathcal{R}(\boldsymbol{\beta}_\star) := \mathbb{E}[\mathbb{1}(\hat{y} \neq y)]. \quad (4)$$

The expectation here is over the data generation model (1) or (2). In particular, here we consider the *excess risk*

$$\mathcal{R}_{\text{excess}}(\boldsymbol{\beta}_\star) := \mathcal{R}(\boldsymbol{\beta}_\star) - \mathcal{R}(\boldsymbol{\eta}_0), \quad (5)$$

where $\mathcal{R}(\boldsymbol{\eta}_0)$ is the risk of the best³ linear classifier, i.e., $\mathcal{R}(\boldsymbol{\eta}_0) := \mathbb{E}[\mathbb{1}(y \neq \text{sign}(\mathbf{x}^T \boldsymbol{\eta}_0))]$. Also of interest is the cosine similarity between the estimate $\boldsymbol{\beta}_\star$ and $\boldsymbol{\eta}_0$:

$$\mathcal{C}(\boldsymbol{\beta}_\star) := \cos(\boldsymbol{\beta}_\star, \boldsymbol{\eta}_0) = \frac{\langle \boldsymbol{\beta}_\star, \boldsymbol{\eta}_0 \rangle}{\|\boldsymbol{\beta}_\star\|_2 \|\boldsymbol{\eta}_0\|_2}. \quad (6)$$

Training error. The training error of $\boldsymbol{\beta}_\star$ is given by

$$\mathcal{R}_{\text{train}}(\boldsymbol{\beta}_\star) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}(\hat{y} \neq y_i).$$

2.3. Convergence behavior of GD iterates

Recent literature studies the behavior of GD iterates for logistic loss. There are two regimes of interest: (i) when data are such that $\widehat{\mathcal{R}}_{\text{emp}}(\boldsymbol{\beta})$ is strongly convex then standard tools show that $\boldsymbol{\beta}$ converges to the unique bounded optimum of $\widehat{\mathcal{R}}_{\text{emp}}(\boldsymbol{\beta})$; (ii) when data are separable then the normalized iterates $\boldsymbol{\beta}_k / \|\boldsymbol{\beta}_k\|_2$ converge to the max-margin solution [JT19, SHN⁺18, Tel13]. These properties guide our approach towards studying the generalization error of GD for logistic loss. Instead of the latter, we study the performance of the max-margin classifier and of the minimum of the empirical logistic risk. We formalize these ideas next.

2.3.1. Separable data

The training set is separable iff there exists a linear classifier achieving zero training error, i.e., $\exists \boldsymbol{\beta} : y_i \mathbf{w}_i^T \boldsymbol{\beta} \geq 1, \forall i \in [n]$. When data are separable, for $k \rightarrow \infty$ it holds

$$\left\| \frac{\boldsymbol{\beta}^{(k)}}{\|\boldsymbol{\beta}^{(k)}\|_2} - \frac{\widehat{\boldsymbol{\beta}}}{\|\widehat{\boldsymbol{\beta}}\|_2} \right\|_2 \rightarrow 0, \text{ where}$$

$$\widehat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \|\boldsymbol{\beta}\|_2 \quad \text{subject to } y_i \mathbf{w}_i^T \boldsymbol{\beta} \geq 1, \quad (7)$$

i.e., the solution to **hard-margin SVM**.

³It assumes knowledge of the entire feature vector \mathbf{x}_i and of $\boldsymbol{\eta}_0$.

2.3.2. Non-separable data

When the separability condition does *not* hold, then $\widehat{\mathcal{R}}_{\text{emp}}(\beta)$ is coercive. Thus, its sub-level sets are closed and GD iterates will converge [JT19] to the minimizer $\widehat{\beta}$ of the empirical loss:

$$\widehat{\beta} = \arg \min_{\beta} \frac{1}{n} \sum_{i=1}^n \ell(y_i \mathbf{w}_i^T \beta), \quad (8)$$

where, $\ell(t) = \log(1 + \exp(-t))$. For the data model in (1), $\widehat{\beta}$ is the **maximum-likelihood (ML)** estimator.

3. SHARP ASYMPTOTICS

3.1. Asymptotic setting

Recall the following notation:

- d : dimension of the ambient space,
- n : training sample size,
- p : number of parameters during training (see (3)).

We study a setting in which d and n are fixed and p varies from 1 to d . Our asymptotic results hold in a *linear asymptotic regime* where $n, d, p \rightarrow +\infty$ such that

$$d/n \rightarrow \zeta > 0 \quad \text{and} \quad p/n \rightarrow \kappa \in (0, \zeta]. \quad (9)$$

We fix $\zeta \in [1, +\infty]$ and derive asymptotic predictions for the generalization error as a function of $\kappa \in (0, \zeta]$, which determines the overparametrization ratio in our model. To quantify the effect of κ on the generalization error, we decompose the feature vector $\mathbf{x}_i \in \mathbb{R}^d$ to its *known* part $\mathbf{w}_i \in \mathbb{R}^p$ and to its *unknown* part \mathbf{z}_i :

$$\mathbf{x}_i := [\mathbf{w}_i^T, \mathbf{z}_i^T]^T.$$

Then, we let $\beta_0 \in \mathbb{R}^p$ (resp., $\gamma_0 \in \mathbb{R}^{d-p}$) denote the vector of weight coefficients corresponding to the known (resp., unknown) features such that

$$\boldsymbol{\eta}_0 := [\beta_0^T, \gamma_0^T]^T.$$

In this notation, we study a sequence of problems of increasing dimensions as in (9) that further satisfy:

$$\begin{aligned} \|\boldsymbol{\eta}_0\|_2 &\rightarrow r, \\ \|\beta_0\|_2 &\rightarrow s := s(\kappa) \quad \text{and} \quad \|\gamma_0\|_2 \rightarrow \sigma := \sigma(\kappa). \end{aligned} \quad (10)$$

The parameters s and σ can be thought of as the useful signal strength and the noise strength, respectively. Our notation specifies that $s(\kappa)$ (hence also, $\sigma(\kappa) = \sqrt{r^2 - s^2(\kappa)}$) is a function of κ . We are interested in functions $s(\kappa)$ that are increasing in κ such that the signal strength increases as more features enter into the training model; Sec. 4.1 for explicit parameterizations.

3.2. Notation

We reserve the following notation for random variables G, H and Y

$$\begin{aligned} H, G, Z &\stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1), \\ \text{and } Y &\sim \text{Rad}(\ell'(sG + \sigma Z)), \end{aligned} \quad (11)$$

for s and σ as defined in (10). All expectations and probabilities are with respect to the randomness of H, G, Z . Also, let $(x)_- = \min\{x, 0\}$. For a sequence of random variables $\mathcal{X}_{n,p,d}$ that converges in probability to a constant c in the limit of (9), we simply write $\mathcal{X}_{n,p,d} \xrightarrow{P} c$. Finally, we denote the proximal operator of the logistic-loss as follows,

$$\text{prox}_{\ell}(x; \lambda) := \arg \min_v \frac{1}{2\lambda}(x - v)^2 + \ell(v).$$

3.3. Regimes of learning: Phase-transition

As discussed in Section 2, the behavior of GD changes depending on whether the training data is separable or not. Under the Gaussian feature model, the following proposition establishes a sharp phase-transition characterizing the separability of the training data. The proposition is an extension of the phase-transition result by Candès and Sur [CS18] for the (noiseless) logistic model. We extend their result to the noisy setting (to accommodate for the feature selection model in Section 2.1) as well as to the Gaussian mixtures model.⁴

Proposition 3.1 (Phase transition). *For a training set \mathcal{T} that is generated by either of the two models in (1) or (2), consider the event*

$$\mathcal{E}_{\text{sep}} := \{ \exists \beta \in \mathbb{R}^p : y_i \mathbf{w}_i^T \beta \geq 1, \forall i \in [n] \},$$

under which the training data is separable. Recall from (3) that for all $i \in [n]$, $\mathbf{w}_i = \mathbf{x}_i(\{1, \dots, p\})$ are the p (out of d) features that are used for training. Let $s = s(\kappa)$ in (10) be an increasing function of κ . Recall the notation in (11) and define a random variable V depending on the data model as follows

$$V := \begin{cases} GY & , \text{ if } \mathcal{T} \sim (1), \\ G + s & , \text{ if } \mathcal{T} \sim (2). \end{cases} \quad (12)$$

Further refine the following threshold function depending on the data generation model:

$$g(\kappa) := \min_{t \in \mathbb{R}} \mathbb{E}(H + tV)_-^2. \quad (13)$$

⁴Our analysis approach is also slightly different to [CS18]. Specifically, we note that data is linearly separable iff the solution to hard-margin SVM (7) is bounded. Then, using the CGMT we are able to show that the latter happens with probability one iff $\kappa > \kappa_*$.

Let $\kappa_\star \in [0, 1/2]$ be the unique solution to the equation $g(\kappa) = \kappa$. Then, the following holds regarding \mathcal{E}_{sep} :

$$\begin{aligned} \kappa > \kappa_\star &\Rightarrow \lim_{n,p,d \rightarrow \infty} \Pr(\mathcal{E}_{\text{sep}}) = 1, \\ \kappa < \kappa_\star &\Rightarrow \lim_{n,p,d \rightarrow \infty} \Pr(\mathcal{E}_{\text{sep}}) = 0. \end{aligned}$$

Put in words: the training data is separable iff $\kappa > \kappa_\star$. When this is the case, then the training error $\mathcal{R}_{\text{train}}$ can be driven to zero and we are in the *interpolating regime*. In contrast, the training error is non-vanishing for smaller values of κ .

Remark 1 (The threshold for Gaussian mixture). Consider the Gaussian mixture model, i.e. $\mathcal{T} \sim (2)$. Substituting the value of $V = G + s$ in (13) and using the fact that G and H are independent Gaussians the threshold function for the GM model simplifies to:

$$\begin{aligned} g(\kappa) &= \min_{t \in \mathbb{R}} \mathbb{E} \left(G \sqrt{1+t^2} + t s \right)_- \\ &= \min_{t \in \mathbb{R}} \left(1 + t^2 + t^2 s^2 \right) \cdot Q \left(\frac{ts}{\sqrt{1+t^2}} \right) \\ &\quad - ts \sqrt{1+t^2} \cdot \psi \left(- \frac{ts}{\sqrt{1+t^2}} \right). \end{aligned}$$

where $\psi(t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$ and $Q(x) = \int_x^\infty \psi(t) dt$ are the density and tail function of the standard normal distribution, respectively.

3.4. High-dimensional asymptotics

Propositions 3.2 and 3.3 characterize the performance of the two optimization-based classifiers in (8) and (7), respectively, under the logistic (cf. (1)) and the GM (cf. (2)) model. When combined with the results of Section 2.3, they characterize the statistical performance of converging points of GD.

3.4.1. Non-separable data

The performance of logistic loss under the logistic model was recently studied in [SC19, SAH19, TPT19]. Here, we follow [TPT19, Thm. III.1]. Specifically, we appropriately modify their proof and result to fit the data model of Section 2 and to obtain a prediction for the classification error (not considered in prior work). Also, we provide extensions to the GM model.

Proposition 3.2 (ML). *Consider a training set \mathcal{T} that is generated by either of the two models in (1) or (2) and fix $\kappa < \kappa_\star$. Recall the notation in (11) and define a random variable V depending on the data model as in (12). Let $(\mu, \alpha > 0, \lambda)$ be the unique solution to the following*

system of three nonlinear equations in three unknowns,

$$\begin{aligned} \mu &= \mathbb{E} [V \cdot \text{prox}_\ell(\alpha H + \mu V; \lambda)], \\ \alpha^2 \kappa &= \lambda^2 \mathbb{E} [(\ell'(\text{prox}_\ell(\alpha H + \mu V; \lambda)))^2], \\ \kappa &= \lambda \mathbb{E} \left[\frac{\ell''(\text{prox}_\ell(\alpha H + \mu V; \lambda))}{1 + \lambda \ell''(\text{prox}_\ell(\alpha H + \mu V; \lambda))} \right]. \end{aligned} \quad (14)$$

Then,

$$\begin{aligned} \mathcal{C}(\hat{\beta}) &\xrightarrow{P} \frac{s \mu}{r \sqrt{\mu^2 + \alpha^2}}, \\ \text{and } \mathcal{R}(\hat{\beta}) &\xrightarrow{P} \mathbb{P}(\mu V + \alpha H < 0). \end{aligned}$$

Remark 2 (Simplifications for the GM model). When data is generated according to (2), then $V = G + s$. Using this and Gaussian integration by parts, the system of three equations in (14) simplifies to the following:

$$\begin{aligned} \mu \kappa &= -s \mathbb{E} \left[\lambda \ell'(\text{prox}_\ell(G_{\mu,\alpha}; \lambda)) \right], \\ \alpha^2 \kappa &= \mathbb{E} [(\lambda \ell'(\text{prox}_\ell(G_{\mu,\alpha}; \lambda)))^2], \\ \kappa &= \lambda \mathbb{E} \left[\frac{\ell''(\text{prox}_\ell(G_{\mu,\alpha}; \lambda))}{1 + \lambda \ell''(\text{prox}_\ell(G_{\mu,\alpha}; \lambda))} \right], \end{aligned} \quad (15)$$

where the expectations are over a *single* Gaussian random variable $G_{\mu,\alpha} \sim \mathcal{N}(\mu s, \alpha^2 + \mu^2)$. We numerically solve the system of equations via a fixed-point iteration method suggested in [TAH18] (see also [SAH19, TPT19]). We empirically observe that reformulating the equations as in (15) is critical for the iteration method to converge. Perhaps an interesting observation is that the performance of logistic regression (the same is true for SVM in Section 3.4.2) does *not* depend on the prior class probability π_{+1} for the GM model with isotropic Gaussian features in (2).

3.4.2. Separable data

Here, we characterize the asymptotic performance of hard-margin SVM under both models (1) and (2).

Proposition 3.3 (SVM). *Consider a training set \mathcal{T} that is generated by either of the two models in (1) or (2) and fix $\kappa > \kappa_\star$. Recall the notation in (11) and define a random variable V depending on the data model as in (12). Further define*

$$\eta(q, \rho) := \mathbb{E} \left(\rho V + H \sqrt{1 - \rho^2} - \frac{1}{q} \right)_-^2 - (1 - \rho^2) \kappa. \quad (16)$$

Let q^\star be the unique solution of the equation

$$\min_{-1 \leq \rho \leq 1} \eta(q, \rho) = 0.$$

Further let ρ^* be the unique minimum of $\eta(q^*, \rho)$ for $\rho \in [-1, 1]$. Then,

$$\mathcal{C}(\hat{\beta}) \xrightarrow{P} \frac{\rho^* s}{r},$$

and $\mathcal{R}(\hat{\beta}) \xrightarrow{P} \mathbb{P}\left(\rho^* V + \sqrt{1 - \rho^{*2}} H < 0\right).$

Remark 3 (Margin). In addition to the stated results in Proposition 3.3, our proof further shows that the optimal cost of the hard-margin SVM (7) converges to q^* . In other words, the margin of the classifier converges to $1/q^*$.

Remark 4 (Simplifications for the GM model). Similar to Remark 2 we note that under the GM model, the function η in Proposition 3.3 simplifies to:

$$\eta(q, \rho) := \mathbb{E}(G + \rho s - 1/q)_-^2 - (1 - \rho^2)\kappa, \quad (17)$$

where, the expectation is over a single Gaussian random variable G . Moreover, the formula predicting the risk simplifies to

$$\mathcal{R}(\hat{\beta}) \xrightarrow{P} Q(\rho^* s).$$

4. NUMERICAL RESULTS

4.1. Feature selection models

Recall that the number of features known at training is determined by κ . Specifically, κ enters the formulae predicting the generalization performance via the signal strength $s = s(\kappa)$. In this section, we specify two explicit models for feature selection and their corresponding functions $s(\kappa)$. Similar models are considered in [BF83, HMRT19, BHX19], albeit in a linear regression setting.

Linear model. We start with a uniform feature selection model characterized by the following parametrization:

$$s^2 = s^2(\kappa) = r^2 \cdot (\kappa/\zeta), \quad \kappa \in (0, \zeta]. \quad (18)$$

for fixed r^2 and $\zeta > 0$. This models a setting where all coefficients of the regressor η_0 have equal contribution. Hence, the signal strength $s^2 = \|\beta_0\|_2^2$ increases linearly with the number p of features considered in training.

Polynomial model. In the linear model, adding more features during training results in a linear increase of the signal strength. In contrast, our second model assumes diminishing returns:

$$s^2 = s^2(\kappa) = r^2 \cdot (1 - (1 + \kappa)^{-\alpha}), \quad \kappa > 0, \quad (19)$$

for some $\alpha \geq 1$. As κ increases, so does the signal strength s , but the increase is less significant for larger values of κ at a rate specified by α .

4.2. Risk curves

Figure 1 assumes the polynomial feature model (19) for $\alpha = 2$ and three values of r . The crosses (' \times ') are simulation results obtained by running GD on synthetic data generated according to (1) and (19). Specifically, the depicted values are averages calculated over 500 Monte Carlo realizations for $p = 150$. For each realization, we ran GD on logistic loss (see Sec. 2.2) with a fixed step size until convergence and recorded the resulting risk. Similarly, the squares (' \square ') are simulation results obtained by solving SVM (7) over the same number of different realizations and averaging over the recorded performance. As expected by [JT19] (also Sec. 2.3), the performance of GD matches that of SVM when data are separable. Also, as predicted by Proposition 3.1, the data is separable with high-probability when $\kappa > \kappa_*$ (the threshold value κ_* is depicted with dashed vertical lines). This is verified by noticing the dotted (' \bullet ') scatter points, which record (averages of) the training error. Recall from Section 2.2 that the training error is zero if and only if the data are separable. Finally, the solid lines depict the theoretical predictions of Proposition 3.2 ($\kappa < \kappa_*$) and Proposition 3.3 ($\kappa > \kappa_*$). The results of Figure 1 validate the accuracy of our predictions: our asymptotic formulae predict the classification performance of GD on logistic loss for all values of κ . Note that the curves correspond to excess risk defined in (5); see also related Footnote 1. Corresponding results for the cosine similarity are presented in Figure 6 in the Appendix.

Under the logistic model (1), Figures 3 and 2 depict the cosine similarity and excess risk curves of GD as a function of κ for the linear and the polynomial model, respectively. Compared to Figure 1, we only show the theoretical predictions. Note that the linear model is determined by parameters (ζ, r) and the polynomial model by (α, r) . Once these values are fixed, we compute the threshold value κ_* . Then, we apply Proposition 3.2 ($\kappa < \kappa_*$) and Proposition 3.3 ($\kappa > \kappa_*$).

Finally, Figure 4 shows risk and cosine-similarity curves for the Gaussian mixture data model (2) with linear feature selection rule. The figures compare simulation results to theoretical predictions similar in nature to Figure 1, thus validating the accuracy of the latter for the GM model. For the simulations, we generate data according to (2) with $\pi_{+1} = 1/2$, $n = 200$ and $d = 600$. The results are averages over 500 Monte Carlo realizations.

4.3. Discussion on double-descent

Here, we discuss the double-descent phenomenon that appears in Figures 1 and 2. First, focus on the under-parameterized regime $\kappa < \kappa_*$ (cf. area on the left of the vertical dashed lines). Here, the number n of training ex-




Fig. 2: Plots of the cosine similarity (left) and of the excess risk (right) as a function of κ for the polynomial feature selection model (cf. (19)) under logistic data. The curves shown are for $r = 10$ and three values of $\alpha = 1, 2$ and 5 .




Fig. 3: Plots of the cosine similarity (left) and of the excess risk (right) as a function of κ for the linear feature selection model (cf. (18)) under logistic data. The curves shown are for $\zeta = 3$ and three values of $r = 2, 5$ and 10 .

amples is large compared to the size p of the unknown weight vector β_0 , which helps learning β_0 . However, since only a (small) fraction p/d of the total number d of features are used for training, the observations are rather noisy. This tradeoff manifests itself with a “U-shaped curve” for $\kappa < \kappa_*$. Such a tradeoff is also present in the overparameterized regime $\kappa > \kappa_*$ (cf. area on the right of the vertical dashed lines) leading to a second “U-shaped curve”. The larger the value of κ the more accurate the model, but also larger the number of unknown parameters to be estimated. As seen in Figure 2, the “U-shape” is more pronounced for larger values of α . Overall the risk curve has two local minima corresponding to the two regimes of learning. Interestingly, the global minimum κ_{opt} appears in the overparameterized regime for all values of r and α depicted here. The value of $\kappa_{\text{opt}} > \kappa_*$ determines the optimal number of features that need to be selected during training to minimize the classification error. Note that for κ_{opt} the training error

is zero, yet the generalization performance is best.

The risk curves for the linear feature selection model in Figure 3 are somewhat different compared to the polynomial model. In the underparameterized regime, we observe a (very shallow) “U-shaped” curve, but the risk is monotonically decreasing for $\kappa > \kappa_*$. For the values of signal strength r considered here the global minimum is always attained at $\kappa = \zeta = 3$, i.e., when all features are used for training.

5. FUTURE WORK

This paper studies the generalization performance of GD on logistic loss under the logistic and GM models with isotropic Gaussian features. In future work, we plan on extending our study to GD on square loss, which is also commonly used in practice. Extensions of the results to multiple classes and non-isotropic features are also of interest.




Fig. 4: Plots of the (absolute) risk (bottom) and of the cosine similarity (up) as a function of κ for the linear feature selection model (cf. (18)) under data from a Gaussian mixture. The curves show $\zeta = 3$ and three signal strength values $r^2 = 1, 4$ and 10 . Please refer to the legend of Figure 1 and to Section 4.2 for further explanations.

6. REFERENCES

- [BF83] Leo Breiman and David Freedman. How many variables should be entered in a regression equation? *Journal of the American Statistical Association*, 78(381):131–136, 1983.
- [BHM18] Mikhail Belkin, Daniel J Hsu, and Partha Mitra. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In *Advances in Neural Information Processing Systems*, pages 2300–2311, 2018.
- [BHMM18] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine learning and the bias-variance trade-off. *arXiv preprint arXiv:1812.11118*, 2018.
- [BHX19] Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *arXiv preprint arXiv:1903.07571*, 2019.
- [BLLT19] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *arXiv preprint arXiv:1906.11300*, 2019.
- [BM12] Mohsen Bayati and Andrea Montanari. The lasso risk for gaussian matrices. *Information Theory, IEEE Transactions on*, 58(4):1997–2017, 2012.
- [BMM18] Mikhail Belkin, Siyuan Ma, and Soumik Mandal. To understand deep learning we need to understand kernel learning. In *International Conference on Machine Learning*, pages 540–548, 2018.
- [BRT18] Mikhail Belkin, Alexander Rakhlin, and Alexandre B Tsybakov. Does data interpolation contradict statistical optimality? *arXiv preprint arXiv:1806.09471*, 2018.
- [CM19] Michael Celentano and Andrea Montanari. Fundamental barriers to high-dimensional regression with convex penalties. *arXiv preprint arXiv:1903.10603*, 2019.
- [CS18] Emmanuel J Candès and Pragya Sur. The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression. *arXiv preprint arXiv:1804.09753*, 2018.

- [DM16] David Donoho and Andrea Montanari. High dimensional robust m-estimation: Asymptotic variance via approximate message passing. *Probability Theory and Related Fields*, 166(3-4):935–969, 2016.
- [DMM11] David L Donoho, Arian Maleki, and Andrea Montanari. The noise-sensitivity phase transition in compressed sensing. *Information Theory, IEEE Transactions on*, 57(10):6920–6941, 2011.
- [DTL18] Oussama Dhifallah, Christos Thrampoulidis, and Yue M Lu. Phase retrieval via polytope optimization: Geometry, phase transitions, and new algorithms. *arXiv preprint arXiv:1805.09555*, 2018.
- [EK18] Noureddine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170(1-2):95–175, 2018.
- [Gor88] Yehoram Gordon. *On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n* . Springer, 1988.
- [HMRT19] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.
- [Hua17] Hanwen Huang. Asymptotic behavior of support vector machine for spiked population model. *The Journal of Machine Learning Research*, 18(1):1472–1492, 2017.
- [JT19] Ziwei Ji and Matus Telgarsky. The implicit bias of gradient descent on nonseparable data. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1772–1798, Phoenix, USA, 25–28 Jun 2019. PMLR.
- [KA19] A. Kammoun and M.-S. Alouini. On the precise error analysis of support vector machines. *Submitted to IEEE Transactions on information theory*, 2019.
- [MLC19a] X. Mai, Z. Liao, and R. Couillet. A large scale analysis of logistic regression: asymptotic performance and new insights. In *ICASSP*, 2019.
- [MLC19b] Xiaoyi Mai, Zhenyu Liao, and Romain Couillet. A large scale analysis of logistic regression: Asymptotic performance and new insights. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3357–3361. IEEE, 2019.
- [MM19] Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355*, 2019.
- [MVS19] Vidya Muthukumar, Kailas Vodrahalli, and Anant Sahai. Harmless interpolation of noisy data in regression. *arXiv preprint arXiv:1903.09139*, 2019.
- [SAH19] Fariborz Salehi, Ehsan Abbasi, and Babak Hassibi. The impact of regularization on high-dimensional logistic regression. *arXiv preprint arXiv:1906.03761*, 2019.
- [SC18] Pragya Sur and Emmanuel J Candès. Additional supplementary materials for: A modern maximum-likelihood theory for high-dimensional logistic regression, 2018.
- [SC19] Pragya Sur and Emmanuel J Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525, 2019.
- [SHN⁺18] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- [Sto13] Mihailo Stojnic. A framework to characterize performance of lasso algorithms. *arXiv preprint arXiv:1303.7291*, 2013.
- [Sun96] R. K. Sundaram. *A first course in optimization theory*. Cambridge University Press, 1996.
- [TAH15] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Lasso with non-linear measurements is equivalent to one with linear measurements. In *Advances in Neural Information Processing Systems*, pages 3420–3428, 2015.

- [TAH18] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Precise error analysis of regularized m -estimators in high dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, 2018.
- [Tel13] Matus Telgarsky. Margins, shrinkage and boosting. In *Proceedings of the 30th International Conference on International Conference on Machine Learning-Volume 28*, pages II–307. JMLR. org, 2013.
- [TOH15] Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. Regularized linear regression: A precise analysis of the estimation error. In *Proceedings of The 28th Conference on Learning Theory*, pages 1683–1709, 2015.
- [TPT19] Hossein Taheri, Ramtin Pedarsani, and Christos Thrampoulidis. Sharp guarantees for solving random equations with one-bit information. *arXiv preprint arXiv:1908.04433*, 2019.
- [WR06] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [WWM19] Shuaiwen Wang, Haolei Weng, and Arian Maleki. Does slope outperform bridge regression? *arXiv preprint arXiv:1909.09345*, 2019.
- [XH19] Ji Xu and Daniel Hsu. How many variables should be entered in a principal component regression equation? *arXiv preprint arXiv:1906.01139*, 2019.

A. PROOFS

The main technical tool that facilitates the analysis of (8) and (7) is the convex Gaussian min-max theorem (CGMT) [Sto13, TOH15], which is an extension of Gordon’s Gaussian min-max inequality (GMT) [Gor88]. More precisely, we utilize the machinery recently put forth in [SAH19, TPT19, KA19], which demonstrate that the CGMT framework can be also applied to classification problems.

Organization. We introduce the necessary background on the CGMT in Section A.1. Then, based on this machinery we study the performance of SVM (cf. (7)) for the logistic model in Section A.4.2. Note that the optimization in (7) is feasible if and only if the training data is linearly separable. Thus, by studying (7) we are also able to show the phase-transition result of Proposition 3.1 in the Section A.4.1. In Section A.5 we briefly explain how to adapt the result of [TPT19, Thm. III.1] for the purpose of Proposition 3.2, again for the logistic model. It turns out that, albeit different, both data models in (1) and (2) naturally fit under the same analysis framework. Specifically the formulae for the GM model can be proved with identical arguments. The proof sketch for the SVM performance under the GM model in Section A.6 aims to showcase this.

A.1. Technical tool: CGMT

A.1.1. Gordon’s Min-Max Theorem (GMT)

The Gordon’s Gaussian comparison inequality [Gor88] compares the min-max value of two doubly indexed Gaussian processes based on how their autocorrelation functions compare. The inequality is quite general (see [Gor88]), but for our purposes we only need its application to the following two Gaussian processes:

$$X_{\mathbf{w}, \mathbf{u}} := \mathbf{u}^T \mathbf{G} \mathbf{w} + \psi(\mathbf{w}, \mathbf{u}), \quad (20a)$$

$$Y_{\mathbf{w}, \mathbf{u}} := \|\mathbf{w}\|_2 \mathbf{g}^T \mathbf{u} + \|\mathbf{u}\|_2 \mathbf{h}^T \mathbf{w} + \psi(\mathbf{w}, \mathbf{u}), \quad (20b)$$

where: $\mathbf{G} \in \mathbb{R}^{m \times n}$, $\mathbf{g} \in \mathbb{R}^m$, $\mathbf{h} \in \mathbb{R}^n$, they all have entries iid Gaussian; the sets $\mathcal{S}_{\mathbf{w}} \subset \mathbb{R}^n$ and $\mathcal{S}_{\mathbf{u}} \subset \mathbb{R}^m$ are compact; and, $\psi : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$. For these two processes, define the following (random) min-max optimization programs, which (following [TAH18]) we refer to as the *primary optimization* (PO) problem and the *auxiliary optimization* AO – for purposes that will soon become clear.

$$\Phi(\mathbf{G}) = \min_{\mathbf{w} \in \mathcal{S}_{\mathbf{w}}} \max_{\mathbf{u} \in \mathcal{S}_{\mathbf{u}}} X_{\mathbf{w}, \mathbf{u}}, \quad (21a)$$

$$\phi(\mathbf{g}, \mathbf{h}) = \min_{\mathbf{w} \in \mathcal{S}_{\mathbf{w}}} \max_{\mathbf{u} \in \mathcal{S}_{\mathbf{u}}} Y_{\mathbf{w}, \mathbf{u}}. \quad (21b)$$

According to Gordon’s comparison inequality, for any $c \in \mathbb{R}$, it holds:

$$\mathbb{P}(\Phi(\mathbf{G}) < c) \leq 2\mathbb{P}(\phi(\mathbf{g}, \mathbf{h}) < c). \quad (22)$$

In other words, a high-probability lower bound on the AO is a high-probability lower bound on the PO. The premise is that it is often much simpler to lower bound the AO rather than the PO. To be precise, (22) is a slight reformulation of Gordon’s original result proved in [TOH15] (see therein for details).

A.1.2. Convex Gaussian Min-Max Theorem (CGMT)

The asymptotic expressions of this paper build on the CGMT [TOH15]. For ease of reference we summarize here the essential ideas of the framework following the presentation in [TAH18]; please see [TAH18, Section 6] for the formal statement of the theorem and further details. The CGMT is an extension of the GMT and it asserts that the AO in (21b) can be used to tightly infer properties of the original PO in (21a), including the optimal cost and the optimal solution. According to the CGMT [TAH18, Theorem 6.1], if the sets $\mathcal{S}_{\mathbf{w}}$ and $\mathcal{S}_{\mathbf{u}}$ are convex and ψ is continuous *convex-concave* on $\mathcal{S}_{\mathbf{w}} \times \mathcal{S}_{\mathbf{u}}$, then, for any $\nu \in \mathbb{R}$ and $t > 0$, it holds

$$\mathbb{P}(|\Phi(\mathbf{G}) - \nu| > t) \leq 2\mathbb{P}(|\phi(\mathbf{g}, \mathbf{h}) - \nu| > t). \quad (23)$$

In words, concentration of the optimal cost of the AO problem around μ implies concentration of the optimal cost of the corresponding PO problem around the same value μ . Moreover, starting from (23) and under strict convexity

conditions, the CGMT shows that concentration of the optimal solution of the AO problem implies concentration of the optimal solution of the PO to the same value. For example, if minimizers of (21b) satisfy $\|\mathbf{w}^*(\mathbf{g}, \mathbf{h})\|_2 \rightarrow \zeta^*$ for some $\zeta^* > 0$, then, the same holds true for the minimizers of (21a): $\|\mathbf{w}^*(\mathbf{G})\|_2 \rightarrow \zeta^*$ [TAH18, Theorem 6.1(iii)]. Thus, one can analyze the AO to infer corresponding properties of the PO, the premise being of course that the former is simpler to handle than the latter.

A.2. Hard-margin SVM: Identifying the PO and AO problems

The max-margin solution is obtained by solving the following problem:

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \|\boldsymbol{\beta}\|_2 \quad \text{subject to} \quad y_i \mathbf{w}_i^T \boldsymbol{\beta} \geq 1, \quad (24)$$

and is feasible only when the training data is separable. Under such a situation, (24) is equivalent to solving:

$$\min_{\boldsymbol{\beta}} \max_{u_i \leq 0} \|\boldsymbol{\beta}\|_2^2 + \frac{1}{n} \sum_{i=1}^n u_i (y_i \mathbf{w}_i^T \boldsymbol{\beta} - 1). \quad (25)$$

Based on the rotational invariance of the Gaussian measure, we may assume without loss of generality that

$$\boldsymbol{\beta}_0 = [s, 0, \dots, 0]^T,$$

where only the first coordinate is nonzero. For convenience, we also write

$$\mathbf{w}_i = [w_i, \mathbf{v}_i^T]^T \quad \text{and} \quad \boldsymbol{\beta} = [\mu, \tilde{\boldsymbol{\beta}}^T]^T.$$

In this new notation,

$$y_i \sim \text{Rad}(\ell'(s w_i + \sigma z_i)), \quad z_i \sim \mathcal{N}(0, 1). \quad (26)$$

Further using this notation and considering the change of variable $\tilde{u}_i = \frac{u_i}{\sqrt{n}}$, (25) becomes

$$\min_{\mu, \tilde{\boldsymbol{\beta}}} \max_{u_i \leq 0} \mu^2 + \|\tilde{\boldsymbol{\beta}}\|_2^2 + \frac{1}{\sqrt{n}} \sum_{i=1}^n \tilde{u}_i y_i w_i \mu + \frac{1}{\sqrt{n}} \sum_{i=1}^n \tilde{u}_i y_i \mathbf{v}_i^T \tilde{\boldsymbol{\beta}} - \frac{1}{\sqrt{n}} \sum_{i=1}^n \tilde{u}_i. \quad (27)$$

Letting $\mathbf{D}_y = \text{diag}(y_1, \dots, y_n)$, and $\tilde{\mathbf{u}} = [\tilde{u}_1, \dots, \tilde{u}_n]^T$, and writing

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_1^T \\ \vdots \\ \mathbf{w}_n^T \end{bmatrix} = [\mathbf{w}_{n \times 1} | \mathbf{V}_{n \times (p-1)}],$$

leads to the following optimization problem

$$\Phi^{(n)} := \min_{\mu, \tilde{\boldsymbol{\beta}}} \max_{\tilde{\mathbf{u}} \leq 0} \frac{1}{\sqrt{n}} \tilde{\mathbf{u}}^T \mathbf{D}_y \mathbf{V} \tilde{\boldsymbol{\beta}} + \frac{1}{\sqrt{n}} (\tilde{\mathbf{u}}^T \mathbf{D}_y \mathbf{w}) \mu - \frac{1}{\sqrt{n}} \tilde{\mathbf{u}}^T \mathbf{1} + \mu^2 + \|\tilde{\boldsymbol{\beta}}\|_2^2. \quad (28)$$

which has the same form of a primary optimization (PO) problem as required by the CGMT [TOH15], with the single exception that the feasibility sets are not compact. To solve this issue, we pursue the approach developed in [KA19] for the analysis of the hard-margin SVM. We write (28) as follows,

$$\Phi^{(n)} := \inf_{R \geq 0} \sup_{\gamma \geq 0} \min_{\substack{\mu, \tilde{\boldsymbol{\beta}} \\ \mu^2 + \|\tilde{\boldsymbol{\beta}}\|_2^2 \leq R^2}} \max_{\substack{\tilde{\mathbf{u}} \leq 0 \\ \|\tilde{\mathbf{u}}\| \leq \gamma}} \Phi_{R, \gamma}^{(n)}, \quad (29)$$

where

$$\Phi_{R, \gamma}^{(n)} := \min_{\mu, \tilde{\boldsymbol{\beta}}} \max_{\substack{\tilde{\mathbf{u}} \leq 0 \\ \mu^2 + \|\tilde{\boldsymbol{\beta}}\|_2^2 \leq R^2}} \frac{1}{\sqrt{n}} \tilde{\mathbf{u}}^T \mathbf{D}_y \mathbf{V} \tilde{\boldsymbol{\beta}} + \frac{1}{\sqrt{n}} (\tilde{\mathbf{u}}^T \mathbf{D}_y \mathbf{w}) \mu - \frac{1}{\sqrt{n}} \tilde{\mathbf{u}}^T \mathbf{1} + \mu^2 + \|\tilde{\boldsymbol{\beta}}\|_2^2. \quad (30)$$

We identified thus a sequence of (PO) problems $\{\Phi_{R,\gamma}^{(n)}\}$ indexed by (R, γ) , each of which satisfies the compactness conditions on the feasibility sets, as required by the CGMT [TOH15]. We associate each one of them with an auxiliary optimization AO problem that can be written as:

$$\begin{aligned}\phi_{R,\gamma}^{(n)} &:= \min_{\substack{\mu, \tilde{\beta} \\ \mu^2 + \|\tilde{\beta}\|_2^2 \leq R^2}} \max_{\substack{\tilde{\mathbf{u}} \leq 0 \\ \|\tilde{\mathbf{u}}\| \leq \gamma}} \frac{1}{\sqrt{n}} \|\tilde{\beta}\|_2 \mathbf{h}_1^T \mathbf{D}_y \mathbf{u} + \frac{1}{\sqrt{n}} \|\mathbf{D}_y \tilde{\mathbf{u}}\|_2 \mathbf{h}_2^T \tilde{\beta} + \frac{1}{\sqrt{n}} (\tilde{\mathbf{u}}^T \mathbf{D}_y \mathbf{w}) \mu - \frac{1}{\sqrt{n}} \tilde{\mathbf{u}}^T \mathbf{1} + \mu^2 + \|\tilde{\beta}\|_2^2 \\ &= \min_{\substack{\mu, \tilde{\beta} \\ \mu^2 + \|\tilde{\beta}\|_2^2 \leq R^2}} \max_{0 \leq \theta \leq \gamma} \max_{\substack{\tilde{\mathbf{u}} \leq 0 \\ \|\tilde{\mathbf{u}}\| = \theta}} \frac{1}{\sqrt{n}} \|\tilde{\beta}\|_2 \mathbf{h}_1^T \mathbf{D}_y \mathbf{u} + \frac{1}{\sqrt{n}} \|\mathbf{D}_y \tilde{\mathbf{u}}\|_2 \mathbf{h}_2^T \tilde{\beta} + \frac{1}{\sqrt{n}} (\tilde{\mathbf{u}}^T \mathbf{D}_y \mathbf{w}) \mu - \frac{1}{\sqrt{n}} \tilde{\mathbf{u}}^T \mathbf{1} + \mu^2 + \|\tilde{\beta}\|_2^2\end{aligned}\quad (31)$$

where the random vectors $\mathbf{h}_1 \sim \mathcal{N}(0, \mathbf{I}_n)$ and $\mathbf{h}_2 \sim \mathcal{N}(0, \mathbf{I}_{p-1})$, and $\mathbf{1}$ represents the vector whose elements are all 1.

Having identified the AO problem, the next step is to simplify them so as to reduce them to problems involving only optimization problems on only a few number of scalars. This will facilitate in the next step inferring their asymptotic behavior.

A.3. Simplification of the AO problem

To simplify the AO problem, we start by optimizing over the direction of the optimization variable $\tilde{\mathbf{u}}$. In doing so, we obtain:

$$\phi_{R,\gamma}^{(n)} = \min_{\substack{\mu, \tilde{\beta} \\ \mu^2 + \|\tilde{\beta}\|_2^2 \leq R^2}} \max_{0 \leq \theta \leq \gamma} \theta \left\| \left(\frac{1}{\sqrt{n}} \|\tilde{\beta}\|_2 \mathbf{D}_y \mathbf{h}_1 + \frac{\mu}{\sqrt{n}} \mathbf{D}_y \mathbf{w} - \frac{\mathbf{1}}{\sqrt{n}} \right) \right\|_2 + \theta \frac{1}{\sqrt{n}} \mathbf{h}_2^T \tilde{\beta} + \mu^2 + \|\tilde{\beta}\|_2^2, \quad (32)$$

where we used the fact that $\|\mathbf{D}_y \mathbf{u}\|_2 = \|\mathbf{u}\|_2$. We note that for any θ , the direction of $\tilde{\beta}$ that minimizes the objective in (32) is given by $\frac{-\mathbf{h}_2}{\|\mathbf{h}_2\|}$. Using [KA19, Lemma 8], $\phi_{R,\gamma}^{(n)}$ becomes:

$$\phi_{R,\gamma}^{(n)} = \min_{\substack{\mu, \tilde{\beta} \\ \mu^2 + \|\tilde{\beta}\|_2^2 \leq R^2}} \max_{0 \leq \theta \leq \gamma} \theta \left\| \left(\frac{1}{\sqrt{n}} \|\tilde{\beta}\|_2 \mathbf{D}_y \mathbf{h}_1 + \frac{\mu}{\sqrt{n}} \mathbf{D}_y \mathbf{w} - \frac{\mathbf{1}}{\sqrt{n}} \right) \right\|_2 - \theta \frac{\alpha}{\sqrt{n}} \|\mathbf{h}_2\| + \mu^2 + \|\tilde{\beta}\|_2^2.$$

In the above optimization problem, it is easy to see that $\tilde{\beta}$ appears only through its norm. Let $\alpha = \|\tilde{\beta}\|_2$. Optimizing over θ , we obtain the following scalar optimization problem:

$$\phi_{R,\gamma}^{(n)} = \min_{\substack{\mu, \alpha \geq 0 \\ \mu^2 + \alpha^2 \leq R^2}} \gamma \max \left(\left\| \left(\frac{\alpha}{\sqrt{n}} \mathbf{D}_y \mathbf{h}_1 + \frac{\mu}{\sqrt{n}} \mathbf{D}_y \mathbf{w} - \frac{\mathbf{1}}{\sqrt{n}} \right) \right\|_2 - \frac{\alpha}{\sqrt{n}} \|\mathbf{h}_2\|, 0 \right) + \mu^2 + \alpha^2 \quad (33)$$

It is important to note that the new formulation of the AO problem is obtained from a deterministic analysis that did not involve any asymptotic approximation. However, this new formulation is more suitable to understand their asymptotic behavior.

A.4. Asymptotic behavior of the AO problem

A.4.1. Data separability (Proof of Proposition 3.1)

Define the sequence of functions

$$L_n(\alpha, \mu) := \left\| \left(\frac{\alpha}{\sqrt{n}} \mathbf{D}_y \mathbf{h}_1 + \frac{\mu}{\sqrt{n}} \mathbf{D}_y \mathbf{w} - \frac{\mathbf{1}}{\sqrt{n}} \right) \right\|_2 - \frac{\alpha}{\sqrt{n}} \|\mathbf{h}_2\|_2,$$

when $\alpha \geq 0$ and $\alpha^2 + \mu^2 \leq r^2$. It is easy to see that L_n is jointly convex in its arguments (α, μ) and converges almost surely to:

$$\bar{L} : (\alpha, \mu) \mapsto \sqrt{\mathbb{E}(\alpha H + \mu G Y - 1)_-^2} - \alpha \sqrt{\kappa}, \quad (34)$$

where, as in (11),

$$H, G, Z \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad Y \sim \text{Rad}(\ell'(sG + \sigma Z)).$$

At this point, to see the relevance of (34) to our end result, note how the form of the function \bar{L} resembles (34) after substituting:

$$\rho = \frac{\mu}{\sqrt{\mu^2 + \alpha^2}} \quad \text{and} \quad q = \sqrt{\mu^2 + \alpha^2}. \quad (35)$$

Since the convergence of convex functions is uniform over compacts, for R fixed, there exists ϵ such that for all $\alpha \geq 0$ and $\alpha^2 + \mu^2 \leq R^2$,

$$L_n(\alpha, \mu) \geq \bar{L}(\alpha, \mu) - \epsilon.$$

Hence,

$$\phi_{R,\gamma}^{(n)} \geq \min_{\substack{\mu, \alpha \geq 0 \\ \mu^2 + \alpha^2 \leq R^2}} \gamma \max(\bar{L}(\alpha, \mu) - \epsilon, 0) + \mu^2 + \alpha^2 \geq \min_{\substack{\mu, \alpha \geq 0 \\ \alpha \leq R}} \gamma \max(\bar{L}(\alpha, \mu) - \epsilon, 0)$$

From the above inequality, it follows that if

$$\inf_{\substack{\alpha \geq 0 \\ \mu \in \mathbb{R}}} \bar{L}(\alpha, \mu) > 0, \quad (36)$$

then for ϵ chosen sufficiently small (concretely: smaller for instance than $\min_{\mu, \alpha \geq 0} 0.5\bar{L}(\alpha, \mu)$), there exists a constant C independent of γ such that:

$$\phi_{R,\gamma}^{(n)} \geq C\gamma. \quad (37)$$

Next, we prove that (37) holds for n sufficiently large when the following condition is satisfied:

$$\kappa > \inf_{t \in \mathbb{R}} \mathbb{E} (H + tGY)_-^2. \quad (38)$$

In the right-hand side of the equation above, recognize the threshold function $g(\kappa)$ defined in Proposition 3.1. To show the desired, we prove that the optimization over μ can be assumed over a compact. This is because for $0 \leq \alpha \leq R$,

$$\bar{L}(\alpha, \mu) \geq \sqrt{\mathbb{E} (\alpha H + \mu GY - 1)_-^2 \mathbf{1}_{\{H < 0\}}} - \alpha\sqrt{\kappa} \geq \sqrt{\mathbb{E} (\mu GY - 1)_-^2 \mathbf{1}_{\{H < 0\}}} - R\sqrt{\kappa}.$$

As $\sqrt{\mathbb{E} (\mu GY - 1)_-^2 \mathbf{1}_{\{H < 0\}}}$ grows unboundedly large as $\mu \rightarrow \infty$ or $\mu \rightarrow -\infty$, we have:

$$\lim_{|\mu| \rightarrow \infty} \min_{0 \leq \alpha \leq R} \bar{L}(\alpha, \mu) = \infty,$$

thus proving that the minimum over μ is bounded. We can thus consider that μ belongs to a given compact set \mathcal{A} of \mathbb{R} . For fixed $\mu \in \mathcal{A}$,

$$\lim_{\alpha \downarrow 0} \bar{L}(\alpha, \mu) = \sqrt{\mathbb{E} (\mu GY - 1)_-^2}.$$

Furthermore, this convergence is uniform over compact sets due to the convexity of function \bar{L} , hence,

$$\lim_{\alpha \downarrow 0} \inf_{\mu \in \mathcal{A}} \bar{L}(\alpha, \mu) = \inf_{\mu \in \mathcal{A}} \sqrt{\mathbb{E} (\mu GY - 1)_-^2} > 0.$$

There exists thus $\eta > 0$, such that for any $0 \leq \alpha \leq \eta$,

$$\min_{0 \leq \alpha \leq \eta} \inf_{\mu \in \mathcal{A}} \bar{L}(\alpha, \mu) > 0. \quad (39)$$

On the other hand,

$$\min_{\eta \leq \alpha \leq R} \inf_{\mu \in \mathcal{A}} \bar{L}(\alpha, \mu) \geq \inf_{\substack{t_2 \geq 0 \\ t_1 \in \mathbb{R}}} \eta \left(\sqrt{\mathbb{E} \left(H + t_1 GY - \frac{1}{t_2} \right)_-^2} - \sqrt{\kappa} \right) \geq \eta \inf_{t_1 \in \mathbb{R}} \left(\sqrt{\mathbb{E} (H + t_1 GY)_-^2} - \sqrt{\kappa} \right). \quad (40)$$

Combining (39) and (40), we conclude the desired, i.e. if (38) holds, $\phi_{r,\gamma}^{(n)} \geq C\gamma$.

We will now prove that (37) implies that the data are not separable or equivalently that $\Phi^{(n)} = \infty$ with probability 1. Denote by \mathcal{E} the following event:

$$\mathcal{E} := \left\{ \bigcup_{k=1}^{\infty} \bigcap_{m=1}^{\infty} \left\{ \phi_{k,m}^{(n)} \geq Cm \right\}, \text{ for } n \text{ and } p \text{ sufficiently large} \right\}.$$

Obviously $\mathbb{P}[\mathcal{E}] = 1$ as \mathcal{E} is the countable union and intersection of events with probability 1. This allows us to show that $\lim_{n \rightarrow \infty} \mathbb{P} \left[\bigcup_{d=1}^{\infty} \left\{ \Phi^{(n)} \leq d \right\} \right] = 0$, which translates into the unboundedness of the primary optimization problem. To see this, take $x > 0$. If $\Phi^{(n)} \neq \infty$, for $\epsilon > 0$ sufficiently small, there exists $k \in \mathbb{N}$ such that $\Phi^{(n)} \geq \Phi_{k,m}^{(n)} - \epsilon$ for all $m \in \mathbb{N}$. Hence,

$$\mathbb{P} \left[\left\{ \Phi^{(n)} \neq \infty \right\} \cup \left\{ \Phi^{(n)} \leq x \right\} \right] \leq \mathbb{P} \left[\bigcup_{k=1}^{\infty} \bigcap_{m=1}^{\infty} \left\{ \Phi_{k,m}^{(n)} \leq x + \epsilon \right\} \right]. \quad (41)$$

(42)

For $m \in \mathbb{N}^*$, the events $\mathcal{E}_k = \left\{ \bigcap_{m=1}^{\infty} \left\{ \Phi_{k,m}^{(n)} \leq x \right\} \right\}$ form an increasing sequence of events, hence,

$$\mathbb{P} \left[\bigcup_{k=1}^{\infty} \mathcal{E}_k \right] = \lim_{k \rightarrow \infty} \mathbb{P}[\mathcal{E}_k].$$

In a similar way, as $\Phi_{k,m}^{(n)} \geq \Phi_{k,m-1}^{(n)}$, for k a strictly positive integer, the sequence of events

$$\mathcal{E}_{k,m} = \left\{ \Phi_{k,m}^{(n)} \leq x \right\}$$

is decreasing, thus yielding,

$$\mathbb{P} \left[\bigcap_{m=1}^{\infty} \mathcal{E}_{k,m} \right] = \lim_{m \rightarrow \infty} \mathbb{P}[\mathcal{E}_{k,m}].$$

Hence, we have,

$$\mathbb{P} \left[\Phi^{(n)} \leq x \right] \leq \lim_{k \rightarrow \infty} \lim_{m \rightarrow \infty} \mathbb{P} \left[\Phi_{k,m}^{(n)} \leq x + \epsilon \right].$$

Using now the CGMT, we have the following tail inequality:

$$\mathbb{P} \left[\Phi_{k,m}^{(n)} \leq x + \epsilon \right] \leq 2 \mathbb{P} \left[\phi_{k,m}^{(n)} \leq x + \epsilon \right].$$

Hence,

$$\lim_{k \rightarrow \infty} \lim_{m \rightarrow \infty} \mathbb{P} \left[\Phi_{k,m}^{(n)} \leq x + \epsilon \right] \leq \lim_{k \rightarrow \infty} \lim_{m \rightarrow \infty} \mathbb{P} \left[\phi_{k,m}^{(n)} \leq x + \epsilon \right] = \mathbb{P} \left[\bigcup_{k=1}^{\infty} \left\{ \bigcap_{m=1}^{\infty} \left\{ \phi_{k,m}^{(n)} \leq x + \epsilon \right\} \right\} \right].$$

As $\mathbb{P}[\mathcal{E}] = 1$, the sequence of events $A_n := \left\{ \bigcup_{k=1}^{\infty} \bigcap_{m=1}^{\infty} \left\{ \phi_{k,m}^{(n)} \leq x + \epsilon \right\} \right\}$ does not occur infinitely often. Since $\{A_n\}$ are independent, each event being generated by independent realizations \mathbf{h}_1 and \mathbf{h}_2 , by the converge of Borel Cantelli Lemma, $\sum_{n=1}^{\infty} \mathbb{P}(A_n) < \infty$ or equivalently:

$$\mathbb{P} \left[\Phi^{(n)} \leq x \right] < \infty.$$

Using Borel-Cantelli Lemma, we thus obtain for any $x > 0$,

$$\Phi^{(n)} > x$$

for sufficiently large n . Hence, $\Phi^{(n)} = \infty$ with probability 1 for n large. This completes the proof of Proposition 3.1.

A.4.2. Asymptotic behavior of the (AO) problem when data are separable (Proof of Proposition 3.3)

Unique minimizer of $q \mapsto \min_{-1 \leq \rho \leq 1} \eta(\rho, q)$: To begin with, we start by showing that function $q \mapsto \min_{-1 \leq \rho \leq 1} \eta(\rho, q)$ defined in admits a unique zero q^* . Note that necessarily a zero of function $\min_{-1 \leq \rho \leq 1} \eta(\rho, q)$ exists, since from the maximum Theorem [Sun96, Theorem 9.7], the function $q \mapsto \min_{-1 \leq \rho \leq 1} \eta(\rho, q)$ is continuous and tends to 0 when $q \rightarrow \infty$ and to a negative value when $q \rightarrow -\infty$ as per the separability condition.

We will now prove that such a zero must be unique. For that we assume that there exists q_1^* and q_2^* such that:

$$\min_{-1 \leq \rho \leq 1} \eta(\rho, q_1^*) = \min_{-1 \leq \rho \leq 1} \eta(\rho, q_2^*).$$

Denote by ρ_1 and ρ_2 the real values in the interval $[-1, 1]$ such that $\eta(\rho_1, q_1^*) = \min_{-1 \leq \rho \leq 1} \eta(\rho, q_1^*) = 0$, and $\eta(\rho_2, q_2^*) = \min_{-1 \leq \rho \leq 1} \eta(\rho, q_2^*) = 0$. Then,

$$0 = \eta(\rho_1, q_1^*) \leq \eta(\rho_2, q_1^*).$$

Since $0 = \eta(\rho_2, q_2^*)$ and $q \mapsto \eta(\rho_2, q)$ is decreasing, $q_1^* \geq q_2^*$. Similarly, we can use the same reasoning to prove that $q_2^* \geq q_1^*$. We thus necessarily have $q_1^* = q_2^*$, which proves the uniqueness of q^* .

Unique minimizer of $\rho \mapsto \eta(\rho, q^*)$: Next, we show that the function $\rho \mapsto \eta(\rho, q^*)$ has a unique minimizer. Towards this goal, we define the function

$$\tilde{\eta} : [-1, 1] \times [0, 1] \rightarrow \mathbb{R}, (\rho_1, \rho_2) \mapsto \mathbb{E} \left[\left(\rho_1 GY + \rho_2 H - \frac{1}{q^*} \right)_-^2 \right] - (1 - \rho_1^2) \kappa.$$

We consider minimization of this when $-1 \leq \rho_1 \leq 1$ and $\rho_2 \in [0, 1]$ with $\rho_1^2 + \rho_2^2 \leq 1$. Obviously, $\rho_1 \mapsto \min_{0 \leq \rho_2 \leq 1} \tilde{\eta}(\rho_1, \rho_2)$ is strictly convex and thus has a unique minimizer ρ_1 which we denote by ρ^* . We will prove now that necessarily the minimizer of function $\min_{0 \leq \rho_2 \leq 1} \tilde{\eta}(\rho^*, \rho_2)$ is $\rho_2^* = \sqrt{1 - (\rho^*)^2}$. It suffices for that to check that function $\rho_2 \mapsto \tilde{\eta}(\rho^*, \rho_2)$ is increasing on $(0, 1]$. This is because:

$$\begin{aligned} \frac{d}{d\rho_2} \tilde{\eta}(\rho_2, \rho^*) &= \frac{d}{d\rho_2} \mathbb{E} \left[\left(\rho^* GY - \frac{1}{q^*} + \rho_2 H \right)_-^2 \right] - (1 - (\rho^*)^2) \kappa \\ &= \frac{d}{d\rho_2} \mathbb{E}_H \left[\int_{\frac{1}{q^*} - \rho_2 H}{\frac{1}{q^*} - \rho_1} (\rho^* x - \frac{1}{q^*})^2 f_{GY}(x) dx \right] \\ &= \mathbb{E} \left[\frac{H}{\rho_1} \int_{\frac{1}{q^*} - \rho_2 H}{\frac{1}{q^*} - \rho_1} (\rho^* x - \frac{1}{q^*})^2 f_{GY}(x) dx \right] \\ &= \frac{\rho_2}{\rho_1^2} \mathbb{E}_H \left[\int_{\frac{1}{q^*} - \rho_2 H}{\frac{1}{q^*} - \rho_1} (\rho^* x - \frac{1}{q^*})^2 f_{GY}(x) dx \right] > 0, \end{aligned}$$

where f_{GY} is the probability density function associated with the random variable GY .

Convergence: With the uniqueness of the parameters at hand, we will now prove the convergence result. In a similar manner to the proof of the separability condition and along the same lines as [KA19], it suffices to check the following convergence for any $\epsilon > 0$ and k sufficiently large:

$$\mathbb{P} \left[\lim_{m \rightarrow \infty} \phi_{k,m}^{(n)} \leq (q^*)^2 - \epsilon, \text{ i.o.} \right] = 0, \quad (43)$$

$$\mathbb{P} \left[\lim_{m \rightarrow \infty} \phi_{k,m}^{(n)} \geq (q^*)^2 + \epsilon, \text{ i.o.} \right] = 0. \quad (44)$$

We will only prove (43) as (44) follows in a similar manner. Recall from Section A.4.1 that

$$\phi_{k,m}^{(n)} = \min_{\substack{\alpha \geq 0, \mu \\ \alpha^2 + \mu^2 \leq k^2}} m \max(L_n(\alpha, \mu), 0) + \alpha^2 + \mu^2.$$

Hence,

$$\phi_{k,m}^{(n)} \geq \min_{\substack{\alpha \geq 0, \mu \\ \alpha^2 + \mu^2 \leq k^2 \\ L_n(\alpha, \mu) \leq 0}} \alpha^2 + \mu^2.$$

The function $(\alpha, \mu) \mapsto L_n(\alpha, \mu)$ is jointly convex in (α, μ) and converges pointwise to $\bar{L}(\alpha, \mu)$ defined in (34). The convergence is thus uniform over the compact $\mathcal{C} := \{\alpha \geq 0, \mu \in \mathbb{R} \mid \alpha^2 + \mu^2 \leq k^2\}$. For any δ sufficiently small, we thus have for n sufficiently large, $(\alpha, \mu) \in \mathcal{C}$,

$$L_n(\alpha, \mu) \geq \bar{L}(\alpha, \mu) - \delta.$$

Therefore,

$$\phi_{k,m}^{(n)} \geq \min_{\substack{\alpha \geq 0, \mu \\ \alpha^2 + \mu^2 \leq k^2 \\ \bar{L}(\alpha, \mu) \leq \delta}} \alpha^2 + \mu^2.$$

To conclude, it suffices to prove that for all δ sufficiently small,

$$\min_{\substack{\alpha \geq 0, \mu \\ \alpha^2 + \mu^2 \leq k^2 \\ \bar{L}(\alpha, \mu) \leq \delta}} \alpha^2 + \mu^2 \geq \left[\min_{\substack{\alpha \geq 0, \mu \\ \alpha^2 + \mu^2 \leq k^2 \\ \bar{L}(\alpha, \mu) \leq 0}} \alpha^2 + \mu^2 \right] - \frac{\epsilon}{2},$$

a fact that follows directly from [KA19, Lemma 9]. We finish the proof of (43) by recalling the mapping in (35) and noting that,

$$\min_{\substack{\alpha \geq 0, \mu \\ \alpha^2 + \mu^2 \leq k^2 \\ \bar{L}(\alpha, \mu) \leq 0}} \alpha^2 + \mu^2 \geq \min_{\substack{\alpha \geq 0, \mu \\ \bar{L}(\alpha, \mu) \leq 0}} \alpha^2 + \mu^2 = \min_{\substack{\rho \in [-1, 1], q \geq 0 \\ \eta(\rho, q) \leq 0}} q^2.$$

A.5. Proof sketch of Proposition 3.2

Proposition 3.2 essentially follows from [TPT19, Thm. III.1]. Specifically, [TPT19] provides asymptotic characterizations for the cosine similarity of ERM optimization problems for general loss and link (cf. [TPT19, Eqn. (36)]) functions. Minor adaptations are needed to account for the specific feature selection model of Section 2.1 and to yield results for the risk $\mathcal{R}(\hat{\beta})$.

Recall from (26) that the responses y_i depend on the feature vectors \mathbf{w}_i used for training as follows:

$$y_i \sim \text{Rad} \left(\ell'(\beta_0^T \mathbf{w}_i + \sigma z_i) \right), \quad z_i \sim \mathcal{N}(0, 1). \quad (45)$$

This is a special case of [TPT19, Eqn. (36)] for the following choice of a (random) link function therein:

$$f(t) = \text{Rad}(\ell'(st + \sigma Z)), \quad Z \sim \mathcal{N}(0, 1).$$

Also, by rotational invariance of the Gaussian distribution of the feature vectors, we can assume without loss of generality that $\beta_0 = [s, 0, \dots, 0]^T$. Further denote

$$\mathbf{w}_i = [w_i, \tilde{\mathbf{w}}_i^T]^T \quad \text{and} \quad \hat{\beta} = [\hat{\beta}_1, \tilde{\beta}_*^T]^T.$$

Then, it follows from the analysis in [TPT19] that (see also Figure 7)

$$\hat{\beta}_1 \xrightarrow{P} \mu \quad \text{and} \quad \|\tilde{\beta}_*\|_2 \xrightarrow{P} \alpha, \quad (46)$$

where $\alpha > 0$ and μ are solutions of the following system of equations:

$$0 = \mathbb{E} [Y G \mathcal{M}'_{\rho,1}(\alpha H + \mu Y G; \lambda)], \quad (47a)$$

$$\alpha \kappa = \lambda \mathbb{E} [H \mathcal{M}'_{\rho,1}(\alpha H + \mu Y G; \lambda)], \quad (47b)$$

$$\alpha^2 \kappa = \lambda^2 \mathbb{E} [(\mathcal{M}'_{\rho,1}(\alpha H + \mu Y G; \lambda))^2]. \quad (47c)$$

Using differentiability properties of the Moreau envelope function, the system of equations above can be equivalently written as follows

$$\begin{aligned}\mu &= \mathbb{E} [Y G \cdot \text{prox}_\rho(\alpha H + \mu Y G; \lambda)], \\ \alpha \kappa &= \lambda \mathbb{E} [H \cdot \rho'(\text{prox}_\rho(\alpha H + \mu Y G; \lambda))], \\ \alpha^2 \kappa &= \lambda^2 \mathbb{E} [(\rho'(\text{prox}_\rho(\alpha H + \mu Y G; \lambda)))^2].\end{aligned}$$

Uniqueness of the solution to this system of equations for the logistic model is claimed in [SC18].

Finally, applying Gaussian integration by parts and [DM16, Prop. A.1], the second equation above is equivalent to the following

$$\lambda \cdot \mathbb{E} \left[\frac{\rho''(\text{prox}_\rho(\alpha H + \mu Y G; \lambda))}{1 + \lambda \rho''(\text{prox}_\rho(\alpha H + \mu Y G; \lambda))} \right] = \kappa,$$

leading to the expressions in Proposition 3.2 for $V = GY$.

From (46) we have,

$$\mathcal{C}(\hat{\beta}) = \frac{s \hat{\beta}_1}{\|\boldsymbol{\eta}_0\|_2 \sqrt{\hat{\beta}_1^2 + \|\tilde{\boldsymbol{\beta}}_*\|_2^2}} \xrightarrow{P} \frac{s \mu}{r \sqrt{\mu^2 + \alpha^2}}, \quad (49)$$

and

$$\begin{aligned}\mathcal{R}(\hat{\beta}) &= \mathbb{P}(y \mathbf{w}^T \hat{\beta} < 0) = \mathbb{P}\left(\text{Rad}(\ell'(\mathbf{w}^T \boldsymbol{\beta}_0 + \mathbf{z}^T \boldsymbol{\gamma}_0)) \cdot (\hat{\beta}_1 w_1 + \tilde{\boldsymbol{\beta}}_*^T \tilde{\mathbf{w}}_i) < 0\right) \\ &= \mathbb{P}\left(\text{Rad}(\ell'(w_1 s + \mathbf{z}^T \boldsymbol{\gamma}_0)) \cdot (\hat{\beta}_1 w_1 + \tilde{\boldsymbol{\beta}}_*^T \tilde{\mathbf{w}}_i) < 0\right) \\ &\xrightarrow{P} \mathbb{P}\left(\text{Rad}(\ell'(s G + \sigma Z)) \cdot (\mu G + \alpha H) < 0\right) \\ &= \mathbb{P}(Y \cdot (\mu G + \alpha H) < 0) = \mathbb{P}(\mu Y G + \alpha H < 0),\end{aligned} \quad (50)$$

where as always G, H, Z and Y follow our notation in (11) and the last inequality follows since $HY \sim \mathcal{N}(0, 1)$.

A.6. Proof sketch of SVM for GM model

Recall that $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$, are n independent observations drawn from two classes \mathcal{C}_{-1} and \mathcal{C}_{+1} of cardinalities n_{-1}, n_{+1} , such that $n = n_{-1} + n_{+1}$. We assume that the data vector \mathbf{x}_i in class \mathcal{C}_a is given by:

$$\mathbf{x}_i = y_i \boldsymbol{\eta}_0 + \mathbf{z}_i$$

where $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $y_i = +1$ if $\mathbf{x}_i \in \mathcal{C}_{+1}$ or -1 otherwise. We assume that the classifier has access only to the first p components of each observations \mathbf{x}_i that we denote by vector \mathbf{w}_i , where $\mathbf{w}_i = y_i \boldsymbol{\beta}_0 + \mathbf{v}_i$. We consider the regime in which:

$$\frac{d}{n} \rightarrow \zeta > 0 \quad \text{and} \quad \frac{p}{n} \rightarrow \kappa,$$

and $\frac{n_{\pm 1}}{n} \rightarrow \pi_{\pm 1}$ where the prior probabilities are constants in $(0, 1)$. We assume that $\|\boldsymbol{\eta}_0\|_2^2 \xrightarrow{P} r^2$ and $\|\boldsymbol{\beta}_0\|_2^2 \xrightarrow{P} s^2$.

We start the analysis of the the max-margin classifier by (25). Specifically, replacing $\mathbf{w}_i = y_i \boldsymbol{\beta}_0 + \mathbf{v}_i$, (25) becomes:

$$\min_{\boldsymbol{\beta}} \max_{u_i \leq 0} \|\boldsymbol{\beta}\|_2^2 + \frac{1}{\sqrt{n}} \sum_{i=1}^n u_i \boldsymbol{\beta}_0^T \boldsymbol{\beta} + \frac{1}{\sqrt{n}} \sum_{i=1}^n u_i y_i \mathbf{v}_i^T \boldsymbol{\beta} - \frac{1}{\sqrt{n}} \sum_{i=1}^n u_i.$$

Differently from the logistic regression model, in the Gaussian-mixture model, y_i and \mathbf{v}_i are independent. Call $\tilde{\mathbf{v}}_i = y_i \mathbf{v}_i$. Then $\tilde{\mathbf{v}}_i$ has the same distribution as \mathbf{v}_i by symmetry of the Gaussian distribution. Let $\mathbf{V} = [\tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_n]^T$ and $\mathbf{u} = [u_1, \dots, u_n]^T$. Then, in matrix form, the above problem becomes:

$$\min_{\boldsymbol{\beta}} \max_{u_i \leq 0} \frac{1}{\sqrt{n}} \mathbf{u}^T \mathbf{V} \boldsymbol{\beta} - \frac{1}{\sqrt{n}} \mathbf{u}^T \mathbf{1} + \|\boldsymbol{\beta}\|_2^2 + \frac{\mathbf{1}^T \mathbf{u}}{\sqrt{n}} \boldsymbol{\beta}_0^T \boldsymbol{\beta}.$$

We recognize thus a (PO) problem with which we associate the following (AO) problem:

$$\min_{\beta} \max_{\mathbf{u}_i \leq 0} \frac{1}{\sqrt{n}} \|\beta\|_2 \mathbf{h}_1^T \mathbf{u} + \frac{1}{\sqrt{n}} \|\mathbf{u}\|_2 \mathbf{h}_2^T \beta + \frac{\mathbf{1}^T \mathbf{u}}{\sqrt{n}} \beta_0^T \beta - \frac{1}{\sqrt{n}} \mathbf{u}^T \mathbf{1} + \|\beta\|_2^2.$$

We decompose β as:

$$\beta = \alpha_1 \beta_0 + \alpha_2 \beta_{\perp},$$

where β_{\perp} is orthogonal to β_0 and $\|\beta_{\perp}\| = 1$. With this notation at hand, the (AO) becomes:

$$\min_{\beta} \max_{\mathbf{u} \leq 0} \frac{1}{\sqrt{n}} \|\beta\| \mathbf{h}_1^T \mathbf{u} - \frac{1}{\sqrt{n}} \|\mathbf{u}\| \mathbf{h}_2^T \beta + \frac{\mathbf{1}^T \mathbf{u}}{\sqrt{n}} \beta_0^T \beta - \frac{1}{\sqrt{n}} \mathbf{u}^T \mathbf{1} + \alpha_1^2 \|\beta_0\|_2^2 + \alpha_2^2.$$

Optimizing over β_{\perp} , we obtain that at optimum $\beta_{\perp} = \frac{\mathbf{h}_2}{\|\mathbf{h}_2\|}$ and using the fact that $\frac{1}{\sqrt{n}} \mathbf{h}_2^T \beta_0 \xrightarrow{P} 0$ the AO problem becomes:

$$\min_{\alpha_1, \alpha_2} \max_{\mathbf{u} \leq 0} \alpha_1^2 \|\beta_0\|_2^2 + \alpha_2^2 + \frac{1}{\sqrt{n}} \mathbf{u}^T \left(\sqrt{\alpha_1^2 \|\beta_0\|_2^2 + \alpha_2^2} \mathbf{h}_1 + \alpha_1 \|\beta_0\|_2^2 \mathbf{1} - \mathbf{1} \right) - \alpha_2 \|\mathbf{u}\|_2 \frac{\|\mathbf{h}_2\|}{\sqrt{n}}.$$

Let $\theta = \|\mathbf{u}\|_2$ and set $q^2 = \alpha_1^2 \|\beta_0\|_2^2 + \alpha_2^2$. Optimizing over the direction of \mathbf{u} , the above optimization problem becomes:

$$\min_{\substack{q \geq 0 \\ \alpha_1^2 \|\beta_0\|_2^2 + \alpha_2^2 \leq q^2}} \max_{\theta \geq 0} q^2 + \frac{\theta}{\sqrt{n}} \left\| (q \mathbf{h}_1 + \alpha_1 \|\beta_0\|_2^2 \mathbf{1} - \mathbf{1})_- \right\|_2 - \theta \alpha_2 \frac{\|\mathbf{h}_2\|}{\sqrt{n}}.$$

Let $\rho = \frac{\alpha_1 \|\beta_0\|_2}{q}$. Then, we obtain:

$$\min_{\substack{q \geq 0 \\ |\rho| \leq 1}} \max_{\theta \geq 0} q^2 + q\theta \left\{ \frac{\left\| (\mathbf{h}_1 + s \rho \mathbf{1} - \frac{1}{q} \mathbf{1})_- \right\|}{\sqrt{n}} - \frac{\|\mathbf{h}_2\|}{\sqrt{n}} \sqrt{1 - \rho^2} \right\}. \quad (51)$$

At this point, recognize that the optimization above is of very similar form to the corresponding optimization for the logistic model in (33) (under the mapping $\sqrt{\mu^2 + \alpha^2} \leftrightarrow q$ and $\mu \leftrightarrow \rho q$). Thus, similar to the logistic regression model (see Sec. A.4.1), if κ is such that:

$$\min_{-1 \leq \rho \leq 1} \sqrt{\mathbb{E} (H + \rho s)_-^2} - \sqrt{\kappa} \sqrt{1 - \rho^2} > 0$$

then, the hard-margin SVM will not converge. The above condition is equivalent to:

$$\kappa \leq \min_{-1 \leq \rho \leq 1} \mathbb{E} \left(\frac{1}{\sqrt{1 - \rho^2}} H + \frac{\rho}{\sqrt{1 - \rho^2}} s \right)_-^2$$

or, setting $t_1 = \frac{\rho}{1 - \rho^2} \geq 0$,

$$\kappa \leq \min_{t_1 \geq 0} \mathbb{E} (\sqrt{1 + t_1^2} H + t_1 s)_-^2.$$

On the other hand, if $\kappa > \min_{t_1 \geq 0} \mathbb{E} (\sqrt{1 + t_1^2} H + t_1 s)_-^2$, then in view of (51) define

$$\eta_G(\rho, q) = \mathbb{E} \left(H + \rho s - \frac{1}{q} \right)_-^2 - (1 - \rho^2) \kappa.$$

Then, $q \mapsto \min_{0 \leq \rho \leq 1} \eta_G(\rho, q)$ admits a unique zero, which we denote by q^* [KA19]. Moreover let ρ^* be the unique real in $(0, 1)$ such that $\eta_G(\rho^*, q^*) = 0$. Then, (see also (49) and (50))

$$\mathcal{C}(\widehat{\beta}) \xrightarrow{P} \frac{\rho^* s}{r} \quad \text{and} \quad \|\widehat{\beta}\|_2 \xrightarrow{P} q^* \quad \text{and} \quad \mathcal{R}(\widehat{\beta}) = \mathbb{P}(y \text{sign}(\mathbf{w}^T \widehat{\beta})) \xrightarrow{P} Q(\rho^* s).$$

A.7. On the equation $g(\kappa) = \kappa$

In previous sections we argued that data are separable with probability approaching one if the overparameterization threshold satisfies $\kappa > g(\kappa)$. Recall that

$$g(\kappa) := \min_{t \in \mathbb{R}} \mathbb{E} (H + t G Y)_-^2.$$

To see why this is a function of κ recall the definition of the random variable Y in (11) and note that the signal strength of the known features s is a function of κ : $s = s(\kappa)$. Thus, the RHS above is a function of $s(\kappa)$, or equivalently, of $\varepsilon := s(\kappa)/r \in [0, 1]$. In Lemma A.1 below we show that the equation $g(\kappa) = \kappa$ has a unique solution $\kappa_* \in [0, 1/2]$.

First, we introduce some useful notation. For a (random) variable X we use the following shorthand:

$$Y_X := \text{Rad}(\ell'(X)).$$

Lemma A.1. Fix $r > 0$. Let $G, H, Z \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and (with a slight abuse of notation) consider the function $g : [0, 1] \rightarrow \mathbb{R}$

$$g(\varepsilon) := \min_{t \in \mathbb{R}} \mathbb{E} (H + t G Y_{rQ})_-^2, \quad (52)$$

where the random variable Q is defined as $Q = \varepsilon G + \sqrt{1 - \varepsilon^2} Z$. Then, $g(0) = 1/2$ and g is decreasing in $[0, 1]$.

Proof. First, we rewrite $g(\varepsilon)$ in a more convenient form. Note that

$$H + t G Y_{rQ} = H Y_{rQ}^2 + t G Y_{rQ} = (H Y_{rQ} + t G) Y_{rQ} \sim (H + t G) Y_{rQ} \sim T_1 Y_{rQ},$$

where

$$\begin{bmatrix} T_1 \\ Q \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 + t^2 & t\varepsilon \\ t\varepsilon & 1 \end{bmatrix} \right).$$

Here, we have used the fact that $Y_{rQ} \in \{\pm 1\}$ and Q is independent of H . By further decomposing T_1 on its projection on Q , i.e.,

$$T_1 = t\varepsilon Q + \sqrt{1 + t^2 - \varepsilon^2 t^2} R,$$

with $R \sim \mathcal{N}(0, 1)$, $\mathbb{E}[RQ] = 0$, we find that

$$H + t G Y_{rQ} \sim t\varepsilon Q Y_{rQ} + \sqrt{1 + t^2 - \varepsilon^2 t^2} R.$$

Thus, we have shown that

$$g(\varepsilon) = \min_{t \in \mathbb{R}} \mathbb{E} \left(t\varepsilon Q Y_{rQ} + \sqrt{1 + t^2 - \varepsilon^2 t^2} R \right)_-^2 =: \min_{t \in \mathbb{R}} h_\varepsilon(t).$$

From this, it is easy to check that $g(0) = 1/2$. To prove that the function is decreasing we apply Lemma A.2 below. \square

In what follows, let $Q, R \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

Lemma A.2. Consider the function $h_\varepsilon : \mathbb{R} \rightarrow \mathbb{R}$ parametrized by $\varepsilon \in [0, 1]$ defined as follows

$$h_\varepsilon(t) := \mathbb{E} \left(t\varepsilon Q Y_{rQ} + \sqrt{1 + t^2(1 - \varepsilon^2)} R \right)_-^2.$$

Then, the function $g(\varepsilon) := \inf_{t \in \mathbb{R}} h_\varepsilon(t)$ is decreasing in ε .

Proof. Let $\varepsilon_1 < \varepsilon_2$. It will suffice to show that $g(\varepsilon_1) \geq g(\varepsilon_2)$. Denote $\delta := \frac{\varepsilon_1}{\varepsilon_2} < 1$. For any $t \in \mathbb{R}$, we have the following chain of inequalities

$$\begin{aligned} h_{\varepsilon_2}(t\delta) &= \mathbb{E} \left(t\varepsilon_1 Q Y_{rQ} + \sqrt{1 - t^2\varepsilon_1^2 + t^2\delta^2} R \right)_-^2 \\ &< \mathbb{E} \left(t\varepsilon_1 Q Y_{rQ} + \sqrt{1 - t^2\varepsilon_1^2 + t^2} R \right)_-^2 = h_{\varepsilon_1}(t), \end{aligned} \quad (53)$$

where the inequality follows from Lemma A.3 and the fact that $\delta < 1$. The desired claim follows by minimizing both sides of the inequality in (53) with respect to t and invoking the definition of $g(\varepsilon)$. \square

Lemma A.3. Fix $t_1 \in \mathbb{R}$ and $t_2, r \geq 0, t_3 > 0$. Consider the function $f(x) = \mathbb{E}[(t_1 Q Y_{rQ} + \sqrt{t_2 + t_3 x} R)_-^2]$. Then, the function f is increasing in $[0, 1]$.

Proof. The claim follows by direct differentiation of f . Formally, let $F(Q, R; x) := (t_1 Q Y_{rQ} + \sqrt{t_2 + t_3 x} R)_-^2$. Then, $|\frac{\partial}{\partial x} F(Q, R; x)| := \frac{t_3}{\sqrt{t_2 + t_3 x}} |R(t_1 Q Y_{rQ} + \sqrt{t_2 + t_3 x} R)_-| \leq \frac{t_3}{\sqrt{t_2}} (|t_1| |R| |Q| + \sqrt{t_2 + t_3} R^2) =: M$ a.s. for all $x \in [0, 1]$. Clearly, $\mathbb{E}[M] < \infty$, thus, by dominated convergence theorem, we have

$$f'(x) = \frac{t_3}{\sqrt{t_2 + t_3 x}} \mathbb{E} [R(t_1 Q Y_{rQ} + \sqrt{t_2 + t_3 x} R)_-] = \frac{t_3^2}{t_2 + t_3 x} \mathbb{E} [\mathbf{1}_{\{t_1 Q Y_{rQ} + \sqrt{t_2 + t_3 x} R < 0\}}] > 0,$$

where the second equality follows by Gaussian integration by parts. □




Fig. 5: Monotonicity of the threshold function of Lemma A.1.

B. ADDITIONAL NUMERICAL RESULTS




Fig. 6: Cosine similarity as a function of the overparameterization ratio κ for binary logistic regression under the polynomial feature selection rule. The setting here is the same as in Figure 1. See also Section 4.2




Fig. 7: Simulation results (\times) vs theoretical predictions (---) for the logistic model for signal strength $r = 10$. Without loss of generality, we set $\beta_0 = [s, 0, \dots, 0]$ and also decompose $\hat{\beta} = [\hat{\beta}_1, \hat{\beta}_*^T]^T$. (**Upper left**) Absolute risk $\mathcal{R}(\hat{\beta})$; (**Upper right**) cosine similarity $\mathcal{C}(\hat{\beta})$; (**Lower left**): solution μ to (14) (---) and averages of $\hat{\beta}_1$ (\times); (**Upper right**) solution α to (14) (---) and averages of $\|\hat{\beta}_*\|_2$ (\times). The bottom two figures experimentally validate (46).




Fig. 8: Simulation results (\times) vs theoretical predictions ($—$) for the GM model for signal strength $r^2 = 10$. Without loss of generality, we set $\beta_0 = [s, 0, \dots, 0]$ and also decompose $\hat{\beta} = [\hat{\beta}_1, \tilde{\beta}_*^T]^T$. (**Upper left**) Absolute risk $\mathcal{R}(\hat{\beta})$; (**Upper right**) cosine similarity $\mathcal{C}(\hat{\beta})$; (**Lower left**): solution μ to (15) ($—$) and averages of $\hat{\beta}_1$ (\times); (**Upper right**) solution α to (15) ($—$) and averages of $\|\tilde{\beta}_*\|_2$ (\times).