

PAPER • OPEN ACCESS

How isotropic kernels perform on simple invariants

To cite this article: Jonas Paccolat *et al* 2021 *Mach. Learn.: Sci. Technol.* **2** 025020

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

You may also like

- [Data-driven discovery of Koopman eigenfunctions for control](#)
Eurika Kaiser, J Nathan Kutz and Steven L Brunton
- [Coarse-grain cluster analysis of tensors with application to climate biome identification](#)
Derek DeSantis, Phillip J Wolfram, Katrina Bennett et al.
- [Improving sample and feature selection with principal covariates regression](#)
Rose K Cersonsky, Benjamin A Helfrecht, Edgar A Engel et al.



PAPER

How isotropic kernels perform on simple invariants

OPEN ACCESS

Jonas Paccolat , Stefano Spigler and Matthieu Wyart

Institute of Physics, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

E-mail: jonas.paccolat@outlook.com

RECEIVED

16 November 2020

ACCEPTED FOR PUBLICATION

17 December 2020

PUBLISHED

2 March 2021

Keywords: machine learning, kernel methods, statistical physics, support vector machines

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

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

**Abstract**

We investigate how the training curve of isotropic kernel methods depends on the symmetry of the task to be learned, in several settings. (i) We consider a regression task, where the target function is a Gaussian random field that depends only on d_{\parallel} variables, fewer than the input dimension d . We compute the expected test error ε that follows $\varepsilon \sim p^{-\beta}$ where p is the size of the training set. We find that $\beta \sim 1/d$ independently of d_{\parallel} , supporting previous findings that the presence of invariants does not resolve the curse of dimensionality for kernel regression. (ii) Next we consider support-vector binary classification and introduce the *stripe model*, where the data label depends on a single coordinate $y(\underline{x}) = y(x_1)$, corresponding to parallel decision boundaries separating labels of different signs, and consider that there is no margin at these interfaces. We argue and confirm numerically that, for large bandwidth, $\beta = \frac{d-1+\xi}{3d-3+\xi}$, where $\xi \in (0, 2)$ is the exponent characterizing the singularity of the kernel at the origin. This estimation improves classical bounds obtainable from Rademacher complexity. In this setting there is no curse of dimensionality since $\beta \rightarrow 1/3$ as $d \rightarrow \infty$. (iii) We confirm these findings for the *spherical model*, for which $y(\underline{x}) = y(\|\underline{x}\|)$. (iv) In the stripe model, we show that, if the data are compressed along their invariants by some factor λ (an operation believed to take place in deep networks), the test error is reduced by a factor $\lambda^{-\frac{2(d-1)}{3d-3+\xi}}$.

1. Introduction and related works

Deep neural networks are successful at a variety of tasks, yet understanding why they work remains a challenge. In particular, we do not know *a priori* how many data are required to learn a given rule—not even the order of magnitude. Specifically, let us denote by p the number of examples in the training set. After learning, performance is quantified by the test error $\varepsilon(p)$. Quite remarkably, empirically one observes that $\varepsilon(p)$ is often well fitted by a power-law decay $\varepsilon \sim p^{-\beta}$. The exponent β is found to depend on the task, on the dataset and on the learning algorithm [1, 2]. General arguments would suggest that β should be extremely small—and learning thus essentially impossible—when the dimension D of the data is large, which is generally the case in practice (e.g. in images where D is the number of pixels multiplied by the number of color channels). For example in a regression task, if the only assumption on the target function is that it is Lipschitz continuous, then the test error cannot be guaranteed to decay faster than with an exponent $\beta \sim 1/D$ [3]. This *curse of dimensionality* [4] stems from the geometrical fact that the distance δ among nearest-neighbor data points decays extremely slowly in large d as $\delta \sim p^{1/D}$, so any interpolation method is very imprecise. The mere observation that deep learning works in large dimension implies that data are very structured [5]. Yet how to describe mathematically this structure and to build a quantitative theory for β remains a challenge. Our present goal is to study the relationship between β and symmetries in the data in simple models.

Recently there has been a considerable interest in studying the infinite-width limit of neural networks, motivated by the observation that performance generally improves with the number of parameters [6–11]. That limit depends on how the weights at initialization scale with the width. For a specific choice, similar to the LeCun initialization often used in practice, deep learning becomes equivalent to a kernel method [12], which has been coined the *neural tangent kernel*. In kernel methods, the learned function $Z(\underline{x})$ is a linear combination of the functions $K(\underline{x}, \underline{x}_{\mu})$, where \underline{x}_{μ} are the training data and K is the kernel. These methods

achieve performances somewhat inferior but still comparable to the more refined deep networks [13, 14], and are often used for both regression and classification. In this work we study the learning curves of isotropic kernels for which $K(\underline{x}, \underline{y}) = K(\|\underline{x} - \underline{y}\|)$, that include the popular Gaussian and Laplace kernels.

When these kernels are used on the image datasets MNIST and CIFAR-10, one finds that the learning curves decay with respective exponents $\beta_{\text{MNIST}} \approx 0.4$ and $\beta_{\text{CIFAR-10}} \approx 0.1$ that are much larger than $1/D_{\text{MNIST}} \approx 10^{-3}$ and $1/D_{\text{CIFAR10}} \approx 3 \times 10^{-4}$ [2]. Several aspects of the data could together explain these findings that β is much larger than $1/D$.

(a) In the kernel literature, upper bounds on the test error with β independent of D are obtained assuming that the target function lies in the *reproducing-kernel Hilbert space* of the kernel¹. However for these kernels this assumption is rather extreme: it supposes that the number of derivatives of the target function that are smooth is proportional to the dimension itself [4, 15]; see [2] for a precise statement for Gaussian random functions.

(b) The data live on a manifold \mathfrak{M} of lower dimensionality $d \leq D$. This is indeed the case for MNIST, where $D \approx 15$ [16–20], and CIFAR-10, where $D \approx 35$ [2]. This effect is presumably important, yet by itself it may not be the resolution of the problem, since the exponents β are significantly larger than $1/D$. Unless stated otherwise, in this work the data manifold extends to the whole space, namely $d = D$.

(c) The function to be learned presents many invariants. It can be expressed in terms of just $d_{\parallel} < d$ spatial components. For example in the context of classification, some pixels at the edge of the image may be unrelated to the class label. Likewise, smooth deformations of the image may leave the class unchanged. It has been argued that the presence of these invariants is central to the success of deep learning [5]. In that view, neural networks correspond to a succession of non-linear and linear operations where invariant directions are compressed [21]. It is supported by the observations that kernels designed to perform such compression perform well [5] and that compression can indeed occur at intermediate layers of deep networks [22]. Yet, relating these views quantitatively to the learning-curve exponent β remains a challenge, even for simple isotropic kernels and simple models of data. In [4], it was shown for a specific kernel in the context of regression that the presence of invariants did not improve guarantees for β . It is currently unclear if this results holds more generally to other kernels, beyond worst case analysis, and to classification tasks.

1.1. Our contribution

Our work consists of two parts that can be read independently, studying respectively regression and classification for different models.

The first part is presented in section 2 and focuses on kernel regression. We consider a target function that varies only along a linear manifold of d_{\parallel} directions of the input space, and is invariant along the remaining $d - d_{\parallel}$ directions. Without loss of generality, we consider that this dependence is on $\underline{x}_{\parallel} \equiv (x_1, \dots, x_{d_{\parallel}})^T$, corresponding to the d_{\parallel} first components of the data vectors $\underline{x} = (x_1, \dots, x_d)^T$. The target function is a Gaussian random function $Z_T(\underline{x}) \equiv Z_T(\underline{x}_{\parallel})$ with covariance determined by an isotropic translation-invariant teacher kernel $K_T(\underline{x})$. Kernel ridgeless regression is then performed using a distinct student kernel $K_S(\underline{x})$. Such a teacher–student framework (without invariants) was first introduced in [23, 24] and recently generalized in [25]. In these references it is investigated via an approximate formula based on averaging on the randomness of the data distribution. Here instead we use the methods of [2] inspired by earlier works on kriging [26] to compute the learning curve by calculating the expectation of the mean-squared test error, so as to extract the exponent β . We find and confirm numerically that β is independent of d_{\parallel} and depends only on d : one cannot escape the curse of dimensionality. This result supports the idea that even in a typical, non-worst case analysis, regression using simple kernels does not benefit from invariance in the data. Beyond the dependence on d , the exponent β is determined by the teacher and student kernels only through two exponents $\alpha_T(d)$, $\alpha_S(d)$ related to the decay of their Fourier transforms. In section 2, we define these exponents and we show that $\beta = \frac{1}{d} \min(\alpha_T(d) - d, 2\alpha_S(d))$.

In the second part of this work, we study kernel classification with support-vector machines, for which conclusions differ. We focus on simple models of data ($d_{\parallel} = 1$) that are arguably necessary first steps to build quantitative predictions for β in more realistic settings. In section 3, we introduce the *stripe model*, in which the class label $y(\underline{x}) = \pm 1$ only varies in one direction, as illustrated in figure 3. Again without loss of generality, we consider $y(\underline{x}) = y(x_1)$. This model corresponds to parallel interfaces separating regions where the label changes sign. We further consider the case without margin, where the data distribution $\rho(\underline{x})$ is non-zero at interfaces.

The performance of kernel classification depends on the bandwidth σ of the kernel, that is the scale over which it varies significantly. If σ is much smaller than the distance δ between training points, then the

¹ Such a Hilbert space is the set of all functions f with finite K -norm: $\|f\|_K < \infty$; see [43].

support-vector machine is tantamount to a nearest-neighbor algorithm, which inevitably suffers from the curse of dimensionality with an exponent $\beta \sim 1/d$. However in the limit of large σ , we provide scaling (heuristic) arguments that we systematically confirm numerically, showing that $\beta = \frac{d-1+\xi}{3d-3+\xi}$, where ξ is an exponent characterizing the singularity of the kernel at the origin (e.g. $\xi = 1$ for a Laplace kernel). This exponent β stays finite even in large dimension.

In section 4, we show that these results are not restricted to strictly flat interfaces: the same exponent β is found for the *spherical model* in which $y(\underline{x}) = y(\|\underline{x}\|)$. More generally, our analysis suggests that this result will break down if the boundary separating labels shows significant variation below a length scale $r_c \sim p^{-1/(d-1)}$. Avoiding the curse of dimensionality thus requires us to have an increasingly regular boundary separating labels as d increases.

Finally, in section 5, we come back to the stripe model and study how compressing the input data along its invariants (namely all the directions different from x_1) by a factor of λ improves performance—an effect believed to play a key role in the success of deep learning [5]. We argue and confirm empirically that, when mild, such a compression leaves the exponent β unchanged but reduces the test error by a factor of $\lambda^{-\frac{2(d-1)}{3d-3+\xi}}$.

1.2. Related works

1.2.1. Regression

The optimal worst-case performance of kernel regression has been investigated using a *source condition* that constrains the decay of the coefficients of the true function in the eigenbasis of the kernel [27–29]. For isotropic kernels and uniform data distribution, this condition is similar to controlling the decay of the Fourier components of the true function as we do here, and with our notation² the optimal worst-case generalization error is $\epsilon_{\text{wc}} \lesssim p^{-\beta_{\text{wc}}}$ with $\beta_{\text{wc}} = \frac{\alpha_T(d)-d}{\alpha_T}$ that is independent of the student. Yet in our approach we average the mean-squared error on all Gaussian fields with a given covariance, leading to a *typical* (instead of worst-case) exponent $\beta = \frac{1}{d} \min(\alpha_T(d) - d, 2\alpha_S(d))$. As expected, we always have $\beta > \beta_{\text{wc}}$: this follows from the fact that the exponents α_T, α_S must be larger than d for the kernels to be finite at the origin, a condition needed for our results to apply.

1.2.2. Classification

There is a long history of works computing the learning curve exponent β in regression or classification tasks where the true function or label depends on a single direction in input space, starting from the perceptron model [30] and including support vector classification [31]. More recently random feature models have received a lot of attention, and can be analytically resolved in some cases using random matrix or replica theories [9, 32–34]. Yet these results for classification generally consider linearly separable data³, and most importantly for both regression and classification tasks apply in the limit $d \rightarrow \infty$ and $p \rightarrow \infty$ with $\alpha = p/d$ fixed. In [31] for a single interface separating labels and kernels similar to ours, the learning curve of the support vector classifier was shown to decrease as $\epsilon \sim 1/\alpha$, as also found for the perceptron [35]. Here we consider both linearly and non-linearly separable data, and take the limit of large training set size p at fixed dimension d . This is in our view warranted considering data sets commonly used as benchmarks, such as MNIST or CIFAR for which $d_{\text{opt}} \in [15, 35]$ and $p \approx 6 \cdot 10^4$. In simple models for such numbers we do find that the training curves are well described by the limit we study. Specifically, the exponent β we find depends on dimension d and does not converge to 1 as $d \rightarrow \infty$, indicating that the two limits do not commute.

Classical works on kernel classification based on Rademacher complexity lead to lower bounds on $\beta \geq 1/4$ [36, 37] for certain algorithms applied to the stripe and spherical model⁴. Our estimation thus improves on that bound, even in the limit of large dimension where we find $\beta = 1/3$.

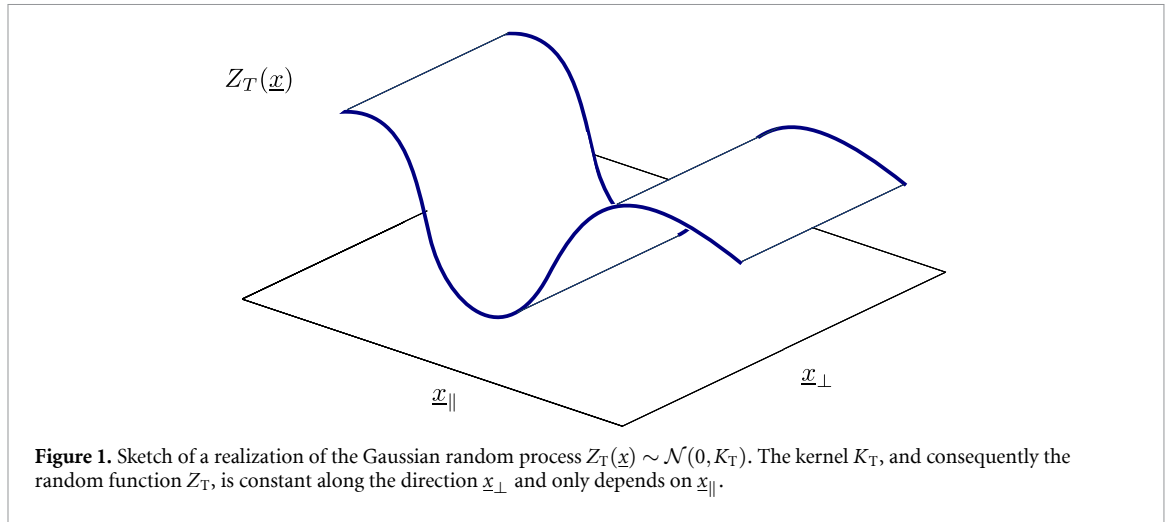
2. Kernel regression: teacher–student framework

We consider kernel ridgeless regression on Gaussian random data that present invariants. Our framework corresponds to a teacher–student setting for supervised learning [35, 38–42], where two variants of the same model (here kernels) are used both to generate the data and to learn them. The target function $Z_T(\underline{x})$ is

² Specifically, this literature introduces an exponent b characterizing the decay of the eigenvalues λ_ρ of the kernel with their rank ρ : $\lambda_\rho \sim \rho^{-b}$. In our setup it is straightforward to show that $b = \alpha_S/d$. Another exponent c (sometimes noted $2r$ [29]) characterizes the smoothness of the target function f^* . It is defined as the largest exponent for which $\langle f^* | K_S^{1-c} f^* \rangle < \infty$. It is straightforward to show that in our case, $c = \frac{\alpha_T - d}{\alpha_S}$. The worst case exponent is $\beta_{\text{wc}} = \frac{bc}{bc+1}$ [27–29] and is expressed in our notations in the main text.

³ See [31] for an example of non-linearly separable data lying on a hypercube.

⁴ For example for a single interface, theorem 21 of [36] bounding the test error can be applied with a linear function $f(\underline{x}) = x_1$ which has a finite RKHS norm. The bound on the test error then behaves as $P^{-1/4}$. An algorithm minimizing the expression for the bound on all functions on the RKHS ball of identical norm must thus lead to $\beta \geq 1/4$.



assumed to be a random Gaussian process $\mathcal{N}(0, K_T)$ with zero mean and covariance determined by a strictly positive-definite isotropic translation-invariant teacher kernel $K_T(\underline{x}, \underline{x}') = K_T(\|\underline{x} - \underline{x}'\|)$, implying that $\mathbb{E}_T Z_T(\underline{x}) = 0$ and $\mathbb{E}_T Z_T(\underline{x}) Z_T(\underline{x}') = K_T(\underline{x}, \underline{x}')$, where we denote by \mathbb{E}_T the expectation over the teacher Gaussian random process⁵. Strictly positive-definiteness is required to generate such a random function.

We further assume that the function $Z_T(\underline{x})$ does not depend on all the variables $\underline{x} = (x_1, \dots, x_d)^T$, but only on the first components $\underline{x}_\parallel \equiv (x_1, \dots, x_{d_\parallel})^T$ for some $d_\parallel \leq d$: $Z_T(\underline{x}) = Z_T(\underline{x}_\parallel)$, as sketched in figure 1. The Gaussian random process $Z_T(\underline{x})$ is constant along the subspace of $\underline{x}_\perp \equiv (x_{d_\parallel+1}, \dots, x_d)^T$ when it is generated by a Teacher kernel that has the same property, namely $K_T(\|\underline{x} - \underline{x}'\|) = K_T(\|\underline{x}_\parallel - \underline{x}'_\parallel\|)$. Indeed, we have that

$$\mathbb{E}_T \left[Z_T(\underline{x}_\parallel + \underline{x}_\perp) - Z_T(\underline{x}_\parallel) \right]^2 = 2K_T(0) - 2K_T(\underline{x}_\perp) = 0. \quad (1)$$

The (finite) training set is made up by the values of the target function $Z_T(\underline{x}^\mu)$ at p points $\{\underline{x}^\mu\}_{\mu=1}^p$. Kernel (ridgeless) regression is performed with a student kernel $K_S(\underline{x}, \underline{x}')$, that we also take to be isotropic and translation invariant and that can be different from the teacher kernel $K_T(\underline{x}, \underline{x}')$. The student has no prior knowledge of the presence of invariants: its kernel is a function of all the spatial components.

Kernel regression consists in writing the prediction for the function $\hat{Z}_S(\underline{x})$ at a generic point \underline{x} as a linear combination of student kernel overlaps on the whole training set, namely

$$\hat{Z}_S(\underline{x}) = \sum_{\mu} a^\mu K_S(\underline{x}^\mu, \underline{x}) \equiv \underline{a} \cdot \underline{k}_S(\underline{x}). \quad (2)$$

The vector of coefficients \underline{a} is determined by minimizing the mean-squared loss on the training set:

$$\underline{a} = \arg \min_{\underline{a}} \sum_{\mu} \left[\hat{Z}_S(\underline{x}^\mu) - Z_T(\underline{x}^\mu) \right]^2. \quad (3)$$

The minimization of such a quadratic loss can be carried out explicitly, and the student prediction can be written as

$$\hat{Z}_S(\underline{x}) = \underline{k}_S(\underline{x}) \cdot \mathbb{K}_S^{-1} \underline{Z}_T, \quad (4)$$

where the vector $\underline{Z}_T \equiv (Z_T(\underline{x}^\mu))_{\mu=1}^p$ contains all the samples in the training set and $\mathbb{K}_S^{\mu\nu} \equiv K_S(\underline{x}^\mu, \underline{x}^\nu)$ is the Gram matrix. By definition, the Gram matrix is always invertible for any training set if the kernel K_S is strictly positive definite. The generalization error is then evaluated as the expected mean-squared error on out-of-sample data that were not used for training: numerically, it is estimated by averaging over a test set composed of p_{test} newly sampled data points:

$$\epsilon_T = \mathbb{E}_{\underline{x}} \left[\hat{Z}_S(\underline{x}) - Z_T(\underline{x}) \right]^2 = \frac{1}{p_{\text{test}}} \sum_{\mu=1}^{p_{\text{test}}} \left[\hat{Z}_S(\underline{x}^\mu) - Z_T(\underline{x}^\mu) \right]^2. \quad (5)$$

⁵ With respect to the kernel literature, note that in our setting Z_T never belongs to the RKHS of K_T , see e.g. [44]. The conditions for it to belong to K_S are discussed in [2].

This quantity is a random variable, and we take the expectation also with respect to the teacher process to define an average test error $\epsilon = \mathbb{E}_{T \in \mathcal{T}}$ —in the numerical simulations that we discuss later, we simply average over several runs of the Teacher Gaussian process.

We study how the expected test error ϵ decays with the size p of the training set. Asymptotically for large p , this decay follows a power law $\epsilon \sim p^{-\beta}$. In [2], β was derived in the absence of invariants ($d_{\parallel} = d$), building on results from the kriging literature [26]. It was found that β depends on three quantities: the dimension d and two exponents $\alpha_T(d), \alpha_S(d)$ related to the two kernels. These exponents describe how the Fourier transform of the kernels decay at large frequencies: $\tilde{K}_T(\underline{w}) \sim \|\underline{w}\|^{-\alpha_T(d)}$, and similarly for the student K_S . Notice that, since the kernels are translation invariant, their Fourier transform is a function of only one frequency vector \underline{w} . Moreover, the exponents $\alpha_T(d), \alpha_S(d)$ depend on the dimension of the space where the Fourier transform is computed.

Our main theorem, formally presented with its proof in appendix A, is as follows:

Theorem 1 (Informal) *Let ϵ be the average mean-squared error of the regression made with a student kernel K_S on the data generated by a teacher kernel K_T , sampled at points taken on a regular d -dimensional square lattice in \mathbb{R}^d with fixed spacing δ . Assume that the teacher kernel only varies in a lower-dimensional space: $K_T(\underline{x}) = K_T(\underline{x}_{\parallel})$, with $\underline{x}_{\parallel} = (x_1, \dots, x_{d_{\parallel}})^T$ a vector in $d_{\parallel} \leq d$ dimensions. The student kernel in contrast varies along all d -dimensional directions in input space. Let the Fourier transforms of the two kernels decay at high frequency with dimension-dependent exponents $\alpha_T(d)$ and $\alpha_S(d)$. Then as $\delta \rightarrow 0$, $\epsilon \sim \delta^{\beta d}$ with*

$$\beta = \frac{1}{d} \min(\alpha_T(d_{\parallel}) - d_{\parallel}, 2\alpha_S(d)). \tag{6}$$

Note 1: We expect that under broad conditions the quantity $\alpha_T(d_{\parallel}) - d_{\parallel} \equiv \theta_T$ (as well as θ_S obviously) does not depend on d_{\parallel} , and that θ_T corresponds to the exponent characterizing the singular behavior of $K_T(\underline{x})$ at the origin:

$$K_T(\underline{x}) = C_0 |\underline{x}|^{\theta_T} + \text{regular terms} \tag{7}$$

as discussed in appendix A. This fact can be shown (see below) for Laplace (where $\theta_T = 1$) and Matérn kernels whose Fourier transform can be computed exactly. Thus we recover the curse of dimensionality since $\beta = \frac{1}{d} \min(\theta_T, 2d + 2\theta_S) \leq \theta_T/d$, which is independent of d_{\parallel} and thus of the presence of invariants.

Note 2: A remark is in order for the case of a Gaussian kernel $K(z) = \exp(-z^2)$, since it is a smooth function and its Fourier transform (being a Gaussian function too) decays faster than any power law at high frequencies. As discussed and verified in the aforementioned paper, this theorem applies also to Gaussian kernels, provided that the corresponding exponent is taken to be $\theta = \infty$. In particular, if the teacher is Gaussian and the student is not, $\beta = 2 + \frac{2\theta_S}{d}$; in the opposite scenario, where the teacher is not Gaussian but the student is, $\beta = \frac{\theta_T}{d}$; if both kernels are Gaussian, $\beta = \infty$ and the test error decays with respect to the training set size faster than a power law.

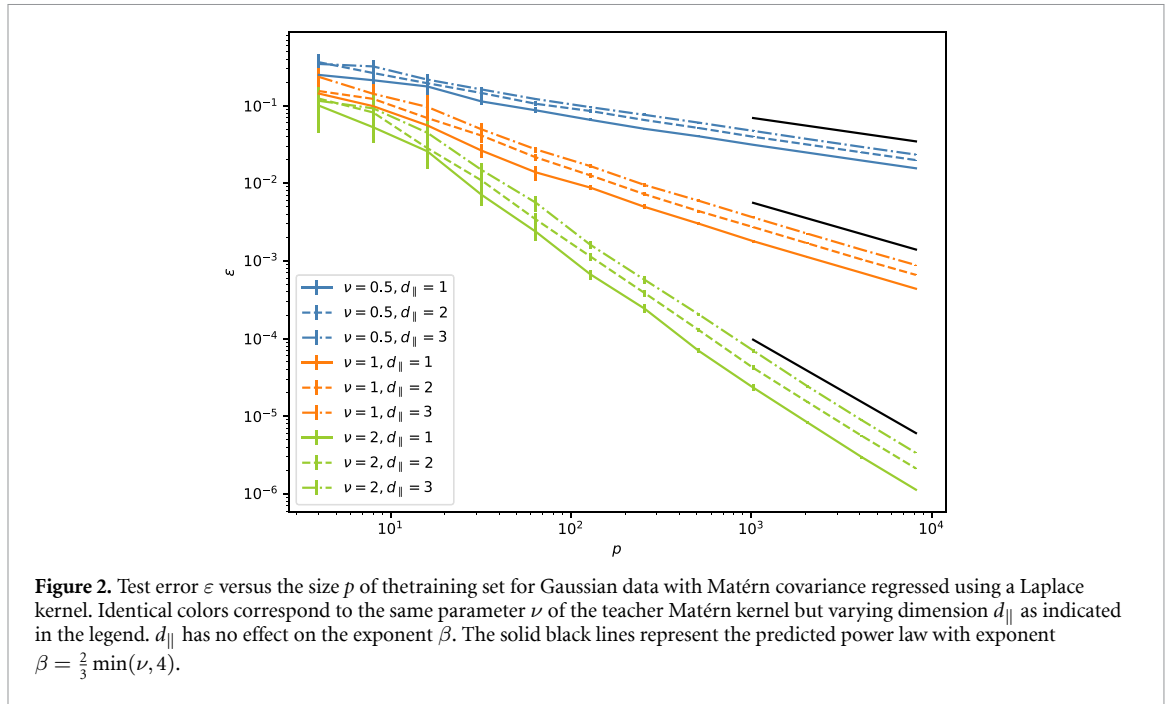
Interpretation: The following interpretation can be given for theorem 1 when α_S is large, leading to $\beta = \frac{\theta_T}{d}$. An isotropic kernel corresponds to a Gaussian prior on the Fourier coefficients of the true function being learned, a prior whose magnitude decreases with wave vectors as characterized by the exponent α_S . Clearly, the number of coefficients that can be correctly reconstructed cannot be larger than the number of observations p . For large α_S , we find that kernel regression does indeed reconstruct well a number of the order of p first Fourier coefficients, which corresponds to wave vectors \underline{w} of norm $\|\underline{w}\| \leq 1/\delta \sim p^{1/d}$. Fourier coefficients of larger wave vectors cannot be reconstructed, however, and the mean-squared error is then simply of order of the sum of the squares of these coefficients:

$$\epsilon \sim \int_{\|\underline{w}\| \geq p^{1/d}} d^{d_{\parallel}} \underline{w} \|\underline{w}\|^{-\alpha(d_{\parallel})} \sim p^{[d_{\parallel} - \alpha(d_{\parallel})]/d} \sim p^{-\theta_T/d}. \tag{8}$$

Numerical test:

We now test numerically that kernel regression is blind to the lower-dimensional nature of the task. We consider a $d = D - 1$ -dimensional sphere of unit radius \mathbb{S}^d embedded in \mathbb{R}^D . To test robustness with respect to our technical assumption of data points lying on an infinite lattice, we consider instead p i.i.d. points sampled uniformly at random. The component x_i^{μ} of each point is generated as a standard Gaussian $\mathcal{N}(0, 1)$ and then the vector \underline{x}^{μ} is normalized by dividing it by its norm. Points belonging to such a training set have a typical nearest-neighbor distance $\delta \sim p^{-1/D}$, and we will show that the test error decays with the predicted scaling $\epsilon \sim \delta^{\beta D} = p^{-\beta}$. For the numerical verification we take the student to be a Laplace kernel:

$$K_S(z) = \exp\left(-\frac{z}{\sigma}\right), \tag{9}$$



that is characterized by $\alpha(d) = d + \theta_S$ with $\theta_S = 1$. As teacher we use Matérn kernels, which are a family of kernels parametrized by one parameter ν :

$$K_{T,\nu}(z) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{z}{\sigma} \right)^\nu \mathcal{K}_\nu \left(\sqrt{2\nu} \frac{z}{\sigma} \right), \quad (10)$$

where $\mathcal{K}_\nu(z)$ is the modified Bessel function of the second kind with parameter ν , and Γ is the Gamma function. Varying ν , one can change the smoothness of the instances of the Gaussian random process, and in particular $\alpha_T(d) = d + \theta_T$ with $\theta_T = 2\nu$. The spatial dimension is $D = 4$ and we vary the number of invariants in the task by taking $d_{\parallel} = 1, 2, 3$. In order to fix d_{\parallel} we simply use $z = \|\underline{x}_{\parallel} - \underline{x}'_{\parallel}\|$ instead of $z = \|\underline{x} - \underline{x}'\|$ when computing the Teacher kernel. The scale of the kernel is fixed by the constant σ , that we have taken equal to 4 for both the teacher and the student. Notice that in theorem 1 the value of σ does not play any role since it does not enter the asymptotic behavior of the test error (at leading order). In figure 2 we show that the numerical simulations match our predictions. Indeed, in this specific case the predicted exponent is

$$\beta = \frac{2}{3} \min(\nu, 4). \quad (11)$$

Notice that the exponent that characterizes the learning curves is indeed independent of d_{\parallel} . Its prefactor may however depend on d_{\parallel} in general.

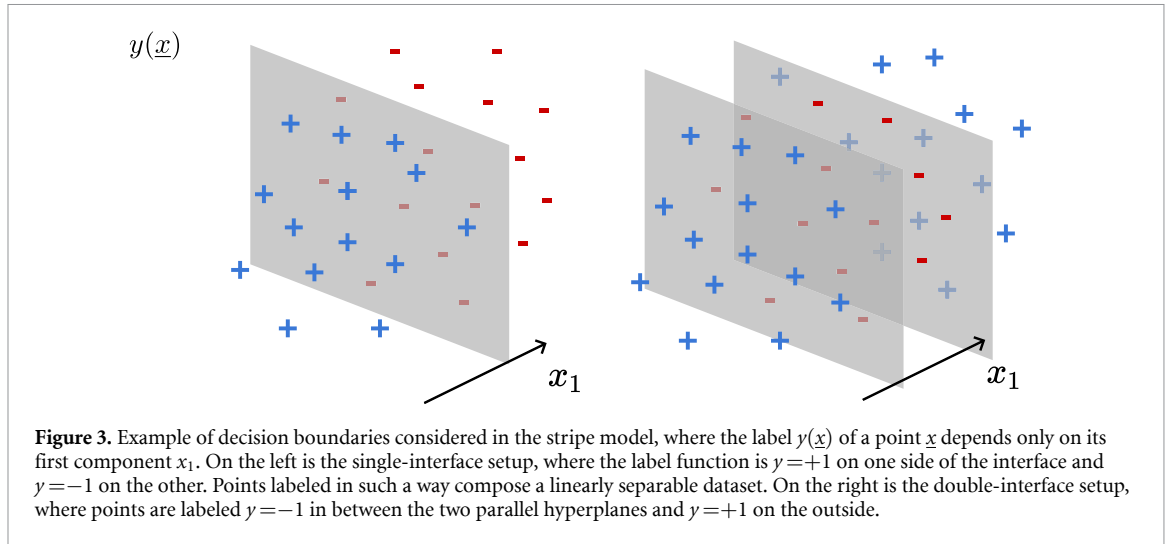
3. Support vector classification and stripe model

3.1. Stripe model

We consider a binary classification task where the labels depend only on one direction in the data space, namely with $y(\underline{x}) = y(x_1)$. Layers of $y = +1$ and $y = -1$ regions alternate along the direction x_1 , separated by parallel planes. Two examples of this setting are sketched in figure 3, corresponding to a single and a double interface. The points \underline{x} that constitute the training and test set are i.i.d. of distribution $\rho(\underline{x})$. To lighten the notation, we assume that $\rho(\underline{x})$ is uniform on a square box Ω of linear extension γ . Yet we expect our arguments to apply more generally if $\rho(\underline{x})$ is continuous and does not vanish at the location of the interfaces (no margin). To confirm this view we will test and confirm below our predictions when $\rho(\underline{x})$ is Gaussian distributed, with each component $x_i \sim \mathcal{N}(0, \gamma^2)$ with some variance γ^2 .

3.2. Definition of margin SVC

In this section we consider margin support-vector classification (margin SVC). This algorithm maximizes the margin between a decision boundary and the points in the training set that are closest to it. The prediction of



the label $\hat{y}(\underline{x})$ of a new point \underline{x} is then made according to the sign of the estimated decision function [43]:

$$f(\underline{x}) = \sum_{\mu=1}^p \alpha^\mu y^\mu K\left(\frac{\|\underline{x}^\mu - \underline{x}\|}{\sigma}\right) + b \rightarrow \hat{y}(\underline{x}) = \text{sign}f(\underline{x}), \tag{12}$$

where the kernel K is conditionally strictly positive definite [45]—a condition defined in appendix C, less stringent than strictly positive definite. In equation (12) we write explicitly the kernel bandwidth σ since it will soon play an important role. The formulation of the margin-SVC algorithm presented below is what is referred to as the *dual formulation*, but it can be equivalently recast as an attempt to maximize a (signed) distance between training points and the decision boundary [43]. In this dual formulation, the variables α^μ are fixed by maximizing

$$\max_{\underline{\alpha}} \mathcal{L}(\underline{\alpha}), \text{ with } \mathcal{L}(\underline{\alpha}) = \sum_{\mu=1}^p \alpha^\mu - \frac{1}{2} \sum_{\mu, \nu=1}^p \alpha^\mu \alpha^\nu y^\mu y^\nu K\left(\frac{\|\underline{x}^\mu - \underline{x}^\nu\|}{\sigma}\right), \tag{13}$$

subject to the constraints

$$\alpha^\mu \geq 0, \tag{14}$$

$$\alpha^\mu > 0 \text{ if and only if } y^\mu f(\underline{x}^\mu) = 1, \tag{15}$$

$$Q \equiv \sum_{\mu=1}^p \alpha^\mu y^\mu = 0 \text{ (charge conservation)}. \tag{16}$$

The bias b is set to satisfy

$$\min_{1 \leq \mu \leq p} |f(\underline{x}^\mu)| = 1 \text{ (canonical condition)}. \tag{17}$$

Equation (15) states that a dual variable α^μ is strictly positive if and only if its associated vector \underline{x}^μ lies on the margin, that is $y^\mu f(\underline{x}^\mu) = 1$, otherwise it is zero. Vectors with $\alpha^\mu > 0$ are called *support vectors* (SVs) and are the only ones that enter the expansion of the decision function equation (12).

3.3. Some limiting cases of SVC

Vanishing bandwidth: If the kernel function $K(z)$ decreases exponentially fast with some power of z , then in the limit $\sigma \ll \delta$, where δ is the average nearest-neighbor distance in the training set, the support-vector machine becomes akin to a nearest-neighbor algorithm. A detailed analysis of this regime for the stripe model is presented in appendix B; here we provide a qualitative argument assuming that the bias b is negligible. If so, as $\sigma \rightarrow 0$ one has for any training point that $f(\underline{x}^\mu) \approx \alpha^\mu y^\mu K(0)$, implying that $\alpha^\mu \neq 0$ to satisfy $|f(\underline{x}^\mu)| \geq 1$: every point is an SV with identical α^μ . $f(\underline{x})$ at a random test point \underline{x} is dominated by the

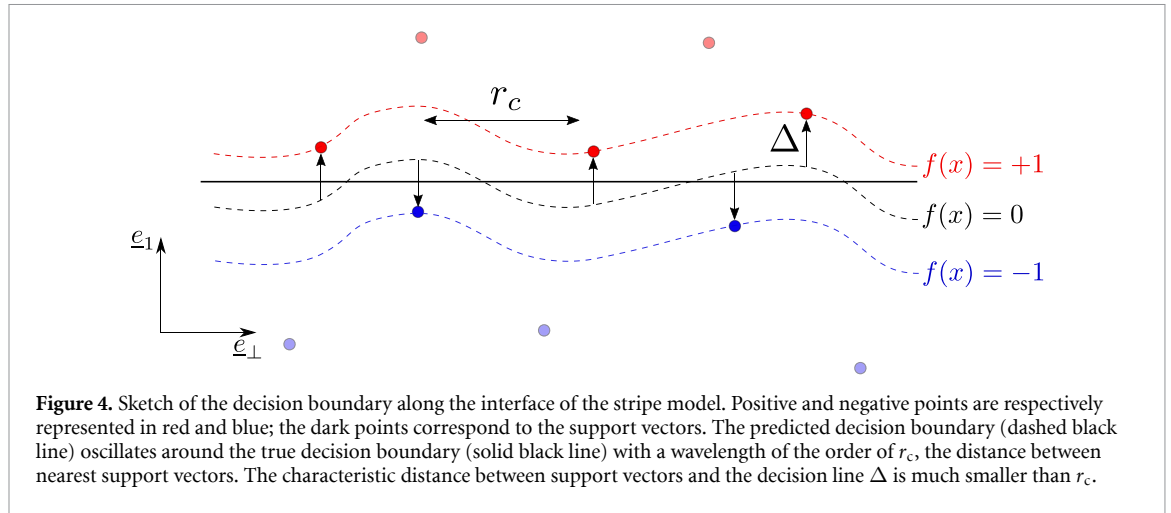


Figure 4. Sketch of the decision boundary along the interface of the stripe model. Positive and negative points are respectively represented in red and blue; the dark points correspond to the support vectors. The predicted decision boundary (dashed black line) oscillates around the true decision boundary (solid black line) with a wavelength of the order of r_c , the distance between nearest support vectors. The characteristic distance between support vectors and the decision line Δ is much smaller than r_c .

closest SV. The classification error is susceptible to the curse of dimensionality for such an algorithm, and one expects generically $\varepsilon \sim \delta \sim p^{-1/d}$, as tested numerically in figure B1 for the stripe model.

Diverging bandwidth: In this work we focus on the other extreme case where the bandwidth is larger than the system size, namely $\sigma \gg \gamma$. In this regime the kernel is always evaluated close to the origin. Assuming that the kernel has a finite derivative in the neighborhood of the origin, we approximate it by its truncated Taylor expansion:

$$K\left(\frac{\|\underline{x} - \underline{x}'\|}{\sigma}\right) \approx K(0) - \text{const.} \left(\frac{\|\underline{x} - \underline{x}'\|}{\sigma}\right)^\xi + o((\gamma/\sigma)^\xi). \tag{18}$$

The exponent ξ is related to the exponent θ introduced in section 2 by $\xi = \min(\theta, 2)$, and varies from kernel to kernel. For instance, we have $\xi = 1$ for Laplace kernels, $\xi = 2$ for Gaussian kernels, $\xi = \tilde{\gamma}$ for $\tilde{\gamma}$ -exponential kernels⁶ and $\xi = \min(2\nu, 2)$ for Matérn kernels. In appendix C we show that for $0 < \xi < 2$ the right-hand side is conditionally strictly positive definite (CSPD), which is the necessary condition for the SVC algorithm to converge. In what follows, we consider $0 < \xi < 2$, which excludes the Gaussian case. A proof that in that case the margin-SVC algorithm with the truncated kernel in equation (18) leads to the same solution as with the full kernel in the limit $\sigma \gg \gamma$ is presented in appendix D. Also, due to the charge conservation in equation (16), the constant term $K(0)$ in equation (18) may safely be ignored.

The decision function equation (12) associated with the considered radial power kernel hence becomes

$$f(\underline{x}) = b - \sum_{\mu=1}^p \alpha^\mu y^\mu \left(\frac{\|\underline{x} - \underline{x}^\mu\|}{\sigma}\right)^\xi. \tag{19}$$

where the positive constant in equation (18) has been removed by rescaling the bias and the α^μ .

3.4. Single interface

We consider a single interface at location $x_1 = 0$, with negative labels for $x_1 < 0$ and positive ones for $x_1 > 0$. Already in this case, computing analytically the test error remains a challenge, and we resort to a scaling (asymptotic) analysis to compute β . As p increases, SVs will be present on a narrower and narrower band around the interface. We denote by Δ the characteristic extension of this band. Δ will depend in general on the position \underline{x}_\perp along the interface. Here we will not study this dependence, as we are interested on its asymptotic behavior with p , γ and σ and only track how quantities depend on these variables. From the canonical condition equation (17) of SVs we have that the function f varies of order one from one side of the band to the other:

$$f(\underline{x}_\perp + \Delta \underline{e}_1) - f(\underline{x}_\perp - \Delta \underline{e}_1) \sim 1, \tag{20}$$

where \underline{e}_1 is the unit vector orthogonal to the interface and \underline{x}_\perp is any vector parallel to the plane.

⁶ We use $\tilde{\gamma}$ to distinguish it from the variance of the data points.

Another useful quantity is the distance r_c between nearest SVs. It can be estimated by counting the number of points lying within a cylinder of height Δ (along x_1) and radius r_c centered on an SV, whose volume follows $\sim \Delta r_c^{d-1}$. Using that the density of data points is $\sim p/\gamma^d$, and imposing that the cylinder contains only one additional SV, yields our first scaling relation:

$$\frac{p}{\gamma^d} \Delta r_c^{d-1} \sim 1 \implies \boxed{p \Delta r_c^{d-1} \sim \gamma^d}. \tag{21}$$

Finally, the last scaling relation results from the function fluctuations being of order one within the band of SVs when moving parallel to the true boundary decision. Indeed, we shall show below that the function gradient along e_\perp is constant at leading order in Δ , and of order $1/\Delta$ following equation (20). Then the facts that (i) on each SV the function is fixed by $f(\underline{x}^\mu) = y^\mu$ and (ii) the distance of the SV with respect to the true boundary fluctuates by a characteristic distance Δ jointly imply that the fluctuations of $f(\underline{x}^\mu)$ as \underline{x}^μ evolves along the true boundary decision must be of the order of unity. This effect is illustrated in figure 4. The characteristic transverse displacement along which these fluctuations decorrelate is simply the distance among SV r_c , thus

$$f(\underline{x}_\perp + r_c e_\perp) - f(\underline{x}_\perp) \sim 1, \tag{22}$$

where e_\perp is any unit vector parallel to the plane. Due to these fluctuations, test points inside the band have a finite probability to be incorrectly classified, and at fixed d^7 the test error must be proportional to the fraction Δ/γ of points falling in that band:

$$\epsilon \sim \Delta/\gamma.$$

We now show that from these considerations alone β can be computed. Starting from equation (19) we estimate the gradient of f along the normal direction e_\perp at any point on the interface:

$$\partial_{x_i} f(\underline{x}_\perp) = \frac{\xi}{\sigma} \sum_{\mu \in \Omega_\Delta} \alpha^\mu y^\mu \left(\frac{\|\underline{x}_\perp - \underline{x}^\mu\|}{\sigma} \right)^{\xi-1} \frac{x_1^\mu}{\|\underline{x}_\perp - \underline{x}^\mu\|} \approx \xi \sigma^{-\xi} p \frac{\Delta}{\gamma} \langle \alpha^\mu y^\mu x_1^\mu \|\underline{x}_\perp - \underline{x}^\mu\|^{\xi-2} \rangle_{\mu \in \Omega_\Delta}, \tag{23}$$

where the sum is over all SVs \underline{x}_μ indicated by the set Ω_Δ . The sum is replaced by its central-limit theorem value valid for large p , and we use that the number of terms in that sum goes as $p\Delta/\gamma$. The average in equation (23) scales as $\bar{\alpha} \Delta \gamma^{\xi-2}$, where $\bar{\alpha}$ is the mean value of the dual variables α^μ . Imposing that $\Delta \partial_{x_i} f(\underline{x}_\perp) \sim 1$ as follows from equation (20) then leads to our second scaling relation:

$$\boxed{p \bar{\alpha} \left(\frac{\Delta}{\gamma} \right)^3 \sim \left(\frac{\sigma}{\gamma} \right)^\xi}. \tag{24}$$

Next we compute the consequences of equation (22), by recasting it in a more suitable format. We define a smoothed function $\bar{f}(\underline{x}_\perp)$ of $f(\underline{x}_\perp)$ on a scale r_c :

$$\bar{f}(\underline{x}_\perp) = \int d^{d-1} \underline{x}'_\perp f(\underline{x}'_\perp) G(\underline{x}_\perp - \underline{x}'_\perp), \tag{25}$$

where the function G is the Fourier transform of $\theta(1/r_c - \|\underline{k}_\perp\|)$ (which is thus small when $\|\underline{x}_\perp - \underline{x}'_\perp\| \gg r_c$):

$$G(\underline{x}_\perp) = \int_{\|\underline{k}_\perp\| < 1/r_c} d^{d-1} \underline{k}_\perp e^{-i \underline{k}_\perp \cdot \underline{x}_\perp}. \tag{26}$$

Thus $\bar{f}(\underline{x}_\perp)$ is obtained by removing from $f(\underline{x}_\perp)$ the Fourier components $\|\underline{k}_\perp\| > 1/r_c$. The constraint of equation (22) is equivalent to imposing that the fluctuations between f and \bar{f} are of the order of unity. Integrated on space this means that

$$\gamma^{-d+1} \int d^{d-1} \underline{x}_\perp [f(\underline{x}_\perp) - \bar{f}(\underline{x}_\perp)]^2 \sim 1, \tag{27}$$

⁷ The value of $f(\underline{x})$ in the band is governed by the neighboring SVs, whose characteristic number is independent of p but should grow with d . We believe this effect to be responsible for the non-commutativity of the limits $\lim_{p/d \rightarrow \infty} \lim_{d \rightarrow \infty}$ and $\lim_{d \rightarrow \infty} \lim_{p \rightarrow \infty}$. Indeed in the former case, it is found [31] that ϵ and the fraction of SVs Δ/γ scale differently with α , unlike what we argue and confirm numerically in the second limit. We have checked numerically (not shown) that the ratio of these two quantities is indeed decaying with d at fixed p .

that can be Fourier transformed as

$$\int \mathbf{d}^{d-1} \underline{k}_\perp \left[\tilde{f}(\underline{k}_\perp) - \tilde{f}(\underline{k}_\perp) \tilde{G}(\underline{k}_\perp) \right]^2 = \int_{\|\underline{k}_\perp\| > 1/r_c} \mathbf{d}^{d-1} \underline{k}_\perp \tilde{f}^2(\underline{k}_\perp) \sim \gamma^{d-1}. \quad (28)$$

The Fourier transform of the decision function along the transverse components can be computed as

$$\tilde{f}(\underline{k}_\perp) = \int \mathbf{d}^{d-1} \underline{x}_\perp e^{-i \underline{k}_\perp \cdot \underline{x}_\perp} f(\underline{x}_\perp) = \sum_{\mu \in \Omega_\Delta} \alpha^\mu \gamma^\mu \int \mathbf{d}^{d-1} \underline{x}_\perp e^{-i \underline{k}_\perp \cdot \underline{x}_\perp} K\left(\frac{\|\underline{x}^\mu - \underline{x}_\perp\|}{\sigma}\right). \quad (29)$$

Using that $\|\underline{x}^\mu - \underline{x}_\perp\| \approx \|\underline{x}_\perp^\mu - \underline{x}_\perp\|$ and changing variables one obtains

$$\tilde{f}(\underline{k}_\perp) \approx \sum_{\mu \in \Omega_\Delta} \alpha^\mu \gamma^\mu e^{-i \underline{k}_\perp \cdot \underline{x}_\perp^\mu} \cdot \int \mathbf{d}^{d-1} \underline{x}_\perp e^{-i \underline{k}_\perp \cdot \underline{x}_\perp} K\left(\frac{\|\underline{x}_\perp\|}{\sigma}\right) \equiv \tilde{Q}(\underline{k}_\perp) \cdot \tilde{K}_\perp(\underline{k}_\perp), \quad (30)$$

where we have defined the kernel (transverse) Fourier transform $\tilde{K}_\perp(\underline{k}_\perp)$ and the ‘charge’ structure factor $\tilde{Q}(\underline{k}_\perp)$. The former can be readily computed for Laplace and Matérn kernels, and at large frequencies it behaves as $\tilde{K}_\perp(\underline{k}_\perp) \sim \sigma^{-\xi} \|\underline{k}_\perp\|^{-(d-1+\xi)}$. Concerning the charge structure factor, for $\|\underline{k}_\perp\| \gg 1/r_c$, the phases associated with each term in the sum defining it vary significantly even between neighboring SVs. From a central-limit argument the factor \tilde{Q} then tends to a random variable with 0 mean and variance $\bar{\alpha}^2 p \Delta / \gamma$. It is verified in appendix E.

We can now estimate the integral in equation (28):

$$\int_{\|\underline{k}_\perp\| > 1/r_c} \mathbf{d}^{d-1} \underline{k}_\perp \tilde{f}^2(\underline{k}_\perp) \sim \bar{\alpha}^2 p \frac{\Delta}{\gamma} \sigma^{-2\xi} \int_{\|\underline{k}_\perp\| > k_c} \mathbf{d}^{d-1} \underline{k}_\perp \|\underline{k}_\perp\|^{-2(d-1+\xi)} \sim \bar{\alpha}^2 p \frac{\Delta}{\gamma} \sigma^{-2\xi} r_c^{d-1+2\xi}. \quad (31)$$

The condition equation (28) leads to the last scaling relation:

$$\bar{\alpha}^2 p \frac{\Delta}{\gamma} \left(\frac{r_c}{\gamma}\right)^{d-1+2\xi} \sim \left(\frac{\sigma}{\gamma}\right)^{2\xi}. \quad (32)$$

Putting all the scaling relations together we find

$$\Delta \sim \gamma p^{-\frac{d-1+\xi}{3d-3+\xi}}, \quad \bar{\alpha} \sim \left(\frac{\sigma}{\gamma}\right)^\xi p^{\frac{2\xi}{3d-3+\xi}}, \quad r_c \sim \gamma p^{-\frac{2}{3d-3+\xi}}. \quad (33)$$

And consequently the asymptotic behavior of the test error is given by

$$\epsilon \sim \frac{\Delta}{\gamma} \sim p^{-\beta}, \quad \text{with } \beta = \frac{d-1+\xi}{3d-3+\xi}. \quad (34)$$

Note 1: The second scaling argument leading to equation (32) can be readily obtained by making a ‘minimal-disturbance hypothesis’. Assuming that adding a new training point \underline{x}^* within the domain Ω_Δ will only affect the dual variables of the few closest SVs, the correction of the decision function on the new SV is given by

$$\sum_{\|\underline{x}^\mu - \underline{x}^*\| \leq r_c} d\alpha^\mu \gamma^\mu \left(\frac{\|\underline{x}^\mu - \underline{x}^*\|}{\sigma}\right)^\xi, \quad (35)$$

where $d\alpha^\mu$ is the charge correction. One must have that $\sum_{\|\underline{x}^\mu - \underline{x}^*\| \leq r_c} d\alpha^\mu \gamma^\mu \approx -\gamma^* \alpha^*$ to ensure that SVs further away are not affected by this perturbation. Thus $d\alpha^\mu \sim \alpha^* \sim \bar{\alpha}$, where the last equivalence stems from the fact that the added SV is statistically identical to any other one. Finally, requiring that the new point \underline{x}^* must also be an SV implies that the correction represented by equation (35) must be of the order of unity to set $|f(\underline{x}^*)| = 1$. Hence, we obtain the scaling relation (that implies equation (32) from equations (21) and (24)):

$$\bar{\alpha} \left(\frac{r_c}{\sigma}\right)^\xi \sim 1. \quad (36)$$

Note 2: The above scaling arguments may also be carried out in the intermediate regime $\delta \ll \sigma < \gamma$. In that case, the kernel equation (18) introduces a cutoff to the volume of interaction in the transverse space. In

particular, the number of terms in the sum of equation (23) now goes as $(\sigma/\gamma)^{d-1}p\Delta/\gamma$ and the average scales as $\bar{\alpha}\Delta\sigma^{\xi-2}$. The discussion on the fluctuations is however unaltered as $r_c \ll \sigma$ by definition. Assembling all the pieces yields the following scaling relations:

$$\Delta \sim \gamma \left(\frac{\sigma}{\gamma}\right)^{-(d-1)\frac{d-3+\xi}{3d-3+\xi}} p^{-\frac{d-1+\xi}{3d-3+\xi}}, \quad \bar{\alpha} \sim \left(\frac{\sigma}{\gamma}\right)^{\frac{2\xi d}{3d-3+\xi}} p^{\frac{2\xi}{3d-3+\xi}}, \quad r_c \sim \gamma \left(\frac{\sigma}{\gamma}\right)^{\frac{d-3+\xi}{3d-3+\xi}} p^{-\frac{2}{3d-3+\xi}} \quad (37)$$

and

$$\epsilon \sim \frac{\Delta}{\gamma} \sim \left(\frac{\sigma}{\gamma}\right)^{-(d-1)\frac{d-3+\xi}{3d-3+\xi}} p^{-\frac{d-1+\xi}{3d-3+\xi}}. \quad (38)$$

Note that when this approach breaks down, namely when $\sigma \sim r_c$, the predictions of the vanishing bandwidth are recovered.

3.5. Multiple interfaces

The scaling analysis considered for the single interface can be directly extended to multiple interfaces. Let us consider the setup of n interfaces separated by a distance w . Because the target function oscillates around the n interfaces, its reproducing kernel Hilbert space (RKHS) norm increases with n leading to a more and more complicated task. In the limit $\Delta \ll w$, the arguments presented between equations (25) and (32) that rely on local considerations apply identically. The computation of the gradient is more subtle as the charges will in general differ in magnitude on each side of interfaces. We discuss in appendix F how the resulting gradient will scale with w . In particular, we identify three regimes on the (n, d) -plane as represented in figure 5. When the dimension is large enough, in the green region, the gradient is dominated by points with large transverse distance, $\|\underline{x}_\perp\| \gg w$. For smaller dimensions, the typical transverse distance decreases so that, in the blue region, the gradient is dominated by points of transverse distance $\|\underline{x}_\perp\| \sim w$. For even smaller dimensions, in the gray region, our description breaks down, because the SVC function is not sufficiently smooth and microscopic effect should be accounted for. The power laws of the three usual observables are shown to be

$$\Delta \sim \gamma \left(\frac{\gamma}{w}\right)^{\frac{(d-1)s}{3d-3+\xi}} p^{-\frac{d-1+\xi}{3d-3+\xi}}, \quad \bar{\alpha} \sim \left(\frac{\sigma}{\gamma}\right)^\xi \left(\frac{\gamma}{w}\right)^{\frac{\xi s}{3d-3+\xi}} p^{\frac{2\xi}{3d-3+\xi}}, \quad r_c \sim \gamma \left(\frac{w}{\gamma}\right)^{\frac{s}{3d-3+\xi}} p^{-\frac{2}{3d-3+\xi}}, \quad (39)$$

with

$$s = \begin{cases} n+1, & \text{if } 3 \leq n \leq d+\xi-4 \\ d+\xi-3, & \text{if } d+\xi-4 \leq n \leq d+\xi-1 \end{cases}, \quad \text{for } n \text{ odd}, \quad (40)$$

or

$$s = \begin{cases} n, & \text{if } 2 \leq n \leq d+\xi-3 \\ d+\xi-3, & \text{if } d+\xi-3 \leq n \leq d+\xi-1 \end{cases}, \quad \text{for } n \text{ even}. \quad (41)$$

The scaling in p is unaltered by the presence of multiple interfaces. However, the increasing complexity of the task is reflected by the large prefactor, which requires exponentially more training points to enter the power-law decay as the width w decreases. Note that, for a given dimension, the task complexity, quantified by $s(d-1)/(3d-3+\xi)$, stops increasing once n is large enough to enter the blue region.

3.6. Numerical results

In this section, we present the numerical simulations with which we verify the scalings predicted in the two previous sections. Both the single- and the double-interface setups have been considered with data points sampled from an isotropic Gaussian distribution of variance $\gamma^2 = 1$ along each component. In the single-interface setup the hyperplane is centered at $x_1 = 0$, while in the double-interface setup one hyperplane is located at $x_{\min} = -0.3$ and the other at $x_{\max} \approx 1.18549^8$. In both setups, the probabilities of positive and negative labels are equal. The margin-SVC algorithm is run using the class `svm.SVC` from the python library `scikitlearn`, which is a soft margin algorithm. To recover the hard margin algorithm presented in section 3.2, the regularization parameter C which bounds the dual variables from above (see for example chapter 7 of [45]) is set to $C = 10^{20}$. All results presented in this section have been obtained with the Laplace kernel of bandwidth $\sigma = 100 \gg \gamma$. Further results with the Matérn kernel are displayed in appendix G.

⁸ The value $x_{\max} = \sqrt{2} \operatorname{erf}^{-1}(1 + \operatorname{erf}(x_{\min})) \approx 1.18549$ is chosen in such a way that the expected number of $y = \pm 1$ points is the same.

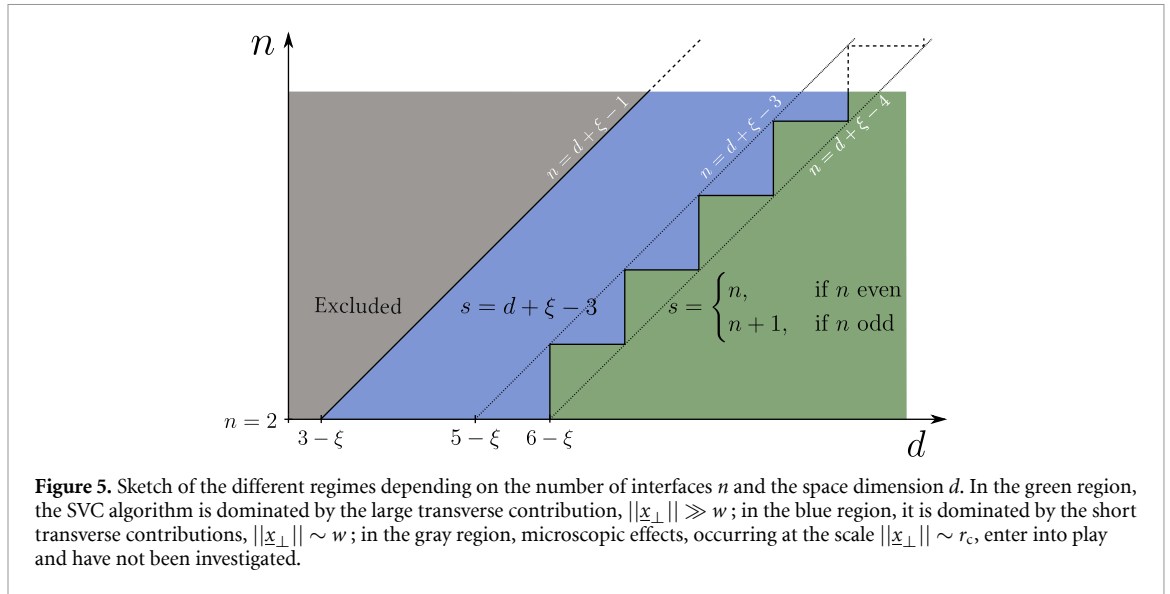


Figure 5. Sketch of the different regimes depending on the number of interfaces n and the space dimension d . In the green region, the SVC algorithm is dominated by the large transverse contribution, $\|x_{\perp}\| \gg w$; in the blue region, it is dominated by the short transverse contributions, $\|x_{\perp}\| \sim w$; in the gray region, microscopic effects, occurring at the scale $\|x_{\perp}\| \sim r_c$, enter into play and have not been investigated.

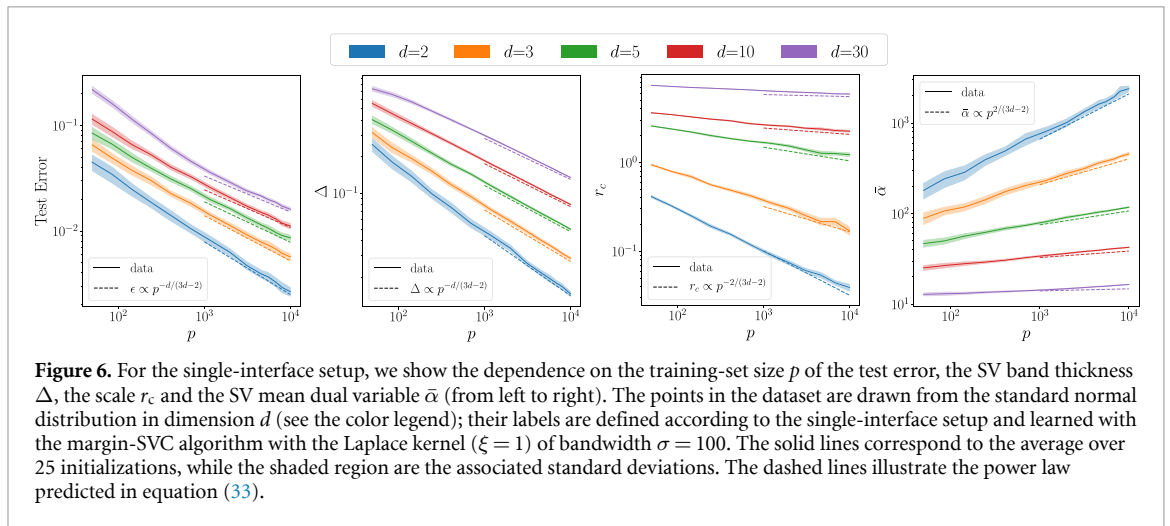


Figure 6. For the single-interface setup, we show the dependence on the training-set size p of the test error, the SV band thickness Δ , the scale r_c and the SV mean dual variable $\bar{\alpha}$ (from left to right). The points in the dataset are drawn from the standard normal distribution in dimension d (see the color legend); their labels are defined according to the single-interface setup and learned with the margin-SVC algorithm with the Laplace kernel ($\xi = 1$) of bandwidth $\sigma = 100$. The solid lines correspond to the average over 25 initializations, while the shaded region are the associated standard deviations. The dashed lines illustrate the power law predicted in equation (33).

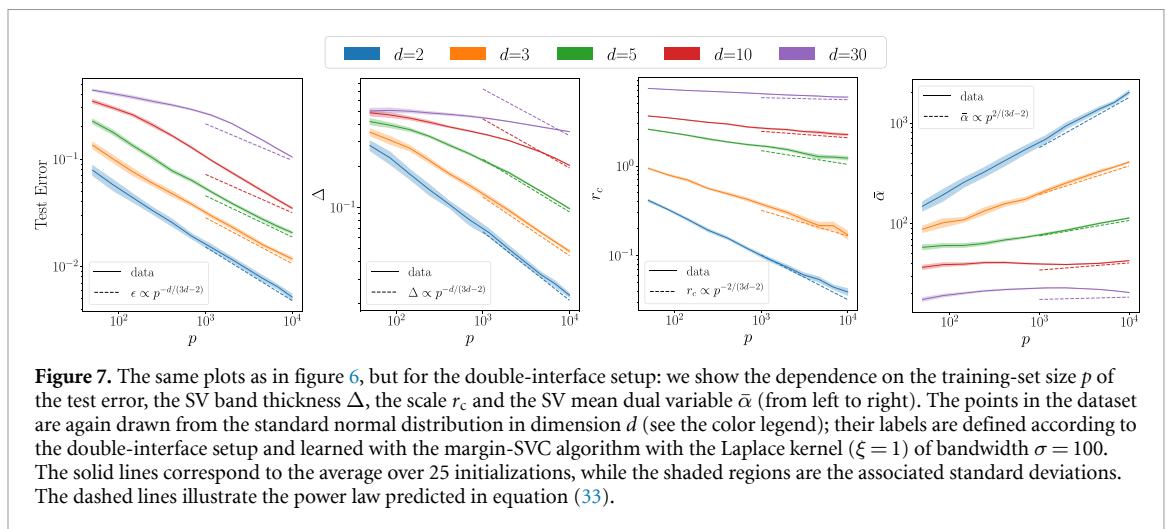
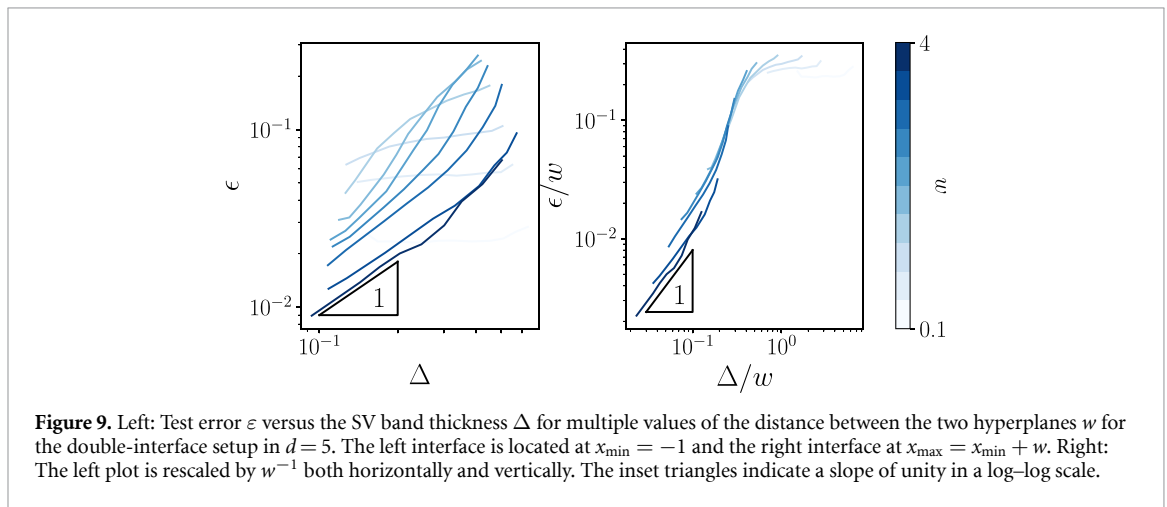
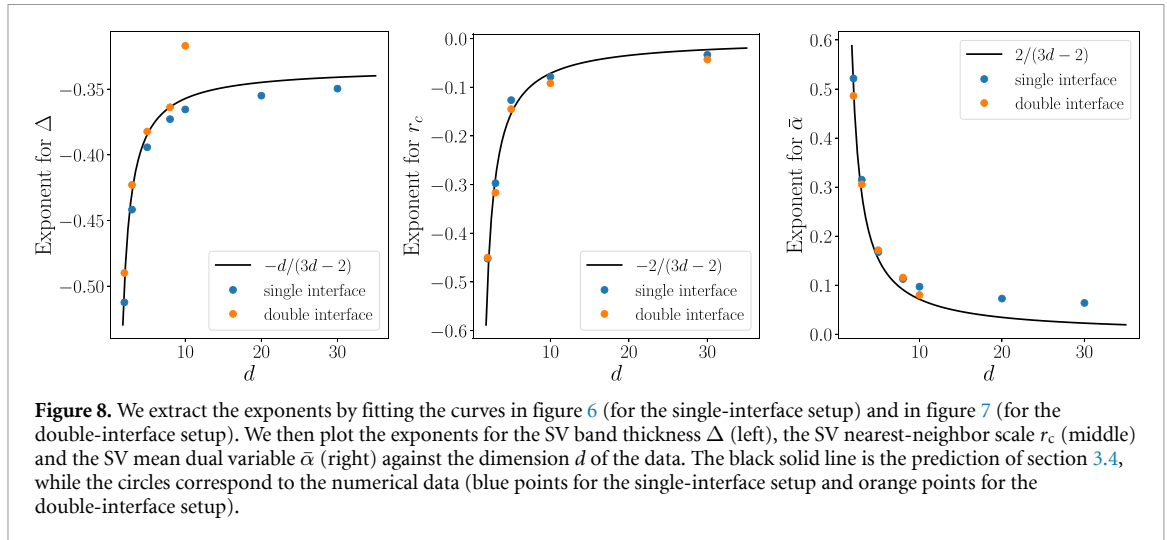


Figure 7. The same plots as in figure 6, but for the double-interface setup: we show the dependence on the training-set size p of the test error, the SV band thickness Δ , the scale r_c and the SV mean dual variable $\bar{\alpha}$ (from left to right). The points in the dataset are again drawn from the standard normal distribution in dimension d (see the color legend); their labels are defined according to the double-interface setup and learned with the margin-SVC algorithm with the Laplace kernel ($\xi = 1$) of bandwidth $\sigma = 100$. The solid lines correspond to the average over 25 initializations, while the shaded regions are the associated standard deviations. The dashed lines illustrate the power law predicted in equation (33).

The power-law predictions of section 3.4 are verified in figure 6 (for the single interface) and figure 7 (for the double interface). The considered numerical observables are defined as follows: the test error is the fraction of mislabeled points in a test set of size $p_{\text{test}} = 10000$; the typical $\bar{\alpha}$ is the average SV dual variable; the band thickness Δ is the average distance of a SV to the closest interface; the procedure to estimate the SV nearest-neighbor scale r_c is described in appendix H. The exponents of the power laws are extracted by fitting



the numerical curves in the asymptotic regime and compared to the theoretical predictions of section 3.4 in figure 8. Note that in large dimensions we observe that the system has not yet fully reached the asymptotic regime in the considered range of training-set sizes p .

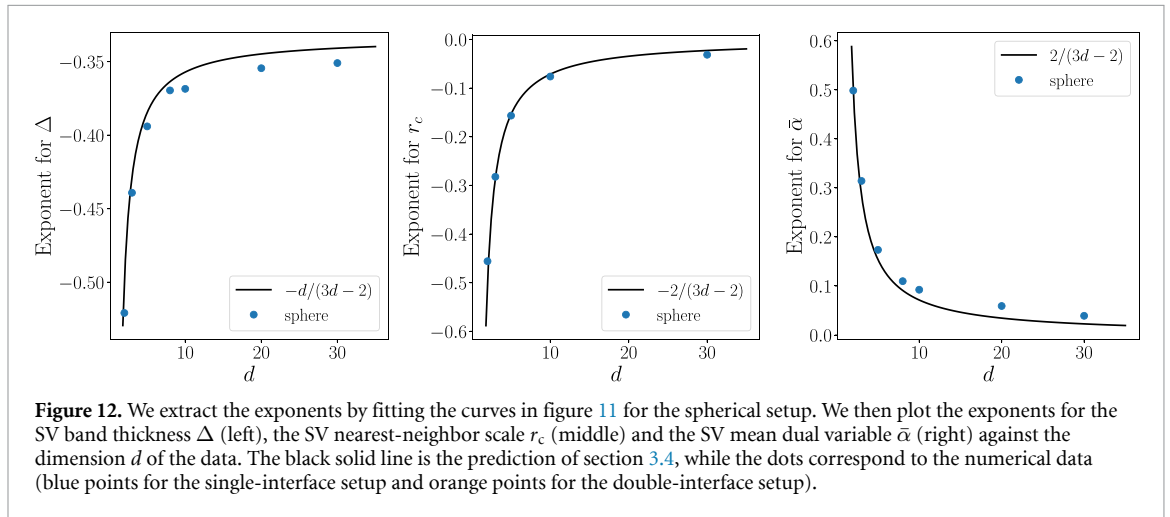
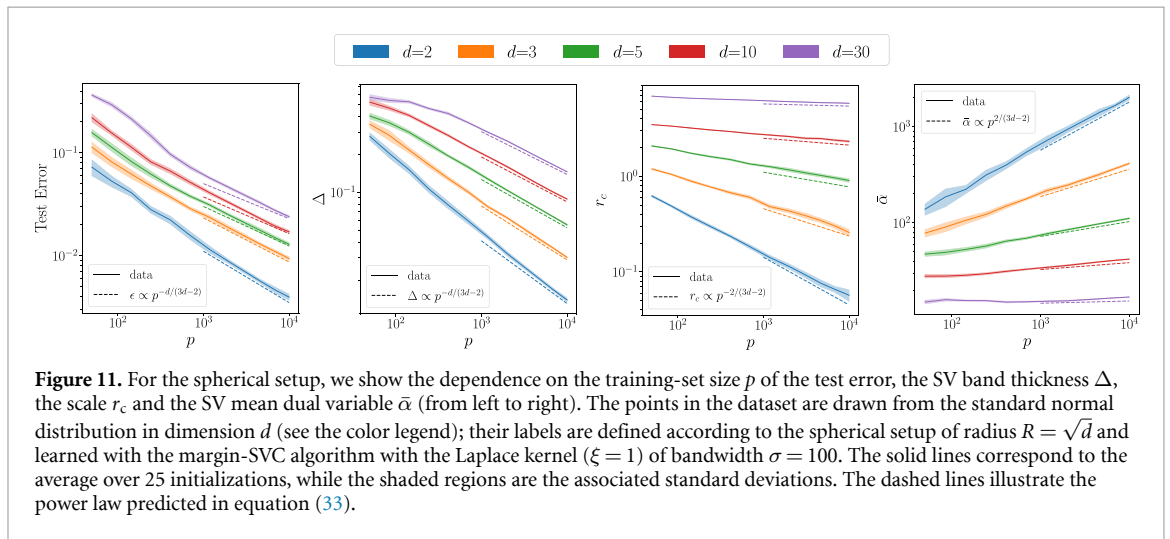

We also observe that in the double-interface setup the system only enters the scaling regime when Δ becomes small enough compared to the distance w between the two hyperplanes, as discussed in section 3.5. The crossover from the interfering-interface regime to the asymptotic regime is illustrated in figure 9. The test error versus Δ displayed in the left panel for multiple values of w confirms that $\varepsilon \sim \Delta$ when $\Delta \ll w$, as expected from the discussion of section 3.5. We show in the right panel that the transition to the asymptotic regime occurs when $\Delta \sim w$ by rescaling the horizontal axis: $\Delta \rightarrow \Delta/w$. Because $\varepsilon \sim \Delta$ in the asymptotic regime, it is necessary to also rescale the vertical axis for the curves to collapse, namely $\varepsilon \rightarrow \varepsilon/w$.

4. Spherical model

We consider a spherical interface separating $y = +1$ points outside a sphere of radius R from $y = -1$ points inside. The relevant direction is therefore $x_{\parallel} = \|\mathbf{x}\|$, and the label is given by $y(\mathbf{x}) = \text{sign}(\|\mathbf{x}\| - R)$. We still assume that the SV are distributed along the interface, thus forming a shell of radius R and thickness Δ . Once again, previous arguments presented between equations (25) and (32) that rely on local considerations apply identically. Furthermore, we compute in appendix I the gradient $\partial f / \partial x_{\parallel}$ and find again the same asymptotic result as for the planar interface specified in equation (23). Thus our predictions for the spherical model are identical to the ones for the stripe model.

We test these results numerically for a sphere of radius $R = \sqrt{d}^9$ with a Laplace kernel of variance $\sigma = 100$. The results displayed in figures 11 and 12 confirm our analysis.

⁹ It guarantees that the fraction of positive and negative labels remain finite. In particular, in the limit $d \rightarrow \infty$, this fraction goes to 1/2.

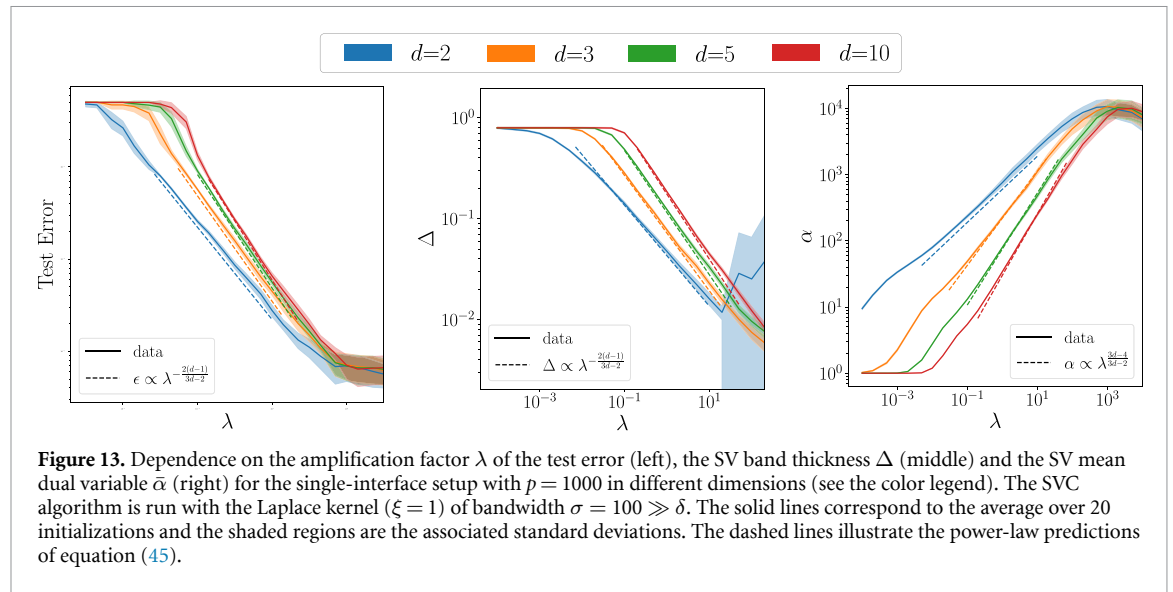


5. Improving kernel performance by compressing invariants

In this section, we investigate how compressing the data along the irrelevant directions x_{\perp} affects the performance of kernel classification. This analysis is of particular interest for neural networks, where it is now argued (see for instance [5]) that a progressive capability to compress invariants in the data is built up moving through the layers of deep networks.

5.1. Stripe model

We consider the stripe model of section 3 with one additional parameter: the amplification factor λ . If the original distribution was characterized by the scales $\gamma_1, \dots, \gamma_d$ along each space direction, we now apply a



contraction in the transverse space: $\gamma_i \rightarrow \gamma_i/\lambda$ for $i = 2, \dots, d$. Following the same reasoning as in section 3.4, we can track the effect of the additional amplification parameter. It is not sufficient to merely rescale γ , since the compression is not isotropic. Nevertheless, it is easy to see that the first scaling becomes

$$\lambda^{d-1} r_c^{d-1} \Delta p \sim \gamma^d, \tag{42}$$

since the density of points inside the SV band is now $\sim p\lambda^{d-1}/\gamma^d$. Then, for the second scaling relation, we need to rescale the gradient $\partial_{x_i} f$ defined in equation (23). The amplification factors only alter the transverse space: when approximating the average by an integral, the boundaries are rescaled to γ/λ in each transverse direction. The second scaling is thus

$$\lambda^{2-\xi} p \bar{\alpha} \left(\frac{\Delta}{\gamma}\right)^3 \left(\frac{\gamma}{\sigma}\right)^\xi \sim 1. \tag{43}$$

Finally, when imposing that the fluctuations between f and its smoothed version \tilde{f} are of the order of unity, one only needs to update the volume of the transverse space in equation (27): $\gamma^{d-1} \rightarrow (\gamma/\lambda)^{d-1}$, which leads to the last scaling,

$$\lambda^{d-1} \bar{\alpha}^2 p \frac{\Delta}{\gamma} \left(\frac{r_c}{\gamma}\right)^{d-1+2\xi} \sim \left(\frac{\sigma}{\gamma}\right)^{2\xi}. \tag{44}$$

Assembling all the scaling relations yields

$$\epsilon \sim \Delta \sim \gamma \lambda^{-\frac{2(d-1)}{3d-3+\xi}} p^{-\frac{d-1+\xi}{3d-3+\xi}}, \quad \bar{\alpha} \sim \left(\frac{\sigma}{\gamma}\right)^\xi \lambda^{\xi \frac{3d-5+\xi}{3d-3+\xi}} p^{\frac{2\xi}{3d-3+\xi}}, \quad r_c \sim \gamma \lambda^{-\frac{3d-5+\xi}{3d-3+\xi}} p^{-\frac{2}{3d-3+\xi}}. \tag{45}$$

These power laws are assessed numerically for the Laplace kernel ($\xi = 1$) of variance $\sigma = 100$ and a training set of size $p = 1000$ generated from the Gaussian distribution of variance $\gamma^2 = 1$. Varying the amplification factor over eight orders of magnitude (see figure 13), our predictions hold in a broad range of λ but break down at large and small values, as we now explain.

In the limit $\lambda \rightarrow 0$, the relevant direction x_1 is negligibly small compared to the other directions; information is thus suppressed and points are classified at random: the test error goes to 1/2. Furthermore, all training points must be SVs, and indeed $\Delta \rightarrow \langle |x| \rangle_{x \sim \mathcal{N}(0,1)} = \sqrt{2/\pi}$ (which is the average distance from any point in the dataset to the interface) and $\bar{\alpha} \rightarrow 1$.

In the opposite limit $\lambda \rightarrow \infty$ the setup lives in dimension one (seeing only x_1) and all curves converge independently of the space dimension d . These relations allow us to identify a critical scale λ_c at which the multidimensional system reduces effectively to a one-dimensional system. It occurs when the test error of the compressed multidimensional kernel is equal to the test error of the kernel that only sees the component x_1 . Using our scalings, we find

$$\lambda^{-\frac{2(d-1)}{3d-3+\xi}} p^{-\frac{d-1+\xi}{3d-3+\xi}} \sim p^{-1} \implies \lambda_c \sim p. \tag{46}$$

5.2. Cylinder model

We now consider a cylinder model in $d = d_{\parallel} + d_{\perp}$ dimensions. A point $\underline{x} = (\underline{x}_{\parallel}, \underline{x}_{\perp}) \in \mathbb{R}^d$ (with $\underline{x}_{\parallel} \in \mathbb{R}^{d_{\parallel}}$ and $\underline{x}_{\perp} \in \mathbb{R}^{d_{\perp}}$) has a positive label if $\|\underline{x}_{\parallel}\| > R \sim \gamma$ and a negative label otherwise. Such a model is also characterized by the asymptotic scalings in p specified in equation (23).

As in the previous section, we compress the perpendicular directions by the amplification factor λ : $\underline{x}_{\perp} \rightarrow \underline{x}_{\perp}/\lambda$. The derivations of the scaling relations equations (42) and (44) hold equally. However, the scaling relation equation (43) is now independent of the amplification factor: the characteristic size of the transverse space occurring in the gradient integral equation (23) remains of the order of the system size γ . Assembling the different scalings yields

$$\epsilon \sim \Delta \sim \gamma \lambda^{-\frac{\xi d_{\perp}}{3d-3+\xi}} p^{-\frac{d-1+\xi}{3d-3+\xi}}, \quad \bar{\alpha} \sim \left(\frac{\sigma}{\gamma}\right)^{\xi} \lambda^{\frac{3\xi d_{\perp}}{3d-3+\xi}} p^{\frac{2\xi}{3d-3+\xi}}, \quad r_c \sim \gamma \lambda^{-\frac{3d_{\perp}}{3d-3+\xi}} p^{-\frac{2}{3d-3+\xi}}. \quad (47)$$

6. Conclusion

We have studied the learning curve exponent β of an isotropic kernel in the presence of invariants, improving on worst case bounds previously obtained in the literature. For regression on Gaussian fields, we find that invariants do not increase β , which behaves as $\sim d^{-1}$ in large dimension: methods based on isotropic kernels suffer from the curse of dimensionality, as already argued in [4]. Our analysis also suggests a simple estimate (8) for the performance of regression beyond the Gaussian fields considered here. For a binary classification and simple models of invariants we find the opposite result. For a planar interface separating labels, $\beta \geq 1/3$ for all dimensions, improving on previous bounds.

Note that the striking difference between classification and regression does not stem from the distinct models considered in each case. Indeed, following equation (8) we expect that performing mean-square ridgeless regression on the stripe model leads to the curse of dimensionality with $\beta = 1/d$, as we have checked on a few examples (data not shown). In the classification problem instead, due to the fact that only a tiny band of data are SVs, the output function ends up being much smoother (i.e. with more rapidly decaying Fourier components) than a step function, leading to better performance.

This success of classification holds when several interfaces are present, or in the spherical case where the interface continuously bends. Thus, isotropic kernels can beat the curse of dimensionality even for non-planar boundaries between labels. For which class of boundaries is this result true? The geometry of the spatial distribution of SVs suggests an intuitive answer. The curse of dimensionality is beaten because a very narrow (i.e. rapidly decaying with p) layer of width Δ is sufficient to fit all data, despite the fact that the distance between SVs r_c is much larger (and indeed subjected to the curse of dimensionality). Thus if the boundary displays significant variations below the scale r_c , it presumably cannot be detected by isotropic kernels. In this view, beating the curse of dimensionality is only possible if the boundary becomes more and more regular as the dimension increases. This geometrical view is consistent with the more abstract kernel literature in which the curse is lifted if labels correspond to the sign of a regular function (in the sense of belonging to the RKHS of the kernel) [36]. Empirically, sufficient regularity may be achieved in practical settings at least along some invariants, such as completely uninformative pixels near the boundaries of images. Under which conditions other invariants, e.g. related to translation, can be exploited by isotropic kernels remains to be understood.

Note added: In [46], these results were extended beyond kernels, to the case of a wide one-hidden layer net. In the lazy training regime, results are identical to those presented here, but more favorable exponents β are found in the feature learning regime.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://gitlab.com/jonas.paccolat/svc-for-simple-invariants>.

Acknowledgments

We acknowledge L Chizat, M Geiger, P Loucas, L Petrini and C Pehlevan for discussions and L Chizat for pointing out several important references. This work was partially supported by a grant from the Simons Foundation (454 953 Matthieu Wyart). M W thanks the Swiss National Science Foundation for support under Grant No. 200 021-165 509.

Appendix A. Kernel regression with invariant dimensions

Theorem. Let $K_T(\underline{x})$ and $K_S(\underline{x})$ be two translation-invariant kernels (called the teacher and student respectively) defined on $\mathcal{V}_d \equiv \mathbb{R}^d$, and let $\tilde{K}_T(\underline{w})$ and $\tilde{K}_S(\underline{w})$ be their Fourier transforms in \mathcal{V}_d . Assume that

- $K_T(\underline{x}), K_S(\underline{x})$ are continuous everywhere and differentiable everywhere except at the origin $\underline{x} = 0$;
- $K_T(\underline{x})$ and $K_S(\underline{x})$ are positive definite and isotropic, that is, they only depend on $\|\underline{x}\|$;
- $K_T(\underline{x})$ and $K_S(\underline{x})$ have a cusp at the origin and their d -dimensional Fourier transform decays at high frequencies with dimensional-dependent exponents $\alpha_T(d_{\parallel})$ and $\alpha_S(d)$, respectively (we will evaluate them at d_{\parallel} for the teacher and at d for the student);
- $\lim_{\underline{x} \rightarrow 0} K_T(0) < \infty$ and $\lim_{\underline{x} \rightarrow 0} K_S(0) < \infty$;
- $\lim_{\underline{w} \rightarrow 0} \tilde{K}_T(\underline{w}) < \infty$ and $\lim_{\underline{w} \rightarrow 0} \tilde{K}_S(\underline{w}) < \infty$.

Assume furthermore that the teacher kernel lives in a reduced space of dimension $d_{\parallel} \leq d$, in the sense that

- $K_T(\underline{x}) \equiv K_T(x_1, \dots, x_d) = K_T(x_1, \dots, x_{d_{\parallel}}) \equiv K_T(\|\underline{x}_{\parallel}\|)$ (where we have defined $\underline{x}_{\parallel} \equiv (x_1, \dots, x_{d_{\parallel}})^T$).

We use the teacher kernel to sample a Gaussian random field $Z_T(\underline{x}) \sim \mathcal{N}(0, K_T)$ at points that lie on a d -dimensional regular lattice in \mathcal{V}_d , with fixed spacing δ , and we use the student kernel to infer $\hat{Z}_S(\underline{x})$ at a new point $\underline{x} \in \mathcal{V}_d$ via regression, and performance is then evaluated by computing the expected mean-squared error on points independent from those used for training. Then, as $\delta \rightarrow 0$,

$$\mathbb{E} \text{MSE} \sim \delta^{\beta d} \quad \text{with} \quad \beta = \frac{1}{d} \min(\alpha_T(d_{\parallel}) - d_{\parallel}, 2\alpha_S). \quad (\text{A1})$$

Proof.

(i) *Setup.*

We first consider a finite number of points p in a box $\mathcal{V}_d = [-L/2, L/2]^d$ and then take the limit $p, L \rightarrow \infty$, keeping the spacing $\delta = Lp^{-1/d}$ fixed. Regression is done by minimizing the mean-squared error on the p points:

$$\sum_{\mu=1}^p \left[Z_T(\underline{x}_{\mu}) - \hat{Z}_S(\underline{x}_{\mu}) \right]^2, \quad (\text{A2})$$

and the generalization error is defined as

$$\mathbb{E} \text{MSE} = L^{-d} \mathbb{E} \int_{\mathcal{V}_d} d^d \underline{x} \left[Z_T(\underline{x}) - \hat{Z}_S(\underline{x}) \right]^2. \quad (\text{A3})$$

(The expectation value is taken with respect to the teacher random process.)

Given a function $F(\underline{x})$ on the d -dimensional box $\mathcal{V}_d = [-L/2, L/2]^d$, we denote its Fourier transform (series) and antitransform by

$$\tilde{F}(\underline{w}) \equiv \mathcal{F}_d[F(\underline{x})](\underline{w}) = L^{-d/2} \int_{\mathcal{V}_d} d^d \underline{x} e^{-i\underline{w} \cdot \underline{x}} F(\underline{x}), \quad \text{where } \underline{w} \in \mathbb{L}_d \equiv \frac{2\pi}{L} \mathbb{Z}^d, \quad (\text{A4})$$

$$F(\underline{x}) \equiv \mathcal{F}_d^{-1}[\tilde{F}(\underline{w})](\underline{x}) = L^{-d/2} \sum_{\underline{w} \in \mathbb{L}} e^{i\underline{w} \cdot \underline{x}} \tilde{F}(\underline{w}). \quad (\text{A5})$$

Given the structure of the teacher kernel we can write

$$\begin{aligned} \tilde{K}_T(\underline{w}) &= L^{-d_{\parallel}/2} \int_{[-L/2, L/2]^{d_{\parallel}}} d^{d_{\parallel}} \underline{x}_{\parallel} e^{-i\underline{w}_{\parallel} \cdot \underline{x}_{\parallel}} K_T(\|\underline{x}_{\parallel}\|) \cdot L^{-d_{\perp}/2} \int_{[-L/2, L/2]^{d_{\perp}}} d^{d_{\perp}} \underline{x}_{\perp} e^{-i\underline{w}_{\perp} \cdot \underline{x}_{\perp}} \\ &= \mathcal{F}_{d_{\parallel}} \left[K_T(\|\underline{x}_{\parallel}\|) \right] (\underline{w}_{\parallel}) \cdot L^{d_{\perp}/2} \delta_{\underline{w}_{\perp}}. \end{aligned} \quad (\text{A6})$$

This formula states that the Fourier transform of the teacher kernel has frequencies that also live in the corresponding d_{\parallel} -dimensional subspace in the frequency domain. The term $\delta_{\underline{w}_{\perp}}$ is a discrete delta (not a Dirac delta): this will be important later because it implies that it is scale invariant: $\delta_{a\underline{w}_{\perp}} = \delta_{\underline{w}_{\perp}}$. The first

term, that is the Fourier transform of the teacher kernel restricted to the d_{\parallel} -dimensional space, decays at large frequencies with an exponent $\alpha_T(d_{\parallel})$ that depends on the intrinsic dimension d_{\parallel} :

$$\tilde{K}_T^{\parallel}(\underline{w}_{\parallel}) \equiv \mathcal{F}_{d_{\parallel}} \left[K_T \left(\|\underline{x}_{\parallel}\| \right) \right] (\underline{w}_{\parallel}) = c_T(d_{\parallel}) \|\underline{w}_{\parallel}\|^{-\alpha_T(d_{\parallel})} + o \left(\|\underline{w}_{\parallel}\|^{-\alpha_T(d_{\parallel})} \right). \tag{A7}$$

(ii) *Regression.*

The solution to the regression problem can be computed in closed form:

$$\hat{Z}_S(\underline{x}) = \underline{k}_S(\underline{x}) \cdot \mathbb{K}_S^{-1} Z_T, \tag{A8}$$

where $Z_T = (Z_T(\underline{x}_{\mu}))_{\mu=1}^p$ are the training data (the points \underline{x}_{μ} lie on the regular lattice), $\underline{k}_S(\underline{x}) = (K_S(\underline{x}_{\mu}, \underline{x}))_{\mu=1}^p$ and $\mathbb{K}_S = (K_S(\underline{x}_{\mu}, \underline{x}_{\nu}))_{\mu, \nu=1}^p$ is the Gram matrix, that is invertible since the kernel K_S is assumed to be positive definite. This formula can be written in Fourier space as

$$\tilde{Z}_S(\underline{w}) = \tilde{Z}^*(\underline{w}) \frac{\tilde{K}_S(\underline{w})}{\tilde{K}_S^*(\underline{w})}, \tag{A9}$$

where we have defined $F^*(\underline{w}) \equiv \sum_{\underline{n} \in \mathbb{Z}^d} F \left(\underline{w} + \frac{2\pi \underline{n}}{\delta} \right)$ for a generic function F .

The mean-squared error can then be written using the Parseval–Plancherel identity. After some calculations we find

$$\begin{aligned} \mathbb{E} \text{MSE} &= L^{-d} \mathbb{E} \int_{\mathcal{V}_d} d^d \underline{x} [Z_T(\underline{x}) - \hat{Z}_S(\underline{x})]^2 = L^{-d} \mathbb{E} \sum_{\underline{w} \in \mathbb{L}_d} \left| \tilde{Z}_T(\underline{w}) - \tilde{Z}^*(\underline{w}) \frac{\tilde{K}_S(\underline{w})}{\tilde{K}_S^*(\underline{w})} \right|^2 \\ &= L^{-d/2} \sum_{\underline{w} \in \mathbb{L}_d \cap \mathcal{B}_d} \tilde{K}_T^*(\underline{w}) - 2 \frac{[\tilde{K}_T \tilde{K}_S]^*(\underline{w})}{\tilde{K}_S^*(\underline{w})} + \frac{\tilde{K}_T^*(\underline{w}) [\tilde{K}_S^*]^*(\underline{w})}{\tilde{K}_S^*(\underline{w})^2}, \end{aligned} \tag{A10}$$

where $\mathbb{L}_d = \frac{2\pi}{L} \mathbb{Z}^d$ and $\mathcal{B}_d = [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^d$ is the Brillouin zone.

In order to simplify this expression in the case where $d_{\parallel} \leq d$, let us also introduce

$$F^{*\parallel}(\underline{w}_{\parallel}) \equiv \sum_{\underline{n}_{\parallel} \in \mathbb{Z}^{d_{\parallel}}} F \left(\underline{w}_{\parallel} + \frac{2\pi \underline{n}_{\parallel}}{\delta} \right). \tag{A11}$$

Using equation (A6) it follows that

$$\tilde{K}_T^*(\underline{w}) \propto \delta_{\underline{w}_{\perp}} \tilde{K}_T^{*\parallel}(\underline{w}_{\parallel}), \tag{A12}$$

$$[\tilde{K}_T \tilde{K}_S]^*(\underline{w}) \propto \delta_{\underline{w}_{\perp}} [\tilde{K}_T \tilde{K}_S]^{*\parallel}(\underline{w}_{\parallel}). \tag{A13}$$

Plugging the last two equations into equation (A10) we see that, because of the terms $\delta_{\underline{w}_{\perp}}$, we have

$$\mathbb{E} \text{MSE} \propto \sum_{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel}} \tilde{K}_T^{*\parallel}(\underline{w}_{\parallel}) \left\{ 1 + \frac{[\tilde{K}_S^*]^*(\underline{w}_{\parallel})}{\tilde{K}_S^*(\underline{w}_{\parallel})} \right\} - 2 \frac{[\tilde{K}_T \tilde{K}_S]^{*\parallel}(\underline{w}_{\parallel})}{\tilde{K}_S^*(\underline{w}_{\parallel})}. \tag{A14}$$

Notice that \tilde{K}_S^* and $[\tilde{K}_S^*]^*$ do not turn into $[\tilde{K}_S]^{*\parallel}$ and $[\tilde{K}_S^*]^{\parallel}$: this is because the student kernel does not have the same invariants as the teacher, and it depends on all the components. Here $\mathbb{L}_{\parallel} = \frac{2\pi}{L} \mathbb{Z}^{d_{\parallel}}$, $\mathcal{B}_{\parallel} = [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^{d_{\parallel}}$.

(iii) *Expansion.*

Using the high-frequency behavior of the Fourier transforms of the two kernels we can write

$$\tilde{K}_T^{*\parallel}(\underline{w}_{\parallel}) \sim \tilde{K}_T(\underline{w}_{\parallel}) + \delta^{\alpha_T(d_{\parallel})} c_T(d_{\parallel}) \psi_{\alpha_T(d_{\parallel})}^{\parallel}(\underline{w}_{\parallel} \delta), \tag{A15}$$

$$[\tilde{K}_T \tilde{K}_S]^{*\parallel}(\underline{w}_{\parallel}) \sim \tilde{K}_T(\underline{w}_{\parallel}) \tilde{K}_S(\underline{w}_{\parallel}) + \delta^{\alpha_T(d_{\parallel}) + \alpha_S} c_T(d_{\parallel}) c_S \psi_{\alpha_T(d_{\parallel}) + \alpha_S}^{\parallel}(\underline{w}_{\parallel} \delta), \tag{A16}$$

$$\tilde{K}_S^*(\underline{w}_{\parallel}) \sim \tilde{K}_S(\underline{w}_{\parallel}) + \delta^{\alpha_S} c_S \psi_{\alpha_S}(\underline{w}_{\parallel} \delta). \tag{A17}$$

We have introduced the functions

$$\psi_{\alpha}(\underline{w}_{\parallel}) = \sum_{\underline{n} \in \mathbb{Z}^d \setminus \{0\}} \|\underline{w}_{\parallel} + 2\pi \underline{n}\|^{-\alpha}, \tag{A18}$$

$$\psi_{\alpha}^{\parallel}(\underline{w}_{\parallel}) = \sum_{\underline{n}_{\parallel} \in \mathbb{Z}^d \setminus \{0\}} \|\underline{w}_{\parallel} + 2\pi \underline{n}_{\parallel}\|^{-\alpha}. \tag{A19}$$

The hypothesis $K_T(\underline{0}) \propto \int d\underline{w} \tilde{K}_T(\underline{w}) < \infty$ and $K_S(\underline{0}) < \infty$ imply $\alpha_T(d_{\parallel}) > d_{\parallel}$ and therefore $\sum_{\underline{n}_{\parallel} \in \mathbb{Z}^d} \|\underline{n}_{\parallel}\|^{-\alpha_T(d_{\parallel})} < \infty$. We can argue similarly that $\psi_{\alpha_T(d_{\parallel})}^{\parallel}(\underline{0}), \psi_{\alpha_T(d_{\parallel})+\alpha_S}^{\parallel}(\underline{0}), \psi_{\alpha_S}(\underline{0})$ are finite. Furthermore, the $\underline{w}_{\parallel}$'s in the sums are at most of order δ^{-1} , therefore the terms $\psi_{\alpha}(\underline{w}\delta)$ are δ^0 and do not influence how equation (A10) scales with δ .

Expanding equation (A10) and keeping only the highest orders we find

EMSE

$$\sim \sum_{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel}} \left[2c_T(d_{\parallel}) \psi_{\alpha_T(d_{\parallel})}^{\parallel}(\underline{w}_{\parallel} \delta) \delta^{\alpha_T(d_{\parallel})} + c_S^2 \psi_{2\alpha_S}(\underline{w}_{\parallel} \delta) \frac{\tilde{K}_T^{\parallel}(\underline{w}_{\parallel})}{\tilde{K}_S^2(\underline{w}_{\parallel})} \delta^{2\alpha_S} \right] + o\left(\|\underline{w}_{\parallel}\|^{-\alpha_T(d_{\parallel})-d_{\parallel}}\right). \tag{A20}$$

We have neglected terms proportional to, for instance, $\delta^{\alpha_T(d_{\parallel})+\alpha_S}$, since they are subleading with respect to $\delta^{\alpha_T(d_{\parallel})}$, but we must keep both $\delta^{\alpha_T(d_{\parallel})}$ and δ^{α_S} since we do not know *a priori* which one is dominant. The additional term δ^{-d} in the subleading terms comes from the fact that $|\mathbb{L} \cap \mathcal{B}| \sim \delta^{-d}$.

The first term in equation (A20) is the simplest to deal with: since $\|\underline{w}_{\parallel} \delta\|$ is smaller than some constant for all $\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel}$ and the function $\psi_{\alpha_T(d_{\parallel})}^{\parallel}(\underline{w}_{\parallel} \delta)$ has a finite limit, we have

$$\delta^{\alpha_T(d_{\parallel})} \sum_{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel}} 2c_T(d_{\parallel}) \psi_{\alpha_T(d_{\parallel})}^{\parallel}(\underline{w}_{\parallel} \delta) \sim \delta^{\alpha_T(d_{\parallel})} |\mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel}| \sim \delta^{\alpha_T(d_{\parallel})-d_{\parallel}}. \tag{A21}$$

We then split the second term in equation (A20) in two contributions.

Small $\|\underline{w}_{\parallel}\|$. We consider 'small' all the terms $\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel}$ such that $\|\underline{w}_{\parallel}\| < \Gamma$, where $\Gamma \gg 1$ is of order δ^0 but large. As $\delta \rightarrow 0$, $\psi_{2\alpha_S}(\underline{w}_{\parallel} \delta) \rightarrow \psi_{2\alpha_S}(\underline{0})$ which is finite because $K_S(\underline{0}) < \infty$. Therefore

$$\delta^{2\alpha_S} \sum_{\substack{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel} \\ \|\underline{w}_{\parallel}\| < \Gamma}} c_S^2 \psi_{2\alpha_S}(\underline{w}_{\parallel} \delta) \frac{\tilde{K}_T^{\parallel}(\underline{w}_{\parallel})}{\tilde{K}_S^2(\underline{w}_{\parallel})} \rightarrow \delta^{2\alpha_S} c_S^2 \psi_{2\alpha_S}(\underline{0}) \sum_{\substack{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel} \\ \|\underline{w}_{\parallel}\| < \Gamma}} \frac{\tilde{K}_T^{\parallel}(\underline{w}_{\parallel})}{\tilde{K}_S^2(\underline{w}_{\parallel})}. \tag{A22}$$

The summand is real and strictly positive because the positive definiteness of the kernels implies that their Fourier transforms are strictly positive. Moreover, as $\delta \rightarrow 0$, $\mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel} \cap \{\|\underline{w}_{\parallel}\| < \Gamma\} \rightarrow \mathbb{L}_{\parallel} \cap \{\|\underline{w}_{\parallel}\| < \Gamma\}$, which contains a finite number of elements, independent of δ . Therefore

$$\delta^{2\alpha_S} \sum_{\substack{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel} \\ \|\underline{w}_{\parallel}\| < \Gamma}} c_S^2 \psi_{2\alpha_S}(\underline{w}_{\parallel} \delta) \frac{\tilde{K}_T^{\parallel}(\underline{w}_{\parallel})}{\tilde{K}_S^2(\underline{w}_{\parallel})} \sim \delta^{2\alpha_S}. \tag{A23}$$

Large $\|\underline{w}_{\parallel}\|$. 'Large' \underline{w} are those with $\|\underline{w}_{\parallel}\| > \Gamma$: we recall that $\Gamma \gg 1$ is of order δ^0 but large. This allows us to approximate $\tilde{K}_T^{\parallel}, \tilde{K}_S$ in the sum with their asymptotic behavior:

$$\delta^{2\alpha_S} \sum_{\substack{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel} \\ \|\underline{w}_{\parallel}\| > \Gamma}} c_S^2 \psi_{2\alpha_S}(\underline{w}_{\parallel} \delta) \frac{\tilde{K}_T^{\parallel}(\underline{w}_{\parallel})}{\tilde{K}_S^2(\underline{w}_{\parallel})} \propto \delta^{2\alpha_S} \sum_{\substack{\underline{w}_{\parallel} \in \mathbb{L}_{\parallel} \cap \mathcal{B}_{\parallel} \\ \|\underline{w}_{\parallel}\| > \Gamma}} \|\underline{w}_{\parallel}\|^{-\alpha_T(d_{\parallel})+2\alpha_S}$$

$$\approx \delta^{2\alpha_S} \int_{\Gamma}^{1/\delta} dw_{\parallel} w_{\parallel}^{d_{\parallel}-1-\alpha_T(d_{\parallel})+2\alpha_S} \sim \delta^{\min(\alpha_T(d_{\parallel})-d_{\parallel}, 2\alpha_S)}. \tag{A24}$$

Therefore in the end

$$\mathbb{E} \text{MSE} \sim \delta^{\min(\alpha_T(d_{\parallel})-d_{\parallel}, 2\alpha_S)} \equiv \delta^{\beta d}. \tag{A25}$$

The kernels K that we consider in the present article, namely Laplace and Matérn, share the property that the respective exponents take the form $\alpha_K(d) = d + \theta_K$, θ_K being a dimension-independent constant that only depends on the isotropic function that defines the kernel. For instance, we have $\alpha = d + 1$ for Laplace and $\alpha(d) = d + 2\nu$ for Matérn (with parameter ν). Consequently, for these kernels the term $\alpha(d_{\parallel}) - d_{\parallel}$ that appears in the last equation is actually independent of d_{\parallel} , and therefore so is the exponent β . We believe that this structure of the exponent $\alpha(d)$ is more general. Signals that point in this direction can be found in several papers. In [47] they show that (with our notation), for functions $K(\|\underline{x}\|)$ that are integrable in \mathbb{R}^d and \mathbb{R}^{d+2} ,

$$\mathcal{F}_{d+2}[K(\|\underline{x}\|)](w) \propto w^{-1} \partial_w \mathcal{F}_d[K(\|\underline{x}\|)](w), \tag{A26}$$

and so if the Fourier transform in dimension d decays as $w^{-\alpha(d)}$, in dimension $d + 2$ it decays with an exponent $\alpha(d + 2) = \alpha(d) + 2$. In [48] they prove a result for functions belonging to the Schwartz space (rapidly decreasing functions). This result implies that if the Fourier transform in dimension $d + 1$ decays with an exponent $\alpha(d + 1)$, then in dimension d the function decays with the exponent $\alpha(d) = \alpha(d + 1) - 1$.

These results offer a link between the exponents in different dimensions. In [49] the author computes the asymptotic behavior of the one-dimensional Fourier transform of functions with a singularity. In particular, it follows that if $K(x) = |x|^{\theta_K} K_{\infty}(x)$, with $-1 < \theta_K \leq 0$ and $K_{\infty} \in C^{\infty}(\mathbb{R})$, then its Fourier transform at the leading order decays with an exponent $\alpha(d = 1) = 1 + \theta_K$. There is a similarity with the value of the exponents for the Laplace and Matérn kernels that we use: the value of θ_K is linked to the exponent of the cusp $|x|^{\theta_K}$ that appears in the Taylor expansion of the Kernel at the origin. We expect that this fact, namely that the exponent $\alpha_K(d)$ is the sum of spatial dimension d and of the cusp exponent θ_K , is more generic and applies to most of the kernels that are used in practice.

Appendix B. Regime $\sigma \ll \delta$: curse of dimensionality

We consider here the case where the kernel bandwidth σ is much smaller than the nearest-neighbor distance δ . In this limit the contributions in the expansion of the decision boundary in equation (12) are significantly suppressed because the kernel is supposed to decay when its argument is large, and the decision boundary is dominated by the charge of training pattern \underline{x}^{μ} that is closest to \underline{x} . The sign of the decision function is thus fixed by the sign of the nearest neighbor's charge and the accuracy is driven by the nearest neighbor distance, which is susceptible to the curse of dimensionality.

We can see this more precisely if we approximate the kernel interaction between two points \underline{x} and \underline{x}' as

$$K\left(\frac{\|\underline{x} - \underline{x}'\|}{\sigma}\right) \approx \begin{cases} a_0 = K(0) & \text{if } \underline{x} = \underline{x}', \\ a_1 = K\left(\frac{\delta}{\sigma}\right) \ll a_0 & \text{if } \underline{x}' \text{ is one of the nearest neighbors of } \underline{x},^{10} \\ 0 & \text{otherwise.} \end{cases} \tag{B1}$$

Hence, the decision function at a point \underline{x}^{μ} reads

$$f(\underline{x}^{\mu}) \approx a_0 \alpha^{\mu} y^{\mu} + a_1 \sum_{\nu \in \partial \underline{x}^{\mu}} \alpha^{\nu} y^{\nu} + b \approx (a_0 + a_1') (\alpha_0 + y^{\mu} \Delta \alpha) y^{\mu} + b, \tag{B2}$$

where the sum runs over the nearest neighbors of \underline{x}^{μ} . We use that all points are SV, which results from the hierarchy $a_1 \ll a_0$. Indeed, the interaction term alone is never sufficient for $\|f(\underline{x}^{\mu})\|$ to exceed one. The second equality is justified by the following reasoning. First, in the limit $\delta \rightarrow 0$, the nearest neighbors typically share the same sign, so that all the y^{ν} in the sum can be replaced by y^{μ} . a_1' is thus a_1 multiplied by the number of terms in the sum. Then, because the distribution is assumed smooth and the kernel is blind to the data structure coming from distant patterns, the SV charge may only depend on its label: $\alpha^{\mu} = \alpha_0 + y^{\mu} \Delta \alpha$. $\Delta \alpha$ is taken independent of the associated label y^{μ} , as we assume the labels to be balanced. The charge

¹⁰ For the derivation of the following scalings the notion of 'nearest neighbors' could be relaxed to include points that lie in a thin shell. In any case we assume that the number of nearest neighbors of a given point is finite.

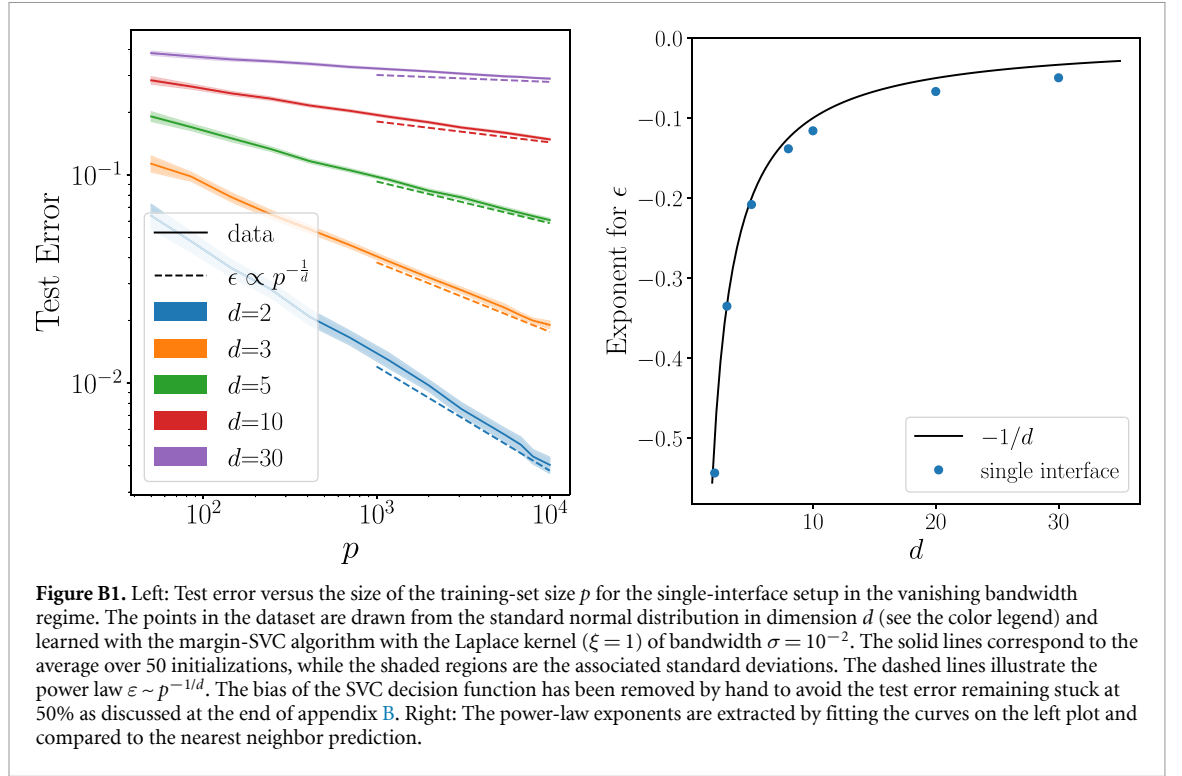


Figure B1. Left: Test error versus the size of the training-set size p for the single-interface setup in the vanishing bandwidth regime. The points in the dataset are drawn from the standard normal distribution in dimension d (see the color legend) and learned with the margin-SVC algorithm with the Laplace kernel ($\xi = 1$) of bandwidth $\sigma = 10^{-2}$. The solid lines correspond to the average over 50 initializations, while the shaded regions are the associated standard deviations. The dashed lines illustrate the power law $\epsilon \sim p^{-1/d}$. The bias of the SVC decision function has been removed by hand to avoid the test error remaining stuck at 50% as discussed at the end of appendix B. Right: The power-law exponents are extracted by fitting the curves on the left plot and compared to the nearest neighbor prediction.

conservation equation (16) implies immediately that $\Delta\alpha = -\alpha_0 \langle y \rangle$, where $\langle y \rangle = \frac{1}{p} \sum_{\mu} y^{\mu} \sim p^{-1/2}$ and imposing the condition $y^{\mu} f(\underline{x}^{\mu}) = 1$ on each points \underline{x}^{μ} yields $\alpha_0 = 1/(a_0 + a'_1)$ and $b = \langle y \rangle$.

We can now compute the test error of the SVC in the limit $\sigma \ll \delta$. The prediction on a test point \underline{x} is

$$\hat{y}(\underline{x}) \approx \text{sign} \left(a_1 \sum_{\nu \in \partial \underline{x}} \alpha^{\nu} y^{\nu} + b \right) \approx \text{sign} \left[\frac{a'_1}{a_0} y_{\text{NN}} + b \right], \tag{B3}$$

where with a slight abuse of notation we take the sum over the points \underline{x}^{ν} in the training set that are nearest neighbors of the test point \underline{x} , and y_{NN} is their label (as before, assumed to be constant among nearest neighbors). We observe two distinct behaviors according to the ratio between the bias $b = \langle y \rangle$ and the nearest-neighbor contribution a'_1 . If $\langle y \rangle \sim p^{-1/2}$ is much larger than a'_1 , the above prediction yields $\hat{y}(\underline{x}) = \text{sign}(\langle y \rangle)$ (for any \underline{x}): this estimator cannot beat a 50% accuracy. In contrast, if $\langle y \rangle$ is much smaller than a'_1 , the prediction yields $\hat{y}(\underline{x}) = \text{sign}(y_{\text{NN}})$: the classifier acts as a nearest-neighbor algorithm, and consequently its test error scales as the nearest-neighbor distance, $\epsilon \sim \delta \sim p^{-1/d}$ —namely, it is susceptible to the curse of dimensionality—as we show in figure B1.

Appendix C. Proof that power kernels are CSPD

The margin-SVC algorithm presented in section 3.4 relies on the assumption that the Gram matrix is conditionally strictly positive definite (CSPD). In this appendix, we prove that the power kernel $K(\underline{x}, \underline{x}') = -\left(\frac{\|\underline{x} - \underline{x}'\|}{\sigma}\right)^{\xi}$ indeed belongs to the CSPD class for $0 < \xi < 2$ and for any space dimension, by introducing the following definitions and theorems.

Definition: A real function k is called conditionally strictly positive definite (CSPD) in \mathbb{R}^d , if

$$\sum_{\mu=1}^p \sum_{\nu=1}^p c_{\mu} c_{\nu} k(\|\underline{x}_{\mu} - \underline{x}_{\nu}\|) > 0, \tag{C1}$$

for any set of p distinct points $\underline{x}_1, \dots, \underline{x}_p \in \mathbb{R}^d$ and any choice of p variables c_1, \dots, c_p , satisfying

$$\sum_{\mu=1}^p c^{\mu} = 0. \tag{C2}$$

Definition: A function ϕ is said to be completely monotonic in $(0, \infty)$ if it satisfies $\phi \in C^\infty(0, \infty)$ and $(-1)^n \partial^{(n)} \phi(r) \geq 0$, for all $n \in \mathbb{N}_0$ and all $r > 0$.

Theorem: Let $\phi \in C[0, \infty) \cap C^\infty(0, \infty)$. The function $k(\bullet) = \phi(\|\bullet\|^2)$ is CSPD in \mathbb{R}^d for all d , if and only if its negative derivative $-\phi'$ is completely monotonic on $(0, \infty)$ and ϕ is not a polynomial of degree at most one. A proof can be found in chapter 8 of [50].

The introductory statement arises naturally when considering the univariate function $\phi(r) = -r^{\xi/2}$ defined on \mathbb{R}_+ . Following the theorem and the definitions, one can easily show that the function $-\phi'(r) = \frac{\xi}{2} r^{\xi/2-1}$ is completely monotonic on $(0, \infty)$ for $0 \leq \xi \leq 2$. The condition that ϕ be not a polynomial of degree at most one excludes further the cases $\xi = 0$ and $\xi = 2$, which proves that the function $k(r) = -r^\xi$ is CSPD for $0 < \xi < 2$. Note that a radial kernel is defined as the multivariate function $K(\underline{x}, \underline{x}') = k(\|\underline{x} - \underline{x}'\|)$, and that if the kernel generator k is CSPD, the kernel K is also called CSPD.

Appendix D. Large σ convergence of the SVC algorithm

In section 3.4, it is loosely argued that in the limit of large σ one could replace the actual kernel $K(r/\sigma)$ by its truncated Taylor expansion $\hat{K}(r/\sigma)$. Here, we prove that in the limit $\sigma \rightarrow \infty$ the SVC solution with the truncated kernel converges to the actual SVC solution: $\{\hat{\alpha}^\mu\} \xrightarrow{\sigma \rightarrow \infty} \{\alpha^\mu\}$.

We assume that the kernel K can be written as

$$K\left(\frac{r}{\sigma}\right) = \hat{K}\left(\frac{r}{\sigma}\right) + o(\sigma^{-\xi}), \text{ with } \hat{K}\left(\frac{r}{\sigma}\right) = c_0 + c_1 \left(\frac{r}{\sigma}\right)^\xi. \quad (\text{D1})$$

For a given classification problem $\{(\underline{x}^\mu, y^\mu)\}$, the SVC algorithm converges to a set of dual variables $\{\alpha^\mu\}$, respectively $\{\hat{\alpha}^\mu\}$, provided that the associated kernel is CSPD. \hat{K} is proved to be CSPD in appendix C if $c_1 < 0$ and $0 < \xi < 2$, while K is assumed to be CSPD from the start. This condition guarantees that the Lagrangian in equation (13) defines a strictly convex problem. Rescaling the dual variables $\alpha^\mu \rightarrow \alpha^\mu / \sigma^\xi$ yields the following rescaled Lagrangians:

$$\hat{\mathcal{L}}(\alpha) = \sum_{\mu=1}^p \alpha^\mu - \frac{c_1}{2} \sum_{\mu, \nu=1}^p \alpha^\mu \alpha^\nu y^\mu y^\nu \|\underline{x}^\mu - \underline{x}^\nu\|^\xi \text{ and } \mathcal{L}(\alpha) = \hat{\mathcal{L}}(\alpha) + \epsilon(\sigma). \quad (\text{D2})$$

The rescaled solution $\{\hat{\alpha}^\mu\}$ of the maximizing problem with the Lagrangian $\hat{\mathcal{L}}$ is well defined in the limit $\sigma \rightarrow \infty$, hence the strict convexity of both Lagrangian ensures that $\{\hat{\alpha}^\mu\} \rightarrow \{\alpha^\mu\}$, when the perturbation $\epsilon(\sigma)$ vanishes.

Appendix E. The charge structure factor

The charge structure factor \tilde{Q} introduced in equation (30) is a good measure of the fluctuations in the system and, in particular, of the cutoff occurring at the scale r_c . It is argued in section 3.4 that $\tilde{Q}^2(\underline{k}_\perp) \sim \bar{\alpha}^2 p \Delta / \gamma$ at high frequencies, namely $\|\underline{k}_\perp\| > r_c^{-1}$. This scaling is verified numerically in figure E2.

The data are obtained as follows: for each $\|\underline{k}_\perp\|$, a set of $N = 2000$ random wave vectors is generated on the interface; the associated factor is computed by summing over the SV of the considered setup and then averaged. The fluctuations observed at large $\|\underline{k}_\perp\|$ decrease when N increases. The insets illustrate the expected asymptotic behavior $\tilde{Q}_\infty^2 \approx \bar{\alpha}^2 p \Delta / \gamma$, while the vertical dotted lines correspond to the typical nearest-neighbor distance r_c .

Appendix F. SVC gradient of the multiple-interface setup

Consider a setup of n (odd) interfaces separated by a distance w . We count the interfaces with the index $I = -\frac{n-1}{2}, \dots, \frac{n-1}{2}$ and set the middle interface at $x_1 = 0$, so that $x_{1,I} = Iw$. We call Δ_I the band thickness on each side of the I th interface and denote the mean SV dual variable on its left, respectively on its right, by $\bar{\alpha}_I$, respectively $\bar{\alpha}'_I$. Without loss of generality, we fix the sign of the setup by setting $y(x_1) = +1$, for $x_1 \in [0, w]$. The symmetry of the system imposes that $\Delta_{-I} = \Delta_I$ and $\bar{\alpha}'_I = \bar{\alpha}_{-I}$ for all $I > 0$, as well as that $\bar{\alpha}'_0 = \bar{\alpha}_0$.

Following the same construction as in section 3.4, in the central-limit approximation the SVC function on the point $\underline{x} = (x_1, \underline{0})$ is given by

$$f(x_1) = b - p \sigma^{-\xi} \gamma^{-d} \sum_{I=-\frac{n-1}{2}}^{\frac{n-1}{2}} (-1)^I \int_{-\Delta_I}^{\Delta_I} du \operatorname{sgn}(u) \bar{\alpha}_I(u) g(x_1 - x_{1,I} - u), \quad (\text{F1})$$

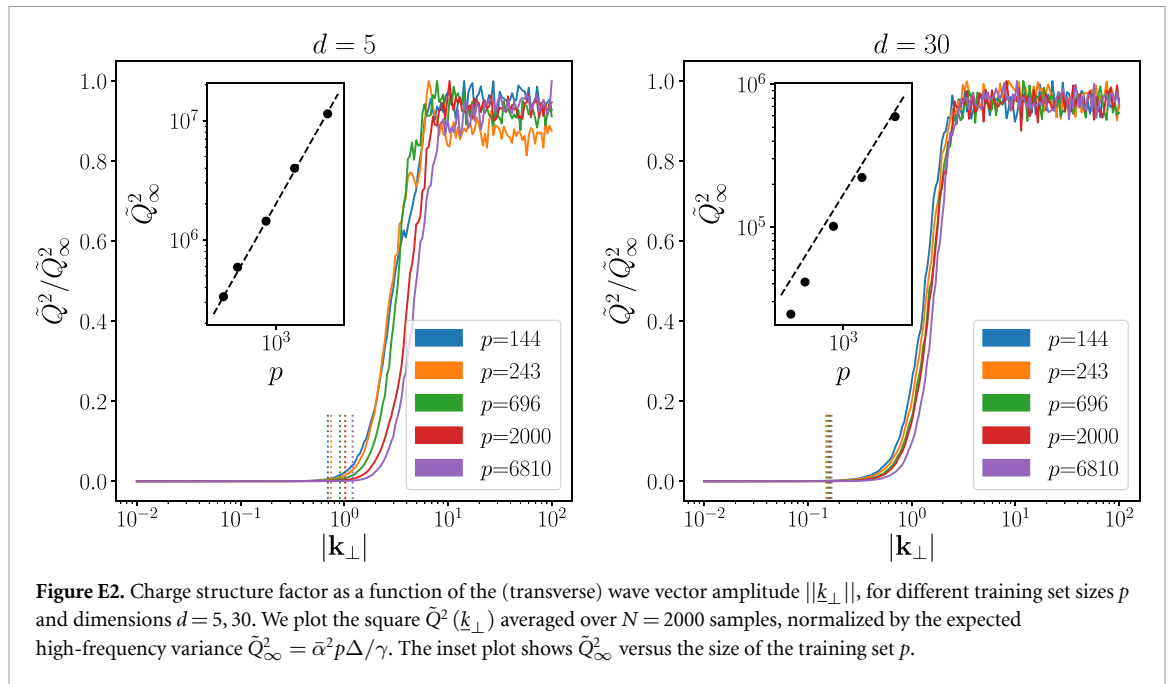


Figure E2. Charge structure factor as a function of the (transverse) wave vector amplitude $\|\mathbf{k}_\perp\|$, for different training set sizes p and dimensions $d = 5, 30$. We plot the square $\tilde{Q}^2(\mathbf{k}_\perp)$ averaged over $N = 2000$ samples, normalized by the expected high-frequency variance $\tilde{Q}_\infty^2 = \bar{\alpha}^2 p \Delta / \gamma$. The inset plot shows \tilde{Q}_∞^2 versus the size of the training set p .

where

$$g(x) = \int d\mathbf{x}_\perp (x^2 + \|\mathbf{x}_\perp\|^2)^{\xi/2} \sim \underbrace{\int_0^w dr r^{d-2} (x^2 + r^2)^{\xi/2}}_{g^s(x)} + \underbrace{\int_w^\gamma dr r^{d-2} (x^2 + r^2)^{\xi/2}}_{g^l(x)}, \quad (\text{F2})$$

and $\bar{\alpha}_I(u) = \alpha_I$, respectively $\bar{\alpha}_I(u) = \alpha'_I$, for $u < 0$, respectively $u > 0$. By symmetry, the target function is of the form¹¹

$$f(x_1) = \beta_1 x_1 + \dots + \beta_i x_1^i + \dots + \beta_n x_1^n + O(x_1)^{n+2}, \quad (\text{F3})$$

with i only running over odd indices. Imposing that the target function is zero on each interface, all coefficients can be expressed in terms of β_n : $\beta_i = b_i w^{n-i} \beta_n$, where $b_i \sim \mathcal{O}(1)$. Similarly, the SVC condition that $\partial_{x_1} f(x_I) \Delta_I$ is identical on each interface allows us to relate all band thicknesses to Δ_0 : $\Delta_I = d_I \Delta_0$, with $d_I \sim \mathcal{O}(1)$. Denote by α , $\delta\alpha$ and Δ respectively the typical values of $(\alpha'_I + \alpha_I)/2$, $\|\alpha'_I - \alpha_I\|/2$ and Δ_I . One can obtain the β coefficients associated with equation (F1) by differentiating it, namely


$$\beta_i = \frac{f^{(i)}(0)}{i!} = \underbrace{\sum_{j=0}^{\infty} g_{i+1+2j}^L T_j}_{\beta_i^L} + \underbrace{g_{i+1}^S w^{d+\xi-2-i} \Delta^2 \alpha}_{\beta_i^S}, \quad (\text{F4})$$

where $g_i^L \sim \gamma^{\xi-1-i} \sigma^{-\xi}$, $g_i^S \sim \gamma^{-d} \sigma^{-\xi}$ and $T_j \sim \mathcal{O}(\Delta^2 w^{2j} \alpha) + \mathcal{O}(\Delta w^{2j+1} \delta\alpha)$. The constrained scaling between the β coefficients forces the terms of index $j = 0, \dots, n-i$ in the sum defining β_i^L to cancel each other up to higher order. In particular, $\delta\alpha \sim \alpha \Delta / w$, and $\beta_n^L \sim w^2 \Delta^2 \alpha$. Eventually, the scaling of the gradient depends on the hierarchy between β_n^L and β_n^S :

$$1 \sim \Delta \partial_{x_1} f \sim \begin{cases} p \alpha \left(\frac{\gamma}{\sigma}\right)^\xi \left(\frac{\Delta}{\gamma}\right)^3 \left(\frac{w}{\gamma}\right)^{n+1}, & \text{if } n \geq d + \xi - 4, \\ p \alpha \left(\frac{w}{\sigma}\right)^\xi \left(\frac{\Delta}{w}\right)^3 \left(\frac{w}{\gamma}\right)^d, & \text{if } n \leq d + \xi - 4. \end{cases} \quad (\text{F5})$$

Also, if $n > d + \xi - 1$, when computing β_n^S , divergences will occur while differentiating g^S . This sets an upper bound on the number of interfaces we can consider without considering microscopic effects on the gradient. For an even number of interfaces, a similar discussion holds with the difference that n should be replaced by $n - 1$ in the above expressions. Finally, the resulting scaling of the usual observables are given in section 3.5.

¹¹ The shift constant β_0 is discarded because of the bias freedom in equation (F1).



Appendix G. SVC scaling with the Matérn kernel

Results of simulations on the single-interface setup with Matérn kernels are shown in figure G3, for several parameters ν and several dimensions d . All the curves follow the scalings predicted in section 3.4.

Appendix H. Numerical definition of the scale r_c

In section 3.4, the scale r_c is defined geometrically as the distance between nearest support vectors. The numerical definition of r_c is different as it aims to confirm the ‘minimal disturbance hypothesis’ presented in the note at the end of section 3.4. From this point of view, the scale r_c is also the scale below which the charges of two SVs are not correlated. To test this idea, the solution of the margin-SVC problem is computed once for a benchmark training set and a second time for the same training set with one additional point close enough to the interface to be a SV. We then calculate the cumulative distribution of the charge variations $d\alpha^\mu = |\alpha^\mu - \alpha'^\mu|$ as function of their distance from the additional point r^μ . The resulting distribution is displayed in figure H4 for multiple realizations of the single-interface setup with $d = 5$ and $p = 6810$. The scale r_c is then defined as the distance for which the cumulative distribution reaches a given value $C < 1$. The particular choice of C does not alter the power-law behavior.

Appendix I. Scaling arguments for the spherical setup

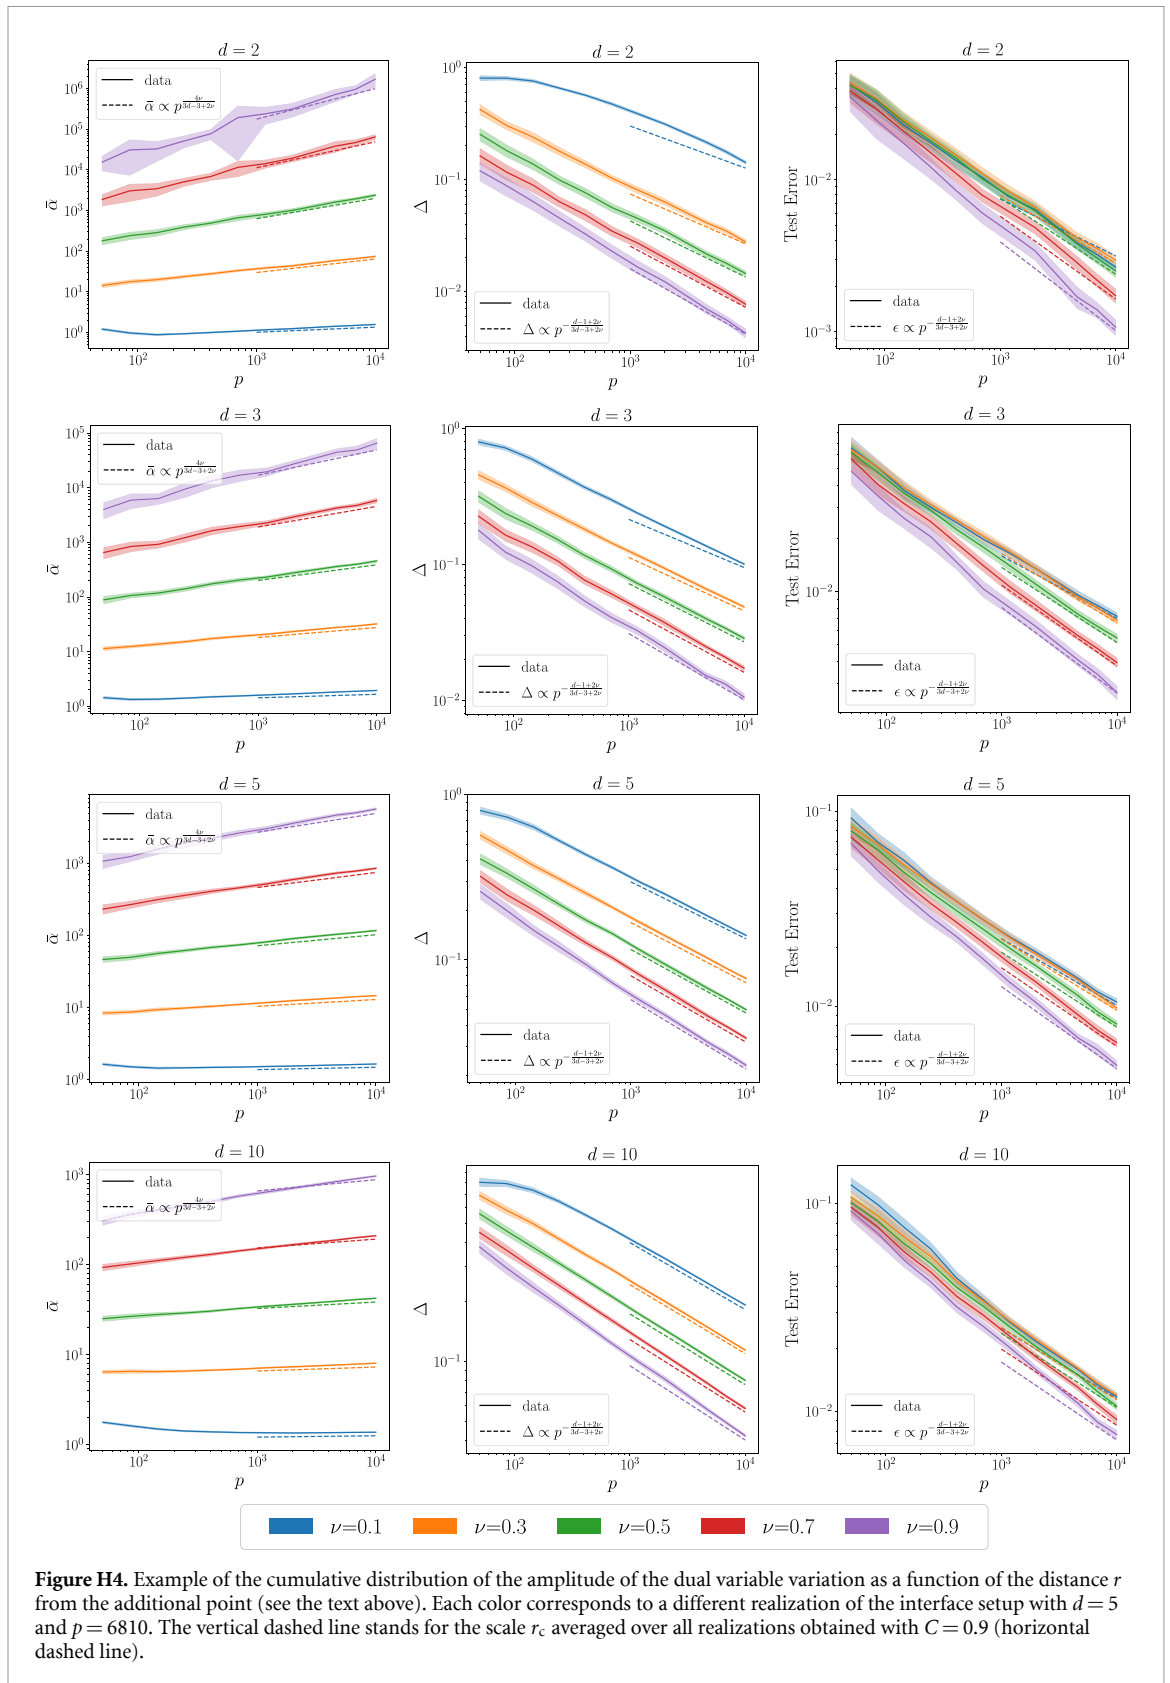
In this appendix, we sketch how the scaling relations in section 3.4 may be derived for the spherical interface setup discussed in section 4, where the label only depends on the norm of the vector: $y(\underline{x}) = \text{sign}(\|\underline{x}\| - R)$, with R the radius of the sphere. In the same line as for the linear interface, it is assumed that all SVs lie within a shell of thickness $\Delta \ll R$ around the interface. The decision function on the vector \underline{x} ,

$$f(\underline{x}) = b - \sum_{\mu=1}^p \alpha^\mu y^\mu \left(\frac{\|\underline{x} - \underline{x}^\mu\|}{\sigma} \right)^\xi, \tag{11}$$

is better apprehended in a Cartesian frame such that $\underline{x} = (x_1 = \|\underline{x}\|, \underline{0})$, which requires us to rotate all SVs: $\underline{x}^\mu \rightarrow x'^\mu = \mathcal{R}\underline{x}^\mu$. In the large p limit, the charge conservation, $Q = \sum_{\mu=1}^p \alpha^\mu y^\mu = 0$, reads

$$0 = \int d^d \underline{x} \rho(\underline{x}) \alpha(\underline{x}) y(\underline{x}) = S_{d-1} \int_{-\Delta}^{\Delta} du (R+u)^{d-1} \rho(R+u) \alpha(R+u) y(R+u). \tag{12}$$

Spherical coordinates are used in the second equality: the angular variables trivially integrate to the unit $(d-1)$ -sphere surface, S_{d-1} , and the variable $u = r - R$ is used instead of the radius $r = \|\underline{x}\|$. For simplicity, we assume that the population distribution is radial: $\rho(\underline{x}) = \rho(r)$. Were this not the case, the angular integral would merely yield a different finite factor.



As for the linear interface, the first scaling relation stems from the condition $\Delta \cdot \partial_{x_{||}} f(\underline{x}^*) \sim 1$, for any \underline{x}^* lying on the spherical interface. According to the change of frame introduced above, the relevant direction correspond to the first coordinate, namely $x_{||} = x_1$. The gradient expression (23) can thus be expressed as an integral in spherical coordinate with the north pole $\underline{x}^* = (R, 0)$:

$$\partial_{x_{\parallel}} f(\underline{x}^*) = \xi \sigma^{-\xi} p S_{d-2} \int_{-\Delta}^{\Delta} du (R+u)^{d-1} \int_0^{\pi} d\phi \sin^{d-2} \phi \rho(R+u) \alpha(R+u) \gamma(R+u) I(u, \phi), \quad (13)$$

where the vector of integration norm is $r = R + u$ and its angle with respect to the north pole is ϕ . All other angles simply integrate to the $(d-2)$ -sphere surface, S_{d-2} , since they do not contribute to the integrand

$$I(u, \phi) = (x_1 - x_1^*) \|\underline{x} - \underline{x}^*\|^{\xi-2} = a_0(\phi) + a_1(\phi)u + \mathcal{O}(u^2), \quad (14)$$

with

$$a_0(\phi) = \frac{1}{2R} \left[2R^2(1 - \cos \phi) \right]^{\xi/2} \quad \text{and} \quad a_1(\phi) = \left[1 - \frac{\xi}{2}(1 - \cos \phi) \right] \left[2R^2(1 - \cos \phi) \right]^{\xi/2-1}. \quad (15)$$

The leading order contribution a_0 vanishes because of the charge conservation (equation (12)), so the gradient reads

$$\partial_{x_{\parallel}} f(\underline{x}^*) \sim p \int_{-\Delta}^{\Delta} du (R+u)^{d-1} \rho(R+u) \alpha(R+u) \gamma(R+u) u \int_0^{\pi} d\phi \sin^{d-2} \phi a_1(\phi) \sim p \Delta^2 \bar{\alpha} \quad (16)$$

and the second scaling relation $p \bar{\alpha} \Delta^3 \sim 1$ is identical as for the stripe model. Since the other relations are obtained from local arguments, they are independent of the global shape of the classification task.

The scaling laws for the spherical model are thus also given by equations (33) and (34).

ORCID iD

Jonas Paccolat  <https://orcid.org/0000-0002-0348-9986>

References

- [1] Hestness J, Narang S, Ardalani N, Diamos G F, Jun H, Kianinejad H, Ali Patwary M M, Yang Y and Zhou Y (arXiv:1712.00409)
- [2] Spigler S, Geiger M and Wyart M 2019 Asymptotic learning curves of kernel methods: empirical data vs teacher-student paradigm (arXiv:1905.10843)
- [3] von Luxburg U and Bousquet O 2004 Distance-based classification with Lipschitz functions *J. Mach. Learn. Res.* **5** 669–95
- [4] Bach F 2017 Breaking the curse of dimensionality with convex neural networks *J. Mach. Learn. Res.* **18** 629–81
- [5] Mallat S 2016 Understanding deep convolutional networks *Phil. Trans. R. Soc. A* **374** 20150203
- [6] Neyshabur B, Tomioka R, Salakhutdinov R and Srebro N 2017 Geometry of optimization and implicit regularization in deep learning (arXiv:1705.03071)
- [7] Neyshabur B, Zhiyuan Li, Bhojanapalli S, LeCun Y and Srebro N 2018 Towards understanding the role of over-parametrization in generalization of neural networks (arXiv:1805.12076)
- [8] Bansal Y, Advani M, Cox D D and Saxe A M 2018 Minnorm training: an algorithm for training over-parameterized deep neural networks (arXiv:1806.00730v2)
- [9] Advani M S and Saxe A M 2017 High-dimensional dynamics of generalization error in neural networks (arXiv:1710.03667)
- [10] Spigler S, Geiger M, d’Ascoli S, Sagun L, Biroli G and Wyart M 2019 A jamming transition from under- to over-parametrization affects generalization in deep learning *J. Phys. A: Math. Theor.* **52** 474001
- [11] Geiger M, Jacot A, Spigler S, Gabriel F, Sagun L, d’Ascoli S, Biroli G, Hongler C and Wyart M 2020 Scaling description of generalization with number of parameters in deep learning *J. Stat. Mech.: Theory Exp.* **2020** 023401
- [12] Jacot A, Gabriel F and Hongler C 2018 Neural tangent kernel: convergence and generalization in neural networks. *Advances in Neural Information Processing Systems* pp 8571–80
- [13] Bruna J and Mallat S 2013 Invariant scattering convolution networks *IEEE Trans. Pattern Anal. Mach. Intell.* **35** 1872–86
- [14] Arora S, Du S S, Wei H, Zhiyuan Li, Salakhutdinov R and Wang R 2019 On exact computation with an infinitely wide neural net (arXiv:1904.11955)
- [15] Maierov V 2006 Approximation by neural networks and learning theory *J. Complexity* **22** 102–17
- [16] Costa J A and Hero A O 2004 Learning intrinsic dimension and intrinsic entropy of high-dimensional datasets *2004 12th European Signal Conf. (IEEE)* pp 369–72
- [17] Hein M and Audibert J-Y 2005 Intrinsic dimensionality estimation of submanifolds in \mathbb{R}^d *Proc. 22nd Int. Conf. on Machine Learning (ACM)* pp 289–96
- [18] Rozza A, Lombardi G, Ceruti C, Casiraghi E and Campadelli P 2012 Novel high intrinsic dimensionality estimators *Mach. Learn.* **89** 37–65
- [19] Facco E, d’Errico M, Rodriguez A and Laio A 2017 Estimating the intrinsic dimension of datasets by a minimal neighborhood information *Sci. Rep.* **7** 12140
- [20] Goldt S, Mézard M, Krzakala F and Lenka Z 2019 Modelling the influence of data structure on learning in neural networks (arXiv:1909.11500)
- [21] Paccolat J, Petrini L, Geiger M, Tyloso K and Wyart M 2020 Geometric compression of invariant manifolds in neural nets (<https://arxiv.org/pdf/2007.11471.pdf>)
- [22] Shwartz-Ziv R and Tishby N 2017 Opening the black box of deep neural networks via information (arXiv:1703.00810)
- [23] Sollich P 1999 Learning curves for Gaussian processes *Advances in Neural Information Processing Systems* pp 344–50
- [24] Sollich P 2002 Gaussian process regression with mismatched models *Advances in Neural Information Processing Systems* pp 519–26
- [25] Bordelon B, Canatar A and Pehlevan C 2020 Spectrum dependent learning curves in kernel regression and wide neural networks (arXiv:2002.02561)

- [26] Stein M L 1999 Predicting random fields with increasing dense observations *Ann. Appl. Probab.* **9** 242–73
- [27] Fischer S and Steinwart I 2017 Sobolev norm learning rates for regularized least-squares algorithm (arXiv:1702.07254)
- [28] Caponnetto A and Ernesto D V 2007 Optimal rates for the regularized least-squares algorithm *Found. Comput. Math.* **7** 331–68
- [29] 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* pp 8114–24
- [30] Gardner E 1988 The space of interactions in neural network models *J. Phys. A: Math. Gen.* **21** 257
- [31] Dietrich R, Opper M and Sompolinsky H 1999 Statistical mechanics of support vector networks *Phys. Rev. Lett.* **82** 2975
- [32] Montanari A, Ruan F, Sohn Y and Yan J 2019 The generalization error of max-margin linear classifiers: high-dimensional asymptotics in the overparametrized regime (arXiv:1911.01544)
- [33] Gerace F, Loureiro B, Krzakala F, Mézard M and Lenka Z 2020 Generalisation error in learning with random features and the hidden manifold model (arXiv:2002.09339)
- [34] d’Ascoli S, Refinetti M, Biroli G and Krzakala F 2020 Double trouble in double descent: bias and variance(s) in the lazy regime (arXiv:2003.01054)
- [35] Engel A and Christian V den B 2001 *Statistical Mechanics of Learning* (Cambridge: Cambridge University Press)
- [36] Bartlett P L and Mendelson S 2002 Rademacher and gaussian complexities: risk bounds and structural results *J. Mach. Learn. Res.* **3** 463–82
- [37] Boucheron S, Bousquet O and Lugosi G 2005 Theory of classification: a survey of some recent advances *ESAIM: Probab. Stat.* **9** 323–75
- [38] Saad D and Solla S A 1995 On-line learning in soft committee machines *Phys. Rev. E* **52** 4225
- [39] Monasson R and Zecchina R 1995 Weight space structure and internal representations: a direct approach to learning and generalization in multilayer neural networks *Phys. Rev. Lett.* **75** 2432
- [40] Opper M and Saad D 2001 *Advanced Mean Field Methods: Theory and Practice* (Cambridge, MA: MIT Press)
- [41] Aubin B, Maillard A, Krzakala F, Macris N and Zdeborová L et al 2018 The committee machine: computational to statistical gaps in learning a two-layers neural network *Advances in Neural Information Processing Systems* pp 3223–34
- [42] Franz S, Hwang S and Urbani P 2018 Jamming in multilayer supervised learning models (arXiv:1809.09945)
- [43] Scholkopf B and Smola A J 2001 *Learning with Kernels: Support Vector Machines, Regularization, Optimization and Beyond* (Cambridge, MA: MIT Press)
- [44] Kanagawa M, Hennig P, Sejdinovic D and Sriperumbudur B K 2018 Gaussian processes and kernel methods: a review on connections and equivalences (arXiv:1807.02582)
- [45] Smola A J, Schölkopf B and Müller K-R 1998 The connection between regularization operators and support vector kernels *Neural Netw.* **11** 637–49
- [46] Paccolat J, Petrini L, Geiger M, Tyloo K and Wyart M 2020 Compressing invariant manifolds in neural nets (arXiv:2007.11471)
- [47] Grafakos L and Teschl G 2013 On fourier transforms of radial functions and distributions *J. Fourier Appl.* **19** 167–79
- [48] Estrada R 2014 On radial functions and distributions and their fourier transforms *J. Fourier Anal. Appl.* **20** 301–20
- [49] Erdélyi A 1955 Asymptotic representations of fourier integrals and the method of stationary phase *J. Soc. Ind. Appl. Math.* **3** 17–27
- [50] Wendland H 2004 *Scattered Data Approximation* vol 17 (Cambridge: Cambridge University Press)