

Statistical learning theory of structured dataMauro Pastore , Pietro Rotondo , Vittorio Erba , and Marco Gherardi **Dipartimento di Fisica, Università degli Studi di Milano and INFN, Via Celoria 16, I-20133 Milan, Italy*

(Received 24 May 2020; accepted 13 August 2020; published 14 September 2020)

The traditional approach of statistical physics to supervised learning routinely assumes unrealistic generative models for the data: Usually inputs are independent random variables, uncorrelated with their labels. Only recently, statistical physicists started to explore more complex forms of data, such as equally labeled points lying on (possibly low-dimensional) object manifolds. Here we provide a bridge between this recently established research area and the framework of statistical learning theory, a branch of mathematics devoted to inference in machine learning. The overarching motivation is the inadequacy of the classic rigorous results in explaining the remarkable generalization properties of deep learning. We propose a way to integrate physical models of data into statistical learning theory and address, with both combinatorial and statistical mechanics methods, the computation of the Vapnik-Chervonenkis entropy, which counts the number of different binary classifications compatible with the loss class. As a proof of concept, we focus on kernel machines and on two simple realizations of data structure introduced in recent physics literature: k -dimensional simplexes with prescribed geometric relations and spherical manifolds (equivalent to margin classification). Entropy, contrary to what happens for unstructured data, is nonmonotonic in the sample size, in contrast with the rigorous bounds. Moreover, data structure induces a transition beyond the storage capacity, which we advocate as a proxy of the nonmonotonicity, and ultimately a cue of low generalization error. The identification of a synaptic volume vanishing at the transition allows a quantification of the impact of data structure within replica theory, applicable in cases where combinatorial methods are not available, as we demonstrate for margin learning.

DOI: [10.1103/PhysRevE.102.032119](https://doi.org/10.1103/PhysRevE.102.032119)**I. INTRODUCTION**

The idea of investigating machine learning within the tools provided by the statistical physics of disordered system is more than 30 years old, starting with the seminal papers by Amit, Gutfreund and Sompolinsky [1,2] on the Hopfield model, and with Gardner's replica analysis of the Perceptron architecture [3,4]. Many of the results produced in this field have been obtained under the restrictive and unrealistic hypothesis that the inputs of the training set were independent identically distributed random variables with no correlation with their labels. Only quite recently, physicists working in this field are starting to probe the impact of more realistic generative models of synthetic data on the available theoretical frameworks. Sompolinsky and collaborators investigated the problem of the linear classification of perceptual manifolds [5,6] and provided a first quantitative measurement of the ability to support the classification of object manifolds in deep neural networks [7]. Mézard suggested that hierarchical architectures with hidden layers naturally emerge in the context of Hopfield models, assuming that the training patterns are structured as superpositions of a given set of random features [8], a common property of empirical data [9,10]. Zdeborová and collaborators provided exact results for the generalization error within the replica approach for two different scenarios of synthetic data: Random features and the hidden manifold

model [11,12]. One of the motivations behind these choices is the observation that many machine learning data sets, or their representations within deep networks, lie on the surface of low-dimensional manifolds, as also verified often in practice by measuring their so-called intrinsic dimension [7,13–16]. More in general, a generative model with a factorized joint probability distribution of the inputs and their corresponding labels is expected to be unrealistic, with respect to the benchmark data sets commonly used in machine learning (e.g., MNIST, CIFAR-10, or Imagenet). Intuitively, one expects there to be a notion of similarity among inputs that constrains similar inputs to have the same label. This regularity is expected to be related to the problem of generalization, i.e., the ability of a classifier to correctly classify inputs beyond the data set used for training.

The results obtained in the statistical physics framework address the typical case performance. In contrast, statistical learning theory (SLT) [17], a successful mathematical framework in the theory of machine learning, follows the tradition of computer science of establishing worst-case bounds. This difference in scope made it difficult, for physicists and computer scientists alike, to work towards inter-disciplinary results, and few examples of cross-fertilization are found in the literature. Statistical learning theory is the branch of mathematics and computer science that studies inference, or the problem of generating models starting from data [18]. It provides formal definitions for words like “generalization” or “overfitting,” and it is ultimately designed to evaluate the performance of learning algorithms. As such, it represents

*marco.gherardi@mi.infn.it

the ideal framework to study the problem of generalization in deep learning. Unfortunately, in spite of its elegance, the insight it provides into the impressive generalization abilities of present deep learning models is poor. The main product of the theory in this setting is a set of upper bounds on the generalization error (which roughly counts the average number of errors made on the test set). These upper bounds in many cases turn out to be too loose to be useful [19–22].

The main drawback of this class of bounds is generally recognized to be their being distribution independent, meaning that they hold for any probability distribution over inputs and labels of the data set, and for all models of the chosen hypothesis class. Substantial effort is being put, within statistical learning theory, to overcome these shortcomings and formulate rigorous data-dependent results [21,23–25]. The Vapnik-Chervonenkis (VC) entropy is a way to establish distribution-dependent, and hopefully tighter, bounds to the generalization error [18]. Informally, the VC entropy measures the number of different ways a given class of functions can classify the inputs of the training set. Unfortunately it is usually very difficult to compute explicitly. Kernel architectures represent a notable exception. Their VC entropy has been evaluated analytically in a remarkable paper by Cover long ago [26], under very mild hypotheses on the probability distribution of the inputs. The explicit calculation shows, however, that knowing the VC entropy does not improve significantly the standard bound obtained using the growth function [27]. In fact, both quantities scale logarithmically with the size of the training set, and depend linearly on the VC dimension, a well known measure of model complexity.

Our goal here is to show how the concept of data structure, as it is emerging in the physics literature, can be addressed within statistical learning theory, thereby providing a bridge between the two viewpoints. This bridge immediately allows a quantification of the generalization capabilities of simple hypothesis classes, which shows how severely loose the classic rigorous bounds in SLT are. Concretely, we investigate the finite-size and asymptotic behavior of the VC entropy of kernel machines by using both combinatorial and replica techniques. While replica theory is well established in the statistical mechanics of neural networks, combinatorial tools, though certainly not foreign to modern statistical mechanics [28–31], have been developed only very recently for what concerns the role of data structure in machine learning [5,32]. We concentrate on two simple models of data structure, or “object manifolds”: (1) k -dimensional simplexes with prescribed geometric relations and (2) spherical manifolds, which are equivalent to classify unstructured data points with margin (and are related to support vector machines [33]). These models are not new and have already received attention for their being general enough to provide insight, but simple enough to allow full analytical treatment.

The manuscript is organised as follows: In Sec. II A we review the main definitions and basic results of statistical learning theory. In Sec. II B we recall Cover’s combinatorial result on the VC entropy of kernel architectures. In Sec. II C we discuss how data structure can be taken into account in SLT and define the two synthetic data ensembles we use in the following. In Sec. III we first recall the extension of Cover’s combinatorial technique to structured data that was introduced

in Ref. [32], and then we establish the asymptotic behavior of the VC entropy for the first ensemble (simplexes). These results substantially expand those of Ref. [32], where the link with SLT was not discussed, as well as those of Ref. [34], where the discussion on generalization was heuristic and mainly a motivation for considering structured data. After introducing an asymptotic method based on analytic combinatorics in Secs. III A and III B, we use it to show that the VC entropy is nonmonotonic in the load in Sec. III C. In Sec. III D we describe a satisfiability transition that is brought about by data structure [35], and we further analyze it in Secs. III E and III F. In Sec. IV we introduce a synaptic volume that monitors the behavior of the VC entropy in the thermodynamic limit. This is particularly useful for those object manifolds for which the combinatorial method is not yet available, which includes our second data ensemble (spherical manifolds). We present calculations for the annealed (in Sec. IV A) and quenched (in Sec. IV B) averages (in the replica symmetric and one-step replica symmetry breaking ansätze) for the case of two-dimensional simplexes, and for spherical manifolds in Sec. IV C.

II. TAKING DATA STRUCTURE INTO ACCOUNT IN STATISTICAL LEARNING THEORY

A. Basic results in statistical learning theory

In this section we recall the basic facts of SLT, mostly following the exposition of Ref. [18]. We restrict ourselves to binary classification problems, in which the goal is to find a function g mapping the input space \mathcal{X} to the output space $\mathcal{Y} = \{+1, -1\}$. Each pair $Z^\mu = (X^\mu, Y^\mu)$ (with $\mu = 1, \dots, p$) in the training set $Z_p = (Z^1, \dots, Z^p)$ is drawn by the unknown joint probability distribution $P_{\mathcal{X}, \mathcal{Y}}(X, Y)$. A map g between the set of inputs $X_p = \{X^\mu\}_\mu$ and $\{+1, -1\}$ is called a dichotomy of X_p . The criterion to choose g is the minimization of the risk

$$R(g) = \langle \mathbb{1}_{g(X) \neq Y} \rangle_p, \quad (1)$$

which is the probability of error. Ideally, we should look for $\inf_g R(g)$ over all the possible g ’s. Since P is unknown, the best we can do is to consider the empirical risk

$$R_p(g) = \frac{1}{p} \sum_{\mu=1}^p \mathbb{1}_{g(X^\mu) \neq Y^\mu} \quad (2)$$

and limit the search within a specific hypothesis class \mathcal{G} to prevent overfitting. A dichotomy $g \in \mathcal{G}$ is called realizable. The output of a learning algorithm is a function g_p that depends on the data Z_p . The goodness of the choice of g_p can be measured by its generalization error $\epsilon_{\text{gen}}(g_p)$, where

$$\epsilon_{\text{gen}}(g) = R(g) - R_p(g). \quad (3)$$

Notice that $\epsilon_{\text{gen}}(g) \leq 1$. In practice, R_p is evaluated on the training set and R is estimated on a test set [36]. One of the primary goals of SLT is to establish rigorous bounds on the generalization error.

A complementary description of risk minimization within a class \mathcal{G} is given through the definition of the loss class \mathcal{L} :

$$\mathcal{L} = \{\ell_g: (X, Y) \mapsto \mathbb{1}_{g(X) \neq Y}, g \in \mathcal{G}\}. \quad (4)$$

To each $g \in \mathcal{G}$, we associate a function ℓ_g such that $\ell_g((x, y)) = 1$ if $g(x) \neq y$, and is zero otherwise. While elements of \mathcal{G} take values in $\{+1, -1\}$, those of \mathcal{L} have range $\{0, 1\}$. ℓ_{g_p} can be used to count the number of errors made on the training set by the function g_p . Given a loss class \mathcal{L} , we can consider its projection on the sample Z_p , by defining

$$\mathcal{L}_{Z_p} = \{(\ell(Z^1), \ell(Z^2), \dots, \ell(Z^p)) : \ell \in \mathcal{L}\}. \quad (5)$$

This is the set of all possible ways that a function in \mathcal{G} can correctly or incorrectly classify each sample in Z_p . Importantly, \mathcal{L}_{Z_p} can be interpreted as the set of all classifications of the points in X_p that can be realized by the model, i.e., the set of all (Y^1, \dots, Y^p) such that there exists a $g \in \mathcal{G}$ such that $Y^\mu = g(X^\mu)$ for all μ . This representation reveals a useful bijection between \mathcal{L}_{Z_p} and the set of realizable dichotomies.

A key quantity in SLT is the Vapnik-Chervonenkis (VC) entropy $\mathcal{H}_{\mathcal{L}}(Z_p)$, which measures the size of \mathcal{L}_{Z_p} :

$$\mathcal{H}_{\mathcal{L}}(Z_p) = \log |\mathcal{L}_{Z_p}|. \quad (6)$$

By virtue of the bijection discussed above, valid when $\mathcal{Y} = \{+1, -1\}$, $\mathcal{H}_{\mathcal{L}}(Z_p)$ can be defined equivalently as

$$\mathcal{H}_{\mathcal{L}}(Z_p) = \log \mathcal{N}_{\mathcal{G}}(X_p), \quad (7)$$

where $\mathcal{N}_{\mathcal{G}}(X_p)$ is the number of dichotomies of the set X_p realizable by \mathcal{G} . The VC entropy controls a rigorous upper bound to the generalization error (see Ref. [18] for a pedagogical proof):

Theorem 1. For any $0 < \delta \leq \delta_{\max} = \min(1, 2e^{\mathcal{H}_{\mathcal{L}}(2p)})$, with probability at least $1 - \delta$,

$$\forall g \in \mathcal{G}, \quad \epsilon_{\text{gen}}(g) \leq 2\sqrt{2 \frac{\mathcal{H}_{\mathcal{L}}(2p) + \log \frac{2}{\delta}}{p}}, \quad (8)$$

where the annealed VC entropy $\mathcal{H}_{\mathcal{L}}(p)$ is defined as

$$\mathcal{H}_{\mathcal{L}}(p) = \log \langle \mathcal{N}_{\mathcal{L}}(Z_p) \rangle \quad (9)$$

and $\langle \cdot \rangle$ is the average over the joint probability distribution $P_{\mathcal{X}, \mathcal{Y}}$ of the training set.

In the following, we will be referring to the annealed VC entropy, Eq. (9), simply as ‘‘VC entropy,’’ the average over the distribution of the training set being understood. It is worth noticing that this version of the theorem is slightly different than that presented in Ref. [18], where the upper bound δ_{\max} on δ is not mentioned. The reason is that, in that context, the VC entropy is tacitly assumed to be larger than zero. However, as we will show in the case of structured data, this is not necessarily true. It is instructive to retrace the very last part of the proof of Theorem 1. The statement, that $\epsilon_{\text{gen}}(g) \leq 2\sqrt{2 \frac{\mathcal{H}_{\mathcal{L}}(2p) + \log \frac{2}{\delta}}{p}}$ with probability at least $1 - \delta$, follows from the standard form of Hoeffding’s inequality, which bounds the probability of having generalization error larger than ϵ as $P[\epsilon_{\text{gen}}(g) > \epsilon] \leq 2e^{\mathcal{H}_{\mathcal{L}}(2p) - p\epsilon^2/8}$. The statement of Theorem 1 is obtained by using a basic result, known as inversion property, which holds for a generic random variable X : If $P(X > t) \leq F(t)$, then with probability at least $1 - \delta$, $X \leq F^{-1}(\delta)$. In our case, this amounts to define $\delta = 2e^{\mathcal{H}_{\mathcal{L}}(2p) - p\epsilon^2/8}$. For negative values of the VC entropy, the only admissible values for δ are those for which $\epsilon^2 = 8 \frac{\mathcal{H}_{\mathcal{L}}(p) + \log(2/\delta)}{p} \geq 0$, thus providing the upper bound $\delta_{\max} =$

$\min(1, 2e^{\mathcal{H}_{\mathcal{L}}(2p)})$. In principle one could state the theorem directly in the form of Hoeffding’s inequality, but we rather keep it in its classic form for recognizability and restrict the domain of δ .

Unfortunately, direct computation of the annealed VC entropy is unfeasible in most cases. For this reason, a main goal of SLT is to construct more tractable upper bounds to the VC entropy. The classic example is based on the Vapnik-Chervonenkis dimension, which is a scalar metric of the expressivity of a given hypothesis class \mathcal{G} . More formally, the VC dimension d_{VC} of a class \mathcal{G} is the largest integer such that there exists at least one set of d_{VC} inputs $X_{d_{\text{VC}}}$ such that

$$\mathcal{N}_{\mathcal{G}}(X_{d_{\text{VC}}}) = 2^{d_{\text{VC}}} \quad (10)$$

(i.e., the class \mathcal{G} realizes all possible dichotomies of the inputs). With this definition, it can be proved (see again Ref. [18]) that

$$\mathcal{H}_{\mathcal{L}}(p) \leq d_{\text{VC}} \log \left(\frac{ep}{d_{\text{VC}}} \right). \quad (11)$$

Hence, a corollary of Theorem 1 is the well-known upper bound first obtained by Vapnik: If the class \mathcal{G} has finite VC dimension d_{VC} , then, with probability at least $1 - \delta$,

$$\forall g \in \mathcal{G}, \quad \epsilon_{\text{gen}}(g) \leq 2\sqrt{2 \frac{d_{\text{VC}} \log \left(\frac{2ep}{d_{\text{VC}}} \right) + \log \frac{2}{\delta}}{p}}. \quad (12)$$

A crucial property of this elegant result is its being distribution independent, meaning that the bound is uniform in the function g , and does not depend on the particular problem at hand. Owing to its universality, the bound is often too loose for most practical applications [21]. Let us consider for instance a deep neural network with a number of weights $w = 10^6 - 10^9$. In this case the VC dimension is of order $d_{\text{VC}} \sim w \log w$ [37]. When the typical size of the data set is $p = 10^4 - 10^6$, as is often the case in practice, it is evident that bounds such as the one in Eq. (12) do not offer any insight on the generalization performance of deep neural networks. Indeed, one of the main pursuits of contemporary SLT is to provide better results on the generalization error, going beyond distribution independent bounds. Several strategies have been proposed, advocating the importance of considering data-dependent hypothesis classes [25] and data-dependent measures of complexity (such as the Rademacher complexity [38], which was recently connected to the statistical mechanics of disordered systems [39]), also in relation to the original concept of VC entropy itself [40].

B. Vapnik-Chervonenkis entropy of kernel machines

As mentioned above, in most cases it is not possible to compute the VC entropy directly. However, kernel machines are a notable exception: Their VC entropy was computed a half century ago by Cover [26]. Kernel architectures provide a special realization of one-hidden layer neural networks and are at the core of the idea of support vector machines. In these machines, one defines *a priori* a kernel function $\phi: \mathbb{R}^n \rightarrow \mathbb{R}^d$ that maps n -dimensional inputs to a d -dimensional feature space. One of the simplest realizations of such maps is a quadratic polynomial kernel, such that each input X is mapped

on a $d(d+1)/2$ -dimensional feature space via a kernel $\phi^{(2)}$ with components $\phi_{ij}^{(2)} = X_i X_j$, $\forall i \leq j$. The map from feature space to the space of labels is realized by a linear separator:

$$Y = \text{sgn}(W \cdot X), \quad (13)$$

where the weight vector, $W \in \mathbb{R}^n$, is the set of learnable parameters.

Cover's theorem is a function counting theorem: It computes the number of dichotomies $\mathcal{N}_\phi(X_p)$ of this function class, the logarithm of which is the VC entropy. It is simpler to state Cover's theorem for linear separators, i.e., for $d = n$ and $\phi = 1$; the realizable dichotomies in this case are called linearly realizable. We comment below on the extension to general ϕ . The key idea behind the theorem is twofold: (1) under a weak condition on the inputs X_p , the number of dichotomies $\mathcal{N}_1(X_p)$ is a function solely of the dimension n and the number of points p ; (2) it is possible to write a solvable recurrence relation, in n and p , for this function. Following Cover's original paper, we denote the (data-independent) number of dichotomies $\mathcal{N}_1(X_p)$ by $C_{n,p}$, and the corresponding VC entropy by $\mathcal{H}_{n,p} = \log C_{n,p}$.

Theorem 2 (Cover, 1965). Let X_p be a set of p points in \mathbb{R}^n . If the points are in general position, i.e., if the points in X' are linearly independent for all subsets $X' \subseteq X_p$ such that $|X'| \leq n$, then $\mathcal{N}_1(X_p) = C_{n,p}$, where

$$C_{n,p} = 2 \sum_{j=0}^{n-1} \binom{p-1}{j}. \quad (14)$$

The proof of Theorem 2 is based on a simple recurrence relation for $C_{n,p}$:

$$C_{n,p+1} = C_{n,p} + C_{n-1,p}, \quad (15)$$

with boundary conditions

$$C_{n \geq 1,1} = 2, \quad C_{0,p} = 0. \quad (16)$$

Equation (15) states that adding the $(p+1)$ th point X to X_p increases the number of dichotomies by $C_{n-1,p}$, which is the number of dichotomies of X_p that are realizable by a vector W such that $W \cdot X = 0$. Cover actually proved a more general statement. Informally, if one maps all elements of X_p by the nonlinear kernel function ϕ from \mathbb{R}^n to \mathbb{R}^d with d larger than n , then, under mild assumptions on ϕ , Eq. (14) holds with d in place of n .

It is straightforward to see how the storage capacity defined in statistical mechanics, α_c (recall that $\alpha = p/n$), can be obtained from $C_{n,p}$. The number of dichotomies is a combinatorial quantity and is expected to scale exponentially in n , at least for small α . Thus, an intensive quantity can be defined by normalizing $C_{n,p}$ with the total number of dichotomies of p points. The fraction of dichotomies $c_{n,p} \equiv C_{n,p}/2^p$ is bounded, $0 \leq c_{n,p} \leq 1$, and has a nontrivial thermodynamic limit $c_\infty(\alpha)$. The thermodynamic limit is defined by taking both $n, p \rightarrow \infty$, with fixed $\alpha = p/n$. It is not hard to see directly from Eq. (14) that

$$c_\infty(\alpha) = \theta(\alpha_c - \alpha), \quad (17)$$

with $\alpha_c = 2$. The expression in Eq. (17) takes the value 1 for $\alpha < \alpha_c$, the value 0 for $\alpha > \alpha_c$, and the value 1/2 for $\alpha = \alpha_c$ (θ is the Heaviside step function). Qualitatively, $c_{n,\alpha n}$ as

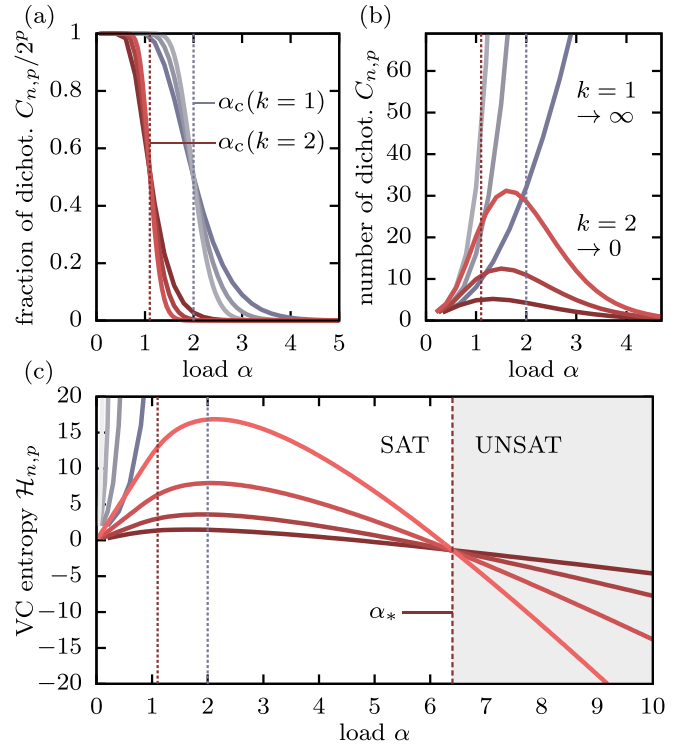


FIG. 1. While the fraction of admissible dichotomies (a) has qualitatively similar behavior for unstructured (gray curves, $k = 1$) and structured (red curves, $k = 2$) data, the absolute number of dichotomies (b) has different limit behaviors. As a consequence, the VC entropy (c) diverges to $+\infty$ for unstructured data and to $-\infty$ for structured data. Curves of the VC entropy at different values of n intersect, for large n , at the same critical value α_* of the load. Vertical dotted lines in all panels are the storage capacities. The dashed line in (c) is the transition caused by data structure. $n = 5, 10, 20$ in (a), $n = 3, 4, 5$ in (b), $n = 5, 10, 20, 40$ in (c).

a function of α is a decreasing sigmoid, which is steeper for larger values of n [see Fig. 1(a)]. This allows the definition of a notion of capacity at finite dimension n , as the value $\tilde{\alpha}_c(n)$ such that $c_{n,\tilde{\alpha}_c(n)n} = 1/2$, or

$$C_{n,\tilde{\alpha}_c(n)n} = 2^{p-1}. \quad (18)$$

Another notable value of p can be read off of $c_{n,p}$: It is the Vapnik-Chervonenkis dimension d_{VC} , equal to the maximum p such that $c_{n,p} = 1$. For a linear separator, $d_{VC} = n$. Notice that one cannot use the asymptotic form Eq. (17) to this aim, since the thermodynamic limit pushes $c_{n,\alpha n}$ to 1 for all values of α up to α_c .

Notice that Eq. (14) implies that the VC entropy grows asymptotically as $\mathcal{H}_{n,p} \sim (n-1) \log p$ for large number of inputs p (see Sec. III B for a derivation). This is the same behavior as that obtained by bounding the VC entropy as in Eq. (11).

Two remarks can be made, concerning the generality of Cover's theorem. First, the general position is a rather weak condition. For instance, we mention three examples of distributions of the points $\xi_\mu \in X_p$ under which the general position holds with probability 1: (1) $\xi^\mu \in X_p$ are i.i.d. variables with the uniform measure on the sphere S^{n-1} ; (2) $\xi^\mu \in X_p$ are

i.i.d. variables with marginal probability distribution $P(\xi)$, and the support of P is \mathbb{R}^n ; (3) the coordinates of each $\xi^\mu \in X_p$ are i.i.d. variables, with discrete probability distribution $p(x) = (1+m)/2\delta_{x,1} + (1-m)/2\delta_{x,-1}$, for any $m \in [-1, 1]$. Clearly, there are trivial ways to violate general position: For instance, if the probability distribution of (1) or (2) above is conditioned to assigning the same value to a fixed subset of size $k < n$ of the coordinates of all inputs. Then Cover's theorem still applies in the subspace, with $n - k$ in place of n .

Second, the condition that ϕ must satisfy for the theorem to apply to the kernel machine specified by ϕ is essentially that the vectors $\phi(\xi^\mu)$ must be in general position in the feature space \mathbb{R}^d . This again is a very mild condition. Starting with a set of inputs X_p in general position in the original n -dimensional space, most interesting mappings satisfy the condition. This includes polynomial kernels, but also more complex functions, such as those of the form $\phi_i(\xi) = g(\sum_j W_{ij}\xi_j)$, where g is an activation function (e.g., ReLU or tanh) and W is any rectangular random matrix. The latter case is relevant for the theory of extreme learning machines [41].

C. Constrained models of structured data

The discussion above suggests that, in order to go beyond the prediction of Cover's theorem, one needs a way of introducing statistical dependence between the inputs X_p and their labels $Y_p = (Y^1, \dots, Y^p)$. This reflects a simple observation that can be made on empirical data sets of images: Similar inputs tend to be classified similarly. For instance, one expects that there exists an (unknown) set of transformations on an input image X , possibly including some translations, dilations, and rotations, that leave the classification of X invariant. Such intuition agrees with the concepts, put forward in neuroscience and gaining momentum in physics, of invariant recognition (the similar neural representation of the same object in different conditions) and object manifolds (sets of input stimuli giving rise to the same neural representation) [5–7,42,43].

Integrating data structure within the framework of statistical mechanics is relatively straightforward and usually follows two steps: (1) define a generative model for the data, given in terms of a nonfactorized joint probability distribution $P(X_p, Y_p)$ and (2) compute averages over the measure P (the “disorder”); this is what was done, for instance, in Refs. [5,12,34]. How to best address data dependence in the SLT formalism, instead, is a debated issue. Here we follow a simple strategy inspired by recent literature in statistical physics: We change the input space \mathcal{X} . Each input X^μ is now an object manifold, i.e., a (possibly countably or uncountably infinite) set of points that, by definition, are to be classified coherently.

We focus on two simple realizations of data structure, the first motivated by the availability of analytical results and the second motivated by its connection to the well-known framework of margin learning.

a. Simplex learning. Inputs are “multiplets” of k points with fixed geometric interrelations. The input set is $X_p = \{X^\mu\}_{\mu=1,\dots,p}$, where each $X^\mu = \{\xi_a^\mu\}_{a=1,\dots,k}$ is a set of k points on the unit $(n - 1)$ -sphere, $\xi_a^\mu \in S^{n-1}$. The $k(k - 1)/2$ overlaps within each multiplet are fixed: $\xi_a^\mu \cdot \xi_b^\mu = \rho_{ab}$ for all $\mu =$

$1, \dots, p$. We assume the uniform probability measure on each point ξ_a^μ , conditioned on the constraint on the overlaps [32]. The usual unconstrained ensemble is recovered for $k = 1$, or at any k if $\rho_{ab} = 1$ for all a, b . The name “simplex” is justified by the fact that, since linear classification is a projective problem, if $Y = g(X)$ for each X in a set of points X^μ , then $Y = g(X)$ for all X in the convex hull of X^μ . The input space $\mathcal{X}_S(\{\rho_{ab}\})$ depends on k and ρ_{ab} , and is the set of all multiplets with the given constraints.

b. Margin learning. Given a kernel machine with kernel $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^d$ and inputs $X \in \mathcal{X} = \mathbb{R}^n$, learning with margin κ is defined by the class $\mathcal{G}_M(\kappa)$ of all functions

$$g_\kappa(X) = \begin{cases} +1 & W \cdot \phi(X) > \kappa \\ -1 & W \cdot \phi(X) < -\kappa. \end{cases} \quad (19)$$

Cases falling within the margin $(-\kappa, \kappa)$ can be defined with a third value, for instance, 0, or left undefined. Hence, the corresponding loss class projected on a sample (X_p, Y_p) , Eqs. (4) and (5), contains all the dichotomies of X_p that can be realized by an element of $\mathcal{G}_M(\kappa)$. An alternative representation of margin learning can be given via the definition of appropriate object manifolds. In fact, linear separation of points with margin κ is equivalent to zero-margin linear separation of spherical object manifolds with radius κ [5]. Thus, $Y^\mu = g_\kappa(X^\mu)$ for all μ if and only if $Y^\mu = g_0(Q^\mu)$ for all μ and all Q^μ such that $|Q^\mu - \phi(X^\mu)|^2 < \kappa^2$. The input space $\mathcal{X}_M(\kappa)$ is the set of the preimages, via ϕ , of all spheres of radius κ in \mathbb{R}^d . Note that, while margin learning has a natural description in terms of the original space $\mathcal{X} = \mathbb{R}^n$, through the hypothesis class $\mathcal{G}_M(\kappa)$, simplex learning does not have such a straightforward representation and is defined directly by means of the object space $\mathcal{X}^S(\{\rho_{ab}\})$.

The VC entropy for margin learning, \mathcal{H}_κ , obeys the general bound of Eq. (11), in terms of the corresponding VC dimension d_{VC} . In turn, an upper bound of d_{VC} exists for points lying on the d -dimensional sphere of radius R [27]:

$$d_{VC}(\kappa) \leq \min \left[\frac{R^2}{\kappa^2}, d \right]. \quad (20)$$

The standard bound is therefore again logarithmic in the sample size p . In the following, we set out to investigate the behavior of the VC entropy of kernel machines for the two data structures defined above, in order to quantify how loose these logarithmic upper bounds are. We do so by means of two complementary approaches: The combinatorial framework and the theory of disordered systems.

III. COMBINATORIAL APPROACH

Very recently the combinatorial approach introduced by Cover was extended to formulate a mean field theory of simplex learning [32]. In this section we focus on this model of data structure. The definition we have given above of simplex learning specifies the ensemble of the sets X_p . It remains to define the hypothesis class $\mathcal{G}_S(\{\rho_{ab}\})$. This is straightforward: One starts from the class \mathcal{G} of linear separators in \mathbb{R}^n and restricts it to the class $\hat{\mathcal{G}}(\{\rho_{ab}\})$ of those functions $h \in \mathcal{G}$ that assign the same label to all points in each multiplet X^μ (i.e., those that are constant on each multiplet). Then the restricted

hypothesis class is defined as

$$\begin{aligned} \mathcal{G}_S(\{\rho_{ab}\}) &= \{g \text{ s.t. } \exists h \in \hat{\mathcal{G}}(\{\rho_{ab}\}) \text{ s.t. } \forall X^\mu \in X_p, g(X^\mu) \\ &= h(\xi \in X^\mu)\}. \end{aligned}$$

The functions in $\hat{\mathcal{G}}(\{\rho_{ab}\})$ are called admissible. The mean-field combinatorial theory allows the computation of the average $\langle \mathcal{N}_\perp(X_p) \rangle_{X_p}$, i.e., the average number of admissible dichotomies of simplexes that can be realized linearly. We will still denote this number with $C_{n,p}$, although it depends on the parameters k and $\{\rho_{ab}\}$ of the ensemble.

The quantities $C_{n,p}$ satisfy the recurrence relation, proven in Ref. [32],

$$C_{n,p+1} = \sum_{l=0}^k \theta_l^k C_{n-l,p}, \quad (21)$$

where the constant coefficients θ_l^k are fixed in turn by the recurrence relation

$$\theta_l^k = \psi_k \theta_l^{k-1} + (1 - \psi_k) \theta_{l-1}^{k-1}, \quad (22)$$

with boundary conditions

$$\theta_0^1 = \theta_1^1 = 1, \quad \theta_{-1}^k = \theta_{k+1}^k = 0. \quad (23)$$

The boundary conditions for Eq. (21) are difficult to express precisely for generic k . Here we will assume the boundary conditions in Eq. (16) for all k . This approximation is expected to have a negligible effect for the asymptotic analysis presented in the following; we checked the validity of this approximation numerically for the first nontrivial cases $k = 2$ and $k = 3$. Each coefficient θ_l^k in Eq. (21) depends on $k - 1$ numbers $\{\psi_m\}_{m=2,\dots,k}$, with $0 \leq \psi_m \leq 1$, having the following geometric-probabilistic interpretation. Let $w \in S^{n-1}$ be a random vector with the flat measure on the unit sphere. Consider any multiplet X^μ and a subset $X' \subseteq X^\mu$ of $m \leq k$ points. Then ψ_m is the symmetrized probability that the scalar product $w \cdot \xi$ has the same sign for all $\xi \in X'$, conditioned on it having the same sign for all $\xi \in X' \setminus \{\xi_\star\}$:

$$\psi_m = 2 \langle \text{P}[(w \cdot \xi_\star) > 0 \mid (w \cdot \xi) > 0 \forall \xi \in X' \setminus \{\xi_\star\}] \rangle_{\text{sym}},$$

where the symmetrization $\langle \cdot \rangle_{\text{sym}}$ is performed by averaging over all subsets X' and over all choices of $\xi_\star \in X'$. These quantities can be expressed in terms of the overlaps ρ_{ab} , e.g.,

$$\psi_2(\rho) = \frac{2}{\pi} \arctan \sqrt{\frac{1+\rho}{1-\rho}}. \quad (24)$$

(More information on ψ_m can be found in Ref. [32].)

The notion of storage capacity α_c can be defined for structured data, similarly to Cover's unstructured case, and coherently with the thermodynamic limit addressed in statistical mechanical computations. The combinatorial theory yields

$$\alpha_c(k) = \left(k - \frac{1}{2} - \sum_{l=2}^k \psi_l \right)^{-1}. \quad (25)$$

A. Asymptotic analysis via analytic combinatorics

In the case of unstructured data, we know that the growth of $C_{n,p}$ as a function of p is exponential up to the capacity

$p_c = 2n$ and subexponential afterwards. Due to this change of behavior, the fraction of linearly realizable dichotomies, $c_{n,p} = C_{n,p}/2^p$, has a discontinuous transition from 1 to 0 in the thermodynamic limit [see Fig. 1(a)]. What is the asymptotic growth rate of $C_{n,p}$? This question can be answered by inspecting the explicit solution (14). However, we construct a different method here, based on the techniques of analytic combinatorics. Our method has the crucial advantage of being applicable to cases where (1) the solution $C_{n,p}$ is not known explicitly and (2) the recurrence equation is given implicitly, as a relation between its coefficients.

Let $g_n(z)$ be the ordinary generating function of $C_{n,p}$ with respect to the variable p :

$$g_n(z) = \sum_{p=1}^{\infty} C_{n,p} z^p. \quad (26)$$

Formally, the coefficient $C_{n,p}$ can be obtained by derivation as

$$C_{n,p} = \frac{1}{p!} \left. \frac{d^p}{dz^p} g_n(z) \right|_{z=0}. \quad (27)$$

When it is unfeasible to compute the p th derivative explicitly, one can extract information on the asymptotic behavior of $C_{n,p}$ for large p by means of analytic techniques (see, for instance, Ref. [44]).

Whenever the generating function (26) is a rational function analytic in $z = 0$, it admits a partial fraction expansion

$$g_n(z) = Q_n(z) + \sum_s \sum_{r=1}^{r_s} \frac{a_{s,r}}{(z - z_s)^r}, \quad (28)$$

where Q_n is a polynomial, s ranges over the poles of g_n , and r_s is the multiplicity of the pole s . Then the asymptotic form of the coefficients of $g_n(z)$ can be read off the series expansion of $(z - z_s)^{-r}$:

$$(z - z_s)^{-r} = \frac{(-1)^r}{z_s^r} \sum_{p=0}^{\infty} \binom{p+r-1}{r-1} z_s^{-p} z^p. \quad (29)$$

By substituting (29) in Eq. (28) one obtains r_s different contributions for each pole s . The overall leading term corresponds to the dominant singularity z_0 of $g_n(z)$, i.e., the one with smallest modulus $|z_0|$. This is due to the term z_s^{-p} in (29) that suppresses the subdominant poles exponentially. Among the contributions due to z_0 , the leading one is that with $r = r_s$, because the binomial coefficient in (29) is a polynomial of degree $r - 1$ in p . Putting it all together, if the dominant singularity is a pole of order r , then

$$C_{n,p} \sim R z_0^{-p-r} \binom{p+r-1}{r-1}, \quad (30)$$

where the constant R can be obtained by factoring out the singularity:

$$R = \lim_{z \rightarrow z_0} (z_0 - z)^r g_n(z). \quad (31)$$

Equation (30) shows that if $|z_0| < 1$ (respectively, > 1), $C_{n,p}$ increases (respectively, decreases) exponentially with p at fixed n ; if $|z_0| = 1$ then the asymptotic behavior is polynomial (of order $r - 1$).

In simple cases, when it is possible to obtain $g_n(z)$ in closed form, this method can be applied straightforwardly. However, this set up allows to probe the asymptotics of $C_{n,p}$ even in more complicated scenarios, where $g_n(z)$ cannot be solved for explicitly, or when even the recurrence relation for $g_n(z)$ is not specified completely. Section III E shows how to tackle this more general problem. Before that, we consider the simpler cases $k = 1$ and $k = 2$.

B. Asymptotics for unstructured data

As a “warm-up exercise,” we use the combinatorial method described above to explore the asymptotics of $C_{n,p}$ in the well-understood unstructured case.

By multiplying both sides of Eq. (15) by z^p and summing over p one obtains

$$\frac{1}{z}g_n(z) - 2 = g_n(z) + g_{n-1}(z), \tag{32}$$

where the constant term 2 comes from the initial condition (16). It is useful to rewrite the equation as

$$g_n(z) = \frac{z}{1-z}[g_{n-1}(z) + 2]. \tag{33}$$

The boundary condition is $g_0(z) = 0$, due to every $C_{0,p}$ being zero. The relation (33) is a linear (nonhomogeneous) first-order recurrence with constant coefficients, whose solution is

$$g_n(z) = \frac{2z}{2z-1} \left[\left(\frac{z}{1-z} \right)^n - 1 \right]. \tag{34}$$

Equation (34) shows that $g_n(z)$ has a single pole at $z_0 = 1$, of order n , with finite part $R = 2$. Therefore, the corresponding asymptotic form has no exponential factor and is purely polynomial:

$$C_{n,p} \sim 2 \binom{p+n-1}{n-1} = \frac{2}{(n-1)!} p^{n-1} + O(p^{n-2}). \tag{35}$$

Note that the right-hand side of Eq. (34) has a removable discontinuity in $z_1 = 1/2$, where the apparent pole in the first term gets canceled by a zero in the numerator (the term in square brackets). The corresponding exponential asymptotic growth, 2^p , is present in $C_{n,p}$ only transiently, for $p < n$.

C. Nonmonotonicity of the VC entropy

The behavior of $C_{n,p}$, and therefore of the VC entropy, changes dramatically when data structure is present, already in the simplest case where the training data are structured as pairs of points, i.e., $k = 2$. Figure 1 shows the fraction of dichotomies, $C_{n,p}/2^p$, and the number of dichotomies, $C_{n,p}$, as functions of α for increasing values of the dimension n , for $k = 1$ and $k = 2$ with $\rho = 0.3$. The fraction of dichotomies is qualitatively similar in the two scenarios, the only apparent difference being the expected decrease in the storage capacity. A remarkable divergence appears instead in the asymptotic behavior of $C_{n,p}$. The absolute number of dichotomies is nonmonotonic for simplex learning already in the simplest nondegenerate case $k = 2$ with $\rho < 1$. What is also evident in Fig. 1(b) is the fact that the storage capacity $\alpha_c(k)$ does not

pinpoint any qualitatively special point for the unnormalized $C_{n,p}$, and therefore for the VC entropy.

Since the two-point case $k = 2$ is the simplest case where the nonmonotonicity of the VC entropy arises, we work it out in detail, before showing the general k -point case below. The geometry of the problem is fixed by the single quantity ψ_2 . The recurrence equation is obtained by specializing Eqs. (21), (22), and (23) to the case $k = 2$:

$$C_{n,p+1} = \psi_2 C_{n,p} + C_{n-1,p} + (1 - \psi_2) C_{n-2,p}, \tag{36}$$

with boundary conditions $C_{0,p} = 0, C_{n,1} = 2\{1 - [1 - \psi_2(d)]\delta_{n,1}\}$. In order to simplify the computations, we will use the same boundary conditions as for $k = 1$, i.e., $C_{0,p} = 0$ and $C_{n \geq 1,1} = 2$. This approximation has negligible effects in the large- n limit [32].

Equation (36) fixes the recurrence relation satisfied by the generating function $g_n(z)$:

$$g_n(z) = \frac{z}{1 - \psi_2 z} [g_{n-1}(z) + (1 - \psi_2)g_{n-2}(z) + 2], \tag{37}$$

with boundary condition $g_{n \leq 0}(z) = 0$. The solution, which can be found by means of the characteristic polynomial method, reads

$$g_n(z) = \left[\frac{z - \sqrt{\Delta(z)}}{2(1 - \psi_2 z)} \right]^n \frac{z}{2z-1} \left[1 + z \frac{2\psi_2 - 3}{\sqrt{\Delta(z)}} \right] + \left[\frac{z + \sqrt{\Delta(z)}}{2(1 - \psi_2 z)} \right]^n \frac{z}{2z-1} \left[1 - z \frac{2\psi_2 - 3}{\sqrt{\Delta(z)}} \right] - \frac{2z}{2z-1},$$

where $\Delta(z) = z[4(1 - \psi_2) + z(1 - 2\psi_2)^2]$. The explicit solution has a pole of order n in $z_0 = 1/\psi_2$, with finite part

$$R = 2\psi_2^{-2n}. \tag{38}$$

Similarly to the unstructured case, the singularity in $z = 1/2$ is removable, which signals that the initial exponential increase of the number of dichotomies must be superseded eventually by the asymptotic behavior due to z_0 . Altogether, the large- p form of $C_{n,p}$ is

$$C_{n,p} \sim 2 \binom{p+n-1}{n-1} \psi_2^{p-n}. \tag{39}$$

The crucial difference between the results for $k = 1$, Eq. (35), and $k = 2$, Eq. (39), lies in the fact that while the first is asymptotically increasing, the second is exponentially decreasing whenever $\psi_2 < 1$, i.e., when the two partner points are distinct. Observe that $C_{n,p}$ always increases for small p ; this is a consequence of the fact that the Vapnik-Chervonenkis dimension of a linear classifier in n dimensions is $d_{VC} = n$, therefore all dichotomies of kp points can be realized when $p \leq n/k$, meaning that $C_{n,p \leq n/k} = 2^p$. The decreasing asymptotic form then proves that the Vapnik-Chervonenkis entropy $\mathcal{H}_{n,p}$ is nonmonotonic in p (and therefore in α) for fixed n .

Intuitively, the nonmonotonicity is due to the competition of two opposing effects. On one hand, the addition of a new pair of points $\{\xi, \bar{\xi}\}$ to a set of p existing pairs entails a combinatorial increase in the total number of linearly realizable dichotomies. On the other hand, some of the $C_{n,p}$ admissible dichotomies can become invalid if they are realizable only by hyperplanes intersecting the segment connecting ξ and $\bar{\xi}$.

D. Emergence of a data-driven satisfiability transition

A nontrivial consequence of the nonmonotonic VC entropy can be observed in Fig. 1(c). Consider the VC entropy $\mathcal{H}_{n,\alpha n}$ as a function of α . The curves $\mathcal{H}_{n,\alpha n}$ at different values of n intersect each other roughly around the same point α_* . More precisely, if $\mathcal{H}_{n,\alpha n}$ and $\mathcal{H}_{n-1,\alpha(n-1)}$ intersect at $\alpha_*(n)$, then $\alpha_* = \lim_{n \rightarrow \infty} \alpha_*(n)$ is finite. This empirical observation can be clarified analytically.

As a function of the load $\alpha = p/n$, Eq. (39) becomes

$$C_{n,\alpha n} \sim C(\alpha; n) \equiv 2 \frac{\Gamma(\alpha n + n)}{\Gamma(n)\Gamma(\alpha n + 1)} \psi_2^{(\alpha-1)n} \quad (40)$$

(Γ is the Euler gamma function), or $\mathcal{H}_{n,\alpha n} \sim \mathcal{H}(\alpha; n)$ with

$$\mathcal{H}(\alpha; n) \equiv \log \left[2 \frac{\Gamma(\alpha n + n)}{\Gamma(n)\Gamma(\alpha n + 1)} \right] + (\alpha - 1)n \log \psi_2. \quad (41)$$

In the nondegenerate case (whenever $\psi_2 < 1$) the second term in (41) is negative for $\alpha > 1$, while the first term is always positive. This competition gives rise to a transition at $\alpha = \alpha_* > 1$, where the asymptotic limit of the VC entropy changes:

$$\lim_{n \rightarrow \infty} \mathcal{H}(\alpha; n) = \begin{cases} -\infty & \alpha < \alpha_* \\ \infty & \alpha > \alpha_* \end{cases} \quad (42)$$

The transition point is pinpointed by the condition

$$\lim_{n \rightarrow \infty} \frac{d}{dn} \mathcal{H}(\alpha_*; n) = 0. \quad (43)$$

With $\mathcal{H}(\alpha; n)$ given by Eq. (41), the condition reads

$$\lim_{n \rightarrow \infty} [(\alpha_* - 1) \log \psi_2 + (\alpha_* + 1) \Psi(\alpha_* n + n) - \Psi(n) - \alpha_* \Psi(\alpha_* n + 1)] = 0, \quad (44)$$

where $\Psi(z) \equiv \partial_z \log \Gamma(z)$ is the poly-gamma function, whose asymptotic behavior is $\Psi(z) = \log(z) + O(1/z)$. Sending n to infinity then gives the transcendental equation

$$(\alpha_* + 1) \log(\alpha_* + 1) - \alpha_* \log \alpha_* + (\alpha_* - 1) \log \psi_2 = 0, \quad (45)$$

which has two solutions: α_* is the larger. As a function of ψ_2 , the transition point α_* has limits

$$\lim_{\psi_2 \rightarrow 0} \alpha_* = 1, \quad \lim_{\psi_2 \rightarrow 1} \alpha_* = \infty. \quad (46)$$

As expected, when ψ_2 goes to 1, the problem reduces to that of classifying unstructured data, and the transition runs to infinity.

The phase transition at α_* can be rationalized as the SAT-UNSAT transition of a random constraint satisfaction problem (CSP). First, we recall that the storage capacity α_c itself corresponds to the transition between the satisfiable and the unsatisfiable phase of an appropriate satisfiability problem. The CSP relevant to α_c can be stated as follows:

Constraint satisfaction Problem 1. Given a set of kn input-label pairs $\{\xi_a^\mu, \sigma^\mu\}$ (with $a = 1, \dots, k$ and $\mu = 1, \dots, p$), find a vector w such that $\text{sgn}(w \cdot \xi_a^\mu) = \sigma^\mu$ for all μ and a .

The input data of this problem satisfy the admissibility constraints by construction. A corresponding random constraint satisfaction problem (rCSP) is an ensemble of CSPs, specified

by a probability measure on the input data. The rCSP is in the SAT (respectively UNSAT) phase when the satisfiability problem admits a solution with probability one (respectively zero) in the thermodynamic limit. The storage capacity (25) marks the transition between the SAT and the UNSAT phases of the rCSP corresponding to Problem 1 with the probability measure of simplex learning described in Sec. II C

A different problem can be constructed by moving the admissibility property from the definition of the input data to the conditions defining the solution:

Constraint satisfaction Problem 2. Given a set of kn input points $\{\xi_a^\mu\}$ (with $a = 1, \dots, k$ and $\mu = 1, \dots, p$), find a set of labels $\{\sigma^\mu\}$ and a vector w such that $\text{sgn}(w \cdot \xi_a^\mu) = \sigma^\mu$ for all μ and a .

Notice that this problem is trivially satisfiable for unstructured data, i.e., it is satisfied by almost all vectors w when the constraint of admissibility is irrelevant (i.e., when $k = 1$). A solution to Problem 2 is given by specifying an admissible dichotomy $\{\sigma^\mu\}$ and a vector w . In this framework, the VC entropy counts the (logarithm of the) number of distinct dichotomies $\{\sigma^\mu\}$ that can appear in such a solution. This means that the corresponding rCSP is in the UNSAT phase when $\mathcal{H}(\alpha; n) \rightarrow -\infty$ and in the SAT phase otherwise.

E. Transition point for generic k

Now we address the more general case where the number of partners in a multiplet is k . The generating function $g_n(z)$ satisfies the recurrence equation

$$g_n(z) = \frac{z}{1 - z\theta_0^k} \left[2 + \sum_{l=1}^k \theta_l^k g_{n-l}(z) \right], \quad (47)$$

as can be obtained from Eq. (21). Solving for $g_n(z)$ from Eqs. (47) and (22) would be hopeless. However, the asymptotic analysis discussed above only needs three pieces of information about $g_n(z)$: (1) the location z_0 of the dominant singularity, (2) its order r , and (3) its finite part R . These can be extracted from the recurrence relations without solving them.

The right-hand side of Eq. (47) has a singularity in $z = 1/\theta_0^k$. The boundary condition is $g_{n \leq 0}(z) = 0$, therefore the first nonzero function is $g_1(z) = 2\sigma(z)$, where

$$\sigma(z) = \frac{z}{1 - z\theta_0^k} \quad (48)$$

encapsulates the singularity. Since the number of terms in the sum in Eq. (47) is finite, no other singularity can appear at finite n . Therefore

$$z_0 = \frac{1}{\theta_0^k}. \quad (49)$$

Now consider one iteration of Eq. (47): The singularity with largest order in the right-hand side comes from $g_{n-1}(z)$, and the singular term gets multiplied by $\theta_1^k \sigma(z)$. Indeed, it is easy to see by induction that the leading term $\hat{g}_n(z)$ in the Laurent expansion of $g_n(z)$ around z_0 is

$$\hat{g}_n(z) = 2(\theta_1^k)^{n-1} \sigma(z)^n. \quad (50)$$

Therefore, the order of the singularity is $r = n$. The constant R [Eq. (31)] can be obtained by multiplying Eq. (50) by

$(1/\theta_0^k - z)^n$ and evaluating it at $z = 1/\theta_0^k$:

$$R = 2(\theta_1^k)^{n-1}(\theta_0^k)^{-2n}. \quad (51)$$

Finally, the asymptotic behavior of $C_{n,p}$ is

$$C_{n,p} \sim 2 \binom{p+n-1}{n+1} (\theta_1^k)^{n-1} (\theta_0^k)^{p-n}, \quad (52)$$

from which one readily obtains the asymptotic form $C(\alpha; n)$ for the number of dichotomies,

$$C(\alpha; n) = 2 \frac{\Gamma(\alpha n + n)}{\Gamma(n)\Gamma(\alpha n + 1)} (\theta_1^k)^{n-1} (\theta_0^k)^{(\alpha-1)n}, \quad (53)$$

and the corresponding one for the VC entropy,

$$\mathcal{H}(\alpha; n) = \log \left[2 \frac{\Gamma(\alpha n + n)}{\Gamma(n)\Gamma(\alpha n + 1)} \right] + (n-1) \log \theta_1^k + (\alpha-1)n \log \theta_0^k. \quad (54)$$

As above, the existence of a critical value α_* can be established by finding the zeros of the derivative of $\mathcal{H}(\alpha; n)$ with respect to n , in the large- n limit. One finds

$$(\alpha_* + 1) \log(\alpha_* + 1) - \alpha_* \log \alpha_* + (\alpha_* - 1) \log \theta_0^k + \log \theta_1^k = 0. \quad (55)$$

The two coefficients θ_0^k and θ_1^k can be obtained from Eq. (22) as functions of the ψ . By solving the recurrence equation, specialized to $l = 0$, one has

$$\theta_0^k = \prod_{m=2}^k \psi_m. \quad (56)$$

Then, by substituting expression (56) into Eq. (22) with $l = 1$, one obtains the recurrence relation

$$\theta_1^k = \psi_k \theta_1^{k-1} + (1 - \psi_k) \prod_{m=2}^{k-1} \psi_m, \quad (57)$$

with boundary condition $\theta_1^1 = 1$. The solution is

$$\theta_1^k = \left(2 - k + \sum_{m=2}^k \frac{1}{\psi_m} \right) \prod_{m=2}^k \psi_m. \quad (58)$$

Specializing to $k = 3$, for instance, yields

$$\theta_0^3 = \psi_3 \psi_2, \quad \theta_1^3 = \psi_3 + \psi_2 - \psi_3 \psi_2. \quad (59)$$

Because of the way θ_0^k and θ_1^k are constructed via the geometric quantities $\psi_m \in [0, 1]$, they are not independent. The range of θ_0^k is $[0, 1]$, as can be seen from Eq. (56). The sup and inf of θ_1^k at fixed θ_0^k can be obtained by considering the two extremal cases:

$$\begin{aligned} \text{(i)} \quad \{\psi_m\}_m &= \{1, \dots, 1, \theta_0^k, 1, \dots, 1\}, \\ \text{(ii)} \quad \{\psi_m\}_m &= \{(\theta_0^k)^{1/(k-1)}, \dots, (\theta_0^k)^{1/(k-1)}\}. \end{aligned} \quad (60)$$

The fact that the evaluation on the two extremal cases gives the appropriate bounds is not obvious: It can be proved by induction using Lagrange's theorem for constrained optimization

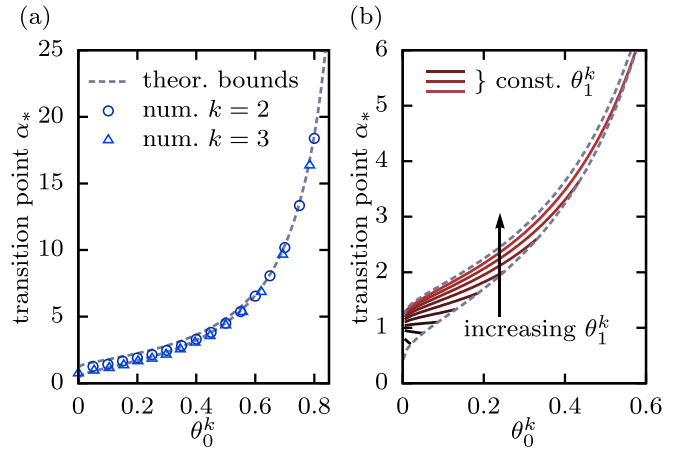


FIG. 2. (a) Numerical estimates of α_* at varying θ_0^k for two different geometries: $k = 2$ (where θ_0^2 is just ψ_2) and $k = 3$. In the latter case we fix $\{\rho_{ab}\}$ by requiring that the three points in the simplex form an equilateral triangle of varying sizes. (b) Theoretical results (red curves) for α_* as a function of θ_0^k for increasing values of θ_1^k , within its allowed range given by Eqs. (61) and (62). Dashed lines in both panels are the k -independent upper and lower bounds for α_* .

(taking care to consider the boundary of the domain as well); see Appendix A. From (i) and (ii), respectively, one gets

$$\begin{aligned} \text{(i)} \quad \sup \theta_1^k &= 1, \\ \text{(ii)} \quad \inf \theta_1^k &= (k-1)(\theta_0^k)^{1-k} + (2-k)\theta_0^k. \end{aligned} \quad (61)$$

The inf is monotonically decreasing with k ; therefore, by letting $k \rightarrow \infty$ one obtains a global lower bound independent of k :

$$\theta_1^k > \theta_1^\infty = \theta_0^k [1 - \log \theta_0^k]. \quad (62)$$

The upper bound (i) is already k -independent.

Figure 2 summarizes the results concerning the value of α_* for generic k . It also shows a comparison with numerical results obtained for $k = 3$, (with $\{\rho_{ab}\}$ given by the equilateral geometry). The theoretical bounds in the figure (dashed lines) are obtained by substituting the k -independent bounds above into Eq. (55).

We point out that there are two sources of approximation in the computations above: (1) the modified boundary conditions and (2) the perturbative nature of the asymptotic analysis. Concerning (1), we remark that the numerical results were obtained by using the correct boundary conditions. However, using the modified conditions does not change the numerical results appreciably. The small discrepancies apparent in Fig. 2 are therefore due almost entirely to (2).

F. Finite-size scaling at the critical point

In the vicinity of the transition point α_* , the quantity $C(\alpha; n)$ satisfies finite-size scaling, as happens for other random satisfiability problems [45,46]. In this section we compute the scaling form and its critical exponents.

Let us define a scaling variable y as n times the reduced load $(\alpha - \alpha_*)/\alpha_*$ around α_* :

$$y = n \frac{\alpha - \alpha_*}{\alpha_*}. \quad (63)$$

By inserting $\alpha = \alpha_* y/n + \alpha_*$ in Eq. (53), and using the asymptotic expansion of the Γ function,

$$\Gamma(x) = e^{x \log x - x} [\sqrt{2\pi} x^{-1/2} + O(x^{-3/2})], \quad (64)$$

one obtains in the large- n limit

$$C(\alpha; n) = e^{nA+B} \left[\frac{\sqrt{2/\pi}}{\sqrt{\alpha_* (1 + \alpha_*)}} n^{-1/2} + O(n^{-3/2}) \right],$$

with

$$\begin{aligned} A &= (\alpha_* + 1) \log(\alpha_* + 1) - \alpha_* \log \alpha_* \\ &\quad + (\alpha_* - 1) \log \theta_0^k + \log \theta_1^k, \\ B &= -\log \theta_1^k + \alpha_* y \log(\alpha_* + 1) - \alpha_* y \log \alpha_* + \alpha_* y \log \theta_0^k. \end{aligned}$$

The linear term nA in the exponential vanishes by Eq. (55). Hence,

$$\begin{aligned} C(\alpha; n) &= n^{-1/2} \frac{1}{\theta_1^k} \frac{\sqrt{2/\pi}}{\sqrt{\alpha_* (1 + \alpha_*)}} \left(\frac{\alpha_* + 1}{\alpha_*} \theta_0^k \right)^{\alpha_* y} \\ &\quad \times [1 + O(n^{-3/2})], \end{aligned} \quad (65)$$

which shows that in the thermodynamic limit $C(\alpha; n)$ obeys the scaling form

$$C(\alpha; n) = n^{-1/2} F \left(\frac{\alpha - \alpha_*}{\alpha_*} n \right) \quad (66)$$

with the exponential scaling function

$$F(y) = \frac{1}{\theta_1^k} \frac{\sqrt{2/\pi}}{\sqrt{\alpha_* (1 + \alpha_*)}} \left(\frac{\alpha_* + 1}{\alpha_*} \theta_0^k \right)^{\alpha_* y}. \quad (67)$$

Equation (66) shows that, within the approximation of our asymptotic analysis, the number of dichotomies satisfies the finite-size scaling form

$$C_{n,\alpha n} \sim n^{-\beta/\nu} F \left(\frac{\alpha - \alpha_*}{\alpha_*} n^{1/\nu} \right) \quad (68)$$

(where F is regular), with critical exponents

$$\beta = 1/2, \quad \nu = 1. \quad (69)$$

Let $h(\alpha)$ be the VC entropy density in the thermodynamic limit:

$$h(\alpha) = \lim_{n \rightarrow \infty} \frac{1}{n} \mathcal{H}(\alpha; n). \quad (70)$$

The condition $h(\alpha) = 0$, satisfied by α_* , can be written from Eq. (67) as

$$(\alpha - \alpha_*) \log \left(\frac{\alpha_* + 1}{\alpha_*} \theta_0^k \right) = 0. \quad (71)$$

Curiously, Eq. (71) is satisfied identically in α if $\alpha_* = \theta_0^k / (1 - \theta_0^k)$. By plugging this value of α_* into Eq. (55), one obtains the simple condition $\theta_1^k = \theta_0^k (1 - \theta_0^k)$. For data structure with θ_0^k and θ_1^k satisfying this relation, one therefore

expects that $C(\alpha; n)$ is constant in α in the large- n limit; equivalently, the VC entropy will be approximately independent of the load, $\mathcal{H}_{n,p} \sim \mathcal{H}_n$.

IV. REPLICA APPROACH

The discussion in the foregoing sections shows that (1) the VC entropy has nonmonotonic behavior for simplex learning, (2) the hallmark of the nonmonotonicity is the existence of a phase transition, and (3) the transition can be framed as the SAT-UNSAT transition of a constraint satisfaction problem, which is different from the one that defines the storage capacity. Since it is often challenging to deal with the combinatorics of complex data structures, our goal in this section is to identify an appropriate synaptic volume that provides access to the transition. Once this observable is identified, we will be able to pinpoint the existence of the phase transition without direct access to the VC entropy, in the same spirit of the original work by Gardner [3], by using disordered systems techniques.

We define the synaptic volume by leveraging on the definition of the CSP corresponding to the transition. As already noted, in looking for a solution to the constraint satisfaction Problem 2 (defined in Sec. III D), we have the freedom to adjust both the synaptic weights W and the outputs σ . This means that the outputs are promoted to be dynamical variables and should be treated at the same level of the synaptic weights. This suggests that the relevant synaptic volume for identifying the corresponding phase transition is the following:

$$\begin{aligned} V(X_p) &= \sum_{\{\sigma^\mu = \pm 1\}} \int \left[\prod_{j=1}^n dW_j \right] \delta \left(\sum_{j=1}^n W_j^2 - n \right) \\ &\quad \times \prod_{\mu=1}^p \prod_{a=1}^k \theta \left(\frac{\sigma^\mu}{\sqrt{n}} \sum_{j=1}^n W_j \xi_{a,j}^\mu \right), \end{aligned} \quad (72)$$

where $\theta(\cdot)$ is the Heaviside theta, $\xi_{a,j}^\mu$ denotes the j th component of the a th element of the μ th multiplet and the weights lie on the surface of a n -dimensional sphere of radius \sqrt{n} (note that this is different from the convention used in the preceding sections). The inputs, constituting the set X_p , are chosen randomly according to the distribution

$$\begin{aligned} dP(X_p) &= v^{-1} \prod_{\mu=1}^p \prod_{a=1}^k \prod_{b=1}^{a-1} \delta \left(\rho_{ab} - \frac{1}{n} \sum_{j=1}^n \xi_{a,j}^\mu \xi_{b,j}^\mu \right) \\ &\quad \times \prod_{j=1}^n [\delta(\xi_{a,j}^\mu - 1) + \delta(\xi_{a,j}^\mu + 1)] d\xi_{a,j}^\mu, \end{aligned} \quad (73)$$

where $-1 \leq \rho_{ab} \leq 1$ are the overlaps, v is a normalization factor, and the inputs lie on the vertices of a n -dimensional hypercube.

Note that data structure is implemented in Eq. (72) by asking that each point of the μ th simplex be labeled by σ^μ . Moreover, this synaptic volume differs from the ordinary Gardner volume by the integration over the labels σ , considered dynamical variables on the same foot of the weights W . Intuitively, an exponential growth of $V(X)$ with n at fixed

load α means that, in the thermodynamic limit, at least one classification compatible with the input-label constraints can be expressed by the model; on the contrary, when $V(X)$ decreases exponentially in n then no such classification exists for $n \rightarrow \infty$. Thus, the logarithm of $V(X)$ is a suitable observable to assess the nonmonotonic behavior of the VC entropy for a given data structure.

We will apply replica theory to compute the averaged (over the inputs positions) logarithm of the synaptic volume defined in Eq. (72), in order to identify the transition. The goal will be the evaluation of the critical value of $\alpha = p/n$ where this volume changes regime, as a function of the overlaps. In the following, we will restrict to the case $k = 2$, i.e., to data organized in doublets, so that the geometry of the simplex is fully specified by a single parameter ρ ; to lighten the notation, we will omit the index $a = 1, 2$, simply denoting the doublets as $(\xi, \bar{\xi})$. Using standard integral representations for the delta and theta functions, we can write the volume of interest as

$$\begin{aligned}
 V &= \sum_{\{\sigma^\mu = \pm 1\}} \int \left[\prod_{j=1}^n dW_j \right] \int_0^{+\infty} \left[\prod_{\mu=1}^p \frac{d\lambda^\mu d\bar{\lambda}^\mu}{(2\pi)^2} \right] \\
 &\times \int_{-\infty}^{+\infty} \left[\prod_{\mu=1}^p dx^\mu d\bar{x}^\mu \right] \int_{-\infty}^{+\infty} \frac{dE}{2\pi} e^{iE(\sum_j W_j^2 - n)} \\
 &\times e^{i \sum_\mu x^\mu (\lambda^\mu - \frac{\sigma^\mu}{\sqrt{n}} \sum_j W_j \xi_j^\mu) + i \sum_\mu \bar{x}^\mu (\bar{\lambda}^\mu - \frac{\sigma^\mu}{\sqrt{n}} \sum_j W_j \bar{\xi}_j^\mu)}, \quad (74)
 \end{aligned}$$

where the auxiliary variable E enforces the spherical constraint, while the standard integral representation of the theta function is obtained via the auxiliary variables λ, x .

We dedicate the following sections to the calculation of the averaged logarithm of this volume in the annealed, replica symmetric (RS) and one-step replica symmetry-breaking (1RSB) approximations. The main results of this section, to which we address the reader not interested in the details, are Eqs. (80), (97), and (109).

A. Annealed computation

The annealed calculation is based on the substitution $\overline{\log V} \rightarrow \log \bar{V}$, so we simply need to average the volume (74) with respect to the input distribution (indicated by the overline); the details are reported in Appendix B. After a large- n expansion and the average over the inputs, the integrals in x and λ can be solved explicitly:

$$\begin{aligned}
 &\left[\sum_{\{\sigma = \pm 1\}} \int_0^{+\infty} \frac{d^2\lambda}{(2\pi)^2} \int_{-\infty}^{+\infty} d^2x e^{-\frac{1}{2}x^T \mathcal{R}x + ix^T \lambda} \right]^p \\
 &= \left[2 \int_0^{+\infty} \frac{d^2\lambda}{(2\pi)^2} \frac{2\pi}{\sqrt{1-\rho^2}} e^{-\frac{1}{2}\lambda^T \mathcal{R}^{-1}\lambda} \right]^p \\
 &= \left[\frac{1}{2} + \frac{1}{\pi} \arcsin \rho \right]^p, \quad (75)
 \end{aligned}$$

where we introduced the notation

$$\mathbf{x} = \begin{pmatrix} x \\ \bar{x} \end{pmatrix}, \quad \boldsymbol{\lambda} = \begin{pmatrix} \lambda \\ \bar{\lambda} \end{pmatrix}, \quad \mathcal{R} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \quad (76)$$

and we used the known formula for the quadrant probability of a bivariate normal distribution; see Ref. [47]. The remaining integrals can be performed: The one over the weights is Gaussian

$$\int \left[\prod_{j=1}^n dW_j \right] e^{iE \sum_j W_j^2} = e^{n[\frac{1}{2} \log \pi - \frac{1}{2} \log(-iE)]}, \quad (77)$$

while the one over E can be performed via a saddle-point method for large n :

$$\int_{-\infty}^{+\infty} \frac{dE}{2\pi} e^{-inE - \frac{n}{2} \log(-iE)} \sim \frac{1}{2\sqrt{\pi n}} e^{n[\frac{1}{2} + \frac{\log 2}{2}]}. \quad (78)$$

Assembling everything, and ignoring inessential factors, we find

$$\bar{V} = \exp \left\{ n \left[\frac{p}{n} \log \left(\frac{1}{2} + \frac{1}{\pi} \arcsin \rho \right) + \frac{1 + \log 2\pi}{2} \right] \right\}. \quad (79)$$

Defining the critical value of $\alpha = p/n$ as the one where the exponent changes sign, we find

$$\alpha_*^A(\rho) = -\frac{1 + \log 2\pi}{2 \log \left(\frac{1}{2} + \frac{1}{\pi} \arcsin \rho \right)}. \quad (80)$$

A comparison of the annealed approximation and of the result obtained with combinatorics in Eq. (45) is shown in Fig. 3. Although the annealed approximation fails in reproducing quantitatively the behavior of $\alpha_*(\rho)$, it bounds the combinatorial result from below, and qualitatively recovers the expected divergence for $\psi_2 \rightarrow 1$.

B. Quenched computation

The quenched calculation of $\overline{\log V}$ is performed via the replica trick. First, we replicate t times the volume (74),




FIG. 3. Critical value of the load α as a function of the overlap ρ for $k = 2$ (data in pairs). Circles represent the combinatorial result, which is in agreement with numerical simulations. All the different approximation schemes used for the replica computations display the same qualitative shape. However, the annealed and RS ansatz fail in reproducing quantitatively the combinatorial result. Using a 1RSB ansatz we obtain a one-parameter expression for α_* [Eq. (109)] that fits the combinatorial result tightly.

obtaining

$$\begin{aligned}
 V^t = & \sum_{\{\sigma_a^\mu = \pm 1\}} \int \left[\prod_{a=1}^t \prod_{j=1}^n dW_{j,a} \right] \int_{-\infty}^{+\infty} \left[\prod_{a=1}^t \frac{dE_a}{2\pi} \right] \\
 & \times \int_{-\infty}^{+\infty} \left[\prod_{a<b} \frac{dF_{ab} dQ_{ab}}{2\pi} \right] e^{i \sum_a E_a (\sum_j W_{j,a}^2 - n)} \\
 & \times e^{i \sum_{a<b} F_{ab} (\sum_j W_{j,a} W_{j,b} - n Q_{ab})} \int_0^{+\infty} \left[\prod_{a=1}^t \prod_{\mu=1}^p \frac{d\lambda_a^\mu d\bar{\lambda}_a^\mu}{(2\pi)^2} \right] \\
 & \times \int_{-\infty}^{+\infty} \left[\prod_{a=1}^t \prod_{\mu=1}^p dx_a^\mu d\bar{x}_a^\mu \right] e^{i \sum_{a,\mu} x_a^\mu (\lambda_a^\mu - \frac{\sigma_a^\mu}{\sqrt{n}} \sum_j W_{j,a} \xi_j^\mu)} \\
 & \times e^{i \sum_{a,\mu} \bar{x}_a^\mu (\bar{\lambda}_a^\mu - \frac{\sigma_a^\mu}{\sqrt{n}} \sum_j W_{j,a} \bar{\xi}_j^\mu)}, \quad (81)
 \end{aligned}$$

where $1 \leq a, b \leq t$ are replica indices (not to be confused with the indices running inside the multiplets, a notation we abandoned at the beginning of this section, when we specialized our calculation to doublets), Q_{ab} is the replica matrix (with $Q_{aa} = 1$), and F_{ab} are the Lagrange multipliers enforcing the constraint

$$Q_{ab} = \frac{1}{n} \sum_{j=1}^n W_{j,a} W_{j,b}. \quad (82)$$

Now we can perform the average over the input ensemble. With the same steps we used to get Eq. (75) (see Appendix B), we obtain, for the x and λ integrals,

$$\begin{aligned}
 & \left\{ \sum_{\{\sigma_a = \pm 1\}} \int_0^{+\infty} \left[\prod_{a=1}^t \frac{d^2 \lambda_a}{(2\pi)^2} \right] \int_{-\infty}^{+\infty} \left[\prod_{a=1}^t d^2 x_a \right] \right. \\
 & \left. \times e^{-\frac{1}{2} \sum_{a,b} Q_{ab} x_a^T \mathcal{R} x_b + i \sum_a \sigma_a x_a^T \lambda_a} \right\}^p, \quad (83)
 \end{aligned}$$

where we already inserted the replica matrix using (82) and we isolated the outputs σ in the source term via the transformation $\mathbf{x} \rightarrow \sigma \mathbf{x}$. The remaining integral over the weights is Gaussian:

$$\begin{aligned}
 & \int \left[\prod_{a,j} dW_{j,a} \right] e^{i \sum_a E_a \sum_j W_{j,a}^2 + i \sum_{a<b} F_{ab} \sum_j W_{j,a} W_{j,b}} \\
 & = e^{-\frac{n}{2} \log \det(-iG) + \frac{m}{2} \log(2\pi)}, \quad (84)
 \end{aligned}$$

where G is the symmetric matrix with elements

$$G_{ab} = 2E_a \delta_{ab} - (1 - \delta_{ab}) F_{ab}. \quad (85)$$

The integral over the elements of G is performed via a saddle point: Ignoring all the inessential factors,

$$\begin{aligned}
 & \int_{-\infty}^{+\infty} \left[\prod_{a=1}^t \frac{dG_{aa}}{4\pi} \right] \int_{-\infty}^{+\infty} \left[\prod_{a<b} \frac{dG_{ab}}{2\pi} \right] \\
 & \times e^{-\frac{n}{2} \sum_{a,b} iG_{ab} Q_{ab} - \frac{n}{2} \log \det(-iG)} \sim e^{\frac{m}{2} + \frac{n}{2} \log \det(Q)}, \quad (86)
 \end{aligned}$$

where we used

$$\frac{\partial}{\partial G_{ab}} \left[\sum_{c,d} iG_{cd} Q_{cd} + \log \det(-iG) \right] = iQ_{ab} + (G^{-1})_{ba}. \quad (87)$$

Finally, the resulting averaged replicated volume to be evaluated is

$$\begin{aligned}
 \bar{V}^t = & \int_{-\infty}^{+\infty} \left[\prod_{a<b} dQ_{ab} \right] e^{\frac{m}{2} + \frac{n}{2} \log \det(Q)} \\
 & \times \left\{ \sum_{\{\sigma_a = \pm 1\}} \int_0^{+\infty} \left[\prod_{a=1}^t \frac{d^2 \lambda_a}{(2\pi)^2} \right] \int_{-\infty}^{+\infty} \left[\prod_{a=1}^t d^2 x_a \right] \right. \\
 & \left. \times e^{-\frac{1}{2} \sum_{a,b} Q_{ab} x_a^T \mathcal{R} x_b + i \sum_a \sigma_a x_a^T \lambda_a} \right\}^p. \quad (88)
 \end{aligned}$$

We cannot proceed further, in taking the limit $t \rightarrow 0$ as prescribed by the replica approach, without making an ansatz on the form of the replica matrix Q_{ab} .

1. RS ansatz

In the RS ansatz, the replica matrix has the form

$$Q_{ab} = (1 - q)\delta_{ab} + q, \quad 0 \leq q \leq 1, \quad (89)$$

so that

$$\log \det(Q) \xrightarrow{t \rightarrow 0} t \log(1 - q) + \frac{tq}{1 - q}. \quad (90)$$

The quadratic form at the exponent of Eq. (88) reads

$$\begin{aligned}
 & \sum_{a,b} Q_{ab} x_a^T \mathcal{R} x_b \\
 & = (1 - q) \sum_a x_a^T \mathcal{R} x_a + q \left(\sum_a x_a \right)^T \mathcal{R} \left(\sum_b x_b \right). \quad (91)
 \end{aligned}$$

The last term can be linearized with a Hubbard-Stratonovich transformation:

$$\begin{aligned}
 & e^{-\frac{q}{2} [\sum_a x_a]^T \mathcal{R} [\sum_b x_b]} \\
 & = \int_{-\infty}^{+\infty} \frac{d^2 \mathbf{y}}{2\pi \sqrt{1 - \rho^2}} e^{-\frac{1}{2} \mathbf{y}^T \mathcal{R}^{-1} \mathbf{y} + i\sqrt{q} \sum_a x_a^T \mathbf{y}}, \quad (92)
 \end{aligned}$$

so that replica indices factorize, to get, after an integration over \mathbf{x} ,

$$\begin{aligned}
 & \left\{ \int_{-\infty}^{+\infty} \frac{d^2 \mathbf{y}}{2\pi \sqrt{1 - \rho^2}} e^{-\frac{1}{2} \mathbf{y}^T \mathcal{R}^{-1} \mathbf{y}} \left[\frac{2\pi}{(1 - q)\sqrt{1 - \rho^2}} \sum_{\{\sigma = \pm 1\}} \right. \right. \\
 & \left. \left. \times \int_0^{+\infty} \frac{d^2 \lambda}{(2\pi)^2} e^{-\frac{1}{2(1-q)} (\lambda + \sigma \sqrt{q} \mathbf{y})^T \mathcal{R}^{-1} (\lambda + \sigma \sqrt{q} \mathbf{y})} \right]^t \right\}^p. \quad (93)
 \end{aligned}$$

Defining $L_{RS}(\mathbf{y})$ the quantity in square brackets, the limit $t \rightarrow 0$ gives

$$\begin{aligned} & p \log \left\{ \int_{-\infty}^{+\infty} \frac{d^2 \mathbf{y}}{2\pi\sqrt{1-\rho^2}} e^{-\frac{1}{2} \mathbf{y}^T \mathcal{R}^{-1} \mathbf{y}} [L_{RS}(\mathbf{y})]^t \right\} \\ & \rightarrow pt \int_{-\infty}^{+\infty} \frac{d^2 \mathbf{y}}{2\pi\sqrt{1-\rho^2}} e^{-\frac{1}{2} \mathbf{y}^T \mathcal{R}^{-1} \mathbf{y}} \log[L_{RS}(\mathbf{y})]. \end{aligned} \quad (94)$$

Since we are looking for the critical value of α of the SAT-UNSAT transition of our CSP, we can just apply the standard argument by Gardner [3]: Starting with a load below the critical value and increasing the number of patterns, the set of solutions in the space of weights shrinks down to a single configuration at the transition (in the thermodynamic limit). This means that, approaching the critical point, the replicas of the vector W must be more and more correlated and therefore $q \rightarrow 1$ at the transition. In this limit, the factor $(1-q)^{-1}$ is large and the integrals in $L_{RS}(\mathbf{y})$ can be evaluated with a saddle point: We need to find the stationary points of the exponent in the integrands as a function of λ . According to the position of the vector \mathbf{y} on the plane, the saddle is in one of the three following spots: (1) inside the region of integration over λ , (2) at one of its boundaries, (3) at the origin. We obtain

$$\begin{aligned} & \int_0^{+\infty} \frac{d^2 \lambda}{2\pi(1-q)\sqrt{1-\rho^2}} e^{-\frac{1}{2(1-q)}(\lambda+\sigma\sqrt{q}\mathbf{y})^T \mathcal{R}^{-1}(\lambda+\sigma\sqrt{q}\mathbf{y})} \\ & \sim \theta(-\sigma y)\theta(-\sigma\bar{y}) + \theta(\sigma y)\theta[\sigma(\rho y - \bar{y})] \frac{e^{-\frac{y^2}{2(1-q)}}}{y} \sqrt{\frac{1-q}{8\pi}} \\ & + \theta[\sigma(\rho\bar{y} - y)]\theta(\sigma\bar{y}) \frac{e^{-\frac{\bar{y}^2}{2(1-q)}}}{\bar{y}} \sqrt{\frac{1-q}{8\pi}} + e^{-\frac{1}{2(1-q)}\mathbf{y}^T \mathcal{R}^{-1} \mathbf{y}} \\ & \times \theta[\sigma(y - \rho\bar{y})]\theta[\sigma(\bar{y} - \rho y)] \frac{1}{2\pi} \frac{(1-q)(1-\rho^2)^{3/2}}{(\bar{y} - \rho y)(y - \rho\bar{y})}, \end{aligned} \quad (95)$$

with the theta functions selecting in turn one of the above cases. In the summation over $\sigma = \pm 1$, in each domain of \mathbf{y} survives only the dominant addend in $(1-q)$: This is the finite term in the first and third quadrant, and the terms proportional to $\exp\{-y^2/[2(1-q)]\}$ or $\exp\{-\bar{y}^2/[2(1-q)]\}$ in the second and forth quadrant (the quadrants bisectors discriminating the larger). In the end, using the obvious symmetry between y and \bar{y} as integration variables and ignoring suppressed factors in $(1-q)$, we get

$$\begin{aligned} & \int_{-\infty}^{+\infty} \frac{d^2 \mathbf{y}}{2\pi\sqrt{1-\rho^2}} e^{-\frac{1}{2} \mathbf{y}^T \mathcal{R}^{-1} \mathbf{y}} \log[L_{RS}(\mathbf{y})] \\ & = \int_0^{+\infty} \frac{dy}{\pi\sqrt{1-\rho^2}} \frac{-y^2}{1-q} \int_{-\infty}^{-y} d\bar{y} e^{-\frac{1}{2}(y,\bar{y})\mathcal{R}^{-1}\begin{pmatrix} y \\ \bar{y} \end{pmatrix}} \\ & = \frac{1}{4(1-q)} \left(\frac{2}{\pi} \sqrt{1-\rho^2} - \frac{4}{\pi} \arctan \frac{\sqrt{1-\rho}}{\sqrt{1+\rho}} \right). \end{aligned} \quad (96)$$

Selecting only the most divergent terms in $(1-q)$ from (90) and (96), we have all the ingredients to evaluate the replica limit of $(\bar{V}^t - 1)/t$ for $t \rightarrow 0$. The result is zero when the load

α assumes the critical value

$$\alpha_*^{\text{RS}}(\rho) = \frac{\pi}{2 \arctan \sqrt{(1-\rho)/(1+\rho)} - \sqrt{1-\rho^2}}. \quad (97)$$

The result is reported in Fig. 3: The RS curve presents the expected limits (46), but again we do not observe quantitative agreement with the combinatorial curve. We are therefore led to conjecture that we need at least one step of replica symmetry breaking (RSB). We work out the derivation of α_* within the 1RSB ansatz in the next section.

2. 1RSB ansatz

In the 1RSB ansatz the replica matrix has the form

$$Q_{ab} = (1-q_1)\delta_{ab} + (q_1-q_0)\varepsilon_{ab} + q_0, \quad (98)$$

where $\varepsilon_{ab} = 1$ if a, b belongs to a diagonal block $m \times m$, 0 otherwise, so that

$$\begin{aligned} \log \det(Q) & \rightarrow t \left\{ \frac{m-1}{m} \log(1-q_1) \right. \\ & \left. + \frac{1}{m} \log[1-q_1+m(q_1-q_0)] + \frac{q_0}{1-q_1+m(q_1-q_0)} \right\}. \end{aligned} \quad (99)$$

From (88), we get

$$\begin{aligned} \sum_{a,b} Q_{ab} \mathbf{x}_a^T \mathcal{R} \mathbf{x}_b & = (1-q_1) \sum_a \mathbf{x}_a^T \mathcal{R} \mathbf{x}_a \\ & + (q_1-q_0) \sum_{B=0}^{t/m-1} \left(\sum_{a=1}^m \mathbf{x}_{mB+a} \right)^T \mathcal{R} \left(\sum_{b=1}^m \mathbf{x}_{mB+b} \right) \\ & + q_0 \left(\sum_a \mathbf{x}_a \right)^T \mathcal{R} \left(\sum_b \mathbf{x}_b \right), \end{aligned} \quad (100)$$

where B is a block index. We now need $2(t/m+1)$ auxiliary Hubbard-Stratonovich variables to linearize the sums over replica indices: To get, after the usual factorizations and the integration over \mathbf{x} ,

$$\begin{aligned} & \left(\int \frac{d^2 \mathbf{y}}{2\pi\sqrt{1-\rho^2}} e^{-\frac{1}{2} \mathbf{y}^T \mathcal{R}^{-1} \mathbf{y}} \left\{ \int \frac{d^2 \mathbf{z}}{2\pi\sqrt{1-\rho^2}} e^{-\frac{1}{2} \mathbf{z}^T \mathcal{R}^{-1} \mathbf{z}} \right. \right. \\ & \times \left[\sum_{\{\sigma=\pm 1\}} \int_0^{+\infty} \frac{d^2 \lambda}{2\pi(1-q_1)\sqrt{1-\rho^2}} \right. \\ & \left. \left. \left. \times e^{-\frac{[\sigma(\sqrt{q_1-q_0}\mathbf{z}+\sqrt{q_0}\mathbf{y})+\lambda]^T \mathcal{R}^{-1}[\sigma(\sqrt{q_1-q_0}\mathbf{z}+\sqrt{q_0}\mathbf{y})+\lambda]}{2(1-q_1)}} \right] \right\}^m \right)^{\frac{t}{m} P}. \end{aligned} \quad (101)$$

Defining $L_{1RSB}(\mathbf{y})$ the argument of the square brackets, we know that the logarithm of the above quantity for $t \rightarrow 0$ gives

$$\frac{pt}{m} \int \frac{d^2 \mathbf{y}}{2\pi\sqrt{1-\rho^2}} e^{-\frac{1}{2} \mathbf{y}^T \mathcal{R}^{-1} \mathbf{y}} \log[L_{1RSB}(\mathbf{y})]. \quad (102)$$

To simplify $L_{1\text{RSB}}(\mathbf{y})$ and to get an expression similar to the one we studied before, we can shift the \mathbf{z} variables to

$$\mathbf{z} \rightarrow \mathbf{z} - \frac{\sqrt{q_0}}{\sqrt{q_1 - q_0}} \mathbf{y}, \quad (103)$$

obtaining

$$\begin{aligned} L_{1\text{RSB}}(\mathbf{y}) &= \int \frac{d^2 \mathbf{z} e^{-\frac{1}{2} \left(\mathbf{z} - \frac{\sqrt{q_0}}{\sqrt{q_1 - q_0}} \mathbf{y} \right)^T \mathcal{R}^{-1} \left(\mathbf{z} - \frac{\sqrt{q_0}}{\sqrt{q_1 - q_0}} \mathbf{y} \right)}}{2\pi \sqrt{1 - \rho^2}} \\ &\times \left[\sum_{\{\sigma = \pm 1\}} \int_0^{+\infty} \frac{d^2 \lambda e^{-\frac{[\sigma \sqrt{q_1 - q_0} \mathbf{z} + \lambda]^T \mathcal{R}^{-1} [\sigma \sqrt{q_1 - q_0} \mathbf{z} + \lambda]}}{2\pi (1 - q_1) \sqrt{1 - \rho^2}} \right]^m. \end{aligned} \quad (104)$$

In order to find the critical load, we investigate the behavior of the 1RSB parameters close to the transition: It turns out that q_1 has to be sent to 1 (in analogy with the RS case) and m to 0 [48] as

$$q_1 \rightarrow 1, \quad m \rightarrow (1 - q_1)w, \quad (105)$$

with w a finite parameter. In this limit we can evaluate the integral over λ with a saddle point. We get

$$\theta(z)\theta(\bar{z}) + \theta(-z)\theta(-\bar{z}) + 4\theta(z)\theta(-\bar{z} - z)e^{-\frac{w(1-q_0)z^2}{2}}. \quad (106)$$

Analytical computations are rather cumbersome after this point. However, the result simplifies a lot if we take $q_0 = 0$. Then the integral over \mathbf{y} decouples and simply gives 1, while the one over \mathbf{z} breaks into the regions

$$\int_0^{+\infty} \frac{d^2 \mathbf{z}}{\pi \sqrt{1 - \rho^2}} e^{-\frac{1}{2} \mathbf{z}^T \mathcal{R}^{-1} \mathbf{z}} = \frac{1}{2} + \frac{1}{\pi} \arcsin(\rho) \quad (107)$$

and

$$\begin{aligned} \int_0^{+\infty} \frac{2dz}{\pi \sqrt{1 - \rho^2}} \int_{-\infty}^{-z} d\bar{z} e^{-\frac{1}{2} \mathbf{z}^T \mathcal{R}^{-1} \mathbf{z} - \frac{wz^2}{2}} \\ = \frac{2 \arctan \left[\sqrt{(1+w) \frac{1-\rho}{1+\rho}} \right]}{\pi \sqrt{1+w}}. \end{aligned} \quad (108)$$

In the end, we find

$$\begin{aligned} \alpha_*^{\text{1RSB}}(\rho; q_0 = 0, w) \\ = \frac{-\log[1+w]}{2 \log \left[\frac{1}{2} + \frac{1}{\pi} \arcsin(\rho) + \frac{2 \arctan \left(\sqrt{(1+w) \frac{1-\rho}{1+\rho}} \right)}{\pi \sqrt{1+w}} \right]}. \end{aligned} \quad (109)$$

We stress that this last result is not the optimal 1RSB solution: In principle we should consider the full expression of $\alpha_*^{\text{1RSB}}(\rho; q_0, w)$ and optimize upon the remaining parameters q_0 and w . However, this is beyond the scope of this section: Here we simply verify that the functional form $\alpha_*^{\text{1RSB}}(\rho; q_0 = 0, w)$ allows to fit nicely the combinatorial result, by adjusting the parameter w (see Fig. 3). This simple observation strongly supports our conjecture that this SAT-UNSAT transition exhibits at least one step of RSB, but it does not rule out a full-RSB scenario.

C. Margin learning

Replica theory turns out to be essential to explore the role of data structure whenever alternative, *ad hoc* methods (such as the combinatorial one) are not available. Here we apply it to identify the SAT-UNSAT transition occurring in margin learning. The synaptic volume relevant to this case is

$$\begin{aligned} V_\kappa &= \sum_{\{\sigma^\mu = \pm 1\}} \int \left[\prod_{j=1}^n dW_j \right] \delta \left(\sum_{j=1}^n W_j^2 - n \right) \\ &\times \prod_{\mu=1}^p \theta \left(\frac{\sigma^\mu}{\sqrt{n}} \sum_{j=1}^n W_j \xi_j^\mu - \kappa \right), \end{aligned} \quad (110)$$

where κ is the margin. Note again that here, as in the case of Eq. (72), the outputs σ^μ are dynamical variables, at variance with the usual Gardner's volume. We skip the details on the annealed and quenched calculations, which are in spirit very similar to those of the previous sections. Nonetheless, it is worth to point out that the tricky multivariate integrals in the auxiliary variable, are now replaced by Gaussian integrals, with the margin κ appearing as an integration limit. The annealed approximation leads to

$$\alpha_*^\Lambda(\kappa) = -\frac{1 + \log(2\pi)}{2 \log[2 \operatorname{erfc}(\kappa)]}. \quad (111)$$

In the quenched calculation, the RS ansatz is again implemented by requiring $q \rightarrow 1$; one obtains the critical threshold

$$\alpha_*^{\text{RS}}(\kappa) = \frac{1}{2} \left[\int_0^\kappa \operatorname{Dy} (\kappa - y)^2 \right]^{-1}, \quad (112)$$

where Dy is the Gaussian measure. Note the difference with Gardner's result [3] for the storage capacity,

$$\alpha_c(\kappa) = \left[\int_{-\kappa}^{+\infty} \operatorname{Dy} (\kappa + y)^2 \right]^{-1}. \quad (113)$$

The one-step RSB ansatz again depends on the parameters q_0 and w , which should be investigated numerically. However, in the special case $q_0 = 0$ we find the simpler expression

$$\begin{aligned} \alpha_*^{\text{1RSB}}(\kappa; q_0 = 0, w) \\ = \frac{-\log[1+w]}{2 \log \left\{ 2 \left[\operatorname{erfc}(\kappa) + \int_0^\kappa \operatorname{Dz} e^{-w \frac{(z-\kappa)^2}{2}} \right] \right\}}. \end{aligned} \quad (114)$$

These results essentially share the same features of those for the simplexes computed above: In particular, at variance with the usual storage capacity (113), α_* computed in all the different approximation schemes diverges in the limit $\kappa \rightarrow 0^+$, when the problem reduces to a standard classification of points (or equivalently, in the object manifold description, when the radius of the spheres shrinks to zero). Even in absence of a closed expression for the VC entropy of margin classification, the existence of the phase transition at a finite load is a clear indication of its nonmonotonicity.

V. DISCUSSION

Understanding how data specificities impact the performance of machine learning models and algorithms can be considered one of the major challenges for contemporary statistical physics. Here we have shown how to deal with data structure, as it is being established in physics, within the framework of the statistical theory of learning. The presence of input-output correlations in a data set suggests constraints to be applied to the hypothesis class under consideration. As a result, the corresponding VC entropy, deeply connected to the generalization capabilities of the model, is considerably lower than in the unstructured case.

For simple models of data structure we have observed two striking phenomena that take place above the VC dimension. First, the VC entropy becomes nonmonotonic. This is a strong indication that the rigorous bounds in SLT may be substantially improved by taking data structure into account. Second, a transition appears beyond the well-known storage capacity, at the onset of unsatisfiability for a data-related constraint satisfaction problem. When available, a combinatorial theory *à la Cover* allows one to compute the VC entropy of a finite-size system, and to reveal explicitly its nonmonotonic behavior. However, this is not always feasible, such as for spherical object manifolds and margin learning. In these cases, we showed how the phase transition can be probed with the standard tools of statistical physics, thus allowing an indirect quantification of the data-dependent behavior.

Our satisfiability transition is due to a competition between the combinatorial expansion, with sample size, of the space of possible functions and the reduction due to the constraints [35]. We believe, as this observation suggests, that the emergence of the data-driven transition, as well as the nonmonotonic VC entropy it entails, is not specific to the two models of data that we have studied here, but is more generally present whenever the constraints imposed on the hypothesis class by data structure are strong enough. On a more quantitative level, notice that the upper and lower bounds obtained for α_* in Sec. III E are very close to one another. The bounds are independent of the particular choice of simplexes, i.e., they do not depend on k or on $\{\rho_{ab}\}$. This is a clue pointing to the robustness of the phenomenology for disparate data structures. We remark that the combinatoric analysis was done at leading order in α ; thus, it remains to assess how much the bounds are affected by perturbative corrections.

An ambitious and pressing goal concerns the generalization of our results to other architectures, notably deep neural networks, in the same spirit of what was achieved in SLT regarding the VC dimension.

ACKNOWLEDGMENTS

The authors would like to thank Enrico Malatesta and Marco Cosentino Lagomarsino for useful discussions and suggestions. P.R. acknowledges funding from the Fellini program under the H2020-MSCA-COFUND action, Grant Agreement No. 754496, INFN (IT).

APPENDIX A: BOUNDS ON θ_1^k

In this Appendix we report the details of the calculation of the bounds on θ_1^k given in Sec. III E. To briefly recall the

definitions, we have

$$\theta_0^k = \prod_{m=2}^k \psi_m, \quad \theta_1^k = \left(2 - k + \sum_{m=2}^k \frac{1}{\psi_m} \right) \prod_{m=2}^k \psi_m \quad (A1)$$

for some $0 \leq \psi_m \leq 1, \forall m \geq 2$. We want to compute the infimum and the supremum of θ_1^k at fixed θ_0^k , as a function of the ψ variables. First of all, let us simplify the notation. Define

$$x_m := \psi_{m+1}, \quad \forall m \geq 1, \\ f_{(k)}(x_1 \dots x_k) := \theta_1^{k+1}(\psi_2 \dots \psi_{k+1}), \quad \forall k \geq 1. \quad (A2)$$

Explicitly,

$$f_{(k)}(\vec{x}) = \left(1 - k + \sum_{m=1}^k \frac{1}{x_m} \right) \prod_{m=1}^k x_m, \quad (A3)$$

where $\vec{x} = (x_1 \dots x_k)$.

Our problem is to optimize (i.e., to find the infimum and the supremum) $f_{(k)}(\vec{x})$ in the hypercube $\vec{x} \in [0, 1]^k$, subject to the constraint

$$\prod_{m=1}^k x_m = t \in [0, 1]. \quad (A4)$$

We will prove by induction that

$$\sup_{\vec{x} \in [0, 1]^k} f_{(k)}(\vec{x}) = 1 - \delta_{t,0}, \quad \forall k \geq 1, \\ \inf_{\vec{x} \in [0, 1]^k} f_{(k)}(\vec{x}) = \phi(k, t), \quad \forall k \geq 1, \quad (A5)$$

where $\phi(k, t) = (1 - k)t + kt^{1-\frac{1}{k}}$. Notice that $\phi(k, t)$ is a monotone decreasing function of k , and it is always less than 1.

The case $t = 0$ is special, as the constraint restricts the domain to the origin and $f_{(k)}$ is null; in the following, suppose that $t > 0$.

If $k = 1$, the constraint implies that $x_1 = t$, so that $f_{(1)}(x_1) = f_{(1)}(t) = 1$. The fact that $\phi(1, t) = 1$ proves that the proposed bounds are indeed true.

If $k > 1$, we first look for critical points inside $[0, 1]^k$ using Lagrange's theorem for constrained optimization; then we optimize our function on the boundary of $[0, 1]^k$ to look for noncritical extrema:

(1) Inside the domain, Lagrange's theorem gives that $\vec{x}_* = (t^{\frac{1}{k}} \dots t^{\frac{1}{k}})$ is the only critical point, and $f_{(k)}(\vec{x}_*) = \phi(k, t)$;

(2) On the boundary, we have that at least one of the x variables (without loss of generality, let us take x_k to be this boundary variable) must be either 0 or 1; the former is not compatible with the constraint as $t > 0$, so $x_k = 1$. But $f_{(k)}(x_1 \dots x_{k-1}, 1) = f_{(k-1)}(x_1 \dots x_{k-1})$, and $t = \prod_{m=1}^k x_m = \prod_{m=1}^{k-1} x_m$, so that the constrained optimization of $f_{(k)}(\vec{x})$ on the boundary of the domain is equivalent to the constrained optimization of $f_{(k-1)}(\vec{x})$ on the full domain $[0, 1]^{k-1}$.

Thus, the candidates for the infimum and the supremum of $f_{(k)}(\vec{x})$ are given by $\phi(k)$ (inside the domain, by Lagrange's theorem) and 1, $\phi(k - 1, t)$ (on the boundary of the domain, by induction hypothesis). The properties of ϕ imply that 1 is the supremum and $\phi(k, t)$ is the infimum of $f_{(k)}(\vec{x})$.

Finally, again by induction, we see that the supremum is realized on the point $(t, 1, \dots)$ and by all the distinct permutations of its coordinates, and that the infimum is realized by $(t^{\frac{1}{k}} \dots t^{\frac{1}{k}})$.

APPENDIX B: AVERAGING OVER THE INPUT DISTRIBUTION

In this Appendix we report the details of the calculation of the averages over the input ensemble, performed in Sec. IV. From Eq. (73), specialized for $k = 2$, we observe that at fixed overlap ρ , given $c, d \in \mathbb{N}$ the numbers of concordant and discordant signs of the components of the pair for each μ , then $c - d = \rho n, c + d = n$, so

$$c = (1 + \rho)n/2, \quad d = (1 - \rho)n/2. \quad (\text{B1})$$

For each μ , we can freely choose in 2^n different ways the components of ξ^μ , but then for $\bar{\xi}^\mu$ we must take c components with the same sign of their counterparts and d with the opposite. We can do that in $\binom{n}{c}$ different ways, so the normalization factor is

$$v = 2^m \binom{n}{\frac{(1+\rho)n}{2}}. \quad (\text{B2})$$

However, the order of the components of the vectors $\xi, \bar{\xi}$ is completely irrelevant, because they appear only in scalar products, among themselves (in the overlap constraint) and with the same vector W , whose components again we are free to relabel. This means that we can choose as a representative of the vector $\bar{\xi}$, for example, the one with the concordant components at the beginning. We can write the ensemble measure as

$$dP_\rho(\Xi) = \prod_{\mu=1}^p \left[\prod_{j=1}^c dP(\xi_j^\mu) \delta(\xi_j^\mu - \bar{\xi}_j^\mu) d\bar{\xi}_j^\mu \right] \times \left[\prod_{j=c+1}^n dP(\xi_j^\mu) \delta(\xi_j^\mu + \bar{\xi}_j^\mu) d\bar{\xi}_j^\mu \right], \quad (\text{B3})$$

where

$$dP(\xi_j^\mu) = \frac{1}{2} [\delta(\xi_j^\mu - 1) + \delta(\xi_j^\mu + 1)] d\xi_j^\mu. \quad (\text{B4})$$

Note that with the choice of a representative we are explicitly breaking the invariance of the original expression under

permutation (relabeling) of the indices j , a symmetry we will reintroduce by hand in the following calculation.

We can now perform the averages of the volumes (74) and (81). We report only the annealed calculation, the quenched one being a straightforward variation. Isolating the only part depending on the inputs in the integrand of Eq. (74), we find

$$\begin{aligned} \int dP_\rho(\Xi) e^{-i \sum_\mu \sigma^\mu x^\mu \sum_j \frac{\xi_j^\mu W_j}{\sqrt{n}} - i \sum_\mu \sigma^\mu \bar{x}^\mu \sum_j \frac{\bar{\xi}_j^\mu W_j}{\sqrt{n}}} \\ = \prod_{\mu=1}^p \prod_{j=1}^c \cos \left[\frac{1}{\sqrt{n}} (x^\mu + \bar{x}^\mu) \sigma^\mu W_j \right] \\ \times \prod_{j=c+1}^n \cos \left[\frac{1}{\sqrt{n}} (x^\mu - \bar{x}^\mu) \sigma^\mu W_j \right] \\ \approx \prod_{\mu=1}^p e^{-\frac{1}{2} (x^\mu)^2 \sum_{j=1}^n \frac{W_j^2}{n} - \frac{1}{2} (\bar{x}^\mu)^2 \sum_{j=1}^n \frac{W_j^2}{n}} \\ \times e^{-x^\mu \bar{x}^\mu (\sum_{j=1}^c - \sum_{j=c+1}^n) \frac{W_j^2}{n}}, \end{aligned} \quad (\text{B5})$$

where, in the final step, a large n expansion is performed. The exponent of the last term, consisting in a sum over j that does not extend over all the n components, cannot be readily solved using the spherical constraint, but we can write it as

$$\left(\sum_{j=1}^c - \sum_{j=c+1}^n \right) \frac{W_j^2}{n} = \left(2 \sum_{j=1}^c - \sum_{j=1}^n \right) \frac{W_j^2}{n}. \quad (\text{B6})$$

Now, only the first sum is not invariant under permutations of the components. However, since the starting point was symmetric, we can also multiply this expression by similar ones obtained with other choices of the vector $\bar{\xi}^\mu$, and then take the corresponding root of the result, obtaining an equivalent formula. The trick to restore a complete sum over the n components, is to multiply by all the c -permutations of n , and then take the $n!/(n-c)!$ -th root of the result. The only nontrivial term at the exponent during this procedure is indeed the partial sum, which reads

$$\frac{(n-c)!}{n!} \sum_{j=1}^c \sum_{\substack{\pi_1 \neq \pi_2 \neq \dots \neq \pi_c \\ \forall i, 1 \leq \pi_i \leq n}} \frac{W_{\pi_j}^2}{n} = \frac{c}{n} \sum_{i=1}^n \frac{W_i^2}{n}. \quad (\text{B7})$$

Now the spherical constraint can be invoked on all terms. Using $(2c/n - 1) = \rho$, and factorizing the p integrals over the auxiliary variables x and λ , we obtain Eq. (75).

[1] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Phys. Rev. A* **32**, 1007 (1985).
 [2] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Phys. Rev. Lett.* **55**, 1530 (1985).
 [3] E. Gardner, *Europhys. Lett.* **4**, 481 (1987).
 [4] E. Gardner and B. Derrida, *J. Phys. A: Math. Gen.* **21**, 271 (1988).

[5] S. Y. Chung, D. D. Lee, and H. Sompolinsky, *Phys. Rev. X* **8**, 031003 (2018).
 [6] S. Y. Chung, D. D. Lee, and H. Sompolinsky, *Phys. Rev. E* **93**, 060301(R) (2016).
 [7] U. Cohen, S. Chung, D. D. Lee, and H. Sompolinsky, *Nat. Commun.* **11**, 746 (2020).
 [8] M. Mézard, *Phys. Rev. E* **95**, 022117 (2017).

- [9] A. Mazzolini, M. Gherardi, M. Caselle, M. Cosentino Lagomarsino, and M. Osella, *Phys. Rev. X* **8**, 021023 (2018).
- [10] A. Mazzolini, J. Grilli, E. De Lazzari, M. Osella, M. C. Lagomarsino, and M. Gherardi, *Phys. Rev. E* **98**, 012315 (2018).
- [11] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová, [arXiv:1909.11500](https://arxiv.org/abs/1909.11500) (2019).
- [12] F. Gerace, B. Loureiro, F. Krzakala, M. Mézard, and L. Zdeborová, [arXiv:2002.09339](https://arxiv.org/abs/2002.09339) (2020).
- [13] V. Erba, M. Gherardi, and P. Rotondo, *Sci. Rep.* **9**, 17133 (2019).
- [14] A. Ansuini, A. Laio, J. H. Macke, and D. Zoccolan, [arXiv:1905.12784](https://arxiv.org/abs/1905.12784).
- [15] E. Facco, M. d'Errico, A. Rodriguez, and A. Laio, *Sci. Rep.* **7**, 12140 (2017).
- [16] V. Erba, S. Ariosto, M. Gherardi, and P. Rotondo, *Phys. Rev. E* **102**, 012306 (2020).
- [17] V. Vapnik, *The Nature of Statistical Learning Theory* (Springer-Verlag, New York, 2000).
- [18] O. Bousquet, S. Boucheron, and G. Lugosi, in *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2-14, 2003, Tübingen, Germany, August 4-16, 2003, Revised Lectures*, edited by O. Bousquet, U. von Luxburg, and G. Rätsch (Springer, Berlin, 2004), pp. 169–207.
- [19] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, in *Proceedings of the 5th International Conference on Learning Representations* (OpenReview.net, Toulon, France, 2017).
- [20] C. H. Martin and M. W. Mahoney, [arXiv:1710.09553](https://arxiv.org/abs/1710.09553) (2017).
- [21] L. Bottou, *Making Vapnik–Chervonenkis Bounds Accurate* (Springer International Publishing, Cham, 2015), pp. 143–155.
- [22] D. Cohn and G. Tesauro, *Neural Comput.* **4**, 249 (1992).
- [23] A. Antos, B. Kégl, T. Linder, and G. Lugosi, *J. Mach. Learn. Res.* **3**, 73 (2003).
- [24] B. Kégl, T. Linder, and G. Lugosi, in *Computational Learning Theory*, edited by D. Helmbold and B. Williamson (Springer, Berlin, 2001), pp. 368–384.
- [25] J. Shawe-Taylor, P. L. Bartlett, R. C. Williamson, and M. Anthony, *IEEE Trans. Inf. Theory* **44**, 1926 (1998).
- [26] T. M. Cover, *IEEE Trans. Electron. Comput.* **EC-14**, 326 (1965).
- [27] V. N. Vapnik, *IEEE Trans. Neural Netw.* **10**, 988 (1999).
- [28] B. M. McCoy, *Advanced Statistical Mechanics* (Oxford University Press, Oxford, 2010).
- [29] S. Caracciolo, A. Di Gioacchino, M. Gherardi, and E. M. Malatesta, *Phys. Rev. E* **97**, 052109 (2018).
- [30] P. Rotondo, L. G. Molinari, P. Ratti, and M. Gherardi, *Phys. Rev. Lett.* **116**, 256803 (2016).
- [31] S. Caracciolo and A. Sportiello, *J. Phys. A: Math. Gen.* **35**, 7661 (2002).
- [32] P. Rotondo, M. C. Lagomarsino, and M. Gherardi, *Phys. Rev. Res.* **2**, 023169 (2020).
- [33] C. Cortes and V. Vapnik, *Mach. Learn.* **20**, 273 (1995).
- [34] F. Borra, M. C. Lagomarsino, P. Rotondo, and M. Gherardi, *J. Phys. A: Math. Theor.* **52**, 384004 (2019).
- [35] P. Rotondo, M. Pastore, and M. Gherardi, *Phys. Rev. Lett.* **125**, 120601 (2020).
- [36] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, *Phys. Rep.* **810**, 1 (2019).
- [37] E. D. Sontag, NATO ASI Ser. F Comput. Sys. Sci. **168**, 69 (1998).
- [38] P. L. Bartlett and S. Mendelson, *J. Mach. Learn. Res.* **3**, 463 (2003).
- [39] A. Abbata, B. Aubin, F. Krzakala, and L. Zdeborová, [arXiv:1912.02729](https://arxiv.org/abs/1912.02729) (2019).
- [40] D. Anguita, A. Ghio, L. Oneto, and S. Ridella, *IEEE Trans. Neural Netw. Learn. Syst.* **25**, 2202 (2014).
- [41] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew, *Neurocomputing* **70**, 489 (2006).
- [42] F. Anselmi, J. Z. Leibo, L. Rosasco, J. Mutch, A. Tacchetti, and T. Poggio, *Theor. Comput. Sci.* **633**, 112 (2016).
- [43] H. S. Seung and D. D. Lee, *Science* **290**, 2268 (2000).
- [44] P. Flajolet and R. Sedgewick, *Analytic Combinatorics* (Cambridge University Press, Cambridge, 2009).
- [45] S. Kirkpatrick and B. Selman, *Science* **264**, 1297 (1994).
- [46] M. Leone, F. Ricci-Tersenghi, and R. Zecchina, *J. Phys. A: Math. Gen.* **34**, 4615 (2001).
- [47] S. S. Gupta, *Ann. Math. Stat.* **34**, 792 (1963).
- [48] C. Baldassi, E. M. Malatesta, and R. Zecchina, *Phys. Rev. Lett.* **123**, 170602 (2019).