
Are Gaussian Data All You Need? The Extents and Limits of Universality in High-Dimensional Generalized Linear Estimation

Luca Pesce¹ Florent Krzakala¹ Bruno Loureiro² Ludovic Stephan¹

Abstract

In this manuscript we consider the problem of generalized linear estimation on Gaussian mixture data with labels given by a single-index model. Our first result is a sharp asymptotic expression for the test and training errors in the high-dimensional regime. Motivated by the recent stream of results on the Gaussian universality of the test and training errors in generalized linear estimation, we ask ourselves the question: “*when is a single Gaussian enough to characterize the error?*”. Our formulas allow us to give sharp answers to this question, both in the positive and negative directions. More precisely, we show that the sufficient conditions for Gaussian universality (or lack thereof) crucially depend on the alignment between the target weights and the means and covariances of the mixture clusters, which we precisely quantify. In the particular case of least-squares interpolation, we prove a strong universality property of the training error and show it follows a simple, closed-form expression. Finally, we apply our results to real datasets, clarifying some recent discussions in the literature about Gaussian universality of the errors in this context.

1. Introduction

It is commonsense in machine learning that structure in the data is an important ingredient for successful learning. Quantifying this statement and in particular, how structure in the features impact the training and generalization errors the most, is an important endeavor in the broad program of “seeing through” the modern machine learning black box.

¹Ecole Polytechnique Fédérale de Lausanne, Information, Learning and Physics lab. CH-1015 Lausanne, Switzerland.
²Département d’Informatique, École Normale Supérieure - PSL & CNRS, 45 rue d’Ulm, F-75230 Paris cedex 05, France. Correspondence to: Luca Pesce <luca.pesce@epfl.ch>.

Proceedings of the 40th International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

On the theoretical side, there have been some important recent progress in this direction in the context of generalized linear estimation. For instance, a recent line of work on linear regression trained on Gaussian data has shown that good generalization can arise even in the “overparametrized regime” where the training error is exactly zero (Bartlett et al., 2020; Hastie et al., 2022; Wu & Xu, 2020). This *benign overfitting* property crucially depends on the covariance structure, occurring when the signal components of the target align with a lower-dimensional of the data, leaving space for the noise to spread along the higher-dimensional orthogonal subspace (Bartlett et al., 2020). Analogous conclusions hold, under similar conditions, to generalized linear tasks (Muthukumar et al., 2021; Wang & Thrampoulidis, 2021; Shamir, 2022). Indeed, this is only one example of many surprising insights learned from the study of generalized linear models on Gaussian data over the past few years (Gardner & Derrida, 1988; Krogh & Hertz, 1992; Donoho & Tanner, 2009; Candès et al., 2020). Despite the seemingly constraining assumption on the distribution of the features, a recent line of work provides strong evidence for the *Gaussian universality* of the training and generalization errors in generalized estimation in different settings. These includes rigorous results for non-Gaussian designs (Montanari & Nguyen, 2017; Panahi & Hassibi, 2017), random feature maps (Mei & Montanari, 2019; Gerace et al., 2020; Goldt et al., 2022; Hu & Lu, 2022), neural tangent features (Montanari & Saeed, 2022), Gaussian mixtures with random labels (Gerace et al., 2023), as well as extensive numerical evidence for other feature maps (Goldt et al., 2022; Loureiro et al., 2021a) and even real datasets (Loureiro et al., 2021a; Jacot et al., 2020; Bordelon et al., 2020). These works beg the question “*When are Gaussian features a good model for learning?*”. Our aim is to give precise answers to this question in the context of generalized linear estimation on a popular model for multi-modal data, known to be able to approximate any distribution (Stone, 1948): the Gaussian mixture model.

1.1. Main Results

Our **main contributions** in this work are as follows:

- **Exact asymptotics of GLMs:** We provide the exact

asymptotic limit for the training and test errors of a generalized linear model with convex loss in high dimensions, when the data is drawn from a Gaussian mixture model with a single-index target. These asymptotics are based on the so-called *replica method* from statistical physics, and follow the same line as (Loureiro et al., 2021b), which considered instead the task of learning the mixture labels.

- **Universality of training and test errors:** We provide a set of sufficient conditions on the target weights θ_0 such that both the asymptotic training and test errors for a Gaussian mixture model are independent from the cluster means. In particular, these conditions are satisfied by a target whose direction is uniform on \mathbb{S}^{d-1} . In the case of ridge regression, we show an even stronger result: namely, the training loss is also independent from the cluster covariances, and reduces to that of a single Gaussian with identity covariance.

- **The importance of Homoscedasticity:** In the particular case of a homoscedastic Gaussian mixture (a mixture of Gaussians that share the same covariance matrix), we further demonstrate universality results that can actually be observed on real data after a random feature map (see e.g. Fig. 1 and 2). We also unveil the universal behavior of the linear separability transition, a phenomenon studied in detail for pure Gaussian data in (Candès et al., 2020) and that appears to be universal for a homoscedastic mixture.

- **Breaking universality:** In contrast to the results of the previous paragraph, we show that there are two ways to break Gaussian universality. First, strong heteroscedasticity can break the universal behavior. Second, in the homoscedastic case, we show that the correlation between the data and the task matters: even a small correlation between the target weights and the cluster means suffice to break universality, in the sense that the asymptotic errors of a model trained in a Gaussian mixture differ from the one of a model trained on Gaussian data. Rather than the structure of the data itself, what appears to matter is thus the correlation between this structure and the task to be learned.

1.2. Related Works

Exact asymptotics: An appealing feature of Gaussian data is that the asymptotic performance of different models can be sharply characterized in the proportional high-dimensional limit where $n, d \rightarrow \infty$ at fixed *sample complexity* $\alpha := n/d$. This is particularly the case for ridge regression, because of the close connection to a random matrix theory problem; see e.g. (Dobriban & Wager, 2018; Hastie et al., 2022; Wu & Xu, 2020). Those results are related to the concept of *deterministic equivalents* around which the regression matrices concentrate (Couillet & Liao, 2022).

Beyond the quadratic case, there exist many asymptotic rig-

orous studies: (Lei et al., 2018; Thrampoulidis et al., 2018) studied m-estimators, (Celentano et al., 2020) the performance of the LASSO estimator, (Liang & Sur, 2022) minimum ℓ^1 -norm interpolators, while (Loureiro et al., 2021a) provided a general result for convex losses and penalties with arbitrary covariances.

In the case of Gaussian mixtures, most of the effort has been geared towards classification, i.e. recovering the cluster label instead of teacher-generated ones. For binary classification, examples include (Mai et al., 2019; Mignacco et al., 2020; Deng et al., 2022), the latter of which also shows an equivalence between classification and a single-index model. In the multi-class setting, (Thrampoulidis et al., 2020) studied the performance of ridge regression classifiers; the most general result in this line is (Loureiro et al., 2021b), which considers any convex (not necessarily separable) loss. On the other hand, the exact asymptotics for teacher-generated binary labels has been studied in (Mannelli et al., 2023) for 2-classes GMMs.

Gaussian universality: The works presented in the previous paragraphs share one common point: the techniques used to show those results, either the Gordon min-max identity (Gordon, 1985) or the rigorous study of an Approximate Message Passing (AMP) sequence (Donoho et al., 2009), are only valid for Gaussian (or Gaussian mixture) distributions. Remarkably, however, those results are also able to capture the errors of particular classes of non-Gaussian features. This *Gaussian universality property* (GEP) (Goldt et al., 2020) was proven to hold for generalized linear estimation with random features (Karoui, 2008; Mei & Montanari, 2019; Gerace et al., 2020; Goldt et al., 2022; Hu & Lu, 2022), neural tangent features (Montanari & Saeed, 2022) and kernel features (Bordelon et al., 2020; Mei et al., 2022; Cui et al., 2021; 2022). (Ba et al., 2022) showed that Gaussian universality is preserved on the random features model when the weights are trained at order one steps, but break if an extensive number of steps are taken. Beyond the realm of theorems, (Loureiro et al., 2021a) provided numerical evidence of Gaussian universality for a broader class of realistic features from trained neural networks. On a close line, (Gerace et al., 2023) studied the Gaussian universality for pure random labels, while (Bordelon et al., 2020; Spigler et al., 2020; Jacot et al., 2020; Loureiro et al., 2021a; Cui et al., 2021; Wei et al., 2022) has numerically shown that for ridge regression in particular, the Gaussian formula captured the learning curves of some simple real datasets.

2. Setting & Motivation

Let $(x^\nu, y^\nu) \in \mathbb{R}^d \times \mathcal{Y}$ denote $\nu = 1, \dots, n$ pairs of independently sampled training points. We shall be interested in studying the properties of generalized linear estimation

$\hat{y}(x) = \hat{f}(\boldsymbol{\theta}^\top \mathbf{x})$ with weights $\boldsymbol{\theta} \in \mathbb{R}^d$ learned from the training data by minimizing the following empirical risk:

$$\hat{\mathcal{R}}_n^\lambda(\boldsymbol{\theta}) = \frac{1}{n} \sum_{\nu=1}^n \ell \left(y^\nu, \frac{\boldsymbol{\theta}^\top \mathbf{x}^\nu}{\sqrt{d}} \right) + \lambda r(\boldsymbol{\theta}) \quad (1)$$

where $\ell : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}_+$ is a convex loss function and $r : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex penalty. For example, this includes the particular case of ridge regression where $\ell(y, \hat{y}) = (y - \hat{y})^2$ and $r(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_2^2$. The key quantities of interest in the following will be the training and generalization error, defined as:

$$\varepsilon_{\text{tr}}(\hat{\boldsymbol{\theta}}) = \frac{1}{n} \sum_{\nu=1}^n \ell \left(y^\nu, \frac{\hat{\boldsymbol{\theta}}^\top \mathbf{x}^\nu}{\sqrt{d}} \right) \quad (2)$$

$$\varepsilon_{\text{gen}}(\hat{\boldsymbol{\theta}}) = \mathbb{E} \left[g \left(y_{\text{new}}, \hat{f} \left(\frac{\hat{\boldsymbol{\theta}}^\top \mathbf{x}_{\text{new}}}{\sqrt{d}} \right) \right) \right] \quad (3)$$

where $g : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}_+$ is a performance metric of the choice of the statistician, not necessarily equal to the loss function $\ell(y, \hat{y})$. For example, in the case of binary classification with $\mathcal{Y} = \{-1, +1\}$, we can take $\ell(y, \hat{y}) = \log(1 + e^{-y\hat{y}})$ to be the logistic loss, while taking $g(y, \hat{y}) = \mathbb{I}[y \neq \hat{y}]$ to be the classification error. Note that in eq. (3) the expectation is taken over a new data pair $(\mathbf{x}_{\text{new}}, y_{\text{new}})$ which we assume is independently drawn from the same distribution of the training data.

In particular, we will be interested in characterizing these errors under the assumption that the labels have been generated by the following target distribution:

$$y^\nu \sim P_0 \left(\cdot \mid \frac{\boldsymbol{\theta}_0^\top \mathbf{x}^\nu}{\sqrt{d}} \right) \quad (4)$$

for some fixed vector $\boldsymbol{\theta}_0 \in \mathbb{R}^d$ and distribution P_0 . A common choice for P_0 is $y^\nu = f_0 \left(\frac{\boldsymbol{\theta}_0^\top \mathbf{x}^\nu}{\sqrt{d}} + \xi \right)$ for additive Gaussian noise $\xi \sim \mathcal{N}(0, \Delta)$. This setting is sometimes referred to as a *teacher-student setting*. We will sometimes adopt this convenient terminology, referring to the target distribution P_0 as the *teacher* and $\boldsymbol{\theta}_0 \in \mathbb{R}^d$ as the *teacher weights*. Similarly, we will sometimes refer to the model \hat{y} as the *student* and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^d$ the *student weights*. As previously mentioned, we shall be considering both regression $\mathcal{Y} = \mathbb{R}$ and binary classification $\mathcal{Y} = \{-1, +1\}$.

This model has been the subject of a plethora of works in the high-dimensional statistics literature over the past few years, in particular under the *Gaussian design* assumption:

Model 2.1 (Gaussian covariate model). In the Gaussian covariate model (GCM), we assume the inputs are independently drawn from a Gaussian distribution:

$$\mathbf{x}^\nu \sim \mathcal{N} \left(\frac{\boldsymbol{\mu}}{\sqrt{d}}, \Sigma \right), \quad \nu = 1, \dots, n. \quad (5)$$

We denote $\mathcal{G}_{\boldsymbol{\theta}_0, \Sigma, \boldsymbol{\mu}}$ the teacher-student problem under this data assumption.

A key motivation for this work is the common intuition that data from standard classification tasks such as MNIST are closer to multi-modal distributions than to a single-mode Gaussian. Therefore, in this manuscript, we ask ourselves the question: “*when is Gaussian data all you need?*”. In particular, we focus our attention on a prototypical distribution to model multi-modal data (and a universal approximator of densities): the K cluster Gaussian mixture:

Model 2.2 (Gaussian mixture model). In the Gaussian mixture model (GMM), we assume the inputs are independently drawn from a mixture of K Gaussians:

$$\mathbf{x}^\nu \sim \sum_{c \in \mathcal{C}} p_c \mathcal{N} \left(\frac{\boldsymbol{\mu}_c}{\sqrt{d}}, \Sigma_c \right), \quad \nu = 1, \dots, n. \quad (6)$$

where $\mathcal{C} := \{1, \dots, K\}$ is the set of possible clusters, and $p_c \in [0, 1]$ is the probability of belonging to cluster $c \in \mathcal{C}$, whose means and covariance are given by $(\boldsymbol{\mu}_c, \Sigma_c)$. We assume the following scaling for the means and covariances of the clusters, as the dimension diverges:

$$\|\boldsymbol{\mu}_c\|_2^2 = O(d), \quad \text{tr} \Sigma_c = O(d), \quad \forall c \in \mathcal{C}. \quad (7)$$

We note that despite Model 2.1 being a special case $K = 1$ of Model 2.2 when the labels y^ν do not depend on the input cluster, it will be instructive to treat the $K = 1$ and $K > 1$ case as two different models.

3. Main Theoretical Results

In this section, we introduce our main theoretical results concerning universality of high-dimensional generalized linear estimation of GMMs. Our result builds on a long line of works providing an exact asymptotic characterization of empirical risk minimizers (2) on the proportional high-dimensional limit for Model 2.1 (Dobriban & Wager, 2018; Hastie et al., 2022; Thrampoulidis et al., 2018; Loureiro et al., 2021a). In particular, closer to our derivation are the rigorous results in (Loureiro et al., 2021b;a). We prove that the training and generalization errors concentrate in high-dimensions in a deterministic expression given by the solution of a set of self-consistent equations. Then we analyze Gaussian universality, provably characterizing a set of sufficient conditions the learning task must respect such that the test and training errors of Model 2.2 asymptotically agree with the ones from Model 2.1.

Since we deal with sequences of random variables, we will need a rigorous definition of convergence. For two sequences of numbers $(a_n), (b_n)$, we write

$$a_n \simeq b_n \quad \text{iff} \quad \lim_{n \rightarrow \infty} a_n - b_n = 0. \quad (8)$$

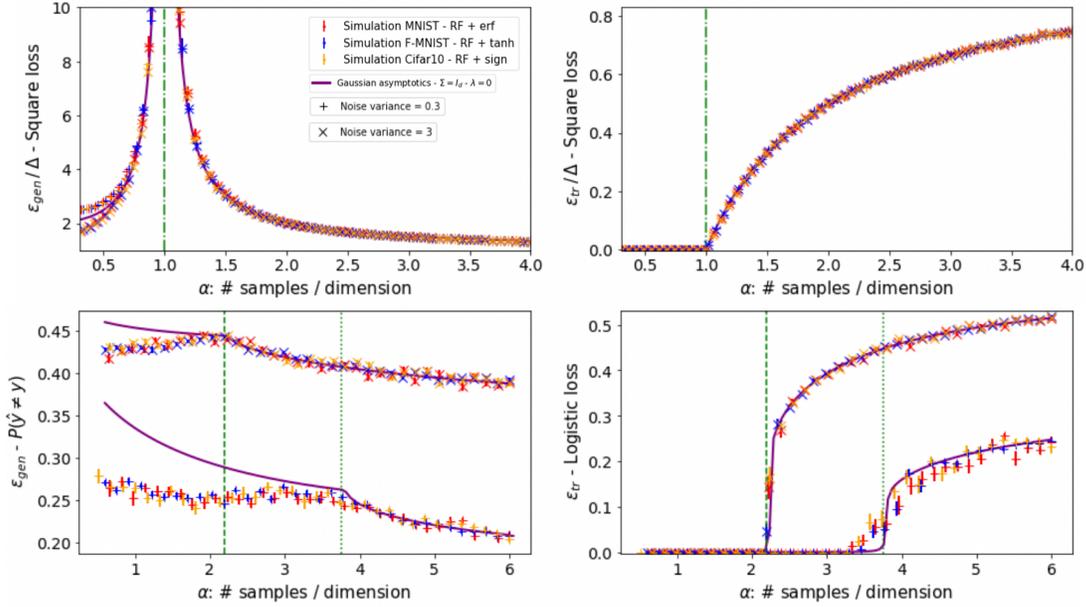


Figure 1. An illustration of Gaussian universality with vanishing regularization $\lambda = 0^+$ for a selection of datasets (MNIST, Fashion-MNIST, Cifar10) with a random teacher function, after a random feature map: Generalization (left) and training (right) errors as a function of the number of samples per dimension $\alpha = n/d$ for ridge regression (top panels) and logistic classification (bottom panels). The solid purple line is the exact asymptotics formula for Gaussian covariates with identity covariance, while the vertical green lines are the threshold values α^* (Δ) at which a unique minimizer of the loss starts to exist. Dots show numerical simulations for different real datasets with random features maps. All data follows the Gaussian predictions for the training loss, illustrating Theorems 3.4 and 3.6. In particular, the separability, or interpolation, threshold is the one of the Gaussian model (see corollary 3.5). Each learning task was run for two different values for the noise variance corrupting the labels - crosses are associated to $\Delta = 3$ while pluses to $\Delta = 0.3$. Error bars are built using standard deviation over 30 runs.

Accordingly, for two sequences of random variables $(X_n), (Y_n)$, we define closeness in probability by

$$X_n \stackrel{P}{\simeq} Y_n \quad \text{iff} \quad X_n - Y_n \xrightarrow{P} 0, \quad (9)$$

where \xrightarrow{P} denotes convergence in probability.

3.1. Exact Asymptotics

Our first result is to give closed-form asymptotic characterization of the performance of the minimizer of 2 for the Gaussian Mixture model 2.2, generalizing the rigorous results of (Loureiro et al., 2021b):

Proposition 3.1. (Exact asymptotics, informal statement)

Consider the empirical risk minimization problem introduced in eq. (1) under Gaussian mixture data given by Model 2.2. For any pseudo-Lipschitz performance metric g , the training and generalization errors (2) converge in the high-dimensional limit of $n, d \rightarrow \infty$ with fixed ratio $\alpha = n/d$ to deterministic expressions which are entirely determined by the solution of a set of self-consistent replica saddle-point equations (36).

$$\begin{aligned} \epsilon_{tr}(\hat{\theta}) &\stackrel{P}{\simeq} \epsilon_{tr}(\theta_0, \{\mu_c\}_{c \in \mathcal{C}}, \{\Sigma_c\}_{c \in \mathcal{C}}) \\ \epsilon_{gen}(\hat{\theta}) &\stackrel{P}{\simeq} \epsilon_{gen}(\theta_0, \{\mu_c\}_{c \in \mathcal{C}}, \{\Sigma_c\}_{c \in \mathcal{C}}) \end{aligned} \quad (10)$$

In order to emphasize the data model considered, we refer to the errors in the following as:

$$\epsilon_{gen/tr}(\hat{\theta}) \stackrel{P}{\simeq} \epsilon_{gen/tr}^{\text{GMM}}(\theta_0, \{\mu_c\}_{c \in \mathcal{C}}, \{\Sigma_c\}_{c \in \mathcal{C}}), \quad \text{if } |\mathcal{C}| > 1. \quad (11)$$

$$\epsilon_{gen/tr}(\hat{\theta}) \stackrel{P}{\simeq} \epsilon_{gen/tr}^{\text{GCM}}(\theta_0, \{\mu_c\}_{c \in \mathcal{C}}, \{\Sigma_c\}_{c \in \mathcal{C}}), \quad \text{if } |\mathcal{C}| = 1. \quad (12)$$

The key difference between Prop. 3.1 and Thm. 1 from (Loureiro et al., 2021b) is the distribution of the labels. While the labels in (Loureiro et al., 2021b) are given by the GMM cluster index, in Prop. 3.1 we consider labels generated by the target function (4). The generic proof scheme from (Loureiro et al., 2021b), which maps the solution to the study of a so-called approximate message passing algorithm (Donoho et al., 2009), can be readily adapted to our setting. Indeed, the approximate message passing scheme in our scenario is the same, and the only difference in the proof is to include a teacher in its asymptotic analysis, similarly as was done in (Cornacchia et al., 2022). While we do not provide a formal proof of these results, we provide a derivation of the exact asymptotics based on the replica method. The full proof of Thm. 3.1 relies on encoding the information obtained in the replica computation into an AMP sequence,

and using existing results about the convergence of such sequences. This is a tedious process that yields no new insights compared to the replica computation, which we present in detail in Appendix A.

Once we specified the GMM, and properly defined the ERM in eq. (1) we can evaluate the expressions in eqs. (2),(3) as a function of low-dimensional quantities which define the sufficient statistics of the asymptotic errors, and are also known as *order parameters*. We refer to Appendix A for a detailed discussion.

3.2. Uncorrelated Teachers

An in-depth comparison of the asymptotic expressions for the errors of Models 2.1 & 2.2 reveals that their key difference lies in the way the leading target direction θ_0 correlate with the cluster means and covariances. A first step towards universality is therefore to characterize under which conditions the asymptotic errors are independent of the means. We make the following assumptions:

Assumption 1. The teacher θ_0 respects, $\forall (c, c') \in \mathcal{C} \times \mathcal{C}$:

$$\lim_{n, d \rightarrow \infty} \frac{\theta_0^\top \mu_c}{d} = 0 \quad (13)$$

$$\lim_{n, d \rightarrow \infty} \frac{1}{d} \theta_0^\top \Sigma_{c'} \left(\lambda + \sum_{c \in \mathcal{C}} \hat{V}_c^* \Sigma_c \right)^{-1} \mu_c = 0 \quad (14)$$

where $\{\hat{V}_c^*\}_{c=1}^K$ are the fixed points of the (replica) saddle point equations describing the centered GMM problem. While the condition in eq. (13) clearly characterizes a measure of uncorrelation between the teacher vector and the data distribution, this is not evident for eq. (14). However, in the particular case of ridge regression, $\left(\lambda + \sum_{c \in \mathcal{C}} \hat{V}_c^* \Sigma_c \right)^{-1}$ corresponds exactly to the deterministic equivalent of the Stieltjes transform of XX^T (Couillet & Liao, 2022). Therefore, it makes sense for this term to appear in the computations, see Appendix B for more details.

Assumption 2. The loss function and the teacher distribution are both symmetric:

$$\ell(x, y) = \ell(-x, -y) \quad (15)$$

$$P_0(y|\tau) = P_0(-y|-\tau), \quad (16)$$

and the regularization is an ℓ^2 penalty $\lambda/2 \|\cdot\|_2^2$.

Theorem 3.2. (Mean Universality)

Under Assumptions 1 and 2, the cluster means $\{\mu_c\}_{c \in \mathcal{C}}$ are not relevant in high-dimensional ERM estimation:

$$\varepsilon_{gen}^{GMM}(\theta_0, \{\mu_c\}_{c=1}^K, \{\Sigma_c\}_{c=1}^K) \simeq \varepsilon_{gen}^{GMM}(\theta_0, \mathbf{0}, \{\Sigma_c\}_{c=1}^K) \quad (17)$$

$$\varepsilon_{tr}^{GMM}(\theta_0, \{\mu_c\}_{c=1}^K, \{\Sigma_c\}_{c=1}^K) \simeq \varepsilon_{tr}^{GMM}(\theta_0, \mathbf{0}, \{\Sigma_c\}_{c=1}^K) \quad (18)$$

Therefore, intuitively mean universality can be achieved when the label generation process is uncorrelated with the data structure. In order to prove this result, we exploit the symmetries induced by Assumptions 1 and 2 to remove the means from the set of self-consistent replica equations (36), we refer to Appendix B for a detailed discussion. The assumptions on the target and loss function are not restrictive and are easily satisfied by odd target activation f_0 and margin-based losses of the form $\ell(y, z) = \ell(yz)$. Stronger universality can be shown by simplifying the data structure. Indeed, if the mixture is homogeneous (in the sense that all the covariance are identical, a condition often called homoscedasticity in statistics) we have Gaussian universality:

Theorem 3.3. (Gaussian Universality of homoscedastic GMMs) *Under the assumption of 3.2, consider an homoscedastic GMM:*

$$\Sigma_c = \Sigma \quad \forall c \in \mathcal{C}$$

Then for all α and Δ , the errors of the GMM are asymptotically equal to those of a GCM:

$$\varepsilon_{gen}^{GMM}(\theta_0, \{\mu_c\}_{c=1}^K, \Sigma) \simeq \varepsilon_{gen}^{GCM}(\theta_0, \mathbf{0}, \Sigma) \quad (19)$$

$$\varepsilon_{tr}^{GMM}(\theta_0, \{\mu_c\}_{c=1}^K, \Sigma) \simeq \varepsilon_{tr}^{GCM}(\theta_0, \mathbf{0}, \Sigma) \quad (20)$$

These results follow in a straightforward way from Theorem 3.2. We have mapped under controllable assumptions a GMM problem to a simpler Gaussian one.

Additionally, if we consider vanishing regularization, we can prove that errors are independent of the shape of the covariance, and therefore we can take an isotropic mixture:

Theorem 3.4. (Covariance universality for $\lambda = 0^+$) *Under the assumption of Thm. 3.3, assume that it exists a unique minimizer of the empirical risk (1) with zero regularization. Then, the test & training errors a homoscedastic GMM (2.2) estimation problem asymptotically coincide with those of a centered, isotropic Gaussian model $\mathcal{G}_{(\theta_0, \mathbf{I}, 0)}$ (2.1).*

The proof of this general result is given in Appendix B.

Since the training error for losses such as logistic or hinge characterizes the separability transition at $\lambda = 0^+$, an interesting consequence of Thm. 3.4 is the following:

Corollary 3.5. (Universality of linear separability for homoscedastic GMMs) *The location of the separability/interpolation transition α_c above which the data stop to be linearly separable is the same for homoscedastic GMMs and the Gaussian model.*

This universality is particularly interesting in light of the detailed study of the separability transition for random teacher weights and *Gaussian data* in (Candès et al., 2020). Similar universality phenomena were observed for the reconstruction transition in linear estimation in (Abbasi et al., 2019).

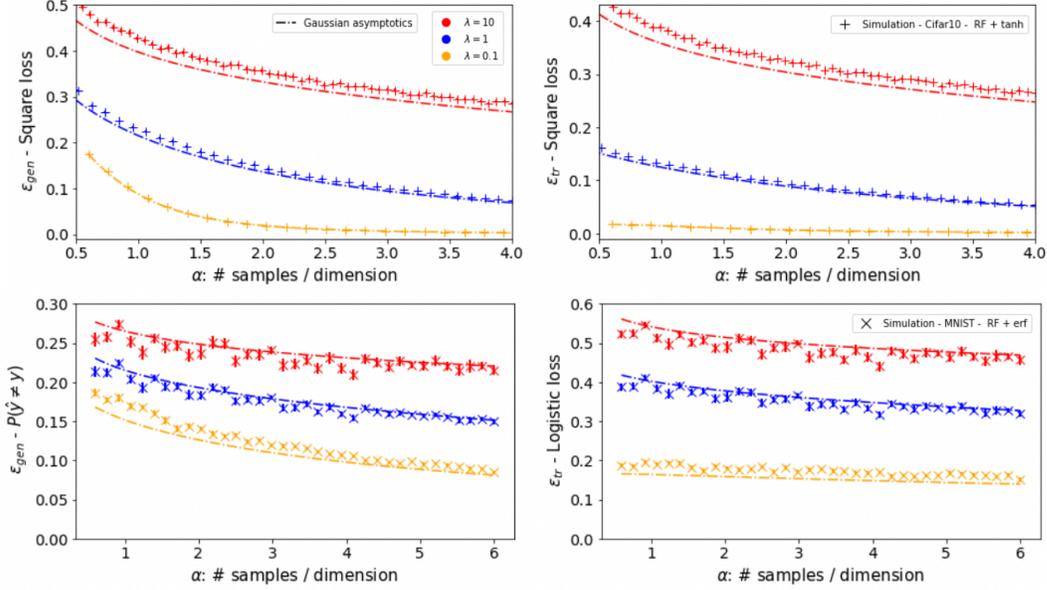


Figure 2. An illustration of Gaussian universality with finite regularization λ for a selection of datasets (MNIST, Fashion-MNIST, Cifar10) with a random teacher function, after a random feature map: Generalization (left) and training (right) errors as a function of the number of samples per dimension $\alpha = n/d$. In the upper panel, we show ridge regression on Cifar10 preprocessed with RF and tanh activation, while the lower one is logistic regression on MNIST with RF and erf activation. The dashed curves are the exact asymptotics prediction of the Gaussian theory by matching covariance. The good agreement illustrates the property of theorem 3.3. The different colors represent different regularization strength $\lambda \in \{0.1, 1, 10\}$ respectively in yellow, blue, and red. Error bars are built over 30 runs.

Surprisingly, if we further consider a square loss minimization, the estimation of *any* GMM (homoscedastic or not!) under mean universality condition can be mapped to those of a trivial Gaussian problem:

Theorem 3.6. (Strong universality of the square loss for $\lambda = 0^+$) Consider underparametrized learning of GMMs with general means and covariance respecting the assumptions of Theorem 3.2. Set in eq. (1): $l(y, \hat{y}) = (y - \hat{y})^2$ and consider a generic noisy linear teacher activation $P_0(\tau; \Delta) = \mathcal{N}(\tau, \Delta)$, with $\xi \sim \mathcal{N}(0, 1)$. Then, the training error for a GMM estimation problem is given by

$$\varepsilon_{tr}^{\text{GMM}}(\boldsymbol{\theta}_0, \{\boldsymbol{\mu}_c\}_{c=1}^K, \{\Sigma_c\}_{c=1}^K) \simeq \frac{(\alpha - 1)\Delta}{\alpha} \quad (21)$$

Note that the strong universality statement does not hold for the test error (a counterexample is discussed in Appendix C using a strongly heteroscedastic case). However, as we see from Fig. 1, it seems surprisingly true that random teacher regression on real data follows the Gaussian asymptotic prediction in the underparametrized region. Moreover, we observe in the lower panel of Fig. 1 that even for non-quadratic losses we can draw a similar conclusion, and we investigate this further in the next section.

Finally, we note that in the limit of infinite noise and binary labels, our results give back the ones observed for purely random Rachemacher labels proven in (Gerace et al., 2023).

3.3. Correlated Teachers

We now consider the general case where the target weights correlate with the structure in the data, relaxing the assumptions in Theorem 3.2. Can we still say something? Our first result shows that the answer is yes. We focus on a simple controlled setting in which we can express the ERM performance in a closed form for *any* teacher vector:

Theorem 3.7. (Exact asymptotics for isotropic covariance) Consider a ridge regression task with a 2-clusters GMM (2.2). Note that, without loss of generality we can take $\boldsymbol{\mu}_+ = -\boldsymbol{\mu}_- = \boldsymbol{\mu}$. Assume isotropic covariances:

$$\Sigma_+ = \Sigma_- = \Sigma = I_d, \quad (22)$$

and denote:

$$\rho = \lim_{d \rightarrow \infty} \frac{1}{d} \boldsymbol{\theta}_0^\top \Sigma \boldsymbol{\theta}_0, \quad \gamma = \lim_{d \rightarrow \infty} \frac{1}{d} \|\boldsymbol{\mu}\|_2^2 \quad (23)$$

$$\pi = \lim_{d \rightarrow \infty} \frac{1}{d} \boldsymbol{\mu}^\top \boldsymbol{\theta}_0 \quad P_0(\tau; \Delta) = \mathcal{N}(\tau, \Delta). \quad (24)$$

Defining:

$$A(\eta) = \frac{(\eta - 1)^2 (\gamma \eta^2 + 2\eta - 1)}{(\Delta + (\eta - 1)^2 \rho) (1 + \gamma \eta)^2} \quad \text{with} \quad (25)$$

$$\eta = 1 + \frac{1}{2} \left(\alpha - 1 + \lambda - \sqrt{4\lambda + (\alpha - 1 + \lambda)^2} \right)$$

The asymptotic errors admit a closed-form expression in terms of η :

$$\varepsilon_{gen}^{GMM}(\theta_0, \mu_{\pm}, \Sigma) \simeq g(\alpha, \Delta, \rho, \eta)(1 - \pi^2 A(\eta)) \quad (26)$$

$$\varepsilon_{tr}^{GMM}(\theta_0, \mu_{\pm}, \Sigma) \simeq t(\alpha, \Delta, \rho, \eta)(1 - \pi^2 A(\eta)) \quad (27)$$

where g and t are the asymptotic limits of the generalization $\varepsilon_{gen}^{GCM}(\theta_0, \mathbf{0}, I)$ and training $\varepsilon_{tr}^{GCM}(\theta_0, \mathbf{0}, I)$ error for a single Gaussian model with unit covariance and uncorrelated teacher:

$$g(\alpha, \Delta, \rho, \eta) = \frac{\alpha(\Delta + (\eta - 1)^2 \rho)}{(\alpha - \eta^2)} \quad (28)$$

$$t(\alpha, \Delta, \rho, \eta) = \frac{(\alpha - \eta)^2(\Delta + (\eta - 1)^2 \rho)}{\alpha(\alpha - \eta^2)} \quad (29)$$

The full closed-form expressions are derived in Appendix C. Note that the correction factor to the Gaussian performance scales with π^2 , a measure of the correlation between teacher vector and data structure. Hence, we would be tempted to state that targets that correlate with the data structure, i.e. $\pi \neq 0$, always break Gaussian universality. Although this is usually the case, in the limit of vanishing regularization we are in the position to present an interesting corollary of Theorem 3.7:

Corollary 3.8. (Restoration of universality for $\lambda = 0^+$) Consider the same setting of Theorem 3.7, further consider underparametrized learning in the limit of vanishing regularization. Then, the test and training errors for a GMM estimation are equal to a Gaussian one for any teacher vector (that is eqs (26) and (27) at $\eta = 0$):

$$\varepsilon_{gen}^{GMM}(\theta_0, \mu_{\pm}, \Sigma) = \Delta \frac{\alpha}{\alpha - 1} \quad (30)$$

$$\varepsilon_{tr}^{GMM}(\theta_0, \mu_{\pm}, \Sigma) = \Delta \frac{(\alpha - 1)}{\alpha} \quad (31)$$

We thus restore Gaussian universality for correlated teachers in the underparametrized regime, in a similar fashion to what Theorem 3.4 is stating for general convex losses and covariances under the mean universality condition. The universality property for correlated teachers is valid also for more general homoscedastic mixtures with identity covariance. For the sake of brevity, we refer to Appendix C where we discuss in detail this interesting extension to more general homoscedastic mixtures.

4. Illustration on Synthetic and Real Datasets

In this section, we investigate the consequences of our main theoretical results. First, we consider the case of real data, illustrating the applicability of our universality theorems in cases in which the target is not correlated with the data structure. Based on these observations we move on to studying simple synthetic settings described in Theorem 3.7 and

for which we can derive analytical results and systematically probe Gaussian universality. All the codes used in our experiments are available in a [GitHub repository](#).

4.1. Random Teacher Universality

We analyze Gaussian universality of real data under the random teacher function assumptions of Theorem 3.2. We consider three standard datasets: MNIST, Fashion-MNIST, and grayscale CIFAR-10, as well as synthetic data.

Universality with random feature maps — To go beyond simple linear fit, we pre-process the data with a random feature map (Rahimi & Recht, 2007) as follows: take an image $\omega_\nu \in \mathbb{R}^{d'}$ and map it to $x_\nu = \sigma(F\omega_\nu)$, where the elements $F \in \mathbb{R}^{d \times d'}$ are i.i.d drawn from $\mathcal{N}(0, 1)$ and $\sigma(\cdot)$ is an activation function. As shown in (Rahimi & Recht, 2007), for $d \rightarrow \infty$ this converges to a kernel method. For each dataset we take a different non-linearity (erf, tanh, and sign as described in Fig. 1). We then build new labels by plugging in eq. (4) $\theta_0 \sim \mathcal{N}(\mathbf{0}, I_d)$, and we set for regression tasks $f_0(\tau) = \tau + \Delta$, while for classification tasks $f_0(\tau) = \text{sign}(\tau + \Delta)$.

First, we analyze the vanishing regularization case in Fig. 1: random teacher regression on pre-processed real datasets respects the strong universality of the training loss as stated in Theorem 3.6. Interestingly, we also observe a perfect match between the Gaussian asymptotics prediction and the simulations in Fig. 1 for the generalization error. This suggests that after the random features map the data is sufficiently close to a homoscedastic mixture. Further, the lower panel of Fig. 1 shows that the strong universality property seems to hold beyond square loss minimization.

We analyze as well the finite regularization setting in Fig. 2: we compare the simulations on real data with the prediction of the exact Gaussian asymptotics: we match the covariance of each pre-processed real dataset and compute the performance of the ERM estimator thanks to the deterministic replica formula. 3.1. The predictions of the Gaussian theory match the numerical simulations. As discussed for Fig. 1, it seems that the homoscedasticity assumption in Theorem 3.3 can be sometimes relaxed.

Universality of the double descent phenomena — One finding in modern machine learning that goes against the classical statistical theory wisdom is the double descent behavior of learning curves (Oppen & Kinzel, 1996; Belkin et al., 2019; Spigler et al., 2019; Mei & Montanari, 2019), see the upper panel of Fig. 1: the test error does not deteriorate as the number of parameters is increased with a characteristic divergence at $\alpha = 1$, known as the interpolation peak. As shown in Fig. 1 and Fig. 2, the generalization error and its characteristic “double descent” behavior for

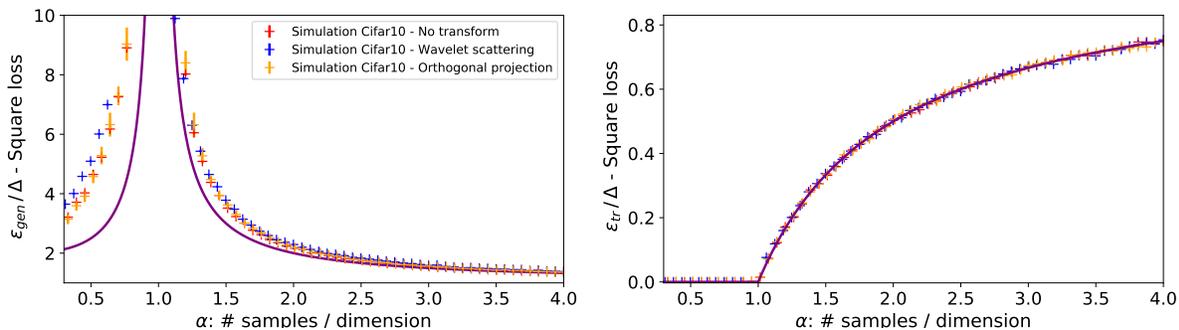


Figure 3. An illustration of Gaussian universality for Training, and lack thereof for Generalization at vanishing regularization. Instead of a random feature map as in Fig. 1, we pre-processed grayscale Cifar10 with wavelet scattering (blue dots) or orthogonal Hadamard projections (yellow dots), or used directly the raw data (red dots). Using again a random teacher function, the training error is found to agree perfectly with the universal Gaussian asymptotics predictions with identity covariance and $\lambda = 0^+$, as for theorem 3.6, but the generalization is found to show clear deviation. This highlights the roles of heteroscedasticity in the data. Error bars are built over 30 runs.

$\alpha > 1$ is universal for homoscedastic data, while the training error appears universal even for heteroscedastic data (see Fig. 3).

Universality of linear separability — Recently, (Candès et al., 2020) investigated the linear separability of Gaussian data with a random teacher and noise Δ , generalizing the classical results by (Cover, 1965) to a single-index target. (Candès et al., 2020) has shown the existence of a critical phase transition $\alpha_c(\Delta)$ that goes continuously from $\alpha_c(\infty) = 2$ (for infinite noise) to $\alpha_c > 2$ for finite Δ (with $\alpha_c \rightarrow \infty$ as $\Delta \rightarrow 0$). As discussed in corollary 3.5, this transition is universal for homoscedastic mixtures. The interest of this transition thus extends way beyond Gaussian data and in fact agrees with the real data experiment in Figs. 1 and 2. As for the double descent phenomena, we believe it is a very interesting consequence of our theorems that the theoretical works mentioned above are valid way beyond the simple Gaussian assumption.

Non-universal behavior for generalization — Finally, we implemented different transforms beyond random features for the pre-processing step: we considered the wavelet scattering transform (Andreux et al., 2020), orthogonal random projections (Choromanski et al., 2017), and even no transform at all. Without the shuffling of the random projection, the fact that the data are complex and probably rather heteroscedastic than homoscedastic, we expected a weaker form of universality. We present the results of ridge regression with labels generated by a random teacher vector in Fig. 3. The strong universality statement in Theorem 3.6, which did not require any assumption on the data, remains valid and we observed a perfect collapse of the training data. However, as expected, the generalization error shows clear

deviations with respect to the Gaussian behavior. Presumably, these transforms do not homogenize the data enough such that Gaussian universality for the test error to hold.

4.2. Correlated Teachers

Previously, we showed that ridge regression with a random teacher model, even when the data is structured, can lead to universal behavior. We now consider the dual task: take a simple homogeneous model for the data and study correlated target weights. Indeed, some recent works have studied examples of the lack of universality between Gaussian mixtures and Gaussian models (Tomasini et al., 2022; Ingresso & Goldt, 2022); we want to follow this direction and use the setting described in Theorem 3.7: we consider a 2-cluster GMM with opposite means of norm μ and same covariance Σ . We build a series of learning tasks at fixed (α, λ) varying the overlap between the teacher vector and the cluster means as follows:

$$\theta_0(\Omega) = \Omega \mu_{\perp} + \sqrt{1 - \Omega^2} \mu \quad (32)$$

$$\text{such that } \mu^{\top} \mu_{\perp} = 0 \quad (33)$$

The results are presented in Fig. 4 for both ridge and logistic regression. They clearly show that a small correlation between the target weights and the mixture means breaks Gaussian Universality. A neat geometrical intuition of the result is given in Fig. 4: as we decrease Ω , we generate labels that are more and more correlated with the data distribution, and consequently there is a correction factor to the Gaussian prediction as Theorem 3.7 predicts. We refer to Appendix. C for a more detailed discussion on correlated teachers. We conclude that universality is broken in tasks where the labels are correlated with the structure. Note, however, that as proven in Corollary. 3.8: the discrepancy

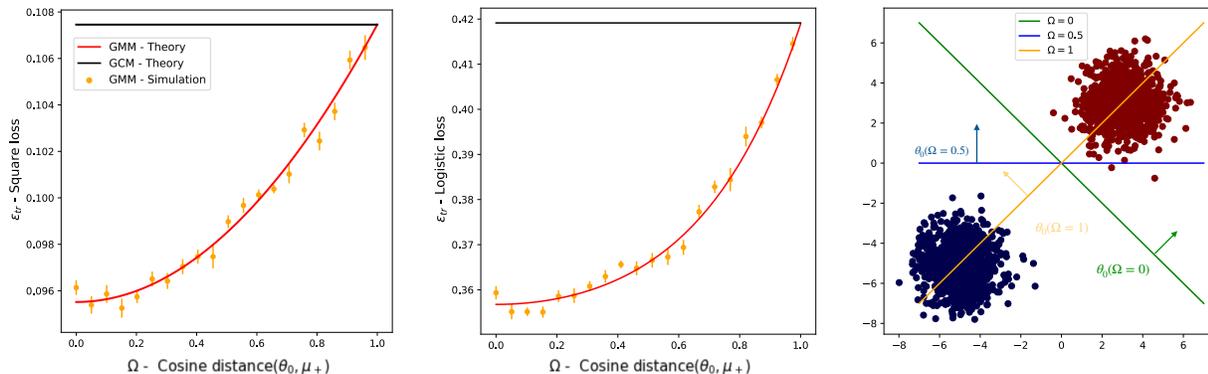


Figure 4. On the importance of the correlation of the task and the structure of the data: **Left & Center** : Training errors achieved in the ERM for estimation of 2-cluster GMM with $\mu_+ = -\mu_-$ and $\Sigma_+ = \Sigma_- = I_d$, plotted versus a correlation parameter Ω : we build a series of learning tasks with teacher weights dependent on Ω , namely we take $\theta_0(\Omega) = \Omega\mu_{\perp} + \sqrt{1 - \Omega^2}\mu_+$, with $\mu_+^{\top}\mu_{\perp} = 0$. We fix the dimension to be $d = 500$, and $(\alpha, \lambda) = (1.2, 0.7)$. The solid black line is the theoretical prediction coming from Gaussian asymptotics. The solid red line is the theoretical prediction for the GMM performance, while the orange dots are the numerical simulations that agree as expected with the theoretical prediction. In the left panel, we have real labels and perform ridge regression, while on the central one, we consider binary labels and perform logistic regression. The error bars are built using standard deviation over 30 runs. **Right**: Geometrical intuition plotted for $d = 2$. Three hyperplanes (lines in 2D) are displayed correspondent to $\Omega = \{0, 0.5, 1\}$, respectively in green, blue, and orange. As Ω increase the labels become more uncorrelated with the data structure.

between the GMM prediction and the Gaussian one goes to zero for any correlation measure Ω as $\lambda \rightarrow 0^+$, where the universality is restored as suggested by Thm. 3.6.

To conclude, and to give a short answer to the question asked in the title, we find that while Gaussian data are not all you need, they seem more valuable, and certainly more relevant, than would be naively thought.

5. Acknowledgements

We thank Yatin Dandi, Federica Gerace, Sebastian Goldt, Denny Wu & Lenka Zdeborová for useful discussions. We acknowledge funding from the Swiss National Science Foundation grant SNFS OperaGOST, 200021_200390 and the Choose France - CNRS AI Rising Talents program.

References

Abbasi, E., Salehi, F., and Hassibi, B. Universality in learning from linear measurements. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Andreux, M., Angles, T., Exarchakis, G., Leonarduzzi, R., Rochette, G., Thiry, L., Zarka, J., Mallat, S., Andén, J., Belilovsky, E., Bruna, J., Lostanlen, V., Chaudhary, M., Hirn, M. J., Oyallon, E., Zhang, S., Cella, C., and Eickenberg, M. Kymatio: Scattering transforms in python. *Journal of Machine Learning Research*, 21(60):1–6, 2020.

Ba, J., Erdogdu, M. A., Suzuki, T., Wang, Z., Wu, D., and

Yang, G. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 37932–37946. Curran Associates, Inc., 2022.

Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, apr 2020. doi: 10.1073/pnas.1907378117.

Belkin, M., Hsu, D., Ma, S., and Mandal, S. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, jul 2019. doi: 10.1073/pnas.1903070116.

Bordelon, B., Canatar, A., and Pehlevan, C. Spectrum dependent learning curves in kernel regression and wide neural networks. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1024–1034. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/bordelon20a.html>.

Candès, E. J., Sur, P., et al. The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression. *The Annals of Statistics*, 48(1):27–42, 2020.

Celentano, M., Montanari, A., and Wei, Y. The lasso with

- general gaussian designs with applications to hypothesis testing, 2020.
- Choromanski, K. M., Rowland, M., and Weller, A. The unreasonable effectiveness of structured random orthogonal embeddings. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Cornacchia, E., Mignacco, F., Veiga, R., Gerbelot, C., Loureiro, B., and Zdeborová, L. Learning curves for the multi-class teacher-student perceptron. *Machine Learning: Science and Technology*, 2022.
- Couillet, R. and Liao, Z. *Random Matrix Methods for Machine Learning*. Cambridge University Press, 2022. doi: 10.1017/9781009128490.
- Cover, T. M. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE transactions on electronic computers*, (3):326–334, 1965.
- Cui, H., Loureiro, B., Krzakala, F., and Zdeborová, L. Generalization error rates in kernel regression: The crossover from the noiseless to noisy regime. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 10131–10143. Curran Associates, Inc., 2021.
- Cui, H., Loureiro, B., Krzakala, F., and Zdeborová, L. Error rates for kernel classification under source and capacity conditions, 2022.
- Deng, Z., Kammoun, A., and Thrampoulidis, C. A model of double descent for high-dimensional binary linear classification. *Information and Inference: A Journal of the IMA*, 11(2):435–495, June 2022. ISSN 2049-8772. doi: 10.1093/imaiai/iaab002.
- Dobriban, E. and Wager, S. High-Dimensional Asymptotics of Prediction: Ridge Regression and Classification. *The Annals of Statistics*, 46(1):247–279, 2018. ISSN 0090-5364.
- Donoho, D. and Tanner, J. Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906):4273–4293, 2009.
- Donoho, D. L., Maleki, A., and Montanari, A. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- Gardner, E. and Derrida, B. Optimal storage properties of neural network models. *Journal of Physics A: Mathematical and general*, 21(1):271, 1988.
- Gerace, F., Loureiro, B., Krzakala, F., Mezard, M., and Zdeborova, L. Generalisation error in learning with random features and the hidden manifold model. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3452–3462. PMLR, 13–18 Jul 2020.
- Gerace, F., Krzakala, F., Loureiro, B., Stephan, L., and Zdeborová, L. Gaussian universality of perceptrons with random labels, 2023.
- Goldt, S., Mézard, M., Krzakala, F., and Zdeborová, L. Modeling the influence of data structure on learning in neural networks: The hidden manifold model. *Phys. Rev. X*, 10:041044, Dec 2020. doi: 10.1103/PhysRevX.10.041044.
- Goldt, S., Loureiro, B., Reeves, G., Krzakala, F., Mezard, M., and Zdeborova, L. The gaussian equivalence of generative models for learning with shallow neural networks. In Bruna, J., Hesthaven, J., and Zdeborova, L. (eds.), *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, volume 145 of *Proceedings of Machine Learning Research*, pp. 426–471. PMLR, 16–19 Aug 2022.
- Gordon, Y. Some inequalities for Gaussian processes and applications. *Israel Journal of Mathematics*, 50(4):265–289, December 1985. ISSN 1565-8511. doi: 10.1007/BF02759761.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. Surprises in high-dimensional ridgeless least squares interpolation. *Annals of Statistics*, 50(2):949–986, April 2022. ISSN 0090-5364. doi: 10.1214/21-aos2133.
- Hu, H. and Lu, Y. M. Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, pp. 1–1, 2022. doi: 10.1109/TIT.2022.3217698.
- Ingrosso, A. and Goldt, S. Data-driven emergence of convolutional structure in neural networks. *Proceedings of the National Academy of Science*, 119(40):e2201854119, October 2022. doi: 10.1073/pnas.2201854119.
- Jacot, A., Simsek, B., Spadaro, F., Hongler, C., and Gabriel, F. Kernel alignment risk estimator: Risk prediction from training data. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 15568–15578. Curran Associates, Inc., 2020.

- Karoui, N. E. Spectrum estimation for large dimensional covariance matrices using random matrix theory. *The Annals of Statistics*, pp. 2757–2790, 2008.
- Krogh, A. and Hertz, J. A. A simple weight decay can improve generalization. In *Advances in neural information processing systems*, pp. 950–957, 1992.
- Lei, L., Bickel, P. J., and El Karoui, N. Asymptotics for high dimensional regression M-estimates: fixed design results. *Probability Theory and Related Fields*, 172(3): 983–1079, December 2018. ISSN 1432-2064. doi: 10.1007/s00440-017-0824-7.
- Liang, T. and Sur, P. A precise high-dimensional asymptotic theory for boosting and minimum- ℓ_1 -norm interpolated classifiers. *The Annals of Statistics*, 50(3):1669 – 1695, 2022. doi: 10.1214/22-AOS2170.
- Loureiro, B., Gerbelot, C., Cui, H., Goldt, S., Krzakala, F., Mezard, M., and Zdeborová, L. Learning curves of generic features maps for realistic datasets with a teacher-student model. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 18137–18151. Curran Associates, Inc., 2021a.
- Loureiro, B., Sicuro, G., Gerbelot, C., Pocco, A., Krzakala, F., and Zdeborová, L. Learning gaussian mixtures with generalized linear models: Precise asymptotics in high-dimensions. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 10144–10157. Curran Associates, Inc., 2021b.
- Mai, X., Liao, Z., and Couillet, R. 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, May 2019. doi: 10.1109/ICASSP.2019.8683376. ISSN: 2379-190X.
- Mannelli, S. S., Gerace, F., Rostamzadeh, N., and Saglietti, L. Unfair geometries: exactly solvable data model with fairness implications, 2023.
- Mei, S. and Montanari, A. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75, 2019.
- Mei, S., Misiakiewicz, T., and Montanari, A. Generalization error of random feature and kernel methods: Hypercontractivity and kernel matrix concentration. *Applied and Computational Harmonic Analysis*, 59:3–84, 2022. ISSN 1063-5203. doi: <https://doi.org/10.1016/j.acha.2021.12.003>. Special Issue on Harmonic Analysis and Machine Learning.
- Mezard, M., Parisi, G., and Virasoro, M. A. *Spin Glass Theory And Beyond: An Introduction To The Replica Method And Its Applications*. World Scientific Publishing Company, November 1987. ISBN 9789813103917. Google-Books-ID: DwY8DQAAQBAJ.
- Mignacco, F., Krzakala, F., Lu, Y., Urbani, P., and Zdeborova, L. The role of regularization in classification of high-dimensional noisy Gaussian mixture. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6874–6883. PMLR, 13–18 Jul 2020.
- Montanari, A. and Nguyen, P.-M. Universality of the elastic net error. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 2338–2342, 2017. doi: 10.1109/ISIT.2017.8006947.
- Montanari, A. and Saeed, B. N. Universality of empirical risk minimization. In Loh, P.-L. and Raginsky, M. (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 4310–4312. PMLR, 02–05 Jul 2022.
- Muthukumar, V., Narang, A., Subramanian, V., Belkin, M., Hsu, D., and Sahai, A. Classification vs regression in overparameterized regimes: Does the loss function matter? *Journal of Machine Learning Research*, 22(222): 1–69, 2021.
- Opper, M. and Kinzel, W. Statistical mechanics of generalization. In *Models of neural networks III*, pp. 151–209. Springer, 1996.
- Panahi, A. and Hassibi, B. A universal analysis of large-scale regularized least squares solutions. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Rahimi, A. and Recht, B. Random Features for Large-Scale Kernel Machines. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- Shamir, O. The implicit bias of benign overfitting. In Loh, P.-L. and Raginsky, M. (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of

- Proceedings of Machine Learning Research*, pp. 448–478. PMLR, 02–05 Jul 2022.
- Spigler, S., Geiger, M., d’Ascoli, S., Sagun, L., Biroli, G., and Wyart, M. A jamming transition from under- to overparametrization affects generalization in deep learning. *Journal of Physics A: Mathematical and Theoretical*, 52(47):474001, oct 2019. doi: 10.1088/1751-8121/ab4c8b.
- Spigler, S., Geiger, M., and Wyart, M. Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12):124001, 2020.
- Stone, M. H. The generalized weierstrass approximation theorem. *Mathematics Magazine*, 21(4):167–184, 1948. ISSN 0025570X, 19300980.
- Talagrand, M. *What Is a Quantum Field Theory?* Cambridge University Press, Cambridge, 2022. ISBN 9781316510278. doi: 10.1017/9781108225144.
- Thrampoulidis, C., Abbasi, E., and Hassibi, B. Precise Error Analysis of Regularized ℓ_2 -Estimators in High Dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, August 2018. ISSN 1557-9654. doi: 10.1109/TIT.2018.2840720.
- Thrampoulidis, C., Oymak, S., and Soltanolkotabi, M. Theoretical insights into multiclass classification: a high-dimensional asymptotic view. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, pp. 8907–8920, Red Hook, NY, USA, December 2020. Curran Associates Inc. ISBN 9781713829546.
- Tomasini, U. M., Sclocchi, A., and Wyart, M. Failure and success of the spectral bias prediction for Laplace Kernel Ridge Regression: the case of low-dimensional data. pp. 21548–21583. PMLR, June 2022.
- Wang, K. and Thrampoulidis, C. 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, 2021. doi: 10.1109/ICASSP39728.2021.9413946.
- Wei, A., Hu, W., and Steinhardt, J. More than a toy: Random matrix models predict how real-world neural representations generalize. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 23549–23588. PMLR, 17–23 Jul 2022.
- Wu, D. and Xu, J. On the optimal weighted ℓ_2 regularization in overparameterized linear regression. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 10112–10123. Curran Associates, Inc., 2020.

A. Replica Computation

A.1. Formal Statement of the Theorem

We first provide the full statement of Proposition 3.1. Consider a minimization problem of the form

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \frac{1}{n} \sum_{\nu=1}^n \ell \left(y^\nu, \boldsymbol{\theta}^\top \mathbf{x}^\nu \right) + r(\boldsymbol{\theta}), \quad (34)$$

where the data (\mathbf{x}^ν, y^ν) is generated according to the following Gaussian mixture model:

$$\mathbf{x}^\nu \sim \sum_{c \in \mathcal{C}} p_c \mathcal{N} \left(\frac{\boldsymbol{\mu}_c}{\sqrt{d}}, \Sigma_c \right) \quad \text{and} \quad y^\nu \sim P_0(\cdot | \boldsymbol{\theta}_0^\top \mathbf{x}^\nu) \quad (35)$$

and assume the following:

1. The functions $\ell(y, \cdot)$ and r are continuous and coercive, and the function $\ell(y, \cdot) + r(\cdot)$ is strongly convex,
2. The covariance matrices Σ_c are positive definite, and their spectral norms are uniformly bounded,
3. The means $\boldsymbol{\mu}_c/\sqrt{d}$ and the teacher vector $\boldsymbol{\theta}_0$ are uniformly bounded,
4. the number of clusters $|\mathcal{C}|$ is finite,
5. the distribution P_0 is sub-gaussian with uniformly bounded norm.

Then, as $n, d \rightarrow \infty$ with $n/d \rightarrow \alpha > 0$, we have

$$\begin{aligned} \varepsilon_{\text{tr}}(\hat{\boldsymbol{\theta}}) &\stackrel{P}{\simeq} \sum_{c \in \mathcal{C}} p_c \mathbb{E}_{\omega_c^{(s)}, \omega_c^{(t)}, y} \left[\ell \left(y, \text{prox}_{V_c^* \ell(y, \cdot)}(\omega_c^{(s)}) \right) \right] &= \varepsilon_{\text{tr}}(\boldsymbol{\theta}_0, \{\boldsymbol{\mu}_c\}_{c \in \mathcal{C}}, \{\Sigma_c\}_{c \in \mathcal{C}}) \\ \varepsilon_{\text{gen}}(\hat{\boldsymbol{\theta}}) &\stackrel{P}{\simeq} \sum_{c \in \mathcal{C}} p_c \mathbb{E}_{\omega_c^{(s)}, \omega_c^{(t)}, y} \left[\ell \left(y, \omega_c^{(s)} \right) \right] &= \varepsilon_{\text{gen}}(\boldsymbol{\theta}_0, \{\boldsymbol{\mu}_c\}_{c \in \mathcal{C}}, \{\Sigma_c\}_{c \in \mathcal{C}}) \end{aligned} \quad (36)$$

where

$$\begin{pmatrix} \omega_c^{(s)} \\ \omega_c^{(t)} \end{pmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \pi_c \\ h_c^* \end{bmatrix}, \begin{bmatrix} \rho_c & m_c^* \\ m_c^* & q_c^* \end{bmatrix} \right) \quad \text{and} \quad y \sim P_0(\cdot | \omega_c^{(t)}). \quad (37)$$

and the prox function is

$$\text{prox}_f(x) = \min_{z \in \mathbb{R}} \left[\frac{1}{2} (z - x)^2 + f(z) \right] \quad (38)$$

The overlaps used in the equation are defined as follows:

$$\rho_c = \frac{1}{d} \boldsymbol{\theta}_0^\top \Sigma_c \boldsymbol{\theta}_0, \quad \pi_c = \frac{1}{d} \boldsymbol{\theta}_0^\top \boldsymbol{\mu}_c, \quad (39)$$

and $(V_c^*, q_c^*, m_c^*, h_c^*)_{c \in \mathcal{C}}$ are the unique fixed point of the following set of self-consistent *replica saddle-point equations*:

$$\begin{cases} \hat{V}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) \partial_\omega f \ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \\ \hat{q}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f \ell(y, h_c + \sqrt{q_c} \xi_c, V_c)^2 \right] \\ \hat{m}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \partial_\omega \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f \ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \\ \hat{h}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f \ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \end{cases} \quad (40)$$

$$\begin{cases} V_c &= \mathbb{E}_{\{\boldsymbol{\xi}_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\hat{\boldsymbol{\eta}}^\top \hat{q}_c^{-1/2} \Sigma_c^{1/2} \boldsymbol{\xi}_c \right] \\ q_c &= \mathbb{E}_{\{\boldsymbol{\xi}_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\hat{\boldsymbol{\eta}}^\top \Sigma_c \hat{\boldsymbol{\eta}} \right] \\ m_c &= \mathbb{E}_{\{\boldsymbol{\xi}_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\boldsymbol{\theta}_0^\top \Sigma_c \hat{\boldsymbol{\eta}} \right] \\ h_c &= \mathbb{E}_{\{\boldsymbol{\xi}_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\boldsymbol{\mu}_c^\top \hat{\boldsymbol{\eta}} \right] \end{cases}$$

and we have defined the following auxiliary functions:

$$\mathcal{Z}_0(y, \omega, V) = \int \frac{d\lambda}{\sqrt{2\pi V}} P_0(y|\lambda) e^{-\frac{(\lambda-\omega)^2}{2V}} \quad (41)$$

$$f_\ell(y, \omega, V) = -V^{-1} \partial_\omega \mathcal{M}_{V\ell(y, \cdot)}(\omega) = V^{-1} \left(\text{prox}_{V\ell(y, \cdot)}(\omega) - \omega \right). \quad (42)$$

Finally, the auxiliary variable $\hat{\eta}$ is a function of $\{\xi_c, \hat{V}_c, \hat{m}_c, \hat{q}_c, \hat{h}_c\}$:

$$\hat{\eta} = \text{prox}_{r(\hat{\Sigma}^{-1/2, \cdot})} \left(\hat{\Sigma}^{-1/2} \left(\sum_{c \in \mathcal{C}} \hat{\mu}_c + \sqrt{\hat{q}_c} \Sigma_c^{1/2} \xi_c \right) \right) \quad (43)$$

with

$$\hat{\mu}_c = \hat{h}_c \mu_c + \hat{m}_c \Sigma_c \theta_0 \quad \text{and} \quad \hat{\Sigma} = \sum_{c \in \mathcal{C}} \hat{V}_c \Sigma_c \quad (44)$$

We will not prove this theorem, since the proof is virtually equivalent to the one in (Loureiro et al., 2021b; Cornacchia et al., 2022). Instead, the next sections are dedicated to the derivation of these equations, using the so-called *replica method* from statistical physics (Mezard et al., 1987). Simpler particular cases of these equations can be found in Appendix A.9.

A.2. Gibbs Measure and Free Energy

The starting point for our replica computation is to define the Gibbs measure over the weights $W \in \mathbb{R}^{K \times d}$:

$$\mu_\beta(\theta) = \frac{1}{\mathcal{Z}_\beta} e^{-\beta \left[\sum_{\nu=1}^n \ell(y^\nu, \theta^\top \mathbf{x}^\nu) + \lambda r(\theta) \right]} = \frac{1}{\mathcal{Z}_\beta} e^{-\beta r(\theta)} \prod_{\nu=1}^n e^{-\beta \ell(y^\nu, \theta^\top \mathbf{x}^\nu)} \quad (45)$$

where $\beta > 0$ is a parameter we eventually want to send to $\beta \rightarrow 0^+$, and $\mathcal{Z}_\beta \in \mathbb{R}$ is the partition function (the normalization of the Gibbs measure). For convenience, we will define the following useful notation:

$$P_\theta(\theta) \equiv e^{-\beta r(\theta)}, \quad P_\ell(y^\nu | \theta^\top \mathbf{x}^\nu) \equiv e^{-\beta \ell(y^\nu, \theta^\top \mathbf{x}^\nu)} \quad (46)$$

which allows us to write the Gibbs measure as:

$$\mu_\beta(\theta) = \frac{1}{\mathcal{Z}_\beta} P_\theta(\theta) \prod_{\nu=1}^n P_\ell(y^\nu | \theta^\top \mathbf{x}^\nu) \quad (47)$$

When $\beta \rightarrow \infty$, the Gibbs measure μ_β will concentrate around the θ that minimizes the empirical risk (1). As a result, the free energy density is defined as

$$-\beta f_\beta = \lim_{d \rightarrow \infty} \frac{1}{d} \mathbb{E} \log \mathcal{Z}_\beta \quad (48)$$

where the limit is taken with $n/d = \alpha$ fixed, and the expectation is over the distribution of the data, will concentrate around the minimum risk. In order to take the expectation explicitly, we use the replica trick:

$$\log \mathcal{Z}_\beta = \lim_{s \rightarrow 0^+} \frac{1}{s} \partial_s \mathcal{Z}_\beta^s |_{s=0} \quad (49)$$

Therefore, the replica computation boils down to the computation of the averaged replicated partition function:

$$\mathbb{E} \mathcal{Z}_\beta^s = \mathbb{E} \left[\int_{\mathbb{R}^d} d\theta P_\theta(\theta) \prod_{\nu=1}^n P_\ell(y^\nu | \theta^\top \mathbf{x}^\nu) \right]^s \quad (50)$$

$$= \prod_{\nu=1}^n \mathbb{E}_{(\mathbf{x}^\nu, y^\nu)} \int \prod_{a=1}^s d\theta^a P_\theta(\theta^a) P_\ell(y^\nu | \theta^a \top \mathbf{x}^\nu) \quad (51)$$

A.3. Taking the Average Over the Data

The main difference in this computation with respect to the usual binary Gaussian mixture is that the labels are generated by a teacher target function. In other words, the joint distribution of the data is given by:

$$P_{\theta_0}(\mathbf{x}, y) = P_0(y|\tau) \sum_{c \in \mathcal{C}} \rho_c \mathcal{N}\left(\mathbf{x}; \frac{\boldsymbol{\mu}_c}{\sqrt{d}}, I_d\right) \quad (52)$$

where $P_0(y|\tau)$ is the probability induced by the teacher activation $f_0(\tau)$. The expression in (51) can then be written as:

$$\mathbb{E}Z_\beta^s = \int \prod_{a=1}^s d\boldsymbol{\theta}^a P_\theta(\boldsymbol{\theta}^a) \left(\int_{\mathbb{R}} dy \int_{\mathbb{R}^d} d\mathbf{x} P_0(y|\boldsymbol{\theta}_0^\top \mathbf{x}) P_{\mathbf{x}}(\mathbf{x}) \prod_{a=1}^s P_\ell(y^\nu | \boldsymbol{\theta}^{a\top} \mathbf{x}^\nu) \right)^n \quad (53)$$

We make a change of variables by introducing the *local fields* $\tau = \boldsymbol{\theta}_0^\top \mathbf{x}$ and $\lambda^a = \boldsymbol{\theta}^{a\top} \mathbf{x}$. The distribution of the local fields will be a Gaussian mixture itself, and we can compute the moments for each mode:

$$\begin{aligned} \pi_c &\equiv \mathbb{E}_{\mathbf{x}^\nu} [\tau_c] = \frac{1}{d} \boldsymbol{\theta}_0^\top \boldsymbol{\mu}_c \\ \rho_c &\equiv \text{Var}_{\mathbf{x}^\nu} [\tau_c] = \frac{1}{d} \boldsymbol{\theta}_0^\top \Sigma_c \boldsymbol{\theta}_0 \\ h_c^a &\equiv \mathbb{E}_{\mathbf{x}^\nu} [\lambda_c^a] = \frac{1}{d} \boldsymbol{\theta}^{a\top} \boldsymbol{\mu}_c \\ q_c^{ab} &\equiv \text{Cov}_{\mathbf{x}^\nu} [\lambda_c^a, \lambda_c^b] = \frac{1}{d} \boldsymbol{\theta}^{a\top} \Sigma_c \boldsymbol{\theta}^b \\ m_c^a &\equiv \text{Cov}_{\mathbf{x}^\nu} [\tau_c, \lambda_c^a] = \frac{1}{d} \boldsymbol{\theta}_0^\top \Sigma_c \boldsymbol{\theta}^a \end{aligned} \quad (54)$$

Equivalently, the local fields distribution is the following low-dimensional Gaussian mixture distribution:

$$P(\tau, \boldsymbol{\lambda}) = \sum_{c \in \mathcal{C}} p_c \mathcal{N}\left(\begin{bmatrix} \tau_c \\ \mathbf{h}_c \end{bmatrix}, \begin{bmatrix} \rho_c & \mathbf{m}_c^\top \\ \mathbf{m}_c & Q_c \end{bmatrix}\right), \quad (55)$$

where we have defined the vectors $\mathbf{m}, \mathbf{h} \in \mathbb{R}^s$ with entries m^a, h^a and the matrix $Q \in \mathbb{R}^{s \times s}$ with entries q^{ab} . Therefore, we can factorize the partition function by only integrating over the local field distribution:

$$\mathbb{E}Z_\beta^s = \int \prod_{a=1}^s d\boldsymbol{\theta}^a P_\theta(\boldsymbol{\theta}^a) \left[\int dy \int d\tau P_0(y|\tau) \int \left(\prod_{a=1}^s d\lambda^a P_\ell(y|\lambda^a) \right) P(\tau, \boldsymbol{\lambda} | \mathbf{m}, \mathbf{h}, q) \right]^n \quad (56)$$

A.4. Writing as a Saddle-Point Problem

Note that (π_c, ρ_c) are fixed inputs in the problem. By Fourier transform arguments, we can write

$$\delta\left(m_c^a - \frac{1}{d} \boldsymbol{\theta}_0^\top \Sigma_c \boldsymbol{\theta}^a\right) = \int_{i\mathbb{R}} \frac{d\hat{m}_c^a}{2\pi} e^{\hat{m}_c^a (dm_c^a - \boldsymbol{\theta}_0^\top \Sigma_c \boldsymbol{\theta}^a)}, \quad (57)$$

where the integral is on the imaginary line. By doing the same arguments for the q_c^{ab} and h_c^a , it ensures that:

$$\mathbb{E}Z_\beta^s = \int \prod_{c \in \mathcal{C}} \prod_{a=1}^s \frac{dm_c^a \hat{m}_c^a}{2\pi} \frac{dh_c^a \hat{h}_c^a}{2\pi} \prod_{1 \leq a \leq b \leq s} \frac{dq_c^{ab} \hat{q}_c^{ab}}{2\pi} e^{d\Phi_{\text{rs}}^{(s)}(\mathbf{m}, \hat{\mathbf{m}}, q, \hat{q})} \quad (58)$$

where we have defined the free energy potential:

$$\begin{aligned} \Phi_{\text{rs}}^{(s)}(\mathbf{m}, \hat{\mathbf{m}}, q, \hat{q}) &= \sum_{c \in \mathcal{C}} \left[\sum_{a=1}^s \hat{m}_c^a m_c^a + \sum_{a=1}^s \hat{h}_c^a h_c^a + \sum_{1 \leq a \leq b \leq s} \hat{q}_c^{ab} q_c^{ab} \right] - \alpha \Psi_g^{(s)}(\mathbf{m}, \mathbf{h}, q) - \Psi_\theta^{(s)}(\hat{\mathbf{m}}, \hat{\mathbf{h}}, \hat{q}) \\ \Psi_\ell^{(s)}(\mathbf{m}, \mathbf{h}, q) &\equiv \log \int dy \int d\tau P_0(y|\tau) \int \prod_{a=1}^s d\lambda^a P_\ell(y|\lambda^a) P(\tau, \boldsymbol{\lambda} | \mathbf{m}, \mathbf{h}, q) \\ \Psi_\theta^{(s)}(\hat{\mathbf{m}}, \hat{\mathbf{h}}, \hat{q}) &\equiv \frac{1}{d} \log \int \prod_{a=1}^s d\boldsymbol{\theta}^a P_\theta(\boldsymbol{\theta}^a) \prod_{c \in \mathcal{C}} e^{\sum_{a=1}^s [\hat{h}_c^a \boldsymbol{\theta}^{a\top} \boldsymbol{\mu}_c + \hat{m}_c^a \boldsymbol{\theta}_0^\top \Sigma_c \boldsymbol{\theta}^a] + \sum_{1 \leq a \leq b \leq s} \hat{q}_c^{ab} \boldsymbol{\theta}^{a\top} \Sigma_c \boldsymbol{\theta}^b} \end{aligned} \quad (59)$$

Therefore, as we take the $d \rightarrow \infty$ limit, we can apply the saddle-point method (Mezard et al., 1987) to compute f_β :

$$-\beta f_\beta = \operatorname{extr}_{\mathbf{m}, \hat{\mathbf{m}}, \mathbf{h}, \hat{\mathbf{h}}, q, \hat{q}} \lim_{s \rightarrow 0^+} \frac{\Phi^{(s)}(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{h}, \hat{\mathbf{h}}, q, \hat{q})}{s} \quad (60)$$

Remark A.1. We have introduced the parameters $\hat{q}, \hat{\mathbf{m}}, \hat{\mathbf{h}}$ as pure imaginary numbers, but the optimization problem considers them as real numbers. This stems from the fact that the function Φ is holomorphic, so the integral is independent of the contour of integration. More details can be found in (Talagrand, 2022).

A.5. Replica Symmetric Ansatz

In order to make progress with the $s \rightarrow 0^+$ limit, we make the following replica symmetric ansatz:

$$\begin{aligned} m_c^a &= m_c & \hat{m}_c^a &= \hat{m}_c, & a &= 1, \dots, s \\ h_c^a &= h_c & \hat{h}_c^a &= \hat{h}_c, & a &= 1, \dots, s \\ q_c^{aa} &= r_c, & \hat{q}_c^{aa} &= -\frac{1}{2} \hat{r}_c, & a &= 1, \dots, s \\ q_c^{ab} &= q_c, & \hat{q}_c^{ab} &= \hat{q}_c, & 1 \leq a < b \leq s & \end{aligned} \quad (61)$$

for all $c \in \mathcal{C}$. Since ℓ and r are convex, the function $\Phi^{(s)}$ has a unique saddle point, which must therefore coincide with the replica symmetric ansatz. We also define

$$V_c = r_c - q_c, \quad \text{and} \quad \hat{V}_c = \hat{r}_c - \hat{q}_c \quad (62)$$

By inserting this ansatz above, we can take the $s \rightarrow 0^+$ limit. This is done through a classical but computationally heavy method known as the *Hubbard-Stratonovich transform*; details can be found in (Gerace et al., 2020), Appendix C. We simply reproduce the final results here:

$$-\beta f_\beta = \operatorname{extr}_{\{m_c, \hat{m}_c, h_c, \hat{h}_c, q_c, \hat{q}_c, V_c, \hat{V}_c\}} \Phi_{\text{rs}}(\{m_c, \hat{m}_c, h_c, \hat{h}_c, q_c, \hat{q}_c, V_c, \hat{V}_c\}) \quad (63)$$

The replicated free energy Φ_{rs} is given by the following formula:

$$\Phi_{\text{rs}}(\{m_c, \hat{m}_c, h_c, \hat{h}_c, q_c, \hat{q}_c, V_c, \hat{V}_c\}) = \sum_{c \in \mathcal{C}} \left[\frac{1}{2} (\hat{V}_c q_c - \hat{q}_c V_c) - \frac{1}{2} \hat{V}_c V_c + \hat{m}_c m_c + \hat{h}_c h_c \right] \quad (64)$$

$$- \alpha \Psi_\ell(\{m_c, h_c, q_c, V_c\}) - \Psi_\theta(\{\hat{m}_c, \hat{h}_c, \hat{q}_c, \hat{V}_c\}) \quad (65)$$

where we have decomposed the contributions coming from the loss (Ψ_ℓ) and the regularization (Ψ_θ):

$$\Psi_\ell = \sum_{c \in \mathcal{C}} \mathbb{E}_{\xi_c \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)} \int dy \mathcal{Z}_0 \left(y, \frac{m_c}{\sqrt{q_c}} \xi + \pi_c, \rho - \frac{m_c^2}{q_c} \right) \log \mathcal{Z}_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \quad (66)$$

$$\Psi_\theta = \frac{1}{d} \mathbb{E}_{\xi_c \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \log \mathcal{Z}_\theta \left(\sum_{c \in \mathcal{C}} \hat{V}_c \Sigma_c, \sum_{c \in \mathcal{C}} \hat{h}_c \boldsymbol{\mu}_c + \hat{m}_c \Sigma_c \boldsymbol{\theta}_0 + \sqrt{\hat{q}_c} \Sigma_c^{1/2} \boldsymbol{\xi}_c \right), \quad (67)$$

and defined the following auxiliary free energies:

$$\mathcal{Z}_{\ell/0}(y, \omega, V) = \int \frac{d\lambda}{\sqrt{2\pi V}} P_{\ell/0}(y|\lambda) e^{-\frac{(\lambda-\omega)^2}{2V}} \quad (68)$$

$$\mathcal{Z}_\theta(A, \mathbf{b}) = \int d\boldsymbol{\theta} e^{-\beta r(\boldsymbol{\theta})} e^{-\frac{1}{2} \boldsymbol{\theta}^\top A \boldsymbol{\theta} + \mathbf{b}^\top \boldsymbol{\theta}} \quad (69)$$

A.6. Taking the Zero Temperature Limit

In order to take the $\beta \rightarrow \infty$ limit, we make the following rescalings:

$$V_c \rightarrow \beta^{-1} V_c \quad \hat{V}_c \rightarrow \beta \hat{V}_c \quad \hat{q}_c \rightarrow \beta^2 \hat{q}_c \quad \hat{m}_c \rightarrow \beta \hat{m}_c. \quad (70)$$

It is easy to check how this rescaling affects $\beta^{-1} \Phi^{(\text{rs})}$, so we only need to consider Ψ_ℓ and Ψ_θ .

We start with the latter: letting

$$\mathcal{L}_\theta(\boldsymbol{\theta}) = \frac{1}{2} \boldsymbol{\theta}^\top A \boldsymbol{\theta} - \mathbf{b}^\top \boldsymbol{\theta} + r(\boldsymbol{\theta})$$

by the Laplace method for any A, \mathbf{b} ,

$$\lim_{\beta \rightarrow \infty} \frac{1}{\beta} \log Z_\theta(\beta A, \beta \mathbf{b}) = - \inf_{\boldsymbol{\theta}} \mathcal{L}_\theta(\boldsymbol{\theta}) = -\mathcal{L}_\theta(\boldsymbol{\eta}) \quad (71)$$

where

$$\boldsymbol{\eta} = \text{prox}_{r(A^{1/2}, \cdot)}(A^{1/2} \mathbf{b}). \quad (72)$$

As a result, every integral involved in Ψ_θ (and, later, its partial derivatives) concentrates around its value at $\hat{\boldsymbol{\eta}}$ defined in (43). For the term Ψ_ℓ , the term \mathcal{Z}_0 is left unchanged, and by the same reasoning as above

$$\lim_{\beta \rightarrow \infty} \frac{1}{\beta} \log \mathcal{Z}_\ell(y, \omega, \beta V) = -V^{-1} \mathcal{M}_{V\ell(y, \cdot)}(\omega). \quad (73)$$

A.7. Saddle-Point Equations

The saddle-point equations are obtained by taking the derivatives of the free energy potential with respect to the overlap parameters. We obtain a set of self-consistent equations which we should solve in order to find a fixed point:

$$\begin{cases} \hat{V}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) \partial_\omega f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \\ \hat{q}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c)^2 \right] \\ \hat{m}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \partial_\omega \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \\ \hat{h}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \end{cases} \quad (74)$$

$$\begin{cases} V_c &= \mathbb{E}_{\{\xi_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\hat{\boldsymbol{\eta}}^\top \hat{q}_c^{-1/2} \Sigma_c^{1/2} \boldsymbol{\xi}_c \right] \\ q_c &= \mathbb{E}_{\{\xi_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\hat{\boldsymbol{\eta}}^\top \Sigma_c \hat{\boldsymbol{\eta}} \right] \\ m_c &= \mathbb{E}_{\{\xi_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\boldsymbol{\theta}_0^\top \Sigma_c \hat{\boldsymbol{\eta}} \right] \\ h_c &= \mathbb{E}_{\{\xi_c\} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, I_d)} \left[\boldsymbol{\mu}_c^\top \hat{\boldsymbol{\eta}} \right] \end{cases} \quad (75)$$

where all the relevant quantities have been defined in Appendix A.1.

A.8. Training and Generalization Errors

It now remains to compute the training and generalization errors from the free energy. Recalling that

$$\varepsilon_{\text{tr}}(\hat{\boldsymbol{\theta}}) = \frac{1}{n} \sum_{\nu=1}^n \ell(y^\nu, \hat{\boldsymbol{\theta}}^\top \mathbf{x}^\nu),$$

we can write using the definition of the free energy (48):

$$\lim_{n \rightarrow \infty} \varepsilon_{\text{tr}} = \lim_{\beta \rightarrow \infty} \partial_\beta f_\beta - r(\hat{\boldsymbol{\theta}}).$$

Computing explicitly the derivative and averaging again over the data yields

$$\lim_{n \rightarrow \infty} \varepsilon_{\text{tr}} = - \lim_{\beta \rightarrow \infty} \partial_\beta \Psi_\ell,$$

where Ψ_ℓ is the free energy contribution of the loss defined in (66). Writing explicitly this derivative,

$$-\partial_\beta \Psi_\ell = \sum_{c \in \mathcal{C}} \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \int_{\mathbb{R}} dy \frac{\mathcal{Z}_0(y, \omega_c^{(0)}, \rho - \frac{m_c^2}{q_c})}{\mathcal{Z}_\ell(y, \omega_c^{(\ell)}, V_c)} \underbrace{\int_{\mathbb{R}} \frac{d\lambda}{\sqrt{2\pi V_c}} e^{-(\lambda - \omega_c^{(\ell)})V_c^{-1}(\lambda - \omega_c^{(\ell)}, V_c) - \beta \ell(y, \lambda)} \ell(y, \lambda)}_{\tilde{\mathcal{Z}}_\ell(y, \omega_c^{(\ell)}, V_c)} \quad (76)$$

where

$$\omega_c^{(0)} = \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c \quad \text{and} \quad \omega_c^{(\ell)} = h_c + \sqrt{q_c} \xi_c.$$

We can now make the change of variables in (70), and use again Laplace's approximation: if

$$\eta_c = \text{prox}_{V_\ell(y, \cdot)}(\omega_c^{(\ell)}) \quad \text{and} \quad \mathcal{L}(\lambda) = -(\lambda - \omega_c^{(\ell)})V_c^{-1}(\lambda - \omega_c^{(\ell)}, V_c) - \beta \ell(y, \lambda),$$

then

$$\mathcal{Z}_\ell(y, \omega_c^{(\ell)}, V_c) \sim e^{-\beta \mathcal{L}(\eta_c)} \quad \text{and} \quad \tilde{\mathcal{Z}}_\ell(y, \omega_c^{(\ell)}, V_c) \sim e^{-\beta \mathcal{L}(\eta_c)} \ell(y, \eta_c),$$

and all of the overlaps will concentrate around the solutions of (74), (75). Finally, the term containing \mathcal{Z}_0 is an expectation of y according to $P_0(y \mid \tau_c)$, where

$$\tau_c = \omega_c^{(0)} + \sqrt{\rho - \frac{m_c^2}{q_c}} \xi_c^{(0)}$$

and $\xi_c^{(0)}$ is a Gaussian variable independent from everything else. Putting all together, we can write

$$\lim_{n \rightarrow \infty} \varepsilon_{\text{tr}}(\hat{\boldsymbol{\theta}}) = \sum_{c \in \mathcal{C}} p_c \mathbb{E}_{\nu_c, \tau_c, y} \left[\ell \left(y, \text{prox}_{V_c^* \ell(y, \cdot)}(\nu_c) \right) \right] \quad (77)$$

with

$$\begin{pmatrix} \nu_c \\ \tau_c \end{pmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \pi_c \\ h_c \end{bmatrix}, \begin{bmatrix} \rho_c & m_c^* \\ m_c^* & q_c^* \end{bmatrix} \right) \quad \text{and} \quad y \sim P_0(\cdot \mid \tau_c). \quad (78)$$

The generalization error is much simpler to obtain: since $x_{\text{new}}, y_{\text{new}}$ are independent from the estimator $\hat{\boldsymbol{\theta}}$, we simply have

$$\lim_{n \rightarrow \infty} \varepsilon_{\text{gen}}(\hat{\boldsymbol{\theta}}) = \sum_{c \in \mathcal{C}} p_c \mathbb{E}_{\nu_c, \tau_c, y} [\ell(y, \nu_c)] \quad (79)$$

where ν_c, τ_c, y follow the same distribution as above.

A.9. Examples

Ridge penalty Consider a particular case of the general equations reported above: the case of a ridge penalty

$$r(\boldsymbol{\theta}) = \frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2.$$

We then have:

$$\text{prox}_r(\mathbf{x}) = (1 + \lambda)^{-1} \mathbf{x} \quad (80)$$

which simplifies the prior equations considerably. Indeed, we can now compute every expectation in (75), which yields the following fixed-point equations:

$$\begin{cases} \hat{V}_c &= -\alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) \partial_\omega f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \\ \hat{q}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c)^2 \right] \\ \hat{m}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \partial_\omega \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \\ \hat{h}_c &= \alpha p_c \mathbb{E}_{\xi_c \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi_c + \frac{m_c}{\sqrt{q_c}} \xi_c, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, h_c + \sqrt{q_c} \xi_c, V_c) \right] \end{cases} \quad (81)$$

$$\begin{cases} V_c &= \frac{1}{d} \text{tr} \left[\Sigma_c \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \\ q_c &= \frac{1}{d} \text{tr} \left[\left(\sum_{c' \in \mathcal{C}} \hat{q}_{c'} \Sigma_{c'} + \sum_{c', c'' \in \mathcal{C}} \hat{\boldsymbol{\mu}}_{c'} \hat{\boldsymbol{\mu}}_{c''}^\top \right) \Sigma_c \left(\lambda I_d + \hat{\Sigma} \right)^{-2} \right] \\ m_c &= \frac{1}{d} \text{tr} \left[\left(\sum_{c' \in \mathcal{C}} \hat{\boldsymbol{\mu}}_{c'} \boldsymbol{\theta}_0^\top \right) \Sigma_c \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \\ h_c &= \frac{1}{d} \text{tr} \left[\left(\sum_{c' \in \mathcal{C}} \hat{\boldsymbol{\mu}}_{c'} \boldsymbol{\mu}_c^\top \right) \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \end{cases} \quad (82)$$

Gaussian covariate model The equations for the Gaussian covariate model can be found in (Loureiro et al., 2021a); they also correspond to taking $|\mathcal{C}| = 1$ in the ones above. We reproduce them here for completeness:

$$\begin{cases} \hat{V} &= -\alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi + \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) \partial_\omega f_\ell(y, h + \sqrt{q} \xi, V) \right] \\ \hat{q} &= \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi + \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) f_\ell(y, h + \sqrt{q} \xi, V)^2 \right] \\ \hat{m} &= \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \partial_\omega \mathcal{Z}_0 \left(y, \pi + \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) f_\ell(y, h + \sqrt{q} \xi, V) \right] \\ \hat{h} &= \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \pi + \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) f_\ell(y, h + \sqrt{q} \xi, V) \right] \end{cases} \quad (83)$$

$$\begin{cases} V &= \frac{1}{d} \text{tr} \left[\Sigma \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \\ q &= \frac{1}{d} \text{tr} \left[\left(\hat{q} \Sigma + \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top \right) \Sigma \left(\lambda I_d + \hat{\Sigma} \right)^{-2} \right] \\ m &= \frac{1}{d} \text{tr} \left[\hat{\boldsymbol{\mu}} \boldsymbol{\theta}_0^\top \Sigma \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \\ h &= \frac{1}{d} \text{tr} \left[\hat{\boldsymbol{\mu}} \boldsymbol{\mu}^\top \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \end{cases} \quad (84)$$

where this time

$$\hat{\boldsymbol{\mu}} = \hat{h} \boldsymbol{\mu} + \hat{m} \Sigma \boldsymbol{\theta}_0 \quad \text{and} \quad \hat{\Sigma} = \hat{V} \Sigma. \quad (85)$$

The errors are given by

$$\lim_{n \rightarrow \infty} \varepsilon_{\text{tr}}(\hat{\boldsymbol{\theta}}) = \mathbb{E}_{\nu, \tau, y} \left[\ell \left(y, \text{prox}_{V^* \ell(y, \cdot)}(\nu) \right) \right] \quad (86)$$

$$\lim_{n \rightarrow \infty} \varepsilon_{\text{gen}}(\hat{\boldsymbol{\theta}}) = \mathbb{E}_{\nu, \tau, y} \left[\ell(y, \nu) \right] \quad (87)$$

with

$$\begin{pmatrix} \nu \\ \tau \end{pmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \pi \\ h \end{bmatrix}, \begin{bmatrix} \rho & m^* \\ m^* & q^* \end{bmatrix} \right) \quad \text{and} \quad y \sim P_0(\cdot | \tau). \quad (88)$$

Ridge regression We place ourselves in the ridge regression case, where

$$P_0(\tau | \Delta) = \mathcal{N}(\tau, \Delta) \quad \text{and} \quad \ell(y, \hat{y}) = (y - \hat{y})^2. \quad (89)$$

In this case, the equations in (82) simplify even further, yielding

$$\begin{cases} \hat{V}_c &= \frac{\alpha p_c}{1+V_c} \\ \hat{q}_c &= \frac{\alpha p_c}{(1+V_c)^2} (\rho_c + \Delta + q_c - 2m_c + (h_c - \pi_c)^2) \\ \hat{m}_c &= \frac{\alpha p_c}{1+V_c} \\ \hat{h}_c &= \frac{\alpha p_c (\pi_c - h_c)}{1+V_c} \end{cases} \quad (90)$$

$$\begin{cases} V &= \frac{1}{d} \text{tr} \left[\Sigma \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \\ q &= \frac{1}{d} \text{tr} \left[\left(\hat{q} \Sigma + \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top \right) \Sigma \left(\lambda I_d + \hat{\Sigma} \right)^{-2} \right] \\ m &= \frac{1}{d} \text{tr} \left[\hat{\boldsymbol{\mu}} \boldsymbol{\theta}_0^\top \Sigma \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \\ h &= \frac{1}{d} \text{tr} \left[\hat{\boldsymbol{\mu}} \boldsymbol{\mu}^\top \left(\lambda I_d + \hat{\Sigma} \right)^{-1} \right] \end{cases} \quad (91)$$

The training and generalization error also benefit from a very simple expression

$$\varepsilon_{\text{tr}} = \sum_c \frac{\hat{q}_c^*}{\alpha} \quad \varepsilon_{\text{gen}} = \sum_c \frac{\hat{q}_c^* (1 + V_c^*)^2}{\alpha}. \quad (92)$$

B. Universality of Gaussian Mixture Models

In this section, we prove the main results shown in Sec. 3 regarding the universality properties of Gaussian Mixture Models.

B.1. Mean Universality : Proof of Theorem 3.2

Let us rewrite the assumptions here. Let $\{\hat{V}_c^*\}_{c=1}^K$ be the fixed points of the (replica) saddle point equations describing the centered Gaussian mixture problem, i.e. the solutions of eqs. (81), (82). We assume that the teacher vector and the data structure respect the following:

$$\lim_{n,d \rightarrow \infty} \frac{\boldsymbol{\theta}_0^\top \boldsymbol{\mu}_c}{d} = 0 \quad \forall c \in \mathcal{C} \quad (93)$$

$$\lim_{n,d \rightarrow \infty} \frac{1}{d} \boldsymbol{\theta}_0^\top \Sigma_{c'} \left(\lambda + \sum_{c \in \mathcal{C}} \hat{V}_c^* \Sigma_c \right)^{-1} \boldsymbol{\mu}_c = 0 \quad \forall (c, c') \in \mathcal{C} \times \mathcal{C} \quad (94)$$

and the loss and teacher are both symmetric:

$$\ell(x, y) = \ell(-x, -y), \quad P_0(y|\tau) = P_0(-y|-\tau). \quad (95)$$

In the saddle-point equations (81), (82), the mean vectors appear always coupled with the overlaps $\{h_c, \hat{h}_c\}_{c \in \mathcal{C}}$. Hence, to prove Prop. 3.2 it suffices to prove the following result:

Lemma B.1. *If eqs. (93), (94), (95) hold, then $\{h_c = \hat{h}_c = 0\}_{c \in \mathcal{C}}$ is a fixed point for the problem.*

First consider the update equations for $\{h_c\}_{c \in \mathcal{C}}$:

$$h_c = \frac{1}{d} \sum_l \text{tr} \left[\left(\hat{h}_l \boldsymbol{\mu}_l \boldsymbol{\mu}_c^\top + \hat{m}_l \Sigma_l \boldsymbol{\theta}_0 \boldsymbol{\mu}_c^\top \right) \left(\lambda I_d + \sum_{l'} \hat{V}_{l'} \Sigma_{l'} \right)^{-1} \right] \quad (96)$$

$$\stackrel{\text{eqs. (93), (95)}}{\simeq} \frac{1}{d} \sum_l \text{tr} \left[\left(\hat{h}_l \boldsymbol{\mu}_l \boldsymbol{\mu}_c^\top \right) \left(\lambda I_d + \sum_{l'} \hat{V}_{l'} \Sigma_{l'} \right)^{-1} \right] \quad (97)$$

By continuity of the saddle-point equations, if at the fixed point $\{\hat{h}_c^* \rightarrow 0\}_{c \in \mathcal{C}}$ holds, we also easily have that $\{h_c^* \rightarrow 0\}_{c \in \mathcal{C}}$. Now assume $\{h_c^* = 0\}_{c \in \mathcal{C}}$ at the fixed point. We exploit the symmetry argument inherent to the update functions to show that $\{\hat{h}_c^* = 0\}_{c \in \mathcal{C}}$ under weak assumptions on the teacher and loss functions. We write the updates using assumption c):

$$\hat{h}_c = \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \frac{m_c}{\sqrt{q_c}} \xi + \pi_c, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, h_c + \sqrt{q_c} \xi, V) \right] \quad (98)$$

$$\underset{h_c \rightarrow 0}{\simeq} \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \frac{m_c}{\sqrt{q_c}} \xi, \rho_c - \frac{m_c^2}{q_c} \right) f_\ell(y, \sqrt{q_c} \xi, V_c) \right] \quad (99)$$

Reminding the definition of the teacher measure term:

$$Z_0(y, \omega, V) = \mathbb{E}_{\lambda \sim \mathcal{N}(\omega, V)} [P_0(y|\lambda)] \quad \lambda \equiv \frac{\mathbf{x}^\top \boldsymbol{\theta}_0}{\sqrt{d}} \quad (100)$$

we see that the symmetry conditions impose basic restrictions on the label generation:

$$f) \quad \mathbb{E}_{\lambda \sim \mathcal{N}(\omega, V)} [P_0(y|\lambda)] = \mathbb{E}_{\lambda \sim \mathcal{N}(-\omega, V)} [P_0(-y|\lambda)] \quad (101)$$

and hence \mathcal{Z}_0 is even in its second argument. Additionally, we have

$$f_\ell(y, \omega, V) = \frac{1}{V} (\text{prox}_{V\ell(y, \cdot)}(\omega) - \omega) \quad \text{prox}_{V\ell(y, \cdot)}(\omega) = \arg \min_z \left[\frac{1}{2V} (z - \omega)^2 + \ell(y, z) \right] \quad (102)$$

The symmetry condition on ℓ then implies that

$$\text{prox}_{V\ell(-y, \cdot)}(-\omega) = - \text{prox}_{V\ell(y, \cdot)}(\omega), \quad (103)$$

so f_ℓ is odd in its second argument. All that's left to notice is that (99) is a Gaussian integral of an odd function, hence it is equal to 0.

B.2. Gaussian Universality : Proof of Theorem 3.3

We proved that the fixed point respects $\{h_c^*, \hat{h}_c^* = 0, 0\}_{c \in \mathcal{C}}$ under the hypothesis of Prop. 3.2. Assume now that the mixture is homogeneous:

$$\Sigma_c = \Sigma \quad \forall c \in \mathcal{C} \quad (104)$$

Then we have

$$\varepsilon_{gen}^{\text{GMM}}(\boldsymbol{\theta}_0, \{\boldsymbol{\mu}_c\}_{c=1}^K, \{\Sigma\}_{c=1}^K) \simeq \varepsilon_{gen}^{\text{GMM}}(\boldsymbol{\theta}_0, \mathbf{0}, \{\Sigma\}_{c=1}^K) \quad (105)$$

$$\varepsilon_{tr}^{\text{GMM}}(\boldsymbol{\theta}_0, \{\boldsymbol{\mu}_c\}_{c=1}^K, \{\Sigma\}_{c=1}^K) \simeq \varepsilon_{tr}^{\text{GMM}}(\boldsymbol{\theta}_0, \mathbf{0}, \{\Sigma\}_{c=1}^K) \quad (106)$$

But the right-hand side of those equations corresponds to a distribution of the form

$$\mathbf{x} \sim \sum_{c \in \mathcal{C}} \mathcal{N}(\mathbf{0}, \Sigma) = \mathcal{N}(\mathbf{0}, \Sigma)!$$

This proves Theorem 3.3

B.3. Covariance Universality

Surprisingly in the limit of vanishing regularization, $\lambda \rightarrow 0^+$, the covariance is irrelevant for the high dimensional learning problem. We show that when a unique minimizer of the loss exists, and we can safely take the limit $\lambda \rightarrow 0^+$ in the saddle point equations, the covariance Σ disappears entirely from the overlap expression. Indeed assuming the minimizer $\hat{\boldsymbol{\theta}}$ is unique, we can safely simplify expressions of the type:

$$\lim_{\lambda \rightarrow 0^+} \Sigma^n (\lambda I_d + \hat{V} \Sigma)^{n'} = \frac{1}{\hat{V}} \Sigma^{n-n'} \quad (107)$$

Then by plugging in the $\lambda \rightarrow 0^+$ simplification in eqs. (83), (84) we have complete independence from the covariance matrix Σ :

$$\begin{cases} \hat{V} &= -\alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \frac{m_c}{\sqrt{q_c}} \xi, \rho_c - \frac{m_c^2}{q_c} \right) \partial_\omega f_\ell(y, \sqrt{q_c} \xi, V_c) \right] \\ \hat{q} &= \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \frac{m_c}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) f_\ell(y, \sqrt{q} \xi, V)^2 \right] \\ \hat{m} &= \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \partial_\omega \mathcal{Z}_0 \left(y, \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) f_\ell(y, \sqrt{q} \xi, V) \right] \end{cases} \quad (108)$$

$$\begin{cases} V &= 1 \\ q &= \hat{q} + \rho \hat{m}^2 \\ m &= \hat{m} \rho \end{cases} \quad (109)$$

which concludes the proof.

B.4. Strong Universality

We never made any specific assumption on the loss up to now, apart from the very general symmetry condition in Assumption 2. In this section, we show strong universality for square loss regression for *any* GMM estimation problem respecting the mean universality property (Prop. 3.2). We assume to consider a ridge regression problem in the underparametrized regime $\alpha > 1$:

$$P_0(\tau|\Delta) = \mathcal{N}(\tau, \Delta) \quad \ell(y, \hat{y}) = (y - \hat{y})^2 \quad (110)$$

the replicas then correspond to equations (90), (91). In the $\lambda \rightarrow 0^+$ limit, the equation simplify greatly:

$$\begin{cases} \hat{V}_c &= \frac{\alpha p_c}{1+V_c} \\ \hat{q}_c &= \frac{\alpha p_c}{(1+V_c)^2} (\rho_c + \Delta + q_c - 2m_c) \\ \hat{m}_c &= \frac{\alpha p_c}{1+V_c} \end{cases} \quad (111)$$

$$\begin{cases} V_c &= \frac{1}{d} \text{tr} \left[\Sigma_c (\sum_{l \in \mathcal{C}} \hat{V}_l \Sigma_l)^{-1} \right] \\ q_c &= \frac{1}{d} \sum_{l \in \mathcal{C}} \text{tr} \left[\hat{q}_l \Sigma_l \Sigma_c (\sum_{l' \in \mathcal{C}} \hat{V}_{l'} \Sigma_{l'})^{-2} \right] + \frac{1}{d} \sum_{(l,l') \in \mathcal{C} \times \mathcal{C}} \text{tr} \left[\hat{m}_l \hat{m}_{l'} \Sigma_l \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^T \Sigma_{l'} \Sigma_c (\sum_{l'' \in \mathcal{C}} \hat{V}_{l''} \Sigma_{l''})^{-2} \right] \\ m_c &= \frac{1}{d} \sum_{l=1}^K \text{tr} \left[\hat{m}_l \Sigma_l \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^T \Sigma_c (\sum_{l'} \hat{V}_{l'} \Sigma_{l'})^{-1} \right] \end{cases} \quad (112)$$

Consider the equation for the overlaps $\{q_c\}_{c \in \mathcal{C}}$:

$$q_c = \frac{1}{d} \sum_{l=1}^K \text{tr} \left[\hat{q}_l \Sigma_l \Sigma_c (\sum_{l'} \hat{V}_{l'} \Sigma_{l'})^{-2} \right] + \frac{1}{d} \sum_{l,k=1}^K \text{tr} \left[\hat{m}_l \hat{m}_k \Sigma_l \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^T \Sigma_k \Sigma_c (\sum_{l'} \hat{V}_{l'} \Sigma_{l'})^{-2} \right] \quad (113)$$

$$\equiv R_c + G_c \quad (114)$$

The error metrics can be decomposed as:

$$\varepsilon_{\text{tr}} = \sum_{c \in \mathcal{C}} \frac{p_c}{(1+V_c)^2} (R_c + \Delta) + \sum_{c \in \mathcal{C}} \frac{p_c}{(1+V_c)^2} (\rho_c + G_c - 2m_c) \quad (115)$$

$$\varepsilon_{\text{gen}} = \sum_{c \in \mathcal{C}} p_c (R_c + \Delta) + \sum_{c \in \mathcal{C}} p_c (\rho_c + G_c - 2m_c) \quad (116)$$

We focus on the second term and show that it is equal to zero. Looking at eqs. (112),(111), we note that:

$$\hat{m}_c = \hat{V}_c \quad \forall c \in \mathcal{C} \quad (117)$$

By using eq. (117) we can simplify the equations for $\{\hat{m}_c, G_c\}$:

$$G_c = \rho_c, \quad m_c = \rho_c \quad \forall c \in \mathcal{C} \quad (118)$$

Plugging these relations in the saddle point equations we obtain:

$$\begin{cases} \hat{V}_c &= \frac{\alpha p_c}{1+V_c} = \hat{m}_c \\ \hat{q}_c &= \frac{\alpha p_c}{(1+V_c)^2} (\Delta + R_c) \end{cases} \quad (119)$$

$$\begin{cases} V_c &= \frac{1}{d} \text{tr} \left[\Sigma_c (\sum_{l \in \mathcal{C}} \hat{V}_l \Sigma_l)^{-1} \right] \\ q_c &= \frac{1}{d} \sum_{l=1}^K \text{tr} \left[\hat{q}_l \Sigma_l \Sigma_c (\sum_{l'} \hat{V}_{l'} \Sigma_{l'})^{-2} \right] + \rho_c \equiv R_c + \rho_c \\ m_c &= \rho_c \end{cases} \quad (120)$$

The fixed point of the equations does not depend on $\{\rho_c\}_{c \in \mathcal{C}}$ anymore and we are left with equations only for $\{\hat{V}_c, \hat{q}_c, V_c, R_c\}_{c \in \mathcal{C}}$:

$$\begin{cases} \hat{V}_c &= \frac{\alpha p_c}{1+V_c} \\ \hat{q}_c &= \frac{\alpha p_c}{(1+V_c)^2} (\Delta + R_c) \end{cases} \quad (121)$$

$$\begin{cases} V_c &= \frac{1}{d} \text{tr} \left[\Sigma_c (\sum_{l \in \mathcal{C}} \hat{V}_l \Sigma_l)^{-1} \right] \\ R_c &= \frac{1}{d} \sum_{l=1}^K \text{tr} \left[\hat{q}_l \Sigma_l \Sigma_c (\sum_{l'} \hat{V}_{l'} \Sigma_{l'})^{-2} \right] \end{cases} \quad (122)$$

The errors can be computed in a closed form with a series of algebraic manipulation. One can verify by algebraic manipulation the following relations:

$$\sum_c V_c \hat{V}_c = \sum_c \frac{1}{d} \text{tr} \left[\hat{V}_c \Sigma_c (\sum_{l \in \mathcal{C}} \hat{V}_l \Sigma_l)^{-1} \right] = 1 \quad (123)$$

$$\sum_c R_c \hat{V}_c - V_c \hat{q}_c = \frac{1}{d} \sum_{l=1}^K \text{tr} \left[\hat{q}_l \Sigma_l (\sum_c \hat{V}_c \Sigma_c) (\sum_{l'} \hat{V}_{l'} \Sigma_{l'})^{-2} \right] - \frac{1}{d} \text{tr} \left[\sum_c \hat{q}_c \Sigma_c (\sum_{l \in \mathcal{C}} \hat{V}_l \Sigma_l)^{-1} \right] = 0 \quad (124)$$

Now we express everything in eq. (123) in terms of the non-hatted overlaps to obtain:

$$\star) \sum_c \frac{\alpha p_c}{1+V_c} V_c = 1 \quad (125)$$

$$\#) \sum_c \frac{\alpha p_c}{(1+V_c)^2} R_c = \sum_c \frac{\alpha p_c}{(1+V_c)^2} V_c \Delta \quad (126)$$

We remark that we can write the training and generalization error can be written as:

$$\varepsilon_{\text{tr}} = \sum_c \frac{\hat{q}_c}{\alpha} \quad \varepsilon_{\text{gen}} = \sum_c \frac{\hat{q}_c^* (1+V_c^*)^2}{\alpha} \quad (127)$$

So if we compute the quantity $\sum_c \hat{q}_c$ we conclude. By plugging in (#) into the expression above we obtain:

$$\sum_c \hat{q}_c = \sum_c \frac{\alpha p_c}{(1+V_c)^2} (R_c + \Delta) \quad (128)$$

$$\stackrel{(\#)}{=} \Delta \sum_c \frac{\alpha p_c}{(1+V_c)^2} + \Delta \sum_c \frac{\alpha p_c}{(1+V_c)^2} V_c \quad (129)$$

$$= \Delta \sum_c \frac{\alpha p_c}{1+V_c} \quad (130)$$

and finally using relation (\star) we prove the theorem:

$$\varepsilon_{\text{tr}} = \sum_c \frac{\hat{q}_c^*}{\alpha} \quad (131)$$

$$= \sum_c \frac{\hat{q}_c^*}{\alpha} \pm \frac{\Delta}{\alpha} \quad (132)$$

$$\stackrel{(\star)}{=} \Delta \left(1 - \frac{1}{\alpha}\right) \quad (133)$$

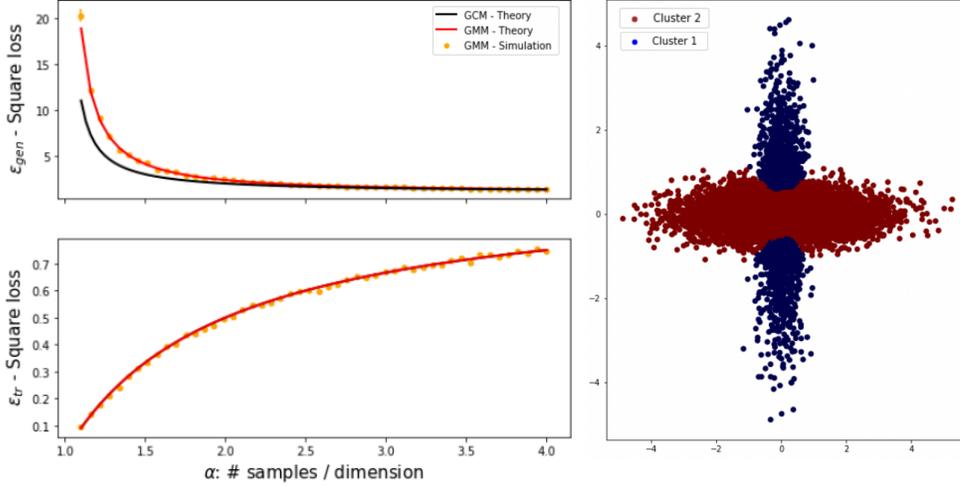


Figure 5. Training and generalization error for ridge regression on the GMM defined in eq. (134). **Left:** We compare the performance of the Gaussian asymptotics (solid black line) with the GMM one (solid red line), while the orange dots represent simulations that agree with the theoretical predictions predicted by Theorem 3.1. **Right:** Two dimensional toy plot of 10000 samples coming from the heterogeneous Gaussian mixture.

On the other hand, the generalization error does not respect the strong universality statement as we analyze in the next section (See Fig. 5).

C. Non-universality of Gaussian Mixture Models

In the previous section, we enumerated a series of results unveiling universality of GMM. Now we want to study the dual task: when GMMs *are not* universal? We have two ways to break universality: a) allow strong heterogeneity in the data structure; b) consider labels that are strongly correlated with the data structure. The plan for this section is to review more in detail these two processes for universality breaking.

C.1. Strongly Heterogeneous Mixtures

We proved in Theorem 3.6 a strong universality statement for the training loss of ridge regression. However, in this section, we want to clarify that the theorem is not valid beyond its assumption for general mixtures. We prove this by analyzing a counterexample: a strongly heterogeneous 2-cluster Gaussian Mixture:

$$\Sigma_+ = \text{diag}(0.1, \dots, 0.1, 1.9, \dots, 1.9) \quad p_+ = 0.8 \quad (134)$$

$$\Sigma_- = \text{diag}(1.9, \dots, 1.9, 0.1, \dots, 0.1) \quad p_- = 0.2 \quad (135)$$

We present in Fig. 5 the comparison of the heterogeneous GMM performance with the Gaussian theory: although the training errors coincide as predicted by Theorem 3.6, the generalization errors are different. However, we remark that real data after preprocessing seem homogeneous enough to obtain a good agreement with the exact Gaussian asymptotics as we see in Fig. 1.

C.2. Correlated Teachers

In this section, we investigate in deeper detail the controlled setting of Theorem 3.7. We do not necessarily consider estimation problems that respect the assumption of mean universality property (Prop. 3.2), and in fact, we precisely analyze the consequences if we relax these conditions. We perform ridge regression on a two-mixture GMM with opposite means and the same covariance matrix proportional to I_d . In this scenario, the exact asymptotics for the performance of the ERM estimator admits a closed-form expression. We remind here the main parameters for the theoretical analysis are:

$$\rho = \frac{1}{d} \boldsymbol{\theta}_0^\top \Sigma \boldsymbol{\theta}_0 = \frac{1}{d} \|\boldsymbol{\theta}_0\|_2^2 \quad \gamma = \frac{1}{d} \|\boldsymbol{\mu}\|_2^2 \quad \Delta = \mathbb{E}_\xi \xi^2 \quad \hat{\boldsymbol{\mu}}_\pm = \hat{h}_\pm \boldsymbol{\mu}_\pm + \hat{m}_\pm \Sigma \boldsymbol{\theta}_0 \quad (136)$$

In this simple setting we can simplify the general mixture equations in eqs. (83),(84) without doing any assumption on the teacher. We first map the equation to a single Gaussian problem $\mathcal{G}_{\theta, I_d, \mu}$. We have to slightly modify the proof of Thm. 3.3 presented in Sec. B.2, indeed we cannot blindly assume that the mean overlaps $\{h_+, h_-\}$ are zero but one can show that at the fixed point the overlaps respect:

$$\hat{h}_+^* = -\hat{h}_-^*, \quad \hat{m}_+^* = \hat{m}_-^* \quad \rightarrow \quad \hat{\mu}_+^* = \hat{\mu}_-^* \quad (137)$$

$$\hat{V} = \hat{V}_+ + \hat{V}_- \quad \hat{q} = \hat{q}_+ + \hat{q}_- \quad \hat{m} = \hat{m}_+ + \hat{m}_- \quad \hat{h} = \hat{h}_+ - \hat{h}_- \quad (138)$$

$$V = V_+ = V_- \quad q = q_+ = q_- \quad m = m_+ = m_- \quad h = h_+ = -h_- \quad (139)$$

the fixed point of the replica equations for the mixture defined above are mapped to the one of a single Gaussian problem for the overlaps $\{V, m, q, h, \hat{V}, \hat{m}, \hat{q}, \hat{h}\}$. Now we can simplify them even further by plugging in the assumption on the covariance, all the traces simplify and we get:

$$\begin{cases} V &= \frac{1}{\lambda + \hat{V}} \\ q &= V^2 \left(\hat{q} + \rho \hat{m}^2 + 2\pi \hat{m} \hat{h} + \gamma \hat{h}^2 \right) \\ m &= V \left(\rho \hat{m} + \pi \hat{h} \right) \\ h &= V \left(\gamma \hat{h} + \pi \hat{m} \right) \end{cases} \quad (140)$$

These equations are actually solvable! Define

$$\eta := V\hat{V} = \frac{\alpha V}{1 + V};$$

Now, we know that the generalization error satisfies:

$$V^2 \hat{q} = \frac{V^2 \alpha}{(1 + V)^2} \varepsilon_{\text{gen}} = \frac{\eta^2}{\alpha} (\varepsilon_{\text{gen}}).$$

We can plug this into the equation for q to get

$$\varepsilon_{\text{gen}} = \rho + \underbrace{\left(\frac{\eta^2}{\alpha} (\varepsilon_{\text{gen}}) + \rho (V\hat{m})^2 + 2\pi (V\hat{m})(V\hat{h}) + \gamma (V\hat{h})^2 \right)}_q - 2 \underbrace{(\rho V\hat{m} + \pi V\hat{h})}_m + (\pi - h)^2 \quad (141)$$

$$\varepsilon_{\text{tr}} = \frac{\varepsilon_{\text{gen}}}{(1 + V)^2} \quad (142)$$

with the relations:

$$V = \frac{\eta}{\alpha - \eta}, \quad V\hat{m} = \eta, \quad V\hat{h} = \eta(\pi - h)$$

and the expression of $\pi - h$ as a function of η :

$$\pi - h = \pi - V(\gamma\hat{h} + \pi\hat{m}) = \pi - \gamma\eta(\pi - h) - \pi\eta \iff \pi - h = \pi \frac{1 - \eta}{1 + \gamma\eta} \quad (143)$$

we simplify everything in terms only of η to get:

$$\varepsilon_{\text{gen}} = \frac{\alpha (\Delta + (\eta - 1)^2 \rho)}{(\alpha - \eta^2)} \left(1 - \pi^2 \frac{(\eta - 1)^2 (\gamma\eta^2 + 2\eta - 1)}{(\Delta + (\eta - 1)^2 \rho) (1 + \gamma\eta)^2} \right) \quad (144)$$

$$\varepsilon_{\text{tr}} = \frac{(\alpha - \eta)^2 (\Delta + (\eta - 1)^2 \rho)}{\alpha (\alpha - \eta^2)} \left(1 - \pi^2 \frac{(\eta - 1)^2 (\gamma\eta^2 + 2\eta - 1)}{(\Delta + (\eta - 1)^2 \rho) (1 + \gamma\eta)^2} \right) \quad (145)$$

All that remains is to solve for η using the equations for V and \hat{V} , which yields

$$\eta = 1 + \frac{1}{2} \left(\alpha - 1 + \lambda - \sqrt{4\lambda + (\alpha - 1 + \lambda)^2} \right). \quad (146)$$

Vanishing regularization We can take $\lambda = 0$ inside (146) even when $\alpha < 1$, and we find

$$\eta(\lambda = 0) = \min(\alpha, 1) \quad (147)$$

Finally, this yields

$$\varepsilon_{gen} = \begin{cases} \frac{\alpha\Delta}{\alpha-1} & \alpha \geq 1 \\ \frac{\pi^2(\alpha-1)(\alpha^2\gamma+2\alpha-1)}{(\alpha\gamma+1)^2} - \frac{(\alpha-1)^2\rho+\Delta}{\alpha-1} & \text{else} \end{cases} \quad (148)$$

$$\varepsilon_{train} = \begin{cases} \frac{(\alpha-1)\Delta}{\alpha} & \alpha \geq 1 \\ 0 & \text{else} \end{cases} \quad (149)$$

We retrieve the results of Corollary. 3.8: even with correlated teachers we show that in the underparametrized regime we have Gaussian universality.

Extension of Corollary. 3.8: The Gaussian universality result for correlated teachers at vanishing regularization can be extended as well to more general balanced mixtures with homoscedastic covariance, i.e. the case where

$$p_c = \frac{1}{|\mathcal{C}|}, \quad \Sigma_c = \Sigma, \quad \lambda = 0.$$

In order to do so, we consider the general saddle point equation for ridge regression in eq. (90),(91), and plug in the assumptions:

$$\begin{cases} \hat{V} &= \frac{\alpha}{|\mathcal{C}|(1+V)} \\ \hat{q}_c &= \frac{\alpha}{|\mathcal{C}|(1+V)^2} (\rho + \Delta + q_c - 2m + (h_c - \pi_c)^2) \\ \hat{m} &= \frac{\alpha}{|\mathcal{C}|(1+V)} \\ \hat{h}_c &= \frac{\alpha(\pi_c - h_c)}{|\mathcal{C}|(1+V)} \end{cases} \quad (150)$$

$$\begin{cases} V &= \frac{1}{d} \text{tr} \left[\Sigma \hat{\Sigma}^{-1} \right] \\ q_c &= \frac{1}{d} \text{tr} \left[\left(\hat{q}_c \Sigma + \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top \right) \Sigma \hat{\Sigma}^{-2} \right] \\ m &= \frac{1}{d} \text{tr} \left[\hat{\boldsymbol{\mu}} \boldsymbol{\theta}_0^\top \Sigma \hat{\Sigma}^{-1} \right] \\ h_c &= \frac{1}{d} \text{tr} \left[\hat{\boldsymbol{\mu}} \boldsymbol{\mu}_c^\top \hat{\Sigma}^{-1} \right] \end{cases} \quad (151)$$

In particular, these assumptions imply that $V_c, \hat{V}_c, m_c, \hat{m}_c$ do not depend on the class labels c , and hence

$$\hat{\Sigma} = \hat{V} \Sigma.$$

Now, consider the assignment

$$\hat{h}_c = 0 \quad \text{and} \quad h_c = \pi_c; \quad (152)$$

it is easy to check that they are a fixed point of the above saddle-point equations, since then

$$h_c = \frac{\hat{m} \pi_c}{\hat{V}} = \pi_c$$

Plugging in this relation in the expression of the update for $\{\hat{q}_c\}_{c \in \mathcal{C}}$ we obtain:

$$\hat{q}_c^* = \frac{\alpha p_c}{(1+V^*)^2} (\rho_c + \Delta + q_c^* - 2m_c^* + (h_c^* - \pi_c)^2) = \frac{\alpha p_c}{(1+V^*)^2} (\rho_c + \Delta + q_c^* - 2m_c^*) \quad (153)$$

Hence by looking at the expression of generalization and training error for ridge regression in eq. (127) we conclude that they will not depend on the values of the mean overlaps $\{h_c^*\}_{c \in \mathcal{C}}$ and we can prove Gaussian universality in the same way as we did in the proof of Theorem 3.4.

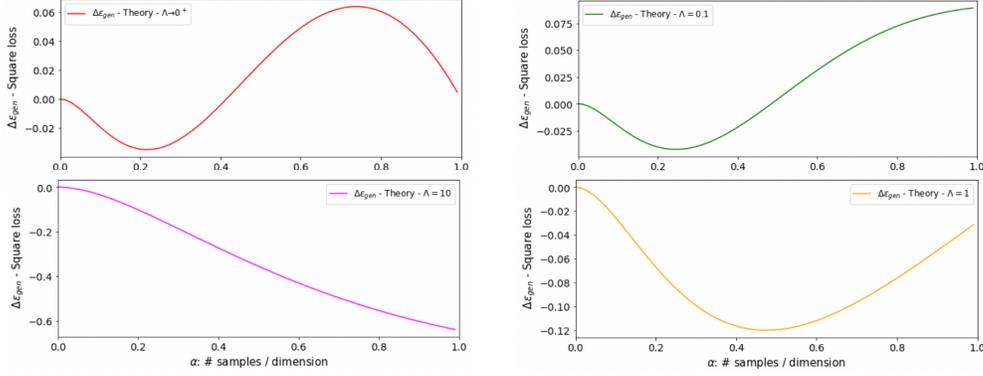


Figure 6. Theoretical prediction from eq. (160) for the difference between the generalization error computed from the Gaussian theory and GMM one plotted as a function of the number of samples used in the fit in Algorithm. C.3. We use different regularization strengths $\Lambda \in \{1e - 10, 0.1, 1, 10\}$, increasing clockwise in the figure.

C.3. Interpolating Teachers

Different works observed over different datasets that if we find an interpolating teacher vector, in such a way that we can keep the real labels to study ERM performance, the theoretical predictions coming from Gaussian asymptotics would agree with simulations, (Loureiro et al., 2021a) among others. This observation apparently contradicts the result of Theorem 3.7 in which we show that correlated labels (as we expect real labels to be) break Gaussian universality. We try to motivate this in the same setting as Theorem 3.7:

Theorem C.1. (Real labels universality) Consider the same setting of Theorem 3.7. Fix now the teacher vector to be the maximally correlated one $\theta = \mu$ and find an interpolating teacher $\hat{\theta}$ using \tilde{n} samples coming from the GMM. If we compare the exact asymptotics of Gaussian and GMM for regularization Λ at fixed $\alpha = \frac{\tilde{n}}{d}$ we obtain a discrepancy:

$$\Delta \varepsilon_{gen} = \frac{4(E - 1)\alpha^2 (E^2\gamma + 2e - 1)}{(\alpha + 1)^2(E\gamma + 1)^2} \quad (154)$$

with $E(\alpha, \Lambda)$ solution of:

$$E = 1 + \frac{1}{2} \left(\alpha - 1 + \Lambda - \sqrt{4\Lambda + (\alpha - 1 + \Lambda)^2} \right). \quad (155)$$

The result above solves the apparent contradiction: if we strongly overparametrize the fitting model, interpolating teachers can be uncorrelated with the data structure.

Algorithm 1 Get theoretical learning curves keeping real labels

Data generation: Generate the data from the 2-clusters GMM described in C.1 and set:

$$\mu_+ = -\mu_- = \sqrt{\gamma}(1, \dots, 1)^\top, \quad \Sigma_+ = \Sigma_- = I_d, \quad \theta_0 = \mu_+ \quad (156)$$

Fit interpolating teacher Perform ridge regression for $\lambda \rightarrow 0$. Take $\alpha = \frac{n}{d} < 1$ so that the ERM estimator will interpolate the data.

Run exact asymptotics Compute the theoretical performance of the associated 2-GMM and equivalent GCM from the estimated $\hat{\theta}$ for the same α with a fixed regularization parameter Λ .

The proof of the results relies strongly on the fact that we can analyze analytically the performance in this scenario, being the same as the previous section. Reminding from eq. (147) that in the limit $\lambda \rightarrow 0$ we have in the overparametrized setting

$\eta(\alpha, \lambda = 0) = \alpha$, we can write:

$$h_{erm} = \frac{\boldsymbol{\mu}^\top \hat{\boldsymbol{\theta}}}{d} = \gamma\alpha \frac{\gamma + 1}{1 + \gamma\alpha} \quad (157)$$

$$q_{erm} = \frac{\hat{\boldsymbol{\theta}}^\top \Sigma \hat{\boldsymbol{\theta}}}{d} = \alpha \left(-\frac{\Delta}{\alpha - 1} - \frac{(\alpha - 1)\pi^2}{\alpha\gamma + 1} + \rho \right) \quad (158)$$

We summarize the algorithmic procedure in Algorithm. C.3. In order to compute the generalization error from eq. (144) we need the quantities $\{\rho, \pi\}$. Note that in Algorithm. C.3 we compute the theoretical curve for the GMM and the GCM from the estimated $\hat{\boldsymbol{\theta}}$. This translates into the fact that the estimated overlaps from numerical simulations in the previous step become the equivalent to:

$$\pi = h_{erm} \qquad \rho = q_{erm} \quad (159)$$

We fix the regularization for the theoretical prediction to be Λ , and for simplicity, we use the same sample complexity parameter α . We are now in a position to compare Gaussian and GMM theoretical prediction, more precisely we analyze the difference in the generalization errors for the two data models. Recalling the expression for the generalization error in eq. (144), we just need to plug the estimated overlaps in eq. (159) to obtain:

$$\Delta\varepsilon_{\text{gen}} = \frac{4\alpha^2(E - 1)(E^2\gamma + 2e - 1)}{(\alpha + 1)^2(E\gamma + 1)^2} \quad (160)$$

where $E(\alpha, \Lambda)$ is the solution of eq. (146), retrieving the result in Theorem C.1. For the sake of clarity, we stated the theorem in a simple setting, indeed once we estimated $\hat{\boldsymbol{\theta}}$ we could have decided to change the value of the sample complexity, as we are now interested in running theory curves. However, to simplify the equations we assumed to run the theoretical prediction for the same value of α as in the ERM fit. We present the results in Fig. 6 for $\gamma = 1$: although the true labels are maximally correlated with the data structure, we see that the Gaussian theoretical predictions with the interpolating teacher $\hat{\boldsymbol{\theta}}$ are very close to the GMM one. We remark that the upper-left plot in Fig. 1 is a nice characterization of Theorem (3.8): as we reach the underparametrized regime we restore Gaussian universality for $\Lambda \rightarrow 0^+$ even for correlated teachers.

D. Details on Real Dataset Simulations

In this section, we report the procedure to create the random regression task on real data, see Algorithm. D. The implementation of the different numerical simulations described in this work is available in the [GitHub repository](#).

Algorithm 2 Random teacher regression on real data

Data processing: Load data and perform the preprocessing step with a transform matrix $F \in \mathbb{R}^{d \times d'}$, with d' dimension of the images in the chosen dataset. We choose for all the figures in this manuscript $d = 2000$.

Match covariance Compute $\hat{\Sigma} = \frac{1}{n} X X^\top$

Create new labels Forget the real labels associated with the dataset and create new label according to:

$$y_\nu = \begin{cases} \frac{\theta_0^\top \mathbf{x}_\nu}{\sqrt{d}} + \sqrt{\Delta} \xi & \text{Ridge regression} \\ \text{sign} \left(\frac{\theta_0^\top \mathbf{x}_\nu}{\sqrt{d}} + \sqrt{\Delta} \xi \right) & \text{Logistic regression} \end{cases} \quad (161)$$

$$\theta_0 \sim \mathcal{N}(\mathbf{0}, I_d) \quad \xi \sim \mathcal{N}(0, 1) \quad (162)$$

Run learning curves Fix the regularization parameter λ .

for α in a given range **do**

Simulation Solve the ERM in eq. (1) using sklearn package *LogisticRegression* (Pedregosa et al., 2011) or the Moore-Penrose pseudo-inverse for the ridge estimator:

$$\hat{\theta} = \begin{cases} X^\top (X X^\top + \lambda I_n)^{-1} Y & n < d \\ (X^\top X + \lambda I_d)^{-1} X^\top y & n > d \end{cases} \quad (163)$$

Gaussian theory Solve saddle point equations in eqs. (83),(84) defining Gaussian asymptotics for the model. 2.1.

Compute errors Compute the training loss and generalization error using the metrics:

$$g(y, \hat{y}) = \begin{cases} (y - \hat{y})^2 & \text{Ridge regression} \\ \mathbb{P}(y \neq \hat{y}) & \text{Logistic regression} \end{cases} \quad (164)$$

end for
