Decomposing neural networks as mappings of correlation functions

Kirsten Fischer,[1, 2, *] Alexandre René,[1, 3, 4] Christian Keup,[1, 2] Moritz Layer,[1, 2] David Dahmen,[1] and Moritz Helias[1, 4]

[1]*Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6) and JARA-Institute Brain Structure-Function Relationships (INM-10), Jülich Research Centre, Jülich, Germany*
[2]*RWTH Aachen University, Aachen, Germany*
[3]*Department of Physics, University of Ottawa, Ottawa, Canada*
[4]*Department of Physics, Faculty 1, RWTH Aachen University, Aachen, Germany*
(˙Dated: February 11, 2022)

Understanding the functional principles of information processing in deep neural networks continues to be a challenge, in particular for networks with trained and thus non-random weights. To address this issue, we study the mapping between probability distributions implemented by a deep feed-forward network. We characterize this mapping as an iterated transformation of distributions, where the non-linearity in each layer transfers information between different orders of correlation functions. This allows us to identify essential statistics in the data, as well as different information representations that can be used by neural networks. Applied to an XOR task and to MNIST, we show that correlations up to second order predominantly capture the information processing in the internal layers, while the input layer also extracts higher-order correlations from the data. This analysis provides a quantitative and explainable perspective on classification.

I. INTRODUCTION

Recent years have shown a great success of deep neural networks in solving a wide range of tasks, from image recognition [1] to playing Go [2]. One major branch is supervised learning, where input-output mappings are learned from examples. In many common problems the target output values are given by a finite set, defining a classification task [3]. The objective then is to minimize an error measure between the correct class label and the prediction made by the neural network with respect to the joint probability distribution of data samples and class labels [4]. Thus, training dynamics, and consequently the solution strategy implemented by the network, depend on this probability distribution and the information it encodes. In this view, a network implements a transformation of the input distribution with the objective to concentrate the output distribution around the assigned target values for each class. How such a transformation is achieved and how the network training depends on the statistics of the presented data is, however, still mostly unknown.

To render the decision-making process of neural networks transparent, a profound understanding regarding their functional principles and extraction of meaningful features from given data is required. Over the past years the discrepancy between success in applications and limited understanding has lead to an increased interest also in the theoretical community [4–9]. An important line of theoretical research investigates ensembles of neural networks in the limit of infinite width for which the central limit theorem implies an exact equivalence to Gaussian processes (GPs) [10–13]. While this approach is informative with respect to how relations between data samples

are transformed by the network, it does not reveal how the internal structure of data samples is processed. As an example, for image classification the Gaussian process view takes into account the relation between all corresponding pixels x_i^α, x_i^β of any pair of images α, β in the form of a scalar product $\sum_i x_i^\alpha x_i^\beta$ – consequently, this approach is insensitive to the structure within individual images determined by e.g. correlations between pixel values x_i^α and x_j^α. In particular, due to the rotational symmetry of the scalar product, the GP view gives identical results when image pixels are shuffled consistently across all images. However, clearly the internal structure of data samples also contains important information that may be employed to solve a given task. The focus of the present study is to investigate how this information is extracted from the data and utilized by the network to perform classification.

Other approaches [14–16], similarly as GPs, focus on ensembles of networks with randomly drawn weights. In contrast, here we study how particular realizations of trained and untrained networks process different statistical features of the data. Thus, we shift the perspective from distributions over network parameters to distributions over the data. In particular, we describe the input-output mapping implemented by deep neural networks in terms of correlation functions. To trace the transformation of correlation functions across layers of neural networks, we make use of methods from statistical field theory [17–20] and obtain recursive relations in a perturbative manner by means of Feynman diagrams. Our results yield a characterization of the network as a non-linear mapping of correlation functions, where each layer exchanges information between different statistical orders. Re-expressing the loss function in terms of data correlations allows us to study their role in the training process, and to link the transformation of data correlations to the solution strategies found by the network. For the particular example of the mean-squared error loss

function, we show that network training relies exclusively on the first two cumulants of the output (mean and covariance), while these, in turn, are predominantly determined by means and covariances of network activations in previous layers. Furthermore, we show that corrections from higher-order correlations to mean and covariance, which are readily computable with the proposed generic field-theoretical framework, are of greatest importance in the first layer, where these corrections effectuate the information flow from higher-order correlations to mean and covariance.

The structure of this study is as follows: Section II provides theoretical background on the definition and architecture of deep neural networks (Section II A), on empirical risk minimization in the context of classification (Section II B), and on field-theoretical descriptions of probability distributions in terms of cumulants and their generating function (Section II C). In Section III we decompose the network mapping into correlation functions, tracing their transformations backwards through the network. We start by relating the loss to the first- and second-order correlations of the network outputs (Section III A), then discuss the mapping of correlations by individual hidden layers (Section III B), and end with the extraction of data correlations by the input layer (Section III C). Section IV applies these theoretical tools to several example datasets. We start with an adaptation of the XOR problem, where the input statistics are fully known and selectively presented to the network to study different encoding and processing schemes of class identities (Section IV A). We proceed with an application to the MNIST dataset [21], where we show that classification performance is largely based on the transformation of means and covariances across layers (Section IV B). Finally, we showcase the importance of higher-order correlations and their extraction in the input layer by constructing a dataset where information on class identity is only encoded in correlations of third and higher order (Section IV C). In Section V we discuss our results and provide an outlook.

FIG. 1. **(a) Network analysis based on data samples** considers each data sample $x^{(d)}$ separately as it passes through the network, producing a single corresponding output $y^{(d)}$. Each layer consists of an affine transformation (W^l, b^l) followed by a non-linearity ϕ applied componentwise. **(b) Network analysis based on data statistics** considers how the entire data distribution $p(x)$ is transformed by the network. At each step, the intermediate distribution is parameterized by its cumulants, the most important of which are the mean μ and the covariance Σ. The affine step transforms μ and Σ independently, while the non-linearity ϕ causes a non-trivial interaction of the two.

II. THEORETICAL BACKGROUND

A. Feed-forward network architecture

We consider fully-connected neural networks with L layers of N_l neurons each, and one additional linear readout layer, as shown in Fig. 1a. Each layer $l = 1, \ldots, L$ consists of an affine transformation

$$z_i^l = \sum_{j=1}^{N_{l-1}} W_{ij}^l \, y_j^{l-1} + b_i^l \tag{1}$$

parameterized by a weight matrix $W^l \in \mathbb{R}^{N_l \times N_{l-1}}$ and bias vector $b^l \in \mathbb{R}^{N_l}$. This step is followed by the pointwise
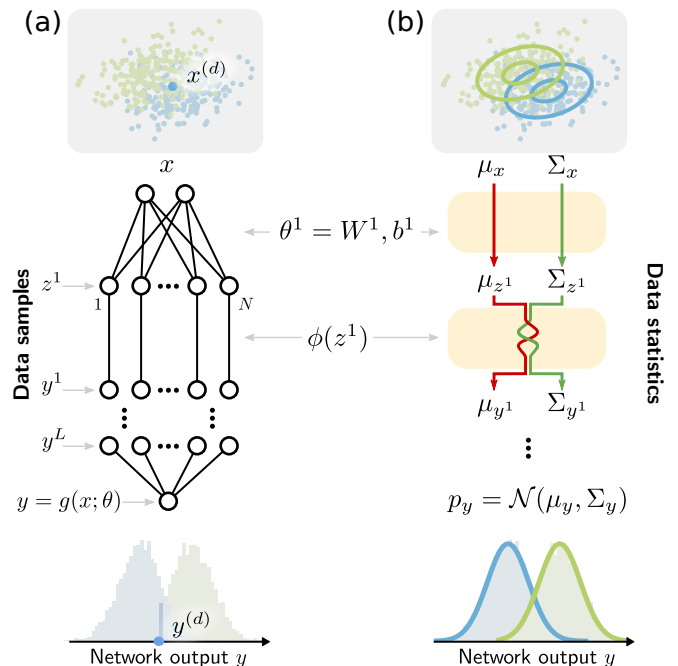
application of a non-linear activation function ϕ, yielding

$$y_i^l = \phi\big(z_i^l\big) = \phi\left(\sum_{j=1}^{N_{l-1}} W_{ij}^l \, y_j^{l-1} + b_i^l \right). \tag{2}$$

Here $y^0 = x \in \mathbb{R}^{N_0}$ denotes the input data of dimension N_0. The readout layer produces the network output $y \in \mathbb{R}^{d_{\mathrm{out}}}$, specifically $y_i = z_i^{L+1}$. The network mapping $y = g(x; \theta)$ is given by iterating over network layers and characterized by parameters $\theta := \{W^l, b^l\}_{l=1, \ldots, L+1}$.

We initialize all network parameters randomly from i.i.d. centered Gaussians $W_{ij}^l \overset{\text{i.i.d.}}{\sim} \mathcal{N}\big(0, \sigma_w^2/N_{l-1}\big)$ and $b_i^l \overset{\text{i.i.d.}}{\sim} \mathcal{N}\big(0, \sigma_b^2\big)$ before training.

B. Learning theory: empirical risk minimization

The fundamental assumption underlying classification[1] is the existence of a joint distribution $p(x, t)$ of data samples x and class labels t that is the same for training and evaluation [3]. By Bayes' theorem, the distribution of the input data can be treated as a mixture model $p(x) = \sum_t p(t) p(x|t)$. The network's task is then to implement a mapping $g : x \mapsto y$ that minimizes the expectation of a loss $\ell(y, t)$ between the network outputs $y = g(x; \theta)$ and the labels t.

This mapping, in turn, induces a mapping of the probability distributions

$$p(x|t) \mapsto p(y|t; \theta) = \int \delta(y - g(x; \theta)) p(x|t) \, dx \quad (3)$$

for each label t, where $\delta(.)$ refers to the Dirac delta distribution. The unconditioned output distribution is then the weighted sum $p(y) = \sum_t p(t) p(y|t; \theta)$. Ideally, the network output y matches the true label t so that the target distribution is given by $p(y|t) = \delta(y - t)$.

Training algorithms seek to minimize the expected loss or risk functional [22]

$$R(\theta) = \langle \ell(y, t) \rangle_{y|\theta} = \sum_t p(t) \langle \ell(y, t) \rangle_{y|t; \theta}, \quad (4)$$

where the expectation value $\langle \cdot \rangle_{y|t; \theta}$ is taken with regard to the class-conditional output distributions $p(y|t; \theta)$. In general, neither the mixture components of the input distribution $p(x|t)$ nor the induced class-conditional output distributions $p(y|t; \theta)$ are known. Instead, the expected loss is replaced by the empirical loss or risk

$$R_{\mathrm{emp}}(\theta) = \frac{1}{D} \sum_{d=1}^{D} \ell(g(x^{(d)}; \theta), t^{(d)}), \quad (5)$$

evaluated for a training set $\{(x^{(d)}, t^{(d)})\}_d$, with D being its size and d the respective sample index. The *empirical risk minimization principle* then assumes the following: the mapping $g(\cdot; \theta^*)$ that minimizes the empirical risk $\theta^* = \arg\min_\theta R_{\mathrm{emp}}(\theta)$ yields an expected risk $R(\theta^*)$ that is close to its minimum $\min_\theta R(\theta)$ [22].

C. Parameterization of probability distributions in terms of cumulants

Neural networks can be regarded as complex systems with many interacting components. A common approach to investigate such systems is by studying generating functions of moments or cumulants rather than the probability distributions themselves. Cumulants often provide a more convenient parameterization of probability

––––––––––

[1] Since this work focuses on classification tasks, we tailor the presentation of empirical risk minimization to that context.

distributions as they are additive with respect to addition of independent variables, leading to simpler expressions for the transformation of statistics across layers. Such approaches are common in statistical physics and in mathematical statistics.

The network mapping $g : x \mapsto y$ relates the cumulant generating function of network outputs y to the statistics of the input x:

$$\mathcal{W}_{y|t; \theta}(j) = \ln \left\langle \exp \left(j^{\mathrm{T}} y \right) \right\rangle_{y|t; \theta} \quad (6)$$

$$= \ln \left\langle \exp \left(j^{\mathrm{T}} g(x; \theta) \right) \right\rangle_{x|t}. \quad (7)$$

The class-conditional output cumulant $G_{y|t; \theta}^{(n)}$ is then defined as

$$G_{y|t; \theta}^{(n)} = \left. \frac{d^n \mathcal{W}_{y|t; \theta}(j)}{d j^n} \right|_{j=0}.$$

Evaluating Eq. (7) would in principle allow one to relate $G_{y|t; \theta}^{(n)}$ to the input cumulants $G_{x|t; \theta}^{(n')}$. However, one intricacy is that the network mapping $g(x; \theta)$ is given via the iterations in Eq. (2). Their iterative nonlinear nature makes deep neural networks powerful as universal function approximators, but complicates their analysis in terms of data processing. Yet, we can study the transformation of cumulants from input to output by considering layers individually.

Since pre-activations z^l are determined by affine linear transformations, the cumulant generating function of pre-activations z^l in layer l is trivially related to the cumulant generating function of post-activations y^{l-1} of layer $l - 1$ as

$$\begin{aligned}\mathcal{W}_{z^l}(j) &= \ln \left\langle \exp \left(j^{\mathrm{T}} z^l \right) \right\rangle_{z^l} \\ &= \ln \left\langle \exp \left(j^{\mathrm{T}} W^l y^{l-1} + j^{\mathrm{T}} b^l \right) \right\rangle_x \\ &= \mathcal{W}_{y^{l-1}} \left(\left(W^l \right)^{\mathrm{T}} j \right) + j^{\mathrm{T}} b^l, \quad (8)\end{aligned}$$

yielding for the first order cumulant ($n = 1$)

$$G_{z^l}^{(1)} = W^l G_{y^{l-1}}^{(1)} + b^l, \quad (9)$$

and for second- and higher-order cumulants ($n \geq 2$)

$$G_{z^l, (r_1, \dots, r_n)}^{(n)} = \sum_{s_1, \dots, s_n} W_{r_1 s_1}^l \dots W_{r_n s_n}^l G_{y^{l-1}, (s_1, \dots, s_n)}^{(n)}. \quad (10)$$

Each index s_i is hence contracted with one factor $W_{r_k s_i}^l$ to produce the index r_k of the resulting cumulant. Consequently, cumulants of pre-activations z^l are linear tensor transformations of cumulants of post-activations y^{l-1} of the same order.

The non-linear activation function ϕ in each layer l then relates the pre-activations z^l to the corresponding post-activations y^l:

$$\begin{aligned}\mathcal{W}_{y^l}(j) &= \ln \left\langle \exp \left(j^{\mathrm{T}} y^l \right) \right\rangle_{y^l} \\ &= \ln \left\langle \exp \left(j^{\mathrm{T}} \phi(z^l) \right) \right\rangle_{z^l}. \quad (11)\end{aligned}$$

Meaning	Algebraic term	Graphical representation
External line	$j_r \, \delta_{rs}$	$j_r \quad z_s^l$
Cumulant vertex with n internal lines	$G^{(n)}_{z^l, (r_1, \ldots, r_n)}$	$\begin{array}{c} z_{r_1}^l \\ \\ z_{r_n}^l \quad \ldots \end{array}$
ϕ-vertex with m internal lines and one external line	$j_r \dfrac{\phi^{(m)}\big\vert_{z^l=0}}{m!} \\ \times \delta_{ri_1} \ldots \delta_{ri_m}$	$j_r \begin{array}{c} z_{i_1}^l \\ \ldots \\ z_{i_m}^l \end{array}$

TABLE I. Diagrammatic elements for the perturbative expansion of $\mathcal{W}_{y^l}(j)$ for $y^l = \phi(z^l)$.

This cumulant generating function of the post-activations y^l cannot, in general, be computed exactly. One common approximation technique is a perturbative expansion [20], which we here recast in the following way: by replacing $\phi(z^l)$ with its Taylor expansion $\sum_m \frac{\phi^{(m)}}{m!}(z^l)^m$ in Eq. (11) and treating nonlinear terms ($m > 1$) as perturbations, we can construct cumulants $G^{(n)}_{y^l}$ as series of Feynman diagrams composed of the graphical elements shown in Table I. For example, for the mean of the first layer we get

$$G^{(1)}_{y^1, i} = \frac{\phi^{(1)}\big\vert_{x=0}}{1!} G^{(1)}_{x, i} + \frac{\phi^{(2)}\big\vert_{x=0}}{2!} G^{(2)}_{x, ii} + \ldots$$

$$= \text{—}\bigoplus\bigcirc \; + \; \text{—}\bigoplus\bigcirc \; + \ldots \tag{12}$$

We find that in general these expressions involve two types of terms, which we represent with two types of vertices: empty circles with internal lines, representing cumulants $G^{(n)}_{z^l}$ of pre-activations z^l, and hatched circles with one external line j that stem from Taylor coefficients $\frac{\phi^{(m)}\big\vert_{z^l=0}}{m!}$ of the non-linearity.

For constructing a cumulant $G^{(n)}_{y^l}$ of the post-activations y^l of order n, we need to determine all diagrams with n external lines. External lines occur either independently or originate from a hatched vertex. Furthermore, they always need to be connected to a cumulant vertex, but cannot be connected to one another. Finally, due to the *linked cluster theorem*, only connected diagrams need to be considered, since others do not contribute to cumulants [20].

III. DECOMPOSING DEEP NEURAL NETWORKS INTO CORRELATION FUNCTIONS

Analyzing how deep networks process data is difficult due to their iterative, parameter-dependent definition.

Statistical learning theory studies the expected error [23], thus shifting from the transformation of data samples to that of data distributions. We follow this idea here by studying how data correlations of the input are iteratively transformed by deep networks, as illustrated in Fig. 1, and how they shape the expected loss.

A. Data correlations drive network training

We here discuss the dependence of the expected loss in Eq. (4) on the data correlations. In general, the expected risk is a function of the class labels t and the class-conditional cumulants $G^{(n)}_{y|t;\theta}$ of arbitrary orders n:

$$R(\theta) = \sum_t \int \mathrm{d}y \, \ell(y, t) \, p(y|t; \theta)$$

$$= \sigma(\{\{G^{(n)}_{y|t;\theta}\}_n; t\}_t).$$

However, for the often employed mean-squared error $\ell_{\mathrm{MSE}}(y, t) = \|y - t\|^2$, $R(\theta)$ only depends on the mean μ_y^t and variance Σ_y^t of outputs of each class t as

$$R_{\mathrm{MSE}}(\{\mu_y^t, \Sigma_y^t; t\}_t) = \sum_t p(t) \left(\mathrm{tr}\, \Sigma_y^t + \|\mu_y^t - t\|^2 \right). \tag{13}$$

Training therefore aims to match class means and labels, while minimizing the variance of each class' output. [2]

In this case, the first- and second-order cumulants (mean and covariance) of the last layer alone drive network training, thus singling these out as the relevant statistics. This result has two implications: 1.) In deep feed-forward networks, only non-Gaussian statistics that appear in network layers before the final layer can contribute to the learned information processing by influencing the first two cumulants in the final layer. 2.) If networks produce non-Gaussian statistics in the final layer, these do not serve a functional role per se; rather they may arise as a by-product of earlier layers operating on higher-order statistics.

Thus, understanding the network mapping reduces to understanding how the Gaussian statistics (μ_y^t, Σ_y^t) of the output arise from the presented data distribution across multiple network layers. Network training and the resulting information processing within the network is therefore directly linked to how data correlations are transformed by the network.

B. Propagation of data correlations within the network

To understand how the extraction of information from the input and its internal processing shape the first- and

---

second-order cumulants of the output, we follow these two quantities backwards through the network. According to Eq. (8)-Eq. (10), the affine transformation in each layer implies for the pre-activations z^l:

$$\mu_{z^l} = W^l \mu_{y^{l-1}} + b^l, \quad \Sigma_{z^l} = W^l \Sigma_{y^{l-1}} (W^l)^\mathsf{T}, \qquad (14)$$

showing that the two quantities are transformed independently of each other in this step (Fig. 1).

In general, the non-linear activation function $\phi : z^l \mapsto y^l$ makes the statistics of the post-activations y^l dependent on cumulants of arbitrary orders in the pre-activations z^l through (cf. Section II C)

$$\mu_{y^l} = \left. \frac{d\mathcal{W}_{y^l}(j)}{dj} \right|_{j=0} = \langle \phi(z^l) \rangle_{z^l}, \qquad (15)$$

$$\Sigma_{y^l} = \left. \frac{d^2\mathcal{W}_{y^l}(j)}{dj \, dj^\mathsf{T}} \right|_{j=0} = \langle \phi(z^l) \, \phi(z^l)^\mathsf{T} \rangle_{z^l} - \mu_{y^l} \mu_{y^l}^\mathsf{T} \qquad (16)$$

However, due to the central limit theorem, initializing the weights independently causes the affine transformation $y^{l-1} \mapsto z^l$ to mainly pass on the Gaussian part of the statistics, since higher-order cumulants $G^{(n)}_{z^l,(i_1,\ldots,i_n)} = \langle\!\langle z^l_{i_1} z^l_{i_2} \ldots z^l_{i_n} \rangle\!\rangle \sim \mathcal{O}((N_{l-1})^{-\frac{n}{2}+1})$ are suppressed by the layer width N_{l-1} for $n > 2$ (see Appendix A). Therefore, in the limit of infinitely wide networks, expectations over pre-activations $\langle \cdot \rangle_{z^l}$ can be taken with respect to Gaussian distributions $z^l \sim \mathcal{N}(\mu_{z^l}, \Sigma_{z^l})$, and we obtain that the mean and covariance of post-activations are non-linear functions of only mean and covariance of pre-activations

$$\mu_{y^l} = f_\mu(\mu_{z^l}, \Sigma_{z^l}), \quad \Sigma_{y^l} = f_\Sigma(\mu_{z^l}, \Sigma_{z^l}). \qquad (17)$$

These functions mediate interactions between first- and second-order cumulants.

Applying this argument iteratively to the network layers $l = L + 1, L, \ldots, 2$, it follows that the information processing in the *internal* network layers is largely determined by an iterated, non-linear mapping of mean and covariance. The interaction functions f_μ and f_Σ can be calculated analytically for various activation functions ϕ; we provide expressions for $\phi = \mathrm{ReLU}$ and $\phi(z) = z + \alpha \, z^2$ in Appendix B Table II. The latter, minimally nonlinear activation function yields especially interpretable interaction functions that are constructed from the following diagrams:

$$\mu_{y^l,i} = \mu_{z^l,i} \qquad + \alpha \, (\mu_{z^l,i})^2$$
$$+ \alpha \, \Sigma_{z^l,ii}$$

$$\Sigma_{y^l,ij} = \Sigma_{z^l,ij} \qquad + 4\alpha^2 \mu_{z^l,i} \Sigma_{z^l,ij} \mu_{z^l,j}$$
$$+ 2\alpha^2 (\Sigma_{z^l,ij})^2 + 2\alpha \Sigma_{z^l,ij} (\mu_{z^l,i} + \mu_{z^l,j})$$

Although training introduces correlations between weights, thus violating the independence assumption of the central limit theorem, we will show that in practice the first- and second-order cumulants provide a useful approximation for the information propagation *within* the network.

C. Information extraction in the input layer

So far we have studied the internal network layers. Here, we discuss the role of the input layer in extracting information from higher-order correlations of the input data. Since the weights in this layer do not scale with the network width N but the input dimension N_0, higher-order input cumulants $G^{(n>2)}_x$ contribute to the mean and covariance of the post-activations y^1, $\mu_{y^1} = h_\mu(\{G^{(n)}_x\}_n)$ and $\Sigma_{y^1} = h_\Sigma(\{G^{(n)}_x\}_n)$. These mean and covariance are then passed on through the entire network.

The interaction functions h_μ and h_Σ can be systematically approximated for any activation function by the diagrammatic techniques discussed in Section II C or alternatively by a Gram-Charlier expansion (see Appendix B). Analytically simple and exact expressions can be computed for a quadratic non-linearity (see Appendix B); in this case, the expression for the mean does not get any contribution from $G^{(n>2)}_x$, while the covariance gets additional contributions from third- and fourth-order input

correlations:

$$
\begin{aligned}
\Sigma_{y^l,ij}\big|_{\text{add.}} = \ & \alpha\left(G^{(3)}_{z^l,(i,j,j)} + G^{(3)}_{z^l,(j,i,i)}\right) \\
& + 2\alpha^2\left(G^{(1)}_{z^l,(i)}\,G^{(3)}_{z^l,(i,j,j)} + G^{(1)}_{z^l,(j)}\,G^{(3)}_{z^l,(j,i,i)}\right) \\
& + \alpha^2\,G^{(4)}_{z^l,(i,i,j,j)}
\end{aligned}
$$

When evaluating the presented diagrams, all permutations of indices $(r_1,\ldots,r_n)$ for both internal and external legs need to be taken into account, leading to the above expressions of similar structure but interchanged indices. Any repeating expressions are absorbed into the corresponding prefactor.

By iterating (14) and (17) across layers, one obtains the mean and covariance of the network output $y = g(x;\theta)$ as functions of the statistics of x:

$$
\mu_y = g_\mu(\{G^{(n)}_x\}_n;\theta,\phi)\,, \qquad \Sigma_y = g_\Sigma(\{G^{(n)}_x\}_n;\theta,\phi)\,.
\tag{18}
$$

These two Gaussian statistics are sufficient to evaluate the MSE loss Eq. (13). Note that these are generally not exact due to the approximation of pre-activations z^l at each layer as Gaussians. Since the network decomposes into a mapping for each class label t, one obtains the distribution of the network output as a Gaussian mixture $p(y) = \sum_t p(t)\,\mathcal{N}(\mu_y^t,\Sigma_y^t)(y)$, whose parameters (μ_y^t,Σ_y^t) are determined by the propagation of data correlations through the network (18). In the following, we call the mapping

$$
g_{\text{stat}} : (\{G^{(n)}_x\}_n,\theta,\phi) \mapsto p(y)
\tag{19}
$$

the *statistical model* of the network. For further details, see Appendix C.

IV. EXPERIMENTAL RESULTS

We now apply the developed methods to the XOR problem and the MNIST dataset. We use the network architecture defined in Section II A with fixed network width $N_l = N$ for $l \geq 1$ and either the ReLU activation function or a minimal non-linearity, namely the quadratic activation function $\phi(z) = z + \alpha z^2$ with $\alpha = 0.5$. For initialization of network parameters θ, we use $\sigma_w^2 = \sigma_b^2 = 0.75$. Following standard procedure, networks are trained by optimizing the empirical risk per data batch $\{(x^{(b)}, t^{(b)})\}_b$ of the expected MSE loss:

$$
R_{\text{emp, MSE}}(\theta) = \frac{1}{B}\sum_{b=1}^{B}\ell_{\text{MSE}}(g(x^{(b)};\theta), t^{(b)})\,.
\tag{20}
$$

The batch size B is set to 10 on XOR and 100 on MNIST. For optimization, we use ADAM [25, 26] with learning rate 10^{-3}, momenta $\beta_1 = 0.9$ and $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and $\lambda = 0$. Network implementations were done in PYTORCH [27].

A. Multiple information encodings of the XOR problem

We first study an adaptation of the XOR problem as a non-linearly separable Gaussian mixture distribution. We make use of two conceptual advantages of this XOR task: First, knowing the exact input distribution allows us to focus on the internal information processing within the network. Second, the fact that each class is itself a mixture distribution allows us to trace the class-conditional correlations in two alternative forms, corresponding to two different statistical representations of class membership, which isolate different statistics of the input – respectively the mean and the covariance. We find that while the task can be solved for both representations, they yield different local minima of the empirical loss landscape.

1. Problem setup as a Gaussian mixture

Our adaptation of the XOR problem uses real-valued instead of binary inputs and describes the input distribution as a Gaussian mixture of four components, illustrated in Fig. 2**a**. For the class label $t = +1$, we choose the mean values of its two components $\pm$ as $\mu_x^{t=+1,\pm} = \pm(0.5, 0.5)^\top$; for $t = -1$, we use $\mu_x^{t=-1,\pm} = \pm(-0.5, 0.5)^\top$. Covariances are isotropic throughout with $\Sigma_x^{t,\pm} = 0.05\,\mathbb{I}$ and the input distribution

$$
p(x,t) = p(t)\sum_{\pm} p_\pm\,\mathcal{N}(\mu_x^{t,\pm},\Sigma_x^{t,\pm})(x)
$$

weighs all components equally $p(t) = p_\pm = \frac{1}{2}$. A data sample $x^{(d)}$ is assigned a target label $t^{(d)} \in \{\pm 1\}$ based on the mixture component it is drawn from. We use training and test datasets of sizes $n_{\text{train}} = 10^5$ and $n_{\text{test}} = 10^4$, respectively.

2. Accuracy of internal information processing in terms of correlation functions

Given the exact input distribution for this problem, we trace the transformation of mean and covariance predicted by Eq. (18) for each mixture component $(t, \pm)$ separately, obtaining

$$
p_{\text{theo.}}(y) = \sum_t p(t)\sum_{\pm} p_\pm\,\mathcal{N}(\mu_y^{t,\pm},\Sigma_y^{t,\pm})(y)\,,
$$

where $\mu_y^{t,\pm} = g_\mu(\mu_x^{t,\pm},\Sigma_x^{t,\pm};\theta,\phi)$ and $\Sigma_y^{t,\pm} = g_\Sigma(\mu_x^{t,\pm},\Sigma_x^{t,\pm};\theta,\phi)$ are functions of the input statistics, the network parameters, and depend on the
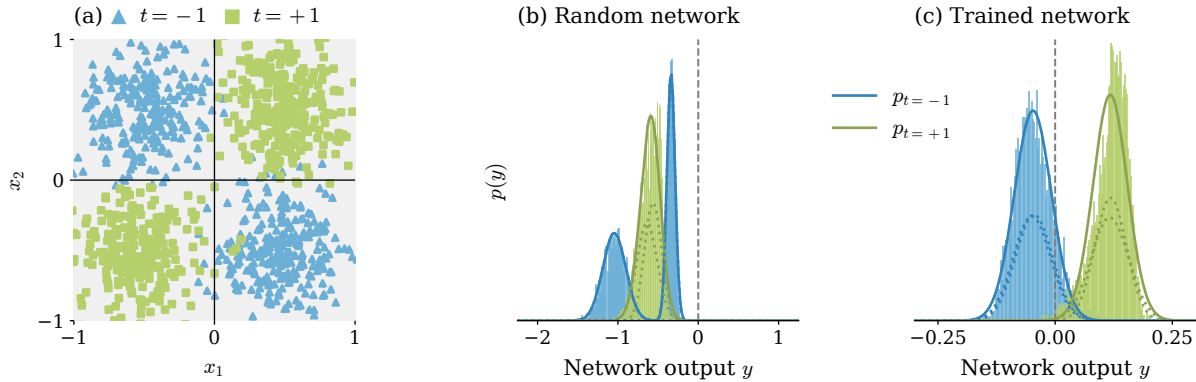
FIG. 2. **Information propagation in** ReLU **networks for the XOR problem.** **(a)** The distribution of input data is modeled as a Gaussian mixture. Data samples $x^{(d)}$ (blue and pink dots) are assigned to class labels $t = \pm 1$ based on the respective mixture component. **(b,c)** Distribution of the network output for random (b) and trained (c) parameters. Class-conditional distributions (solid curves) are determined as a superposition of the propagated mixture components (dashed curves) and empirical estimates (blue and pink histograms) are obtained from the test data. Since networks are trained on class labels $t = \pm 1$, the classification threshold is set to $y = 0$ (gray lines). Other parameters: $\phi = $ ReLU, depth $L = 1$, width $N = 10$; trained network in (c) achieves $P = 93.82\%$ performance.
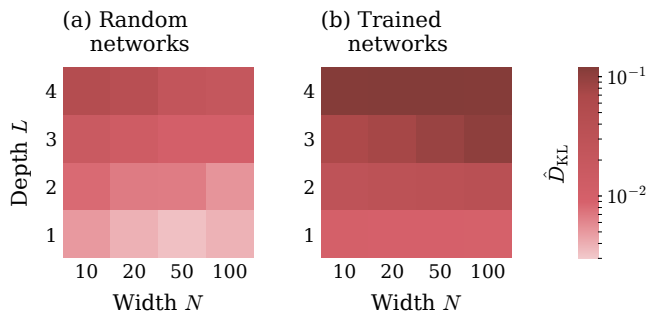
FIG. 3. **Deviation between theoretical and empirical output distribution for (a) random and (b) trained networks**, measured across 10^2 different network realizations using the normalized Kullback-Leibler divergence $\hat{D}_{\mathrm{KL}}(p_{\mathrm{emp.}}\|p_{\mathrm{theo.}})$. On average, the trained networks achieve performance values of $P = 97.21\% \pm 0.36\%$. Networks were trained to perform the XOR task described in Section IV A 1. Other parameters: $\phi(z) = z + \alpha z^2$.

activation function. In Fig. 2 we compare this theoretical result to an empirical estimate of the output distribution $p_{\mathrm{emp.}}(y)$, given as a histogram obtained from the test data. We test the validity of the statistical model for both, an untrained network with random weight initialization (Fig. 2**b**) and a trained network (Fig. 2**c**).

The untrained network produces an output distribution of complex shape composed of superimposed close-to-Gaussian distributions, each corresponding to one component, as shown in Fig. 2**b**. Training the network reshapes the output distribution such that the class-conditional distributions $p(y|t)$ become well separated by the threshold at $y = 0$, as shown in Fig. 2**c**. The overlap between these two distributions around the thresh-

old corresponds to the classification error. Qualitatively, theory and simulation agree well for both random and trained networks. These results apply for different activation functions ϕ (see Fig. 7 in Appendix D).

To quantify the alignment of theory and simulation, we compute the Kullback-Leibler divergence D_{KL} between the empirical estimate $p_{\mathrm{emp.}}(y)$ and the theoretical result $p_{\mathrm{theo.}}(y)$, considering the empirical distribution $p_{\mathrm{emp.}}(y)$ as the reference. To account for the variability of output distributions across different network realizations, this quantity is normalized by the entropy H of the empirical distribution $p_{\mathrm{emp.}}(y)$, yielding $\hat{D}_{\mathrm{KL}}(p_{\mathrm{emp.}}\|p_{\mathrm{theo.}}) = D_{\mathrm{KL}}(p_{\mathrm{emp.}}\|p_{\mathrm{theo.}})/H(p_{\mathrm{emp.}})$.

We average $\hat{D}_{\mathrm{KL}}(p_{\mathrm{emp.}}\|p_{\mathrm{theo.}})$ across 10^2 different network realizations, for random (Fig. 3**a**) and trained (Fig. 3**b**) networks. In both cases, the deviation between theory and simulation is generally small, but increases mildly with the network depth L as approximation errors accumulate across network layers. For random networks, we observe a decrease of the deviation for wider networks, in agreement with the central limit theorem as discussed in Section III B. This effect vanishes for trained and thus correlated parameters, resulting in an overall increase of deviations between theory and simulation. Nonetheless, this increase remains modest, showing that the theory continues to be applicable for networks with trained, and thus non-random, parameters.

3. Different information coding paradigms and their relations

In the previous section, we have shown that the mapping implemented by the network can be described as a mapping of correlation functions (see Eq. (19)). On the
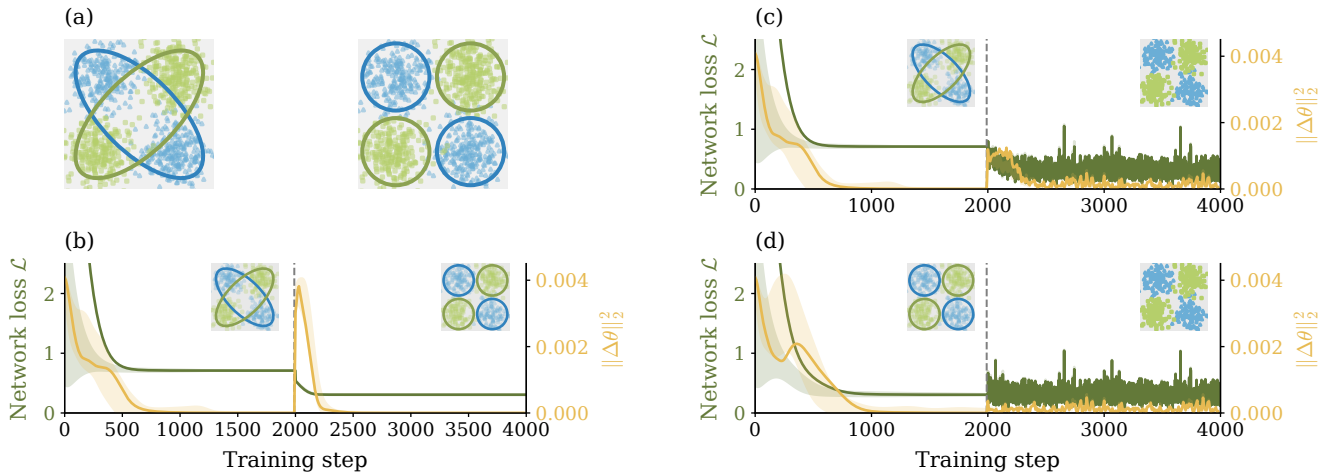
FIG. 4. **Mean and covariance coding. (a)** Data distributions of XOR task. Left: **Covariance coding**; class membership (pink and blue ellipses) is encoded in the covariance alone. Right: **Mean coding**; the two individual Gaussian components of each class (both pink circles / both blue circles) differ in their means, while having the same covariance. **(b-d)** Evolution of network loss R ($R_{\mathrm{emp,MSE}}$ Eq. (20) for training the network model, R_{MSE} Eq. (13) for the statistical model) and change of network parameters $\|\Delta\theta\|_2^2$ in each training step: The first $T_1 = 2000$ training steps train one model, starting from random parameters θ. The model representation is changed at T_1, starting from parameters $\theta(T_1)$ obtained in the preceding period. The change of network parameters is evaluated every 10 training steps $\|\Delta\theta(T)\|_2^2 = \|\theta(T) - \theta(T-10)\|_2^2$. Shaded areas show the typical range, based on mean and one standard deviation across 10^2 network realizations. Solid curves show the behavior of a single network realization. **(b)** First period: statistical model with covariance coding; second period: statistical model with mean coding. **(c)** First period: statistical model with covariance coding; second period: network model. **(d)** First period: statistical model with mean coding; second period: network model. Training parameters: $n_{\mathrm{train}} = 10^4$, 2 epochs. Other parameters: $\phi(z) = z + \alpha z^2$, depth $L = 1$, width $N = 10$.

level of data correlations, it directly follows that the network's expressivity with respect to a given task depends on two properties: (1) the ability of the network architecture to implement a desired mapping of data correlations from its input to its output; (2) the way in which information about class membership is represented by data correlations in the input.

A complete study of the first property would be provided by fully describing the space of possible mappings, which is challenging in general. However, the forward mapping of cumulants we have obtained in Section III B allows us to probe this space experimentally, and it provides a path to more systematic studies of network expressivity – see our remarks on statistical receptive fields in Section V.

In this section, we study the second property by investigating two different information representations: (A) the class membership is represented by different means, while the covariances and all higher-order cumulants are identical; (B) the class membership is represented by different covariances, while the means and all higher-order cumulants are identical. Accordingly, these two representations are called *mean coding* (A) and *covariance coding* (B) in the following. While each of these two settings confines the class membership to one particular cumulant order, the more general case is that class membership is represented by various orders of statistical moments. In that case, the network may make use of this redundant

information to maximize performance.

To be able to compare these settings, in either case we draw the input data from the same data distribution. We use the statistical model corresponding to the network described in Section III C, limiting input correlations to mean and covariance. We take the binary XOR problem (see Section IV A 1) which can be cast into either information representation in a natural way: For mean coding (A), we provide to the network both the class labels t and the specific mixture component $\pm$ from which a sample was drawn, yielding four sets of statistics $\{\mu_x^m, \Sigma_x^m\}_{m=(t,\pm)}$ with different means but identical covariances. For covariance coding (B), only the class label t is provided to the network, yielding two sets of statistics $\{\mu_x^m, \Sigma_x^m\}_{m=(t)}$, for which the covariances $\Sigma_x^{t=\pm 1} = \left(\begin{smallmatrix} 0.3 & \pm 0.25 \\ \pm 0.25 & 0.3 \end{smallmatrix}\right)$ differ between the two classes, while their means are the same (see Fig. 4a).

We compare these two statistical representations of the network to network models trained on batches of samples; this we refer to as *sample coding* in the following. Sample coding can be considered as the case where potentially all statistical moments of the data are accessible to the network. Our goal is to address the following questions: First, which statistical representation most closely matches the information representation used by a network trained on data samples? Second, is there a difference in performance between information representations; in particular, can the network equivalently use the

information provided by either mean or covariance coding? Finally, does the network make use of redundant information to improve performance?

To answer these questions, we optimize models until convergence using either representation. We then switch to a different representation, continuing optimization for the same number of steps, and observe the stability of the previously found solution. Each experimental setup is repeated with 10^2 different weight initializations. Results are shown in Fig. 4 for three different coding combinations, where we plot both the loss and the magnitude of change $\|\Delta\theta\|_2^2$ of the model parameters.

We find that, after initial convergence, all three models correspond to trained networks with at least $P = 91.94\%$ performance. Furthermore, we observe that immediately after the switch from covariance to mean coding, $\|\Delta\theta\|_2^2$ jumps to values similar to the initial training steps (Fig. 4b). This indicates a near complete change of the model, which suggests that mean and covariance coding induce fundamentally different solutions. In contrast, the jump is modest when switching from covariance to sample coding (Fig. 4c), and non-existent when switching from mean to sample coding (Fig. 4d) – suggesting that those different solutions coexist in the true loss landscape of the network model. Thus, we find that the network utilizes the presented information in different ways for the two representations, as expected based on the information flow in these networks, derived in Section III B.

In particular, the case of covariance coding highlights the importance of a non-linear activation function when the discriminating information is not contained in the class means. Since classification is based on different mean values in the network output, the difference in covariance for each class needs to be transferred to the mean. This information transfer is mediated by the non-linearity ϕ; for the case $\phi(z) = z + \alpha z^2$ used in Fig. 4, we have the particularly simple transfer function

$$\mu_{y^l,i} = \mu_{z^l,i} + \alpha\left(\mu_{z^l,i}\right)^2 + \alpha\,\Sigma_{z^l,ii} \qquad (21)$$

from covariances to means. In this way, we can track how information flows into the mean as it is transformed by successive network layers.

In summary, we find that for this task the network can effectively utilize the information presented by either mean or covariance coding, both representations leading to different solutions with comparable performance. Sample coding tends to yield similar solutions as mean coding, implying that the network makes negligible use of redundant information present in higher-order moments of the data samples.

B. Essential data correlations of the MNIST data set

We consider in this section the MNIST data set [21], consisting of 10 classes of 28×28 images. This dataset is highly structured: if one approximates each class by a
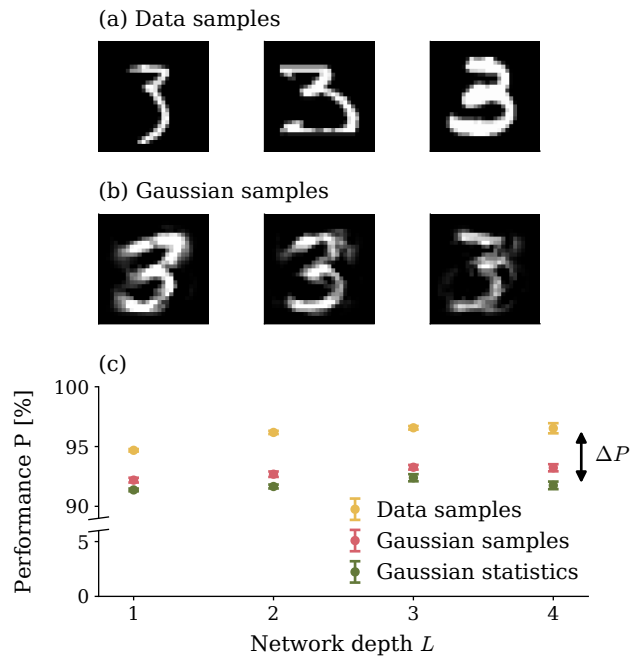

FIG. 5. **First- and second-order correlations of MNIST. (a)** Three example data samples showing the digit 3 from the MNIST training dataset. **(b)** Data samples showing the digit 3, drawn from the Gaussian approximation of the input distribution. **(c)** Classification performance on the MNIST test dataset for different input encodings. Network training consistently achieves performance values of $P \approx 95\%$ or more (yellow), while performance of the corresponding statistical model is $4.2\% \pm 0.3\%$ lower (green). Training networks on Gaussian input samples yields a comparable performance difference (red). In all cases, performance is evaluated on the MNIST test dataset. Error bars show the standard deviation across 10 different network realizations, excluding 3 performance values for the statistical model due to the optimization diverging. Other parameters: $\phi(z) = z + \alpha z^2$, depth $L \in [1, 2, 3, 4]$, width $N = 100$.

multivariate Gaussian, the resulting samples are already visually recognizable (Fig. 5a,b; see Appendix E for further details). Our goal is to use the theory developed in previous sections to quantify this observation, in a matter which can be generalized to different data sets and different sets of input cumulants. We also argue that truncation of cumulants in the input layer has the largest impact, and in the process validate our theory on a non-trivial task.

Concretely, we proceed as follows: when optimizing the parameters θ^* of the statistical model, we restrict the data statistics to a particular set of cumulants $\{G_x^{(n)}\}_{n=1,\dots,\hat{n}}$ and compare the achieved performance to that of the network trained on samples $y = g(x; \theta^*)$. The difference in performance is then indicative of the importance of the cumulants we kept. We employ one-hot encoding, making the network output $d_{\mathrm{out}} = 10$ dimensional.

As a baseline, we first train network models on both the

MNIST data set (Fig. 5**a**) and the corresponding Gaussian samples (Fig. 5**b**). The latter case limits the information that can be extracted by the input layer to the class-conditional means μ_x^t and covariances Σ_x^t. In both cases, networks are trained with the standard empirical loss (Eq. (20)); in particular, this allows inner network layers to make use of cumulants of any order. With respect to classification performance, we find that training on Gaussian samples yields a performance that is lower by $\Delta P \simeq 3.1\% \pm 0.3\%$ (Fig. 5**c**): a difference we can ascribe to the removal of higher-order cumulants in the data distribution. Based on the modest magnitude of this difference, we conclude that data mean and covariance are already highly informative for these data and account for about $P \approx 91\%$.

We next train the corresponding statistical model (Eq. (18)) on the Gaussian approximation of MNIST (Fig. 5**b**). Compared to the network model trained on the same data distribution, we find only slightly lower performance – by about $1.0 \pm 0.3\%$ (Fig. 5**c**) – suggesting that the statistical model given by Eq. (18) is a good representation for the information processing in internal network layers. The fact that most of the performance drop with respect to standard training on MNIST is due to the Gaussian approximation of the input data indicates the importance of processing higher-order cumulants by the input layer. In the next section, we show with an illustrative example how these can be included into the theory.

C. Including higher-order correlation functions in the input layer

So far we have studied class-conditional means μ_x^t and covariances Σ_x^t of the input data; however, these two statistics may not always be informative. It is in fact easy to construct a low-dimensional task with two classes $t = 0, 1$, where both class-conditional means and covariances of the data are identical – $\mu_x^{t=0} = \mu_x^{t=1}$, $\Sigma_x^{t=0} = \Sigma_x^{t=1}$ – thereby conveying no information regarding the class membership (Fig. 6**a**,**b**). Classification in such cases must therefore rely on higher-order statistics. For the example in Fig. 6, since third-order cumulants differ between classes ($G_x^{(3),t=0} = -G_x^{(3),t=1}$), we expect their inclusion into the statistical model to be sufficient for solving the task. We here demonstrate that such higher-order cumulants can indeed be treated by our approach – in particular, we validate the statement made in Section III C that it suffices to consider higher-order cumulants in only the first layer.

The input distribution for this task is defined as a Gaussian mixture of four components, illustrated in Fig. 6**a**,**b** (details in Appendix F). As expected, training the network model yields near-optimal performance values, while a statistical model $g_{\text{stat}}(\{G_x^{(n)}\}_{n=1,2}, \theta)$ that considers only the class-conditional means μ_x^t and covariances Σ_x^t fails to solve the task, yielding chance-level
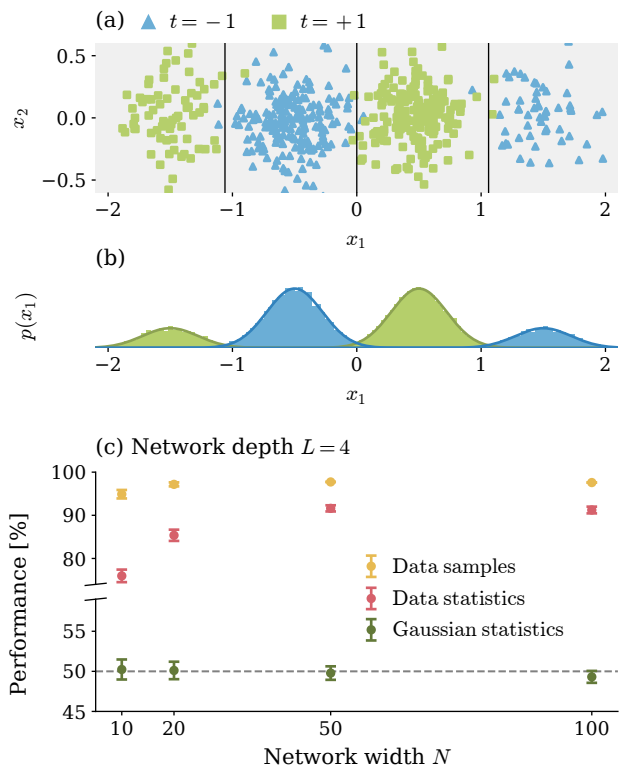

FIG. 6. **Information extracted from higher-order correlations.** **(a)** The distribution of input data is modeled as a Gaussian mixture. Data samples $x^{(d)}$ (blue and pink dots) are assigned to class labels $t = 0, 1$ based on their respective mixture component. The two classes have zero mean and the same covariance. **(b)** Projection of data samples to the x_1-axis (histograms), which corresponds to the marginalization of the input distribution with respect to x_2 (solid lines), illustrating the different weighing of the mixture components. **(c)** Classification performance for different model choices. Network training consistently achieves performance values of $P \approx 94\%$ or more (yellow). Optimizing the statistical model $g_{\text{stat}}(\{G_x^{(n)}\}_{n=1,2}, \theta)$ that considers only the first- and second-order correlations (green) results in performance values corresponding to chance level (dotted line). However, including the third-order correlations into the statistical model $\tilde{g}_{\text{stat}}(\{G_x^{(n)}\}_{n=1,2,3}, \theta)$ nearly bridges this gap (red). In all cases, performance is evaluated on a test dataset. Error bars show the standard deviation of performance across 10^2 different network realizations. Other parameters: $\phi(z) = z + \alpha z^2$, depth $L = 4$, width $N = [10, 20, 50, 100]$.

performance (Fig. 6**c**). This performance gap is nearly bridged when we include the third-order input cumulants $G_x^{(3)}$ (via Eq. (B85)) in the first layer of the statistical model $\tilde{g}_{\text{stat}}(\{G_x^{(n)}\}_{n=1,2,3}, \theta)$. The activation function ϕ allows information in $G_x^{(3)}$ to be transferred to lower-order cumulants, which are then processed by subsequent layers in the manner described in previous sections – facilitating different means $G_y^{(1)}$ in the output of the statistical model.

V. DISCUSSION

The question of how neural networks process data is fundamentally the question of how information is encoded in the data distribution and subsequently transformed by the network. We here present an analytical approach, based on methods from statistical physics, to study the mapping of data distributions implemented by deep feed-forward neural networks: we parameterize the data distribution in terms of correlation functions and derive their successive transformations across layers. We show that the initial network layer effectuates the extraction of information from higher-order correlations in the data; for subsequent layers, a restriction to first- and second-order correlation functions (mean and covariance) already captures the main properties of the network computation. This reduction of the bulk of the network to a non-linear mapping of a few correlation functions provides an attractive view for further analyses. It relies on the assumption of sufficiently wide layers to apply the central limit theorem, but, in practice, we find that the approximations are useful even for narrow networks.

We validate these results for different datasets. We first investigate an adaptation of the XOR problem that is purely based on first- and second-order cumulants. Despite the non-linear transformations in each layer giving rise to higher-order correlations, the network solutions to this task can largely be described in terms of transformations solely between mean and covariance of each class. We then consider the MNIST database: we show that network solutions based on empirical estimates for mean and covariance of each class capture a large amount of the variability within the dataset, but still exhibit a non-negligible performance gap in comparison to solutions based on the actual dataset. We discuss how this performance difference results from the omission of higher-order correlations. Lastly, we introduce an example task where higher-order correlations exclusively encode class membership, which allows us to explore their role in isolation.

Limitations The dimensionality N_0 of the data may limit the applicability of the presented approach to low orders n, since cumulants of order n are tensors with N_0^n entries. We note however that there exist methods to ease the computational cost of higher-order cumulants in large dimensions: for example, one can make use of the inherent symmetries in these tensors, as well as in the theory itself. The application of such methods to our framework remains a point for future work. A parametrization of a probability distribution in terms of cumulants, moreover, needs to be chosen such that it maintains positivity of the probability density function. Conserving this property implies constraints for truncating cumulant orders, which require further investigations.

The presented framework and its perturbative methods naturally apply to polynomial approximations of activation functions. Although networks with polynomial nonlinearity are, in principle, not capable of universal function approximation [28–30], this is not an issue for the classification tasks we consider. To obtain illustrative analytical expressions for the mixing of correlation functions, we chose to demonstrate the approach with a quadratic activation function. Non-polynomial and even non-differentiable activation functions can, however, also be dealt with in our framework using Gram-Charlier expansions that are detailed for the example of the ReLU activation in the Appendix B. While we here mostly focus on the mean and covariance, we also show how to generalize the results to higher-order cumulants.

Relation to kernel limit of deep networks In this paper we study *individual* networks with specific parameters θ. There is a complementary approach that studies *ensembles* of (infinitely) large networks with random parameters: Poole *et al.* [14] expose a relation between the Lyapunov exponents and the depth to which information propagates in randomly initialized deep networks. They find the regime close to chaos beneficial for information propagation. Furthermore, random network parameters are central to studying training as Bayesian inference [31]: independent Gaussian priors on the network parameters render Bayesian inference exact on the resulting Gaussian process [7, 12, 32, 33]. The works [9, 34–36] use methods similar to ours to compute finite-width corrections and corrections arising from training with stochastic gradient descent. While these approaches consider distributions over network parameters θ, we here study distributions of the input data x and how these are mapped to distributions of outputs y for given network parameters θ. The focus on individual networks rather than ensembles thereby allows us to take into account the internal statistical structure of data samples, for example covariances $\Sigma_{x,ij}$ between individual pixels i and j in images. Bayesian inference on Gaussian processes, in contrast, is only sensitive to the overall similarity between different images.

Related works Describing data and network activity in terms of correlations was initially explored by Deco and Brauer [37] on the particular architecture of volume-preserving networks. They derived expressions of the output in terms of its correlations as well as training rules that aim to decorrelate given input data. The work we present here differs in that our goal is not to impose a specific statistical structure on the network output, but to relate the correlations of the input and output distributions and thereby obtain a description of the information processing within the network.

While we do show that these distributions are not exactly Gaussian, that the networks can utilize higher-order correlations in the hidden layers, and how these contributions could in principle be computed, we focus mostly on self-consistently tracking the distributions in Gaussian approximation. This is because, as we show, the Gaussian approximation is tractable while staying accurate also for trained networks and capturing the majority of the test accuracy in our examples. That a Gaussian approximation is surprisingly effective has also been argued

in a recent line of works using teacher-student models with realistic data structure [38–40]. Other works which are based on a Gaussian approximation of the representation in each layer are: [41] using general deep networks, [42] focusing on Res-nets, and [43] considering the case of GANs. Finally, a pedagogical text focusing on field-theory for deep neural networks has recently been published [44].

Outlook Tracing transformations of data correlations through layers of a neural network allows the investigation of mechanisms for both information encoding and processing; in this manner, it presents a handle towards interpretability of deep networks. The availability of tractable expressions describing the transformations of data correlations within neural networks is therefore an interesting prospect for future work seeking to dissect how networks learn and perform tasks. In this context, the theory we propose assumes data statistics of the input distribution $p(x)$ to be known and exposes how statistical features of the data are transformed to generate the output, with the goal of shedding light onto the networks' functioning principles.

Another natural application of the proposed framework is the identification of essential correlations in the data. In that scenario, we do not need the exact distribution $p(x)$, but only sufficiently accurate estimates of some statistics of x that can be obtained from the training data. By manipulating the information available to the model during training, we expose different information encodings the network can employ to solve the same task. In this work, we mostly focus on first- and second-order cumulants as a proof-of-concept. In particular for more complex data sets such as CIFAR-10 [45], we expect higher-order cumulants to be of larger importance and essential to network performance. Further steps towards applying the presented theory to state-of-the-art networks, such as ResNet-50 [46], will require the extension to convolutional network layers. However, since these are equivalent to linear layers with weight matrices of a particular shape [13], they can straightforwardly be included in the framework.

Another future direction targets expressivity of deep networks: by reversely tracing the data correlations through the network, from target to data, one may ask which input distributions are mapped to a given output distribution – in effect constructing layer-resolved, statistical receptive fields for each target. Expressing these receptive fields in terms of data correlations may also be useful for studying how the complexity of data distributions is reduced by deep neural networks.

ACKNOWLEDGMENTS

APPENDIX A DISORDER AVERAGE FOR WIDE NETWORKS

The application of a non-linear activation function ϕ in each network layer generates higher-order cumulants from Gaussian distributed layer input, as discussed in Section III B. For wide networks we can, however, perform a disorder average with regard to the network parameters θ, by which follows that higher-order cumulants $G_{z^l}^{(n)}$ of the pre-activations z^l beyond mean and covariance on expectation scale down with the layer width N_{l-1}. Consequently, these become negligible for wide networks where $N_{l-1} \gg 1$.

We assume the network parameters $\theta = \{W^l, b^l\}_{l=1,\ldots,L+1}$ to be Gaussian distributed with $W_{rs}^l \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(0, \sigma_w^2/N_{l-1}\right)$ and $b_r^l \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(0, \sigma_b^2\right)$. This choice of initialization ensures that the magnitude of the covariance is preserved within the network as it does not scale with the layer width

$$\langle \Sigma_{z^l, rs}\rangle_\theta = \left\langle \left(W^l \Sigma_{y^l} (W^l)^{\text{T}}\right)_{rs}\right\rangle_\theta \tag{A22}$$

$$= \left\langle \sum_{t,u=1}^{N_{l-1}} W_{rt}^l W_{su}^l \Sigma_{y^l, tu}\right\rangle_\theta \tag{A23}$$

$$= \delta_{rs} \frac{\sigma_w^2}{N_{l-1}} \sum_{t=1}^{N_{l-1}} \Sigma_{y^l, tt} \tag{A24}$$

$$= \delta_{rs} \mathcal{O}_{N_{l-1}}(1). \tag{A25}$$

According to the central limit theorem, a properly normalized sum of independent random variables tends towards a Gaussian distribution [47]. We apply this idea to the affine linear transformation in each layer, which is given by $z_r^l = \sum_{s=1}^{N_{l-1}} W_{rs}^l y_s^l + b_r^l$.

We investigate the dependence of the corresponding cumulants $G_{z^l}^{(n)}$ on the layer width N_{l-1} by taking the average with regard to the network parameters θ. As these are independently distributed, we can perform the average layerwise. For Gaussian input data, we have in the first network layer

$$\langle \mu_{z^1, r}\rangle_\theta = \sum_{t=1}^{d_{\text{in}}} \langle W_{rt}^0\rangle_{W^0} \mu_{x,t} + \langle b_r^0\rangle_{b^0} = 0, \tag{A26}$$

$$\langle \Sigma_{z^1, rs}\rangle_\theta = \delta_{rs} \mathcal{O}_{d_{\text{in}}}(1). \tag{A27}$$

In consequence, only diagrammatic contributions that contain cumulant vertices with the same indices ($s_1 = s, \ldots, s_n = s$) on all internal lines remain non-zero after applying the disorder average $\langle . \rangle_\theta$. Since the non-linear activation is applied elementwise, it generates only ϕ-vertices with the same indices ($s_1 = s, \ldots, s_n = s$) on all

internal lines. Taken together, all diagrammatic contributions but the ones where all internal lines, and thus also all external lines, come with the same indices, yield zero in the disorder average:

$$\left\langle G^{(n)}_{y^l,(r_1,\dots,r_n)}\right\rangle_\theta \propto \delta_{r_1\dots r_n}. \tag{A28}$$

After applying the affine linear transformation in the next layer, cumulants of odd order vanish since

$$\left\langle G^{(n)}_{z^l,(r_1,\dots,r_n)}\right\rangle_\theta \tag{A29}$$

$$=\left\langle \sum_{s_1,\dots,s_n=1}^{N_{l-1}} W^l_{r_1 s_1}\dots W^l_{r_n s_n}\, G^{(n)}_{y^l,(s_1,\dots,s_n)} + \delta_{1n} b^l_{r_1}\right\rangle_\theta \tag{A30}$$

$$\propto \langle W^l_{r_1 s_1}\rangle_{W^l} + \delta_{1n}\langle b^l_{r_1}\rangle_{b^l} = 0. \tag{A31}$$

For even orders, it follows with Eq. (A28) that

$$\left\langle G^{(n)}_{z^l,(r_1,\dots,r_n)}\right\rangle_\theta \tag{A32}$$

$$=\left\langle \sum_{s_1,\dots,s_n=1}^{N_{l-1}} W^l_{r_1 s_1}\dots W^l_{r_n s_n}\, G^{(n)}_{y^l,(s_1,\dots,s_n)}\right\rangle_\theta \tag{A33}$$

$$=\sum_{s=1}^{N_{l-1}} \left\langle W^l_{r_1 s}\dots W^l_{r_n s}\right\rangle_{W^l} \left\langle G^{(n)}_{y^l,(s,\dots,s)}\right\rangle_{\theta'} \tag{A34}$$

$$=\mathcal{O}\left(\left(\frac{\sigma_w^2}{N_{l-1}}\right)^{\frac{n}{2}} N_{l-1}\right) = \mathcal{O}\left((N_{l-1})^{-\frac{n}{2}+1}\right). \tag{A35}$$

Thus, cumulants beyond second order become negligible for $N_{l-1} \gg 1$ and the distribution at internal network layers can be approximated as a Gaussian.

APPENDIX B   INTERACTION FUNCTIONS FOR DIFFERENT ACTIVATION FUNCTIONS

In Section III B, we derived the interaction functions resulting from the non-linearity ϕ,

$$f_\mu(\mu_{z^l}, \Sigma_{z^l}) = \langle\phi(z^l)\rangle_{z^l},$$
$$f_\Sigma(\mu_{z^l}, \Sigma_{z^l}) = \langle\phi(z^l)\phi(z^l)^\mathsf{T}\rangle_{z^l} - \mu_{y^l}\mu_{y^l}^\mathsf{T}.$$

Table II gives these expressions for the ReLU and quadratic nonlinearities.

Derivations for ReLU activations

We here consider networks with the ReLU activation function $\phi(z) = \max(0, z)$. Taking the distribution of pre-activations z^l to be Gaussian distributed with mean μ_{z^l} and covariance Σ_{z^l}, the mean post-activations are

given by

$$\mu_{y^l,i} = \langle\max(0, z_i^l)\rangle_{z^l\sim\mathcal{N}(\mu_{z^l},\Sigma_{z^l})} \tag{B36}$$

$$=\frac{1}{\sqrt{2\pi\Sigma_{z^l,ii}}}\int_0^\infty \mathrm{d}z_i^l\, z_i^l\, \exp\left(-\frac{(z_i^l-\mu_{z^l,i})^2}{2\Sigma_{z^l,ii}}\right) \tag{B37}$$

$$=-\frac{\sqrt{\Sigma_{z^l,ii}}}{\sqrt{2\pi}}\int_{-\mu_{z^l,i}}^\infty \mathrm{d}z_i^l\, \frac{-z_i^l}{\Sigma_{z^l,ii}}\, \exp\left(-\frac{(z_i^l)^2}{2\Sigma_{z^l,ii}}\right) \tag{B38}$$

$$+\mu_{z^l,i}\frac{1}{\sqrt{2\pi\Sigma_{z^l,ii}}}\int_{-\mu_{z^l,i}}^\infty \mathrm{d}z_i^l\, \exp\left(-\frac{(z_i^l)^2}{2\Sigma_{z^l,ii}}\right) \tag{B39}$$

$$=\frac{\sqrt{\Sigma_{z^l,ii}}}{\sqrt{2\pi}}\exp\left(-\frac{\mu_{z^l,i}^2}{2\Sigma_{z^l,ii}}\right) \tag{B40}$$

$$+\frac{\mu_{z^l,i}}{2}\left(1+\mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right). \tag{B41}$$

For the covariance of post-activations, we distinguish the cases $i = j$ and $i \neq j$, starting with the former by calculating its second moment as

$$\langle\phi(z_i^l)\phi(z_i^l)^\mathsf{T}\rangle_{z^l\sim\mathcal{N}(\mu_{z^l},\Sigma_{z^l})} \tag{B42}$$

$$=\frac{1}{\sqrt{2\pi\Sigma_{z^l,ii}}}\int_0^\infty \mathrm{d}z_i^l\,(z_i^l)^2\, \exp\left(-\frac{1}{2\Sigma_{z^l,ii}}(z_i^l-\mu_{z^l,i})^2\right) \tag{B43}$$

$$=-\frac{\sqrt{\Sigma_{z^l,ii}}}{\sqrt{2\pi}}\mu_{z^l,i}\exp\left(-\frac{\mu_{z^l,i}^2}{2\Sigma_{z^l,ii}}\right) \tag{B44}$$

$$+\frac{\Sigma_{z^l,ii}}{2}\left(1+\mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right) \tag{B45}$$

$$+\sqrt{\frac{2}{\pi}}\mu_{z^l,i}\sqrt{\Sigma_{z^l,ii}}\exp\left(-\frac{\mu_{z^l,i}^2}{2\Sigma_{z^l,ii}}\right) \tag{B46}$$

$$+\frac{\mu_{z^l,i}^2}{2}\left(1+\mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right) \tag{B47}$$

$$=\frac{\sqrt{\Sigma_{z^l,ii}}\,\mu_{z^l,i}}{\sqrt{2\pi}}\exp\left(-\frac{\mu_{z^l,i}^2}{2\Sigma_{z^l,ii}}\right) \tag{B48}$$

$$+\frac{\Sigma_{z^l,ii}+\mu_{z^l,i}^2}{2}\left(1+\mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right). \tag{B49}$$

Non-linearity	Interaction function
$\phi(z) = \mathrm{ReLU}(z)$	$f_{\mu,i} = \frac{\sqrt{\Sigma_{z^l,ii}}}{\sqrt{2\pi}} \exp\left(-\frac{\mu_{z^l,i}^2}{2\Sigma_{z^l,ii}}\right) + \frac{\mu_{z^l,i}}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right)$

$$f_{\Sigma,ii} = \frac{\Sigma_{z^l,ii}}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right) + \frac{\mu_{z^l,i}^2}{4} - \left(\frac{\sqrt{\Sigma_{z^l,ii}}}{\sqrt{2\pi}}\exp\left(-\frac{\mu_{z^l,i}^2}{2\Sigma_{z^l,ii}}\right) + \frac{\mu_{z^l,i}}{2}\mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right)^2$$

$$
\begin{aligned}
f_{\Sigma,ij} =\; & \frac{\sqrt{\det(\tilde{\Sigma}_{z^l})}}{2\pi}\exp\left(-\tfrac{1}{2}\tilde{\mu}_{z^l}^{\mathsf{T}}\tilde{\Sigma}_{z^l}^{-1}\tilde{\mu}_{z^l}\right) \\
& + \frac{\sqrt{\det(\tilde{\Sigma}_{z^l})}}{2\pi}\mu_{z^l,j}\frac{\sqrt{\pi\,\tilde{\Sigma}_{z^l,jj}^{-1}}}{\sqrt{2}}\exp\left(-\tfrac{1}{2}\tilde{\Sigma}_{z^l,ii}^{-1}\mu_{z^l,i}^2\right)\exp\left(\frac{\left(\tilde{\Sigma}_{z^l,ji}^{-1}\mu_{z^l,i}\right)^2}{2\tilde{\Sigma}_{z^l,jj}^{-1}}\right)\left(1 + \mathrm{erf}\left(\frac{\left(\tilde{\Sigma}_{z^l}^{-1}\tilde{\mu}_{z^l}\right)_j}{\sqrt{2\tilde{\Sigma}_{z^l,jj}^{-1}}}\right)\right) \\
& + \frac{\sqrt{\det(\tilde{\Sigma}_{z^l})}}{2\pi}\mu_{z^l,i}\frac{\sqrt{\pi\,\tilde{\Sigma}_{z^l,ii}^{-1}}}{\sqrt{2}}\exp\left(-\tfrac{1}{2}\tilde{\Sigma}_{z^l,jj}^{-1}\mu_{z^l,j}^2\right)\exp\left(\frac{\left(\tilde{\Sigma}_{z^l,ij}^{-1}\mu_{z^l,j}\right)^2}{2\tilde{\Sigma}_{z^l,ii}^{-1}}\right)\left(1 + \mathrm{erf}\left(\frac{\left(\tilde{\Sigma}_{z^l}^{-1}\tilde{\mu}_{z^l}\right)_i}{\sqrt{2\tilde{\Sigma}_{z^l,ii}^{-1}}}\right)\right) \\
& + \left[\mu_{z^l,i}\,\mu_{z^l,j} - \tilde{\Sigma}_{z^l,ij}^{-1}\det\left(\tilde{\Sigma}_{z^l}\right)\right]\left[\frac{1}{2}\mathrm{erf}\left(\frac{\sqrt{2}\,\mu_{z^l,i}}{\sqrt{\Sigma_{z^l,ii}}}\right) + \frac{1}{2}\mathrm{erf}\left(\frac{\sqrt{2}\,\mu_{z^l,j}}{\sqrt{\Sigma_{z^l,jj}}}\right) + F_{\tilde{\mu}_{z^l},\tilde{\Sigma}_{z^l}}(0,0)\right] \\
& - \left[\frac{\sqrt{\Sigma_{z^l,ii}}}{\sqrt{2\pi}}\exp\left(-\frac{\mu_{z^l,i}^2}{2\Sigma_{z^l,ii}}\right) + \frac{\mu_{z^l,i}}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right)\right] \\
& \quad \times \left[\frac{\sqrt{\Sigma_{z^l,jj}}}{\sqrt{2\pi}}\exp\left(-\frac{\mu_{z^l,j}^2}{2\Sigma_{z^l,jj}}\right) + \frac{\mu_{z^l,j}}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_{z^l,j}}{\sqrt{2\Sigma_{z^l,jj}}}\right)\right)\right]
\end{aligned}
$$

$\phi(z) = z + \alpha\,z^2$	$f_{\mu,i} = \mu_{z^l,i} + \alpha\left(\mu_{z^l,i}\right)^2 + \alpha\,\Sigma_{z^l,ii}$

$$f_{\Sigma,ij} = \Sigma_{z^l,ij} + 2\alpha\,\Sigma_{z^l,ij}\left(\mu_{z^l,i} + \mu_{z^l,j}\right) + 2\alpha^2\left(\Sigma_{z^l,ij}\right)^2 + 4\alpha^2\,\mu_{z^l,i}\,\Sigma_{z^l,ij}\,\mu_{z^l,j}$$

TABLE II. Interaction functions for different non-linearities ϕ. We assume $z^l \sim \mathcal{N}(\mu_{z^l}, \Sigma_{z^l})$ in both examples. For ReLU, $\tilde{\mu}_{z^l}$ and $\tilde{\Sigma}_{z^l}$ denote the marginalized mean and covariance with respect to $\tilde{z}^l = (z_i^l, z_j^l)^{\mathsf{T}}$, and $F_{\tilde{\mu}_{z^l},\tilde{\Sigma}_{z^l}}(x,y)$ denotes the corresponding cumulative distribution function.

Combining with the expression for the mean μ_{y^l} then yields the diagonal terms of the covariance:

$$\Sigma_{y^l,ii} \tag{B50}$$

$$= \langle \phi(z_i^l)\,\phi(z_i^l)\rangle_{z^l \sim \mathcal{N}(\mu_{z^l},\Sigma_{z^l})} - \left(\mu_{y^l,i}\right)^2 \tag{B51}$$

$$= \frac{\Sigma_{z^l,ii}}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right) + \frac{\mu_{z^l,i}^2}{4} \tag{B52}$$

$$- \left(\frac{\sqrt{\Sigma_{z^l,ii}}}{\sqrt{2\pi}}\exp\left(-\frac{1}{2\Sigma_{z^l,ii}}\mu_{z^l,i}^2\right) + \frac{\mu_{z^l,i}}{2}\,\mathrm{erf}\left(\frac{\mu_{z^l,i}}{\sqrt{2\Sigma_{z^l,ii}}}\right)\right)^2 . \tag{B53}$$

In the case $i \neq j$, we look at the joint distribution of (z_i, z_j) and denote the marginalized mean and covariance by $\tilde{\mu}_z = (\mu_{z,i}, \mu_{z,j})^{\mathsf{T}}$ and $\tilde{\Sigma}_z = \begin{pmatrix} \Sigma_{z,ii} & \Sigma_{z,ij} \\ \Sigma_{z,ji} & \Sigma_{z,jj} \end{pmatrix}$. For the second moment, we obtain

$$\langle \phi(z_i^l)\,\phi(z_i^l)^{\mathsf{T}}\rangle_{z^l \sim \mathcal{N}(\mu_{z^l},\Sigma_{z^l})} \tag{B54}$$

$$= \frac{1}{\sqrt{(2\pi)^2 \det(\tilde{\Sigma}_z)}}\int_0^\infty \mathrm{d}z_i^l \int_0^\infty \mathrm{d}z_j^l\, z_i^l\, z_j^l \tag{B55}$$

$$\times \exp\left(-\frac{1}{2}(\tilde{z}^l - \mu_z)^{\mathsf{T}}\tilde{\Sigma}_z^{-1}(\tilde{z}^l - \mu_z)\right) \tag{B56}$$

$$= \frac{\sqrt{\det(\tilde{\Sigma}_z)}}{2\pi}\exp\left(-\frac{\tilde{\mu}_z^{\mathsf{T}}\tilde{\Sigma}_z^{-1}\tilde{\mu}_z}{2}\right) \tag{B57}$$

$$+ \frac{\sqrt{\det(\tilde{\Sigma}_z)}}{2\pi}\tilde{\Sigma}_{z^l,jj}^{-1}\mu_{z^l,j}\frac{\sqrt{\pi}}{\sqrt{2\tilde{\Sigma}_{z^l,jj}^{-1}}} \tag{B58}$$

$$\times \exp\left(-\frac{\tilde{\Sigma}_{z^l,ii}^{-1}\mu_{z^l,i}^2}{2}\right)\exp\left(\frac{\left(\tilde{\Sigma}_{z^l,ji}^{-1}\mu_{z^l,i}\right)^2}{2\tilde{\Sigma}_{z^l,jj}^{-1}}\right) \tag{B59}$$

$$\times \left[1 + \mathrm{erf}\left(\frac{\left(\tilde{\Sigma}_{z^l}^{-1}\tilde{\mu}_{z^l}\right)_j}{\sqrt{2\tilde{\Sigma}_{z^l,jj}^{-1}}}\right)\right] \tag{B60}$$

$$+ \frac{\sqrt{\det(\tilde{\Sigma}_{z^l})}}{2\pi} \tilde{\Sigma}_{z^l, ii}^{-1} \mu_{z^l, i} \frac{\sqrt{\pi}}{\sqrt{2\tilde{\Sigma}_{z^l, ii}^{-1}}} \tag{B61}$$

$$\times \exp\left(-\frac{\tilde{\Sigma}_{z^l, jj}^{-1} \mu_{z^l, j}^2}{2}\right) \exp\left(\frac{\left(\tilde{\Sigma}_{z^l, ij}^{-1} \mu_{z^l, j}\right)^2}{2\tilde{\Sigma}_{z^l, ii}^{-1}}\right) \tag{B62}$$

$$\times \left[1 + \mathrm{erf}\left(\frac{(\tilde{\Sigma}_{z^l}^{-1} \tilde{\mu}_{z^l})_i}{\sqrt{2\tilde{\Sigma}_{z^l, ii}^{-1}}}\right)\right] \tag{B63}$$

$$+ \left(\mu_{z^l, i} \mu_{z^l, j} - \tilde{\Sigma}_{z^l, ij}^{-1} \det(\tilde{\Sigma}_{z^l})\right) \tag{B64}$$

$$\times \left[\frac{1}{2}\mathrm{erf}\left(\frac{\sqrt{2}\,\mu_{z^l, i}}{\sqrt{\Sigma_{z^l, ii}}}\right) + \frac{1}{2}\mathrm{erf}\left(\frac{\sqrt{2}\,\mu_{z^l, j}}{\sqrt{\Sigma_{z^l, jj}}}\right) + F_{\tilde{\mu}_{z^l}, \tilde{\Sigma}_{z^l}}(0,0)\right], \tag{B65}$$

where $\tilde{\mu}_{z^l}$ and $\tilde{\Sigma}_{z^l}$ denote the marginalized mean and covariance with respect to $\tilde{z}^l = (z_i^l, z_j^l)^\mathsf{T}$, and $F_{\tilde{\mu}_{z^l}, \tilde{\Sigma}_{z^l}}(x,y)$ denotes the corresponding cumulative distribution function. $F_{\tilde{\mu}_{z^l}, \tilde{\Sigma}_{z^l}}(0,0)$ is also known as the *quadrant probability*. By substracting $\mu_{y^l, i}\mu_{y^l, j}$, we obtain the expression for the cross-covariances given in Table II.

Contributions from higher-order correlations

Using the Gram-Charlier expansion [**?**] of the probability density function $p_{z_i^l}(z_i^l)$, we can derive approximate expressions for the interaction of higher-order correlations of the pre-activations z^l. As an example, we derive contributions to the mean of the post-activations y^l up to linear order in $G_{z^l, (i,i,i)}^{(3)}$ for ReLU activations:

$$\mu_{y^l, i} = \langle \phi(z_i^l) \rangle_{z^l \sim \mathcal{N}(\mu_{z^l}, \Sigma_{z^l})} \tag{B66}$$

$$= \int_0^\infty \mathrm{d}z_i^l\, z_i^l\, p_{z_i^l}(z_i^l) \tag{B67}$$

$$\approx \int_0^\infty \mathrm{d}z_i^l\, z_i^l\, \frac{1}{\sqrt{2\pi\,\Sigma_{z^l, ii}}} \exp\left(-\frac{\left(z_i^l - \mu_{z^l, i}\right)^2}{2\Sigma_{z^l, ii}}\right) \tag{B68}$$

$$+ \frac{G_{z^l, (i,i,i)}^{(3)}}{3!\sqrt{\Sigma_{z^l, ii}^3}} \int_0^\infty \mathrm{d}z_i^l\, z_i^l \tag{B69}$$

$$\times \left[\left(\frac{z_i^l - \mu_{z^l, i}}{\sqrt{\Sigma_{z^l, ii}}}\right)^3 - 3\frac{\left(z_i^l - \mu_{z^l, i}\right)}{\sqrt{\Sigma_{z^l, ii}}}\right] \tag{B70}$$

$$\times \frac{\exp\left(-\left(z_i^l - \mu_{z^l, i}\right)^2 / 2\Sigma_{z^l, ii}\right)}{\sqrt{2\pi\,\Sigma_{z^l, ii}}} \tag{B71}$$

$$= \frac{\sqrt{\Sigma_{z^l, ii}}\,\mu_{z^l, i}}{\sqrt{2\pi}} \exp\left(-\frac{\mu_{z^l, i}^2}{2\,\Sigma_{z^l, ii}}\right) \tag{B72}$$

$$+ \frac{\mu_{z^l, i}}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_{z^l, i}}{\sqrt{2\,\Sigma_{z^l, ii}}}\right)\right) \tag{B73}$$

$$- \frac{G_{z^l, (i,i,i)}^{(3)}}{2\,\Sigma_{z^l, ii}^2}\left(\Sigma_{z^l, ii}^2 - 1\right)\frac{1}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_{z^l, i}}{\sqrt{2\,\Sigma_{z^l, ii}}}\right)\right) \tag{B74}$$

$$+ \frac{G_{z^l, (i,i,i)}^{(3)}}{3!\,\Sigma_{z^l, ii}^3}\left(3\,\mu_{z^l, i}\Sigma_{z^l, ii}^2 + 2\,\Sigma_{z^l, ii} + \mu_{z^l, i}^3 + \mu_{z^l, i}^2 - 3\right) \tag{B75}$$

$$\times \frac{\sqrt{\Sigma_{z^l, ii}}\,\mu_{z^l, i}}{\sqrt{2\pi}} \exp\left(-\frac{\mu_{z^l, i}^2}{2\Sigma_{z^l, ii}}\right) \tag{B76}$$

Alternatively, if one wishes to compute higher-order cumulants of y^l, this can be done by first evaluating the integrals for higher-order moments, analogously to the computations above for the first and second moment. Cumulants can then be obtained via the relations given by Gardiner [47].

Derivations for quadratic activations

We here consider networks with a quadratic activation function $\phi(z) = z + \alpha z^2$. For any distribution of pre-activations z^l with mean μ_{z^l} and covariance Σ_{z^l}, the mean post-activations are given by

$$\mu_{y^l, i} = \langle z_i^l + \alpha\,(z_i^l)^2 \rangle_{z^l} \tag{B77}$$

$$= \mu_{z^l, i} + \alpha\,(\mu_{z^l, i})^2 + \alpha\,\Sigma_{z^l, ii}. \tag{B78}$$

For the covariance of post-activations, we first calculate the second moment

$$\langle \phi(z_i^l)\,\phi(z_j^l) \rangle_{z^l} = \left\langle \left[z_i^l + \alpha\,(z_i^l)^2\right]\left[z_j^l + \alpha\,(z_j^l)^2\right]\right\rangle_{z^l} \tag{B79}$$

$$= \Sigma_{z^l, ij} + \mu_{z^l, i}\,\mu_{z^l, j} + \alpha\,M_{z^l, (i,j,j)}^{(3)} \tag{B80}$$

$$+ \alpha\,M_{z^l, (j,i,i)}^{(3)} + \alpha^2\,M_{z^l, (i,i,j,j)}^{(4)}, \tag{B81}$$

where $M_{z^l}^{(n)}$ denotes the n-th moment of pre-activations z^l. Combining the expression Eq. (B78) for the mean μ_{y^l} then yields the covariance

$$\Sigma_{y^l, ij} = \langle \phi(z_i^l)\,\phi(z_j^l) \rangle_{z^l} - \langle \phi(z_i^l) \rangle_{z^l}\,\langle \phi(z_j^l) \rangle_{z^l} \tag{B82}$$

$$= \Sigma_{z^l, ij} + 2\,\alpha\,\Sigma_{z^l, ij}\left(\mu_{z^l, i} + \mu_{z^l, j}\right) \tag{B83}$$

$$+ 2\,\alpha^2\left(\Sigma_{z^l, ij}\right)^2 + 4\,\alpha^2\,\mu_{z^l, i}\,\Sigma_{z^l, ij}\,\mu_{z^l, j} + \Sigma_{y^l, ij}\big|_{n>2}, \tag{B84}$$

where $\Sigma_{y^l, ij}\big|_{n>2}$ contains all terms involving cumulants of order $n > 2$. It is given by

$$\Sigma_{y^l, ij}\big|_{n>2} = \alpha\,(1 + 2\alpha\,\mu_{z^l, i})\,G_{z^l, (i,j,j)}^{(3)} \tag{B85}$$

$$+ \alpha\,(1 + 2\alpha\,\mu_{z^l, j})\,G_{z^l, (j,i,i)}^{(3)} \tag{B86}$$

$$+ \alpha^2\,G_{z^l, (i,i,j,j)}^{(4)}. \tag{B87}$$

In these expressions, $G^{(n)}_{z^l, (i_1, i_2, ..., i_n)}$ denotes the n-th cumulant of pre-activations given by $\langle\langle z^l_{i_1} z^l_{i_2} \cdots z^l_{i_n} \rangle\rangle$. Keeping terms in linear order of the quadratic coefficient α, only the third-order cumulant $G^{(3)}_{z^l}$ contributes (see Section IV C of the main text). If the pre-activations z^l are Gaussian distributed $z^l \sim \mathcal{N}(\mu_{z^l}, \Sigma_{z^l})$, all cumulants beyond second order vanish, $G^{(n>2)}_{z^l} = 0$, yielding $\Sigma_{y^l, ij}|_{n>2} = 0$ and consequently the result in Table II.

APPENDIX C DETAILS ON THE OPTIMIZATION OF THE STATISTICAL MODEL

In Section III C of the main text, we introduce the statistical model $g_{\text{stat}} : (\{G^{(n)}_x\}_n, \theta, \phi) \mapsto p(y)$ corresponding to a given network model $g : (x; \theta) \mapsto y$. This model follows from the approximate expressions in Eq. (18) of the main text, which give the mean μ_y and covariance Σ_y of the network output purely in terms of the cumulants $\{G^{(n)}_x\}_n$. Since the expected mean-squared error loss $R_{\text{MSE}}(\{\mu^t_y, \Sigma^t_y\}_t)$ depends solely on μ_y and Σ_y, it can therefore be approximated as a function of the input statistics:

$$R_{\text{MSE}}\big(\{\mu^t_y, \Sigma^t_y\}_t\big) \tag{C88}$$

$$\approx R_{\text{MSE}}\big(\big\{g_\mu\big(\{G^{(n),t}_x\}_n; \theta, \phi\big), g_\Sigma\big(\{G^{(n),t}_x\}_n; \theta, \phi\big)\big\}_t\big) \tag{C89}$$

$$= R_{\text{MSE}}\big(\{G^{(n),t}_x\}_{n,t}; \theta\big). \tag{C90}$$

Minimizing this loss yields optimal parameters θ^* for the statistical model. The corresponding network model $g(x; \theta^*)$ is then dependent to the given set of data correlations $\{G^{(n),t}_x\}_{n,t}$, allowing to investigate their relevance in solving a particular network task.

APPENDIX D INFORMATION PROPAGATION IN NETWORKS WITH QUADRATIC ACTIVATIONS

Fig. 2 in the main text illustrates information propagation in networks with ReLU activations. For completeness, we include here as Supplemental Fig. 7 the analogous illustration for a network with quadratic activations.

APPENDIX E DATA SAMPLE GENERATION FOR MNIST BASED ON GAUSSIAN APPROXIMATION OF INPUT DISTRIBUTION

In Section IV B of the main text, we discuss training networks to solve MNIST using $R_{\text{emp, MSE}}$ (Eq. (20)) with data samples drawn from the Gaussian approximation of the input distribution. For this Gaussian approximation,

means $\hat{\mu}^t_x$ and covariances $\hat{\Sigma}^t_x$ for each class t are estimated empirically from the training data set where we flattened the 28×28 images into 784-dimensional vectors. Due to lack of variability in some pixel values at the image edges, the resulting covariances $\hat{\Sigma}^t_x$ are not positive definite, but only positive semi-definite.

To account for the zero eigenvalues of the covariance, data samples are generated based on a principal component analysis of the covariance matrix. For each class t, we decompose the covariance matrix as

$$\hat{\Sigma}^t_x = VDV^{\mathsf{T}} \tag{E91}$$

with $V = (v_1 | \dots | v_{N_0})$ containing the unit-length eigenvectors v_i and $D = \text{diag}(\lambda_1, \dots, \lambda_{N_0})$ containing the corresponding eigenvalues λ_i of $\hat{\Sigma}^t_x$, which we assume to be ordered according to their size, $\lambda_1 \geq \cdots \geq \lambda_{N_0} \geq 0$. We set a threshold $\vartheta_{\text{PCA}} > 0$ that defines a subspace U spanned by the eigenvectors $\{v_i\}_{i=1,\dots,N_{\text{PCA}}}$ for which $\lambda_i > \vartheta_{\text{PCA}}$. Data samples $\hat{x}^{(d)}|_U$ are then generated with respect to this subspace U and projected back to the input space $\mathbb{R}^{N_0}$ according to

$$\hat{x}^{(d)}|_U \sim \mathcal{N}(0, \text{diag}(\lambda_1, \dots, \lambda_{N_{\text{PCA}}})), \tag{E92}$$

$$\hat{x}^{(d)} = \hat{\mu}^t_x + V \begin{pmatrix} \hat{x}^{(d)}|_U \\ 0 \end{pmatrix}. \tag{E93}$$

For all experiments in Section IV B of the main text, we choose $\vartheta_{\text{PCA}} = 10^{-2}$, corresponding to N_{PCA} between 103 and 234 for the different classes t. Since for all classes t the magnitude of the largest eigenvalue is of order 1, this choice of ϑ_{PCA} ensures including relevant eigenvectors while excluding noise due to finite numerical precision. Since the MNIST training dataset contains 60,000 samples, to allow for a fair comparison between training on Gaussian samples and on the original images, we generated a similarly-sized training dataset of $D = 60,000$ Gaussian samples.

APPENDIX F PROBLEM SETUP FOR INCLUSION OF HIGHER-ORDER STATISTICS IN THE MAIN TEXT

The problem studied in Section IV C of the main text is constructed as follows. We define two classes, $t = \pm 1$, each composed of two Gaussian components + and −, with the following means:

$$\mu^{t=+1, -}_x = (-0.5, 0)^{\mathsf{T}}, \qquad \mu^{t=-1, -}_x = (-1.5, 0)^{\mathsf{T}}; \tag{F94}$$

$$\mu^{t=+1, +}_x = (1.5, 0)^{\mathsf{T}}, \qquad \mu^{t=-1, +}_x = (0.5, 0)^{\mathsf{T}}. \tag{F95}$$

Covariances are isotropic throughout with

$$\Sigma^{t, \pm}_x = 0.05 \, \mathbb{I}. \tag{F96}$$

The outer components $(t = -1, -)$ and $(t = +1, +)$ are weighed by $p_{\text{outer}} = \frac{1}{8}$, while the inner components
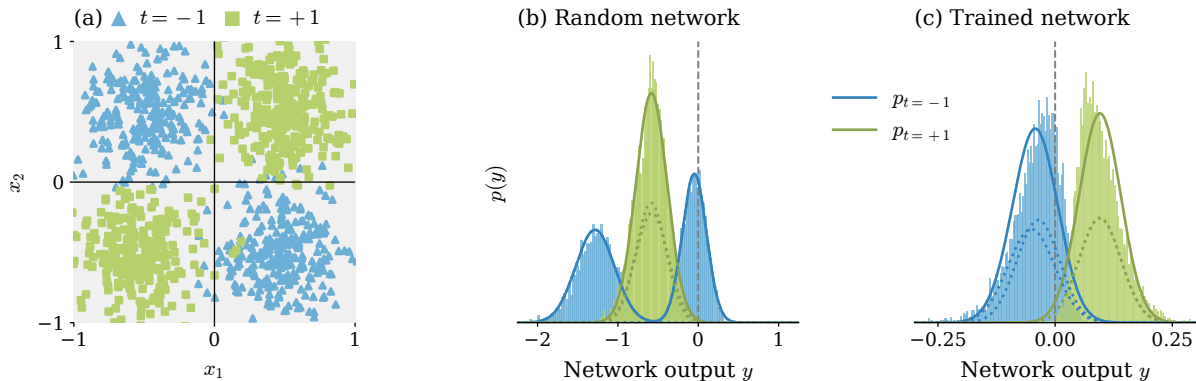
FIG. 7. **Information propagation in networks with quadratic activation function for the XOR problem.** **(a)** The distribution of input data is modeled as a Gaussian mixture. Data samples $x^{(d)}$ (blue and pink dots) are assigned to class labels $t = \pm 1$ based on the respective mixture component. **(b,c)** Distribution of the network output for random (b) and trained (c) parameters. Class-conditional distributions (solid curves) are determined as a superposition of the propagated mixture components (dashed curves) and empirical estimates (blue and pink histograms) are obtained from the test data. Since networks are trained on class labels $t = \pm 1$, the classification threshold is set to $y = 0$ (gray lines). Other parameters: $\phi(z) = z + \alpha z^2$, network depth $L = 1$, width $N = 10$; trained network in (c) achieves $P = 90.46\%$ performance.

$(t = -1, +)$ and $(t = +1, -)$ are weighed by $p_{\text{inner}} = \frac{3}{8}$, as illustrated in Fig. 6**a,b** of the main text. A data sample $x^{(d)}$ is assigned a target label $t^{(d)} \in \{\pm 1\}$ based on the mixture component it is drawn from. Distribution parameters are chosen such that the class-conditional means and covariances of the data are identical,

$$\mu_x^{t=\pm 1} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \qquad \Sigma_x^{t=\pm 1} = \begin{pmatrix} 0.8 & 0 \\ 0 & 0.5 \end{pmatrix}, \qquad \text{(F97)}$$

while the third-order correlations differ

$$G_{x,(i,j,k)}^{(3),t=\pm 1} = \pm 0.75 \, \delta_{ij} \delta_{jk} \delta_{ki} \delta_{i1} . \qquad \text{(F98)}$$

We use training and test datasets of size $D = 10^4$.

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks, in *Advances in Neural Information Processing Systems*, Vol. 25, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Curran Associates, Inc., 2012) pp. 1097–1105.

[2] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, *et al.*, Mastering the game of go with deep neural networks and tree search, Nature **529**, 484 (2016).

[3] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006).

[4] Y. Bahri, J. Kadmon, J. Pennington, S. S. Schoenholz, J. Sohl-Dickstein, and S. Ganguli, Statistical mechanics of deep learning, Annual Review of Condensed Matter Physics **11**, 501 (2020).

[5] H. W. Lin, M. Tegmark, and D. Rolnick, Why does deep and cheap learning work so well?, J. Stat. Phys. **168**, 1223 (2017).

[6] R. Shwartz-Ziv and N. Tishby, Opening the black box of deep neural networks via information, (2017), arXiv:1703.00810 [cs].

[7] A. Jacot, F. Gabriel, and C. Hongler, Neural tangent kernel: Convergence and generalization in neural networks, in *Advances in Neural Information Processing Systems*, Vol. 31 (2018) pp. 8580–8589.

[8] A. M. Saxe, J. L. McClelland, and S. Ganguli, A mathematical theory of semantic development in deep neural networks, Proc. Nat. Acad. Sci. USA **116**, 11537 (2019).

[9] O. Cohen, O. Malka, and Z. Ringel, Learning curves for overparametrized deep neural networks: A field theory perspective, Phys. Rev. Research **3**, 023034 (2021).

[10] R. M. Neal, *Bayesian Learning for Neural Networks* (Springer New York, 1996).

[11] C. K. Williams, Computation with infinite neural networks, Neural Comput. **10**, 1203 (1998).

[12] J. Lee, J. Sohl-Dickstein, J. Pennington, R. Novak, S. Schoenholz, and Y. Bahri, Deep neural networks as gaussian processes, in *International Conference on Learning Representations* (2018).

[13] A. Garriga-Alonso, C. E. Rasmussen, and L. Aitchison, Deep convolutional networks as shallow gaussian processes, in *International Conference on Learning Representations* (2019).

[14] B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, and S. Ganguli, Exponential expressivity in deep neural net-

works through transient chaos, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 57 Morehouse Lane; Red Hook, NY 12571, 2016) pp. 3360–3368.

[15] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein, On the expressive power of deep neural networks, in *Proceedings of the 34th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 70, edited by D. Precup and Y. W. Teh (PMLR, 2017) pp. 2847–2854.

[16] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein, Deep information propagation, (2017), arXiv:1611.01232 [cs, stat].

[17] H. Kleinert, *Gauge fields in condensed matter, Vol. I , SUPERFLOW AND VORTEX LINES Disorder Fields, Phase Transitions* (World Scientific, 1989).

[18] J. Zinn-Justin, *Quantum field theory and critical phenomena* (Clarendon Press, Oxford, 1996).

[19] J. A. Hertz, Y. Roudi, and P. Sollich, Path integral methods for the dynamics of stochastic and disordered systems, Journal of Physics A: Mathematical and Theoretical **50**, 033001 (2017).

[20] M. Helias and D. Dahmen, *Statistical Field Theory for Neural Networks*, Vol. 970 (Springer International Publishing, 2020) p. 203.

[21] Y. LeCun, C. Cortes, and C. J. Burges, The mnist database of handwritten digits (1998).

[22] V. Vapnik, Principles of risk minimization for learning theory, in *Advances in Neural Information Processing Systems*, Vol. 4, edited by J. Moody, S. Hanson, and R. P. Lippmann (Morgan-Kaufmann, 1992) pp. 831–838.

[23] V. N. Vapnik, Adaptive and learning systems for signal processing communications, and control, Statistical learning theory (1998).

[24] R. Kohavi and D. H. Wolpert, Bias plus variance decomposition for zero-one loss functions, in *Proceedings of the Thirteenth International Conference on Machine Learning*, Vol. 96 (1996) pp. 275–283.

[25] D. P. Kingma and J. L. Ba, Adam: A method for stochastic gradient descent, in *ICLR: International Conference on Learning Representations* (2015) pp. 1–15.

[26] I. Loshchilov and F. Hutter, Decoupled weight decay regularization, in *International Conference on Learning Representations* (2019).

[27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32*, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019) pp. 8024–8035.

[28] G. Cybenko, Approximation by superpositions of a sigmoidal function, Mathematics of Control, Signals, and Systems **2**, 303 (1989).

[29] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken, Multilayer feedforward networks with a nonpolynomial activation function can approximate any function, Neural networks **6**, 861 (1993).

[30] A. Pinkus, Approximation theory of the mlp model in neural networks, Acta numerica **8**, 143 (1999).

[31] D. J. MacKay, *Information theory, inference and learning algorithms* (Cambridge university press, 2003).

[32] C. K. I. Williams and D. Barber, Bayesian classification with gaussian processes, IEEE Transactions on Pattern Analysis and Machine Intelligence **12** (1998).

[33] C. K. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning*, 1st ed. (MIT Press, Cambridge, 2006).

[34] E. Dyer and G. Gur-Ari, Asymptotics of wide networks from feynman diagrams, in *International Conference on Learning Representations* (2020).

[35] G. Naveh, O. Ben David, H. Sompolinsky, and Z. Ringel, Predicting the outputs of finite deep neural networks trained with noisy gradients, Phys. Rev. E **104**, 064301 (2021).

[36] S. Yaida, Non-Gaussian processes and neural networks at finite widths, in *Proceedings of The First Mathematical and Scientific Machine Learning Conference*, Proceedings of Machine Learning Research, Vol. 107, edited by J. Lu and R. Ward (PMLR, Princeton University, Princeton, NJ, USA, 2020) pp. 165–192.

[37] G. Deco and W. Brauer, Higher order statistical decorrelation without information loss, in *Proceedings of the 7th International Conference on Neural Information Processing Systems*, NIPS'94 (MIT Press, Cambridge, MA, USA, 1994) pp. 247–254.

[38] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová, Modeling the Influence of Data Structure on Learning in Neural Networks: The Hidden Manifold Model, Phys. Rev. X **10**, 041044 (2020).

[39] S. Goldt, G. Reeves, M. Mézard, F. Krzakala, and L. Zdeborová, The Gaussian equivalence of generative models for learning with two-layer neural networks, (2020), arXiv:2006.14709 [stat.ML].

[40] B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mézard, and L. Zdeborová, Capturing the learning curves of generic features maps for realistic data sets with a teacher-student model, (2021), arXiv:2102.08127 [stat.ML].

[41] G. Yang and E. J. Hu, Tensor programs iv: Feature learning in infinite-width neural networks, in *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 139, edited by M. Meila and T. Zhang (PMLR, 2021) pp. 11727–11737.

[42] C. Fang, J. Lee, P. Yang, and T. Zhang, Modeling from features: a mean-field framework for over-parameterized deep neural networks, in *Proceedings of Thirty Fourth Conference on Learning Theory*, Proceedings of Machine Learning Research, Vol. 134, edited by M. Belkin and S. Kpotufe (PMLR, 2021) pp. 1887–1936.

[43] M. E. A. Seddik, C. Louart, M. Tamaazousti, and R. Couillet, Random Matrix Theory Proves that Deep Learning Representations of GAN-data Behave as Gaussian Mixtures, in *International Conference on Machine Learning* (PMLR, 2020) pp. 8573–8582.

[44] D. A. Roberts, S. Yaida, and B. Hanin, *The Principles of Deep Learning Theory: An Effective Theory Approach to Understanding Neural Networks* (Cambridge University Press, 2022) (to be published).

[45] A. Krizhevsky and G. Hinton, *Learning multiple layers of features from tiny images*, Tech. Rep. (2009).

[46] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE*

Conference on Computer Vision and Pattern Recognition (CVPR) (2016).

[47] C. W. Gardiner, *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, 2nd ed., Springer Series in Synergetics No. 13 (Springer-Verlag, Berlin, 1985).