

THE BREAKDOWN OF GAUSSIAN UNIVERSALITY IN CLASSIFICATION OF HIGH-DIMENSIONAL LINEAR FACTOR MIXTURES

Xiaoyi Mai

Institut de Mathématiques de Toulouse (IMT)
University of Toulouse-Jean Jaurès
Toulouse, France
xiaoyi.mai@math.univ-toulouse.fr

Zhenyu Liao*

EIC
Huazhong University of Science & Technology
Wuhan, China
zhenyu_liao@hust.edu.cn

ABSTRACT

The assumption of Gaussian or Gaussian mixture data has been extensively exploited in a long series of precise performance analyses of machine learning (ML) methods, on large datasets having comparably numerous samples and features. To relax this restrictive assumption, subsequent efforts have been devoted to establish “Gaussian equivalent principles” by studying scenarios of Gaussian universality where the asymptotic performance of ML methods on non-Gaussian data remains unchanged when replaced with Gaussian data having the *same mean and covariance*. Beyond the realm of Gaussian universality, there are few exact results on how the data distribution affects the learning performance.

In this article, we provide a precise high-dimensional characterization of empirical risk minimization, for classification under a general mixture data setting of *linear factor models* that extends Gaussian mixtures. The Gaussian universality is shown to break down under this setting, in the sense that the asymptotic learning performance depends on the data distribution *beyond* the class means and covariances. To clarify the limitations of Gaussian universality in the classification of mixture data and to understand the impact of its breakdown, we specify conditions for Gaussian universality and discuss their implications for the choice of loss function.

1 INTRODUCTION

Modern machine learning (ML) is dealing with increasingly larger datasets having high-dimensional features, using large-scale models of increasing complexity. Understanding the generalization ability of these large-scale ML models has become a major focus of research efforts (Bartlett et al., 2020; Loog et al., 2020; Nakkiran et al., 2021). One analysis approach of growing popularity is the high-dimensional asymptotic analysis (Liao & Couillet, 2019; Taheri et al., 2021a; Celentano & Montanari, 2022; Hastie et al., 2022; Loureiro et al., 2022; Celentano et al., 2023), by considering the regime where the number n of samples and the dimension p of feature vectors are commensurately large. Despite its asymptotic nature, this approach turns out to be surprisingly effective in explaining and predicting modern ML practice: the asymptotic performance curves are repetitively observed to closely match the average empirical performance curves on realistic datasets of moderate size and dimension (Couillet & Liao, 2022), and are particularly *different* from those offered by, e.g., classical maximum likelihood theory (Bean et al., 2013; Sur & Candès, 2019; Taheri et al., 2021b).

To analytically characterize the generalization performance of large-scale ML models in the aforementioned high-dimensional regime, advanced statistical tools such as the approximate message passing (Donoho & Montanari, 2016; Barbier et al., 2019), convex Gaussian min-max theorem (Tham-poulidis et al., 2018; Salehi et al., 2019; Deng et al., 2022; Javanmard & Soltanolkotabi, 2022), replica method (Huang, 2017; Gerace et al., 2020; Maillard et al., 2020), and random matrix theory (RMT) (Couillet & Liao, 2022; Mai et al., 2019; Mai & Couillet, 2021) have been carefully adapted to take into account of the nonlinearity of ML models. As these tools apply directly on

*Corresponding author: Zhenyu Liao.

Gaussian data, a majority of high-dimensional asymptotic analyses are performed under Gaussian data models in the context of regression (El Karoui et al., 2013; Donoho & Montanari, 2016; Taheri et al., 2021a; Celentano & Montanari, 2022) or Gaussian mixture models (GMMs) in the context of classification (Mignacco et al., 2020; Thrampoulidis et al., 2020; Refinetti et al., 2021).

Despite this seemingly restrictive assumption of data Gaussianity, the derived high-dimensional asymptotic results have been empirically observed to remain valid on non-Gaussian data, including both synthetic non-Gaussian data and samples drawn from realistic (say image) datasets (Loureiro et al., 2021; Taheri et al., 2021b), hinting at a phenomenon of *Gaussian universality*. This motivated a series of recent works establishing, through, e.g., an one-directional central limit theorem (CLT) argument, a Gaussian equivalent principal (GEP) for high-dimensional ML models ranging from generalized linear models to shallow neural networks (Gerace et al., 2020; Goldt et al., 2022; Hu & Lu, 2022; Montanari & Saeed, 2022; Schröder et al., 2023; Han & Shen, 2023). According to the GEP, the performance of ML methods on non-Gaussian data is asymptotically the same under an equivalent Gaussian model with matching first- and second-order moments. Assuming a conditional one-directional CLT, Dandi et al. (2024) put forward a conditional Gaussian equivalent principle (CGEP) stating that the asymptotic classification error for non-Gaussian mixtures remains unchanged when replaced by a Gaussian mixture model with identical class-conditional means and covariances. This contribution, however, did not specify the conditions required on the mixture data model for this conditional one-directional CLT to hold.

This work is driven by the need to investigate the applicability of CGEP under mixture models and to characterize the impact of non-Gaussian data variations when the CGEP does *not* apply. By considering a more general mixture model (see Definition 1) where classes are described by linear factor models—a fundamental probabilistic framework in statistical inference and generative models (Goodfellow et al., 2016, Chapter 13), our analysis extends a long line of high-dimensional asymptotic performance analyses in classification of Gaussian mixtures (Dobriban & Wager, 2018; Huang, 2017; Liao & Couillet, 2019; Mai & Liao, 2019; Huang & Yang, 2021; Kammoun & Alouini, 2021; Wang & Thrampoulidis, 2021; Pesce et al., 2023). We discuss the validity of CGEP under this linear factor mixture model and specify its conditions. On a technical level, we develop a flexible “leave-one-out” analysis approach, in a similar spirit to the analysis of robust linear regression by El Karoui et al. (2013). The elementary nature of this leave-one-out procedure allows us to extend the approach of high-dimensional asymptotic analysis beyond the realm of Gaussian universality.

Our Contributions. The main findings of this paper are summarized below.

1. We provide in Theorem 1 an asymptotic characterization of ridge-regularized empirical risk minimization (ERM) for classification of data drawn from a linear factor mixture model (LFMM, see Definition 1 below, that generalizes the GMM). This precise characterization gives access to the asymptotic performance on mixture data *beyond Gaussian universality*.
2. Based on the proposed analysis, we study Gaussian universality in the sense of CGEP and provide conditions on LFMM under which the data distribution affects the asymptotic learning behavior *only* via its first two class-conditional moments.
 - We first discuss in Section 5.1 the Gaussian universality on *in-distribution performance* and conclude in Corollary 2 that the training and generalization performances of ERM under a given LFMM remain unchanged under its equivalent GMM (with identical class means and covariances, see Definition 2), if all *informative factors* of the LFMM that significantly correlated with the class label are *class-conditional normal variables*.
 - We then investigate in Section 5.2 the Gaussian universality of *classifier* (see Definition 3) and conclude in Corollary 3 that on a given test set (of arbitrary distribution), the ERM classifier trained on data drawn from an LFMM gives the same asymptotic classification error as the one trained on its equivalent GMM, whenever the square loss is used and/or in the case of class-conditional Gaussian informative factors for LFMM.
3. While it has been known that for high-dimensional GMM, the square loss is optimal in both unregularized (Taheri et al., 2021b) and ridge-regularized (Mai & Liao, 2019) classifications, it is *no longer* the case under the general LFMM due to the breakdown of Gaussian universality. We discuss in Section 5.2 how the suboptimality of square loss under LFMM relates to its effect on the Gaussian universality of the ERM classifier. Our analysis thus opens the door to future investigation on the optimal loss design for *non-Gaussian* data.

2 BACKGROUND ON GAUSSIAN UNIVERSALITY IN HIGH DIMENSIONS

The Gaussian universality phenomenon was observed in many high-dimensional inference or ML problems, where some key statistics such as estimation error or classification accuracy exhibit universal behaviors independent of the data distribution. This phenomenon was extensively exploited to relax the data Gaussianity assumption that underlined many results in high-dimensional statistics, through a universality argument often established with two key ingredients - the law of large numbers (LLN) and the CLT.

Here we briefly review previous findings on Gaussian universality in the high-dimensional regime.

Universality of large sample covariance matrices. It has been long known in RMT that the eigenspectra of large random matrices enjoy universal properties for Gaussian and non-Gaussian entries (Tao et al., 2010; Pastur & Shcherbina, 2011). Fundamentally, Marchenko & Pastur (1967) put forward that for sample covariance matrices of the type $\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top \in \mathbb{R}^{p \times p}$ obtained from n i.i.d. data vectors \mathbf{x}_i of dimension p , the universality on the limiting eigenvalue distribution hinges on the concentration of quadratic forms of \mathbf{x}_i around their expectations, that is

$$\lim_{n,p \rightarrow \infty} (\mathbf{x}_i^\top \mathbf{M} \mathbf{x}_i - \mathbb{E}[\mathbf{x}_i^\top \mathbf{M} \mathbf{x}_i]) / \mathbb{E}[\mathbf{x}_i^\top \mathbf{M} \mathbf{x}_i] = 0, \quad (1)$$

for deterministic $\mathbf{M} \in \mathbb{R}^{p \times p}$. This LLN-type result holds for a wide family of high-dimensional random vectors \mathbf{x}_i . An important example studied by Bai & Silverman (2008) is $\mathbf{x}_i = \Sigma^{\frac{1}{2}} \mathbf{z}_i$ with \mathbf{z}_i of i.i.d. standardized entries with bounded fourth moments and non-negative definite symmetric Σ .

Universality of empirical risk minimization. In line with the universal behavior of large sample covariance matrices, it has been recently demonstrated in a series of works (Gerace et al., 2020; Goldt et al., 2022; Hu & Lu, 2022; Montanari & Saeed, 2022; Schröder et al., 2023) that random (and deterministic under certain conditions) feature maps can produce feature matrices that, when replaced by “equivalent” Gaussian features with the same first- and second-order moments, yield the same training and/or generalization performance for many ML methods. This GEP has also been proven for data vectors of independent entries in the context of regularized regression (Montanari & Nguyen, 2017; Panahi & Hassibi, 2017; Han & Shen, 2023).

In the context of ERM, the GEP traced back to a CLT on the inner product $\mathbf{x}^\top \beta$ for feature vector $\mathbf{x} \in \mathbb{R}^p$ independent of classifier β living in a subspace $\mathcal{B} \subset \mathbb{R}^p$ containing the ERM solution $\hat{\beta}$, i.e.,

$$\lim_{n,p \rightarrow \infty} \sup_{\beta \in \mathcal{B}} (\mathbb{E}[f(\mathbf{x}^\top \beta)] - \mathbb{E}[f(\mathbf{g}^\top \beta)]) = 0, \quad (2)$$

with $\mathbf{g} \sim \mathcal{N}(\mathbb{E}[\mathbf{x}], \text{Cov}[\mathbf{x}])$ being the “equivalent” Gaussian vector, for any bounded Lipschitz function $f: \mathbb{R} \rightarrow \mathbb{R}$. The one-directional CLT in (2) requires the ERM solution $\hat{\beta}$ to not particularly aligned with any non-Gaussian variation in the feature vector \mathbf{x} , in order to ensure the asymptotic normality of $\mathbf{x}^\top \beta$ per a CLT-type argument.

Universality of empirical risk minimization on mixture data. Inspired by the one-directional CLT condition (2) for GEP in ERM, Dandi et al. (2024) demonstrated the Gaussian universality for mixture models under a key assumption that is a conditional version of (2):

$$\lim_{n,p \rightarrow \infty} \sup_{\beta \in \mathcal{B}} \left(\mathbb{E}[f(\mathbf{x}^\top \beta) | y_{\mathbf{x}} = C] - \mathbb{E}[f(\mathbf{g}_{[C]}^\top \beta)] \right) = 0, \quad (3)$$

where $y_{\mathbf{x}}$ is the class label of \mathbf{x} , C a class modality, and $\mathbf{g}_{[C]} \sim \mathcal{N}(\mathbb{E}[\mathbf{x} | y_{\mathbf{x}} = C], \text{Cov}[\mathbf{x} | y_{\mathbf{x}} = C])$. Under this conditional one-directional CLT in (3), Dandi et al. (2024) showed that the asymptotic training and generalization errors only depend on the class-conditional means and covariances of the mixture model, obeying thus a conditional Gaussian equivalent principle (CGEP).

However, the condition in (3) is *not* verifiable from the data distribution, making it essentially different from the ones given in Section 5.

Other related works established equivalences between Gaussian data and Gaussian mixtures. For classification with random labels $y_{\mathbf{x}} \sim \text{Unif}(\{-1, 1\})$ generated independently of \mathbf{x} , Gerace et al. (2024) proved that the training loss on GMM is asymptotically equal to that on a *single* Gaussian. Pesce et al. (2023) considered a teacher-student model and showed that when the target label y is generated by a teacher model *uncorrelated* with cluster means, the same asymptotic performance can be obtained by replacing a homoscedastic (i.e., having identical covariance) Gaussian mixture with a single Gaussian.

Universality under elliptical distributions. For “elliptic-like” data vectors of form $\mathbf{x} = a\mathbf{M}\mathbf{u}$ with $a \in \mathbb{R}$ a random scaling variable, $\mathbf{M} \in \mathbb{R}^{p \times d}$ a deterministic matrix and $\mathbf{u} \in \mathbb{R}^d$ a vector of standardized variables satisfying the concentration of quadratic forms in (1) (e.g., $\mathbf{u} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$), El Karoui (2009) revealed a universal limiting spectrum for the sample covariance matrix that is insensitive to the distribution of \mathbf{u} but depends on the scaling variable a .

Due to the existence of the scaling variable a , the one-directional CLT in (2) can not hold unless when conditioned on a . This remark was confirmed by the findings of El Karoui (2018); Adomaityte et al. (2024). For $\mathbf{M} = \mathbf{I}_p$ and \mathbf{u} of i.i.d. entries, El Karoui (2018) characterized the asymptotic error of ridge-regularized regression, which is universal with respect to the distribution of \mathbf{u} but not a . In other words, the GEP collapses while the CGEP with respect to the scaling factor a can still apply in this setting. Adomaityte et al. (2024) considered a mixture model $\mathbf{x} \sim \mathcal{N}(y\boldsymbol{\mu}, a\mathbf{I}_p)$ with label $y = \pm 1$ and random scaling factor a , under which the asymptotic classification error is *non-universal* with respect to the distribution of a . Here with our analysis under LFMM, we show that Gaussian universality may breakdown even for data vectors of concentrated quadratic forms as described in (1), a condition not satisfied by elliptical data vectors due to the presence of a random scaling factor a .

3 PROBLEM SETUP

For a set of n training samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ with feature vectors $\mathbf{x}_i \in \mathbb{R}^p$ and binary labels $y_i \in \{\pm 1\}$, a classifier is trained by minimizing the following ridge-regularized empirical risk:

$$\hat{\beta}_{\ell, \lambda} = \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_i^\top \beta, y_i) + \frac{\lambda}{2} \|\beta\|^2, \quad (4)$$

for some non-negative loss function $\ell: \mathbb{R} \times \{\pm 1\} \rightarrow \mathbb{R}_+$ that evaluates the difference between the classification score $\hat{y}_i = \beta^\top \mathbf{x}_i$ and the corresponding ground-truth label y_i . Data instances \mathbf{x} with negative scores $\beta^\top \mathbf{x}$ will be assigned to the class of label $y = -1$, and those with positive scores to the class annotated by $y = 1$. The addition of the l_2 regularization term with $\lambda > 0$ can improve the generalization through a better bias-variance trade-off, and also ensures the uniqueness of the solution $\hat{\beta}_{\ell, \lambda}$ in the over-parametrized regime where the feature dimension p is greater than the sample size n .

In this paper, we consider convex and continuously differentiable loss functions.

Assumption 1 (Loss function). *The function $\ell(\cdot, y): \mathbb{R} \rightarrow \mathbb{R}_+$ in (4) is convex and continuously differentiable with its (first) derivative different from 0 at the origin. Its second and third derivatives exist and are bounded, except on a finite set of points.*

Assumption 1 holds for the logistic loss $\ell(\hat{y}, y) = -\ln(1/(1 + e^{-y\hat{y}}))$ used in logistic regression, the square loss $\ell(\hat{y}, y) = (y - \hat{y})^2/2$ for least-squares classifier, and the square hinge loss $\ell(\hat{y}, y) = \max\{0, 1 - y\hat{y}\}^2$. Non-smooth losses such as the hinge loss $\ell(\hat{y}, y) = \max\{0, 1 - y\hat{y}\}$ used in SVMs (Schölkopf & Smola, 2018), and the absolute loss $\ell(\hat{y}, y) = |\hat{y} - y|$, fail to meet Assumption 1¹.

In the following, we focus on the ERM in (4), and use the shorthand notation $\hat{\beta}$ for $\hat{\beta}_{\ell, \lambda}$ in (4) unless there is a risk of confusion. We consider the following linear factor mixture model.

Definition 1 (Linear factor mixture model, LFMM). *We say that a data instance $(\mathbf{x}, y) \sim \mathcal{D}_{(\mathbf{x}, y)}$ with class label $y \in \{\pm 1\}$ and class priors $\Pr(y = -1) = \rho$, $\Pr(y = 1) = 1 - \rho$, follows a linear factor mixture model if the corresponding feature vector $\mathbf{x} \in \mathbb{R}^p$ can be expressed as a linear mapping of p independent factors z_1, \dots, z_p as*

$$\mathbf{x} = \sum_{k=1}^p z_k \mathbf{v}_k = \sum_{k=1}^p (ys_k + e_k) \mathbf{v}_k, \quad (5)$$

for linearly independent deterministic vectors $\mathbf{v}_1, \dots, \mathbf{v}_p \in \mathbb{R}^p$ and standardized independent² noises $e_1, \dots, e_p \in \mathbb{R}$ of symmetric distribution. Among the p factors z_1, \dots, z_p , we have

- q **informative factors** z_1, \dots, z_q with deterministic signals $s_k > 0$, $\forall k \in \{1, \dots, q\}$; and
- $p - q$ **noise factors** z_{q+1}, \dots, z_p with $s_k = 0$, $\forall k \in \{q+1, \dots, p\}$.

¹To study these non-smooth losses, a workaround would be to evaluate instead a series of smooth functions that gradually approach the non-smooth functions, so as to retrieve their performance in some carefully taken limit. Such consideration is however beyond the focus of this paper.

²In other words, $\mathbb{E}[e_k] = 0$, $\text{Var}[e_k] = 1$, $\forall k \in \{1, \dots, p\}$.

Note that (5) can be compactly written as $\mathbf{x} = \mathbf{V}\mathbf{z}$, with $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_p] \in \mathbb{R}^{p \times p}$ and $\mathbf{z} = [z_1, \dots, z_p]^\top = [ys_1 + e_1, \dots, ys_q + e_q, e_{q+1}, \dots, e_p]^\top \in \mathbb{R}^p$. The class-conditional means and covariances of \mathbf{x} are therefore given by

$$\boldsymbol{\mu} \equiv \mathbb{E}[\mathbf{x}|y = 1] = \sum_{k=1}^q s_k \mathbf{v}_k \in \mathbb{R}^p, \quad \mathbb{E}[\mathbf{x}|y = -1] = -\boldsymbol{\mu}, \quad (6)$$

$$\boldsymbol{\Sigma} \equiv \text{Cov}[\mathbf{x}|y = \pm 1] = \mathbf{V}\mathbf{V}^\top = \sum_{k=1}^p \mathbf{v}_k \mathbf{v}_k^\top \in \mathbb{R}^{p \times p}. \quad (7)$$

Notice that GMM of form $\mathbf{x} \sim \mathcal{N}(y\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is a special case of LFMM in Definition 1 with exclusively Gaussian noises e_1, \dots, e_p . See also Definition 2 below for the associated equivalent GMM.

Linear factor models are among the most fundamental probabilistic models with latent variables, which underlie many ML methods such as PCA and ICA, and serve as building blocks of deep generative models (Goodfellow et al., 2016, Chapter 13). They are often expressed as:

$$\mathbf{x} = \mathbf{W}\mathbf{h} + \mathbf{b} + \text{noise},$$

where \mathbf{h} is a vector of latent variables, \mathbf{b} a constant bias, and noise stands for an uninformative term of independent Gaussian noises. The LFMM in Definition 1 can be related to this form minus the bias \mathbf{b} . Our framework requires the clusters to have opposite means (therefore $\mathbf{b} = \mathbf{0}_p$), which can be satisfied through a centering operation on the original data space.

Our analysis applies under the following assumption on the distribution of LFMM.

Assumption 2 (Distribution of LFMM). *We consider, for the LFMM in Definition 1, that (i) the factors z_1, \dots, z_p have bounded fourth moments and (ii) the signal subspace $\text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_q\}$ is orthogonal to the noise subspace $\text{Span}\{\mathbf{v}_{q+1}, \dots, \mathbf{v}_p\}$.*

The condition of bounded fourth moment for z_1, \dots, z_p in Item (i) of Assumption 2 is required for some concentration results in our high-dimensional asymptotic analysis and Item (ii) separates the informative signal subspace from the noise subspace (in which no classifier can achieve better performance than random guesses).

We position ourselves under the following high-dimensional asymptotic setting, where the feature dimension p and sample size n are both large and comparable.

Assumption 3 (High-dimensional regime). *As $n \rightarrow \infty$ with fixed $n/p \in (0, \infty)$, we have, for the LFMM in Definition 1 that (i) $\|\boldsymbol{\mu}\|, \|\boldsymbol{\Sigma}\|, \|\boldsymbol{\Sigma}^{-1}\| = \Theta(1)$ and (ii) $s_1, \dots, s_q = \Theta(1)$ with fixed q .*

In plain words, Assumption 3 says that the ratio n/p , or the number of samples per dimension, remains finite in high dimensions. Item (i) of Assumption 3 ensures, by bounding $\|\boldsymbol{\mu}\|$ and $\|\boldsymbol{\mu}\|^{-1}$, that the distance between the LFMM class centers is comparable to 1. It also guarantees, by controlling $\|\boldsymbol{\Sigma}\|$ and $\|\boldsymbol{\Sigma}^{-1}\|$, that the variation of feature vector \mathbf{x} on any direction in \mathbb{R}^p is also comparable to 1. This implies that the feature vector \mathbf{x} does not live in a subspace of dimension smaller than p . The fixed number q of informative factors in Item (ii) of Assumption 3 is a consequence of $\|\boldsymbol{\mu}\| = \|\sum_{k=1}^q s_k \mathbf{v}_k\| = \Theta(1)$.

4 HIGH-DIMENSIONAL ASYMPTOTIC PERFORMANCE UNDER LFMM

In this section, we present a self-consistent system of equations that gives access to the high-dimensional training and generalization performances of the ERM classifier in (4), under the LFMM in Definition 1. The characterization of high-dimensional asymptotic performance via a system of equations is reminiscent of previous analyses under GMM (Mai & Liao, 2019; Mignacco et al., 2020; Pesce et al., 2023), but our equations are different due to the collapse of the conditional one-dimensional CLT in (3) required for applying the CGEP.

Before presenting our system of equations, let us introduce first some mathematical objects involved in these equations. With the proximal operator $\text{prox}_{\tau, f}(t) = \arg \min_{a \in \mathbb{R}} [f(a) + \frac{1}{2\tau}(a - t)^2]$ for $\tau > 0$ and convex $f: \mathbb{R} \rightarrow \mathbb{R}$, we define the mapping

$$h_\kappa(t, y) = (\text{prox}_{\kappa, \ell(\cdot, y)}(t) - t)/\kappa, \quad (8)$$

for some constant $\kappa > 0$. Let $r \in \mathbb{R}$ be a random variable of form

$$r = ym + \sigma\tilde{e} + \sum_{k=1}^q \psi_k e_k, \quad (9)$$

for constants $m, \sigma, \psi_1, \dots, \psi_q$, with label y and e_1, \dots, e_q the corresponding noise variables in the informative factors z_1, \dots, z_q of the LFMM in Definition 1, as well as $\tilde{e} \sim \mathcal{N}(0, 1)$ independent of y, z_1, \dots, z_q . Remark that the distribution of r is parameterized by $m, \sigma^2, \psi_1, \dots, \psi_q$.

Self-consistent system of equations. Our system of equations is on the $q+3$ deterministic constants $\theta, \eta, \gamma, \omega_1, \dots, \omega_q$ that fully characterize the asymptotic performance of ERM classifier trained on high-dimensional LFMM 3:

$$\begin{aligned} \theta &= -\mathbb{E} \left[\frac{\partial h_\kappa(r, y)}{\partial r} \right], \quad \eta = \mathbb{E}[y h_\kappa(r, y)], \quad \gamma = \sqrt{\mathbb{E}[h_\kappa^2(r, y)]}, \\ \omega_k &= \mathbb{E}[h_\kappa(r, y) e_k] + \theta \cdot \mathbf{v}_k^\top \mathbf{Q} \boldsymbol{\xi}, \quad \forall k \in \{1, \dots, q\}, \end{aligned} \quad (10)$$

where

$$\boldsymbol{\xi} = \eta \boldsymbol{\mu} + \sum_{k=1}^q \omega_k \mathbf{v}_k, \quad \mathbf{Q} = (\lambda \mathbf{I}_p + \theta \boldsymbol{\Sigma})^{-1}, \quad (11)$$

the mapping $h_\kappa(r, y)$ is as defined in (8) for

$$\kappa = \frac{1}{n} \text{tr} \boldsymbol{\Sigma} \mathbf{Q}, \quad (12)$$

and the random variable r as defined in (9) with

$$m = \boldsymbol{\mu}^\top \mathbf{Q} \boldsymbol{\xi}, \quad \sigma^2 = \frac{\gamma^2}{n} \text{tr} (\mathbf{Q} \boldsymbol{\Sigma})^2, \quad \psi_k = \mathbf{v}_k^\top \mathbf{Q} \boldsymbol{\xi}, \quad \forall k \in \{1, \dots, q\}. \quad (13)$$

We are now ready to present our Theorem 1 on the asymptotic distributions of in-sample and out-of-sample predicted scores. The proof of Theorem 1 is provided in Appendix A.1

Theorem 1 (Asymptotic distribution of predicted scores). *Let Assumptions 1, 2, and 3 hold, for $\hat{\beta}$ solution to the ERM problem in (4) on a training set $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ of size n drawn i.i.d. $(\mathbf{x}_i, y_i) \sim \mathcal{D}_{(\mathbf{x}, y)}$ from the LFMM in Definition 1 we have that, for any bounded Lipschitz function $f: \mathbb{R} \rightarrow \mathbb{R}$,*

$$\mathbb{E}[f(\hat{\beta}^\top \boldsymbol{\nu})] - \mathbb{E}[f(\tilde{\beta}^\top \boldsymbol{\nu})] \rightarrow 0, \quad (14)$$

for any deterministic feature vector $\boldsymbol{\nu} \in \mathbb{R}^p$, and

$$\mathbb{E}[f(\hat{\beta}^\top \mathbf{x}_i)] - \mathbb{E}[f(\text{prox}_{\kappa, \ell(\cdot, y_i)}(\tilde{\beta}^\top \mathbf{x}_i))] \rightarrow 0, \quad \forall i \in \{1, \dots, n\}, \quad (15)$$

where

$$\tilde{\beta} = (\lambda \mathbf{I}_p + \theta \boldsymbol{\Sigma})^{-1} \left(\eta \boldsymbol{\mu} + \sum_{k=1}^q \omega_k \mathbf{v}_k + \gamma \boldsymbol{\Sigma}^{\frac{1}{2}} \mathbf{u} \right), \quad (16)$$

for Gaussian vector $\mathbf{u} \sim \mathcal{N}(\mathbf{0}_p, \mathbf{I}_p/n)$ independent of $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ and constants $\theta, \eta, \gamma, \omega_1, \dots, \omega_q$ determined by the self-consistent system of equations in (10), with κ given in (12).

According to (14) in Theorem 1, for a fresh test sample (\mathbf{x}', y') (which might be drawn from a distribution different from $\mathcal{D}_{(\mathbf{x}, y)}$ of training samples, as in the case of transfer learning), the out-of-sample predicted scores $\hat{\beta}^\top \mathbf{x}', \hat{\beta}^\top \mathbf{x}'$ produced by the ERM classifier $\hat{\beta}$ and its high-dimensional “equivalent” $\tilde{\beta}$ given in (16) have asymptotically the same distribution in the sense of (14). Furthermore, (15) tells us that the in-sample predicted score $\hat{\beta}^\top \mathbf{x}_i$ of $\hat{\beta}$ on a training sample (\mathbf{x}_i, y_i) follows asymptotically the same distribution as $\text{prox}_{\kappa, \ell(\cdot, y_i)}(\tilde{\beta}^\top \mathbf{x}_i)$. Since the distribution of $\tilde{\beta}$ is given in (16), we obtain directly from Theorem 1 the asymptotic training and generalization errors of the ERM classifier $\hat{\beta}$.

Furthermore, it follows from LLN and CLT that $(\tilde{\beta}^\top \mathbf{x}, y)$ with $(\mathbf{x}, y) \sim \mathcal{D}_{(\mathbf{x}, y)}$ independent of $\tilde{\beta}$ converges in distribution to (r, y) with r as defined in (9) with $m, \sigma^2, \psi_1, \dots, \psi_q$ given in (13). We thus obtain the following corollary on the asymptotic classification accuracy of $\tilde{\beta}$ on any training sample (\mathbf{x}_i, y_i) and test sample (\mathbf{x}', y') drawn from the same distribution $\mathcal{D}_{(\mathbf{x}, y)}$. The proof of Corollary 1 is deferred to Appendix A.2.1

Corollary 1 (Asymptotic generalization and training performances). *Under the conditions and notations of Theorem 1 we have that, for any bounded Lipschitz function $f: \mathbb{R} \rightarrow \mathbb{R}$,*

$$\mathbb{E}[f(\tilde{\beta}^\top \mathbf{x})|y] - \mathbb{E}[f(r)|y] \rightarrow 0, \quad (17)$$

³According to Assumption 1, $\frac{\partial h(r, y)}{\partial r}$ exists except on a finite set of points. On those points, we use the left derivative of $h(r, y)$ with respect r , i.e., $\lim_{t \rightarrow r-} (h(t, y) - h(r, y)) / (t - r)$.

for $(\mathbf{x}, y) \sim \mathcal{D}_{(\mathbf{x}, y)}$ independent of $\tilde{\beta}$, where r is as defined in (9) with $m, \sigma^2, \psi_1, \dots, \psi_q$ given in (13). Consequently, we have

$$\Pr(y' \tilde{\beta}^\top \mathbf{x}' > 0) - \Pr(yr > 0) \rightarrow 0, \quad (18)$$

for some test sample $(\mathbf{x}', y') \sim \mathcal{D}_{(\mathbf{x}, y)}$ independent of $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, and

$$\Pr(y_i \hat{\beta}^\top \mathbf{x}_i > 0) - \Pr(y \text{prox}_{\kappa, \ell(\cdot, y)}(r) > 0) \rightarrow 0, \quad \forall i \in \{1, \dots, n\}. \quad (19)$$

Remark 1 (On classifier bias under GMM and LFMM). Taking expectation on both sides of (16), we get $\mathbb{E}[\tilde{\beta}] = (\lambda \mathbf{I}_p + \theta \Sigma)^{-1} (\eta \boldsymbol{\mu} + \sum_{k=1}^q \omega_k \mathbf{v}_k)$. It follows from (10) and Stein’s lemma (Ingersoll 1987) that $\omega_1, \dots, \omega_q = 0$ in the case of Gaussian informative factors z_1, \dots, z_q . As GMM is a special case of LFMM with exclusively Gaussian (informative and noise) factors, we have that $\tilde{\beta}$ aligns, in expectation, with $(\lambda \mathbf{I}_p + \theta \Sigma)^{-1} \boldsymbol{\mu}$ under GMM. For non-Gaussian informative factors, we generally have non-zero $\omega_1, \dots, \omega_q$ that account for the non-Gaussian variation in data, making $\hat{\beta}$ more or less aligned with the directions $\mathbf{v}_1, \dots, \mathbf{v}_q$ of the informative factors.

5 CONDITIONS AND IMPLICATIONS OF GAUSSIAN UNIVERSALITY

In this section, we exploit our high-dimensional asymptotic analysis in Section 4 to derive the conditions of Gaussian universality under LFMM. To discuss the Gaussian universality in classification of mixture data, we introduce the notion of equivalent Gaussian mixture model (to a given LFMM), in a similar spirit to Dandi et al. (2024).

Definition 2 (Equivalent Gaussian mixture model). For an LFMM $\mathcal{D}_{(\mathbf{x}, y)}$ as in Definition 1 we define its equivalent Gaussian mixture model (GMM) $\mathcal{D}_{(\mathbf{g}, y)}$ as the GMM with the same class-conditional means and covariances as the LFMM $\mathcal{D}_{(\mathbf{x}, y)}$. Namely,

$$\mathbf{g} \sim \mathcal{N}(y\boldsymbol{\mu}, \Sigma), \quad (20)$$

for $\boldsymbol{\mu}, \Sigma$ given in (6) and (7) of Definition 1, respectively. We denote by $\hat{\beta}^g$ the ERM solution to (4) obtained on n i.i.d. equivalent GMM samples $(\mathbf{g}_1, y_1), \dots, (\mathbf{g}_n, y_n) \sim \mathcal{D}_{(\mathbf{g}, y)}$, and similarly its high-dimensional “equivalent” $\tilde{\beta}^g$ as in (16) of Theorem 1.

Notice importantly from Definition 1 that the equivalent GMM $\mathcal{D}_{(\mathbf{g}, y)}$ to an LFMM $\mathcal{D}_{(\mathbf{x}, y)}$ can be obtained by taking e_1, \dots, e_p of the LFMM $\mathcal{D}_{(\mathbf{x}, y)}$ to be standard Gaussian variables.

We define two types of Gaussian universality considered in this paper as follows.

Definition 3 (Gaussian universality under LFMM). For an ERM solution $\hat{\beta}$ obtained on a general LFMM $\mathcal{D}_{(\mathbf{x}, y)}$ in Definition 1 and an ERM solution $\hat{\beta}^g$ obtained on the equivalent GMM in Definition 2 we say that the Gaussian universality holds

- on **classifier** if $\hat{\beta}$ has asymptotically the same predictive ability as $\hat{\beta}^g$ on a given test set, as a consequence of their high-dimensional equivalents $\tilde{\beta}_{\ell, \lambda}, \tilde{\beta}_{\ell, \lambda}^g$ provided in Theorem 1 following the same distribution;
- on **in-distribution performance** if the respective training and generalization performances under $\mathcal{D}_{(\mathbf{x}, y)}$ are asymptotically the same as under $\mathcal{D}_{(\mathbf{g}, y)}$, that is

$$\Pr(y_i \mathbf{x}_i^\top \hat{\beta} > 0) - \Pr(y_i \mathbf{g}_i^\top \hat{\beta}^g > 0) \rightarrow 0, \quad (21)$$

and

$$\Pr(y' \mathbf{x}'^\top \hat{\beta} > 0) - \Pr(y' \mathbf{g}'^\top \hat{\beta}^g > 0) \rightarrow 0, \quad (22)$$

for $(\mathbf{x}', y') \sim \mathcal{D}_{(\mathbf{x}, y)}$ a test sample independent of $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, and $(\mathbf{g}', y') \sim \mathcal{D}_{(\mathbf{g}, y)}$ independent of $\{(\mathbf{g}_i, y_i)\}_{i=1}^n$.

In the following, we study first in Section 5.1 the Gaussian universality in the sense of in-distribution performance, and discuss our results with respect to the conditional one-directional CLT and the CGEP in Dandi et al. (2024). We then reveal in Section 5.2 the key role of square loss in inducing the Gaussian universality of classifier, and discuss its implication for the choice of loss function.

Throughout this section, our discussions are illustrated through numerical experiments on datasets of moderately large size, with n, p only in hundreds. A close match is consistently observed between the proposed asymptotic analysis and the empirical results.




Figure 1: Theoretical and empirical distribution of predicted scores $\hat{\beta}^T \mathbf{x}'$ for some fresh $(\mathbf{x}', y') \sim \mathcal{D}_{(\mathbf{x}, y)}$ independent of $\hat{\beta}$. The theoretical probability densities (red) are obtained from Theorem 1, and the empirical histograms (blue) are the values of $\hat{\beta}^T \mathbf{x}'$ over 10^6 independent copies of \mathbf{x}' , for three different LFMMs as in Definition 1 with $n = 600$, $p = 200$, $\rho = 0.5$, $s = [\sqrt{2}; \mathbf{0}_{p-1}]$ (so that $q = 1$), and Haar distributed \mathbf{V} . **Left:** normal e_1 and uniformly distributed e_2, \dots, e_p ; **Middle:** normal e_1, \dots, e_p ; **Right:** uniformly distributed e_1 , and normal e_2, \dots, e_p .

5.1 GAUSSIAN UNIVERSALITY OF IN-DISTRIBUTION PERFORMANCE

Notice from (18) in Corollary 1 that the in-distribution generalization performance of $\hat{\beta}$ under an LFMM $\mathcal{D}_{(\mathbf{x}, y)}$ is determined by the random variable r in (9), the distribution of which depends solely on (i) the distributions of y, e_1, \dots, e_q in the LFMM and (ii) the values of $m, \sigma^2, \psi_1, \dots, \psi_q$ given in (13). Remark also that the values of $m, \sigma^2, \psi_1, \dots, \psi_q$ in (13) are determined by the system of equations in (10), which concerns only the distributions of r, y, e_1, \dots, e_q , as well as the deterministic parameters $\mu, \Sigma, \mathbf{v}_1, \dots, \mathbf{v}_q$ of the LFMM.

We thus conclude that the distribution of r is *insensitive* to the distributions of noise factors z_{q+1}, \dots, z_p . In other words, an LFMM with Gaussian noises e_1, \dots, e_q in its informative factors z_1, \dots, z_q has the *same* asymptotic generalization performance as its equivalent GMM in Definition 2 *regardless of* the distributions of the noise factors z_{q+1}, \dots, z_p .

A similar conclusion can be drawn from (19) of Corollary 1 on the asymptotic in-distribution training performance, by studying also the distribution of r but through a proximal mapping $\text{prox}_{\kappa, \ell(\cdot, y)}$. We formalize these conclusions on the universality of in-distribution performance in Corollary 2, the proof of which is given in Appendix A.2.2.

Corollary 2 (Condition of Gaussian universality on in-distribution performance). *Under the settings and notations of Theorem 1 and Definition 2 the Gaussian universality of in-distribution performance in Definition 3 holds if and only if noises e_1, \dots, e_q of LFMM informative factors in (5) are Gaussian.*

Figure 1 provides numerical illustrations of Corollary 2, where we compare the empirical histograms and the asymptotic distributions of the out-of-sample predicted scores $\hat{\beta}^T \mathbf{x}'$ for data drawn from three different LFMMs: an LFMM satisfying the in-distribution performance universality condition in Corollary 2 (left), an LFMM sharing the same parameters (μ, Σ, ρ) with the first but violating the condition in Corollary 2 (right), and their equivalent GMM in the sense of Definition 2 (middle).

Remark 2 (Connection to conditional one-directional CLT in (3)). Our universality results on the in-distribution performance in Corollary 2 are related to the CGEP proven by Dandi et al. (2024) under the presumed validity of a conditional one-directional CLT in (3). Under our notations, the conditional one-directional CLT in (3) translates to the convergence of $y' \hat{\beta}^T \mathbf{x}'$ and $y' \hat{\beta}_g^T \mathbf{x}'$ to the same normal distribution, i.e.,

$$\frac{y' \mathbf{x}'^T \hat{\beta} - m}{\sqrt{\sigma^2 + \sum_{k=1}^q \psi_k^2}} \xrightarrow{d} \mathcal{N}(0, 1), \quad \frac{y' \mathbf{g}'^T \hat{\beta}^g - m}{\sqrt{\sigma^2 + \sum_{k=1}^q \psi_k^2}} \xrightarrow{d} \mathcal{N}(0, 1), \quad (23)$$

as it can be shown from (14), (17) that $y' \hat{\beta}^T \mathbf{x}'$ has asymptotically the same distribution as yr , which is of mean m and variance $\sigma^2 + \sum_{k=1}^q \psi_k^2$. It is easy to see from (9) that yr is normally distributed if and only if e_1, \dots, e_q are Gaussian, which is exactly the condition of universality stated in Corollary 2.




Figure 2: Empirical and theoretical results under an LFMM with $p = 200$, $\rho = 0.5$, $s = [\sqrt{2}; \mathbf{0}_{p-1}]$, Rademacher e_1 , normal e_2, \dots, e_p , and Haar distributed \mathbf{V} . **Top:** scatter plot of 200 independent $[r, h_\kappa(r, \pm 1)]$. **Bottom:** histograms of predicted scores on 10^6 fresh samples $(\mathbf{x}', y') \sim \mathcal{D}_{(\mathbf{x}, y)}$ given by $\hat{\beta}$ and $\hat{\beta}^g$, versus theoretical densities obtained from Theorem 1. **Left:** $n = 100$, square loss $\ell(\hat{y}, y) = (\hat{y} - y)^2/2$. **Right:** $n = 600$, square hinge loss $\ell(\hat{y}, y) = \max\{0, (1 - \hat{y}y)\}^2$.

5.2 GAUSSIAN UNIVERSALITY OF CLASSIFIER AND IMPLICATION FOR CHOICE OF LOSS

As discussed in Section 5.1 the system of equations in (10) does not depend on the distributions of noise factors z_{q+1}, \dots, z_p . As the distribution of the high-dimensional equivalent $\hat{\beta}$ to $\hat{\beta}^g$ given in (16) is controlled by the constants $\theta, \eta, \gamma, \omega_1, \dots, \omega_q$ that are determined by (10), it is therefore also *universal* over the distributions of z_{q+1}, \dots, z_p . Then, by a similar reasoning to Corollary 2, we conclude that the Gaussian universality of classifier in Definition 3 holds in the case of normally distributed e_1, \dots, e_q .

This is however *not* the only case of Gaussian universality on classifier. Note from (9) and (10) that, even though the system of equations in (10) *does* depend on the distributions of e_1, \dots, e_q , it only involves their means and variances if $h_\kappa(r, y)$ is linear in r . Remark also from (8) that $h_\kappa(r, y)$ varies linearly with r if and only if the square loss $\ell(\hat{y}, y) = (\hat{y} - y)^2/2$ (or its rescaled version) is used.

These two conditions for the Gaussian universality of classifier as understood in Definition 3 are made precise in the following result, proven in Appendix A.2.3.

Corollary 3 (Condition of Gaussian universality on classifier). *Under the settings and notations of Theorem 1 and Definition 2, the Gaussian universality of classifier in Definition 3 holds if and only if one of the following two conditions is met: (i) e_1, \dots, e_q in (5) are normally distributed; (ii) $\partial \ell(\hat{y}, y) / \partial \hat{y}$ is a linear function of \hat{y} , e.g., $\ell(\hat{y}, y) = (\hat{y} - y)^2/2$.*

Remark 3 (Limitation of square loss). As an important consequence of Corollary 3, any classifier $\hat{\beta}$ trained using the square loss on generic LFMM samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n \sim \mathcal{D}_{(\mathbf{x}, y)}$ and $\hat{\beta}^g$ trained on equivalent GMM samples $\{(\mathbf{g}_i, y_i)\}_{i=1}^n \sim \mathcal{D}_{(\mathbf{g}, y)}$ have asymptotically the same probability of correctly classifying a fresh LFMM test sample $(\mathbf{x}', y') \sim \mathcal{D}_{(\mathbf{x}, y)}$. That is, ERM classifiers trained with square loss are *unable* to adapt to non-Gaussian informative factors of LFMM, contrarily to other (non-square) losses.

The particular effect of square loss discussed in Remark 3 is numerically demonstrated in Figure 2. On the left hand side, the square loss $\ell_{\text{sq}}(\hat{y}, y) = (\hat{y} - y)^2/2$ is used, and $h_\kappa(r, y)$ varies linearly with r as in the top left plot (the two elongated scatter plots are associated respectively with $y = \pm 1$), so that the distribution of $\mathbf{x}'^T \hat{\beta}_{\text{sq}, \lambda}$ and $\mathbf{x}'^T \hat{\beta}_{\text{sq}, \lambda}^g$ are *indistinguishable* in the bottom left plot of Figure 2. On the right hand, the square hinge loss $\ell_{\text{shg}}(\hat{y}, y) = \max\{0, 1 - \hat{y}y\}^2$ is used, and we observe drastically different behaviors for $\mathbf{x}'^T \hat{\beta}_{\text{shg}, \lambda}$ and $\mathbf{x}'^T \hat{\beta}_{\text{shg}, \lambda}^g$ in the right column of Figure 2 when the points $[r, h_\kappa(r, \pm 1)]$ are highly nonlinear.

Remark 3, supported by the numerically results in Figure 2, points to the insensitivity of least-square classifiers to the distributions of non-Gaussian informative factors, despite their non-universal impact




Figure 3: Empirical classification accuracy of $\hat{\beta}_{\ell, \lambda}$ computed over 10^5 independent copies of $(\mathbf{x}', y') \sim \mathcal{D}_{(\mathbf{x}, y)}$ and averaged over 100 trials with a width of ± 1 standard deviation, versus theoretical performance given in Theorem 1, given by the square loss $\ell(\hat{y}, y) = (y - \hat{y})^2/2$ and the logistic loss $\ell(\hat{y}, y) = -\ln(1/(1 + e^{-y\hat{y}}))$ and on $n = 800$ training samples. **Left:** GMM under Definition 1 with $p = 200$, $\rho = 0.5$, $s = [1, 5; 0.5; \mathbf{0}_{p-2}]$ (so that $q = 2$), and $\mathbf{V} = \text{diag}(2, \mathbf{1}_{p-1})\mathbf{H}$ with Haar distributed \mathbf{H} . **Right:** LFMM identical to the GMM in the left, but with Rademacher e_1 .

on the in-distribution performance as discussed in Section 5.1. The incapacity of square loss to account for non-Gaussian variations in the informative factors sheds light on the suboptimality of square loss observed in the right display of Figure 3, where the logistic loss yields better performance than the square loss with *optimally chosen* regularization λ on LFMM having non-Gaussian informative factors, while the logistic loss *fails* to do better than the square loss under the equivalent GMM in the left-hand figure. Further experiments on real-world datasets are given in Appendix B.

This finding on the suboptimality of square loss under LFMM provide new insights on the impact of loss function beyond previous optimality results of square loss under GMM. For high-dimensional GMM data, the square loss has been proven optimal, see (Taheri et al., 2021b) for the case of unregularized ERM, and (Mai & Liao, 2019) for ridge-regularized ERM, in the $n, p \rightarrow \infty$ limit. That is, the square loss not only gives the best unbiased classifier, but also allows for an *optimal* bias-variance trade-off with well calibrated ridge-regularization. As a result of the Gaussian universality breakdown discussed above, the optimality of square loss is no longer valid under the more general LFMM. This motivates a few open questions on the optimal loss:

- Is the square loss optimal *only* under GMM, or when the Gaussian universality of in-distribution performance in Definition 3 holds?
- In the case of Gaussian universality breakdown, does the optimal loss depend on the sample ratio n/p as in the setting of linear regression in (Bean et al., 2013)?
- Is it possible, in the large n, p regime, to propose an optimal design of classification loss adapted to the data distribution and sample size?

6 CONCLUDING REMARKS

Our analysis considered a basic framework of linear factor mixture models (LFMM) and showed that the Gaussian universality can already break down under this natural extension of GMM. Based on the precise performance characterization, we derived conditions of Gaussian universality to shed light on the limit of the widely observed and extensively studied Gaussian universality phenomenon.

Breaking the Gaussian universality in classification of mixture models allows also deeper insight into the choice of classification loss beyond the optimality of square loss under GMM (Taheri et al., 2021b; Mai & Liao, 2019). The suboptimality of square loss under LFMM can be further investigated in future works, to propose, for instance, an optimal design of loss function that takes into account the data distribution and the sample size, as done by (Bean et al., 2013) for linear regression.

Several simplifications made in our analysis can be removed more or less easily. For instance, while the extension to multi-classification is fairly straightforward, the generalization to non-smooth losses is less direct: even though our system of equations in (10) does not require access to the derivatives of the loss function, they are involved in the establishment of these equations.

ACKNOWLEDGMENTS

Z. Liao would like to acknowledge the National Natural Science Foundation of China (via fund NSFC-62206101 and NSFC-12141107) and the Guangdong Provincial Key Laboratory of Mathematical Foundations for Artificial Intelligence (2023B1212010001) for providing partial support. This work is also supported by the AI Cluster ANITI (ANR-19-PI3A-0004).

REFERENCES

- Urte Adomaityte, Gabriele Sicuro, and Pierpaolo Vivo. Classification of heavy-tailed features in high dimensions: a superstatistical approach. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zhidong D Bai and Jack W Silverstein. Clt for linear spectral statistics of large-dimensional sample covariance matrices. In *Advances In Statistics*, pp. 281–333. World Scientific, 2008.
- Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, 2019.
- Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- Derek Bean, Peter J Bickel, Noureddine El Karoui, and Bin Yu. Optimal m-estimation in high-dimensional regression. *Proceedings of the National Academy of Sciences*, 110(36):14563–14568, 2013.
- Michael Celentano and Andrea Montanari. Fundamental barriers to high-dimensional regression with convex penalties. *The Annals of Statistics*, 50(1):170–196, 2022.
- Michael Celentano, Andrea Montanari, and Yuting Wei. The lasso with general gaussian designs with applications to hypothesis testing. *The Annals of Statistics*, 51(5):2194–2220, 2023.
- Romain Couillet and Zhenyu Liao. *Random matrix methods for machine learning*. Cambridge University Press, 2022.
- Yatin Dandi, Ludovic Stephan, Florent Krzakala, Bruno Loureiro, and Lenka Zdeborová. Universality laws for gaussian mixtures in generalized linear models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zeyu Deng, Abba Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *Information and Inference: A Journal of the IMA*, 11(2): 435–495, 2022.
- Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018. ISSN 0090-5364. doi: 10.1214/17-aos1549.
- David Donoho and Andrea Montanari. High dimensional robust m-estimation: Asymptotic variance via approximate message passing. *Probability Theory and Related Fields*, 166:935–969, 2016.
- Noureddine El Karoui. Concentration of measure and spectra of random matrices: Applications to correlation matrices, elliptical distributions and beyond. 2009.
- 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*, 170:95–175, 2018.
- Noureddine El Karoui, Derek Bean, Peter J Bickel, Chinghway Lim, and Bin Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 110(36): 14557–14562, 2013.

- Federica Gerace, Bruno Loureiro, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Generalisation error in learning with random features and the hidden manifold model. In *International Conference on Machine Learning*, pp. 3452–3462. PMLR, 2020.
- Federica Gerace, Florent Krzakala, Bruno Loureiro, Ludovic Stephan, and Lenka Zdeborová. Gaussian universality of perceptrons with random labels. *Physical Review E*, 109(3):034305, 2024.
- Sebastian Goldt, Bruno Loureiro, Galen Reeves, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. The gaussian equivalence of generative models for learning with shallow neural networks. In *Mathematical and Scientific Machine Learning*, pp. 426–471. PMLR, 2022.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- Qiyang Han and Yandi Shen. Universality of regularized regression estimators in high dimensions. *The Annals of Statistics*, 51(4):1799–1823, 2023.
- Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *Annals of statistics*, 50(2):949, 2022.
- Hong Hu and Yue M Lu. Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, 69(3):1932–1964, 2022.
- Hanwen Huang. Asymptotic behavior of support vector machine for spiked population model. *Journal of Machine Learning Research*, 18(45):1–21, 2017.
- Hanwen Huang and Qinglong Yang. Large dimensional analysis of general margin based classification methods. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(11):113401, 2021.
- Jonathan E Ingersoll. *Theory of financial decision making*, volume 3. Rowman & Littlefield, 1987.
- Adel Javanmard and Mahdi Soltanolkotabi. Precise statistical analysis of classification accuracies for adversarial training. *The Annals of Statistics*, 50(4):2127–2156, 2022.
- Abla Kammoun and Mohamed-Slim Alouini. On the precise error analysis of support vector machines. *IEEE Open Journal of Signal Processing*, 2:99–118, 2021.
- Zhenyu Liao and Romain Couillet. A large dimensional analysis of least squares support vector machines. *IEEE Transactions on Signal Processing*, 67(4):1065–1074, 2019.
- Marco Loog, Tom Viering, Alexander Mey, Jesse H Krijthe, and David MJ Tax. A brief prehistory of double descent. *Proceedings of the National Academy of Sciences*, 117(20):10625–10626, 2020.
- Bruno Loureiro, Gabriele Sicuro, Cédric Gerbelot, Alessandro Pocco, Florent Krzakala, and Lenka Zdeborová. Learning gaussian mixtures with generalized linear models: Precise asymptotics in high-dimensions. *Advances in Neural Information Processing Systems*, 34:10144–10157, 2021.
- Bruno Loureiro, Cédric Gerbelot, Maria Refinetti, Gabriele Sicuro, and Florent Krzakala. Fluctuations, bias, variance & ensemble of learners: Exact asymptotics for convex losses in high-dimension. In *International Conference on Machine Learning*, pp. 14283–14314. PMLR, 2022.
- Xiaoyi Mai and Romain Couillet. Consistent semi-supervised graph regularization for high dimensional data. *Journal of Machine Learning Research*, 22(94):1–48, 2021.
- Xiaoyi Mai and Zhenyu Liao. High dimensional classification via regularized and unregularized empirical risk minimization: Precise error and optimal loss. *arXiv preprint arXiv:1905.13742*, 2019.
- 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)*, pp. 3357–3361. IEEE, 2019.
- Antoine Maillard, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Phase retrieval in high dimensions: Statistical and computational phase transitions. *Advances in Neural Information Processing Systems*, 33:11071–11082, 2020.

- Vladimir Alexandrovich Marchenko and Leonid Andreevich Pastur. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik*, 114(4):507–536, 1967.
- Francesca Mignacco, Florent Krzakala, Yue Lu, Pierfrancesco Urbani, and Lenka Zdeborova. The role of regularization in classification of high-dimensional noisy gaussian mixture. In *International conference on machine learning*, pp. 6874–6883. PMLR, 2020.
- Andrea Montanari and Phan-Minh Nguyen. Universality of the elastic net error. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 2338–2342, June 2017. doi: 10.1109/ISIT.2017.8006947.
- Andrea Montanari and Basil N Saeed. Universality of empirical risk minimization. In *Conference on Learning Theory*, pp. 4310–4312. PMLR, 2022.
- Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12):124003, 2021.
- Ashkan Panahi and Babak Hassibi. A Universal Analysis of Large-Scale Regularized Least Squares Solutions. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Leonid Andreevich Pastur and Mariya Shcherbina. *Eigenvalue distribution of large random matrices*. Number 171. American Mathematical Soc., 2011.
- Luca Pesce, Florent Krzakala, Bruno Loureiro, and Ludovic Stephan. Are gaussian data all you need? the extents and limits of universality in high-dimensional generalized linear estimation. In *International Conference on Machine Learning*, pp. 27680–27708. PMLR, 2023.
- Maria Refinetti, Sebastian Goldt, Florent Krzakala, and Lenka Zdeborová. Classifying high-dimensional gaussian mixtures: Where kernel methods fail and neural networks succeed. In *International Conference on Machine Learning*, pp. 8936–8947. PMLR, 2021.
- Fariborz Salehi, Ehsan Abbasi, and Babak Hassibi. The impact of regularization on high-dimensional logistic regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, 2018. ISBN 978-0-262-25693-3. doi: 10.7551/mitpress/4175.001.0001.
- Dominik Schröder, Hugo Cui, Daniil Dmitriev, and Bruno Loureiro. Deterministic equivalent and error universality of deep random features learning. In *International Conference on Machine Learning*, pp. 30285–30320. PMLR, 2023.
- 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.
- Hossein Taheri, Ramtin Pedarsani, and Christos Thrampoulidis. Fundamental limits of ridge-regularized empirical risk minimization in high dimensions. In *International Conference on Artificial Intelligence and Statistics*, pp. 2773–2781. PMLR, 2021a.
- Hossein Taheri, Ramtin Pedarsani, and Christos Thrampoulidis. Sharp guarantees and optimal performance for inference in binary and gaussian-mixture models. *Entropy*, 23(2):178, 2021b.
- Terence Tao, Van Vu, and Manjunath Krishnapur. Random matrices: Universality of esds and the circular law. 2010.
- 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.
- Christos Thrampoulidis, Samet Oymak, and Mahdi Soltanolkotabi. Theoretical insights into multi-class classification: A high-dimensional asymptotic view. *Advances in Neural Information Processing Systems*, 33:8907–8920, 2020.

Ke Wang and Christos Thrampoulidis. Benign overfitting in binary classification of gaussian mixtures. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4030–4034. IEEE, 2021.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.