
Classification of Heavy-tailed Features in High Dimensions: a Superstatistical Approach

Urte Adomaityte
Department of Mathematics
King's College London
urte.adomaityte@kcl.ac.uk

Gabriele Sicuro
Department of Mathematics
King's College London
gabriele.sicuro@kcl.ac.uk

Pierpaolo Vivo
Department of Mathematics
King's College London
pierpaolo.vivo@kcl.ac.uk

Abstract

We characterise the learning of a mixture of two clouds of data points with generic centroids via empirical risk minimisation in the high dimensional regime, under the assumptions of generic convex loss and convex regularisation. Each cloud of data points is obtained via a double-stochastic process, where the sample is drawn from a Gaussian distribution whose variance is itself a random parameter sampled from a scalar distribution ϱ . As a result, our analysis covers a large family of data distributions, including the case of power-law-tailed distributions with no covariance, and allows us to test recent “Gaussian universality” claims. We study the generalisation performance of the obtained estimator, we analyse the role of regularisation, and we analytically characterise the separability transition.

1 Introduction

Generalised linear models (GLMs) are still ubiquitous in the theoretical research on machine learning, despite their simplicity. Their nontrivial phenomenology is often amenable to complete analytical treatment and has been key to understanding the unexpected behavior of large and complex architectures. Random features models [61, 47, 24], for example, allowed to clarify many aspects of the well-known double-descent phenomenon in neural networks [54, 8, 9]. A line of research spanning more than three decades [68, 74] has considered a variety of GLMs to investigate a number of aspects of learning in high dimensions. Yet, a crucial assumption adopted in many such theoretical models is that the covariates are obtained from a *Gaussian* distribution, or from a mixture of Gaussian distributions [42, 49, 3, 37, 43]. Although such a Gaussian design has served as a convenient working hypothesis (in some cases experimentally and theoretically justified [67, 44]), it is however not obvious how much it limits the scope of the results. It is reasonable to expect structure, heavy tails, and large fluctuations in real (often non-Gaussian) data [1, 62] to play an important role in the learning process, and it would be therefore desirable to include structured and heavy-tailed-distributed covariates in our theoretical toolbox.

This paper presents, to the best of our knowledge for the first time, the exact asymptotics for classification tasks on covariates obtained from a mixture of heavy-tailed distributions. We focus on supervised binary classification assuming that the sample size n and the dimensionality d of the space where the covariates live are both sent to infinity, keeping their ratio $n/d = \alpha$ fixed. The paper fits therefore in the line of works on exact high-dimensional asymptotics for classification learning via a GLM [49, 43], but, crucially, we relax the usual Gaussian data hypothesis by including in our

analysis power-law-tailed distributions with possibly no covariance. The mixture is obtained from two distributions, each centered around a centroid $\mu \in \mathbb{R}^d$ and resulting from a double stochastic process. Namely, each sample point is obtained from a distribution $\mathcal{N}(\mu, \Sigma)$ whose covariance $\Sigma = \Delta \mathbf{I}_d \in \mathbb{R}^{d \times d}$ is itself a random variable so that Δ has density ρ , supported on \mathbb{R}_*^+ . Using, for example, an appropriately parametrised inverse-Gamma distribution for ρ , such a *superstatistical* construction (as known in the physics literature [5, 7]) can provide, for example, a heavy-tailed, Cauchy-like data distribution with infinite covariance. The replica method [48], an analytical tool widely adopted in statistical physics, provides asymptotic formulas for a generic density ρ , allowing us to study the effects of non-Gaussianity on the performance curves, and test Gaussian universality hypotheses [59] in such a classification task.

Motivation and related works Learning a rule to classify data points clustered in clouds in high dimension is a classical problem in statistics [32]. Its ubiquity is exemplified by the recently observed neural collapse phenomenon in deep neural networks, in which the last layer classifier was found to operate on features clustered in clouds around the vertices of a simplex [38, 56]; more recently, Seddik et al. [67] showed that Gaussian mixtures suitably describe the deep learning representation of GAN data. In theoretical models, data points are typically assumed to be organised in K Gaussian clouds, each corresponding to a class. Each class k , $k \in \{1, \dots, K\}$, is centered around a mean vector μ_k and has covariance Σ_k . The binary classification case, $K = 2$, is the simplest, and possibly most studied, setting. For this case, Mai and Liao [46] considered an ERM task with generic convex loss and ridge regularisation in the high-dimensional regime $n, d \rightarrow +\infty$ with $n/d \in (0, +\infty)$. In this setting, they gave a precise prediction of the classification error, showing the optimality of the square loss in the unregularised case. Their results have been extended by Mignacco et al. [49], who showed that the presence of a regularisation can actually drastically affect the performance, improving it. In the same setting, the Bayes optimal estimator has been studied in both the supervised and the semi-supervised setting [49, 41]; almost-Bayes-optimal solutions have been put in relation to wide flat landscape regions [3]. In the context of binary classification, the maximum number n of samples that can be perfectly fitted by a linear model in dimension d [49, 16, 37] has been the topic of investigation since the seminal work of Cover [12] and is related to the classical storage capacity problems on networks [20, 21, 39]. The corresponding separability transition value $\alpha = n/d$ is remarkably associated with the existence transition of the maximum likelihood estimator [70, 76]. The precise asymptotic characterisation of the test error in learning to classify $K \geq 2$ Gaussian clouds with generic means and covariances has been recently obtained by Loureiro et al. [43]. Within this line of research, rigorous results have been obtained by a variety of methods, such as Gordon’s inequality technique [31, 71, 49] or mapping to approximate message passing schemes [4, 36, 10, 43, 27].

As previously mentioned, the working hypothesis of *Gaussian design* is widely adopted in high-dimensional classification problems. On top of being a convenient technical hypothesis, this assumption has been justified in terms of a “Gaussian universality” principle. In other words, in a number of circumstances, non-Gaussian features distributions are found to be effectively described by Gaussian ones with matching first and second moments as far as the asymptotic properties of the estimators are concerned [44]. Such a “universality property” has been rigorously proven for example in compressed sensing [50] or in the case of LASSO with non-Gaussian dictionaries [55]. A Gaussian equivalence principle introduced by Goldt et al. [30] has been proven to hold for a wide class of generalised linear estimation problems [47]. Extending work by Hu and Lu [34], Montanari and Saeed [51] recently proved such a principle in a GLM under the assumption of *pointwise normality* of the distribution of the features. This crucial assumption can be intended, roughly speaking, as the assumption of sub-Gaussian decay of a marginal of the feature distribution in any direction (see also [25, 13] for further developments).

The Gaussian universality principle, however, can break down by relaxing some of the aforementioned assumptions on the feature distribution. Montanari and Saeed [51] for example showed that pointwise normality is a *necessary* hypothesis for universality. Studying regression tasks on an elliptically distributed dataset, El Karoui [18] showed that claims of “universality” obtained in the Gaussian setting require serious scrutiny, as the statistics of the estimators might strongly depend on the non-Gaussianity of the covariates [18, 69, 72]. More recently, Pesce et al. [59] showed that a structured Gaussian mixture cannot be effectively described by a single Gaussian cloud. Building upon a series of contributions related to the asymptotic performance of models in the proportional high-dimensional limit [33, 47, 28, 24, 30, 44], we aim precisely to explore the validity of Gaussian universality within classification problems.

In this paper, we work in a “superstatistical” data setting [5, 7], meaning that we superpose the statistics of Gaussian data distribution with an assumed distribution of its variance. Such a construction is adopted in a number of disciplines and contexts to go beyond Gaussianity, albeit is known under different names. In statistical physics, it is known as “superstatistics” and is employed in the analysis of non-equilibrium and non-linear systems [6]. In Bayesian modeling, it is common to refer to hierarchical priors and models [22, 23], while in probability, statistics, and actuarial sciences such distributions are known as compound probability distributions [63], or doubly-stochastic models [60, 66]. Crucially, this construction allows us to consider a very large family of non-Gaussian distributions, which include, but are not limited to, any power-law decay and Cauchy-like with possible infinite-covariance parametrisations.

Our contributions In this manuscript we provide the following results.

- We study a classification task on a non-Gaussian mixture model (see Eq. (1) below) via a generalised linear model (GLM) and we analytically derive, using the replica method [14, 19, 48], an asymptotic characterisation of the statistics of the empirical risk minimisation (ERM) estimator. Our results go therefore beyond the usual Gaussian assumption for the dataset, and include for example the case of covariates obtained from a mixture of distributions with infinite variance. The analysis is performed in the high-dimensional, proportional limit and for any convex loss and convex regularisation. By using this result, we provide asymptotic formulas for the generalisation, training errors and training loss.
- We analyse the performance of this ERM task on a specific family of dataset distributions by using different convex loss functions (quadratic and logistic) and ridge regularisation. We show in particular that, in the case of two balanced clusters with a specific non-Gaussian distribution, the optimal ridge regularisation strength λ^* is finite, at odds with the Gaussian case, for which $\lambda^* \rightarrow \infty$ [49]. In this respect, by considering distributions with matching first and second moments, we analytically show that the performances of the analysed GLM do depend in general on higher moments, and therefore the “Gaussian universality principle” breaks down when heavy-tailed distributions are considered.
- We derive the separability threshold on a large family of non-Gaussian dataset distributions, possibly with unbounded covariance, generalising the known asymptotics for the separability of Gaussian clouds [49]. The result of Cover [12] is recovered in the case of infinite distribution width.
- Under some moment conditions, we derive the Bayes-optimal performance in the case of binary classifications with symmetric centroids, generalising the argument in Ref. [49].
- We finally extend recent results on Gaussian universality in the Gaussian mixture model with random labels [25, 59] to the case of non-Gaussian distributions. We show that the universal formula for the square training loss at zero regularisation found in Ref. [25] holds in our more general setting as well.

2 Main result

Dataset construction We consider the task of classifying two data clusters in the d -dimensional space \mathbb{R}^d . The dataset $\mathcal{D} := \{(\mathbf{x}^v, y^v)\}_{v \in [n]}$ is obtained by extracting n independent datapoints \mathbf{x}^v , each associated with a label $y^v \in \{-1, 1\}$ (a more general setting, involving a multiclass classification task, is discussed in Appendix A). The data points are correlated with the labels via a law $P(\mathbf{x}, y)$ which we assume to have the form

$$P(\mathbf{x}, y) = \delta_{y,1} \rho P(\mathbf{x}|\boldsymbol{\mu}_+) + \delta_{y,-1} (1 - \rho) P(\mathbf{x}|\boldsymbol{\mu}_-), \quad \rho \in (0, 1), \quad \boldsymbol{\mu}_\pm \in \mathbb{R}^d. \quad (1a)$$

Here $P(\mathbf{x}|\boldsymbol{\mu})$ is a distribution with mean $\boldsymbol{\mu}$, and the mean vectors $\boldsymbol{\mu}_\pm \in \mathbb{R}^d$ are distributed according to some density, such that $\mathbb{E}[\|\boldsymbol{\mu}\|^2] = \Theta(1)$, and correspond to the center of the two clusters. The scalar quantity ρ weighs the relative contribution of the two clusters: in the following, we will denote $\rho_+ = \rho = 1 - \rho_-$. Each cluster distribution $P(\mathbf{x}|\boldsymbol{\mu})$ around a vector $\boldsymbol{\mu}$ is assumed to have the form

$$P(\mathbf{x}|\boldsymbol{\mu}) := \mathbb{E}_\Delta [\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \Delta \mathbf{I}_d)], \quad (1b)$$

where $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is a Gaussian distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$, and Δ is randomly distributed with some density ϱ with support on $\mathbb{R}_*^+ := (0, +\infty)$. The family of “elliptic-like” distributions in Eq. (1b) has been extensively studied, for instance, by the physics community, in the context of *superstatistics* [5, 7]. Mixtures of normals in the form of Eq. 1b are a central tool in Bayesian statistics [65] due to their ability to approximate any distribution given a sufficient number

of components [53, 2, 29]. Although the family in Eq. (1b) is not the most general of such mixtures, it is sufficient to include a large family of power-law-tailed densities and to allow us to go beyond the usual Gaussian approximation. El Karoui [18] considered the statistical properties, in the same asymptotic proportional regime as here, of ridge-regularised regression estimators on datasets with distribution as in Eq. 1b, under the assumption (here relaxed) that $\mathbb{E}[\Delta^2] < +\infty$. This *elliptic* family includes a large class of distributions with properties markedly different from the ones of Gaussians. For example, the inverse-Gamma with density $\varrho(\Delta) = (2\pi\Delta^3)^{-1/2} e^{-\frac{1}{2\Delta}}$ leads to a d -dimensional Cauchy-like distribution $P(\mathbf{x}|\boldsymbol{\mu}) \propto (1 - \|\boldsymbol{\mu} - \mathbf{x}\|^2)^{-\frac{d+1}{2}}$, having as marginals Cauchy distributions and $\mathbb{E}[\|\mathbf{x} - \boldsymbol{\mu}\|^2] = +\infty$. We will perform our classification task by searching for a set of parameters (\mathbf{w}^*, b^*) , called *weights* and *bias*, respectively, that will allow us to construct an estimator via a certain classifier $\varphi: \mathbb{R} \rightarrow \{-1, 1\}$

$$\mathbf{x} \mapsto \varphi \left(\frac{\mathbf{x}^\top \mathbf{w}^*}{\sqrt{d}} + b^* \right). \quad (2)$$

By means of the above law, we will predict the label for a new, unseen data point \mathbf{x} sampled from the same law $P(\mathbf{x}, y)$. Our analysis will be performed in the high-dimensional limit where both the sample size n and dimensionality d are sent to infinity, with $n/d \equiv \alpha$ kept constant.

Learning task In the most general setting, the parameters are estimated by minimising an empirical risk function in the form

$$(\mathbf{w}^*, b^*) \equiv \arg \min_{\substack{\mathbf{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} \mathcal{R}(\mathbf{w}, b) \quad \text{where} \quad \mathcal{R}(\mathbf{w}, b) \equiv \sum_{\nu=1}^n \ell \left(y^\nu, \frac{\mathbf{w}^\top \mathbf{x}^\nu}{\sqrt{d}} + b \right) + \lambda r(\mathbf{w}). \quad (3)$$

Here ℓ is a strictly convex loss function with respect to its second argument, and r is a strictly convex regularisation function with the parameter $\lambda \geq 0$ tuning its strength.

State evolution equations. Let us now present our main result, namely the exact asymptotic characterisation of the distribution of the estimator $\mathbf{w}^* \in \mathbb{R}^d$ and of $1/\sqrt{d} \mathbf{w}^* \top \mathbf{X} \in \mathbb{R}^n$, where $\mathbf{X} \in \mathbb{R}^{d \times n}$ is the concatenation of the n dataset column vectors $\mathbf{x}^\nu \in \mathbb{R}^d$, $\nu \in [n]$. Such an asymptotic characterisation is performed via a set of order parameters satisfying a system of self-consistent “state-evolution” equations, which we will solve numerically. The asymptotic expressions for generalisation and training errors, in particular, are written in terms of their solutions, the order parameters. In the following, we use the shorthand $\ell_\pm(u) \equiv \ell(\pm 1, u)$, and given a function $\Phi: \{-1, 1\} \rightarrow \mathbb{R}$, we write $\Phi_\pm \equiv \Phi(\pm 1)$ and $\mathbb{E}_\pm[\Phi_\pm] := \rho_+ \Phi_+ + \rho_- \Phi_-$.

Result 2.1 Let $\zeta \sim \mathcal{N}(0, 1)$ and $\Delta \sim \varrho$, both independent from other quantities. Let also be $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. In the setting described above, given $\phi_1: \mathbb{R}^d \rightarrow \mathbb{R}$ and $\phi_2: \mathbb{R}^n \rightarrow \mathbb{R}$, the estimator \mathbf{w}^* and the vector $\mathbf{z}^* := \frac{1}{\sqrt{d}} \mathbf{w}^* \top \mathbf{X} \in \mathbb{R}^n$ verify:

$$\phi_1(\mathbf{w}^*) \xrightarrow[n, d \rightarrow +\infty]{\text{P}} \mathbb{E}_\xi [\phi_1(\mathbf{g})], \quad \phi_2(\mathbf{z}^*) \xrightarrow[n, d \rightarrow +\infty]{\text{P}} \mathbb{E}_\zeta [\phi_2(\mathbf{h})], \quad (4)$$

where we have introduced the proximal $\mathbf{g} \in \mathbb{R}^d$, defined as

$$\mathbf{g} := \arg \min_{\mathbf{w}} \left(\hat{v} \frac{\|\mathbf{w}\|^2}{2} - \sqrt{d} \sum_{k=\pm} \hat{m}_k \mathbf{w}^\top \boldsymbol{\mu}_k - \sqrt{\hat{q}} \boldsymbol{\xi}^\top \mathbf{w} + \lambda r(\mathbf{w}) \right) \quad (5)$$

and where we also have introduced the proximal for the loss $\mathbf{h} \in \mathbb{R}^n$, obtained by concatenating $\rho_+ n$ quantities h_+ with $\rho_- n$ quantities h_- , with h_\pm given by

$$h_\pm := \arg \min_u \left[\frac{(u - \omega_\pm)^2}{2\Delta v} + \ell_\pm(u) \right], \quad \text{where } \omega_\pm := m_\pm + b + \sqrt{q} \Delta \zeta. \quad (6)$$

The collection of parameters $(q, m_\pm, v, \hat{q}, \hat{m}_\pm, \hat{v}, b)$ appearing in the equations above is given by the fixed-point solution of the following self-consistent equations:

$$\begin{cases} m_\pm = \frac{1}{\sqrt{d}} \mathbb{E}_\xi [\mathbf{g}^\top \boldsymbol{\mu}_\pm], \\ q = \frac{1}{d} \mathbb{E}_\xi [\|\mathbf{g}\|^2], \\ v = \frac{1}{d} \hat{q}^{-1/2} \mathbb{E}_\xi [\mathbf{g}^\top \boldsymbol{\xi}], \end{cases} \quad \begin{cases} \hat{q} = \alpha \mathbb{E}_{\pm, \zeta, \Delta} [\Delta f_\pm^2], \\ \hat{v} = -\alpha q^{-1/2} \mathbb{E}_{\pm, \Delta, \zeta} [\sqrt{\Delta} f_\pm \zeta], \\ \hat{m}_\pm = \alpha \rho_\pm \mathbb{E}_{\Delta, \zeta} [f_\pm], \end{cases} \quad \begin{cases} f_\pm := \frac{h_\pm - \omega_\pm}{v \Delta}, \\ b = \mathbb{E}_{\pm, \Delta, \zeta} [h_\pm - m_\pm]. \end{cases} \quad (7)$$

The derivation of the result above is given, in a multiclass setting, in Appendix A using the replica method, which can be put on rigorous ground by a mapping of the risk minimisation problem into an approximate message-passing iteration scheme [17, 43, 26]. As in the purely Gaussian case, the state evolution equations naturally split the dependence on the loss function ℓ , which is relevant in the computation of h_{\pm} , and on the regularisation function r , entering in the computation of \mathbf{g} . In the most general setting, we are required to solve two convex minimisation problems, consisting of the computation of the proximals h_{\pm} and \mathbf{g} [57]; the convexity of the problem guarantees the uniqueness of the fixed point solution. Result 2.1 implies the following.

Result 2.2 *Under the assumptions of Result 2.1, the training loss ϵ_{ℓ} , the training error ϵ_t , and the test (or generalisation) error ϵ_g are given, in the proportional asymptotic limit, by*

$$\begin{aligned} \frac{1}{n} \sum_{v=1}^n \ell\left(y^v, \frac{\mathbf{w}^{*\top} \mathbf{x}^v}{\sqrt{d}} + b^*\right) &\rightarrow \mathbb{E}_{\pm, \zeta, \Delta}[\ell_{\pm}(h_{\pm})] =: \epsilon_{\ell}, \\ \frac{1}{n} \sum_{v=1}^n \mathbb{I}\left(\varphi\left(\frac{\mathbf{w}^{*\top} \mathbf{x}^v}{\sqrt{d}} + b^*\right) \neq y^v\right) &\rightarrow \mathbb{E}_{\pm, \zeta, \Delta}[\mathbb{I}(\varphi(h_{\pm}) \neq \pm 1)] =: \epsilon_t, \\ \mathbb{E}_{(y, \mathbf{x})} \left[\mathbb{I}\left(\varphi\left(\frac{\mathbf{w}^{*\top} \mathbf{x}}{\sqrt{d}} + b^*\right) \neq y\right)\right] &= \mathbb{E}_{\pm, \Delta, \zeta}[\mathbb{I}(\varphi(\omega_{\pm}) \neq \pm 1)] =: \epsilon_g. \end{aligned} \quad (8)$$

Result 2.3 *Under the hypothesis that $\mathbb{E}[\Delta] < +\infty$ and $\mathbb{E}[\Delta^{-2}] < +\infty$, the Bayes optimal test error for binary classification is*

$$\epsilon_g^{\text{BO}} = \rho_+ \mathbb{E}[\hat{\Phi}(\kappa_+)] + \rho_- \mathbb{E}[\hat{\Phi}(\kappa_-)], \quad \kappa_{\pm} := \frac{1}{\sqrt{\Delta_0^* \left(1 + \frac{\mathbb{E}[\Delta] \mathbb{E}[\Delta^{-2}]}{\alpha \mathbb{E}[\Delta^{-1}]^2}\right)}} \left(1 \pm \frac{\Delta_0}{2} \left(1 + \frac{1}{\alpha \mathbb{E}[\Delta^{-1}]}\right) \ln \frac{\rho_+}{\rho_-}\right) \quad (9)$$

where the expectation is over the two identically distributed random variables $\Delta_0, \Delta_0^* \sim \varrho$, and $\hat{\Phi}(x)$ is the complementary error function.

Quadratic loss with ridge regularisation. Explicit choices for ℓ or r can highly simplify the fixed-point equations above. As an example, let us consider the case of quadratic loss $\ell(y, x) = \frac{1}{2} (y - x)^2$ with ridge regularisation $r(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|^2$. In this setting, an explicit formula for the proximal h_{\pm} in Eq. (6) and for the proximal \mathbf{g} in Eq. (5) can be easily found. As a result, we can obtain explicit expressions for the corresponding set of state evolution equations which become particularly suitable for a fast numerical solution and explicit the dependence on higher moments of Δ :

$$\begin{cases} m_{\pm} = \frac{\sum_{k=\pm} \hat{m}_k \mu_k^{\top} \mu_{\pm}}{\lambda + \hat{v}}, & \hat{q} = \alpha \mathbb{E}_{\pm}[(\pm 1 - m_{\pm} - b)^2] \mathbb{E}_{\Delta} \left[\frac{\Delta}{(1+v\Delta)^2} \right] + \alpha q \mathbb{E}_{\Delta} \left[\frac{\Delta^2}{(1+v\Delta)^2} \right], \\ q = \frac{\|\sum_{k=\pm} \hat{m}_k \mu_k\|^2 + \hat{q}}{(\lambda + \hat{v})^2}, & \hat{v} = \alpha \mathbb{E}_{\Delta} \left[\frac{\Delta}{1+v\Delta} \right], \\ v = \frac{1}{\lambda + \hat{v}}, & \hat{m}_{\pm} = \alpha \rho_{\pm} (\pm 1 - m_{\pm} - b) \mathbb{E}_{\Delta} \left[\frac{1}{1+v\Delta} \right]. \end{cases} \quad (10)$$

3 Application to synthetic datasets

In this section, we compare our theoretical predictions with the results of numerical experiments for a large family of data distributions. The results have been obtained using ridge ℓ_2 -regularisation, and both quadratic and logistic losses, with various data cluster balances ρ . We will also assume, without loss of generality, that $\mu_{\pm} = \pm \frac{1}{\sqrt{d}} \boldsymbol{\mu}$, where $\boldsymbol{\mu} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. The synthetic data sets will be produced using, an inverse-Gamma-distributed variance Δ , with density parametrised as

$$\varrho(\Delta) \equiv \varrho_{a,c}(\Delta) = \frac{c^a}{\Gamma(a) \Delta^{a+1}} e^{-\frac{c}{\Delta}}, \quad (11a)$$

depending on the shape parameter $a > 0$ and on the scale parameter $c > 0$. As a result, each class will be distributed as

$$P(\mathbf{x}|\boldsymbol{\mu}) = \frac{(2c)^a \Gamma(a + d/2)}{\Gamma(a) \pi^{d/2} (2c + \|\mathbf{x} - \boldsymbol{\mu}\|^2)^{a + d/2}}, \quad (11b)$$

which produces a cloud centered in $\boldsymbol{\mu}$. The reason for our choice is that the distribution $\varrho_{a,c}$ allows us to easily explore different regimes by changing two parameters only. Indeed, the distribution has, for $a > 1$, covariance matrix $\boldsymbol{\Sigma} = \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu}) \otimes (\mathbf{x} - \boldsymbol{\mu})^{\top}] = \sigma^2 \mathbf{I}_d$, where $\sigma^2 = \frac{c}{a-1}$. By fixing σ^2 and taking the limit $a \rightarrow +\infty$, $P(\mathbf{x}|\boldsymbol{\mu}) \rightarrow \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \sigma^2 \mathbf{I}_d)$, i.e., the Gaussian case discussed by Mignacco et al. [49]. For $a \in (0, 1]$, instead, $\sigma^2 = +\infty$. Note that the choice of inverse-Gamma distribution is not that unusual. It has been adopted, for example, to describe non-Gaussian data in quantitative finance [15, 40] or econometrics models [52]. Finally, we will construct our label estimator as in Eq. (2) with $\varphi(x) = \text{sign}(x)$.

3.1 Finite-covariance case

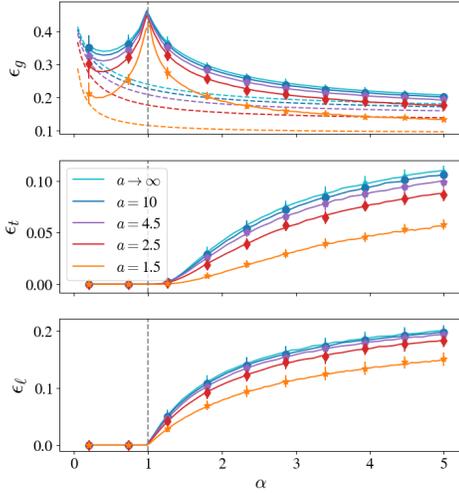


Figure 1: Test error ϵ_g (solid line, top), training error ϵ_t (center) and training loss ϵ_l (bottom) as predicted by Eq. (8) in the balanced $\rho = 1/2$ case. The dataset distribution is parametrised as in Eq. (12). The classification task is solved using a quadratic loss with ridge regularisation with $\lambda = 10^{-5}$. In the top figure, the dashed line corresponds to the Bayes optimal bound. Dots correspond to the average outcome of 50 numerical experiments in dimension $d = 10^3$. In our parametrisation, the population covariance is $\Sigma = \mathbf{I}_d$ for all values of a and moreover, for $a \rightarrow +\infty$, the case of Gaussian clouds with the same centroids and covariance is recovered. For further details on the numerical solutions, see Appendix B.

universality principle holds. This (perhaps) counter-intuitive effect stems from the fact that, although for all classes $\Sigma = \mathbf{I}_d$, in our example the mean absolute deviation is smaller for finite a

$$\lim_{d \rightarrow +\infty} \frac{\mathbb{E}[\|\mathbf{x} - \boldsymbol{\mu}\|]}{\sqrt{\mathbb{E}[\|\mathbf{x} - \boldsymbol{\mu}\|^2]}} = \frac{\Gamma(a-1/2)\sqrt{a-1}}{\Gamma(a)} \xrightarrow{a \rightarrow +\infty} 1^- . \quad (13)$$

The learning process benefits from the fact that as a decreases while keeping the variance fixed, points move on average closer to their means $\boldsymbol{\mu}_\pm$ to compensate heavier tails. Finally, we observed that test errors are systematically smaller for unbalanced clusters than for balanced clusters (see Appendix A.3).

The same numerical experiment was repeated using the logistic loss for training. Once again, in Fig. 2, we focus on $\rho = 1/2$ and show that the theoretical predictions of the generalisation and training errors agree with the results of numerical experiments for a range of sample complexity values α and various values of a . Just as for the square loss, the test error ϵ_g of the Gaussian model ($a \rightarrow +\infty$) is larger than the one observed for power-law distributed clouds at finite values of a . In the $\lambda \rightarrow 0$ limit, the typical interpolation cusp in the generalisation error is observed: the cusp occurs at different values of α as a is changed; interpolation occurs at smaller values of α for larger a (i.e., for “more Gaussian” distributions). A comparison with the test error ϵ_t confirms that this cusp occurs at the value of α where the training error becomes non-zero and the two training data clouds become non-separable [70, 16]. This sharp transition in α also corresponds to the existence transition of maximum-likelihood estimator in high-dimensional logistic regression, analysed for Gaussian data in Refs. [70, 49] and, in our setting, in Section 4 below. For larger regularisation strength, the cusp smoothens, and the training error becomes non-zero at smaller values of α .

Let us start by considering a data distribution as in Eq. (11) with shape parameter $a = c + 1 > 1$, which is therefore given by

$$P(\mathbf{x}|\boldsymbol{\mu}) = \frac{2^a (a-1)^a \Gamma(a+d/2)}{\Gamma(a) \pi^{d/2} (2a-2+\|\mathbf{x}-\boldsymbol{\mu}\|^2)^{a+d/2}} \quad (12)$$

and decays as $\|\mathbf{x}\|^{-2a-1}$ in the radial direction for $\|\mathbf{x} - \boldsymbol{\mu}\| \gg 0$. As a consequence, the distribution has $\lim_{d \rightarrow +\infty} 1/d \mathbb{E}[\|\mathbf{x} - \boldsymbol{\mu}\|^k] = +\infty$ for $k \geq 2a$. With this choice, all elements of the family in Eq. (12) have the same covariance $\Sigma = \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu}) \otimes (\mathbf{x} - \boldsymbol{\mu})^\top] = \mathbf{I}_d$ as a is varied, including the Gaussian limit $P(\mathbf{x}|\boldsymbol{\mu}) \rightarrow \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_d)$ obtained for $a \rightarrow +\infty$.

In Fig. 1 we present the results of our numerical experiments using the square loss and small regularisation. An excellent agreement between the theoretical predictions and the results of numerical experiments is found for a range of values of a and sample complexity α , both for balanced, i.e., equally sized, and unbalanced clusters of data (the plot for this case can be found in Appendix A.3). The test error ϵ_g presents the classical interpolation peak at $\alpha = 1$, smoothed by the presence of a non-zero regularisation strength λ , and the typical double-descent behavior [54]. As a is increased, the results of Mignacco et al. [49] for Gaussian clouds with $P(\mathbf{x}|\boldsymbol{\mu}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \mathbf{I}_d)$ are approached. The $a \rightarrow +\infty$ curves correspond, therefore, to the Gaussian mixture model we would have constructed by simply fitting the first and second moments of each class in the finite- a case. The plot shows that, at given population covariance $\Sigma = \mathbf{I}_d$, the classification of power-law distributed clouds is associated with different, and in particular smaller, errors than the corresponding task on Gaussian clouds with the same covariance: in this sense, no Gaussian

3.2 Infinite-covariance case

The family of distributions specified by Eq. (11) allows us to also consider the case of infinite covariance, i.e., $\sigma^2 = +\infty$. To test our formulas in this setting, we considered the case in which each cloud is obtained by “contaminating” a standard Gaussian with an infinite-covariance distribution as in Eq. (11b) with $c = 1$ and $a < 1$ [35]. In other words, we use the density $\varrho(\Delta) = r\varrho_{a,1}(\Delta) + (1-r)\delta(\Delta - 1)$, with $r \in [0, 1]$ interpolation parameter, for Δ in Eq. (1). Each class has therefore infinite covariance for $0 < r \leq 1$ and the Gaussian case is recovered for $r = 0$. Fig. 3 collects our results for two balanced clouds, analysed using square loss and logistic loss. Good agreement between the theoretical predictions and the results of numerical experiments is found for a range of values of sample complexity and of the ratio r . In this case, the finite-variance case $r = 0$ corresponds, not surprisingly, to the lowest test error with both square loss and logistic loss, which grows with r .

3.3 The role of regularisation in the classification of non-Gaussian clouds

One of the main results in the work of Mignacco et al. [49] is related to the effect of regularisation on the performances in a classification task on a Gaussian mixture. They observed that, for all values of the sample complexity α , the optimal ridge classification performance on a pair of balanced clouds is obtained for an infinite regularisation strength λ , using both hinge loss and square loss. We tested the robustness of this result beyond the purely Gaussian setting. In Fig. 4 (left), we plot the test error obtained using square loss and different regularisation strengths λ on the dataset distribution as in Eq. (12) with $a = 2$. We observe that, both in the balanced case and in the unbalanced case, the optimal regularisation strength λ^* is *finite* and, moreover, α -dependent. On the other hand, if we fix α and let a grow towards $+\infty$, the optimal λ^* grows as well to recover the result of Mignacco et al. [49] of diverging optimal regularisation for *balanced* Gaussian clouds, as shown in Fig. 4 (center)

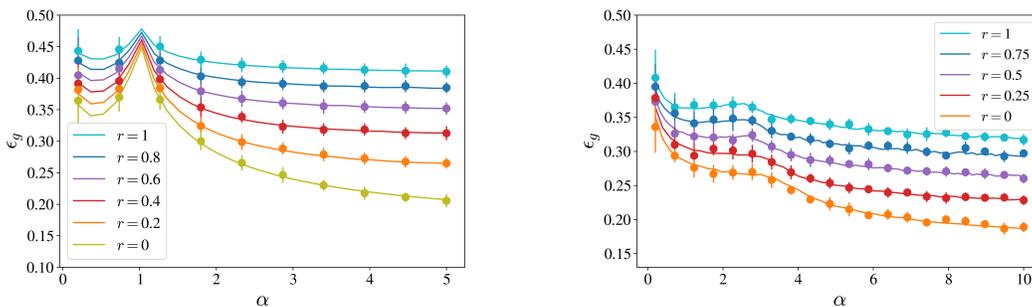


Figure 3: Test error ϵ_g in the classification of two balanced clouds, via quadratic loss (*left*) and logistic loss (*right*). In both cases, ridge regularisation is adopted ($\lambda = 10^{-4}$ for the square loss case, $\lambda = 10^{-3}$ for the logistic loss case). Each cloud is a superposition of a power-law distribution with infinite variance and a Gaussian with covariance $\Sigma = \mathbf{I}_d$. The parameter r allows us to contaminate the purely Gaussian case ($r = 0$) with an infinite-variance contribution ($0 < r \leq 1$) as in Eq. (11) with $c = 1$ and $a = 1/2$ (*left*) or $a = 3/4$ (*right*). Dots correspond to the average test error of 20 numerical experiments in dimension $d = 10^3$. Note that, at a given sample complexity, Gaussian clouds are associated with the lowest test error for both losses.

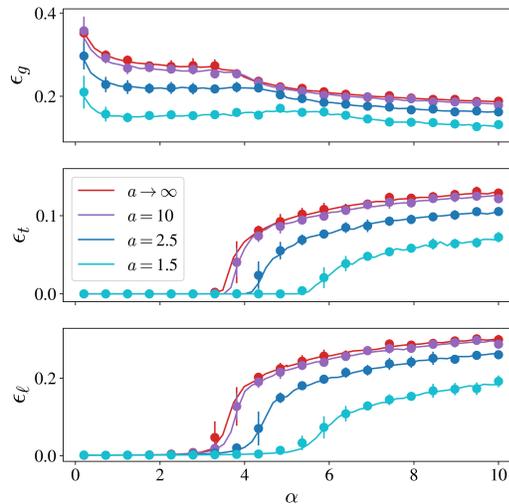


Figure 2: Test error ϵ_g (*top*), training error ϵ_t (*center*) and training loss ϵ_l (*bottom*) via logistic loss training on balanced clusters parametrised as in Eq. (12) ($\Sigma = \mathbf{I}_d$). Ridge regularisation with $\lambda = 10^{-4}$ is adopted. Dots correspond to the average over 20 numerical experiments with $d = 10^3$. The Gaussian limit is recovered for $a \rightarrow +\infty$. Further details on the numerical solutions can be found in Appendix B.

for $\alpha = 2$. This result is represented also in Fig. 4 (right), where it is shown that, in the *unbalanced* case, the optimal regularisation strength λ^* instead saturates to a finite value for $a \rightarrow +\infty$, i.e., for Gaussian clouds [49].

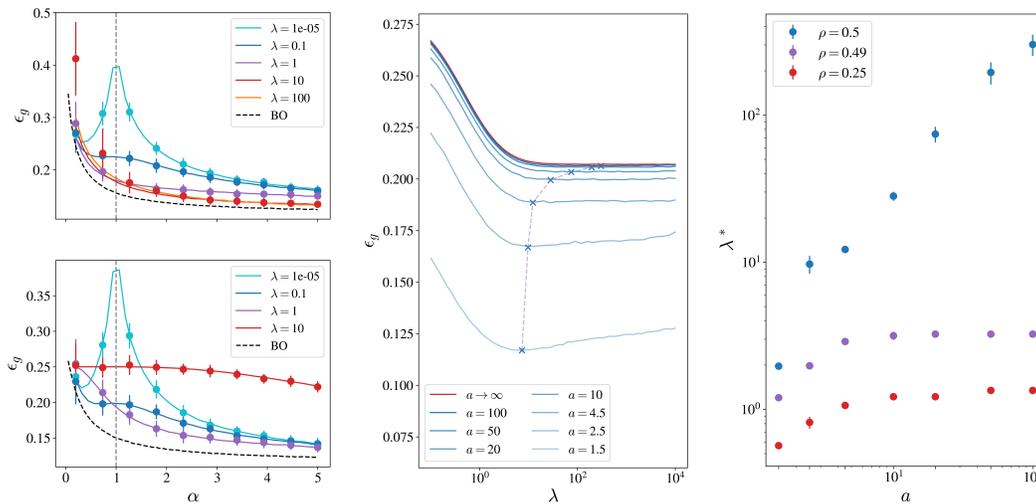


Figure 4: (Left) Test error for ridge regularised quadratic loss for various regularisation strengths. The data points of each cloud in the training set are distributed as in Eq. (12), with shape parameter $a = 2$, for balanced clusters (top) and unbalanced clusters ($\rho = 1/4$, bottom). Points are the results of 50 numerical experiments, and the dashed lines are Bayes-optimal bounds. (Center) Test error for different regularisation strengths λ for two *balanced* clusters with quadratic loss at sample complexity $\alpha = 2$ using the data distribution (12). The optimal regularisation strength value λ^* obtained from averaging 5 runs for each a is marked with a cross. (Right) Optimal regularisation strength λ^* at $\alpha = 2$ for different values of $a \in [1.5, 10^2]$ for both balanced and unbalanced clusters, obtained from averaging 5 runs. Note that, for $\rho = 1/2$, $\lambda^* \rightarrow +\infty$ as $a \rightarrow +\infty$.

4 The separability threshold

By studying the training error ϵ_t with logistic loss at zero regularisation we can obtain information on the boundary between the regime in which the training data are perfectly separable and the regime in which they are not. In other words, we can extract the value of the so-called separability transition complexity α^* [70, 49, 43]. Once again, a complete characterisation of this transition point is available in the Gaussian mixture case, where α^* can be explicitly given as a function of σ^2 [49].

It is not trivial to extend this result to the general case of non-Gaussian mixtures. It is, however, possible to derive α^* within the large family of “superstatistical” models we are considering here under the assumption that Result 2.1 holds. In this case, in Appendix A.3 we prove the following.

Result 4.1 *In the considered double-stochastic model, data are linearly separable for $\alpha < \alpha^*$, where*

$$\alpha^* := \max_{\theta \in (0,1], \gamma} \frac{1-\theta^2}{\mathcal{S}(\theta, \gamma)}, \quad \mathcal{S}(\theta, \gamma) := \int_0^\infty f^2 \mathbb{E}_\Delta \left[\rho_+ \mathcal{N}\left(f + \frac{\theta + \gamma}{\sqrt{\Delta}} \mid 0, 1\right) + \rho_- \mathcal{N}\left(f + \frac{\theta - \gamma}{\sqrt{\Delta}} \mid 0, 1\right) \right] df. \quad (14)$$

Fig. 5 shows the values of α^* for different choices of the shape parameters a and c in the case of two balanced clouds of datapoints distributed as in Eq. (11b) as predicted by Eq. (14). In Fig. 5 (left), we fixed $\sigma^2 = \frac{c}{a-1} < +\infty$, and plotted the separability threshold α^* for a range of values of the shape parameter $a > 1$. As a grows, the known threshold value for the Gaussian mixture case is recovered [49]. The double limit $a \rightarrow \infty$ and $\sigma^2 \rightarrow \infty$, on the other hand, provides the Cover [12] transition value $\alpha^* = 2$, as expected. In Fig. 5 (right), instead, we analyse the case of infinite-variance clusters, by fixing $a \in (0, 1)$ and by varying the scale parameter c in Eq. (11b), which controls the spread (width) of the distribution. The Cover transition is therefore correctly recovered as c diverges for all values of a .

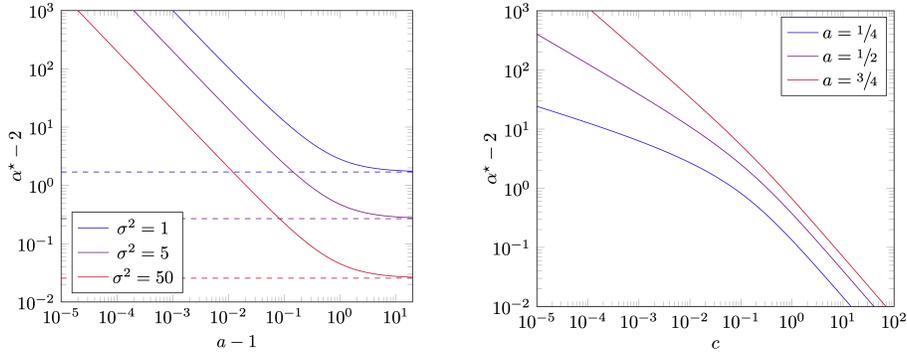


Figure 5: Separability threshold α^* obtained by solving the equations in Eq. (7) with logistic loss, ridge regularisation strength $\lambda = 10^{-5}$ and $\rho = 1/2$. The data points of each cloud are distributed around their mean $\boldsymbol{\mu}$ with a two-parameter distribution as in Eq. (11b). (Left) Finite covariance $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}_d$ case, $\sigma^2 = \frac{c}{a-1}$, as a function of a . Dashed lines are the threshold values of the Gaussian case derived by Mignacco et al. [49]. At large a and large variance σ^2 , the Cover’s transition $\alpha = 2$ for balanced clusters is recovered. (Right) Infinite-variance data clouds obtained by fixing $0 < a < 1$ in Eq. (11b).

5 Random labels in the teacher-student scenario and Gaussian universality

Our analysis can be extended to the case in which a point \mathbf{x} is labeled not according to its class but by a “teacher”, represented by a probabilistic law $P_0(y|\tau)$ and a vector $\boldsymbol{\theta}_0 \in \mathbb{R}^d$ so that each dataset element is independently generated with the joint distribution

$$P(\mathbf{x}, y) = P_0\left(y \left| \frac{\boldsymbol{\theta}_0^\top \mathbf{x}}{\sqrt{d}} \right.\right) \sum_{k=\pm} \rho_k \mathbb{E}[\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \Delta \mathbf{I}_d)]. \quad (15)$$

Let us assume now that $\boldsymbol{\theta}_0$ is such that $\lim_d 1/\sqrt{d} \boldsymbol{\theta}_0^\top \boldsymbol{\mu}_\pm = 0$, i.e., the teacher is “uncorrelated” with the structure of the mixture. In Appendix A.5 we show that, if $\ell(y, \eta) = \ell(-y, -\eta)$ and $P_0(y|\tau) = P_0(-y|-\tau)$, the task is equivalent to a classification problem on a *single* cloud centered in the origin, i.e., to the case $\boldsymbol{\mu}_\pm = \mathbf{0}$, so that $P(\mathbf{x}, y) = P_0(y | 1/\sqrt{d} \boldsymbol{\theta}_0^\top \mathbf{x}) \mathbb{E}[\mathcal{N}(\mathbf{x} | \mathbf{0}, \Delta \mathbf{I}_d)]$. This mean-invariance, recently proven in Ref. [25] in the pure Gaussian setting, holds therefore in our non-Gaussian setting as well.

A special, relevant case of “uncorrelated teacher” is the one of *random labels*, where $P_0(y|\tau) = 1/2(\delta_{y,+1} + \delta_{y,-1})$. This apparently vacuous model is relevant in the study of the storage capacity problem [21, 39, 11] and of the separability threshold [12], and it has also been an effective tool to better understand deep learning [45, 75]. In this case, the following universal result holds with respect to our setup.

Result 5.1 *In the random label setting studied via ridge regression case with zero regularisation, the asymptotic training loss is independent of the distribution of Δ , and given by the universal formula*

$$\frac{1}{2n} \sum_{y=1}^{\infty} \left(y^y - \frac{\mathbf{x}^\top \mathbf{w}^*}{\sqrt{d}} \right)^2 \xrightarrow[n/d=\alpha]{n \rightarrow +\infty} \frac{\alpha-1}{2\alpha} \theta(\alpha-1) =: \epsilon_\ell. \quad (16)$$

The proof of the previous result is given in Appendix A.5. Note that this result does not require $\sigma^2 := \mathbb{E}[\Delta]$ to be finite, and it includes, as a special case, the Gaussian mixture model with random labels considered in Ref. [25, 59]. Therein, Eq. (16) has been obtained under the Gaussian design hypothesis. It turns out therefore that Eq. (16) holds for a much larger ensemble, which includes heavy-tailed distributions with possibly infinite covariance. On the other hand, universality *does not* hold, e.g., for the training error.

In Fig. 6 we exemplify these results by using ridge-regularised square loss. The numerical experiments correspond to a classification task on a dataset of *two* clouds centered in $\boldsymbol{\mu}_\pm, \boldsymbol{\mu}_\pm = -\boldsymbol{\mu}_- \sim \mathcal{N}(\mathbf{0}, 1/d \mathbf{I}_d)$, obtained using the distribution

$$P(\mathbf{x}, y) = \frac{\delta_{y,+1} + \delta_{y,-1}}{4} \sum_{k=\pm} P(\mathbf{x} | \boldsymbol{\mu}_k), \quad (17)$$

where $P(\mathbf{x}|\boldsymbol{\mu}_\pm)$ is the distribution in Eq. (11b) depending on the shape parameters a and c . These results are compared with the theoretical prediction obtained for a *single* cloud, with the same parameters a and c , but centered in the origin, so that $P(\mathbf{x}, y) = \frac{\delta_{y,+1} + \delta_{y,-1}}{2} P(\mathbf{x}|\boldsymbol{\mu} = \mathbf{0})$: the agreement confirms the “mean-invariance” property stated above. As anticipated, as we are considering random labels, the training loss ϵ_ℓ follows Eq. (16) in all cases, whereas the test and training errors are not universal. Note that, in place of the test error (which is trivially $\epsilon_g = 1/2$) we plot the mean square error on the pre-activations $\hat{\epsilon}_g := \mathbb{E}_{(y,\mathbf{x})} [(y - 1/\sqrt{a}\mathbf{x}^\top \mathbf{w}^*)^2]$. Also, observe that this quantity is divergent in the case of infinite covariance, hence is absent from the plot for $a \in (0, 1]$.

6 Conclusions and perspectives

The framework introduced in this work has allowed us to perform an exact asymptotic analysis of the learning of a mixture of two data point clouds with possibly power-law tails. In our dataset, each sample point \mathbf{x} is labeled by $y \in \{-1, 1\}$ and obtained as $\mathbf{x} = y\boldsymbol{\mu} + \sqrt{\Delta}\mathbf{z}$, where $\boldsymbol{\mu} \in \mathbb{R}^d$, $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $\Delta \sim \varrho$. We have derived exact expressions for the test and training errors and training loss achieved by a GLM in the limit of large size n of the dataset and large dimensionality d under the assumption that $\alpha = n/d$ is kept fixed. We have shown that the performance crucially depends on higher-order moments of the data distribution, so that a Gaussian universality assumption is not feasible. Moreover, we have shown the effects of non-Gaussianity on the optimal regularisation strength and derived an exact formula for the critical α^* at which the two clouds become linearly inseparable. Finally, we have considered the so-called random labels setting (in which sample label assignment is randomised) and shown that, in this case, the training square loss takes a universal asymptotic form independent of the distribution of Δ : this result further supports the crucial role of labels’ randomness in some recent universality results [25].

In conclusion, we have provided an exactly solvable toy model exhibiting power-law tails and possibly unbounded covariance, highlighting the fact that a dataset deviation from Gaussianity in high dimensions cannot, in general, be neglected. The application of the introduced theoretical framework to real datasets is an interesting direction to explore in the future: the main difficulty, in this case, is the choice of the best distribution ϱ given the observed dataset, a problem that has a long tradition in the context of Bayesian estimation [2, 23]¹.

Acknowledgements

The authors would like to thank Nikolas Nüsken for stimulating discussions. Also, we would like to thank Bruno Loureiro for his feedback and for his careful reading of the manuscript.

¹In a more simplistic approach, it can be observed that in the case in which the square loss is adopted, equations (10) depend on $\mathbb{E}[(1 + v\Delta)^{-1}]$ and $\mathbb{E}[(1 + v\Delta)^{-2}]$ only. Such quantities can be numerically estimated from the dataset, for example, by empirically evaluating the identity $\mathbb{E}[(1 + v\Delta)^{-1}] = \sqrt{v}\partial_v \int_0^v (v - u)^{-1/2} \mathbb{E}[\exp(-ux^2/2)] du$ when it is assumed that $x \sim \mathbb{E}[\mathcal{N}(0, \Delta)]$.

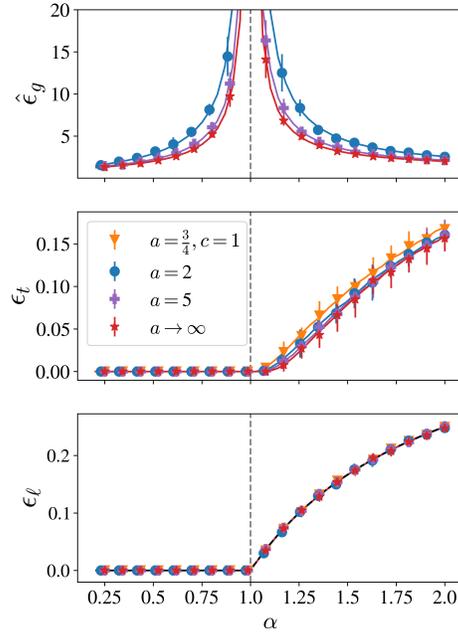


Figure 6: Mean square error $\hat{\epsilon}_g$ (top), training error ϵ_t (center) and training loss ϵ_ℓ (bottom) obtained using square loss on a dataset distributed as in Eq (17) with random labels (dots) compared with the theoretical predictions for the single cluster case (continuous lines). The distributions are parametrised using Eq. (12) for $a > 1$ and Eq. (11b) for $0 < a \leq 1$. A ridge regularisation with $\lambda = 10^{-4}$ is adopted. Note that $\hat{\epsilon}_g$ is absent in the top plot for $a = 1/2$ and $c = 1$ being a divergent quantity in this case.

References

- [1] R.J. Adler, R.E. Feldman, and M.S. Taqqu, editors. *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*. Birkhauser Boston Inc., USA, 1998. ISBN 0817639519.
- [2] D. Alspach and H. Sorenson. Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Trans. Autom. Control*, 17(4):439–448, 1972.
- [3] C. Baldassi, E.M. Malatesta, M. Negri, and R. Zecchina. Wide flat minima and optimal generalization in classifying high-dimensional gaussian mixtures. *J. Stat. Mech.: Theory Exp.*, 2020(12):124012, 2020.
- [4] M. Bayati and A. Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Trans. Inf. Theory*, 57(2):764–785, 2011.
- [5] C. Beck. Superstatistics: Theory and applications. *Contin. Mech. Thermodyn.*, 16, 2003.
- [6] C. Beck. Recent developments in superstatistics. *Braz. J. Phys.*, 39:357–363, 2008.
- [7] C. Beck and E.G.D. Cohen. Superstatistics. *Physica A*, 322:267–275, 2003.
- [8] M. Belkin, D. Hsu, S. Ma, and S. Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proc. Natl. Acad. Sci. U.S.A.*, 116(32):15849–15854, 2019.
- [9] M. Belkin, A. Rakhlin, and A.B. Tsybakov. Does data interpolation contradict statistical optimality? In K. Chaudhuri and M. Sugiyama, editors, *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proc. Mach. Learn. Res.*, pages 1611–1619. PMLR, 2019.
- [10] R. Berthier, A. Montanari, and P.-M. Nguyen. State evolution for approximate message passing with non-separable functions. *Inf. Inference*, 9(1):33–79, 2020.
- [11] N. Brunel, J.-P. Nadal, and G. Toulouse. Information capacity of a perceptron. *JJ. Phys. A: Math. Gen.*, 25(19):5017, 1992.
- [12] T.M. Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Trans. Comput.*, EC-14(3):326–334, 1965.
- [13] Y. Dandi, L. Stephan, F. Krzakala, B. Loureiro, and L. Zdeborová. Universality laws for gaussian mixtures in generalized linear models. *arXiv:2302.08933*, 2023.
- [14] P. del Giudice, S. Franz, and M. Virasoro. Perceptron beyond the limit of capacity. *J. Physique*, 50:121–134, 1989.
- [15] D. Delpini and G. Bormetti. Minimal model of financial stylized facts. *Phys. Rev. E*, 83:041111, 2011.
- [16] A. Deng, Z. Kammoun and C. Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *Inf. Inference*, 11, 2021.
- [17] D.L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proc. Natl. Acad. Sci. U.S.A.*, 106(45):18914–18919, Nov 2009.
- [18] N. El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probab. Theory Relat. Fields*, 170, 2018.
- [19] S. Franz, D.J. Amit, and M.A. Virasoro. Prosopagnosia in high capacity neural networks storing uncorrelated classes. *J. Physique*, 51(5):387–408, 1990.
- [20] E Gardner. The space of interactions in neural network models. *J. Phys. A: Math. Gen.*, 21(1): 257, 1988.
- [21] E. Gardner and B. Derrida. Three unfinished works on the optimal storage capacity of networks. *J. Phys. A: Math. Gen.*, 22(12):1983, 1989.

- [22] A. Gelman and J. Hill. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Analytical Methods for Social Research. Cambridge University Press, 2006.
- [23] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, and D.B. Rubin. *Bayesian Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis, 2013.
- [24] F. Gerace, B. Loureiro, F. Krzakala, M. Mézard, and L. Zdeborová. Generalisation error in learning with random features and the hidden manifold model. In *ICML*, pages 3452–3462, 2020.
- [25] F. Gerace, F. Krzakala, B. Loureiro, L. Stephan, and L. Zdeborová. Gaussian universality of perceptrons with random labels. *arXiv:2205.13303*, 2023.
- [26] C. Gerbelot and R. Berthier. Graph-based approximate message passing iterations. *arXiv preprint arXiv:2109.11905*, 2021.
- [27] C. Gerbelot, A. Abbara, and F. Krzakala. Asymptotic errors for teacher-student convex generalized linear models (or: How to prove kabashima’s replica formula). *IEEE Trans. Inf. Theory*, 69(3):1824–1852, 2023.
- [28] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari. Limitations of lazy training of two-layers neural network. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [29] J.K. Ghosh and R.V. Ramamoorthi. *Bayesian Nonparametrics*. Springer Series in Statistics. Springer New York, 2006.
- [30] 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.
- [31] Y. Gordon. Some inequalities for gaussian processes and applications. *Isr. J. Math.*, 50:265–289, 1985.
- [32] T. Hastie, R. Tibshirani, and J.H. Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.
- [33] T. Hastie, A. Montanari, S. Rosset, and R.J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *Ann. Stat.*, 50(2):949 – 986, 2022.
- [34] H. Hu and Y.M. Lu. Universality laws for high-dimensional learning with random features. *arXiv:2009.07669*, 2022.
- [35] P.J. Huber. Robust estimation of a location parameter. *Ann. Math. Stat.*, 35(1):73–101, 1964.
- [36] A. Javanmard and A. Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Inf. Inference*, 2(2):115–144, 2013.
- [37] G.R. Kini and C. Thrampoulidis. Phase transitions for one-vs-one and one-vs-all linear separability in multiclass gaussian mixtures. In *ICASSP 2021*, pages 4020–4024, 2021.
- [38] V. Kothapalli, E. Rasromani, and V. Awatramani. Neural collapse: A review on modelling principles and generalization. *arXiv:2206.04041*, 2022.
- [39] W. Krauth and M. Mézard. Storage capacity of memory networks with binary couplings. *J. Physique*, 50(20):3057–3066, 1989.
- [40] N. Langrené, G. Lee, and Z. Zhu. Switching to non-affine stochastic volatility: A closed-form expansion for the inverse gamma model. *Int. J. Theor. Appl. Finance*, 19, 2015.
- [41] M. Lelarge and L. Miolane. Asymptotic bayes risk for gaussian mixture in a semi-supervised setting. In *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pages 639–643, 2019.

- [42] M. Lelarge and L. Miolane. Asymptotic bayes risk for gaussian mixture in a semi-supervised setting. In *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 639–643, 2019.
- [43] 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. *Advances in Neural Information Processing Systems*, 34:10144–10157, 2021.
- [44] B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mézard, and L. Zdeborová. Learning curves of generic features maps for realistic datasets with a teacher-student model. *J. Stat. Mech.: Theory Exp.*, 2022(11):114001, 2022.
- [45] H. Maennel, I.M. Alabdulmohsin, I.O. Tolstikhin, R. Baldock, O. Bousquet, S. Gelly, and D. Keysers. What do neural networks learn when trained with random labels? In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 19693–19704. Curran Associates, Inc., 2020.
- [46] X. Mai and Z. Liao. High dimensional classification via regularized and unregularized empirical risk minimization: Precise error and optimal loss. *arXiv:1905.13742*, 2020.
- [47] S. Mei and A. Montanari. The Generalization Error of Random Features Regression: Precise Asymptotics and the Double Descent Curve. *Commun. Pure Appl. Math.*, 75, 2019.
- [48] M. Mézard, G. Parisi, and M.A. Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987.
- [49] F. Mignacco, F. Krzakala, Y. Lu, P. Urbani, and L. Zdeborová. The role of regularization in classification of high-dimensional noisy Gaussian mixture. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proc. Mach. Learn. Res.*, pages 6874–6883. PMLR, 2020.
- [50] A. Montanari and P.-M. Nguyen. Universality of the elastic net error. In *2017 IEEE ISIT*, pages 2338–2342. IEEE, 2017.
- [51] A. Montanari and B.N. Saeed. Universality of empirical risk minimization. In Po-Ling Loh and Maxim Raginsky, editors, *Proceedings of 35th Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 4310–4312. PMLR, 2022.
- [52] D.B. Nelson. Arch models as diffusion approximations. *J. Econom.*, 45(1-2):7–38, 1990.
- [53] V. Nestoridis, S. Schmutzhard, and V. Stefanopoulos. Universal series induced by approximate identities and some relevant applications. *J. Approx. Theory*, 163(12):1783–1797, 2011.
- [54] M. Opper, W. Kinzel, J. Kleinz, and R. Nehl. On the ability of the optimal perceptron to generalise. *J. Phys. A: Math. Gen.*, 23(11):L581, 1990.
- [55] A. Panahi and B. Hassibi. A universal analysis of large-scale regularized least squares solutions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [56] V. Pappayan, X.Y. Han, and D.L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proc. Natl. Acad. Sci. U.S.A.*, 117(40):24652–24663, 2020.
- [57] N. Parikh and S. Boyd. Proximal algorithms. *Found. Trends Optim.*, 1(3):127–239, 2014.
- [58] 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. Mach. Learn. Res.*, 12:2825–2830, 2011.
- [59] L. Pesce, F. Krzakala, B. Loureiro, and L. Stephan. Are Gaussian data all you need? Extents and limits of universality in high-dimensional generalized linear estimation. *arXiv:2302.08923*, 2023.

- [60] M.A. Pinsky and S. Karlin. *An Introduction to Stochastic Modeling: Fourth Edition*, pages 1–563. Elsevier Inc, December 2010. ISBN 9780123814166.
- [61] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [62] W.J. Reed and B.D. Hughes. From gene families and genera to incomes and internet file sizes: Why power laws are so common in nature. *Phys. Rev. E*, 66:067103, 2002.
- [63] H.E. Robbins. Asymptotically subminimax solutions of compound statistical decision problems. 1985.
- [64] S. Rosset, J. Zhu, and T. Hastie. Margin maximizing loss functions. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems*, volume 16. MIT Press, 2003.
- [65] P.E. Rossi. *Bayesian Non- and Semi-parametric Methods and Applications*. Princeton University Press, 2014.
- [66] D. Schnoerr, R. Grima, and G. Sanguinetti. Cox process representation and inference for stochastic reaction-diffusion processes. *Nat. Comm.*, 7:11729, 05 2016.
- [67] M.E.A. Seddik, C. Louart, M. Tamaazousti, and R. Couillet. Random matrix theory proves that deep learning representations of GAN-data behave as Gaussian mixtures. In H.D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proc. Mach. Learn. Res.*, pages 8573–8582. PMLR, 2020.
- [68] H. S. Seung, H. Sompolinsky, and N. Tishby. Statistical mechanics of learning from examples. *Phys. Rev. A*, 45:6056–6091, 1992.
- [69] L. Steinberger and H. Leeb. Conditional predictive inference for stable algorithms. *Ann. Stat.*, 51(1):290 – 311, 2018.
- [70] P. Sur and E.J. Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proc. Natl. Acad. Sci. U.S.A.*, 116(29):14516–14525, 2019.
- [71] C. Thrampoulidis, S. Oymak, and B. Hassibi. Regularized linear regression: A precise analysis of the estimation error. In Peter Grünwald, Elad Hazan, and Satyen Kale, editors, *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proc. Mach. Learn. Res.*, pages 1683–1709, Paris, France, 2015. PMLR.
- [72] C. Thrampoulidis, E. Abbasi, and B. Hassibi. Precise error analysis of regularized m -estimators in high dimensions. *IEEE Trans. Inf. Theory*, 64(8):5592–5628, 2018.
- [73] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- [74] T.L.H. Watkin, A. Rau, and M. Biehl. The statistical mechanics of learning a rule. *Rev. Mod. Phys.*, 65:499–556, 1993.
- [75] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning (still) requires rethinking generalization. *Commun. ACM*, 64(3):107–115, 2021.
- [76] Q. Zhao, P. Sur, and E.J. Candès. The asymptotic distribution of the MLE in high-dimensional logistic models: Arbitrary covariance. *Bernoulli*, 28(3):1835 – 1861, 2022.