

PAPER • OPEN ACCESS

## Error scaling laws for kernel classification under source and capacity conditions

To cite this article: Hugo Cui *et al* 2023 *Mach. Learn.: Sci. Technol.* 4 035033

View the [article online](#) for updates and enhancements.

You may also like

- [Encrypted machine learning of molecular quantum properties](#)

Jan Weinreich, Guido Falk von Rudorff and O Anatole von Lilienfeld

- [Machine learning the computational cost of quantum chemistry](#)

Stefan Heinen, Max Schwilk, Guido Falk von Rudorff et al.

- [A meshfree moving least squares-Tchebychev shape function approach for free vibration analysis of laminated composite arbitrary quadrilateral plates with hole](#)

Songhun Kwak, Kwanghun Kim, Kwangil An et al.



## PAPER

## OPEN ACCESS

RECEIVED  
4 November 2022REVISED  
4 July 2023ACCEPTED FOR PUBLICATION  
4 August 2023PUBLISHED  
29 August 2023

Original Content from  
this work may be used  
under the terms of the  
[Creative Commons  
Attribution 4.0 licence](#).

Any further distribution  
of this work must  
maintain attribution to  
the author(s) and the title  
of the work, journal  
citation and DOI.



# Error scaling laws for kernel classification under source and capacity conditions

Hugo Cui<sup>1,\*</sup> , Bruno Loureiro<sup>2</sup>, Florent Krzakala<sup>2</sup> and Lenka Zdeborová<sup>1</sup><sup>1</sup> Statistical Physics Of Computation lab, Institute of Physics, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland<sup>2</sup> Information Learning and Physics lab, Institute of Electrical Engineering, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

\* Author to whom any correspondence should be addressed.

E-mail: [hugo.cui@epfl.ch](mailto:hugo.cui@epfl.ch)**Keywords:** kernels, statistical physics, generalization error rates

## Abstract

In this manuscript we consider the problem of kernel classification. While worst-case bounds on the decay rate of the prediction error with the number of samples are known for some classifiers, they often fail to accurately describe the learning curves of real data sets. In this work, we consider the important class of data sets satisfying the standard source and capacity conditions, comprising a number of real data sets as we show numerically. Under the Gaussian design, we derive the decay rates for the misclassification (prediction) error as a function of the source and capacity coefficients. We do so for two standard kernel classification settings, namely margin-maximizing support vector machines and ridge classification, and contrast the two methods. We find that our rates tightly describe the learning curves for this class of data sets, and are also observed on real data. Our results can also be seen as an explicit prediction of the exponents of a scaling law for kernel classification that is accurate on some real datasets.

## 1. Introduction and related work

A recent line of work [1–4] has empirically evidenced that the test error of neural networks often obey scaling laws with the number of parameters of the model, training set size, or other model parameters. Because of their implications in terms of relating performance and model size, these findings have been the object of sustained theoretical attention. Authors of [5] relate the decay rate of the test loss with the number of parameters to the intrinsic dimension of the data. This idea is refined by [6] for the case of regression tasks, building on the observation that in a number of settings, the covariance of the learnt features exhibits a power-law spectrum, whose rate of decay controls the scaling of the error. This investigation is actually very closely related to another large body of works. In fact, the study of a power-law features spectrum (and of a target function whose components in the corresponding eigenbasis also decay as a power-law) has a long history in the kernel literature, dating back to the seminal works of [7, 8]. The corresponding rates governing the power-laws are respectively known as the *capacity* and *source* coefficients, and the scaling of the test error with the training set size can be entirely characterized in terms of these two numbers. While the study of kernel ridge regression [7–15] therefore offers a rich viewpoint on the question of neural scaling laws with the training set size, little is so far known for kernel *classification*. Since ascertaining the test error decay under source and capacity conditions would automatically translate into neural scaling laws in *classification* tasks—similarly to [6] for regression—this is a question of sizeable interest addressed in the present work.

### 1.1. Related works

#### 1.1.1. Neural scaling laws

A number of works [1–4] have provided empirical evidence of scaling laws in neural networks, with the number of parameters, training samples, compute, or other observables. These findings motivated theoretical investigations of the underlying mechanisms. Authors of [5] show how the scaling of the test loss with the number of parameters is related to the intrinsic dimension of the data. This dimension is further

tied in with the kernel spectrum by [6], a work that leverages the kernel ridge regression viewpoint to translate, in turn, the decay of the spectrum to test error rates. Authors of [16] similarly study a simple toy model where the power-law data is processed through a random features layer. Finally, [17] investigate a toy model of scalar integer data in the context of classification, and ascertain the corresponding scaling law. Relating in classification settings the rate of decay of the kernel spectrum to the test error, like [6] for regression, is still an open question.

### 1.1.2. Source and capacity conditions

The source and capacity conditions are standard regularity assumptions in the theoretical study of kernel methods, as they allow to subsume a large class of learning setups, see [7, 8, 12, 13, 15, 18].

### 1.1.3. Kernel ridge regression

The error rates for kernel ridge regression have been extensively and rigorously characterized in terms of the source/capacity coefficients in the seminal work of [7, 8], with a sizeable body of work being subsequently devoted thereto [9–14, 19]. In particular, in [15] it was shown that rates derived under worst-case assumptions [7–10, 20] are identical to the typical rates computed under the standard Gaussian design [21–23] assumption. Crucially, it was observed that many real data-sets satisfy the source/capacity conditions, and display learning rates in very good agreement to the theoretical values [15].

### 1.1.4. Worst-case analyses for SVM

The worst-case bounds for support vector machines (SVMs) classification—see e.g. [24, 25] for general introductions thereto—are known from the seminal works of [24, 26, 27]. However, it is not known how tightly the corresponding rates hold for a given realistic data distributions, not even for synthetic Gaussian data. We show that, contrary to the case of ridge regression, for classification the worst case bounds are not tight for Gaussian data. This effectively hinders the ability to predict and understand the error rates for relevant classes of data-sets, and in particular the class of data described by source/capacity conditions, which as mentioned above includes many real data-sets [15]. The key goal of this work is to fill this gap by leveraging the recent work on learning curves for the Gaussian covariate model [28] specified to data satisfying the capacity and source conditions.

## 1.2. Main contribution

In this work, we investigate the decay rate of the misclassification (generalization) error for noiseless kernel classification, under the Gaussian design and source/capacity regularity assumptions with capacity coefficient  $\alpha$  and source coefficient  $r$ . Building on the analytic framework of [28], we consider the two most widely used classifiers: margin-maximizing SVMs and ridge classifiers. We derive in section 3 the error rate (describing the decay of the prediction error with the number of samples) for margin-maximizing SVM:

$$\epsilon_g^{\text{SVM}} \sim n^{-\frac{\alpha \min(r, \frac{1}{2})}{1 + \alpha \min(r, \frac{1}{2})}}.$$

As a consequence, we conclude that the worst-case rates [24, 26, 27] are indeed loose and fail to describe this class of data. This fact alone is not at all surprising. However, it becomes remarkable in the light of the fact that for ridge regression the worst case bounds and the typical case rates do agree [15].

We contrast the SVM rate with the rate for optimally regularized ridge classification, which we establish in section 4 to be

$$\epsilon_g^{\text{ridge}} \sim n^{-\frac{\alpha \min(r, 1)}{1 + 2\alpha \min(r, 1)}}.$$

We argue in the light of these findings that the SVM always displays faster rates than the ridge classifier for the classification task considered.

Finally, we observe that some real data-sets fall in the same universality class as the considered setting, in the sense that, as illustrated in section 5, their error rates are in very good agreement with the ones above. This work is thus a key step for theoretically predicting the error rates of kernel classification for a broad range of real data-sets.

## 2. Setting

### 2.1. Kernel classification

Consider a data-set  $\mathcal{D} = \{(x^\mu, y^\mu)\}_{\mu=1}^n$  with  $n$  independent samples from a probability measure  $\nu$  on  $\mathcal{X} \times \{-1, +1\}$ , with  $\mathcal{X} \subset \mathbb{R}^d$ . We will assume that the labels can be expressed as

$$y^\mu = \text{sign}(f^*(x^\mu)) \tag{1}$$

for some non-stochastic target function  $f^* : \mathcal{X} \rightarrow \mathbb{R}$ . Note that the *noiseless* setting considered here is out of the validity domain of many worst case analyses, whose bounds become void without noise [27], whereas a number of real learning settings are well described by a noiseless setup, see section 5. Learning to classify  $\mathcal{D}$  in the direct space  $\mathcal{X}$  for a *linear*  $f^*$  has been the object of extensive studies. In the present work, we focus on the case where  $f^*$  more generically belongs to the space of square-integrable functions  $L^2(\mathcal{X})$ . To classify  $\mathcal{D}$ , a natural method is then to perform *kernel classification* in a  $p$ - dimensional reproducing kernel Hilbert space (RKHS)  $\mathcal{H}$  associated to a kernel  $K$ , by minimizing the regularized empirical risk:

$$\hat{\mathcal{R}}_n(f) = \frac{1}{n} \sum_{\mu=1}^n \ell(f(x^\mu), y^\mu) + \lambda \|f\|_{\mathcal{H}}^2. \tag{2}$$

The function  $\ell(\cdot)$  is a loss function and  $\lambda$  is the strength of the  $\ell_2$  regularization term. In this paper we shall more specifically consider the losses  $\ell(z, y) = \max(0, 1 - yz)$  (hinge classification) and  $\ell(z, y) = (y - z)^2$  (ridge classification), and the case of an infinite dimensional RKHS ( $p = \infty$ ). The risk (2) admits a dual rewriting in terms of a standard parametric risk. To see this, diagonalize  $K$  in an orthogonal basis of *kernel features*  $\{\psi_k(\cdot)\}_{k=1}^p$  of  $L^2(\mathcal{X})$ , with corresponding eigenvalues  $\{\omega_k\}_{k=1}^p$ :

$$\int_{\mathcal{X}} \nu(dx') K(x, x') \psi_k(x') = \omega_k \psi_k(x). \tag{3}$$

It is convenient to normalize the eigenfunctions to

$$\int_{\mathcal{X}} \nu(dx) \psi_k(x)^2 = \omega_k, \tag{4}$$

so that the kernel  $K$  can be rewritten in simple scalar product form  $K(x, x') = \psi(x)^\top \psi(x')$ , where we named  $\psi(x)$  the  $p$ -dimensional vector with components  $\{\psi_k(x)\}_{k=1}^p$ . Furthermore, note that the covariance  $\Sigma$  of the data in feature space with this choice of feature map is simply diagonal

$$\Sigma = \mathbb{E}_{x \sim \nu} (\psi(x) \psi(x)^\top) = \text{diag}(\omega_1, \dots, \omega_p). \tag{5}$$

Any function  $f \in \mathcal{H}$  can then be expressed as  $f(\cdot) = w^\top \psi(\cdot)$  for a vector  $w$  with square summable components. Using this parametrization, the risk (2) can be rewritten as

$$\hat{\mathcal{R}}_n(w) = \frac{1}{n} \sum_{\mu=1}^n \ell(w^\top \psi(x^\mu), y^\mu) + \lambda w^\top w. \tag{6}$$

Throughout this manuscript we will refer to the components of the target function in the features basis as the *teacher*  $\theta^*$ , so that

$$f^*(\cdot) = \theta^{*\top} \psi(\cdot).$$

Note that any  $f^* \in L^2(\mathcal{X})$  can be formally written in this form with a certain  $\theta^*$  (allowing for non square-summable components if  $f^* \in L^2(\mathcal{X}) \setminus \mathcal{H}$ ). Similarly, the minimizer  $\hat{w}$  of the parametric risk (6) is related to the argmin  $\hat{f}$  of (2) by  $\hat{f}(\cdot) = \hat{w}^\top \psi(\cdot)$ , and will be referred to as the *estimator* in the following. We make two further assumptions : first, we work under the *Gaussian design*, and assume the features  $\psi(x)$  to follow a Gaussian distribution with covariance  $\Sigma$ , i.e.  $\psi(x) \sim \mathcal{N}(0, \Sigma)$ . Note that this assumption might appear constraining, as the distribution of the data in feature space strongly depends on its distribution in the original space, and the feature map associated to the kernel. In fact, for a large class of data distributions and standard kernels, the Gaussian design assumption does not hold. However, rates derived under Gaussian design can hold more broadly. For instance, the rates established by [15] under Gaussian design were later proven by [29] under weaker conditions on the features. We will moreover discuss in section 5 several settings in which our theoretical rates are in good agreement with rates observed for real data.

Second, following [15], we assume that the *regularization strength*  $\lambda$  *decays as a power-law* of the number of samples  $n$  with an exponent  $\ell$ :  $\lambda = n^{-\ell}$ . Note that this form of regularization is natural, since the need for regularizing is lesser for larger training sets. Furthermore, this allows to investigate the classical question of the *asymptotically optimal regularization* [7, 8, 15], i.e. the decay  $\ell$  of the regularization yielding fastest decrease of the prediction error.

### 2.2. Source and capacity conditions

Under the above assumptions of Gaussian design with features covariance  $\Sigma$  and existence of a teacher  $\theta^*$  that generates the labels using equation (1) we can now study the error rates. In statistical learning theory one often uses the *source and capacity conditions*, which assume the existence of two parameters  $\alpha > 1, r \geq 0$  (hereafter referred to as the *capacity coefficient* and the *source coefficient* respectively) so that

$$\text{tr}\Sigma^{\frac{1}{\alpha}} < \infty, \quad \theta^{*\top} \Sigma^{1-2r} \theta^* < \infty. \tag{7}$$

As in [13, 15, 21, 30, 31], we will consider the particular case where both the spectrum of  $\Sigma$  and the teacher components  $\theta_k^*$  have exactly a power-law form satisfying the limiting source/capacity conditions (7):

$$\omega_k = k^{-\alpha}, \quad \theta_k^* = k^{-\frac{1+\alpha(2r-1)}{2}}. \tag{8}$$

The power-law forms (8) have been empirically found in [15] in the context of kernel regression to be a reasonable approximation for a number of real data-sets including MNIST [32] and Fashion MNIST [33] and a number of standard kernels such as polynomial kernels and radial basis functions. Similar observations were also made in the present work and are discussed in appendix F and section 5.

The capacity parameter  $\alpha$  and source parameter  $r$  capture the complexity of the data-set in feature space—i.e. after the data is transformed through the kernel feature map into  $\{\psi(x^\mu), y^\mu\}_{\mu=1}^n$ . A large  $\alpha$ , for example, signals that the spectrum of the data covariance  $\Sigma$  displays a fast decay, implying that the data effectively lies along a small number of directions, and has a low effective dimension. Conversely, a small capacity  $\alpha$  means that the data is effectively large dimensional, and therefore *a priori* harder to learn. Similarly, a large  $r$  signals a good alignment of the teacher  $\theta^*$  with the main directions of the data, and *a priori* an easier learning task. In terms of the target function  $f^*$ , larger  $r$  correspond to smoother  $f^*$ . Note that  $r > 1/2$  implies that  $f^* \in \mathcal{H}$ , while  $r \leq 1/2$  implies  $f^* \in L^2(\mathcal{X}) \setminus \mathcal{H}$ .

Finally, note that while [18] suggested an alternative definition for the source and capacity coefficients in the case of non-square loss functions, their redefinition is not directly applicable for the hinge loss.

### 2.3. Misclassification error

The performance of learning the data-set  $\mathcal{D}$  using kernel classification (6) is quantified by the misclassification (generalization) error

$$\epsilon_g = \frac{1}{2} - \frac{1}{2} \mathbb{E}_{\mathcal{D}} \mathbb{E}_{x, y \sim \nu} (y \text{sign}(\hat{w}^\top \psi(x))), \tag{9}$$

where  $\hat{w}$  is the minimizer of the risk (6). The error (9) corresponds to the probability for the predicted label  $\text{sign}(\hat{w}^\top \psi(x))$  of a test sample  $x$  to be incorrect. The rate at which the error (9) decays with the number of samples  $n$  in  $\mathcal{D}$  depends on the complexity of the data-set, as captured by the source and capacity coefficients  $\alpha, r$  equation (8). To compute this rate, we build upon the work of [28] who, following a long body of work in the statistical physics literature [31, 34–38], provided and proved a mathematically rigorous closed form asymptotic characterization of the misclassification error as

$$\epsilon_g = \frac{1}{\pi} \arccos(\sqrt{\eta}), \quad \eta = \frac{m^2}{\rho q}, \tag{10}$$

where  $\rho$  is the squared  $L^2(\mathcal{X})$  norm of the target function  $f^*$ , i.e.  $\rho = \int_{\mathcal{X}} \nu(dx) f^*(x)^2 = \theta^{*\top} \Sigma \theta^*$ , and  $m, q$  are the solution of a set of self-consistent equations, which are later detailed and analyzed in section 3 for margin-maximizing SVMs and section 4 for ridge classifiers. The order parameters  $m, q$  are known as the *magnetization* and the *self-overlap* in statistical physics and respectively correspond to the target/estimator and estimator/estimator  $L^2(\mathcal{X})$  correlations:

$$m = \mathbb{E}_{\mathcal{D}} \int_{\mathcal{X}} \nu(dx) f^*(x) \hat{f}(x) = \mathbb{E}_{\mathcal{D}} (\hat{w}^\top \Sigma \theta^*), \quad q = \mathbb{E}_{\mathcal{D}} \int_{\mathcal{X}} \nu(dx) \hat{f}(x)^2 = \mathbb{E}_{\mathcal{D}} (\hat{w}^\top \Sigma \hat{w}). \tag{11}$$

It follows from these interpretations that  $\eta$  has to be thought of as the *cosine-similarity* between the teacher  $\theta^*$  and the estimator  $\hat{w}$ , with perfect alignment ( $\eta = 1$ ) resulting in minimal error  $\epsilon_g = 0$  from (9).

Note that while this characterization has formally been proven in [28] in the asymptotic proportional  $n, p \rightarrow \infty, n/p = \mathcal{O}(1)$  limit, we are presently using it in the  $n \ll p = \infty$  limit, thereby effectively working at  $n/p = 0^+$ . The non-asymptotic rate guarantees of [28] are nevertheless encouraging in this respect, although a finer control of the limit would be warranted to put the present analysis on fully rigorous grounds. Further, [15] also build on [28] in the  $n/p = 0^+$  limit, and display solid numerics-backed results, later rigorously

proven by [29]. We thus conjecture that this limit can be taken as well safely in our case. Finally, we mention that a recent line of works [39–42] has explored the connections between kernel regression and Bayesian learning for networks in the  $n/p = \mathcal{O}(1)$  limit, where  $p$  is in this case the width of the network. While the high-dimensional limit is indeed related to the one originally discussed in [28], which we relax here to  $n/p = 0^+$ , the main object of [39–41] was not to study kernel regression per se, but to show how observables in Bayesian regression could be expressed in terms of well-chosen kernels. In the present work, we focus on analyzing kernel classification in the  $n/p = 0^+$  regime.

### 3. Max-margin classification

#### 3.1. Self-consistent equations

In this section we study regression using SVMs. The risk (6) then reads for the hinge loss

$$\hat{\mathcal{R}}_n(w) = \frac{1}{n} \sum_{\mu=1}^n \max(0, 1 - y^\mu w^\top \psi(x^\mu)) + \lambda w^\top w. \tag{12}$$

In the following, we shall focus more specifically on the max-margin limit with  $\lambda = 0^+$ . We show in appendix B that zero regularization is indeed asymptotically optimal for the data following equation (8) when the target function is characterized by a source  $r \leq 1/2$ , i.e.  $f^* \in L^2(\mathcal{X}) \setminus \mathcal{H}$ . We heuristically expect margin maximization to be *a fortiori* optimal also for easier and smoother teachers  $f^* \in \mathcal{H}$ . For the risk (12) at  $\lambda = 0^+$ , the self-consistent equations defining  $m, q$  in (11) read (see appendix A)

$$\left\{ \begin{array}{l} \rho = \sum_{k=1}^{\infty} \theta_k^{*2} \omega_k, \\ m = \hat{r}_1 \frac{n}{z} \sum_{k=1}^{\infty} \frac{\omega_k^2 \theta_k^{*2}}{1 + \frac{n}{z} \omega_k}, \\ q = \hat{r}_1^2 \frac{n^2}{z^2} \sum_{k=1}^{\infty} \frac{\theta_k^{*2} \omega_k^3}{(1 + \frac{n}{z} \omega_k)^2} \\ \quad + \hat{r}_2 \frac{n}{z^2} \sum_{k=1}^{\infty} \frac{\omega_k^2}{(1 + \frac{n}{z} \omega_k)^2} \end{array} \right. , \quad \left\{ \begin{array}{l} \hat{r}_1 = \frac{1}{2\pi\sqrt{\rho}} \frac{\left(\sqrt{2\pi} \left(1 + \operatorname{erf}\left(\frac{1}{\sqrt{2q(1-\eta)}}\right)\right) + 2e^{-\frac{1}{2q(1-\eta)}} \sqrt{q(1-\eta)}\right)}{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx e^{-\frac{1}{2}x^2} \left[1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right)\right]}, \\ \hat{r}_2 = \frac{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx e^{-\frac{1}{2}x^2} \left[1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right)\right] (1 - \sqrt{q}x)^2}{\left(\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx e^{-\frac{1}{2}x^2} \left[1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right)\right]\right)^2}, \\ z = \frac{\sum_{k=1}^{\infty} \frac{\omega_k}{\frac{n}{z} + \omega_k}}{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx e^{-\frac{1}{2}x^2} \left[1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right)\right]}. \end{array} \right. \tag{13}$$

Here  $\hat{r}_1$  should be thought of as the ratio between the norms of the estimator  $\hat{w}$  and the teacher  $\theta^*$ , while  $z$  can be loosely interpreted as an effective regularization. A detailed derivation of these equations can be found in appendix A.

#### 3.2. Decay rates for max-margin

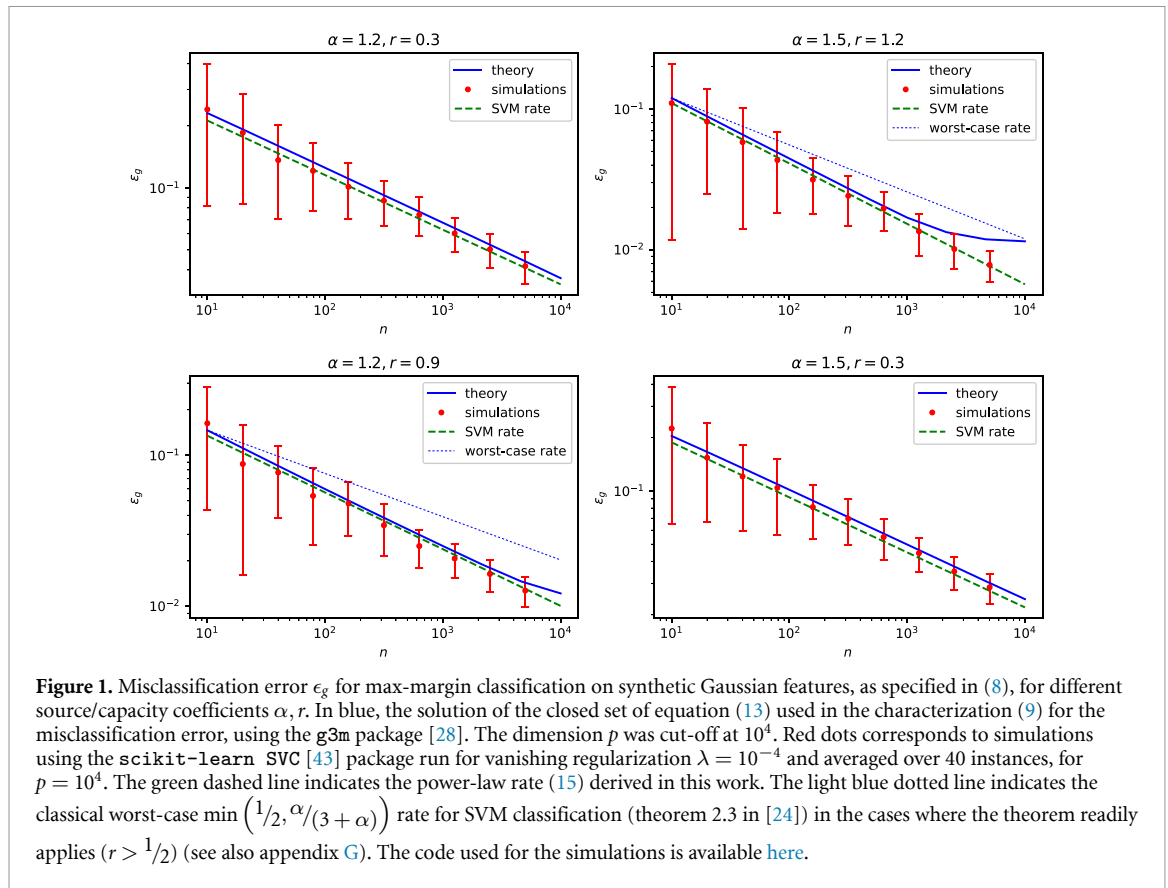
From the investigation of the equation (13), as detailed in appendix A, the following scalings are found to hold between the order parameters:

$$m \sim \sqrt{q} \sim \hat{r}_1 \sim n \left(\frac{z}{n}\right)^{\frac{1}{\alpha}} \sim n^{\frac{\alpha \min(r, \frac{1}{2})}{1 + \alpha \min(r, \frac{1}{2})}}. \tag{14}$$

Note that the mutual scaling between  $m, q$  also follows intuitively from the interpretation of these order parameters—as the overlap of  $\hat{w}$  with the ground truth and itself respectively—see the discussion around equations (11) and (13). Since the width of the margin is generically expected to shrink with the number of samples (as more training data are likely to be sampled close to the separating hyperplane), the increase of the norm of  $\hat{w}$  (as captured by  $q, \hat{r}_1$ ) with  $n$  is also intuitive. Finally, an analysis of the subleading corrections to  $m$  and  $q$ , detailed in appendix A, leads to

$$\epsilon_g \sim n^{-\frac{\alpha \min(r, \frac{1}{2})}{1 + \alpha \min(r, \frac{1}{2})}}. \tag{15}$$

The error rate (15) stands in very good agreement with numerical simulations on artificial Gaussian features generated using the model specification (8), see figure 1. Two observations can further be made on the decay rate (15). First, the rate is as expected an increasing function of  $\alpha$  (low-dimensionality of the



features) and  $r$  (smoothness of the target  $f^*$ ). Second, for a source  $r > \frac{1}{2}$  (corresponding to a target  $f^* \in \mathcal{H}$ ), the rate saturates, suggesting that all functions in  $\mathcal{H}$  are all equally easy to classify, while for rougher target  $f^* \in L^2(\mathcal{X}) \setminus \mathcal{H}$  the specific roughness of the target function, as captured by its source coefficient  $r$ , matters and conditions the rate of decay of the error.

Finally, we briefly discuss for completeness in appendix E the more general case where the label distribution (1) includes data noise, and show that the rates display a crossover from the noiseless value (15) to a noisy value, much like what was reported for kernel ridge regression [15].

### 3.3. Comparison to classical rates

To the best of the authors' knowledge, there currently exists little work addressing the error rates for datasets satisfying source and capacity conditions (7). The closest result is the worst-case bound of [24] for SVM classification, which can be adapted to the present setting provided  $f^* \in \mathcal{H}$  ( $r > \frac{1}{2}$ ). The derivation is detailed in appendix G and results in an upper bound of  $\min\left(\frac{1}{2}, \frac{\alpha}{3+\alpha}\right)$  for the error rate for max-margin classification, which is always slower than (15). This rate [24] is plotted for comparison in figure 1 against numerical simulations and is visibly off, failing to capture the learning curves. It is to be expected that the worst case rates will be loose when compared to rate that assume a specific data distribution. What makes our result interesting is the comparison with the more commonly studied ridge regression where, as discussed already in the introduction, the worst case rates actually match those derived for Gaussian data, see [15].

Importantly, the rates from [24] only hold for capacity  $r > \frac{1}{2}$ , while real datasets are typically characterized by sources  $r < \frac{1}{2}$  (see for instance figure 4). The present work therefore fills an important gap in the literature in providing rates (15) which accurately capture the learning curves of datasets satisfying source and capacity conditions. Further discussions of [24], along with [26, 27, 44] for completeness, are provided in appendix G. Also note that while [45] report  $\frac{\alpha}{1+\alpha}$  rates under Gaussianity assumptions, they rely on very stringent assumptions which are too strong and unfulfilled in our setting.

## 4. Ridge classification

### 4.1. Self-consistent equations

Another standard classification method is the ridge classifier, which corresponds to minimizing

$$\hat{\mathcal{R}}_n(w) = \frac{1}{n} \sum_{\mu=1}^n (y^\mu - w^\top \psi(x^\mu))^2 + \lambda w^\top w. \tag{16}$$

As previously discussed in section 2, we consider a decaying regularization  $\lambda = n^{-\ell}$ . The self-consistent equations characterizing the quantities  $(q, m)$ , read for the ridge risk (16)

$$\begin{cases} \rho = \sum_{k=1}^{\infty} \theta_k^{*2} \omega_k, \\ m = \sqrt{\frac{2}{\pi\rho}} \frac{n}{z} \sum_{k=1}^{\infty} \frac{\omega_k^2 \theta_k^{*2}}{1 + \frac{n}{z} \omega_k}, \\ q = \frac{n^2}{z^2} \sum_{k=1}^{\infty} \frac{\frac{2}{\pi\rho} \theta_k^{*2} \omega_k^3 + \frac{1+q-2m}{n} \sqrt{\frac{2}{\pi\rho}} \omega_k^2}{(1 + \frac{n}{z} \omega_k)^2}, \\ z = n\lambda + \frac{z}{n} \sum_{k=1}^{\infty} \frac{\lambda_k}{\lambda_k + \frac{z}{n}}. \end{cases} \tag{17}$$

Further details on the derivation of (17) from [28] are provided in appendix C. Like (13), equation (17) have been formally proven in the proportional  $n, p \rightarrow \infty, n/p = \mathcal{O}(1)$  limit in [28], but are expected to hold also in the present  $n \ll p = \infty$  setting [15, 29]. Note that comparing to (13), equation (17) correspond to a constant student/teacher norm ratio  $\hat{r}_1 = 2/(\pi\rho)$  and to a simple  $\hat{r}_2 = 1 + q - 2m\sqrt{2/(\pi\rho)}$ .  $\hat{r}_2$  moreover admits a very intuitive interpretation as the prediction mean squared error (MSE) between the true label  $y = \text{sign}(\theta^{*\top} \psi(x))$  and the pre-activation linear predictor  $\hat{w}^\top \psi(x)$ , i.e.  $\hat{r}_2 = \mathbb{E}_{\psi(x)} \left( \text{sign} \left( \theta^{*\top} \psi(x) \right) - \hat{w}^\top \psi(x) \right)^2$ .

### 4.2. Decay rates for ridge classification

Similarly to [15, 31], an analysis of the equation (17) (see appendix C) reveals that, depending on how the rate of decay  $\ell$  of the regularization compares to the capacity  $\alpha$ , two regimes (called *effectively regularized* and *effectively un-regularized* in [15] in the context of ridge regression) can be found:

#### 4.2.1. Effectively regularized regime

$\ell \leq \alpha$ . In this regime, an analysis of the corrections to the self-overlap  $q$  and magnetization  $m$ , presented in appendix C, shows that the misclassification error scales like

$$\epsilon_g \sim n^{-\frac{1}{2} \min(2\ell \min(r, 1), \frac{\alpha - \ell}{\alpha})}. \tag{18}$$

The rate (18) compares very well to numerical simulations, see figure 2. Note that the saturation for ridge happens for  $r = 1$ , rather than  $r = 1/2$  as for max-margin classification (see discussion in section 3): very smooth targets  $f^*$  characterized by a source  $r \geq 1$  are all equally easily classified by ridge. For rougher teachers  $f^*$  characterized by  $r \leq 1$  however, the rate of decay of the error (18) depends on the specific roughness of the target, even if, in contrast to max-margin, the latter belongs to  $\mathcal{H}$  ( $r > 1/2$ ). Two important observations should further be made on the rates (18):

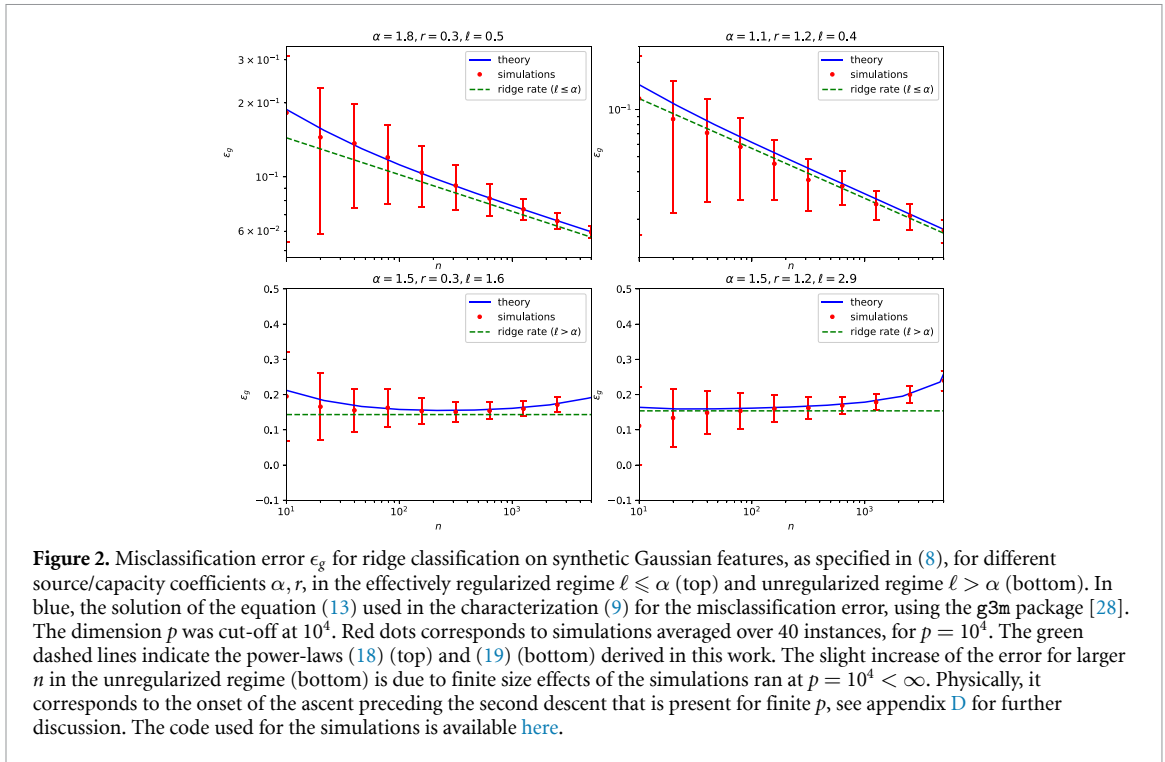
- If the regularization remains small (fast decay  $\alpha > \ell > \alpha/(1 + 2\alpha \min(r, 1))$ ), the decay (18) is determined only by the data capacity  $\alpha$ , while the source  $r$  plays no role. As a matter of fact, with insufficient regularization, the limiting factor to the learning is the tendency to overfit, which depends on the effective dimension of the data as captured by the capacity  $\alpha$ .
- For larger regularizations (slow decays  $\ell < \alpha/(1 + 2\alpha \min(r, 1))$ ), the limiting factor becomes the complexity of the teacher  $\theta^*$ , as captured by the source  $r$ .

#### 4.2.2. Effectively un-regularized regime

$\ell > \alpha$ . As derived in appendix C, the error plateaus and stays of order 1:

$$\epsilon_g = \mathcal{O}(1). \tag{19}$$





This plateau is further elaborated upon in appendix D, and is visible in numerical experiments, see figure 2. It corresponds to the first plateau in a double descent curve, with the second descent never happening since  $p = \infty$ . Intuitively, this phenomenon is attributable to the ridge classifier overfitting the labels using the small-variance directions of the data (8).

Interestingly, all the rates (19) and (18) correspond exactly (up to a factor  $1/2$ ) to those reported in [15] for the MSE of ridge regression, where they are respectively called the *red*, *blue* and *orange* exponents. Notably, the plateau (19) at low regularizations and the  $(\alpha - \ell)/\alpha$  exponent in (18) only appeared in [15] for noisy cases in which the labels are corrupted by an additive noise. The fact that they hold in the present *noiseless* study very temptingly suggests that model mis-specification (trying to interpolate binary labels using a linear model) effectively plays the role of a large noise.

### 4.3. Optimal rates

#### 4.3.1. Optimally regularized ridge classification

In practice, the strength of the regularization  $\lambda$  is a tunable parameter. A natural question to ask is then the one of the *asymptotically optimal* regularization, that is the regularization decay rate  $\ell^*$  leading to fastest decay rates for the misclassification error. From the expressions of (18) (which hold provided  $\ell < \alpha$ ) and (19) (which holds provided  $\ell > \alpha$ ), the value of  $\ell$  maximizing the error rate is found to be

$$\ell^* = \frac{\alpha}{1 + 2\alpha \min(r, 1)}, \tag{20}$$

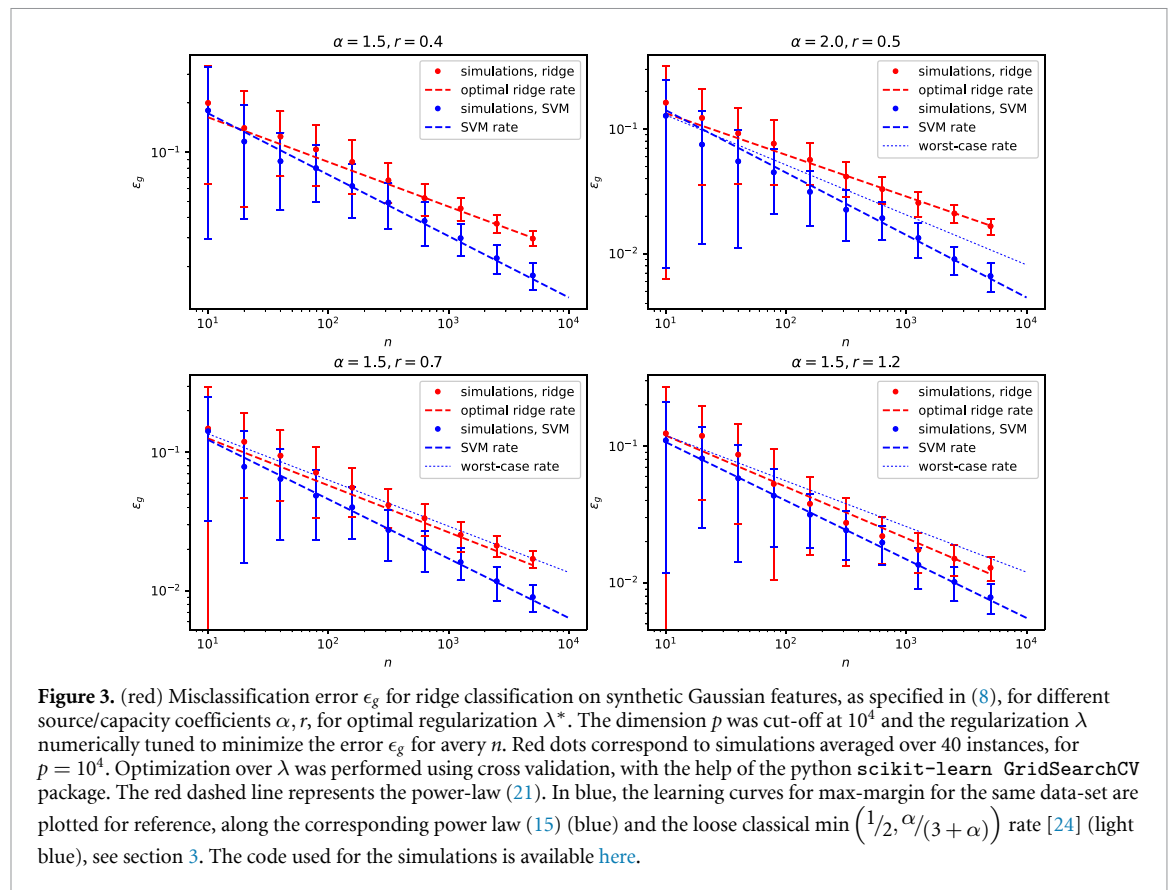
and the corresponding error rate for  $\epsilon_g^* = \epsilon_g(\lambda^* = n^{-\ell^*})$  is

$$\epsilon_g^* \sim n^{-\frac{\alpha \min(r, 1)}{1 + 2\alpha \min(r, 1)}}, \tag{21}$$

see the red dashed lines in figure 3. Coincidentally, the optimal rate (21) is up to a factor  $1/2$  identical to the classical optimal rate known for the rather distinct problem of the MSE of kernel ridge regression on noisy data [7, 8]. Like the max-margin exponent (15), the optimal error rate for ridge (21) is an increasing function of both the capacity  $\alpha$  and the source  $r$ , i.e. of the easiness of the learning task. Note that in contrast to max-margin classification which is insensitive to the specifics of the target function  $f^*$ , provided it is in  $\mathcal{H}$ , ridge is sensitive to the source (smoothness)  $r$  of  $f^*$  up to  $r = 1$ .

#### 4.3.2. Comparison to max-margin

A comparison of the max-margin rate  $a_{\text{SVM}} = \alpha \min(r, \frac{1}{2}) / (1 + \alpha \min(r, \frac{1}{2}))$  (15) and the optimal ridge exponent  $a_r = \frac{\alpha \min(r, 1)}{1 + 2\alpha \min(r, 1)}$  (21) reveals that for any  $\alpha > 1, r \geq 0, a_{\text{SVM}} - a_r > 0$ . In other

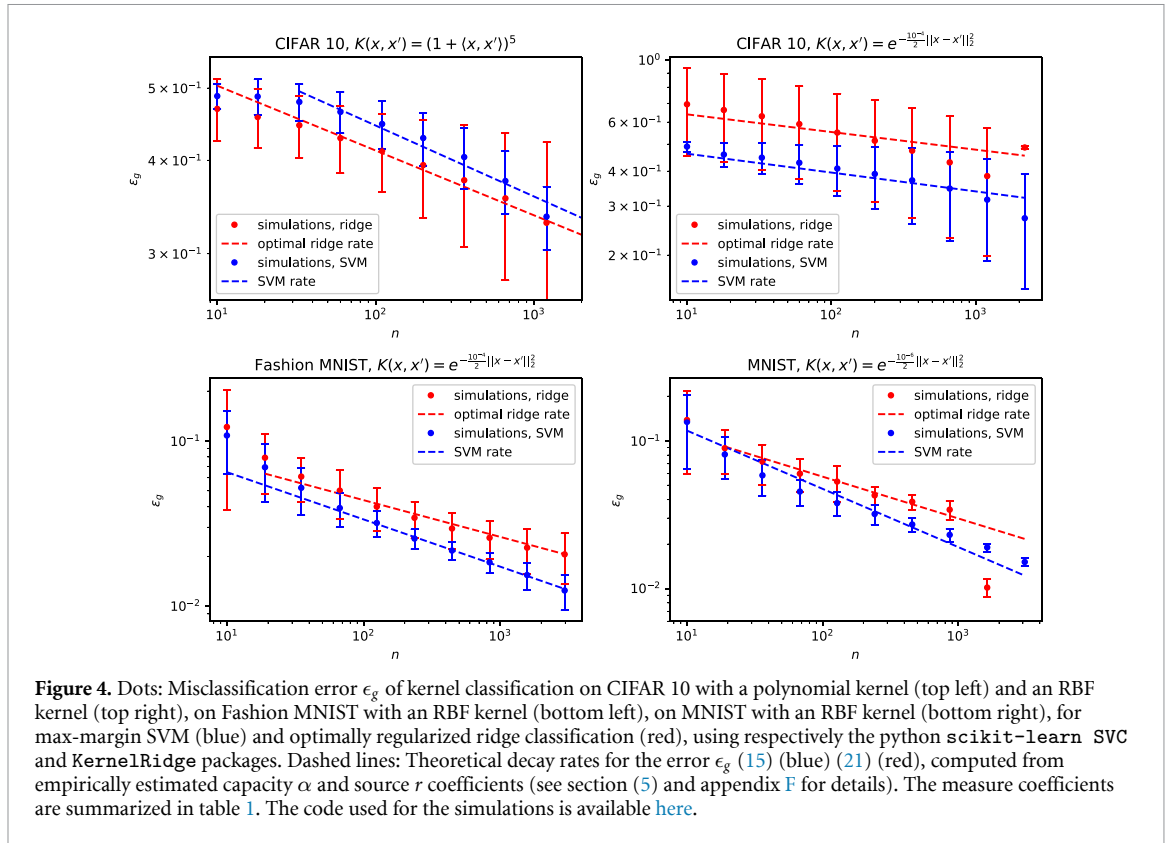


words, the margin-maximizing SVM displays faster rates than the ridge classifier for the class of data studied (8), see figure 3.

We finally briefly comment on support vector proliferation. [46–48] showed that in some settings almost every training sample in  $\mathcal{D}$  becomes a support vector for the SVM. In such settings, the estimators  $\hat{w}$  (and hence the error  $\epsilon_g$ ) consequently coincide for the ridge classifier and the margin-maximizing SVM. In the present setting however, the result  $a_{\text{SVM}} - a_r > 0$  establishes that for features with a power-law decaying spectrum (8), there is no such support vector proliferation. Note that this result does not follow immediately from theorem 3 in [48]. In fact, the spiked covariance (8), with only a small number of important (large variance) directions and a tail of unimportant (low-variance) directions does effectively not offer enough overparametrization [20, 47] for support vector proliferation, and the support consists only of the subset of the training set with weakest alignment with the spike.

### 5. Remarks for real data-sets

The source and capacity condition (8) provide a simple framework to study a large class of structured data-sets. While idealized, we observe, similarly to [15], that many real data-sets seem to fall under this category of data-sets, and hence display learning curves which are to a good degree described by the rates (15) for SVM and (21) for ridge classification. We present here three examples of such data-sets : a data-set of  $10^4$  randomly sampled CIFAR 10 [49] images of animals (labelled +1) and means of transport (labelled -1), a data-set of 14000 FashionMNIST [33] images of t-shirts (labelled +1) and coats (labelled -1), and a data-set of 14702 MNIST [32] images of 8s (labelled +1) and 1s (labelled -1). On the one hand, the learning curves for max-margin classification and optimally regularized ridge classification were obtained using the python `scikit-learn SVC, KernelRidge` packages. On the other hand, the spectrum  $\{\omega_k\}_k$  of the data covariance  $\Sigma$  in feature space was computed, and a teacher  $\theta^*$  providing perfect classification of the data-set was fitted using margin-maximizing SVM. Then, the capacity and source coefficients  $\alpha, r$  (8) were estimated for the data-set by fitting  $\{\omega_k\}_k$  and  $\{\theta_k^*\}_k$  by power laws, and the theoretical rates (15) and (21) computed therefrom. More details on this method, adapted from [30, 31], are provided in appendix F. The results of the simulations are presented in figure 4 and compared to the theoretical rates (15) and (21) computed from the empirically evaluated source and capacity coefficients for a radial basis function (RBF) kernel and a



**Table 1.** Values of the source and capacity coefficients (7) as estimated from the data sets, and the corresponding theoretical error rates for SVM (15) and ridge (21). The details on the estimation procedure can be found in appendix F.

Dataset	Kernel	$\alpha$	$r$	$a_{SVM}$	$a_r$
CIFAR 10	polynomial	1.51	0.07	0.095	0.086
CIFAR 10	RBF	1.005	0.07	0.067	0.063
Fashion MNIST	RBF	1.72	0.23	0.28	0.22
MNIST	RBF	1.65	0.39	0.39	0.28

polynomial kernel of degree 5, with overall very good agreement. We do not compare here with the worst case bounds because the observed values of  $r < 1/2$  in which case we remind the known results do not apply.

## 6. Conclusion

We compute the generalization error rates as a function of the source and capacity coefficients for two standard kernel classification methods, margin-maximizing SVM and ridge classification, and show that SVM classification consistently displays faster rates. Our results establish that known worst-case upper bound rates for SVM classification fail to tightly capture the rates of the class of data described by source/capacity conditions. We illustrate empirically that a number of real data-sets fall under this class, and display error rates which are to a very good degree described by the ones derived in this work.

## Data availability statement

No new data were created or analysed in this study.

## Acknowledgments

We thank Loucas Pillaud-Vivien, Alessandro Rudi, Vittorio Erba and Mauro Pastore for useful discussion. We acknowledge funding from the ERC under the European Union’s Horizon 2020 Research and Innovation Programme Grant Agreement 714608-SMiLe.

## Appendix A. Rates for margin maximizing SVM

In this appendix, we provide some analytical discussion of the equation (13) motivating the scaling (15) and (14). We remind the risk for the hinge loss (12)

$$\hat{\mathcal{R}}_n(w) = \frac{1}{n} \sum_{\mu=1}^n \max(0, 1 - y^\mu w^\top \psi(x)) + \lambda w^\top w, \tag{A.1}$$

with  $\lambda = 0^+$  for max-margin. The predictor is then  $\hat{y} = \text{sign}(\hat{w} \cdot \psi(x))$ .

### A.1. Mapping from [28]

The starting point is the closed-form asymptotic characterization of the misclassification error of [28]. We begin by reviewing the main results of [28], and detail how their setting can be mapped to ours. Consider hinge regression on  $n$  independent  $p$ - dimensional Gaussian samples  $\mathcal{D} = \{x^\mu, y^\mu\}_{\mu=1}^n$ , by minimizing the empirical risk

$$\hat{\mathcal{R}}_n(w) = \sum_{\mu=1}^n \max\left(0, 1 - y^\mu \frac{w^\top x^\mu}{\sqrt{p}}\right) + \frac{\lambda}{2} w^\top w \tag{A.2}$$

for some constant regularization strength  $\lambda \geq 0$ . Suppose that the labels are generated from a teacher/target/oracle  $\theta^* \in \mathbb{R}^p$  as  $y^\mu = \text{sign}(\theta^{*\top} x^\mu)$ . Then provided the assumptions

- **A.1**  $n \gg 1, p \gg 1, \frac{n}{p} = \mathcal{O}(1)$ ,
- **A.2**  $0 < \frac{1}{p} \|\theta^*\|_2^2 < \infty$ ,

are satisfied, there exist constants  $C, c, c' > 0$  so that for all  $0 < \epsilon < c'$ ,

$$\mathbb{P}\left(\left|\epsilon_g - \frac{1}{\pi} \arccos\left(\frac{m^*}{\sqrt{\rho q^*}}\right)\right| > \epsilon\right) < \frac{C}{\epsilon} e^{-c\epsilon^2}, \tag{A.3}$$

where  $\rho = \theta^{*\top} \Sigma \theta^* / p$  and  $m^*, q^*$  are the solutions of the *fixed point equations*

$$\begin{cases} m = \frac{\hat{m}}{p} \text{tr} \frac{\Sigma \theta^* \theta^{*\top} \Sigma}{\lambda + \hat{V} \Sigma} \\ q = \frac{1}{p} \text{tr} \frac{\hat{m}^2 \Sigma \theta^* \theta^{*\top} \Sigma^2 + \hat{q} \Sigma^2}{(\lambda + \hat{V} \Sigma)^2} \\ V = \frac{1}{p} \text{tr} \frac{\Sigma}{\lambda + \hat{V} \Sigma} \end{cases}, \quad \begin{cases} \hat{m} = \frac{\frac{n}{p}}{V \sqrt{\rho}} \frac{1}{2\pi} \left\{ \sqrt{2\pi} \left[ \text{erf}\left(\frac{1}{\sqrt{2q(1-\eta)}}\right) - \text{erf}\left(\frac{1-V}{\sqrt{2q(1-\eta)}}\right) \right] \right. \\ \left. + 2\sqrt{q(1-\eta)} \left( e^{-\frac{1}{2q(1-\eta)}} - e^{-\frac{(1-V)^2}{2q(1-\eta)}} \right) + \sqrt{2\pi} V \left( 1 + \text{erf}\left(\frac{1-V}{\sqrt{2q(1-\eta)}}\right) \right) \right\} \\ \hat{q} = \frac{n}{p} \left[ \int_{-\infty}^{\frac{1-V}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}} x\right) \right] + \frac{1}{V^2} \int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}} x\right) \right] (1 - \sqrt{q}x)^2 \right] \\ \hat{V} = \frac{\frac{n}{p}}{V} \int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}} x\right) \right] \end{cases},$$

with  $\eta = m^2/(\rho q)$ , and  $\Sigma$  the covariance of the samples  $x$ . The limit  $\lambda = 0^+$  can be taken using the rescaling [50]

$$\hat{m} \leftarrow \frac{\hat{m}}{\lambda}, \quad \hat{q} \leftarrow \frac{\hat{q}}{\lambda}, \quad \hat{V} \leftarrow \frac{\hat{V}}{\lambda}, \quad V \leftarrow \lambda V. \tag{A.4}$$

The equation (A.4) simplify in this limit to

$$\begin{cases} m = \frac{\hat{m}}{p} \text{tr} \frac{\Sigma \theta^* \theta^{*\top} \Sigma}{1 + \hat{V} \Sigma} \\ q = \frac{1}{p} \text{tr} \frac{\hat{m}^2 \Sigma \theta^* \theta^{*\top} \Sigma^2 + \hat{q} \Sigma^2}{(1 + \hat{V} \Sigma)^2} \\ V = \frac{1}{p} \text{tr} \frac{\Sigma}{1 + \hat{V} \Sigma} \end{cases}, \begin{cases} \hat{m} = \frac{n}{\sqrt{p}} \frac{1}{2\pi} \left( \sqrt{2\pi} \left( 1 + \text{erf} \left( \frac{1}{\sqrt{2q(1-\eta)}} \right) \right) + 2e^{-\frac{1}{2q(1-\eta)}} \sqrt{q(1-\eta)} \right) \\ \hat{q} = \frac{n}{\sqrt{2}} \int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] (1 - \sqrt{qx})^2 \\ \hat{V} = \frac{n}{\sqrt{p}} \int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] \end{cases}. \quad (\text{A.5})$$

Note that the risk studied by [28] (A.2) differs from the one we consider (12) by the scaling  $1/\sqrt{p}$ , the missing  $1/n$  in front of the sum, and a factor 2 for the regularization strength. All those scalings can be absorbed in  $\lambda \leftarrow 2\lambda/n$  and  $\Sigma \leftarrow \Sigma/p$ , leading to (13):

$$\begin{cases} m = \hat{m} p \text{tr} \frac{\Sigma \theta^* \theta^{*\top} \Sigma}{1 + p \hat{V} \Sigma} \\ q = p \text{tr} \frac{p \hat{m}^2 \Sigma \theta^* \theta^{*\top} \Sigma^2 + \hat{q} \Sigma^2}{(1 + p \hat{V} \Sigma)^2} \\ V = \text{tr} \frac{\Sigma}{1 + p \hat{V} \Sigma} \end{cases}, \begin{cases} \hat{m} = \frac{n}{\sqrt{p}} \frac{1}{2\pi} \left( \sqrt{2\pi} \left( 1 + \text{erf} \left( \frac{1}{\sqrt{2q(1-\eta)}} \right) \right) + 2e^{-\frac{1}{2q(1-\eta)}} \sqrt{q(1-\eta)} \right) \\ \hat{q} = \frac{n}{\sqrt{2}} \int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] (1 - \sqrt{qx})^2 \\ \hat{V} = \frac{n}{\sqrt{p}} \int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] \end{cases}, \quad (\text{A.6})$$

thereby completing the mapping from the setup in [28] to the present setting.

### A.2. Equations for max-margin under source and capacity conditions

In the following, we detail the asymptotic scaling analysis of the equation (13). For a diagonal covariance  $\Sigma = \text{diag}(\omega_1, \dots, \omega_p)$  and  $\theta^* = (\theta_k^*)_k$ , (13) reads

$$\begin{cases} \rho = \sum_{k=1}^p \theta_k^{*2} \omega_k \\ m = \hat{m} p \sum_{k=1}^d \frac{\omega_k^2 \theta_k^{*2}}{1 + \hat{V} p \omega_k} \\ q = p \sum_{k=1}^p \frac{p \hat{m}^2 \theta_k^{*2} \omega_k^3 + \hat{q} \omega_k^2}{(1 + \hat{V} p \omega_k)^2} \\ V = \sum_{k=1}^p \frac{\omega_k}{1 + \hat{V} p \omega_k} \end{cases}, \begin{cases} \hat{m} = \frac{n}{\sqrt{p}} \frac{1}{2\pi} \left( \sqrt{2\pi} \left( 1 + \text{erf} \left( \frac{1}{\sqrt{2q(1-\eta)}} \right) \right) + 2e^{-\frac{1}{2q(1-\eta)}} \sqrt{q(1-\eta)} \right) \\ \hat{q} = \frac{n}{\sqrt{2}} \int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] (1 - \sqrt{qx})^2 \\ \hat{V} = \frac{n}{\sqrt{p}} \int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] \end{cases}. \quad (\text{A.7})$$

To simplify the equation (A.7), we introduce the auxiliary variables  $z = \frac{n}{p} / \hat{V}$ ,  $\hat{r}_1 = \hat{m} / \hat{V}$  and  $\hat{r}_2 = (n\hat{q}) / (p\hat{V}^2)$ . The intuitive meaning of  $z$  is loosely that of an effective regularization. In the context of kernel ridge regression where a similar variable appears, the role of  $z$  as an effective regularizing term is quite clear in [15]. We also refer the reader to the discussion in [28], also for ridge regression, where the role of  $\hat{V}$  as parametrizing the gap between training and test error is mentioned.  $\hat{r}_1$  is to be regarded as the ratio between the norm of the estimator  $\hat{w}$  minimizing the risk (12), and the norm of the teacher  $\theta^*$ . Introducing these variables in (A.7) allows to have a well defined  $p = \infty$  limit, which reads

$$\begin{cases} \rho = \sum_{k=1}^{\infty} \theta_k^{*2} \omega_k \\ m = \hat{r}_1 \frac{n}{z} \sum_{k=1}^{\infty} \frac{\omega_k^2 \theta_k^{*2}}{1 + \frac{n}{z} \omega_k} \\ q = \frac{n^2}{z^2} \sum_{k=1}^{\infty} \frac{\hat{r}_1^2 \theta_k^{*2} \omega_k^3 + \frac{1}{n} \hat{r}_2 \omega_k^2}{(1 + \frac{n}{z} \omega_k)^2} \\ V = \sum_{k=1}^d \frac{\omega_k}{1 + \frac{n}{z} \omega_k} \end{cases}, \begin{cases} \hat{r}_1 = \frac{\hat{m}}{\hat{V}} = \frac{1}{2\pi\sqrt{p}} \frac{\left( \sqrt{2\pi} \left( 1 + \text{erf} \left( \frac{1}{\sqrt{2q(1-\eta)}} \right) \right) + 2e^{-\frac{1}{2q(1-\eta)}} \sqrt{q(1-\eta)} \right)}{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right]} \\ \hat{r}_2 = \frac{n\hat{q}}{\hat{V}^2} = \frac{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] (1 - \sqrt{qx})^2}{\left( \int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \text{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] \right)^2} \end{cases}. \quad (\text{A.8})$$

From (A.7),  $z$  satisfies the self-consistent equation

$$z = \frac{\frac{z}{n} \sum_{k=1}^{\infty} \frac{\omega_k}{\frac{z}{n} + \omega_k}}{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right]} \approx \frac{\left(\frac{z}{n}\right)^{1-\frac{1}{\alpha}} \int_{\left(\frac{z}{n}\right)^{\frac{1}{\alpha}} \frac{dx}{1+x^\alpha}}^{\infty}}{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right]}, \tag{A.9}$$

where a Riemann approximation of the sum was used. We introduce its to-be-determined scaling  $\gamma$ ,

$$z \sim C_z n^{1-\alpha\gamma}, \tag{A.10}$$

where  $C_z$  designates the prefactor. In the following, we discuss the scaling and the corrections of the order parameters and express them as a function of  $\gamma$ , before determining its value using the numerical solution of (A.8). Note that the scaling (A.10) implies in particular the following scaling for the integral

$$\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] \sim n^{\gamma-1}. \tag{A.11}$$

### A.3. First order corrections

We now plug the source and capacity ansatz (7) into the equation (A.8)

$$\omega_k \sim k^{-\alpha}, \quad \theta_k^* \sim k^{-\frac{1+\alpha(2r-1)}{2}}, \tag{A.12}$$

yielding

$$\begin{cases} \rho = \sum_{k=1}^{\infty} k^{-1-2\alpha r} \\ m = \hat{r}_1 \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r}}{1+\frac{z}{n}k^\alpha} \\ q = \hat{r}_1^2 \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r}}{\left(1+\frac{z}{n}k^\alpha\right)^2} + \hat{r}_2 \sum_{k=1}^{\infty} \frac{1}{\left(1+\frac{z}{n}k^\alpha\right)^2} \\ V = \frac{z}{n} \sum_{k=1}^{\infty} \frac{1}{1+\frac{z}{n}k^\alpha} \end{cases}, \quad \begin{cases} \hat{r}_1 = \frac{1}{2\pi\sqrt{\beta}} \frac{\left(\sqrt{2\pi}\left(1+\operatorname{erf}\left(\frac{1}{\sqrt{2q(1-\eta)}}\right)\right)+2e^{-\frac{1}{2q(1-\eta)}}\sqrt{q(1-\eta)}\right)}{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right]} \\ \hat{r}_2 = \frac{\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] (1-\sqrt{q}x)^2}{\left(\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right]\right)^2} \end{cases}.$$

Massaging the expression for  $m$  (and remembering  $z \sim n^{1-\alpha\gamma}$  (A.10))

$$\begin{aligned} m &= \hat{r}_1 \rho - \hat{r}_1 \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r} C_z \left(\frac{k}{n^\gamma}\right)^\alpha}{1 + C_z \left(\frac{k}{n^\gamma}\right)^\alpha} \\ &= \hat{r}_1 \rho - C_z \hat{r}_1 n^{-\gamma\alpha} \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha(r-\frac{1}{2})}}{1 + C_z \left(\frac{k}{n^\gamma}\right)^\alpha} \\ &\approx \hat{r}_1 \rho - \mathbb{1}_{r \geq \frac{1}{2}} \hat{r}_1 C_z \left[ n^{-\gamma\alpha} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-\frac{1}{2})}}{1 + C_z \left(\frac{k}{n^\gamma}\right)^\alpha} + n^{-2\gamma\alpha r} \int_1^{\infty} \frac{x^{-1+\alpha(1-2r)}}{1 + C_z x^\alpha} \right] - \mathbb{1}_{r \leq \frac{1}{2}} C_z \hat{r}_1 n^{-2\gamma\alpha r} \int_0^{\infty} \frac{x^{-1+\alpha(1-2r)}}{1 + C_z x^\alpha} \\ &= \hat{r}_1 \left[ \rho - \mathbb{1}_{r \geq \frac{1}{2}} C_z A_m^n n^{-\gamma\alpha} - C_z n^{-2\gamma\alpha r} \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{1 + C_z x^\alpha} + o\left(\max\left(n^{-\gamma\alpha}, n^{-2\gamma\alpha r}\right)\right) \right], \end{aligned} \tag{A.13}$$

where we used the shorthand

$$A_m^n \stackrel{\text{def}}{=} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-\frac{1}{2})}}{1 + C_z \left(\frac{k}{n^\gamma}\right)^\alpha}. \tag{A.14}$$

Massaging by the same token the equation for  $q$ ,

$$\begin{aligned}
 q &\approx \hat{r}_1^2 \rho - \hat{r}_1^2 \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r} \left( C_z^2 \left( \frac{k}{n^\gamma} \right)^{2\alpha} + 2C_z \left( \frac{k}{n^\gamma} \right)^\alpha \right)}{\left( 1 + C_z \left( \frac{k}{n^\gamma} \right)^\alpha \right)^2} + \hat{r}_2 n^{\gamma-1} \int_0^\infty \frac{1}{(1 + C_z x^\alpha)^2} \\
 &\approx \hat{r}_1^2 \rho + \hat{r}_2 n^{\gamma-1} \int_0^\infty \frac{1}{(1 + C_z x^\alpha)^2} \\
 &\quad - 2\hat{r}_1^2 C_z \left[ \mathbb{1}_{r \geq \frac{1}{2}} \left( n^{-\gamma\alpha} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-\frac{1}{2})}}{\left( 1 + C_z \left( \frac{k}{n^\gamma} \right)^\alpha \right)^2} + n^{-2\gamma\alpha r} \int_1^\infty \frac{x^{-1+\alpha(1-2r)}}{(1 + C_z x^\alpha)^2} \right) \right. \\
 &\quad \left. + \mathbb{1}_{r \leq \frac{1}{2}} n^{-2\gamma\alpha r} \int_0^\infty \frac{x^{-1+\alpha(1-2r)}}{(1 + C_z x^\alpha)^2} \right] \\
 &\quad - \hat{r}_1^2 C_z^2 \left[ \mathbb{1}_{r \geq 1} \left( n^{-2\gamma\alpha} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-1)}}{\left( 1 + C_z \left( \frac{k}{n^\gamma} \right)^\alpha \right)^2} \right. \right. \\
 &\quad \left. \left. + n^{-2\gamma\alpha r} \int_1^\infty \frac{x^{-1+\alpha(2-2r)}}{(1 + C_z x^\alpha)^2} \right) + \mathbb{1}_{r \leq 1} n^{-2\gamma\alpha r} \int_0^\infty \frac{x^{-1+\alpha(2-2r)}}{(1 + C_z x^\alpha)^2} \right] \\
 &= \hat{r}_1^2 \left[ \rho + \frac{\hat{r}_2}{\hat{r}_1^2} n^{\gamma-1} \int_0^\infty \frac{1}{(1 + C_z x^\alpha)^2} - \mathbb{1}_{r \geq \frac{1}{2}} 2C_z A_q^n n^{-\gamma\alpha} - \mathbb{1}_{r \geq 1} C_z^2 B_q^n n^{-2\gamma\alpha} \right. \\
 &\quad \left. - n^{-2\gamma\alpha r} \left( 2C_z \int_{\Theta(r-\frac{1}{2})}^\infty \frac{x^{-1+\alpha(1-2r)}}{(1 + C_z x^\alpha)^2} + C_z^2 \int_{\Theta(r-1)}^\infty \frac{x^{-1+\alpha(2-2r)}}{(1 + C_z x^\alpha)^2} \right) \right. \\
 &\quad \left. + o \left( \max \left( \frac{\hat{r}_2}{\hat{r}_1^2} n^{\gamma-1}, n^{-\gamma\alpha}, n^{-2\gamma\alpha r} \right) \right) \right], \tag{A.15}
 \end{aligned}$$

where we defined

$$A_q^n \stackrel{\text{def}}{=} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-\frac{1}{2})}}{\left( 1 + C_z \left( \frac{k}{n^\gamma} \right)^\alpha \right)^2}, \quad B_q^n \stackrel{\text{def}}{=} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-1)}}{\left( 1 + C_z \left( \frac{k}{n^\gamma} \right)^\alpha \right)^2}. \tag{A.16}$$

An important remark on (A.13) and (A.15) is that  $A_q^m$  and  $A_m^n$  tend to the same limit as  $n \rightarrow \infty$ ,

$$A_{m,q}^\infty = \sum_{k=1}^\infty k^{-1-2\alpha(r-\frac{1}{2})}, \tag{A.17}$$

with (for  $r \geq \frac{1}{2}$  since there is an indicator function in front of  $A_{m,q}^n$ )

$$\begin{cases}
 n^{-\gamma\alpha} |A_m^n - A^\infty| = \frac{n^{-2\alpha\gamma r}}{2\alpha(r-\frac{1}{2})} + \mathbb{1}_{\frac{1}{2} \leq r \leq 1} C_z n^{-2\alpha\gamma r} \int_0^1 \frac{x^{-1-2\alpha(r-1)}}{1+C_z x^\alpha} \\
 \quad + \mathbb{1}_{r \geq 1} C_z n^{-2\gamma\alpha} \sum_{k=1}^\infty k^{-1-2\alpha(r-1)}, \\
 n^{-\gamma\alpha} |A_q^n - A^\infty| = \frac{n^{-2\alpha\gamma r}}{2\alpha(r-\frac{1}{2})} + 2\mathbb{1}_{\frac{1}{2} \leq r \leq 1} C_z n^{-2\alpha\gamma r} \int_0^1 \frac{x^{-1-2\alpha(r-1)}}{(1+C_z x^\alpha)^2} \\
 \quad + 2\mathbb{1}_{r \geq 1} C_z n^{-2\gamma\alpha} \sum_{k=1}^\infty k^{-1-2\alpha(r-1)} + o(n^{-2\gamma\alpha}).
 \end{cases}$$

More straightforwardly,

$$B_q^\infty = \sum_{k=1}^\infty k^{-1-2\alpha(r-1)}. \tag{A.18}$$

Finally, remark that the expansions (A.13) and (A.15) imply that  $m \sim \hat{r}_1 \rho$  and  $q \sim \hat{r}_1^2 \rho$ , which echoes the intuitive meaning of  $m, q$  as  $\langle \hat{w}, \theta^* \rangle_{L^2(\mathcal{X})}$ ,  $\|\hat{w}\|_{L^2(\mathcal{X})}^2$  and of  $\hat{r}_1$  as  $\|\hat{w}\|_{L^2(\mathcal{X})} / \|\theta^*\|_{L^2(\mathcal{X})}$ , see also discussion in section 2 under equation (11).

The expansions (A.13) and (A.15) can be plugged into the expression of the cosine similarity  $\eta$  (9) to access the decay rate for the misclassification error.

$$\begin{aligned} \eta &= \left( 1 - \mathbb{1}_{r \geq \frac{1}{2}} C_z \frac{A_m^n}{\rho} n^{-\gamma\alpha} - n^{-2\gamma\alpha r} \frac{C_z}{\rho} \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{1+C_z x^\alpha} \right)^2 / \\ &\quad \left( 1 + \frac{\hat{r}_2}{\rho \hat{r}_1^2} n^{\gamma-1} \int_0^\infty \frac{1}{(1+x^\alpha)^2} - \frac{\mathbb{1}_{r \geq \frac{1}{2}} 2C_z A_q^n n^{-\gamma\alpha} - \mathbb{1}_{r \geq 1} C_z B_q^n n^{-2\gamma\alpha}}{\rho} \right. \\ &\quad \left. - n^{-2\gamma\alpha r} \frac{C_z}{\rho} \left( 2 \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{(1+C_z x^\alpha)^2} + C_z \int_{\Theta(r-1)}^{\infty} \frac{x^{-1+\alpha(2-2r)}}{(1+C_z x^\alpha)^2} \right) \right)^{-1} \\ &= 1 - \mathcal{O} \left( \frac{\hat{r}_2}{\rho \hat{r}_1^2} n^{\gamma-1} \right) + \mathbb{1}_{r \geq 1} \mathcal{O} (n^{-2\gamma\alpha}) + \mathcal{O} (n^{-2\gamma\alpha r}), \end{aligned} \tag{A.19}$$

where we used  $A_q^\infty = A_m^\infty$ , meaning the leading order of the  $n^{-\alpha\gamma}$  term cancels out. The scaling of the other order parameters  $q, m, \hat{r}_1, \hat{r}_2$  can at this point also be deduced. To see how, notice that

$$\int_{-\infty}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}} x \right) \right] \approx \frac{1}{\sqrt{q}} \int_{-\infty}^1 dx \frac{e^{-\frac{x^2}{2q}}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2q(1-\eta)}} x \right) \right] \sim \frac{1}{\sqrt{q}}, \tag{A.20}$$

which together with the observation (A.11) implies

$$q \sim \hat{r}_1^2 \sim n^{2(1-\gamma)}. \tag{A.21}$$

In parallel, it follows from a similar change of variables in the expression of  $\hat{r}_2$  that

$$\hat{r}_2 \sim \sqrt{q} \sim \hat{r}_1 \sim n^{1-\gamma}. \tag{A.22}$$

Then, returning to (A.19),

$$1 - \eta = \mathcal{O} \left( n^{2(\gamma-1)} \right) + \mathbb{1}_{r \geq 1} \mathcal{O} (n^{-2\gamma\alpha}) + \mathcal{O} (n^{-2\gamma\alpha r}). \tag{A.23}$$

The final step is to determine the value of  $\gamma$ . Suppose self-consistently that the leading rate in (A.15) is  $2(\gamma - 1)$ . Since from the physical meaning of the order parameters (11) one expects the leading corrections of  $m$  and  $q$  to share the same decay rate—since in (11), only  $\hat{w}$  admits a correction with  $n^{-}$ , and since the leading corrections of  $m$  has rate  $-2\gamma\alpha \min(r, 1/2)$  (A.13), it follows from equating the two rates that  $\gamma$  should be

$$\gamma = \frac{1}{1 + \alpha \min(r, \frac{1}{2})}. \tag{A.24}$$

Plugging the expression (A.24) back in (A.15), it can be checked that the  $2(\gamma - 1)$  term is indeed leading. (A.24) compares very well with numerical simulations, and with the numerical solution of (A.8), see figure 1.

Assuming (A.24) to be true, since  $\epsilon_g = \cos^{-1}(\eta) \sim \sqrt{1-\eta}$ ,

$$\epsilon_g \sim n^{-\gamma\alpha \min(r, \frac{1}{2})} = n^{-\frac{\alpha \min(r, \frac{1}{2})}{1 + \alpha \min(r, \frac{1}{2})}}, \tag{A.25}$$

which is the error rate for max-margin classification (15).

### Appendix B. Rates for regularized hinge classification

In this appendix, we show that margin-maximizing SVM is optimal for classifying teachers with source  $r \leq \frac{1}{2}$ , i.e. in  $L^2(\mathcal{X}) \setminus \mathcal{H}$ , and argue why this is also expected to extend to teachers in  $\mathcal{H}$ . In other words, we give evidence that the optimal regularization in the hinge classification risk (12) on the dataset (8) is for vanishing  $\lambda = 0$ . To this end, we derive the rates for the misclassification error of ERM on (12) for generic regularization  $\lambda$ , building upon the results of [28] in similar fashion to appendix A.



**B.1. Regularized hinge classification under source and capacity conditions**

We remind the fixed point equations for the risk (12) (see appendix A):

$$\begin{cases} m = \hat{m}p \operatorname{tr} \frac{\Sigma \theta^* \theta^{*\top} \Sigma}{\frac{n\lambda}{2} + p\hat{V}\Sigma} \\ q = p \operatorname{tr} \frac{p\hat{m}^2 \Sigma \theta^* \theta^{*\top} \Sigma^2 + \hat{q}\Sigma^2}{(\frac{n\lambda}{2} + p\hat{V}\Sigma)^2} \\ V = \operatorname{tr} \frac{\Sigma}{\frac{n\lambda}{2} + p\hat{V}\Sigma} \\ \hat{m} = \frac{\frac{n}{p}}{V\sqrt{\rho}} \frac{1}{2\pi} \left\{ \sqrt{2\pi} \left[ \operatorname{erf} \left( \frac{1}{\sqrt{2q(1-\eta)}} \right) - \operatorname{erf} \left( \frac{1-V}{\sqrt{2q(1-\eta)}} \right) \right] \right. \\ \left. + 2\sqrt{q(1-\eta)} \left( e^{-\frac{1}{2q(1-\eta)}} - e^{-\frac{(1-V)^2}{2q(1-\eta)}} \right) + \sqrt{2\pi}V \left( 1 + \operatorname{erf} \left( \frac{1-V}{\sqrt{2q(1-\eta)}} \right) \right) \right\} \\ \hat{q} = \frac{\frac{n}{p}}{V} \left[ \int_{-\infty}^{\frac{1-V}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right] + \frac{1}{V^2} \int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right] (1 - \sqrt{q}x)^2 \right] \\ \hat{V} = \frac{\frac{n}{p}}{V} \int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right] \end{cases} \quad . \quad (\text{B.1})$$

Introducing the effective regularization  $z = \frac{\alpha n \lambda}{2\hat{V}}$ , satisfying

$$z = \frac{\frac{z}{n} \sum_{k=1}^{\infty} \frac{\lambda_k}{\frac{z}{n} + \lambda_k}}{\int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right]} \approx \frac{\left(\frac{z}{n}\right)^{1-\frac{1}{\alpha}} \int_{\left(\frac{z}{n}\right)^{\frac{1}{\alpha}}}^{\infty} \frac{dx}{1+x^\alpha}}{\int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right]} \quad , \quad (\text{B.2})$$

the equation (B.1) can be rewritten as

$$\begin{cases} \rho = \sum_{k=1}^{\infty} \theta_k^2 \lambda_k \\ m = \hat{r}_1 \frac{n}{z} \sum_{k=1}^{\infty} \frac{\lambda_k^2 \theta_k^2}{1 + \frac{n}{z} \lambda_k} \\ q = \frac{n^2}{z^2} \sum_{k=1}^{\infty} \frac{\hat{r}_1^2 \theta_k^2 \lambda_k^3 + \frac{1}{2} \hat{r}_2 \lambda_k^2}{\left(1 + \frac{n}{z} \lambda_k\right)^2} \\ V = \frac{1}{n\lambda} \sum_{k=1}^{\infty} \frac{\lambda_k}{1 + \frac{n}{z} \lambda_k} \end{cases} \quad (\text{B.3})$$

$$\begin{cases} \hat{r}_1 = \frac{\hat{m}}{\hat{V}} = \frac{1}{2\pi\sqrt{\rho}} \frac{\sqrt{2\pi} \left[ \operatorname{erf} \left( \frac{1}{\sqrt{2q(1-\eta)}} \right) - \operatorname{erf} \left( \frac{1-V}{\sqrt{2q(1-\eta)}} \right) \right]}{\int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right]} \\ \quad + \frac{1}{2\pi\sqrt{\rho}} \frac{2\sqrt{q(1-\eta)} \left( e^{-\frac{1}{2q(1-\eta)}} - e^{-\frac{(1-V)^2}{2q(1-\eta)}} \right) + \sqrt{2\pi}V \left( 1 + \operatorname{erf} \left( \frac{1-V}{\sqrt{2q(1-\eta)}} \right) \right)}{\int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right]} \\ \hat{r}_2 = \frac{\alpha \hat{q}}{\hat{V}^2} = \frac{V^2 \int_{-\infty}^{\frac{1-V}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right] + \int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right] (1 - \sqrt{q}x)^2}{\left( \int_{\frac{1-V}{\sqrt{q}}}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf} \left( \sqrt{\frac{\eta}{2(1-\eta)}}x \right) \right] \right)^2} \end{cases} .$$

Specializing to the source/capacity power-law forms (8)

$$\omega_k \sim k^{-\alpha}, \quad \theta_k \sim k^{-\frac{1+\alpha(2r-1)}{2}}, \quad (\text{B.4})$$

one finally reaches

$$\begin{cases} \rho = \sum_{k=1}^{\infty} k^{-1-2\alpha r} \\ m = \hat{r}_1 \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r}}{1 + \frac{z}{n} k^\alpha} \\ q = \hat{r}_1^2 \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r}}{(1 + \frac{z}{n} k^\alpha)^2} + \frac{\hat{r}_2}{n} \sum_{k=1}^{\infty} \frac{1}{(1 + \frac{z}{n} k^\alpha)^2} \\ V = \frac{1}{n\lambda} \frac{z}{n} \sum_{k=1}^{\infty} \frac{1}{1 + \frac{z}{n} k^\alpha} \end{cases} \tag{B.5}$$

$$\begin{cases} \hat{r}_1 = \frac{\hat{m}}{\hat{V}} = \frac{1}{2\pi\sqrt{\rho}} \frac{\sqrt{2\pi} \left[ \operatorname{erf}\left(\frac{1}{\sqrt{2q(1-\eta)}}\right) - \operatorname{erf}\left(\frac{1-V}{\sqrt{2q(1-\eta)}}\right) \right]}{\int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right]} \\ + \frac{1}{2\pi\sqrt{\rho}} \frac{2\sqrt{q(1-\eta)} \left( e^{-\frac{1}{2q(1-\eta)}} - e^{-\frac{(1-V)^2}{2q(1-\eta)}} \right) + \sqrt{2\pi}V \left( 1 + \operatorname{erf}\left(\frac{1-V}{\sqrt{2q(1-\eta)}}\right) \right)}{\int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right]} \\ \hat{r}_2 = \frac{\alpha \hat{q}}{\hat{V}^2} = \frac{V^2 \int_{-\infty}^{\frac{1-V}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right] + \int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right] (1 - \sqrt{q}x)^2}{\left( \int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right] \right)^2} \end{cases} \tag{B.6}$$

**B.2. Rate analysis for targets outside the Hilbert space**

In the following, we deliver an asymptotic analysis of equation (B.6). To that end, we first ascertain the scaling of the effective regularization  $z$  using the self-consistent equation (B.2). Depending on the scaling of the order parameter  $V$ , two regimes can be distinguished.

*B.2.1. Effectively regularized regime  $V \xrightarrow{n \rightarrow \infty} \infty$*

In this regime the denominator of (B.2) scales like

$$\int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right] \approx \frac{1}{\sqrt{q}} \int_{-\infty}^1 dx \frac{e^{-\frac{1}{2q}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2q(1-\eta)}}x\right) \right] = \mathcal{O}\left(\frac{1}{\sqrt{q}}\right), \tag{B.7}$$

which is exactly the scaling for the corresponding integral in the vanishing regularization case  $\lambda = 0^+$ , see appendix A equation (A.11). In this regime, all the discussion in appendix A then carries through and the error rate coincides with the one for margin-maximizing SVM (15)

$$\epsilon_g \sim n^{-\frac{\alpha \min(r, \frac{1}{2})}{1 + \alpha \min(r, \frac{1}{2})}}. \tag{B.8}$$

*B.2.2. Effectively regularized regime  $V \xrightarrow{n \rightarrow \infty} 0$*

The integral in the denominator in (B.2) then admits the following scaling

$$\int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right] = \frac{1}{\sqrt{q}} \int_{1-V}^1 dx \frac{e^{-\frac{1}{2q}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2q(1-\eta)}}x\right) \right] = \mathcal{O}\left(\frac{V}{\sqrt{q}}\right). \tag{B.9}$$

This means in particular that

$$z = \frac{Vn\lambda}{\int_{\frac{1-V}{\sqrt{q}}}^{\frac{1}{\sqrt{q}}} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \left[ 1 + \operatorname{erf}\left(\sqrt{\frac{\eta}{2(1-\eta)}}x\right) \right]} = \mathcal{O}(n\lambda\sqrt{q}) \stackrel{\text{def}}{\sim} C_2 n^{1-\alpha\gamma}, \tag{B.10}$$

where we introduced similarly to the max-margin case the to-be determined parameter  $\gamma$ , see also appendix A. Note that it follows from equation (B.10) and the definition of  $\lambda = n^{-\ell}$  that

$$q \sim n^{2(\ell-\alpha\gamma)}. \tag{B.11}$$

Since  $n\lambda V \sim (z/n)^{1-\frac{1}{\alpha}}$ , one also has that

$$V \sim n^{-1+\gamma(1-\alpha)+\ell}. \tag{B.12}$$

In particular, it can be seen from (B.12) that the assumption  $V \xrightarrow{n \rightarrow \infty} 0$  only holds for large enough regularizations (slow enough decays  $\ell < \ell^*$  for some limiting value  $\ell^*$ ), satisfying

$$-1 + \gamma(1 - \alpha) + \ell < 0. \tag{B.13}$$

The regularization decay  $\ell = \ell^*$  gives the boundary between the effectively regularized and unregularized regime. Because of this, the rate of  $q$  at  $\ell^*$  should coincide with its max-margin rate (14). This, together with (B.13), allow to determine  $\ell^*$  as the solution of the system (denoting  $\gamma^*$  the value of  $\gamma$  at  $\ell = \ell^*$ )

$$\begin{cases} -1 + \gamma^*(1 - \alpha) + \ell = 0, \\ \ell^* - \alpha\gamma^* = \frac{\alpha \min(r, \frac{1}{2})}{1 + \alpha \min(r, \frac{1}{2})}, \end{cases} \tag{B.14}$$

viz.

$$\ell^* = \alpha \frac{1 + \min(r, \frac{1}{2})}{1 + \alpha \min(r, \frac{1}{2})}. \tag{B.15}$$

Summarizing, for  $\ell > \ell^*$ , in the effectively unregularized regime, the max-margin scalings (14) and (15) hold. In the following, we focus on pursuing the discussion for the new  $\ell < \ell^*$  (effectively regularized) regime.

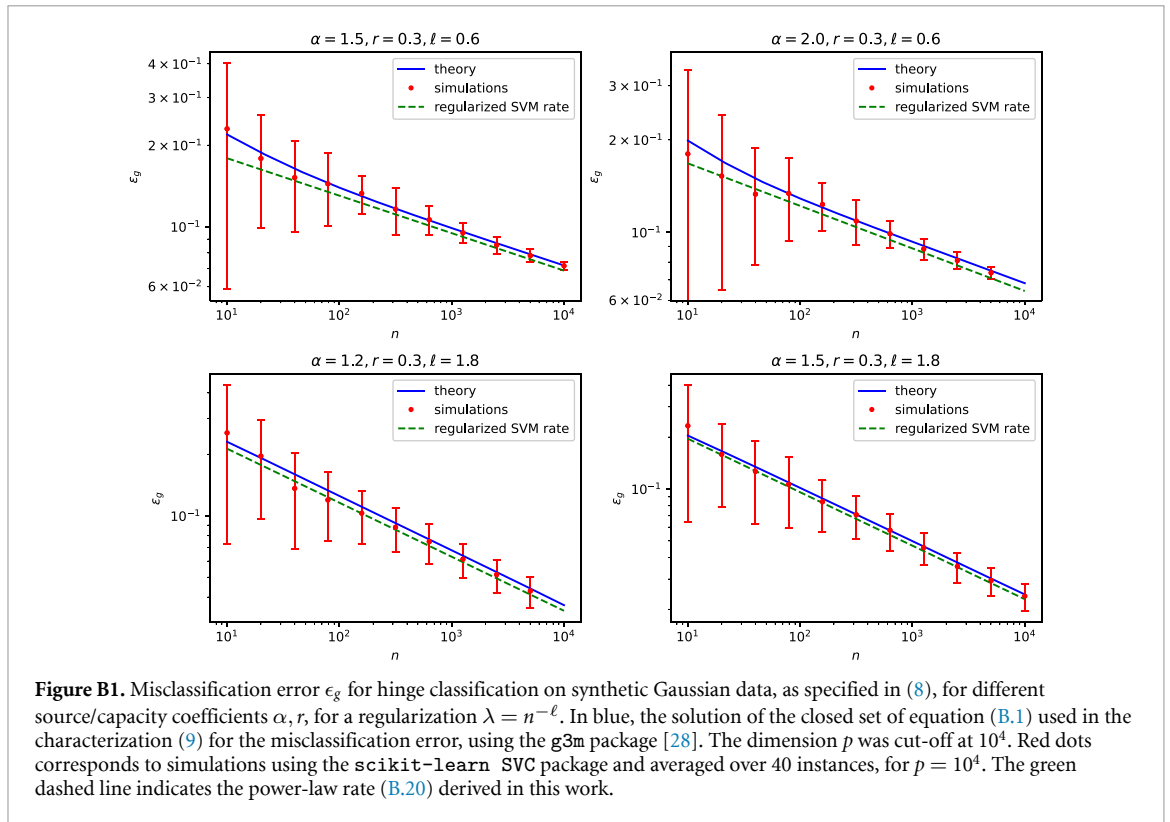
The expansion for  $m$  and  $q$  carry over in similar fashion to max-margin (see appendix A), yielding

$$\begin{cases} m = \hat{r}_1 \left[ \rho - \mathbb{1}_{r \geq \frac{1}{2}} C_z A_m^n n^{-\gamma\alpha} - C_z n^{-2\gamma\alpha} \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{1+C_z x^\alpha} \right] + o(\max(n^{-\gamma\alpha}, n^{-2\gamma\alpha})), \\ q = \hat{r}_1^2 \left[ \rho + \frac{\hat{r}_2}{\hat{r}_1} n^{\gamma-1} \int_0^{\infty} \frac{1}{(1+C_z x^\alpha)^2} - \mathbb{1}_{r \geq \frac{1}{2}} 2C_z A_q^n n^{-\gamma\alpha} - \mathbb{1}_{r \geq 1} C_z^2 B_q^n n^{-2\gamma\alpha} \right. \\ \left. - n^{-2\gamma\alpha} \left( 2C_z \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{(1+C_z x^\alpha)^2} + C_z^2 \int_{\Theta(r-1)}^{\infty} \frac{x^{-1+\alpha(2-2r)}}{(1+C_z x^\alpha)^2} \right) \right] \\ + o\left(\max\left(\frac{\hat{r}_2}{\hat{r}_1} n^{\gamma-1}, n^{-\gamma\alpha}, n^{-2\gamma\alpha}\right)\right). \end{cases}$$

Therefore  $\hat{r}_1 \sim \sqrt{q} \sim n^{\ell-\alpha\gamma}$ , while a rescaling of the integrals in the equations for  $\hat{r}_2$  also reveal that  $\hat{r}_2 \sim \hat{r}_1 \sim \sqrt{q} \sim n^{\ell-\alpha\gamma}$ . Summarizing

$$m \sim \hat{r}_2 \sim \hat{r}_1 \sim \sqrt{q} \sim n^{\ell-\alpha\gamma}. \tag{B.16}$$

Note that the mutual scaling of the order parameters  $m, q, \hat{r}_1$  is the same as for max-margin classification (14), which is the mutual scaling directly following from the physical interpretation of  $m, q$  as  $\langle \hat{w}, \theta^* \rangle_{L^2(\mathcal{X})}$ ,  $\|\hat{w}\|_{L^2(\mathcal{X})}^2$  and of  $\hat{r}_1$  as  $\|\hat{w}\|_{L^2(\mathcal{X})} / \|\theta^*\|_{L^2(\mathcal{X})}$ , see also discussion in section 2 under equation (11). Finally,



$$\begin{aligned}
 \eta &= \left( 1 - \mathbb{1}_{r \geq \frac{1}{2}} C_z \frac{A_m^n}{\rho} n^{-\gamma\alpha} - n^{-2\gamma\alpha r} \frac{C_z}{\rho} \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{1+C_z x^\alpha} \right)^2 / \\
 &\quad \left( 1 + \frac{\hat{r}_2}{\rho \hat{r}_1^2} n^{\gamma-1} \int_0^{\infty} \frac{1}{(1+x^\alpha)^2} - \frac{\mathbb{1}_{r \geq \frac{1}{2}} 2C_z A_q^n n^{-\gamma\alpha} - \mathbb{1}_{r \geq 1} C_z^2 B_q^n n^{-2\gamma\alpha}}{\rho} \right. \\
 &\quad \left. - n^{-2\gamma\alpha r} \frac{C_z}{\rho} \left( 2 \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{(1+C_z x^\alpha)^2} + C_z \int_{\Theta(r-1)}^{\infty} \frac{x^{-1+\alpha(2-2r)}}{(1+C_z x^\alpha)^2} \right) \right)^{-1} \\
 &= 1 - \mathcal{O} \left( \frac{\hat{r}_2}{\rho \hat{r}_1^2} n^{\gamma-1} \right) + \mathbb{1}_{r \geq 1} \mathcal{O} \left( n^{-2\gamma\alpha} \right) + \mathcal{O} \left( n^{-2\gamma\alpha r} \right) \\
 &= 1 - \mathcal{O} \left( n^{\gamma(\alpha+1)-1-\ell} \right) + \mathcal{O} \left( n^{-2\alpha\gamma r} \right), \tag{B.17}
 \end{aligned}$$

where we specialized to  $r \leq \frac{1}{2}$  in the last line, thereby focusing on target functions  $f^* \in L^2(\mathcal{X}) \setminus \mathcal{H}$ . At this point, it remains to heuristically determine  $\gamma$ , a rigorous analytical derivation of our results being left for future work. Making the two (numerically verified) assumptions

**Assumption B.1** Like the max-margin case,  $1 - \eta \sim q^{-1}$ ,

**Assumption B.2** The term of rate  $2\alpha\gamma r$  dominates in (B.17).

$\gamma$  can be guessed as

$$\gamma = \frac{\ell}{\alpha(1+r)}. \tag{B.18}$$

Note that consistently, the value of  $\gamma$  at the boundary  $\ell = \ell^*$  with the unregularized regime  $\gamma^*$  coincides with its max-margin value (A.24). In the effectively regularized regime, under these assumptions, we thus conjecture the error to scale for  $r \leq \frac{1}{2}$  as

$$\epsilon_g \sim \sqrt{1-\eta} \sim n^{\alpha\gamma-\ell} \sim n^{-\frac{\ell}{1+r}}. \tag{B.19}$$

Finally observe that the rates for the effectively unregularized regime  $\ell \geq \ell^*$  (B.8) and regularized regime  $\ell \leq \ell^*$  (B.19) can be subsumed in the more compact form, still for  $r \leq \frac{1}{2}$ :

$$\epsilon_g \sim n^{-\min(\ell, \ell^*) \frac{r}{1+r}} = n^{-\min(\ell, \alpha \frac{1+r}{1+\alpha r}) \frac{r}{1+r}}. \tag{B.20}$$

Figure B1 contrasts the rates (B.20) to the numerical solution of the equation (B.1) and to numerical simulations, and displays a very good agreement. The main conclusion from (B.20) is that for any  $\ell$ , the rate is necessarily slower than  $\ell^* r / (1+r) = \alpha r / (1+\alpha r)$ , which is the max-margin rate (15). Therefore, the max-margin rate (15) is optimal for  $r \leq \frac{1}{2}$ , and is achieved for any regularization  $\lambda$  decaying at least as fast as  $n^{-\ell^*}$ . In particular, it is achieved for the limit case  $\ell = \infty$ , i.e.  $\lambda = 0^+$  a.k.a. the max-margin case, thereby suggesting that no regularization is optimal for rough enough target functions  $f^* \in L^2(\mathcal{X}) \setminus \mathcal{H}$ . Since regularization is not needed for hard teachers (small source  $r$ ), we do not *a fortiori* expect it to help for easier, smoother teachers  $f^* \in \mathcal{H}$  characterized by source  $r \geq \frac{1}{2}$ . This suggests that  $\lambda = 0^+$  (max-margin) should be optimal for all sources  $r$ , i.e. any target  $f^* \in L^2(\mathcal{X})$ . While this conjecture is observed to hold in numerical simulations, a more thorough theoretical analysis of the error rates for  $r \geq \frac{1}{2}$  is nevertheless warranted. We leave this more challenging analysis to future work.

### Appendix C. Rates for ridge classification

In this section we derive the rates (18) for ridge classifiers.

#### C.1. Ridge classification under the source and capacity conditions

The equations on  $m, q$  for ridge classification [28] can as in appendix A be adapted to the  $p = \infty, n \gg 1$  limit as

$$\begin{cases} m = \hat{m} p \operatorname{tr} \frac{\Sigma \theta^* \theta^{*\top} \Sigma}{n\lambda + p\hat{V}\Sigma} \\ q = p \operatorname{tr} \frac{p\hat{m}^2 \Sigma \theta^* \theta^{*\top} \Sigma^2 + \hat{q}\Sigma^2}{(n\lambda + p\hat{V}\Sigma)^2} \\ V = \operatorname{tr} \frac{\Sigma}{n\lambda + p\hat{V}\Sigma} \end{cases}, \quad \begin{cases} \hat{V} = \frac{\frac{n}{p}}{1+V} \\ \hat{q} = \frac{n}{p} \frac{1+q-2m\sqrt{\frac{2}{\pi\rho}}}{(1+V)^2} = \frac{\hat{V}^2}{\frac{n}{p}} \left(1+q-2m\sqrt{\frac{2}{\pi\rho}}\right) \\ \hat{m} = \sqrt{\frac{2}{\pi\rho}} \frac{\frac{n}{p}}{1+V} = \sqrt{\frac{2}{\pi\rho}} \hat{V}. \end{cases} \tag{C.1}$$

Following [15], we introduce the effective regularization  $z$

$$z = \frac{\frac{n}{p} n\lambda}{\hat{V}} = n\lambda + \frac{z}{n} \sum_{k=1}^{\infty} \frac{\lambda_k}{\lambda_k + \frac{z}{n}} \approx n\lambda + \left(\frac{z}{n}\right)^{1-\frac{1}{\alpha}} \int_{\left(\frac{z}{n}\right)^{\frac{1}{\alpha}}}^{\infty} \frac{dx}{1+x^\alpha}. \tag{C.2}$$

Like in appendix A, we denote  $z \sim C_z n^{1-\alpha\gamma}$  the scaling of  $z$ . In contrast to appendices A and B for SVM, the exponent  $\gamma$  for ridge classifier can be straightforwardly determined. Depending on the relative strength of the two terms in (C.2), two possible values for  $\gamma$  exist:

- $\gamma = \frac{\ell}{\alpha}$  if  $n\lambda \gg (z/n)^{1-1/\alpha}$  in (C.2), i.e.  $\ell < \alpha$ . This correspond to the *effectively regularized regime* [15].
- $\gamma = 1$  if  $n\lambda \ll (z/n)^{1-1/\alpha}$  in (C.2), i.e.  $\ell > \alpha$ . In this regime, the regularization  $\lambda$  is negligible and therefore the learning is *effectively un-regularized* [15].

Rewriting the equation (C.1) using (C.2), the equation (17) of the main text are reached (see section 4):

$$\begin{cases} \rho = \sum_{k=1}^{\infty} k^{-1-2\alpha r} \\ m = \sqrt{\frac{2}{\pi\rho}} \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r}}{1+\frac{z}{n}k^\alpha} \\ q = \frac{2}{\pi\rho} \sum_{k=1}^{\infty} \frac{k^{-1-2\alpha r}}{\left(1+\frac{z}{n}k^\alpha\right)^2} + \frac{1+q-2m\sqrt{\frac{2}{\pi\rho}}}{n} \sum_{k=1}^{\infty} \frac{1}{\left(1+\frac{z}{n}k^\alpha\right)^2} \\ V = \frac{z}{n^2\lambda} \sum_{k=1}^{\infty} \frac{1}{1+\frac{z}{n}k^\alpha} \end{cases}. \tag{C.3}$$

### C.2. Scaling analysis

We proceed to deliver an scaling analysis of equation (17). We know from appendix A the following scalings

$$\left\{ \begin{aligned} m &= \sqrt{\frac{2}{\pi\rho}} \left[ \rho - \mathbb{1}_{r \geq \frac{1}{2}} C_z A_m^n n^{-\gamma\alpha} - n^{-2\gamma\alpha r} C_z \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{1+C_z x^\alpha} \right] + o(\max(n^{-\gamma\alpha}, n^{-2\gamma\alpha r})), \\ q &= \frac{2}{\pi\rho} \left[ \rho + \frac{\pi\rho C_z}{2} \left( 1 + q - 2m\sqrt{\frac{2}{\pi\rho}} \right) n^{\gamma-1} \int_0^{\infty} \frac{1}{(1+C_z x^\alpha)^2} - \mathbb{1}_{r \geq \frac{1}{2}} 2C_z A_q^n n^{-\gamma\alpha} - \mathbb{1}_{r \geq 1} C_z^2 B_q^n n^{-2\gamma\alpha} \right. \\ &\quad \left. - n^{-2\gamma\alpha r} \left( 2C_z \int_{\Theta(r-\frac{1}{2})}^{\infty} \frac{x^{-1+\alpha(1-2r)}}{(1+C_z x^\alpha)^2} + C_z^2 \int_{\Theta(r-1)}^{\infty} \frac{x^{-1+\alpha(2-2r)}}{(1+C_z x^\alpha)^2} \right) \right] \\ &\quad + o\left(\max\left(\frac{\hat{r}_2}{\hat{r}_1^2} n^{\gamma-1}, n^{-\gamma\alpha}, n^{-2\gamma\alpha r}\right)\right). \end{aligned} \right.$$

Once again, we defined the shorthands

$$A_m^n \stackrel{\text{def}}{=} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-\frac{1}{2})}}{1+C_z\left(\frac{k}{n^\gamma}\right)^\alpha}, \quad A_q^n \stackrel{\text{def}}{=} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-\frac{1}{2})}}{\left(1+C_z\left(\frac{k}{n^\gamma}\right)^\alpha\right)^2}, \quad B_q^n \stackrel{\text{def}}{=} \sum_{k=1}^{n^\gamma} \frac{k^{-1-2\alpha(r-1)}}{\left(1+C_z\left(\frac{k}{n^\gamma}\right)^\alpha\right)^2}. \quad (\text{C.4})$$

We remind that, similarly to the discussion in appendix A, the sequences  $A_m^n$  and  $A_q^n$  admit identical limits as  $n \rightarrow \infty$ :  $A_q^\infty = A_m^\infty$ . The consequences of this identity are expounded below. We now focus on ascertaining the scalings of the order parameters  $m, q$ . These scalings depend on the regime considered.

#### C.2.1. Regularized regime $\ell < \alpha$

In the case  $\ell < \alpha$ , we have  $\gamma = \ell/\alpha < 1$ . Then the expansion for  $q$  reads

$$q(1 - o(1)) = \frac{2}{\pi\rho} \left( \rho - \mathbb{1}_{r \geq \frac{1}{2}} 2C_z A_q^n n^{-\ell} \right) - \mathcal{O}\left(n^{-2\ell\min(r,1)}\right) - \frac{4-\pi}{\pi} C_z n^{-\frac{\alpha-\ell}{\alpha}} \int_0^{\infty} \frac{1}{(1+C_z x^\alpha)^2}, \quad (\text{C.5})$$

from which it follows that

$$q = \frac{2}{\pi\rho} \left( \rho - \mathbb{1}_{r \geq \frac{1}{2}} 2C_z A_q^n n^{-\ell} \right) - \mathcal{O}\left(n^{-\min(2\ell\min(r,1), \frac{\alpha-\ell}{\alpha})}\right). \quad (\text{C.6})$$

Therefore the cosine similarity  $\eta$  admits the following expression

$$\eta = \frac{m^2}{\rho q} = \frac{\frac{2}{\pi\rho} \rho^2 \left[ 1 - \mathbb{1}_{r \geq \frac{1}{2}} \frac{C_z A_m^n}{\rho} n^{-\ell} + \mathcal{O}\left(n^{-2\ell r}\right) \right]^2}{\frac{2}{\pi\rho} \rho^2 \left[ 1 - \mathbb{1}_{r \geq \frac{1}{2}} \frac{2C_z A_q^n}{\rho} n^{-\ell} - \mathcal{O}\left(n^{-\min(2\ell\min(r,1), \frac{\alpha-\ell}{\alpha})}\right) \right]} = 1 - \mathcal{O}\left(n^{-\min(2\ell\min(r,1), \frac{\alpha-\ell}{\alpha})}\right), \quad (\text{C.7})$$

where we used that  $A_q^n$  and  $A_m^n$  share the same limit. Finally, the scaling for the misclassification error can be accessed:

$$\epsilon_g \sim \sqrt{1-\eta} \sim n^{-\frac{1}{2}\min(2\ell\min(r,1), \frac{\alpha-\ell}{\alpha})}, \quad (\text{C.8})$$

which is equation (18).

**Remark C1.** Note that only the classification error (9) tends to zero, while neither the MSE between the label  $y = \text{sign}(\theta^{*\top} \psi(x))$  and the pre-activation linear predictor  $\hat{w}^\top \psi(x)$  (below denoted  $\text{MSE}_1$ ) nor the MSE between the teacher and student pre-activations  $\theta^{*\top} \psi(x), \hat{w}^\top \psi(x)$  (below denoted  $\text{MSE}_2$ ) tend to zero:

$$\text{MSE}_1 = \mathbb{E}_{\psi(x)} \left( \text{sign}(\theta^{*\top} \psi(x)) - \hat{w}^\top \psi(x) \right)^2 = 1 + q - 2m\sqrt{\frac{2}{\pi\rho}} \xrightarrow{n \rightarrow \infty} 1 - \frac{2}{\pi}, \quad (\text{C.9})$$

$$\text{MSE}_2 = \mathbb{E}_{\psi(x)} \left( \theta^{*\top} \psi(x) - \hat{w}^\top \psi(x) \right)^2 = \rho + q - 2m \xrightarrow{n \rightarrow \infty} \rho \left( 1 - \sqrt{\frac{2}{\pi\rho}} \right)^2. \quad (\text{C.10})$$

Note also the identity between the order parameter  $\hat{r}_2$  and  $\text{MSE}_1$ :  $\hat{r}_2 = \text{MSE}_1$ , discussed in section 4 in the main text.

C.2.2. Unregularized regime  $\ell > \alpha$

If  $\ell > \alpha$ , then  $\gamma = 1$  and

$$q \left( 1 - C_z \int_0^\infty \frac{1}{(1 + C_z x^\alpha)^2} \right) = \frac{2}{\pi} - \frac{4 - \pi}{\pi} C_z \int_0^\infty \frac{1}{(1 + C_z x^\alpha)^2} - \mathcal{O} \left( n^{-2\alpha \min(r, \frac{1}{2})} \right). \tag{C.11}$$

Hence

$$\eta = \frac{m^2}{\rho q} \xrightarrow{n \rightarrow \infty} \frac{1 - C_z \int_0^\infty \frac{1}{(1 + C_z x^\alpha)^2}}{1 - \frac{4 - \pi}{2} C_z \int_0^\infty \frac{1}{(1 + C_z x^\alpha)^2}} \neq 1. \tag{C.12}$$

Since  $\eta$  is now bounded away from 1, the misclassification error  $\epsilon_g$  fails to go to 0 asymptotically, and plateaus to a finite value:

$$\epsilon_g = \mathcal{O}(1), \tag{C.13}$$

which is equation (19). This plateau is attributable to ridge classifiers overfitting the binary labels, see appendix D for further discussion.

C.3. Rates for noisy ridge classification

In this subsection, we provide briefly the derivation of the optimal rate for noisy ridge classification, defined by (E.1), which is plotted in figure E2 of the main text. Introducing the noise (E.1) in the equations of [28] and mapping to our setting along identical lines as appendix A, the self-consistent equations read

$$\begin{cases} m = \hat{m} p \operatorname{tr} \frac{\Sigma \theta^* \theta^{*\top} \Sigma}{n\lambda + p\hat{V}\Sigma} \\ q = p \operatorname{tr} \frac{p\hat{m}^2 \Sigma \theta^* \theta^{*\top} \Sigma^2 + \hat{q} \Sigma^2}{(n\lambda + p\hat{V}\Sigma)^2} \\ V = \operatorname{tr} \frac{\Sigma}{n\lambda + p\hat{V}\Sigma} \end{cases}, \quad \begin{cases} \hat{V} = \frac{\frac{n}{p}}{1 + V} \\ \hat{q} = \frac{n}{p} \frac{1 + q - 2m \sqrt{\frac{2}{\pi(\rho + \sigma^2)}}}{(1 + V)^2} = \frac{\hat{V}^2}{\frac{n}{p}} \left( 1 + q - 2m \sqrt{\frac{2}{\pi(\rho + \sigma^2)}} \right) \\ \hat{m} = \sqrt{\frac{2}{\pi(\rho + \sigma^2)}} \frac{\frac{n}{p}}{1 + V} = \sqrt{\frac{2}{\pi(\rho + \sigma^2)}} \hat{V}. \end{cases} \tag{C.14}$$

This differs from (C.1) by a the replacement of  $\rho$  by  $\rho + \sigma^2$ , simply reflecting the fact that the teacher  $L^2(\mathcal{X})$  squared norm should now include the contribution of the additive noise.

In the *regularized regime*  $\ell < \alpha$ , the expression (C.7) for the cosine-similarity  $\eta$  becomes

$$\eta = \frac{m^2}{\rho q} = \frac{\frac{2}{\pi(\rho + \sigma^2)} \rho^2 \left[ 1 - \mathbb{1}_{r \geq \frac{1}{2}} \frac{C_z A_m^n}{\rho} n^{-\ell} + \mathcal{O}(n^{-2\ell r}) \right]^2}{\frac{2}{\pi(\rho + \sigma^2)} \rho^2 \left[ 1 - \mathbb{1}_{r \geq \frac{1}{2}} \frac{2C_z A_q^n}{\rho} n^{-\ell} - \mathcal{O}(n^{-\min(2\ell \min(r, 1), \frac{\alpha - \ell}{\alpha})}) \right]} = 1 - \mathcal{O} \left( n^{-\min(2\ell \min(r, 1), \frac{\alpha - \ell}{\alpha})} \right). \tag{C.15}$$

Therefore the excess error reads

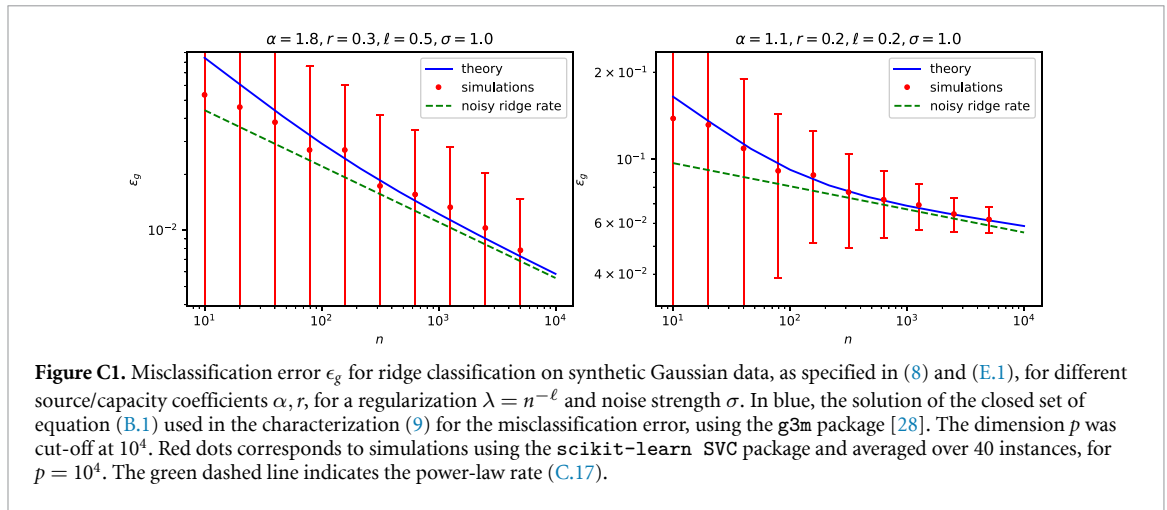
$$\begin{aligned} \epsilon_g - \epsilon_g^\infty &= \frac{1}{\pi} \cos^{-1} \left( \sqrt{\frac{\rho}{\rho + \sigma^2}} \eta \right) - \frac{1}{\pi} \cos^{-1} \left( \sqrt{\frac{\rho}{\rho + \sigma^2}} \right) \\ &= \begin{cases} \frac{1}{\sqrt{2\pi}} \left( \sqrt{\frac{\sigma^2}{\rho + \sigma^2} + (1 - \eta)} - \sqrt{\frac{\sigma^2}{\rho + \sigma^2}} \right) & \text{if } \sigma^2 \ll 1 \\ \frac{1}{2\pi} \frac{\rho + \sigma^2}{\rho} (1 - \eta) & \text{if } \sigma^2 \sim 1 \end{cases} \end{aligned} \tag{C.16}$$

$$\sim 1 - \eta = \mathcal{O} \left( n^{-\min(2\ell \min(r, 1), \frac{\alpha - \ell}{\alpha})} \right). \tag{C.17}$$

Note that the noisy rate (C.17) corresponds to twice the noiseless rate (18), see figure C1.

In the *unregularized regime*  $\ell > \alpha$ , the discussion in the noiseless setting carries through and the excess error also asymptotically saturates to a non-zero limit

$$\epsilon_g - \epsilon_g^\infty = \mathcal{O}(1) \tag{C.18}$$



### Appendix D. Remark on the overfitting regime of ridge classification

In this appendix, we provide further discussion on the saturation of the error (9) to a finite value in the effectively unregularized regime of ridge classification (19), which defies the common intuition that for large enough sample complexity  $n$ , ridge classification should be able to generalize almost perfectly. In this section, we first provide an analytical reminder of why the error does go to zero for finite-dimensional feature spaces ( $p < \infty, n \gg 1$ ), before discussing what differs in the setup at hand ( $1 \ll n \ll p = \infty$ ). We argue that while in the first limit a double-descent phenomenon is observed, the second limit always corresponds to a learning regime located in the plateau/valley following the first descent. In the following, we consider directly the  $\lambda = 0$  case for the sake of simplicity.

#### D.1. The finite dimensional setup

We first sketch up an brief analysis for the limit of finite feature space dimension  $\infty > p \gg 1$  setting, with large sample complexity  $n \gg p$ . Naming  $X \in \mathbb{R}^{n \times p}$  (with the  $\mu$ th row of  $X$  being  $\psi(x^\mu)$ ) the matrix of the data in feature space and  $y = \text{sign}(X\theta^*) \in \mathbb{R}^n$  the corresponding vector of stacked labels, the ridge estimator can be compactly written as

$$\hat{w} = (X^T X)^{-1} X^T \text{sign}(X\theta^*). \tag{D.1}$$

In the  $n \rightarrow \infty$  limit

$$\frac{1}{n} X^T X = \Sigma + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right), \quad \frac{1}{n} X^T \text{sign}(X\theta^*) = \mathbb{E}_x^{\mathcal{N}(0, \Sigma)} [x \text{sign}(x \cdot \theta^*)] + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right). \tag{D.2}$$

Hence

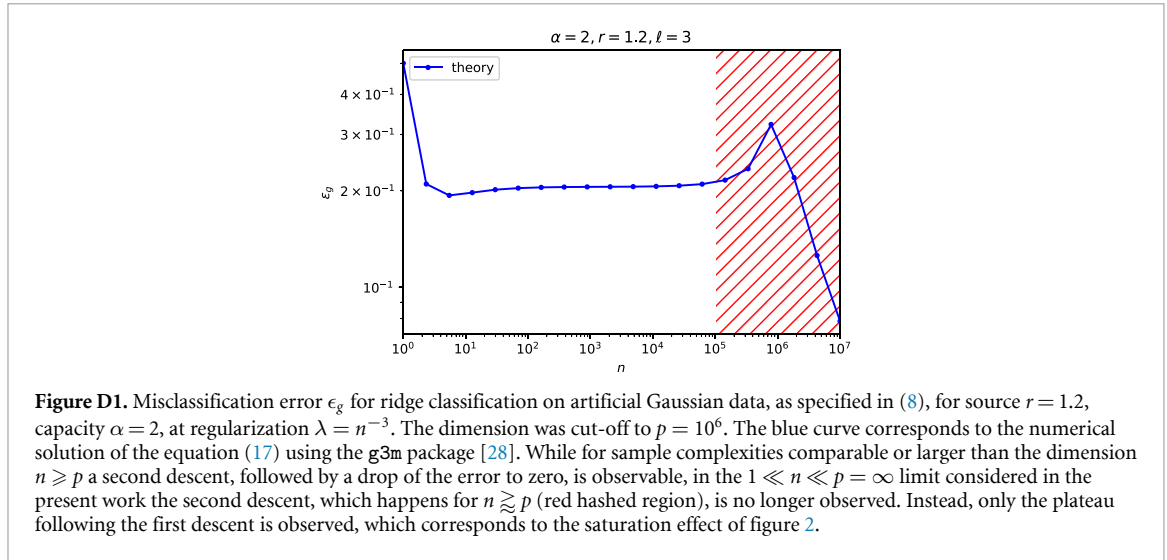
$$\begin{aligned} \hat{w} &\approx \Sigma^{-1} \mathbb{E}_x^{\mathcal{N}(0, \Sigma)} [x \text{sign}(\theta^{*T} x)] \\ &\approx \Sigma^{-1} \Sigma^{\frac{1}{2}} \mathbb{E}_x^{\mathcal{N}(0, 1)} [x \text{sign}(x^T \Sigma^{\frac{1}{2}} \theta^*)] \\ &\approx \Sigma^{-\frac{1}{2}} \mathbb{E}_{x_{\parallel}}^{\mathcal{N}(0, 1)} [x_{\parallel} \text{sign}(x_{\parallel})] \frac{\Sigma^{\frac{1}{2}} \theta^*}{\|\Sigma^{\frac{1}{2}} \theta^*\|_2}, \end{aligned} \tag{D.3}$$

where we noted  $x_{\parallel}$  the component of  $x$  parallel to  $\theta^*$ . So

$$\hat{w} = \sqrt{\frac{2}{\pi}} \frac{\theta^*}{\|\Sigma^{\frac{1}{2}} \theta^*\|_2} \propto \theta^*. \tag{D.4}$$

This means in particular the direction of the teacher is perfectly recovered, and that the prediction error  $\epsilon_g$  (9) goes to zero in the  $p \ll n \rightarrow \infty$  limit.





**Figure D1.** Misclassification error  $\epsilon_g$  for ridge classification on artificial Gaussian data, as specified in (8), for source  $r = 1.2$ , capacity  $\alpha = 2$ , at regularization  $\lambda = n^{-3}$ . The dimension was cut-off to  $p = 10^6$ . The blue curve corresponds to the numerical solution of the equation (17) using the `g3m` package [28]. While for sample complexities comparable or larger than the dimension  $n \geq p$  a second descent, followed by a drop of the error to zero, is observable, in the  $1 \ll n \ll p = \infty$  limit considered in the present work the second descent, which happens for  $n \gtrsim p$  (red hashed region), is no longer observed. Instead, only the plateau following the first descent is observed, which corresponds to the saturation effect of figure 2.

### D.2. Infinite dimensional feature space

We now argue why this discussion ceases to hold in the setup of interest  $n \ll p = \infty$  (plateau reached after the first descent, see figure D1). In this limit, loosely written,

$$\begin{aligned} \hat{w} &\approx \left( \Sigma^{-1} + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \right) \left( \mathbb{E}_x^{\mathcal{N}(0, \Sigma)} [\text{sign}(\theta^{*T} x)] + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \right) \\ &= \sqrt{\frac{2}{\pi}} \frac{\theta^*}{\|\Sigma^{\frac{1}{2}} \theta^*\|_2} + \mathcal{O}\left(\sqrt{\frac{p}{n}}\right) \neq \text{cte.} \times \theta^*. \end{aligned} \tag{D.5}$$

The last term comes from the sum of  $p$  random terms of order  $1/\sqrt{n}$  entailed by the matrix multiplication. In particular, this implies that the teacher fails to be perfectly recovered by ERM in this limit, causing the misclassification error to plateau to a finite limit, see equation (19). This is due to the fact that in an infinite-dimensional space, the ridge ERM (16) always has enough dimensions to overfit the dataset. Another way to look at this limit is in the framework of the well-known *double-descent* phenomenon which occurs in the finite dimensional limit  $p < \infty$  discussed in appendix D.1. In the  $p = \infty$  limit, the second descent, which commences at  $n \approx p$ , happens at infinite sample complexity and is thus not observed for finitely large  $n$ . In fact, the  $1 \ll n \ll p = \infty$  limit always correspond to the plateau following the first descent, see figure D1.

### Appendix E. Crossovers in noisy kernel classification

In this appendix, we briefly discuss the *noisy* setup where the target function  $f^*$  is corrupted by a Gaussian noise of variance  $\sigma^2$ , i.e.

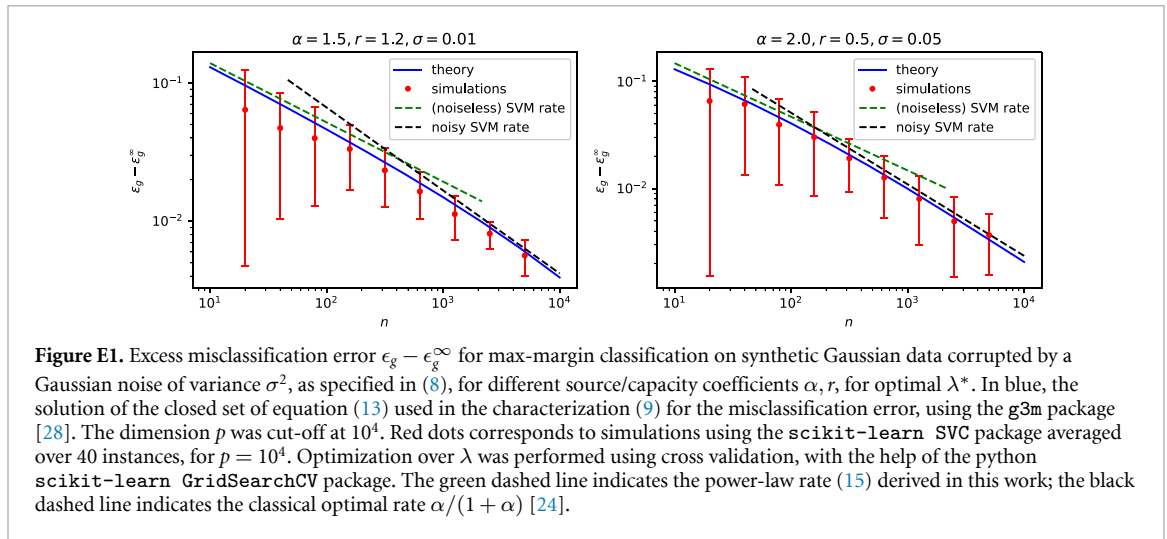
$$y^\mu = \text{sign}(f^*(x^\mu) + \sigma \mathcal{N}(0, 1)). \tag{E.1}$$

In this setting, the misclassification error (9) no longer tends asymptotically to 0, but instead converges to a positive value dictated by the noise strength

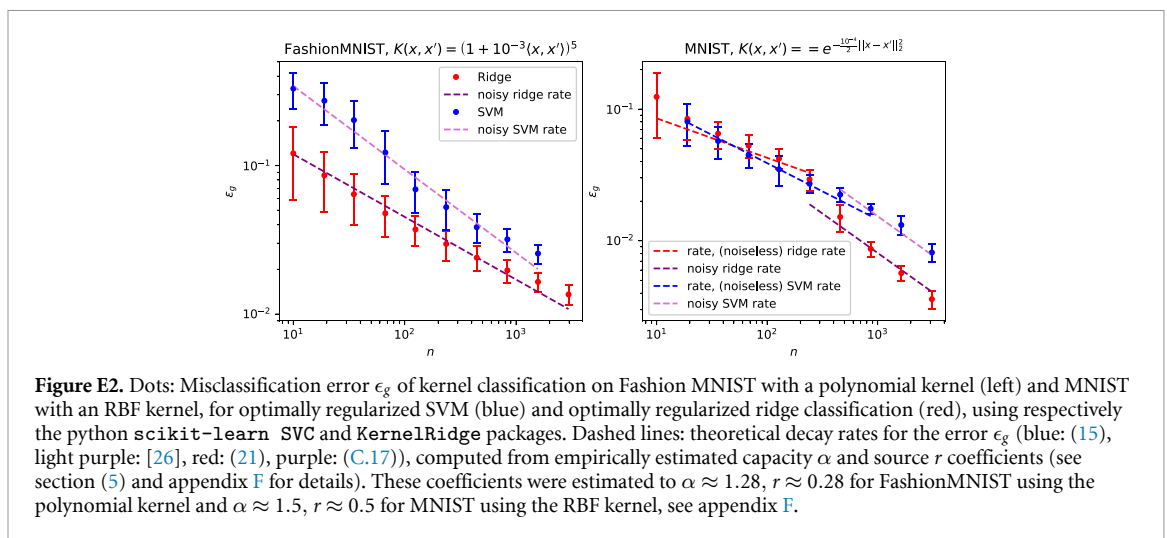
$$\epsilon_g \xrightarrow{n \rightarrow \infty} \epsilon_g^\infty = \frac{1}{\pi} \arccos\left(\sqrt{\frac{\rho}{\rho + \sigma^2}}\right). \tag{E.2}$$

Observe that the residual error  $\epsilon_g^\infty$  corresponds to the smallest achievable 0/1 risk, and is achieved by the teacher  $f^*$ . In contrast to the noiseless setting (1), the stochasticity in the generation of the labels (E.1) induces a non-zero error even for  $f^*$ . Numerical experiments presented in figure E1 show that the rate of decay of the excess misclassification error  $\epsilon_g - \epsilon_g^\infty$ , when optimized over the regularization strength  $\lambda$ , transitions from the noiseless rate (15) to a  $\alpha/(1 + \alpha)$  rate as the sample complexity  $n$  is increased. Note that interestingly the rate  $\alpha/(1 + \alpha)$  has been reported in [24] as an upper bound for the excess misclassification error, although it is possible to show that the noisy setup (E.1) does *not* satisfy the conditions of their corresponding theorem. This crossover phenomenon is highly similar to the one observed in [15] for kernel regression.

Finally, we note that for some real datasets the learning curves are satisfactorily described by the noisy rate  $\alpha/(1 + \alpha)$  [24] (see also the discussion in appendix A). This is for instance the case for Fashion MNIST



**Figure E1.** Excess misclassification error  $\epsilon_g - \epsilon_g^\infty$  for max-margin classification on synthetic Gaussian data corrupted by a Gaussian noise of variance  $\sigma^2$ , as specified in (8), for different source/capacity coefficients  $\alpha, r$ , for optimal  $\lambda^*$ . In blue, the solution of the closed set of equation (13) used in the characterization (9) for the misclassification error, using the `g3m` package [28]. The dimension  $p$  was cut-off at  $10^4$ . Red dots corresponds to simulations using the `scikit-learn` `SVC` package averaged over 40 instances, for  $p = 10^4$ . Optimization over  $\lambda$  was performed using cross validation, with the help of the python `scikit-learn` `GridSearchCV` package. The green dashed line indicates the power-law rate (15) derived in this work; the black dashed line indicates the classical optimal rate  $\alpha/(1 + \alpha)$  [24].



**Figure E2.** Dots: Misclassification error  $\epsilon_g$  of kernel classification on Fashion MNIST with a polynomial kernel (left) and MNIST with an RBF kernel, for optimally regularized SVM (blue) and optimally regularized ridge classification (red), using respectively the python `scikit-learn` `SVC` and `KernelRidge` packages. Dashed lines: theoretical decay rates for the error  $\epsilon_g$  (blue: (15), light purple: [26], red: (21), purple: (C.17)), computed from empirically estimated capacity  $\alpha$  and source  $r$  coefficients (see section (5) and appendix F for details). These coefficients were estimated to  $\alpha \approx 1.28, r \approx 0.28$  for FashionMNIST using the polynomial kernel and  $\alpha \approx 1.5, r \approx 0.5$  for MNIST using the RBF kernel, see appendix F.

when a polynomial kernel SVM is employed, see figure E2. Interestingly, the rate which describes the learning curve of an optimally regularized ridge classifier in this setting is also the *noisy* rate (C.17) (see appendix C), rather than the noiseless rate (21). For MNIST classified using an RBF kernel, both the noiseless rates (15) and (21) and the noisy rates  $\alpha/(1 + \alpha)$ , (C.17) are observed, with a crossover phenomenon during which the learning curve transitions from the former to the latter.

## Appendix F. Details on real data-sets

### F.1. Measuring the capacity and source of real datasets

In this section we provide details on the experiments on real data (section 5, figure 4), which can also be found in [15, 30, 31]. Consider a real dataset  $\mathcal{D}_m = \{x^\mu, y^\mu\}_{\mu=1}^m$  of size  $m$ . The data distribution  $\nu$  defined in section 2 is then the empirical distribution over  $\mathcal{D}_m$ :

$$\nu(\cdot) = \frac{1}{m} \sum_{\mu=1}^m \delta(\cdot - x^\mu). \tag{E.1}$$

Then the definitions of the feature map  $\psi$  (3) and of the feature covariance  $\Sigma$  (5) admit the simple rewriting

$$\frac{1}{m} G \Psi = \Psi \Sigma, \tag{E.2}$$

provided we introduce the Gram matrix  $G_{\mu, \nu} = K(x^\mu, x^\nu) \in \mathbb{R}^{m \times m}$  and the matrix of horizontally stacked features  $\Psi_{\mu k} = \psi(x^\mu)_k \in \mathbb{R}^{m \times m}$ . Note that for real data, one has that the dimension of the feature space  $p$  is equal to the total size of the dataset  $m$ . A teacher  $\theta^*$  providing perfect classification  $y = \text{sign}(\Psi \theta^*)$  can be found e.g. by performing max margin classification on  $(\Psi, y)$ . While the method to fit a perfect classifier  $\theta^*$  is

not unique, we observed that using logistic classifiers or regularized hinge classifiers instead of max-margin classifiers did not significantly impact the measured source coefficient.

From the eigenvalues  $\{\omega_k\}_{k=1}^p$  of the covariance  $\Sigma$  (E.2) and the components of the teacher  $\{\theta_k^*\}_{k=1}^p$ , the capacity and source  $\alpha, r$  (8) can be estimated. From (8), the following scalings should hold:

$$\sum_{k'=k}^p \omega_{k'} \sim k^{1-\alpha}, \quad \sum_{k'=k}^p \omega_{k'} \theta_{k'}^{*2} \sim k^{-2\alpha r}. \quad (\text{E.3})$$

These curves are presented in figure F1 for the dataset comprised of  $10^4$  randomly samples CIFAR ten images, see section 5 of the main text. Note that we estimate the cumulative functions (E.3) following [15, 30] because the summation allows to smoothen out the curves, thus permitting a relatively more precise evaluation. Because the feature space is finite, the functions (E.3) fail to be exact power-laws, and exhibit in particular a sharp drop for  $n$  approaching  $m$ . Nonetheless, we identify for each curve a region of indices  $k$  for which the curve looks qualitatively like a power-law, and fit the curves by a power-law using least-squares linear regression to finally extract the coefficients  $\alpha, r$  therefrom. These fits are shown in figure F1 for the reduced CIFAR ten dataset discussed in the main text (see section 5), for an RBF kernel with inverse variance  $10^{-7}$  and a polynomial kernel of degree 5. The capacity and source were estimated to be  $\alpha \approx 1.16, r \approx 0.10$  for the RBF kernel and  $\alpha \approx 1.51, r \approx 0.07$  for the polynomial kernel.

## F.2. Details on numerical simulations

In this subsection, we provide further details on the simulations presented in figure 4 of the main text. For each sample complexity  $n$ , a subset of size  $n$  was randomly sampled without replacements from  $\mathcal{D}_m$ . Like in [15, 28, 51], the whole dataset  $\mathcal{D}_m$  was used as a test set. The max-margin simulations in figure 4 were performed using the `scikit-learn` SVC package at vanishing regularization  $\lambda = 10^{-5}$ , and averaged over 50 realizations of the training set. The ridge simulations in figure 4 were realized using the `scikit-learn` `KernelRidge` package, with the optimal  $\lambda$  estimated using `scikit-learn` `GridSearchCV`'s default 5-fold cross validation routine over a grid  $\lambda \in \{0\} \cup (10^{-10}, 10^5)$ , with logarithmic step size 0.026. The misclassification error was also averaged over 50 realizations of the training set.

## Appendix G. Comparison to the classical rates

In this appendix, we provide a detailed comparison of the presently reported rate for noiseless max-margin classification (15) and the rate reported in [24] (equation 7.53, following from theorem 7.23 and Lemma A.1.7), in the case of a teacher in the Hilbert space  $f^* \in \mathcal{H}$ , for which the latter result holds. We alternatively point the reader to section 5 of [45], where the same bound, and the underlying assumptions thereof, are reminded in concise fashion. We first provide a brief reminder of this upper bound, before proceeding to evaluate it in the present setting (8), and show that it yields a slower rate than (15), i.e. that the bound in [24] is loose. For completeness, we finally provide a comparison to the rates reported in [27].

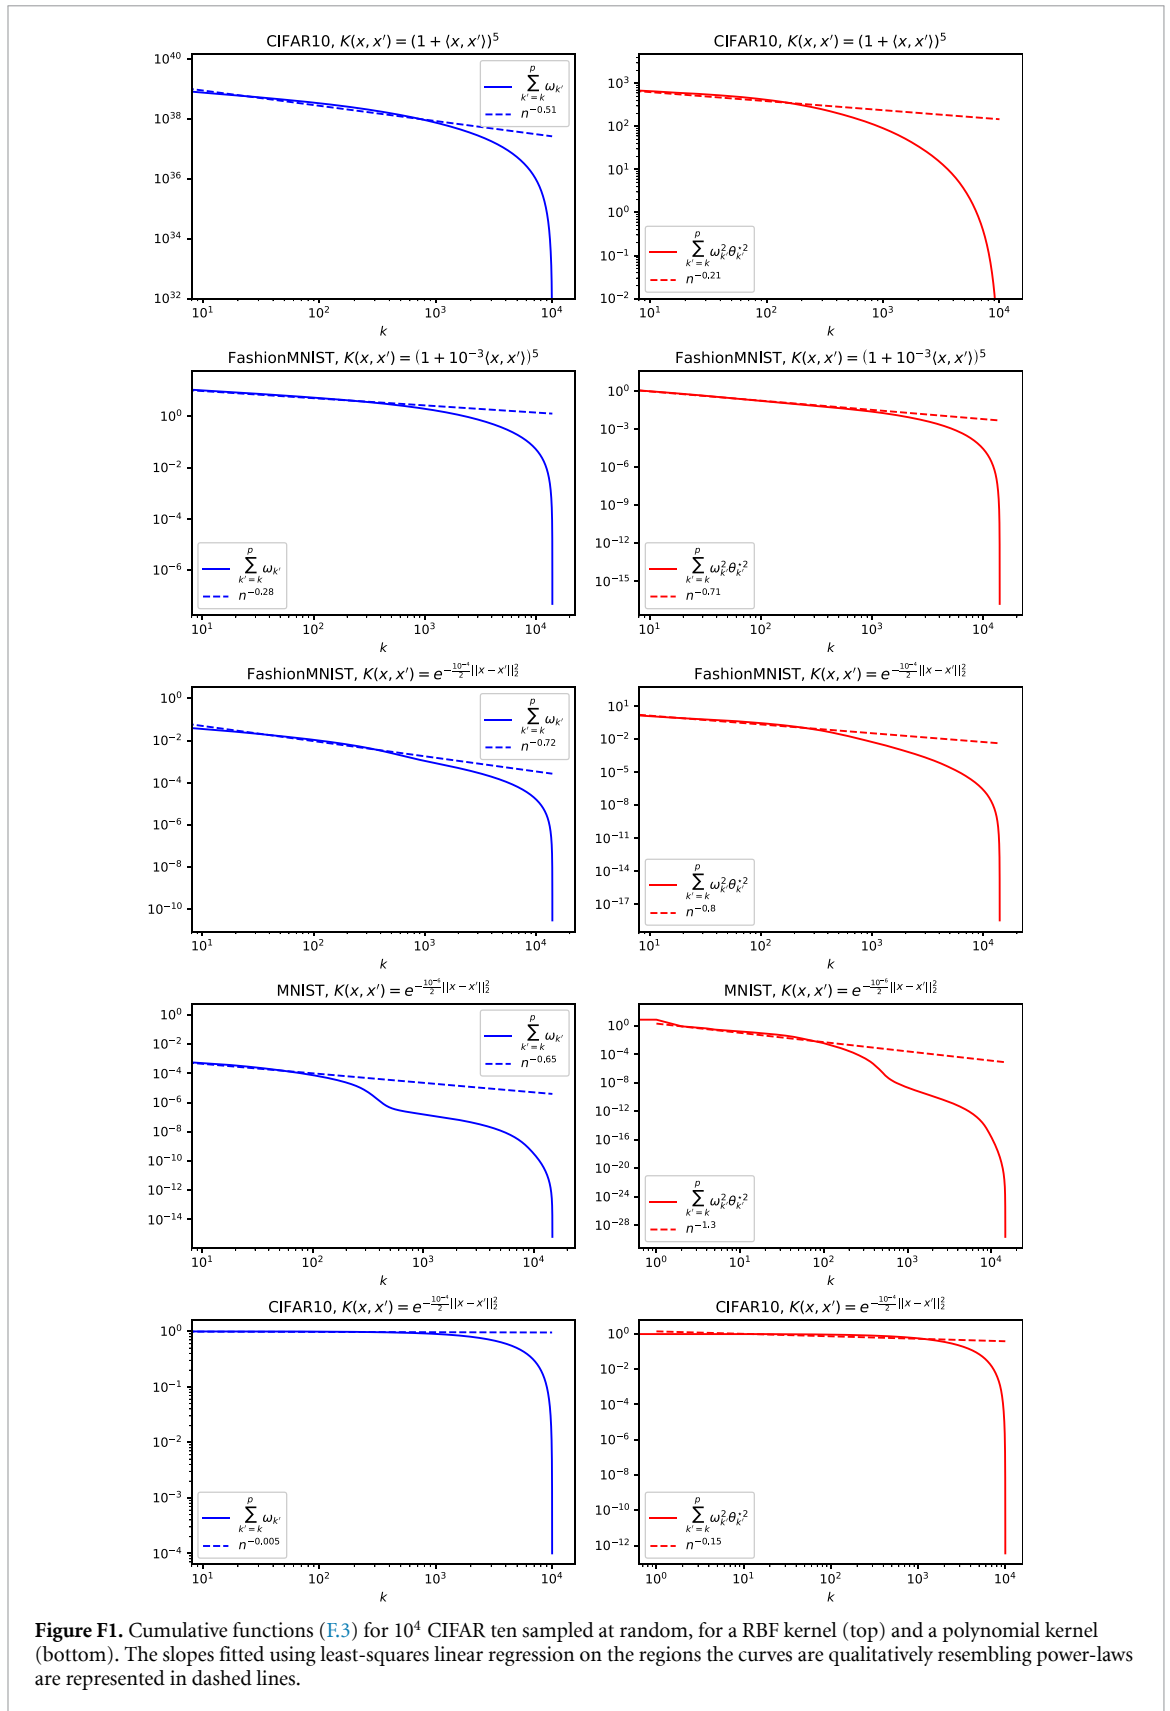
### G.1. Theorem 7.23 in [23]

In the following,  $\ell(y, a) = \max(0, 1 - ya)$  designates the hinge loss, and the script `cl` refers to the clipped value [24]. Suppose that

- A.1  $\inf_w \mathbb{E}_{\psi(x), y} \ell(y, w^\top \psi(x)) = \mathbb{E}_{\psi(x), y} \ell(y, f^*(\psi(x)))$  where the Bayes estimator  $f^*$  is defined as  $f^*(\psi(x)) = \underset{a}{\operatorname{argmin}} \mathbb{E}_{\psi(x), y} \ell(y, a)$ . The expectation refers to the joint distribution of the data  $\psi(x)$  and the labels  $y$ .
- A.2 (*Bernstein condition*) There exist  $B > 0, \theta \in [0, 1], V \geq B^{2-\theta}$  so that for all  $w, \ell(y, (w^\top \psi(x))^{\text{cl}}) \leq B$  a.s. and

$$\mathbb{E}_{\psi(x), y} \left( \left[ \ell(y, (w^\top \psi(x))^{\text{cl}}) - \ell(y, f^*(\psi(x))) \right]^2 \right) \leq V \left[ \mathbb{E}_{\psi(x), y} \left( \ell(y, (w^\top \psi(x))^{\text{cl}}) - \ell(y, f^*(\psi(x))) \right) \right]^\theta. \quad (\text{G.1})$$

- A.3 the spectrum of the covariance  $\Sigma = \mathbb{E}_{\psi(x)} \psi(x) \psi(x)^\top$  has a polynomial decay with rate  $\alpha$ .
- A.4 The approximation error  $\mathcal{A}_2(\lambda) = \min_w \mathbb{E}_{\psi(x), y} \ell(y, w^\top \psi(x)) + \lambda \|w\|^2 - \inf_w \mathbb{E}_{\psi(x), y} \ell(y, w^\top \psi(x))$  is upperbounded by  $\lambda^b$  for  $b \in (0, 1]$ .



**Figure F1.** Cumulative functions (F.3) for  $10^4$  CIFAR ten sampled at random, for a RBF kernel (top) and a polynomial kernel (bottom). The slopes fitted using least-squares linear regression on the regions the curves are qualitatively resembling power-laws are represented in dashed lines.

Then with high probability [24]

$$\epsilon_g - \epsilon_g^* \lesssim n^{-\min\left(\frac{2b}{1+b}, \frac{\alpha b}{b(2\alpha-1-\alpha\theta+\theta)+1}\right)}. \tag{G.2}$$

In (G.2),  $\epsilon^*$  refers to the infimum of the classification error over all measurable functions (as opposed to  $\epsilon_g$  resulting from ERM (12) over linear estimators only).

**G.2. Specialization to the noiseless source/capacity setting**

We proceed to translate theorem (G.2) [24] for our setting (8) (1), in the case where the teacher  $f^*$  is in  $\mathcal{H}$  ( $r \geq \frac{1}{2}$ ), for which assumption (A.1) is readily satisfied. To this end, one must evaluate the parameters  $\theta, b$  entering in (G.2). First note that the Bayes estimator  $f^*$  assumes the simple form  $f^*(x) = \text{sign}(\theta^{*\top} x)$  in the presence of the teacher  $\theta^*$ . As a consequence, in the noiseless setting,

$$\mathbb{E}_{\psi(x), y} \ell(y, f^*(\psi(x))) = \mathbb{E}_{\psi(x)} \ell(\text{sign}(\theta^{*\top} x), \text{sign}(\theta^{*\top} x)) = 0. \tag{G.3}$$

The Bernstein condition is straightforwardly satisfied with  $B = V = 2$  and  $\theta = 1$ , since  $\ell(1, \cdot^{\text{cl}})$  (resp.  $\ell(1, \cdot^{\text{cl}})$ ) is a piecewise continuous function, constant equal to 0 on  $(1, \infty)$  (resp.  $(-\infty, -1)$ ) and to 2 on  $(-\infty, -1)$  (resp.  $(1, \infty)$ ), and linear in between. In our model, the argmin in the definition of  $\mathcal{A}_2(\lambda)$  can be approximated first as  $w \approx c \times \theta^* / \|\theta^*\|$ , i.e. for a  $w$  proportional to the teacher. Note that this is an approximation, and in general there exist a small (shrinking with  $\lambda$  going to 0) angle between the true minimizer and the teacher. Typically,  $c$  will tend to infinity as  $\lambda$  goes to zero. Let us evaluate the proportionality constant  $c$ :

$$\begin{aligned} \mathbb{E}_{\psi(x), y} \ell(y, w^\top \psi(x)) &\propto \int dx e^{-\frac{1}{2} x^\top \Sigma^{-1} x} \max(0, 1 - \text{sign}(\theta^{*\top} x) c \theta^{*\top} x / \|\theta^*\|) \\ &\propto \int_{c \theta^{*\top} x / \|\theta^*\| < 1} dx e^{-\frac{1}{2} x^\top \Sigma^{-1} x} (1 - \text{sign}(\theta^{*\top} x) c \theta^{*\top} x / \|\theta^*\|) \\ &\propto \int_{-\frac{1}{c}}^{\frac{1}{c}} dx_{\parallel} e^{-\frac{x_{\parallel}^2}{2 \theta^{*\top} \Sigma \theta^* / \|\theta^*\|^2}} (1 - c |x_{\parallel}|) \\ &\propto 2 \int_0^1 dx_{\parallel} e^{-\frac{x_{\parallel}^2}{2 c \theta^{*\top} \Sigma \theta^* / \|\theta^*\|^2}} (1 - x) \sim \frac{1}{c} \end{aligned} \tag{G.4}$$

where in the last line we supposed  $c$  large. Then

$$\mathcal{A}(\lambda) \approx \min \left\{ \frac{A}{c} + \lambda c^2 \right\} \tag{G.5}$$

for some constant  $A$ , for  $\lambda$  small ( $c$  large). Minimizing the function leads to  $c \sim \lambda^{-\frac{1}{3}}$ , then to  $\mathcal{A}(\lambda) \sim \lambda^{\frac{1}{3}}$ , i.e.  $b = \frac{1}{3}$ . We also checked this scaling by numerically solving the self-consistent equation (13) in the large sample complexity limit ( $n \gg p$ ), noticing that the train loss admits the following expression in terms of the order parameters  $m, q, \rho, V, \hat{m}, \hat{q}, \hat{V}$  [28]:

$$\begin{aligned} \mathcal{A}_2(\lambda) &\approx \lim_{n \gg p} 2 \int_0^\infty \frac{e^{-\frac{u^2}{2}}}{\sqrt{2\pi}} \int_{-\infty}^\infty \frac{e^{-\frac{h^2}{2}}}{\sqrt{2\pi}} \max \left( 0, 1 - \text{prox}_{V \max(0, 1-\cdot)} \left( \frac{m}{\sqrt{\rho}} s + \sqrt{q(1-\eta)h} \right) \right) \\ &\quad + \lambda p \text{tr} \frac{p \hat{m}^2 \Sigma \theta^* \theta^{*\top} \Sigma + \hat{q} \Sigma}{\left( \frac{n\lambda}{2} + p \hat{V} \Sigma \right)^2}, \end{aligned} \tag{G.6}$$

where the proximal map is given by

$$\text{prox}_{V \max(0, 1-\cdot)}(\omega) = \begin{cases} \omega + \gamma V & \text{if } \gamma \omega \leq 1 - V, \\ \gamma & \text{if } 1 - V \leq \gamma \omega \leq 1, \\ \omega & \text{if } \gamma \omega \geq 1. \end{cases} \tag{G.7}$$

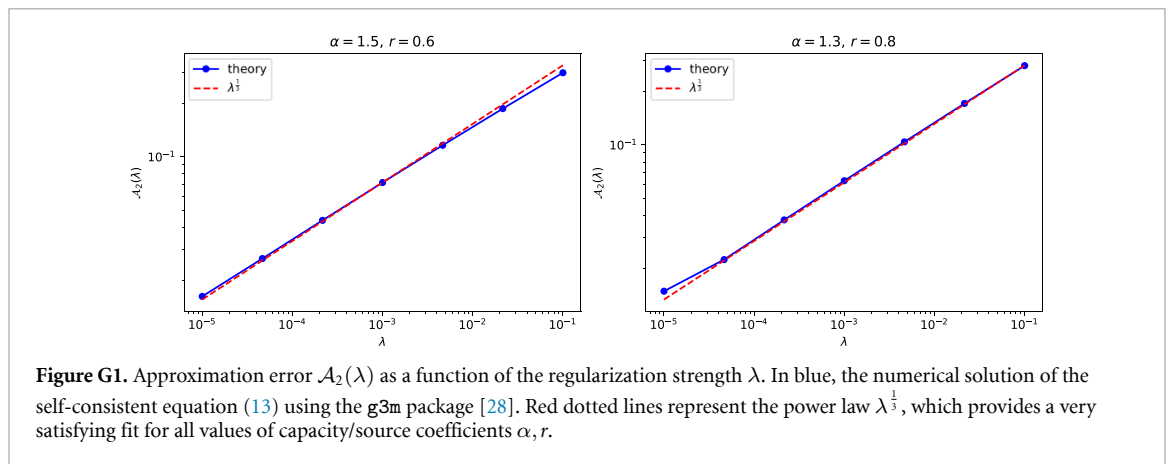
The function  $\mathcal{A}_2(\lambda)$  is plotted for several values of source/capacity coefficients  $\alpha, r$  in figure G1 and is well captured by a power law with rate  $b = \frac{1}{3}$ .

Finally, the bound provided by the theorem (G.2) [24] reads (notice that  $\epsilon_g^* = 0$ )

$$\epsilon_g - \epsilon_g^* \lesssim n^{-\min\left(\frac{1}{2}, \frac{\alpha}{\alpha+1}\right)} \tag{G.8}$$

Note that the rate (15) is always faster:

$$n^{-\frac{\alpha}{1+\frac{\alpha}{2}}} \ll n^{-\min\left(\frac{1}{2}, \frac{\alpha}{\alpha+1}\right)}. \tag{G.9}$$



**Figure G1.** Approximation error  $\mathcal{A}_2(\lambda)$  as a function of the regularization strength  $\lambda$ . In blue, the numerical solution of the self-consistent equation (13) using the `g3m` package [28]. Red dotted lines represent the power law  $\lambda^{\frac{1}{3}}$ , which provides a very satisfying fit for all values of capacity/source coefficients  $\alpha, r$ .

In other words, the upper bound of [24] is loose in the present setting when  $f^* \in \mathcal{H}$ . While numerical investigations suggest this is also true for  $f^* \in L^2(\mathcal{X}) \setminus \mathcal{H}$ , we leave a more detailed comparison to [24] in this case to future work.

### G.3. Theorems 3.3 and 3.5 in [25]

Finally, we discuss the rates reported in [27] under margin assumptions as [26, 44]. Note that the regression function  $\eta(\psi(x)) \equiv \mathbb{P}(y = 1 | \psi(x))$  has in the noiseless setting (1) the compact form  $\eta(\psi(x)) = \text{sign}(\theta^{*\top} \psi(x))$  and is therefore *not* Hölder class for any  $\beta > 0$ , rendering Theorems 3.3 and 3.5 in [27] effectively inapplicable. Note that in addition, the rates would in any case be ambiguous, as both the margin exponent  $\alpha$  and the dimension  $d$  are infinite in our setting. Finally remark that [26, 52] provide rates under similar margin conditions in *direct space*, for the particular case of Gaussian kernels. Relating those conditions to the characterization (8) in *feature space* used in the present work is out of the scope of the present manuscript.

## ORCID iD

Hugo Cui  <https://orcid.org/0000-0003-4648-244X>

## References

- [1] Hestness J, Narang S, Ardalani N, Diamos G, Jun H, Kianinejad H, Patwary M M A, Yang Y and Zhou Y 2017 arXiv:1712.00409
- [2] Kaplan J, McCandlish S, Henighan T, Brown T B, Chess B, Child R, Gray S, Radford A, Wu J and Amodei D 2020 arXiv:2001.08361
- [3] Rosenfeld J S, Rosenfeld A, Belinkov Y and Shavit N 2019 arXiv:1909.12673
- [4] Henighan T et al 2020 arXiv:2010.14701
- [5] Sharma U and Kaplan J 2022 *J. Mach. Learn. Res.* **23** 9–1
- [6] Bahri Y, Dyer E, Kaplan J, Lee J and Sharma U 2021 arXiv:2102.06701
- [7] Caponnetto A and De Vito E 2007 *Found. Comput. Math.* **7** 331–68
- [8] Caponnetto A and Vito E D 2005 Fast rates for regularized least-squares algorithm (available at: <https://dspace.mit.edu/handle/1721.1/30539>)
- [9] Lin J, Rudi A, Rosasco L and Cevher V 2018 *Appl. Comput. Harmon. Anal.* **48** 868–90
- [10] Jun K S, Cutkosky A and Orabona F 2019 Kernel truncated randomized ridge regression: optimal rates and low noise acceleration *Conf. on Neural Information Processing Systems*
- [11] Liu F, Liao Z and Suykens J A 2020 arXiv:2010.02681
- [12] Pillaud-Vivien L, Rudi A and Bach F 2018 Statistical optimality of stochastic gradient descent on hard learning problems through multiple passes *Advances in Neural Information Processing Systems* vol 31 pp 8114–24
- [13] Berthier R, Bach F and Gaillard P 2020 arXiv:2006.08212
- [14] Varre A, Pillaud-Vivien L and Flammarion N 2021 arXiv:2102.03183
- [15] Cui H, Loureiro B, Krzakala F and Zdeborová L 2021 arXiv:2105.15004
- [16] Maloney A, Roberts D A and Sully J 2022 arXiv:2210.16859
- [17] Hutter M 2021 arXiv:2102.04074
- [18] Marteau-Ferey U, Ostrovskii D, Bach F R and Rudi A 2019 Beyond least-squares: fast rates for regularized empirical risk minimization through self-concordance *COLT* (arXiv:1902.03046)
- [19] Steinwart I, Hush D R and Scovel C 2009 Optimal rates for regularized least squares regression *The 22nd Conf. on Learning Theory (COLT 2009) (Montreal, Quebec, Canada, 18–21 June 2009)* pp 79–93
- [20] Bartlett P, Long P M, Lugosi G and Tsigler A 2020 *Proc. Natl Acad. Sci.* **117** 30063–70
- [21] Dobriban E and Wager S 2018 *Ann. Stat.* **46** 247–79
- [22] Dicker L H et al 2016 *Bernoulli* **22** 1–37
- [23] Hsu D, Kakade S M and Zhang T 2012 Random design analysis of ridge regression *Proc. 25th Annual Conf. on Learning Theory (Proc. Machine Learning Research vol 23)* ed S Mannor, N Srebro and R C Williamson (JMLR Workshop and Conference Proceedings) p 9.1–24

- [24] Steinwart I and Christmann A 2008 *Support Vector Machines (Information Science and Statistics)* (Springer)
- [25] Schölkopf B and Smola A J 2002 *Learning With Kernels: Support Vector Machines, Regularization, Optimization and Beyond* (MIT Press)
- [26] Steinwart I and Scovel C 2007 *Ann. Stat.* **35** 575–607
- [27] Audibert J Y and Tsybakov A B 2007 *Ann. Stat.* **35** 608–33
- [28] Loureiro B, Gerbelot C, Cui H, Goldt S, Krzakala F, Mézard M and Zdeborová L 2021 arXiv:2102.08127
- [29] Jin H, Banerjee P K and Montúfar G 2021 Learning curves for gaussian process regression with power-law priors and targets (arXiv:2110.12231)
- [30] Spigler S, Geiger M and Wyart M 2020 *J. Stat. Mech.* **2020** 124001
- [31] Bordelon B, Canatar A and Pehlevan C 2020 Spectrum dependent learning curves in kernel regression and wide neural networks *Int. Conf. on Machine Learning* (PMLR) pp 1024–34
- [32] LeCun Y, Bottou L, Bengio Y and Haffner P 1998 *Proc. IEEE* **86** 2278–324
- [33] Xiao H, Rasul K and Vollgraf R 2017 arXiv:1708.07747
- [34] Mézard M, Parisi G and Virasoro M 1987 *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* vol 9 (World Scientific Publishing Company)
- [35] Dietrich R, Opper M and Sompolinsky H 1999 *Phys. Rev. Lett.* **82** 2975–8
- [36] Engel A and Van den Broeck C 2001 *Statistical Mechanics of Learning* (Cambridge University Press)
- [37] Mézard M and Montanari A 2009 *Information, Physics and Computation* (Oxford University Press)
- [38] Advani M S, Saxe A M and Sompolinsky H 2020 *Neural Netw.* **132** 428–46
- [39] Li Q and Sompolinsky H 2021 *Phys. Rev. X* **11** 031059
- [40] Ariosto S, Pacelli R, Pastore M, Ginelli F, Gherardi M and Rotondo P 2022 arXiv:2209.04882
- [41] Seroussi I, Naveh G and Ringel Z 2023 *Nat. Commun.* **14** 908
- [42] Cui H, Krzakala F and Zdeborová L 2023 arXiv:2302.00375
- [43] Pedregosa F et al 2011 *J. Mach. Learn. Res.* **12** 2825–30
- [44] Boucheron S, Bousquet O and Lugosi G 2005 *ESAIM: Probab. Stat.* **9** 323–75
- [45] Vecchia A D, Mourtada J, de Vito E and Rosasco L 2021 arXiv:2006.10016
- [46] Muthukumar V, Narang A, Subramanian V, Belkin M, Hsu D and Sahai A 2021 *J. Mach. Learn. Res.* **22** 1–69
- [47] Hsu D, Muthukumar V and Xu J 2020 On the proliferation of support vectors in high dimensions (arXiv:2009.10670)
- [48] Ardehshir N, Sanford C and Hsu D 2021 arXiv:2105.14084
- [49] Krizhevsky A et al 2009 Learning multiple layers of features from tiny images (available at: [www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf](http://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf))
- [50] Aubin B, Krzakala F, Lu Y M and Zdeborová L 2020 Generalization error in high-dimensional perceptrons: approaching bayes error with convex optimization *Advances in Neural Information Processing Systems* vol 33
- [51] Canatar A, Bordelon B and Pehlevan C 2021 *Nat. Commun.* **12** 1–12
- [52] Paccolat J, Spigler S and Wyart M 2021 *Mach. Learn.: Sci. Technol.* **2** 025020