Gaussian universality of perceptrons with random labels

Federica Gerace [1,2] Florent Krzakala [3] Bruno Loureiro [3,4] Ludovic Stephan [3] and Lenka Zdeborová [2]

[1] *International School of Advanced Studies (SISSA), Trieste, Via Bonomea, 265, 34136 Trieste, Italy*

[2] *EPFL Statistical Physics of Computation (SPOC) Laboratory, Rte Cantonale, 1015 Lausanne, Switzerland*

[3] *EPFL, Information, Learning and Physics (IdePHICS) Laboratory, Rte Cantonale, 1015 Lausanne, Switzerland*

[4] *Département d'Informatique, École Normale Supérieure (ENS)–PSL & CNRS, F-75230 Paris Cedex 05, France*

(Received 3 March 2023; revised 18 September 2023; accepted 8 December 2023; published 8 March 2024)

While classical in many theoretical settings—and in particular in statistical physics-inspired works—the assumption of Gaussian i.i.d. input data is often perceived as a strong limitation in the context of statistics and machine learning. In this study, we redeem this line of work in the case of generalized linear classification, also known as the perceptron model, with random labels. We argue that there is a large universality class of high-dimensional input data for which we obtain the same minimum training loss as for Gaussian data with corresponding data covariance. In the limit of vanishing regularization, we further demonstrate that the training loss is independent of the data covariance. On the theoretical side, we prove this universality for an arbitrary mixture of homogeneous Gaussian clouds. Empirically, we show that the universality holds also for a broad range of real data sets.

DOI: 10.1103/PhysRevE.109.034305

I. INTRODUCTION

Statistical physics studies of artificial neural networks have a long history, including many works that continue to have an impact on the current investigations of deep neural networks. A large fraction of this ongoing work has focused on Gaussian input data; see [1–3] for some of the earliest and most influential examples. However, the Gaussian data assumption is not limited to works from statistical physics of learning. Indeed, it is a widespread assumption in the high-dimensional statistics literature, where it is also known under the umbrella term of *Gaussian design*; see, for example, Refs. [4–6]. Despite being both common and convenient for doing theory, i.i.d. Gaussian data might come across as a stringent limitation at first glance, out of touch with the real-world practice in which data are structured. Indeed, an important branch of statistical learning theory is data-agnostic and avoids making overly specific assumptions on the data distribution [7]. However, a number of recent observations (both heuristic and rigorous) suggest that the Gaussian assumption is not always that far-fetched for high-dimensional data (see, for instance, [8–12] and references therein). The goal of the present work is to redeem the Gaussian hypothesis for perhaps the simplest, yet deeply fundamental, problem of high-dimensional statistics, namely the perceptron problem, also known as generalized linear classification, with random labels.

Models with random labels are ubiquitous in the theory of machine learning. The problem of how many randomly labeled patterns a perceptron model can fit, known as the *storage capacity problem*, is at the root of the historical interest of the statistical physics community for machine learning problems. Indeed, works on this classical subject span more than four decades [1,2,13–19]. Interestingly, it was noted in these early works that the asymptotic storage capacity of the perceptron is equivalent for both binary and Gaussian input data. This observation is one of the early manifestations of universality in this literature. The interest in random labels is also not bound to the statistical physics of learning community. They appear in several contexts in statistical learning theory, such as in the definition of Rademacher complexities [7,20], in the pioneering studies of Wendel and Cover [21,22], or in thought-provoking numerical experiments with deep learning [23,24],

In this work, we ask the following: How would these theories for random labels change if we used a realistic data set instead of a Gaussian one? We consider the training loss of generalized linear classifiers (perceptrons) trained on random labels, including ridge, hinge, and logistic classification [25], but also kernel methods [26] and neural networks trained in the lazy regime [27] (the so-called neural tangent kernel [28]), as well as with engineered features such as the scattering transform [29]. We focus on the thermodynamic limit (known as the high-dimensional setting in statistics) where both n (the number of training samples) and p (the input dimension) go to infinity at a fixed rate $\alpha = n/p$.

Our main result is to argue that in the aforementioned setting with random labels, many input data distributions actually have the same learning properties as Gaussian data, thus providing a rather surprising Gaussian universality result for this problem. In particular, the minimum training loss for a wide range of settings is the same as that of a corresponding Gaussian problem with matching data covariance. Furthermore, in the limit of vanishing regularization, we show that Gaussian universality is even stronger, as the minimum training loss is independent of the data covariance (and therefore the same as that of i.i.d. Gaussian data). In other words: as far as random labels are concerned, it turns out that the theoretical results derived under the Gaussian data assumptions capture what is actually happening in practice. Certainly, the value of the interpolation (or capacity) threshold

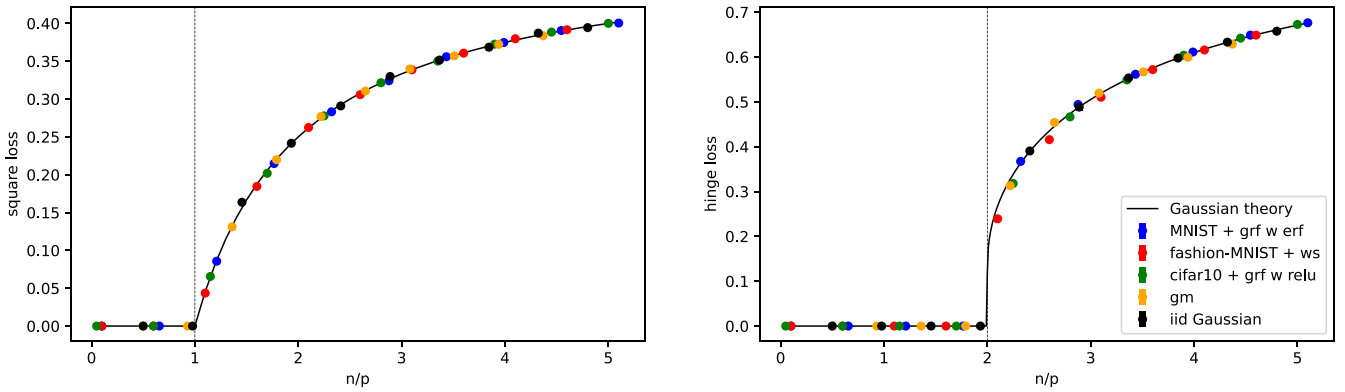

FIG. 1. Training loss as a function of the number of samples n per input dimension p at regularization $\lambda = 10^{-15}$. Left panel: the square loss; right panel: the hinge loss. The black solid line represents the outcome of the replica calculation for i.i.d. Gaussian inputs, namely when the covariance matrix Σ corresponds to the identity matrix. Dots refer to numerical simulations on different full-rank data sets. In particular, blue dots correspond to MNIST with Gaussian random features and error function nonlinearity, red dots correspond to fashion-MNIST with wavelet scattering transform, green dots correspond to CIFAR10 in grayscale with Gaussian random features and ReLU nonlinearity, yellow dots correspond to a mixture of Gaussians, with means $\boldsymbol{\mu}_{\pm} = (\pm 1, 0, \ldots, 0)$, covariances $\Sigma_{\pm}$ both equal to the identity matrix, and relative class proportions $\rho_{\pm} = \frac{1}{2}$. Finally, black dots correspond to i.i.d. Gaussian inputs.

was known to be universal and occurs (for full-rank data) at $n = p$ for ridge regression, and at $n = 2p$ for linear classifiers (perceptrons) [22] (in fact, Cover's result is stronger, since it establishes the universality of the number of hyperplanes separating the inputs). However, the fact that the loss itself is universal is a stronger statement that redeems an entire line of work using the Gaussian data assumption, and in particular a large part of those from statistical physics of learning.

Summary of main results

The main points of the present work can be summarized by Figs. 1 and 2, which show the training loss of real-world data sets trained with random labels and various feature maps, compared with the (analytical) prediction derived for Gaussian data with matching covariance. The code used to run these experiments is publicly available in a GitHub repository [30]. As illustrated in these plots, Gaussian universality seems to hold even for finite-dimensional data, and for actual real data sets. Notably, we observe that when using random labels, the training losses plotted as a function of the ratio between the number of samples and the dimension $\alpha = n/p$ are indistinguishable from results obtained for Gaussian input data when using MNIST [31], fashion-MNIST [32], or CIFAR10 [33] preprocessed through various standard feature maps. This conclusion seems robust and holds for different features of the raw data, such as random features [34] or the convolutional scattering transform [29,35]. It also holds, as we prove, if we simply use a synthetic Gaussian mixture model, a classical model for complex multimodel data. The agreement between the real world and the asymptotic Gaussian theory is striking. While we may expect that such data could be approximated by a multimodal distribution such as a Gaussian mixture with enough modes, it is rather puzzling that they lead to the same loss as a single Gaussian cloud. Our main contribution is to provide a rigorous theoretical foundation for these observations that vindicates the classical line of works on Gaussian design, in particular the one stemming from statistical physics.

Our main results are as follows:

(a) We provide a strong universality theorem for linear interpolators corresponding to ridgeless regression (with vanishing regularization) in high dimensions and random labels, Theorem 5. Informally, we prove that a perceptron trained on randomly labeled Gaussian mixture data (a setting that encompasses complex multimodal distributions) has the same minimum asymptotic loss as a perceptron trained on randomly labeled Gaussian data with isotropic covariance, that is, $\mathcal{E}_\ell(\alpha) = \frac{1}{2}(1 - \frac{1}{\alpha})_+$. This provides a theoretical explanation for the phenomena illustrated in Fig. 1 (left).

(b) Under an additional homogeneity assumption on the different modes of the data, Gaussian universality can be generalized to *any convex loss* (and we conjecture that it is valid for nonconvex losses as well), Theorem 4. This provides a theoretical explanation of the phenomena illustrated in Fig. 1 (right).

(c) At finite regularization and under the same homogeneity assumption, we show that the asymptotic training loss depends solely on the *data covariance matrix*, such that it is, again, the same loss as that of a single Gaussian cloud with matching covariance, Theorem 3. This is illustrated in Fig. 2.

The proof technique used to establish these universality theorems is of interest on its own. It builds on recent progress in high-dimensional statistics and in mathematical insights drawn from the replica method in statistical physics. In particular, we provide an *explicit* matching of the expression (obtained from a rigorous proof of the replica prediction) for the asymptotic minimal loss [12,36–38]. We further demonstrate the strong universality for ridge regression with vanishing regularization, again by showing explicitly that the exact solution [12,39,40] reduces to one of the homogeneous Gaussian cases. These results are obtained through methods that have been developed from statistical physics and mathematical physics-inspired techniques.

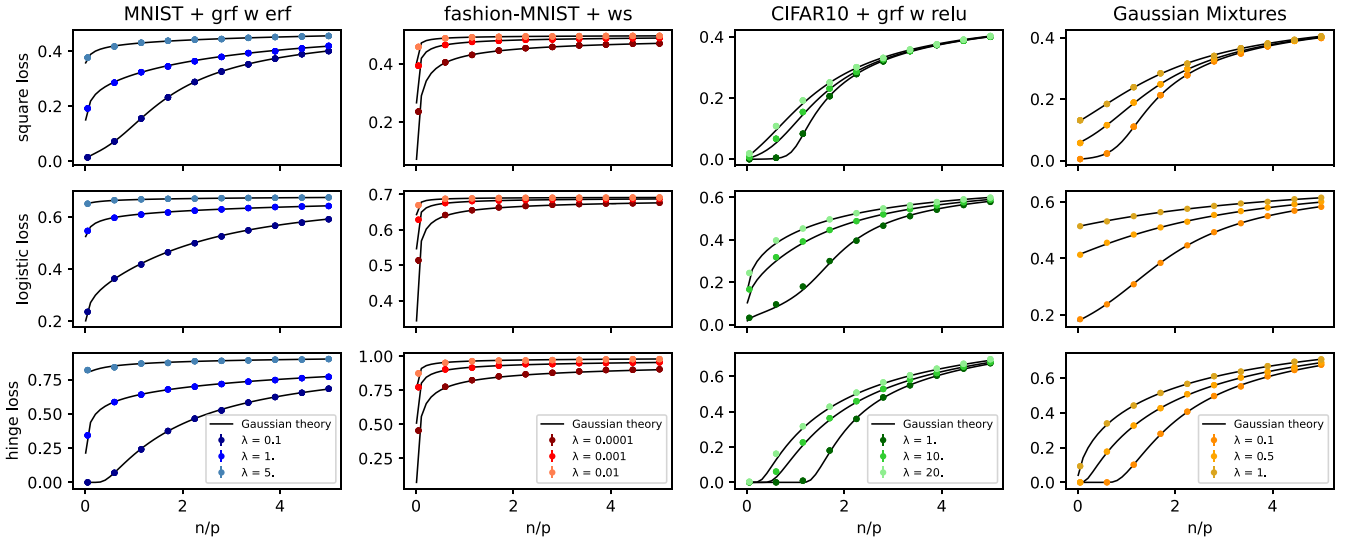

FIG. 2. This figure shows the training loss as a function of the number of samples n per dimension p at finite regularization λ. Top panel: the square loss; bottom panel: the hinge loss. The first column refers to MNIST with Gaussian random features and error function nonlinearity, the second column corresponds to fashion-MNIST with wavelet scattering transform, the third column corresponds to CIFAR10 in grayscale with Gaussian random features and ReLU nonlinearity, and the fourth column corresponds to a mixture of Gaussians, with means $\boldsymbol{\mu}_\pm = (\pm 1, 0, \ldots, 0)$, and covariances $\Sigma_\pm$ both equal to the identity matrix and relative class proportions $\rho_\pm = \frac{1}{2}$. Black solid lines correspond to the outcome of the replica calculation, obtained by assigning to Σ the covariance matrix of each data set plus the corresponding transformation. The colored dots correspond to the simulations for different values of λ, as specified in the plot legend. Simulations are averaged over 10 samples, and the error bars are not visible at the plot scale.

Further related work

The perceptron. The question of how many samples can be perfectly fitted by a linear model is a classical one. For a ridge classifier, it amounts to asking whether a linear system of n equations with p unknowns is invertible so that for full-rank data the transition arises at $n = p$. For the $0/1$ loss or its convex surrogate such as the hinge loss, the question of linear separability was famously discussed by Cover [22], who showed that for full-rank data the transition is given by $n = 2p$. In both cases, the transition is universal and does not depend on details of the data distribution (provided it is full rank, otherwise the rank replaces the dimension). The capacity problem for a multiclass perceptron has been studied in [41], where it was shown that for a specific rule it is independent of the number of classes. For Gaussian data, such questions have gotten a large amount of attention in the statistical physics community [1,2,13,14,42,43] but also recently in theoretical computer science [15–19]. In a complementary line, there has been significant effort to extend these results to structured data models [44–47]. It is one of our goals to attract attention to these works, given that the Gaussian universality we present shows that their relevance is not limited to idealistic Gaussian data.

Random labels. Random labels are a fundamental and useful concept in machine learning. The pioneering work of Ref. [23], for instance, was instrumental in the modern critics of classical measures of model complexity, including the Rademacher complexity or the VC dimension. These considerations have driven an entire line of research aiming to find substantial differences between learning with true and random labels, for instance in training time [48–50], in minima sharpness [51,52], or in what neural networks can actually learn with random labels [24]. It has also been recently claimed [24] that pretraining on random labels under a given initial condition scaling can consistently speed up neural network training on both true and random labels, with respect to training from scratch.

Gaussian universality. There has been much progress on a similar, though more restricted, Gaussian universality for random feature maps on Gaussian input data [34]. Following early insights in Ref. [53], the authors of Refs. [54,55] showed that the empirical distribution of the Gram matrix of random features is asymptotically equivalent to a linear model with matched covariance. This was extended to generic convex losses in Ref. [56] using the heuristic replica method, and proven in Ref. [57]. A specific *Gaussian equivalence principle* [8] for learning with random features has been proven in a succession of works for convex penalties in Refs. [58,59] and some nonconvex ones in Ref. [60]. Early ideas on Gaussian universality have also appeared in the context of signal processing and compressed sensing in Refs. [4,61–64]. These theoretical results, however, fall short when considering realistic data sets as we do in this work. Indeed, these previous works considered only unimodal Gaussian data (observed through random feature maps), a situation far from realistic multimodal, complex, real-world data sets. Instead, the authors of Refs. [9,65,66] argued that real data sets can be efficiently approximated in high dimensions by a finite *mixture of Gaussians*. These, of course, are multimodal distributions that cannot be approximated by a single Gaussian. Gaussian mixtures will be the starting point of our theory.

Finally, we note that the observation that Gaussian data can fit or represent well some real data has been heuristically observed in many situations, but without theoretical

justification and often limited to ridge regression; see, e.g., Refs. [10,12,67–69].

II. SETTING, NOTATION, AND ASYMPTOTIC FORMULAS

The focus of the present work is the analysis of high-dimensional binary linear classification (also known as perceptron) on a data set $\mathcal{D} = \{(\boldsymbol{x}_\mu, y_\mu)\}_{\mu=1}^n$. We shall consider a minimization problem of the form

$$\widehat{\mathcal{R}}_n^*(\boldsymbol{X}, \boldsymbol{y}) = \inf_{\boldsymbol{\theta}\in\mathbb{R}^p} \frac{1}{n}\sum_{\mu=1}^n \ell(\boldsymbol{\theta}^\top \boldsymbol{x}_\mu, y_\mu) + \frac{\lambda}{2}||\boldsymbol{\theta}||_2^2, \tag{1}$$

where the $\boldsymbol{x}_\mu \in \mathbb{R}^p$ are input vectors, and $y_\mu \in \{-1, +1\}$ are binary labels. We assume that the loss ℓ only depends on the inputs $\boldsymbol{x}_\mu$ through a one-dimensional projection $\boldsymbol{\theta}^\top \boldsymbol{x}_\mu$ with θ being the perceptron learning weights, and we work in the so-called *thermodynamic* or *proportional high-dimensional limit*, where n, p go to infinity with

$$\frac{n}{p} \to \alpha > 0.$$

In practice, practitioners seldom use the raw data $\mathbf{x}$ directly in their linear classifiers and usually perform a preprocessing step. For instance, instead of using the raw MNIST, a classical approach is to use a fixed feature map, and to observe the data as $\mathbf{x} = \sigma(F\mathbf{x})$, with F a random matrix. This is called the random feature map [34], and it has the advantage, among others, that the effective data $\mathbf{x}$ are full-rank. One may use a more complicated approach such as the convolutional scattering transform [13,29], or even pretrained neural networks, a situation called transfer learning [70,71]. We shall apply such transforms to our real data as well in order to avoid theoretical pitfalls in direct space (in images some pixels are always zero, for instance, so that the data may not even be full-rank). There is another advantage of using fixed features: they correspond to deep learning (with actual multilayer nets) in the so-called lazy regime [27,28]. In this case, the feature matrix is a random matrix. Therefore, our results go beyond linear models and are also relevant to deep learning in the lazy regime. In our numerical experiments, we shall thus not only work with the original data (see Appendix C, and in particular Fig. 4), but also—and mainly—with random feature maps and fixed feature maps (as in Figs. 1 and 2).

For the labels, we shall focus in this work on the *random label* model, where the y_μ are independent of the inputs $\boldsymbol{x}_\mu$, and they are generated independently according to a Rademacher distribution:

$$y_\mu \sim \mathrm{Unif}(\{-1, 1\}). \tag{2}$$

In our theoretical approach, we shall use mainly two data models:

(i) The simplest one is the Gaussian covariate model (GCM), where the inputs $\boldsymbol{x}_\mu \in \mathbb{R}^p$ are independently drawn from a Gaussian distribution:

$$\boldsymbol{x}_\mu \sim \mathcal{N}(\mathbf{0}_p, \boldsymbol{\Sigma}). \tag{3}$$

The Gaussian covariate model has been the subject of much attention recently [6,12,36,39,40,55,72–76]. In particular, the asymptotic statistics of the minimizer of Eq. (1) for different models for the labels can be computed using the replica method, and rigorously proven as well. In particular, the random label limit relevant to our discussion can be obtained as a limit of the expressions derived using the replica method of statistical physics and mathematically proven in [12]. We shall use the following expressions here (see also Appendix A 3):

Theorem 1. Asymptotics of the GCM for random labels, adapted from [12], informal. Consider the minimization problem in Eq. (1), with the inputs $\boldsymbol{x}_\mu$ generated according to a Gaussian covariate model. Assume that the loss ℓ is strictly convex (or that ℓ is convex and $\lambda > 0$). Under mild regularity conditions on $\boldsymbol{\Sigma}$, as well as the loss and regularizer, we have the asymptotic training performance of the empirical risk minimizer Eq. (A2) for the random label Gaussian mixture model satisfying the scalings (A4) in the proportional high-dimensional limit as $n \to \infty$:

$$\begin{aligned}\widehat{\mathcal{R}}_n^*(\boldsymbol{X}, y(\boldsymbol{X})) &\xrightarrow{\mathbb{P}} \mathcal{E}_\ell^{\mathrm{gcm}}(\alpha, \lambda) \\ &:= \frac{1}{2}\sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[\ell(\mathrm{prox}_{V^\star \ell(\cdot,y)}(\sqrt{q^\star}\xi), y)],\end{aligned} \tag{4}$$

where $\mathrm{prox}_{\tau f(\cdot)}$ is the proximal operator associated with the loss:

$$\mathrm{prox}_{\tau\ell(\cdot,y)}(x) := \arg\min_{z\in\mathbb{R}}\left[\frac{1}{2\tau}(z-x)^2 + \ell(z, y)\right], \tag{5}$$

and the parameters $(V^\star, q^\star)$ are the (unique) fixed point of the following self-consistent equations:

$$\begin{aligned}\hat{V} &= \frac{\alpha}{2}\sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[\partial_\omega f_\ell(y, \sqrt{q}\xi, V)], \\ \hat{q} &= \frac{\alpha}{2}\sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[f_\ell(y, \sqrt{q}\xi, V)^2], \\ V &= \frac{1}{p}\mathrm{tr}\,\Sigma(\lambda \boldsymbol{I}_p + \hat{V}\Sigma)^{-1}, \\ q &= \frac{1}{p}\hat{q}\,\mathrm{tr}\,\Sigma^2(\lambda \boldsymbol{I}_p + \hat{V}\Sigma)^{-2},\end{aligned} \tag{6}$$

where $f_g(y, \omega, V) := V^{-1}(\mathrm{prox}_{V\ell(\cdot,y)}(\omega) - \omega)$.

(ii) A more generic model of data, which has the advantage of being multimodal to befit complex situations, is the Gaussian mixture model (GMM). In this case, the inputs $\boldsymbol{x}_\mu \in \mathbb{R}^p$ are independently generated as

$$\boldsymbol{x}_\mu \sim \sum_{c\in\mathcal{C}} \rho_c \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c), \tag{7}$$

where $\mathcal{C} := \{1, \ldots, K\}$ indexes the K Gaussian clouds, and $\rho_c \in [0, 1]$ is the density of points in each cloud and satisfies $\sum_{c\in\mathcal{C}} \rho_c = 1$. The analysis of Gaussian mixture models in the high-dimensional regime has been the subject of many works. The exact asymptotic expression for the minimum training loss has been derived for a range of particular cases in, among others, Refs. [77–82] and in full generality for arbitrary means and covariances in Ref. [38]. We shall thus use the random label limit of their expression in the binary classification case:

Theorem 2. Asymptotics of the GMM for random labels, adapted from [38], informal. Consider the minimization problem in Eq. (1), with the inputs $\boldsymbol{x}_\mu$ generated according to a

Gaussian mixture as in (7). Assume that the loss ℓ is strictly convex (or that ℓ is convex and $\lambda > 0$). Under mild regularity conditions on the $\boldsymbol{\mu}_c$, $\boldsymbol{\Sigma}_c$, as well as the loss and regularizer, we have the training performance of the empirical risk minimizer Eq. (A2) for the random label Gaussian mixture model satisfying the scalings (A4) given by

$$\begin{aligned}\widehat{\mathcal{R}}_n^*(\boldsymbol{X}, y(\boldsymbol{X})) &\xrightarrow{\mathbb{P}} \mathcal{E}_\ell^{\mathrm{gmm}}(\alpha, \lambda, K)\\ &\coloneqq \frac{1}{2}\sum_{c\in\mathcal{C}} \rho_c \sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}\\ &\quad \times [\ell(\mathrm{prox}_{V_c^\star \ell(\cdot,y)}(m_c^\star + \sqrt{q_c^\star}\xi), y)], \end{aligned} \tag{8}$$

where ℓ is the loss function used in the empirical risk minimization in Eq. (A2), $\mathrm{prox}_{\tau f(\cdot)}$ is the proximal operator associated with the loss:

$$\mathrm{prox}_{\tau\ell(\cdot,y)}(x) \coloneqq \arg\min_{z\in\mathbb{R}}\left[\frac{1}{2\tau}(z-x)^2 + \ell(z,y)\right], \tag{9}$$

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

$$\begin{aligned}\hat{V}_c &= \frac{\alpha}{2}\rho_c \sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[\partial_\omega f_\ell(y, m_c + \sqrt{q_c}\xi, V_c)],\\ \hat{q}_c &= \frac{\alpha}{2}p_c \sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[f_\ell(y, m_c + \sqrt{q_c}\xi, V_c)^2],\\ \hat{m}_c &= \frac{\alpha}{2}p_c \sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[f_\ell(y, m_c + \sqrt{q_c}\xi, V_c)],\\ V_c &= \frac{1}{p}\mathrm{tr}\,\Sigma_c\left(\lambda \boldsymbol{I}_p + \sum_{c'\in\mathcal{C}} \hat{V}_{c'}\Sigma_{c'}\right)^{-1},\\ q_c &= \frac{1}{p}\sum_{c'\in\mathcal{C}}\left[\mathrm{tr}\left(\hat{q}_{c'}\Sigma_{c'} + \hat{m}_c\hat{m}_{c'}\boldsymbol{\mu}_{c'}\boldsymbol{\mu}_c^\top\right)\right.\\ &\quad \left.\times \Sigma_c\left(\lambda \boldsymbol{I}_p + \sum_{c''\in\mathcal{C}} \hat{V}_{c''}\Sigma_{c''}\right)^{-2}\right],\\ m_c &= \frac{1}{p}\sum_{c'\in\mathcal{C}} \hat{m}_c\hat{m}_{c'}\left[\mathrm{tr}\boldsymbol{\mu}_{c'}\boldsymbol{\mu}_c^\top\left(\lambda \boldsymbol{I}_p + \sum_{c''\in\mathcal{C}} \hat{V}_{c''}\Sigma_{c''}\right)^{-1}\right], \end{aligned} \tag{10}$$

where $f_\ell(y, \omega, V) \coloneqq V^{-1}(\mathrm{prox}_{V\ell(\cdot,y)}(\omega) - \omega)$.

Although Eqs. (6) and (10) look cumbersome at first glance, they are simply self-consistent equations for a few scalar variables, also known as *order parameters*, *overlaps*, or *summary statistics*. Differently from the original high-dimensional problem (1) for the weights $\boldsymbol{\theta} \in \mathbb{R}^p$, the self-consistent equations for the overlaps can be efficiently solved numerically by iteration from a chosen initial condition, with strong convexity of the original problem (1) guaranteeing the convergence to the unique minimizer. The reduction of a random high-dimensional optimization problem to a set of deterministic low-dimensional self-consistent equations is a common theme in statistical physics and high-dimensional probability. Indeed, Eqs. (6) and (10) can be independently derived using different techniques from these fields, such as the replica method [83], the leave-one-out or cavity method [84,85], Gordon minimax inequalities [36], and message passing schemes [86]. As we will show in the section that follows, valuable analytical insights can be drawn from a careful analysis of these equations.

III. THE MAIN THEORETICAL RESULTS: FROM MIXTURES TO A SINGLE GAUSSIAN

In this section, we present the main theoretical results of the present work and discuss their consequences: We show that with random labels, GMM models can be reduced to a single GCM model. This provides an explanation of the universality observed in Figs. 1 and 2.

We would like the reader to note that we state our results using theorems because indeed we were able to rigorously establish them mathematically. However, the proofs are deferred to the Appendixes, and the reasoning and derivations presented in this section follow the level of rigour common in theoretical physics. We made this choice to ensure readability to both physics- and mathematics-oriented audiences.

The starting point is the Gaussian mixture model. This is a very generic model of data, and standard approximation results (e.g., the Stone-Weierstrass theorem) show in particular that one can approximate data density to arbitrary precision by Gaussian mixtures. While in the worst case this would require a diverging number of Gaussians in the mixture, it can be shown that (as far as the generalized linear model is concerned) a mixture of a small number of Gaussians is actually able to approximate very complex data set in high dimension [9,65,87]. More precisely, in the proportional high-dimensional regime, data generated by generative adversarial networks (GANs), one of the state-of-the-art techniques to generate realistic looking data, behave as Gaussian mixtures for such classifiers [66]. We shall thus use this model as our benchmark of "complex" data distribution.

If a mixture model is a good approximation of reality in high dimension, the question remains: *Why is it that we can fit real data set with a single Gaussian*? Our main technical question will thus be the following: If we use random labels, what is the difference between a GMM and a single Gaussian model?

A. Mean invariance with random labels

We thus now move to the random label case and show how we can surprisingly use a simple Gaussian distribution instead of the GMM. We are going to use Theorems 1 and 2. Note that the asymptotic value of the energy, or loss, only depends on the probability vector $\boldsymbol{\rho} \in [0,1]^K$ (with entries ρ_c corresponding to the respective sizes of the K clusters), the matrix of averages $\boldsymbol{M} \in \mathbb{R}^{K\times p}$ (with rows $\boldsymbol{\mu_c} \in \mathbb{R}^p$), and the concatenation of covariances $\boldsymbol{\Sigma}^\otimes \in \mathbb{R}^{K\times p\times p}$ (with rows $\boldsymbol{\Sigma}_c \in \mathbb{R}^{p\times p}$), and therefore we denote

$$\mathcal{E}_\ell = \mathcal{E}_\ell^{\mathrm{gmm}}(\boldsymbol{\rho}, \boldsymbol{M}, \boldsymbol{\Sigma}^\otimes).$$

Similarly, for the Gaussian covariate model we define the limiting value

$$\mathcal{E}_\ell = \mathcal{E}_\ell^{\mathrm{gcm}}(\boldsymbol{\Sigma}),$$

where in both cases we omitted the explicit dependence on the parameters (α, λ). We are now in a position to state a lemma crucial to our first main universality result:

Lemma 1. Single mean lemma for random labels. In the random label setting (2), assume that the loss ℓ is symmetric, in the sense that $\ell(x, y) = \ell(-x, -y)$ for $x, y \in \mathbb{R}$. Then, the limiting value $\mathcal{E}_\ell$ of the risk is independent from the means, i.e., for all choices of $\boldsymbol{\rho}$, $\boldsymbol{M}$, and $\boldsymbol{\Sigma}^{\otimes}$ we have

$$\mathcal{E}_\ell^{\text{gmm}}(\boldsymbol{\rho}, \boldsymbol{M}, \boldsymbol{\Sigma}^{\otimes}) = \mathcal{E}_\ell^{\text{gmm}}(\boldsymbol{\rho}, \mathbf{0}_{K\times p}, \boldsymbol{\Sigma}^{\otimes}).$$

The symmetry condition on the loss is not really restrictive and is satisfied by virtually all losses used in binary classification [in particular, margin-based losses of the form $\ell(x, y) = \phi(xy)$]. Since a mixture of Gaussians with equal means and covariances is equivalent to a single Gaussian, we can now write the following theorem:

Theorem 3. Gaussian universality for random labels. Consider the same assumptions as in Lemma 1, and assume further that the data are homogeneous, i.e.,

$$\boldsymbol{\Sigma}_c = \boldsymbol{\Sigma} \quad \text{for all} \quad c \in \mathcal{C}.$$

Then the asymptotic risk is equivalent to that of a single centered Gaussian:

$$\mathcal{E}_\ell^{\text{gmm}}(\boldsymbol{\rho}, \boldsymbol{M}, \boldsymbol{\Sigma}^{\otimes}) = \mathcal{E}_\ell^{\text{gcm}}(\boldsymbol{\Sigma}).$$

This is our first main universality theorem: a mixture of homogeneous Gaussians [88] can be replaced, when using random labels by a single Gaussian.

This surprising fact, alone, explains the empirical observation presented in Figs. 1 and 2, at least if we accept that the different modes are homogeneous (see the discussion in Sec. IV).

Proof sketch. Both Lemma 1 and Theorem 3 stem from the detailed analysis of the replica free energy for the GMM [38]. Indeed, to prove our claims, it suffices to show that the fixed points of the replica equations are the same. This is done in detail in Appendix B, using the replica equation that we provide in Appendix A. In a nutshell, we show that the expression of the GMM reduces to those of the GCM. ■

B. Generic loss with vanishing regularization

Additionally, we note that in Fig. 1 at vanishing regularization, we did not even require a matching covariance, and instead we used a trivial one. This is because of the following consequence of Lemma 1:

Theorem 4. Gaussian universality for vanishing regularization. Consider the same assumptions as in Theorem 3. Then if the minimizer of ℓ is unique and the data covariance full-rank, then the asymptotic minimal loss for Gaussian data does not depend on the covariance when the regularization is absent, $\lambda = 0$.

Proof. Consider the empirical risk minimization problem in Eq. (1) with data from the Gaussian covariate model Eq. (3) with random labels. Without loss of generality, we can write $\boldsymbol{x}_\mu = \boldsymbol{\Sigma}^{1/2}\boldsymbol{z}_\mu$, with $\boldsymbol{z}_\mu \sim \mathcal{N}(\mathbf{0}_p, \boldsymbol{I}_p)$. Then, making a change of variables $\boldsymbol{\theta}' = \boldsymbol{\Sigma}^{1/2}\boldsymbol{\theta}$, we can write

$$\begin{aligned}\widehat{\mathcal{R}}_n^*(\boldsymbol{X}, \boldsymbol{y}) &= \inf_{\boldsymbol{\theta}\in\mathbb{R}^p} \frac{1}{n}\sum_{\mu=1}^{n} \ell(\boldsymbol{\theta}^\top \boldsymbol{x}_\mu, y_\mu) + \frac{\lambda}{2}||\boldsymbol{\theta}||_2^2 \\ &= \inf_{\boldsymbol{\theta}'\in\mathbb{R}^p} \frac{1}{n}\sum_{\mu=1}^{n} \ell(\boldsymbol{\theta}'^\top \boldsymbol{z}_\mu, y_\mu) + \frac{\lambda}{2}||\Sigma^{-1/2}\boldsymbol{\theta}'||_2^2,\end{aligned}$$

where we have used the fact that y_μ are independent of $\boldsymbol{x}_\mu$. Since the minimizer of ℓ is unique, the result follows from taking $\lambda \to 0^+$.

Note that in particular Theorem 4 implies that for random labels, the GCM model with a covariance Σ is equivalent to a Gaussian i.i.d. model with a different regularization given by the norm $||\cdot||_{\Sigma^{-1}}$ induced by the inverse covariance matrix Σ^{-1}. Therefore, in the case in which ℓ has several minima, the $\lambda \to 0^+$ limit will give the performance of the solution with minimum $||\cdot||_{\Sigma^{-1}}$ norm.

Finally, we also note that this analysis also allows us to answer the following important question: *What is being learned with random labels?* This was discussed in particular in the machine learning literature in [24]. For generalized linear models: the model is simply fitting the second-order statistics (the total covariance $\boldsymbol{\Sigma}$).

C. Ridge regression with vanishing regularization

Even though it seems to be well obeyed in practice, one may wonder if we can in some cases get rid of the homogeneity condition. As we shall see, the answer is no: in general, a mixture of *inhomogeneous* Gaussians cannot be strictly replaced by a single one. It turns out, however, that there is one exception, and that the hypothesis can be lifted in one case, namely ridge regression with vanishing regularization with the squared loss $\ell(x, y) = \frac{1}{2}(x - y)^2$:

Theorem 5. Strong universality for ridge loss. In the ridge regression case with vanishing regularization, i.e., when $\lambda \to 0^+$, we have

$$\lim_{\lambda\to 0^+} \mathcal{E}_\ell^{\text{gmm}}(\boldsymbol{\rho}, \boldsymbol{M}, \boldsymbol{\Sigma}^{\otimes}) = \frac{1}{2}\left(1 - \frac{1}{\alpha}\right)_+$$

for any choice of $\boldsymbol{\rho}, \boldsymbol{M}$, or $\boldsymbol{\Sigma}^{\otimes}$.

In particular, it means that in the unregularized limit, any Gaussian mixture behaves in terms of its loss as a single cluster Gaussian model with identity covariance, whose asymptotic training loss is given by $\lim_{\lambda\to 0^+} \mathcal{E}_\ell^{\text{gcm}}(\alpha, \lambda) = \frac{1}{2}(1 - 1/\alpha)_+$.

Proof sketch. The proof of the strong universality, which follows from a rigorous analysis of the replica predictions, amounts to showing that the replica free energy for GMM reduces to that of a single Gaussian. Interestingly, although the fixed points of the replica equations differ between the GMM and Gaussian case, they do give rise to the same free energy. Details can, again, be found in Appendix B 2. ■

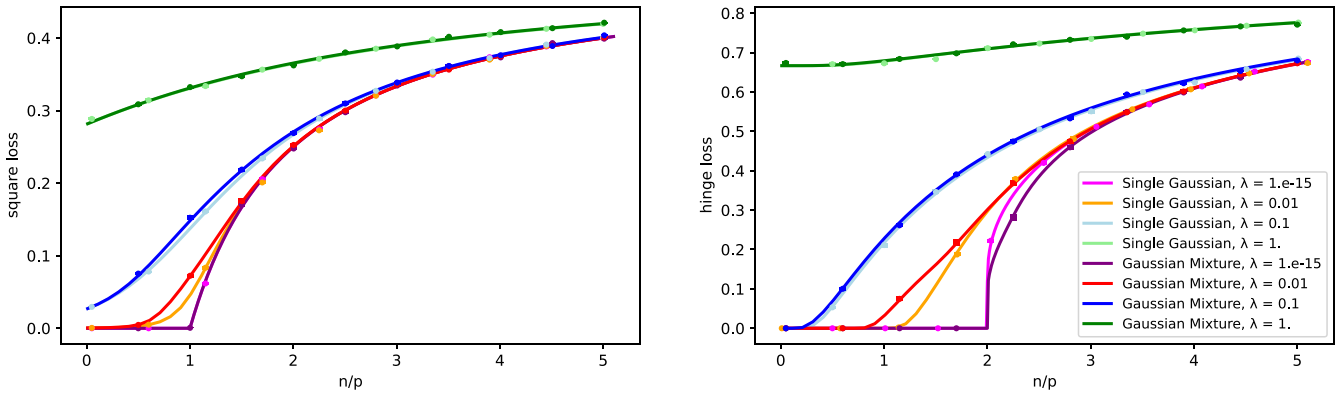

FIG. 3. Ridge/square loss (left) and hinge loss for a single Gaussian vs a mixture of inhomogeneous Gaussians at finite λ. Lines are the asymptotic exact results while dots are simulation ($p = 900$, dark lines for mixture, lighter ones for single Gaussian). When the homogeneity assumption is not obeyed, then a mixture of two Gaussians does not yield results equal to those of a single Gaussian with matching covariance. (Here, a mixture with zero mean and a block covariance with, respectively, diagonal elements equal to 0.01, 0.98, and 0.01 for the first one, and 0.495, 0.01, and 0.495 for the second). Note, however, that the universality is restored in the Ridge case when $\lambda \to 0$, as stated in Theorem 5. It is also very well obeyed with large enough λ, and deviations appear small in general.

IV. NUMERICAL EXPERIMENTS

In this section, we describe more in detail the numerical experiments of Figs. 1 and 2. The colored dots represent the outcome of the simulations on several full-rank data sets. In particular, the blue and the green dots refer to both MNIST and grayscale CIFAR-10 preprocessed with random Gaussian feature maps [34]. In this case, the input data points are constructed as $\boldsymbol{x}_\mu = \sigma(\boldsymbol{z}_\mu \boldsymbol{F})$, with $\boldsymbol{z}_\mu \in \mathbb{R}^d$ being a sample from one of the two data sets, $\boldsymbol{F} \in \mathbb{R}^{d\times p}$ representing the matrix of random features, whose row elements are sampled according to a normal distribution, and σ being some pointwise nonlinearity, namely erf for MNIST and relu for grayscale CIFAR-10. The red dots correspond instead to fashion-MNIST preprocessed with wavelet scattering transform, an ensemble of engineered features producing rotational and translational invariant representations of the input data points [29]. The orange dots correspond to simulations on the synthetic data set built as a mixture of two Gaussians, with data covariance of the two clusters both equal to the identity matrix [$\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \boldsymbol{I}$, $\boldsymbol{\mu}_{1/2} = (\pm 1, 0, \ldots, 0)$, and $\rho_{1/2} = 1/2$]. Further technical details are given in the Appendixes.

Experiments with finite regularization. Figure 2 illustrates the Gaussian universality taking place at finite regularization. The colored dots correspond to the outcome of the simulations for several values of the regularization strength. As we can see from this set of plots, the theoretical learning curve of a single Gaussian with matching covariance perfectly fits the behavior of multimodal and more realistic input data distributions. In fact, even though the experiment is performed for a realistic data set and finite n and p, the asymptotic Gaussian theory gives a perfect fit of the data.

Experiments with vanishing regularization. Figure 1 provides an illustration of the universality effect occurring at vanishing regularization. Here we use $\lambda \to 0$, and following Theorem 4, we observe a collapse on a single curve given by the asymptotic theory for a single Gaussian with unit covariance. It is quite remarkable that our asymptotic theory, which is valid only in the infinite high-dimensional limit, is validated by such experiments with finite dimension, and finite sample size.

Homogeneity assumption. A remarkable point is that the homogeneity assumption (often called homoskedasticity in statistics) we use in Theorem 3, which can be relaxed only for ridge regression, does not seem to be that important in practice, as we observed on our experiments on real data. One may thus wonder if the strong universality of Theorem 5 could be proved in full generality, and not only for the ridge loss. It turns out that the answer is no. Using Theorem 2, we can actually construct an artificial mixture of Gaussians, using *very different* covariances for each individual Gaussians, and we observe small deviations from the strict universality. A mixture of *nonhomogeneous* Gaussians is not strictly equivalent to a single one with random labels (except, as stated in Theorem 5, for the least squares that obey a strong universality). This is illustrated in Fig. 3, where we show the disagreement in the behavior of the training loss between a single Gaussian and a mixture of two nonhomogeneous Gaussians. This is a simple counterexample to the existence of a universal strong form of Gaussian universality, even for ridge regression (see the discussions in, e.g., Refs. [10,58,69,89,90]).

It may thus come as a surprise that real data sets, which certainly will not obey such a strict homogeneity of the different modes, display such a spectacular agreement with the theory. We believe that this is due to two effects: first, the deviations we observed, even in our designed counterexample, are small, so they might not even be seen in practice. Secondly, and especially after observing the data through random or scattering features, it turns out that when we measure the empirical correlation matrix of the different modes, they look quite similar. In fact, it has even been suggested that neural networks are *precisely* learning representations that find such homogeneous Gaussian mixtures [91].

A remark on Rademacher complexity. A final comment is that the discussed universality indicates that, in high dimension, the Rademacher complexity can be effectively replaced

by the one for Gaussian i.i.d. data. *Rademacher complexity* is a key quantity appearing in generalization bounds for binary classification problems [7,20] that measures the ability of estimators in a hypothesis class $\mathcal{H}$ to fit i.i.d. random labels $y^\mu \sim \mathrm{Rad}(\frac{1}{2})$:

$$\mathrm{Rad}_n(\mathcal{H}) = \mathbb{E}\left[\sup_{h\in\mathcal{H}} \frac{1}{n}\sum_{\mu=1}^{n} y_\mu h(\boldsymbol{x}_\mu)\right]. \tag{11}$$

It is explicitly dependent on the specific distribution of the input data points $\boldsymbol{x}_\mu$. As discussed in Ref. [17], there exists a direct mapping between the Rademacher complexity and the minimum $0/1$ training loss—or ground-state energy in the statistical physics parlance. Indeed, for a binary hypothesis class $\mathcal{H} = \{h : \mathbb{R}^p \to \{-1, +1\}\}$ the two are asymptotically related by the following equation:

$$\lim_{n\to\infty} \inf_{h\in\mathcal{H}} \frac{1}{n}\sum_{\mu=1}^{n} \mathbb{P}(h(\boldsymbol{x}_\mu) \neq y_\mu) = \frac{\alpha}{2}[1 - \mathrm{Rad}_n(\mathcal{H})]. \tag{12}$$

Moreover, Ref. [17] discussed how to explicitly compute the Rademacher complexity for Gaussian data using the replica method from statistical physics. This is actually a classical problem, studied by the pioneers of the application of the replica method and spin glass theory to theoretical machine learning [1,2,42,43]. Given the universality advocated in this present work, these Gaussian results thus seem to be of more relevance than previously thought, and in fact they allow us to compute a closed-form asymptotic expression for the Rademacher complexity for realistic data. This is a very interesting outcome of the Gaussian universality with random labels.

However, while we prove universality for convex losses, at this point we only *conjecture it* for nonconvex objectives, such as the ones appearing in the definition of the Rademacher complexity. The proof that a Gaussian mixture approximates well real data sets is still valid for nonconvex losses. The identification of these mixtures with a single Gaussian is, however, using the replica formulas of [12,38], which have been proven only for the case of convex losses. Our conjecture thus depends on proving a similar result for nonconvex (as well as replica symmetry breaking) losses. This (and similar questions on multilayer networks) is left for future work.

V. CONCLUSION

For the classical problem of fitting random labels with perceptrons, also known as generalized linear models in high dimensions, we showed that, far from being only a toy example, the Gaussian i.i.d. assumption is an excellent model of reality. The conclusion extends to deep-learning models in the lazy regimes as those are essentially random feature models. There are a number of potentially interesting extensions of this work, including nonconvex losses and multilayer neural networks, and beyond the random label cases, that should be investigated in the future.

These results, we believe, are of special interest given the number of theoretical studies with the Gaussian design and its variants that are amenable to exact characterization, and that turn out to be less idealistic, and more realistic, than perhaps previously assumed. We believe, in particular, that these strengthen considerably the ensemble of results obtained within the statistical physics community, as well as in the statistical analysis of high-dimensional data. We anticipate that such redemption of the Gaussian assumption will lead to more work in this direction using the Gaussian assumption and for those aiming to extend out universality results.

ACKNOWLEDGMENTS

We acknowledge funding from the ERC under the European Union's Horizon 2020 Research and Innovation Program Grant Agreement 714608-SMiLe, as well as by the Swiss National Science Foundation grant SNFS OperaGOST, 200021_200390 and the *Choose France - CNRS AI Rising Talents* program.

APPENDIX A: EXACT ASYMPTOTIC PERFORMANCES OF GCM AND GMM

In this Appendix, we summarize the exact asymptotic formulas for the performance of the generalized linear classifiers on random labels for the two structured data models studied in the main body of the text: the Gaussian covariate model (GCM) and the Gaussian mixture model (GMM).

1. Preliminaries: The setting

Before moving to the key formulas, let us recap the setting. We are interested in the performance of generalized linear classifiers:

$$\hat{y}(\boldsymbol{x}) = \mathrm{sgn}(\hat{\boldsymbol{\theta}}^\top \boldsymbol{x}), \tag{A1}$$

where $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$ is trained by minimizing the following empirical risk on n independent training samples $(\boldsymbol{x}_\mu, y_\mu)_{\mu\in[n]} \in \mathbb{R}^p \times \{-1, +1\}$:

$$\widehat{\mathcal{R}}_n^*(\boldsymbol{X}, \boldsymbol{y}) = \inf_{\boldsymbol{\theta}\in\mathbb{R}^p} \frac{1}{n}\sum_{\mu=1}^{n} \ell(\boldsymbol{\theta}^\top \boldsymbol{x}_\mu, y_\mu) + \frac{\lambda}{2}||\boldsymbol{\theta}||_2^2 \tag{A2}$$

for a compact subset $\mathcal{S}_p \subset \mathbb{R}^p$ and a convex loss function ℓ. In particular, we are interested in the case in which the labels $y_\mu \in \{-1, +1\}$ are randomized (i.e., not correlated with the inputs $\boldsymbol{x}_\mu$),

$$y_\mu \sim \mathrm{Unif}(\{-1, 1\}), \tag{A3}$$

and the inputs are generated independently from one of the following two structured models: (i) The Gaussian covariate model (GCM): $\boldsymbol{x}_\mu \sim \mathcal{N}(\boldsymbol{0}_p, \boldsymbol{\Sigma})$, and (ii) the Gaussian mixture model (GMM): $\boldsymbol{x}_\mu \sim \sum_{c\in\mathcal{C}} \rho_c \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)$, where $\mathcal{C} = \{1, \ldots, K\}$ is the label set for the Gaussian clouds, and $\rho_c \in [0, 1]$ are the density of points in each class, satisfying $\sum_{c\in\mathcal{C}} \rho_c = 1$. Note that in this random label setting, the GCM model is a special case of the GMM, where $K := |\mathcal{C}| = 1$ and $\boldsymbol{\mu}_1 = \boldsymbol{0}_p$.

In the following, we will be interested in describing the exact asymptotic limit of the following performance metrics in the proportional high-dimensional limit where $n, p \to \infty$ with the ratio $\alpha := \frac{n}{p}$ and the number of clusters K are fixed: (i) Training loss: $\hat{\mathcal{E}}_\ell(\boldsymbol{X}, \boldsymbol{y}) := \frac{1}{n}\sum_{\mu=1}^{n} \ell(\hat{\boldsymbol{\theta}}^\top \boldsymbol{x}_\mu, y_\mu)$, and

(ii) 0/1 training error: $\hat{\mathcal{E}}_{0/1}(\boldsymbol{X},\boldsymbol{y}) \coloneqq \frac{1}{n}\sum_{\mu=1}^{n}\mathbb{P}(\mathrm{sgn}(\hat{\boldsymbol{\theta}}^{\top}\boldsymbol{x}_{\mu}) \neq y_{\mu})$, where we have defined the design matrix $\boldsymbol{X} \in \mathbb{R}^{p\times n}$ and the label vector $\boldsymbol{y} \in \{-1,+1\}^{n}$. Note that for convenience in this Appendix, we will focus the discussion on these two measures. But all results could have been stated for $\hat{\mathcal{R}}_{n}^{\star}$ instead. In particular, the training loss $\hat{\mathcal{E}}_{\ell}$ differs from the empirical risk $\hat{\mathcal{R}}_{n}^{\star}$ by the regularization term.

Note on scalings. Although the model above is well defined for any scaling, in the following we focus in the case in which the means and the covariances satisfy:

$$||\boldsymbol{\mu}_{c}||_{2}^{2} = O(1), \quad \mathrm{tr}\,\boldsymbol{\Sigma}_{c} = O(p). \tag{A4}$$

This scaling of the mean and variance is indeed the natural one (see, e.g., [80,92–95]) as well as the most interesting in high dimensions. If the means have a larger norm, then the problem becomes trivial (i.e., the Gaussians are trivially completely separable), while if the means are smaller, it is impossible to separate them (i.e., they become trivially indistinguishable from a single Gaussian cloud).

Ridge and ordinary least-squares classification. Note that for the special case of the ridge classification in which $\ell(x,y) = \frac{1}{2}(y-x)^{2}$, the empirical risk minimization problem defined in Eq. (A2) admits a closed-form solution:

$$\hat{\boldsymbol{\theta}} = (\lambda \boldsymbol{I}_{p} + \boldsymbol{X}\boldsymbol{X}^{\top})^{-1}\boldsymbol{X}\boldsymbol{y} \tag{A5}$$

and therefore the computation of the asymptotic training error or loss boils down to a random matrix theory problem, with a solution equivalent to the one we will discuss shortly below. However, some qualitative features can be drawn just from this expression. First, note that for $\lambda > 0$, the ridge estimator above will always have a nonzero training loss because of the bias introduced by the regularization term $\frac{1}{2}\lambda||\boldsymbol{\theta}||_{2}^{2}$. This can only be achieved in the limit of vanishing regularization $\lambda \to 0^{+}$, in which case the ridge estimator simplifies to

$$\hat{\boldsymbol{\theta}}_{\mathrm{ols}} \coloneqq (\boldsymbol{X}^{\top})^{\dagger}\boldsymbol{y}, \tag{A6}$$

where $\boldsymbol{X}^{\dagger} \in \mathbb{R}^{n\times p}$ is the Moore-Penrose inverse of $\boldsymbol{X}$. In the simplest case in which $\boldsymbol{X}$ is a full-rank matrix (which ultimately depends on the covariances), it can be explicitly written as

$$\boldsymbol{X}^{\dagger} \coloneqq \begin{cases} (\boldsymbol{X}^{\top}\boldsymbol{X})^{-1}\boldsymbol{X}^{\top} \text{ if } \alpha < 1, \\ \boldsymbol{X}^{\top}(\boldsymbol{X}\boldsymbol{X}^{\top})^{-1} \text{ if } \alpha > 1. \end{cases} \tag{A7}$$

An important property of the estimator in Eq. (A6) is that it corresponds to the least ℓ_{2}-norm interpolator when the system is underdetermined. Indeed, in the strict case when $\lambda = 0$ (i.e., least-squares regression), the ERM problem in Eq. (A2) is equivalent to inverting a linear system:

$$\boldsymbol{y} = \boldsymbol{X}^{\top}\boldsymbol{\theta}, \tag{A8}$$

i.e., to solve a system of n equations for p unknowns. Again, assuming the data are full-rank [96], for $\alpha = \frac{n}{p} < 1$ the system is *underdetermined*, meaning that there are infinitely many solutions that perfectly interpolate the data. Among all of them, $\hat{\boldsymbol{\theta}}_{\mathrm{ols}}$ is the one that has the lowest ℓ_{2}-norm. Instead, when $\alpha > 1$, the system is overdetermined, and no interpolating (zero-loss) solution exists.

2. Gaussian mixture model with general labels

Exact asymptotics of generalized linear classification with Gaussian mixtures in the proportional regime have been derived under different settings in the literature [77–82]. Of particular interest to our work are the formulas proved in [38] under the most general setting of a multiclass learning problem with convex losses and penalties and generic means and covariances. In their work, the asymptotic performance of the minimizer in Eq. (A2) was proven in the case in which the labels are correlated to the mean. The formula we state in the text as Theorem 2 is a straightforward adaptation of their result in the particular case of binary classification with K clusters and randomized labels.

Zero mean limit: Of particular interest for what follows is the zero-mean limit $\boldsymbol{\mu}_{c} = \boldsymbol{0}_{p}$ of the above equations, which is simply given by

$$\begin{aligned}
&\hat{m}_{c} = 0, \\
&\hat{V}_{c} = \frac{\alpha}{2}\rho_{c}\sum_{y\in\{-1,+1\}}\mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[\partial_{\omega}f_{\ell}(y,\sqrt{q_{c}}\xi,V_{c})], \\
&\hat{q}_{c} = \frac{\alpha}{2}p_{c}\sum_{y\in\{-1,+1\}}\mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[f_{\ell}(y,\sqrt{q_{c}}\xi,V_{c})^{2}], \\
&m_{c} = 0, \\
&V_{c} = \frac{1}{p}\mathrm{tr}\,\Sigma_{c}\left(\lambda\boldsymbol{I}_{p} + \sum_{c'\in\mathcal{C}}\hat{V}_{c'}\Sigma_{c'}\right)^{-1}, \\
&q_{c} = \frac{1}{p}\sum_{c'\in\mathcal{C}}\left[\hat{q}_{c'}\mathrm{tr}\,\Sigma_{c'}\Sigma_{c}\left(\lambda\boldsymbol{I}_{p} + \sum_{c''\in\mathcal{C}}\hat{V}_{c''}\Sigma_{c''}\right)^{-2}\right].
\end{aligned} \tag{A9}$$

A particular case: ridge classification. The self-consistent equations above crucially depend on the loss function ℓ. A particular case of interest in this work—and for which the expressions simplify considerably—is the case of ridge regression where $\ell(x,y) = \frac{1}{2}(x-y)^{2}$. In this case, the proximal can be explicitly written as

$$\mathrm{prox}_{\tau\ell(\cdot,y)}(x) = \frac{x+\tau y}{1+\tau} \quad \Leftrightarrow \quad f_{\ell}(y,\omega,V) = \frac{y-\omega}{1+V} \tag{A10}$$

and therefore the asymptotic training loss admits a closed-form expression:

$$\mathcal{E}_{\ell}^{\mathrm{gmm}} = \sum_{c\in\mathcal{C}}\rho_{c}\frac{1+q_{c}^{\star}}{2(1+V_{c}^{\star})^{2}} \tag{A11}$$

for $(V_{c}^{\star},q_{c}^{\star})_{c\in\mathcal{C}}$ solutions of the following simplified self-consistent equations:

$$\begin{aligned}
&\hat{V}_{c} = \frac{\alpha\rho_{c}}{1+V_{c}}, \quad \hat{q}_{c} = \alpha p_{c}\frac{1+q_{c}}{(1+V_{c})^{2}}, \\
&V_{c} = \frac{1}{p}\mathrm{tr}\,\Sigma_{c}\left(\lambda\boldsymbol{I}_{p} + \sum_{c'\in\mathcal{C}}\hat{V}_{c'}\Sigma_{c'}\right)^{-1}, \\
&q_{c} = \frac{1}{p}\sum_{c'\in\mathcal{C}}\left[\hat{q}_{c'}\mathrm{tr}\,\Sigma_{c'}\Sigma_{c}\left(\lambda\boldsymbol{I}_{p} + \sum_{c''\in\mathcal{C}}\hat{V}_{c''}\Sigma_{c''}\right)^{-2}\right].
\end{aligned} \tag{A12}$$

Note that in particular at the fixed point, we can also express the training loss Eq. (A11) as

$$\mathcal{E}_\ell^{\rm gmm} = \sum_{c\in\mathcal{C}} \frac{\hat{q}_c^*}{2\alpha}. \tag{A13}$$

3. Gaussian covariate model

The asymptotic training loss for the Gaussian covariate model for a fairly general teacher-student setting was first proven in [12]. Although the random label limit can be obtained from this work, as discussed in Appendix A 1 the random label Gaussian covariate model can also be seen as a particular case of the general Gaussian mixture model with $K = 1$ and $\boldsymbol{\mu}_1 = \mathbf{0}_p$. Therefore, its asymptotic performance is included in the discussion above. This leads to Theorem 1 in the main text.

It is worth noting that, for the square loss the expressions simplify considerably. The training loss is given by

$$\mathcal{E}_\ell^{\rm gmm} = \frac{1+q^\star}{2(1+V^\star)^2}, \tag{A14}$$

where $(V^\star, q^\star)$ are solutions of the following simplified self-consistent equations:

$$\hat{V} = \frac{\alpha}{1+V}, \quad \hat{q} = \alpha\frac{1+q}{(1+V)^2}, \quad V = \frac{1}{p}\mathrm{tr}\,\Sigma(\lambda \boldsymbol{I}_p + \hat{V}\Sigma)^{-1},$$
$$q = \frac{1}{p}\hat{q}\,\mathrm{tr}\,\Sigma^2(\lambda \boldsymbol{I}_p + \hat{V}\Sigma)^{-2}. \tag{A15}$$

Since the covariance Σ is positive-definite (and therefore invertible), in the overdetermined regime (for which the training loss is nonzero), the limit $\lambda \to 0^+$ can be easily taken, and the equations reduce to

$$\hat{V} = \frac{\alpha}{1+V}, \quad \hat{q} = \alpha\frac{1+q}{(1+V)^2}, \quad V = \frac{1}{\hat{V}}, \quad q = \frac{\hat{q}}{\hat{V}}, \tag{A16}$$

which is completely independent of the covariance matrix Σ [97]. Moreover, it admits a closed-form solution given by

$$V^\star = q^\star = \frac{1}{\alpha - 1}, \quad \hat{V}^\star = \hat{q}^\star = \alpha - 1. \tag{A17}$$

Therefore, the full training loss is given by

$$\lim_{\lambda\to 0^+} \mathcal{E}_\ell^{\rm gcm}(\alpha,\lambda) = \begin{cases} 0 & \text{for } \alpha \leqslant 1, \\ \frac{1}{2}\left(1 - \frac{1}{\alpha}\right) & \text{for } \alpha > 1. \end{cases} \tag{A18}$$

APPENDIX B: FROM GAUSSIAN MIXTURE TO SINGLE GAUSSIAN

1. Mixture of Gaussians with zero means

We first prove Lemma 1 in the main text. First, by Theorem 2, the asymptotic loss $\mathcal{E}_\ell^{\rm gmm}(\boldsymbol{\rho}, \boldsymbol{M}, \boldsymbol{\Sigma}^\otimes)$ $[\mathcal{E}_\ell^{\rm gmm}(\boldsymbol{\rho}, \mathbf{0}, \boldsymbol{\Sigma}^\otimes)]$ is a deterministic function of $(m_c^\star, q_c^\star, V_c^\star)_{c\in\mathcal{C}}$, which are the *unique* fixed points of (10) [(A9)]. Since both saddle point equations differ only by setting $m_c = \hat{m}_c = 0$, Lemma 1 is a consequence of the following:

Lemma 2. Let $(V_c^\star, q_c^\star)_{c\in\mathcal{C}}$ be the solutions of Eqs. (A9). Then, $(0, V_c^\star, q_c^\star)_{c\in\mathcal{C}}$ satisfy the general fixed point equations of (10).

Proof. If we plug in $m_c = \hat{m}_c = 0$ for all $c \in \mathcal{C}$, the equations for $V_c, \hat{V}_c, q_c, \hat{q}_c$ become identical in (10) and (A9). It is also easy to check that $\hat{m}_c = 0$ for all c implies that $m_c = 0$; what remains is to show that the last equation holds, i.e.,

$$\frac{\alpha}{2}\rho_c \sum_{y\in\{-1,+1\}} \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[f_\ell(y, \sqrt{q_c^\star}\xi, V_c^\star)] = 0. \tag{B1}$$

Define the function

$$g(\omega, V) = f_\ell(-1, \omega, V) + f_\ell(+1, \omega, V),$$

so that

$$\hat{m}_c^\star \propto \mathbb{E}_{\xi\sim\mathcal{N}(0,1)}[g(\sqrt{q_c^\star}\xi, V_c^\star)].$$

We shall show that g is odd in ω; since ξ is centered, the lemma will be proven. To do so, we shall show that

$$f_\ell(y, \omega, V) = -f_\ell(-y, -\omega, V)$$

for all $y \in \{-1, +1\}$, $\omega \in \mathbb{R}$, and $V \in \mathbb{R}$. By definition, we have

$$f_\ell(y, \omega, V) = V^{-1}[\mathrm{prox}_{V\ell(\cdot,y)}(\omega) - \omega],$$

and the linear term in ω is immediate. For the proximal operator, we use the symmetry of ℓ and write

$$\begin{aligned}
&\mathrm{prox}_{V\ell(\cdot,y)}(\omega) \\
&\quad = \arg\min_{z\in\mathbb{R}}\left[\frac{1}{2\tau}(z-\omega)^2 + \ell(z,y)\right] \\
&\quad = \arg\min_{z\in\mathbb{R}}\left[\frac{1}{2\tau}[(-z)-(-\omega)]^2 + \ell(-z,-y)\right] \\
&\quad = -\mathrm{prox}_{V\ell(\cdot,-y)}(-\omega),
\end{aligned}$$

which concludes the proof.

2. Strong universality of ordinary least-squares

We now have all the elements we need to establish the universality of the ordinary least-squares estimator stated in Theorem 5 in the main text. Our starting point is the ordinary least-squares problem for the Gaussian mixture model in the overdetermined regime $\alpha > 1$. In this case, the training loss is given by Eq. (A11) with $(V_c^\star, q_c^\star)_{c\in\mathcal{C}}$ unique solutions of the following equations:

$$\hat{V}_c = \frac{\alpha\rho_c}{1+V_c}, \quad \hat{q}_c = \alpha\rho_c\frac{1+q_c}{(1+V_c)^2},$$
$$V_c = \frac{1}{d}\mathrm{tr}\,\Sigma_c\left(\sum_{c'\in\mathcal{C}}\hat{V}_{c'}\Sigma_{c'}\right)^{-1},$$
$$q_c = \frac{1}{d}\sum_{c'\in\mathcal{C}}\left[\hat{q}_{c'}\,\mathrm{tr}\,\Sigma_{c'}\Sigma_c\left(\sum_{c''\in\mathcal{C}}\hat{V}_{c''}\Sigma_{c''}\right)^{-2}\right]. \tag{B2}$$

We shall now show how to reduce these equation to a simple analytical formula, equivalent to that of a single Gaussian. Combining the equations for $\hat{V}_c$ and V_c, one sees that the fixed

Algorithm 1 Generating data set $\mathcal{D} = \{\boldsymbol{x}^\mu, y^\mu\}_{\mu=1}^n$

Input: Integer p, flag *data set*, matrix $F \in \mathbb{R}^{d\times p}$ of random Gaussian features
If the *data set type* is i.i.d. Gaussian:
  Sample each input data-point as $\boldsymbol{x}^\mu \sim \mathcal{N}(0, \boldsymbol{I})$, with $\boldsymbol{I} \in \mathbb{R}^{p\times p}$ the identity matrix;
Else if the *data set type* is a Gaussian Mixture:
  Sample each input data-point as $\boldsymbol{x}^\mu \sim \sum_{k=1}^K \rho_k \mathcal{N}(\boldsymbol{\mu}_k, \Sigma_k)$, with $\boldsymbol{\mu}_k$ being the centroid of the kth cluster and Σ_k the corresponding covariance matrix;
Else if the *data set type* is a real data set preprocessed with random Gaussian features:
  Load the real data-set samples z^μ $\forall \mu = 1, \ldots, n$ with Pytorch dataloaders;
  Assign $\boldsymbol{x}^\mu \to \sigma(\boldsymbol{z}^\mu F)$;
Else if the *data set type* is a real data set preprocessed with wavelet scattering:
  Load the real data-set samples z^μ $\forall \mu = 1, \ldots, n$;
  Apply wavelet scattering transform on z^μ;
Sample the labels according to the Rademacher distribution as $y^\mu \sim \frac{1}{2}(\delta_{+1} + \delta_{-1})$
Return: $\mathcal{D} = \{\boldsymbol{x}^\mu, \boldsymbol{y}^\mu\}_{\mu=1}^n$

point must satisfy the following identity:

$$\sum_{c\in\mathcal{C}} \hat{V}_c^\star V_c^\star = 1. \tag{B3}$$

Similarly, multiplying the equation for q_c by $\hat{V}_c$, summing over $c \in \mathcal{C}$, and doing the same for the equation for $\hat{q}_c$ with V_c, we get a second identity satisfied by the fixed point:

$$\sum_{c\in\mathcal{C}} (\hat{V}_c^\star q_c^\star - V_c^\star \hat{q}_c^\star) = 0. \tag{B4}$$

Note that, at this point these relations could have been derived for any loss functions. For the specific case of the square loss, further substituting the hat variables, these conditions are equivalent to

$$\sum_{c\in\mathcal{C}} \rho_c \frac{V_c^\star}{1+V_c^\star} = \frac{1}{\alpha}, \tag{B5}$$

$$\sum_{c\in\mathcal{C}} \rho_c \frac{V_c - q_c}{(1+V_c)^2} = 0. \tag{B6}$$

We thus find, combining Eq. (B2) for $\hat{q}_c$ with Eq. (B6),

$$\sum_{c\in\mathcal{C}} \hat{q}_c^\star = \sum_{c\in\mathcal{C}} \alpha\rho_c \frac{1+V_c^\star}{(1+V_c^\star)^2} = \sum_{c\in\mathcal{C}} \alpha\rho_c \frac{1}{1+V_c^\star}. \tag{B7}$$

Our goal is to evaluate the loss at the fixed point, which is given by Eq. (A13):

$$\mathcal{E}_\ell^{\rm gmm} = \sum_{c\in\mathcal{C}} \frac{\hat{q}_c^\star}{2\alpha}. \tag{B8}$$

Combining this definition with Eqs. (B6) and (B7), we find that

$$2\mathcal{E}_\ell^{\rm gmm} + \frac{1}{\alpha} = \sum_{c\in\mathcal{C}} \rho_c \frac{1}{1+V_c^\star} + \sum_{c\in\mathcal{C}} \rho_c \frac{V_c^\star}{1+V_c^\star} = 1 \tag{B9}$$

so that finally we reach the promised result:

$$\lim_{\lambda\to 0^+} \mathcal{E}_\ell^{\rm gmm}(\alpha, \lambda, K) = \frac{1}{2}\left(1 - \frac{1}{\alpha}\right)_+ = \lim_{\lambda\to 0^+} \mathcal{E}_\ell^{\rm gcm}(\alpha, \lambda) \tag{B10}$$

as claimed in Theorem 5 in the main text.

APPENDIX C: NUMERICAL SIMULATIONS

In this Appendix, we provide further details concerning the protocol we used to perform the numerical simulations (see Fig. 4), which corroborate the theoretical results exemplified in the main manuscript. All codes are publicly available on the GitHub repository associated with the current paper [30].

1. Datasets

We consider different types of datasets. Two of them are synthetic datasets and correspond to i.i.d Gaussian input data points and Gaussian Mixtures. Four of them are instead real datasets which represent standard benchamarks in machine learning experiments. In particular, the first one is MNIST [31], consisting in a collection of 70.000 images of size 28 × 28 pixels of handwritten digits and ten different classes, namely, the digits from 0 to 10 [31], with 7000 images per class. The second one, is fashion-MNIST, composed of 70.000 images of size 28 × 28 pixels of Zalandos articles and ten different classes corresponding to ten different articles [32], with 7000 images per class. The third one is grayscale CIFAR10, consisting in a collection of 60000 natural images of 32 × 32 pixels and 10 different classes, with 6000 images per class. The fourth is grayscale tiny-Imagenet, which corresponds instead to a smaller version of the well-known Imagenet benchmark [98] and it is made of 100.000 natural images, downsampled to 64 × 64 pixels each and grouped into 200 different classes. Contrary to Fig. 2 in the main text, Fig. 4 shows the predictions of the Gaussian theory with respect to the numerical simulations on MNIST, fashion-MNIST, CIFAR10 and tiny ImageNet when no preprocessing is applied. As can be seen, despite the overall quite good agreement between theory and numerical experiments, we start observing some (very) small deviations from the Gaussian predictions. Indeed, as shown in Sec. D, the covariance matrices associated to the different modes of the underlying real data distribution are, in this case, more heterogeneous than the ones observed when a preprocessing stage is applied. This is consistent with the homogeneous assumption in Theorem 3 and implying Gaussian Universality.

2. Preprocessing

In the numerical experiments shown in Figs. 1 and 2 of the main text, we have both normalized and preprocessed the real data sets with either random feature maps [34] or through wavelet scattering transform [29]. For concreteness and completeness, Fig. 5 illustrates these simulations with yet another data set, namely grayscale tiny-Imagenet, with respect to what was already presented in the main text in Fig. 5. The procedure used to preprocess the real data sets is exemplified in Sec. IV. For the sake of clarity, we summarized it through the pseudocode in Algorithm 1.

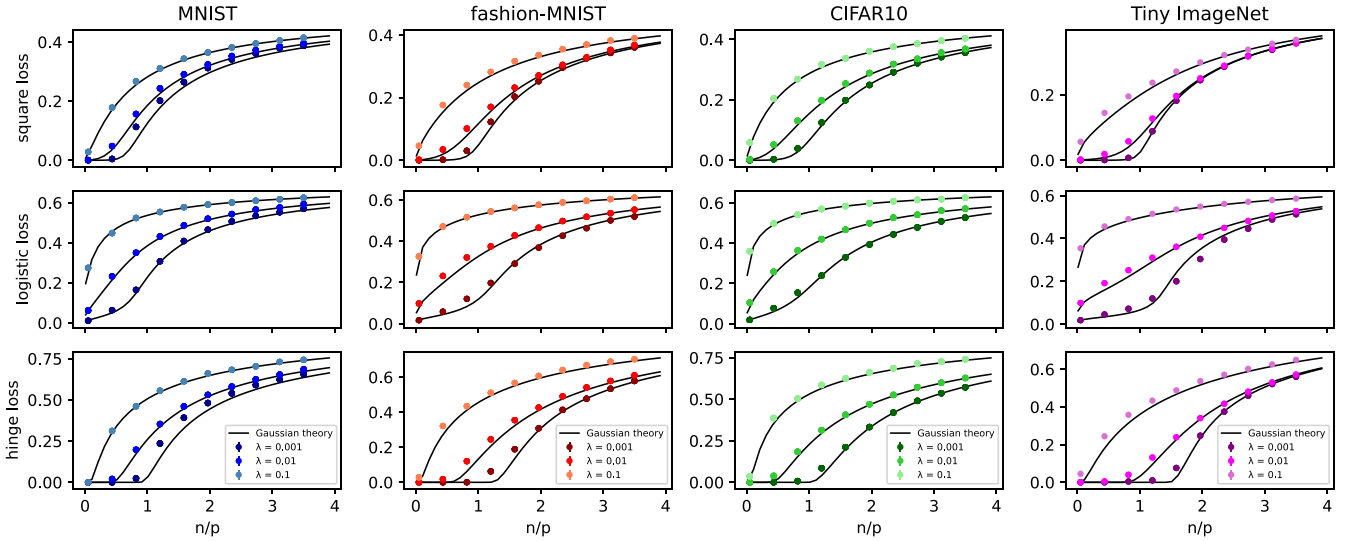

FIG. 4. This figure shows the training loss as a function of the number of samples n per dimension p at finite regularization λ. Top panel: the square loss; bottom panel: the hinge loss. The first column refers to MNIST, the second column corresponds to fashion-MNIST, the third column corresponds to CIFAR10 in grayscale, and the fourth column corresponds to tiny ImageNet in grayscale. Black solid lines correspond to the outcome of the replica calculation, obtained by assigning to Σ the covariance matrix of each data set. The colored dots correspond to the simulations for different values of λ, as specified in the plot legend. Simulations are averaged over 10 samples, and the error bars are not visible at the plot scale.

The real data sets are loaded through Pytorch dataloaders [99]. In particular, the dataloader of CIFAR10 includes a grayscale transformation of the data set in order to reduce the number of input channels of the RGB color encoding scheme to one. The wavelet scattering transform is instead implemented by means of the Kymatio Python library [35]. Note that with the purpose of speeding up the realization of the learning curves and to reduce fluctuations, the preprocessed real data sets are generated once for all through Algorithm 1 and then stored in a hdf5 file.

3. Learning phase

Given the data set generated as in Algorithm 1, the aim is to infer the estimator $\boldsymbol{\theta}$ minimizing the empirical risk as in Eq. (1) of the main paper. In the present work, we consider three distinct kinds of loss functions:

(i) *Square loss.* In this specific case, the goal is to solve the following optimization problem:

$$\widehat{\mathcal{R}}_n^*(\boldsymbol{X},\boldsymbol{y}) = \inf_{\boldsymbol{\theta}\in\mathcal{S}_p} \frac{1}{2n}\sum_{\mu=1}^{n}(\boldsymbol{\theta}^\top \boldsymbol{x}_\mu - y_\mu)^2 + \frac{\lambda}{2}||\boldsymbol{\theta}||_2^2. \tag{C1}$$

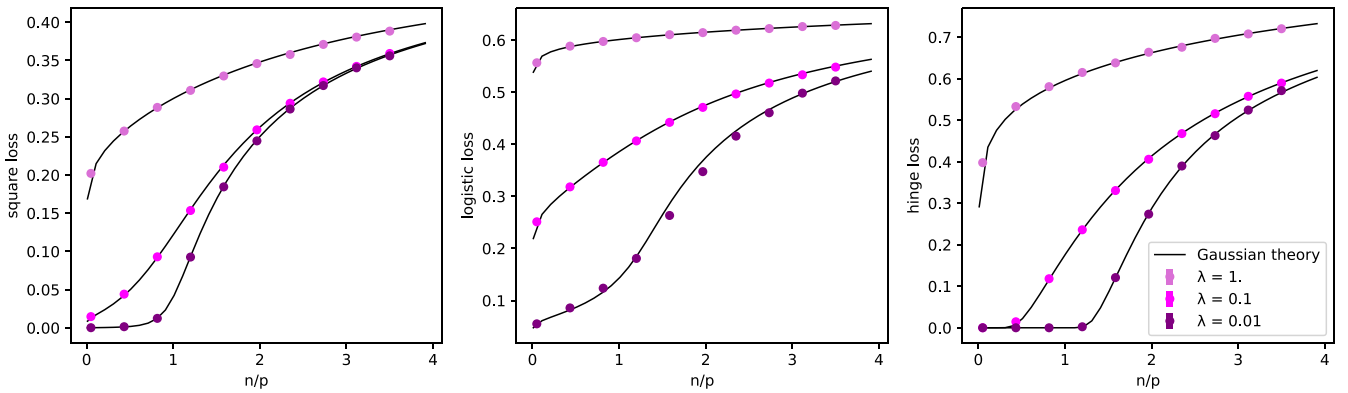

FIG. 5. Numerical simulations of universality: As in Fig. 2, this figure shows the training loss as a function of the number of samples n per dimension p at various values of λ for another data set we used here for completeness. Here we used a grayscale tiny-Imagenet preprocessed with Gaussian random features and tanh nonlinearity. Left panel: the square loss; middle panel: the logistic loss; right panel: the hinge loss. The colored dots refer to numerical simulations while the black solid lines correspond to the theoretical prediction of a single Gaussian with corresponding input covariance matrices. The numerical simulations are averaged over ten different realizations.

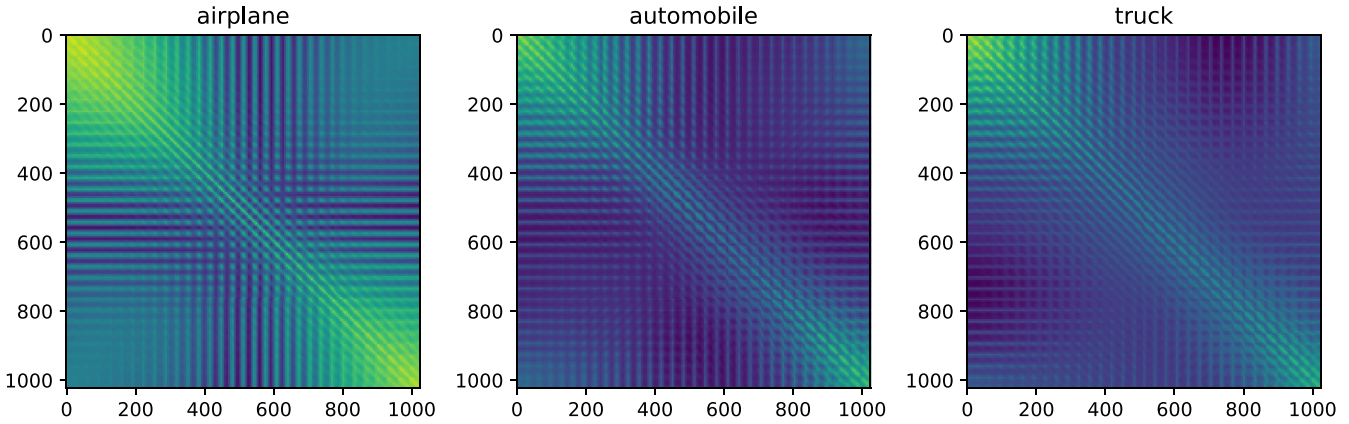

FIG. 6. Input data correlation matrix of grayscale CIFAR10, conditioned on the true labels, e.g., airplane (leftmost), automobile (middle), truck (rightmost). Lighter colors refer to stronger correlation.

The estimator can be determined here through the Moore-Penrose inverse as follows, without relying on any learning algorithm:

$$\boldsymbol{\theta} = \begin{cases} \left(\boldsymbol{X}^\top \boldsymbol{X} + \lambda \boldsymbol{I}_p\right)^{-1} \boldsymbol{X}^\top \boldsymbol{y} & \text{if } n > p, \\ \boldsymbol{X}^\top \left(\boldsymbol{X}\boldsymbol{X}^\top + \lambda \boldsymbol{I}_n\right)^{-1} \boldsymbol{y} & \text{if } p > n. \end{cases} \tag{C2}$$

(ii) *Logistic loss.* In this specific case, the goal is to solve the following optimization problem:

$$\widehat{\mathcal{R}}_n^*(\boldsymbol{X}, \boldsymbol{y}) = \inf_{\boldsymbol{\theta} \in \mathcal{S}_p} \frac{1}{n} \sum_{\mu=1}^{n} \log(1 + \exp(-y_\mu \boldsymbol{\theta}^\top \boldsymbol{x}_\mu)) + \frac{\lambda}{2} ||\boldsymbol{\theta}||_2^2. \tag{C3}$$

Since the estimator of logistic regression cannot be determined through an explicit closed formula, we made use here of the *lbgfs* solver with *penalty* set to ℓ_2. This optimizer corresponds to a gradient descent (GD) -like second-order optimization method, and it is implemented in the LogisticRegression class of the Scikit-Learn Python library [100]. The GD algorithm stops either if a maximum number of iterations has been reached or if the maximum component of the gradient goes below a certain threshold. We fixed this tolerance to 1×10^{-5} and the maximum number of iterations to 1×10^4.

(iii) *Hinge loss.* In this specific case, the goal is to solve the following optimization problem:

$$\widehat{\mathcal{R}}_n^*(\boldsymbol{X}, \boldsymbol{y}) = \inf_{\boldsymbol{\theta} \in \mathcal{S}_p} \frac{1}{n} \sum_{\mu=1}^{n} \max(0, 1 - y_\mu \boldsymbol{\theta}^\top \boldsymbol{x}_\mu)) + \frac{\lambda}{2} ||\boldsymbol{\theta}||_2^2. \tag{C4}$$

As for logistic regression, even in this case where we cannot rely on any explicit formula for the estimator, it is inferred by means of a learning algorithm. In particular, for the simulations at finite regularization strength, we made use of the LinearSVC class provided by Scikit-Learn [100] and implementing the Support Vector Classification (SVC) with linear kernels and L_2 regularization if the *penalty* is set to ℓ_2. In this case, we set the tolerance of convergence to 1×10^{-5} and the maximum number of iterations to 1×10^5. Unfortunately, LinearSVC struggles to converge for vanishing regularization strengths. Therefore, we made use of CVXPY [101,102] in order to perform the simulations at $\lambda = 1 \times 10^{-15}$. CVXPY is an open-source Python-embedded modeling language for convex optimization problems. We set the *solver* option to none; in this way, CVXPY chose automatically the most specialized solver for the optimization problem type. While being slower than LinearSVC, CVXPY guarantees convergence at vanishing regularization strengths.

At the end of the training process, we evaluate the training loss ℓ on the minimizer of the corresponding empirical risk minimization problem. To get the learning curves, we then repeat the whole process for a specified range of n/p and for a certain number of different realizations of the learning problem, as exemplified in Algorithm 2.

Algorithm 2 Learning curve

Input: range of n/p, flag *data set type*, flag *which estimator*
For *seed* in a specified number of seeds **do**:
  For n/p in a specified range **do**:
    Choose the data set according to *data set type*;
    Compute the estimator according to the desired optimization problem as in (i)–(iii);
    Compute the training loss ℓ at fixed n/p;
  Update the mean train loss and its standard deviation with the new contribution from the current seed.
Return: Mean train loss and standard deviation as a function of n/p.

APPENDIX D: EMPIRICAL EVIDENCE OF THE HOMOGENEITY ASSUMPTION

As seen in the counter example illustrated in Fig. 3, in the case of very heterogeneous Gaussian mixtures, we can observe small deviations from universality both at zero and finite regularization. However, this disagreement between a single Gaussian and Gaussian mixtures does not appear in the experiments with real data sets of Figs. 1 and 2, despite their certainly multimodal and mode-heterogeneity nature. First,

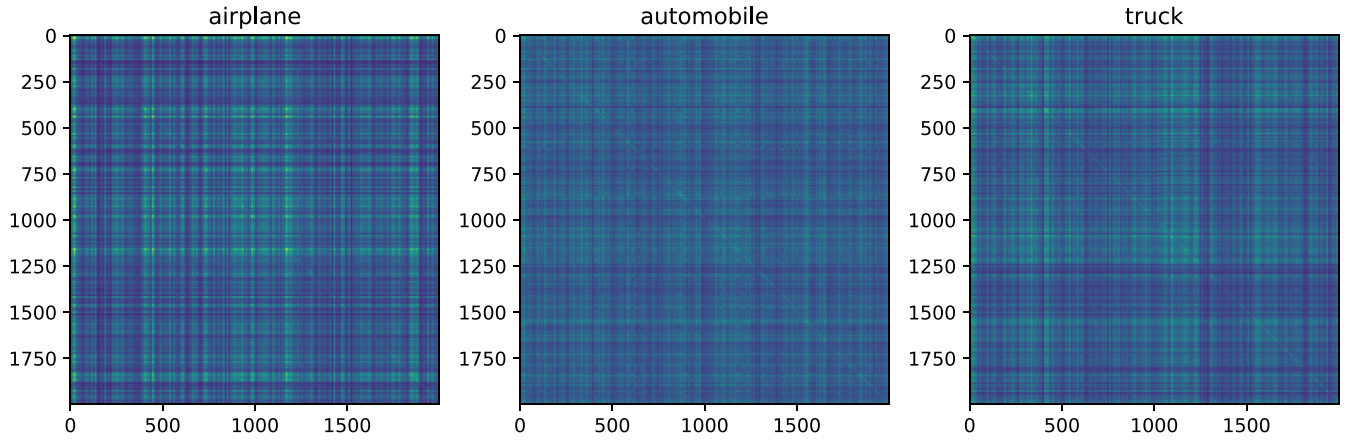

FIG. 7. Input data correlation matrix of grayscale CIFAR10 preprocessed with Gaussian random features and erf nonlinearity. The correlation matrices are conditioned on the true labels, e.g., airplane (leftmost), automobile (middle), truck (rightmost). Lighter colors refer to stronger correlation.

we must acknowledge that deviations are, in general, observed to be small with respect to the homogeneous case, and that the data presented in Fig. 3 were carefully tuned so that the difference is visible.

Additionally, in this Appendix, we also empirically demonstrate the similarity among the empirical correlation matrices of the various modes characterizing real data-set distributions. Figure 6 shows the correlation matrix of all grayscale CIFAR-10 images depicting airplanes (leftmost), automobiles (middle), and trucks (rightmost), respectively. The point we wish to convey in this plot is that, despite the fact that there exist some modes of the CIFAR-10 empirical distribution that display a consistently different correlation structure (airplane mode) with respect to the other modes (automobile and truck mode), there are some others that look more similar to each other (automobile and truck mode).

As can be seen in Figs. 7 and 8, the structure similarity of the covariance matrices of the various mode is further enhanced when preprocessing grayscale CIFAR-10 with both Gaussian random feature maps and wavelet scattering transforms, at the point that even the less similar modes in the raw data set conform to the others (see airplane mode).

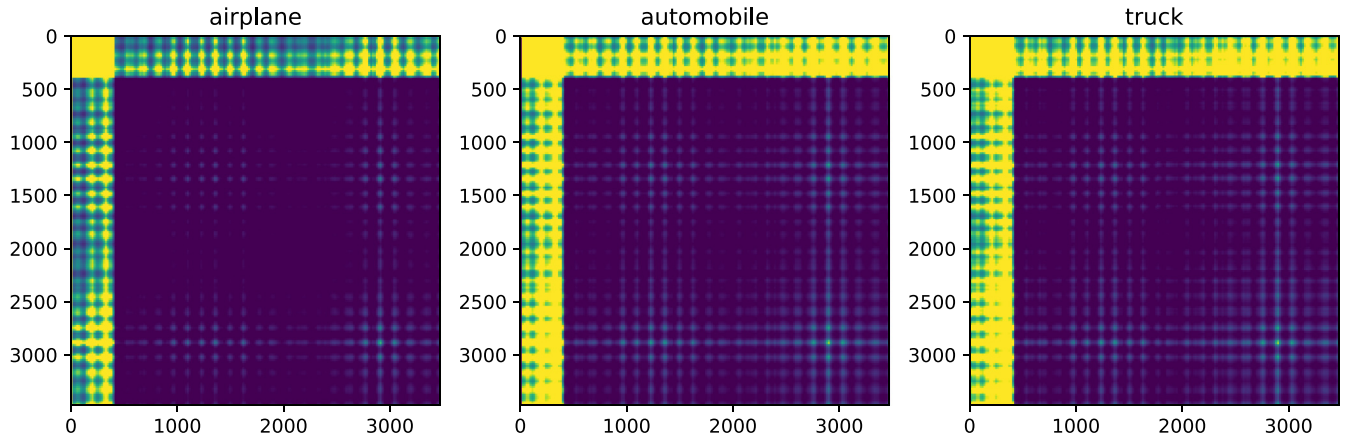

FIG. 8. Input data correlation matrix of grayscale CIFAR10 preprocessed with wavelet scattering transform. The correlation matrices are conditioned on the true labels, e.g., airplane (leftmost), automobile (middle), truck (rightmost). Lighter colors refer to stronger correlation.

[1] E. Gardner and B. Derrida, Three unfinished works on the optimal storage capacity of networks, J. Phys. A **22**, 1983 (1989).

[2] W. Krauth and M. Mézard, Storage capacity of memory networks with binary couplings, J. Phys. **50**, 3057 (1989).

[3] H. S. Seung, H. Sompolinsky, and N. Tishby, Statistical mechanics of learning from examples, Phys. Rev. A **45**, 6056 (1992).

[4] D. Donoho and J. Tanner, Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing, Philos. Trans. R. Soc. London A **367**, 4273 (2009).

[5] E. J. Candès and P. Sur, The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression, Ann. Statist. **48**, 27 (2020).

[6] P. L. Bartlett, P. M. Long, G. Lugosi, and A. Tsigler, Benign overfitting in linear regression, Proc. Natl. Acad. Sci. USA **117**, 30063 (2020).

[7] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press, Cambridge, 2014).

[8] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová, Modeling the influence of data structure on learning in neural networks: The hidden manifold model, Phys. Rev. X **10**, 041044 (2020).

[9] M. El A. Seddik, C. Louart, M. Tamaazousti, and R. Couillet, Random matrix theory proves that deep learning representations of GAN-data behave as Gaussian mixtures, *Proceedings of the 37th International Conference on Machine Learning*, ICML'20 (JMLR.org, 2020), pp. 8573–8582.

[10] B. Bordelon, A. Canatar, and C. Pehlevan, Spectrum dependent learning curves in kernel regression and wide neural networks, *International Conference on Machine Learning* (PMLR, 2020), pp. 1024–1034.

[11] H. Hu and Y. M. Lu, Universality laws for high-dimensional learning with random features, arXiv:2009.07669 (2020).

[12] B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mezard, and L. Zdeborová, Learning curves of generic features maps for realistic datasets with a teacher-student model, J. Stat. Mech. (2022) 114001.

[13] N. Brunel, J.-P. Nadal, and G. Toulouse, Information capacity of a perceptron, J. Phys. A **25**, 5017 (1992).

[14] S. Franz, G. Parisi, M. Sevelev, P. Urbani, and F. Zamponi, Universality of the SAT-UNSAT (jamming) threshold in non-convex continuous constraint satisfaction problems, SciPost Phys. **2**, 019 (2017).

[15] J. Ding and N. Sun, Capacity lower bound for the ising perceptron, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019 (ACM, USA, 2019), pp. 816–827.

[16] B. Aubin, W. Perkins, and L. Zdeborová, Storage capacity in symmetric binary perceptrons, J. Phys. A **52**, 294003 (2019).

[17] A. Abbara, B. Aubin, F. Krzakala, and L. Zdeborová, Rademacher complexity and spin glasses: A link between the replica and statistical theories of learning, *Mathematical and Scientific Machine Learning* (PMLR, 2020), pp. 27–54.

[18] A. Montanari, Y. Zhong, and K. Zhou, Tractability from overparametrization: The example of the negative perceptron, arXiv:2110.15824 (2021).

[19] A. E. Alaoui and M. Sellke, Algorithmic pure states for the negative spherical perceptron, J. Stat. Phys. **189**, 27 (2022).

[20] V. Vapnik, *The Nature of Statistical Learning Theory* (Springer Science & Business Media, New York, 2000).

[21] J. G. Wendel, A problem in geometric probability, Math. Scand. **11**, 109 (1962).

[22] T. M. Cover, Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition, IEEE Trans. Electron. Comput. **EC-14**, 326 (1965).

[23] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, Understanding deep learning (still) requires rethinking generalization, Commun. ACM **64**, 107 (2021).

[24] H. Maennel, I. Alabdulmohsin, I. Tolstikhin, R. J. N. Baldock, O. Bousquet, S. Gelly, and D. Keysers, What do neural networks learn when trained with random labels? *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20 (Curran Associates Inc., 2020), pp. 19693–19704.

[25] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: with Applications in Python* (Springer, New York, 2013).

[26] I. Steinwart and A. Christmann, *Support Vector Machines* (Springer Science & Business Media, 2008).

[27] L. Chizat, E. Oyallon, and F. Bach, On lazy training in differentiable programming, *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (Curran Associates Inc., USA, 2019), pp. 2937–2947.

[28] A. Jacot, F. Gabriel, and C. Hongler, Neural tangent kernel: Convergence and generalization in neural networks, *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18 (Curran Associates Inc., USA, 2018), pp. 8580–8589.

[29] J. Bruna and S. Mallat, Invariant scattering convolution networks, IEEE Trans. Pattern Anal. Machine Intell. **35**, 1872 (2013).

[30] https://github.com/IdePHICS/RandomLabelsUniversality.

[31] L. Deng, The MNIST database of handwritten digit images for machine learning research, IEEE Sign. Proc. Mag. **29**, 141 (2012).

[32] H. Xiao, K. Rasul, and R. Vollgraf, Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms, arXiv:1708.07747 [cs.LG] (2017).

[33] A. Krizhevsky, Learning multiple layers of features from tiny images, Tech. Rep. TR-2009, University of Toronto, Toronto, 2009.

[34] A. Rahimi and B. Recht, Random features for large-scale kernel machines, *Proceedings of the 20th International Conference on Neural Information Processing Systems*, NIPS'07 (Curran Associates Inc., 2007), pp. 1177–1184.

[35] M. Andreux, T. Angles, G. Exarchakis, R. Leonarduzzi, G. Rochette, L. Thiry, J. Zarka, S. Mallat, J. Andén, E. Belilovsky *et al.*, Kymatio: Scattering transforms in python, J. Mach. Learn. Res. **21**, 1 (2020).

[36] C. Thrampoulidis, E. Abbasi, and B. Hassibi, Precise error analysis of regularized m-estimators in high dimensions, IEEE Trans. Inf. Theor. **64**, 5592 (2018).

[37] A. Montanari, F. Ruan, Y. Sohn, and J. Yan, The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime, arXiv:1911.01544 (2019).

[38] B. Loureiro, G. Sicuro, C. Gerbelot, A. Pacco, F. Krzakala, and L. Zdeborová, Learning gaussian mixtures with generalized linear models: Precise asymptotics in high-dimensions, in *Advances in Neural Information Processing Systems*, edited by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. Wortman Vaughan (Curran Associates, USA, 2021), Vol. 34, pp. 10 144–10 157.

[39] E. Dobriban, S. Wager *et al.*, High-dimensional asymptotics of prediction: Ridge regression and classification, Ann. Statist. **46**, 247 (2018).

[40] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, Surprises in high-dimensional ridgeless least squares interpolation, Ann. Statist. **50**, 949 (2022).

[41] P. Del Giudice, S. Franz, and M. A. Virasoro, Perceptron beyond the limit of capacity, J. Phys. France **50**, 121 (1989).

[42] E. Gardner and B. Derrida, Optimal storage properties of neural network models, J. Phys. A **21**, 271 (1988).

[43] B. Derrida, R. B. Griffiths, and A. Prugel-Bennett, Finite-size effects and bounds for perceptron models, J. Phys. A **24**, 4907 (1991).

[44] B. Lopez, M. Schroder, and M. Opper, Storage of correlated patterns in a perceptron, J. Phys. A **28**, L447 (1995).

[45] S. Y. Chung, D. D. Lee, and H. Sompolinsky, Linear readout of object manifolds, Phys. Rev. E **93**, 060301(R) (2016).

[46] P. Rotondo, M. C. Lagomarsino, and M. Gherardi, Counting the learnable functions of geometrically structured data, Phys. Rev. Res. **2**, 023169 (2020).

[47] M. Farrell, B. Bordelon, S. Trivedi, and C. Pehlevan, Capacity of group-invariant linear readouts from equivariant representations: How many objects can be linearly classified under all possible views? *International Conference on Learning Representations* (ICRL, 2022).

[48] D. Arpit, S. Jastrzębski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, and Y. Bengio *et al.*, A closer look at memorization in deep networks, *International Conference on Machine Learning* (PMLR, 2017), pp. 233–242.

[49] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, and M. Sugiyama, Co-teaching: Robust training of deep neural networks with extremely noisy labels, *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18 (Curran Associates Inc., USA, 2018), pp. 8536–8546.

[50] Z. Zhang and M. R. Sabuncu, Generalized cross entropy loss for training deep neural networks with noisy labels, *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18 (Curran Associates Inc., Canada, 2018), pp. 8792–8802.

[51] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, On large-batch training for deep learning: Generalization gap and sharp minima, *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings* (ICLR, 2017).

[52] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro, Exploring generalization in deep learning, *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17 (Curran Associates Inc., USA, 2018), pp. 5949–5958.

[53] N. E. Karoui, The spectrum of kernel random matrices, Ann. Statist. **38**, 1 (2010).

[54] J. Pennington and P. Worah, Nonlinear random matrix theory for deep learning, *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17 (Curran Associates Inc., USA, 2018), pp. 2634–2643.

[55] S. Mei and A. Montanari, The generalization error of random features regression: Precise asymptotics and double descent curve, Commun. Pure Appl. Math. **75**, 667 (2022).

[56] F. Gerace, B. Loureiro, F. Krzakala, M. Mézard, and L. Zdeborová, Generalisation error in learning with random features and the hidden manifold model, J. Stat. Mech. (2021) 124013.

[57] O. Dhifallah and Y. M. Lu, A precise performance analysis of learning with random features, arXiv:2008.11904 (2020).

[58] S. Goldt, B. Loureiro, G. Reeves, F. Krzakala, M. Mezard, and L. Zdeborova, The gaussian equivalence of generative models for learning with shallow neural networks, in *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, edited by J. Bruna, J. Hesthaven, and L. Zdeborova, Vol. 145 of Proceedings of Machine Learning Research (PMLR, 2022), pp. 426–471.

[59] H. Hu and Y. M. Lu, Universality laws for high-dimensional learning with random features, IEEE Trans. Inf. Theory **69**, 1932 (2023).

[60] A. Montanari and B. Saeed, Universality of empirical risk minimization, in *Proceedings of Thirty Fifth Conference on Learning Theory*, edited by L. Po-Ling and R. Maxim, Proceedings of Machine Learning Research Vol. 178 (PMLR, 2022), pp. 4310–4312.

[61] A. Montanari and P.-M. Nguyen, Universality of the elastic net error, *2017 IEEE International Symposium on Information Theory (ISIT)* (IEEE, 2017), pp. 2338–2342.

[62] S. B. Korada and A. Montanari, Applications of the lindeberg principle in communications and statistical learning, IEEE Trans. Inf. Theor. **57**, 2440 (2011).

[63] E. Abbasi, F. Salehi, and B. Hassibi, Universality in learning from linear measurements, *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (Curran Associates, USA, 2019), pp. 12372–12382.

[64] A. Panahi and B. Hassibi, A universal analysis of large-scale regularized least squares solutions, *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17 (Curran Associates, USA, 2017), pp. 3384–3393.

[65] C. Louart, Z. Liao, and R. Couillet, A random matrix approach to neural networks, Ann. Appl. Probab. **28**, 1190 (2018).

[66] Y. Dandi, F. Krzakala, B. Loureiro, L. Stephan, and L. Zdeborova, Universality laws for gaussian mixtures in generalized linear models, arXiv:2302.08933 (2023).

[67] Q. Li and H. Sompolinsky, Statistical mechanics of deep linear neural networks: The backpropagating kernel renormalization, Phys. Rev. X **11**, 031059 (2021).

[68] H. Cui, B. Loureiro, F. Krzakala, and L. Zdeborová, Generalization error rates in kernel regression: The crossover from the noiseless to noisy regime, J. Stat. Mech. (2022) 114004.

[69] A. Ingrosso and S. Goldt, Data-driven emergence of convolutional structure in neural networks, Proc. Natl. Acad. Sci. **119**, e2201854119 (2022).

[70] L. Torrey and J. Shavlik, Transfer learning, *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques* (IGI Global, 2010), pp. 242–264.

[71] F. Gerace, L. Saglietti, S. S. Mannelli, A. Saxe, and L. Zdeborová, Probing transfer learning with a model of synthetic correlated datasets, Machine Learn.: Sci. Technol. **3**, 015030 (2022).

[72] D. Wu and J. Xu, On the optimal weighted ℓ_2 regularization in overparameterized linear regression, *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20 (Curran Associates Inc., Canada, 2020).

[73] A. Jacot, B. Şimşek, F. Spadaro, C. Hongler, and F. Gabriel, Kernel alignment risk estimator: Risk prediction from training data, *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20 (Curran Associates Inc., USA, 2020), pp. 15568–15578.

[74] M. Celentano, A. Montanari, and Y. Wei, The Lasso with general Gaussian designs with applications to hypothesis testing, Ann. Statist. **51**, 2194 (2023).

[75] B. Aubin, F. Krzakala, Y. M. Lu, and L. Zdeborová, Generalization error in high-dimensional perceptrons: Approaching bayes error with convex optimization, *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20 (Curran Associates Inc., Canada, 2020).

[76] B. Loureiro, C. Gerbelot, M. Refinetti, G. Sicuro, and F. Krzakala, Fluctuations, bias, variance & ensemble of learners: Exact asymptotics for convex losses in high-dimension, in *Proceedings of the 39th International Conference on Machine Learning*, edited by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Proceedings of Machine Learning Research Vol. 162 (PMLR, USA, 2022), pp. 14283–14314.

[77] G. R. Kini and C. Thrampoulidis, Phase transitions for one-vs-one and one-vs-all linear separability in multiclass gaussian mixtures, *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, 2021), pp. 4020–4024.

[78] H. Sifaou, A. Kammoun, and M.-S. Alouini, Phase transition in the hard-margin support vector machines, *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)* (IEEE, 2019), pp. 415–419.

[79] X. Mai, Z. Liao, and R. Couillet, A large scale analysis of logistic regression: Asymptotic performance and new insights, *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, 2019), pp. 3357–3361.

[80] F. Mignacco, F. Krzakala, Y. Lu, P. Urbani, and L. Zdeborova, The role of regularization in classification of high-dimensional noisy gaussian mixture, *International Conference on Machine Learning* (PMLR, 2020), pp. 6874–6883.

[81] H. Taheri, R. Pedarsani, and C. Thrampoulidis, Optimality of least-squares for classification in gaussian-mixture models, *2020 IEEE International Symposium on Information Theory (ISIT)* (IEEE, 2020), pp. 2515–2520.

[82] E. Dobriban, H. Hassani, D. Hong, and A. Robey, Provable tradeoffs in adversarially robust classification, IEEE Trans. Inf. Theory **69**, 7793 (2023).

[83] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* (World Scientific, Singapore, 1987), Vol. 9.

[84] M. Mézard, G. Parisi, and M. A. Virasoro, Sk model: The replica solution without replicas, Europhys. Lett. **1**, 77 (1986).

[85] N. El Karoui, D. Bean, P. J. Bickel, C. Lim, and B. Yu, On robust regression with high-dimensional predictors, Proc. Natl. Acad. Sci. USA **110**, 14557 (2013).

[86] M. Bayati and A. Montanari, The dynamics of message passing on dense graphs, with applications to compressed sensing, IEEE Trans. Inf. Theor. **57**, 764 (2011).

[87] M. El A. Seddik, C. Louart, R. Couillet, and M. Tamaazousti, The unexpected deterministic and universal behavior of large softmax classifiers, *International Conference on Artificial Intelligence and Statistics* (PMLR, 2021), pp. 1045–1053.

[88] Also called homoskedastic Gaussians, as opposed to heteroskedastic Gaussians.

[89] U. M. Tomasini, A. Sclocchi, and M. Wyart, Failure and success of the spectral bias prediction for kernel ridge regression: the case of low-dimensional data, in *Proceedings of the 39th International Conference on Machine Learning*, edited by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Proceedings of Machine Learning Research Vol. 162 (PMLR, USA, 2022), pp. 21548–21583.

[90] J. Ba, M. A. Erdogdu, T. Suzuki, Z. Wang, D. Wu, and G. Yang, High-dimensional asymptotics of feature learning: How one gradient step improves the representation, in *Advances in Neural Information Processing Systems*, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Vol. 35 (Curran Associates, Inc., USA, 2022), pp. 37932–37946.

[91] V. Papyan, X. Y. Han, and D. L. Donoho, Prevalence of neural collapse during the terminal phase of deep learning training, Proc. Natl. Acad. Sci. (USA) **117**, 24652 (2020).

[92] N. Barkai and H. Sompolinsky, Statistical mechanics of the maximum-likelihood density estimation, Phys. Rev. E **50**, 1766 (1994).

[93] T. Lesieur, C. De Bacco, J. Banks, F. Krzakala, C. Moore, and L. Zdeborová, Phase transitions and optimal algorithms in high-dimensional gaussian mixture clustering, *54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)* (IEEE, 2016), pp. 601–608.

[94] M. Lelarge and L. Miolane, Asymptotic bayes risk for gaussian mixture in a semi-supervised setting, *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)* (IEEE, 2019), pp. 639–643.

[95] K. Wang and C. Thrampoulidis, Benign overfitting in binary classification of gaussian mixtures, *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, 2021), pp. 4030–4034.

[96] The general case is given by changing p for the rank of the design matrix.

[97] As discussed in Theorem 4, the fact that the loss is independent of Σ in this regime can be directly seen from the optimization.

[98] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, Imagenet: A large-scale hierarchical image database, *2009 IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2009), pp. 248–255.

[99] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai *et al.*, Pytorch: An imperative style, high-performance deep learning library, *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (Curran Associates Inc., USA, 2019). pp. 8026–8037.

[100] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Scikit-learn: Machine learning in Python, J. Machine Learn. Res. **12**, 2825 (2011).

[101] A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd, A rewriting system for convex optimization problems, J. Control Decision **5**, 42 (2018).

[102] S. Diamond and S. Boyd, CVXPY: A Python-embedded modeling language for convex optimization, J. Machine Learn. Res. **17**, 1 (2016).