

Weight fluctuations in deep linear neural networks and a derivation of the inverse-variance flatness relation

Markus Gross,^{*} Arne P. Raulf, and Christoph R ath

*Deutsches Zentrum f ur Luft- und Raumfahrt, Institute for AI Safety and Security,
Rathausallee 12, 53757 Sankt Augustin, Germany and Wilhelm-Runge-Stra e 10, 89081 Ulm, Germany*



(Received 24 November 2023; accepted 3 June 2024; published 25 July 2024)

We investigate the stationary (late-time) training regime of single- and two-layer underparameterized linear neural networks within the continuum limit of stochastic gradient descent (SGD) for synthetic Gaussian data. In the case of a single-layer network in the weakly underparameterized regime, the spectrum of the noise covariance matrix deviates notably from the Hessian, which can be attributed to the broken detailed balance of SGD dynamics. The weight fluctuations, are in this case, generally anisotropic but effectively experience an isotropic loss. For an underparameterized two-layer network, we describe the stochastic dynamics of the weights in each layer and analyze the associated stationary covariances. We identify the interlayer coupling as a distinct source of anisotropy for the weight fluctuations. In contrast to the single-layer case, the weight fluctuations are effectively subject to an anisotropic loss, the flatness of which is inversely related to the fluctuation variance. We thereby provide an analytical derivation of the recently observed inverse-variance flatness relation in a model of a deep linear neural network.

DOI: [10.1103/PhysRevResearch.6.033103](https://doi.org/10.1103/PhysRevResearch.6.033103)

I. INTRODUCTION

A. Motivation

Deep neural networks have proven to be formidable tools in the area of machine learning [1–3]. Their ability to generalize from data often hinges on the dynamics of their training processes, notably the stochastic gradient descent (SGD) method. While the success of these networks in practical applications is well documented, the underlying principles governing their learning dynamics, especially concerning weight fluctuations and loss landscapes, are still not fully understood. In order to address some of these aspects analytically, we focus here on one- and two-layer linear networks trained on synthetic Gaussian data with noisy labels. These simplified models have been proven to share many aspects of the learning dynamics of full nonlinear networks [4–9]. In fact, in the stationary regime at the end of training, a deep nonlinear network typically follows a linearized dynamics (multivariate Ornstein-Uhlenbeck process with correlated noise) [10–12]. We employ here the continuum description of SGD in terms of a Langevin equation, which allows us to analyze training dynamics and steady-state properties from a statistical physics perspective [13–15].

One motivation of the present study stems from the recently observed inverse relation between the variance of a

weight fluctuation mode and the flatness of the loss along the direction of that mode [16]. While this inverse-variance flatness relation (IVFR) has been investigated from various perspectives [16–19], we seek here a first-principles derivation in a neural network model. This requires, as a preparatory step, also a detailed understanding of the various sources of anisotropy of the weight fluctuations, which we address in the first part of this study, focusing on single-layer linear network. Complementing recent papers on this aspect [20,21], we consider in particular the slightly oversampled regime near the interpolation threshold (i.e., the point where the sample number equals the input dimension) and investigate the dependencies on the parameter count, sample number, and label noise. We find that arguably the simplest model showing an IVFR is a two-layer linear network. In fact, the interlayer coupling of the weights is a crucial ingredient responsible for the loss anisotropy, which therefore does not occur in the case of a single-layer network. However, even for a single-layer network, the weight fluctuations are anisotropic except in the highly oversampled regime. In a physical sense, the IVFR indicates that the effective noise temperature is higher in regions where the loss valleys are narrow, thus effectively biasing the model towards flat regions of the loss. Notably, wide loss valleys are typically associated with improved generalization ability of the network [11,22,23].

Another motivation of our study stems from the observation that, while the (deterministic) gradient flow dynamics of deep linear networks has been addressed in several studies [4,9,24], the corresponding layerwise stochastic gradient flow dynamics has received less attention so far. A notable exception is the online learning in two-layer nets, which can be assessed analytically in certain regimes [25–28]. Recently, effects of stochasticity in two-layer linear nets have

^{*}Contact author: markus.gross@dlr.de

been analyzed to explain the observed bias towards sparse weight distributions [29]. In the present paper, we focus on supervised learning of minibatches of data and study the associated stochastic dynamics of the weights in each layer within a Langevin approach.

B. Related work

Continuum modeling of SGD. The modeling of SGD in terms of a stochastic differential equation (SDE)—in the sense that the former emerges from a Euler-Maruyama time-discretization of the latter—is an established approach in the literature [10–12,30–35]. In particular, the steady-state dynamics at the end of training can be described as a multivariate Ornstein-Uhlenbeck process [10,11,20,36,37]. Despite being formally derived in the limit of infinitely small learning rate [38], the continuum approximation to SGD is typically found to be meaningful also for finite-learning rates [39] and it has been successfully used to study various aspects of the training of neural networks. The noise due to the minibatch sampling of gradients in SGD plays an important role for the generalization ability of a model as it directs the weights towards wider minima of the loss, which improves generalization [11,18,22,40–48] and helps escaping from saddle points and sharp minima [33,44,49,50].

For sufficiently large minibatch sizes, the central limit theorem implies a Gaussian distribution of the gradient noise [10,11,51,52], although deviations towards more heavy-tailed distributions have been observed [53–55]. Moreover, the gradient noise is assumed to be uncorrelated in time (see, however, Ref. [56]). A crucial property of the noise is its anisotropy, i.e., its dependence on the direction in parameter space [16,21,42,44,50]. In the simplest case, this anisotropy is modeled by taking its covariance to be proportional to the Hessian [20,51,57]. However, this approximation becomes exact only in the infinite-sample limit (see discussion in Sec. II B 2 below). In fact, the discrepancy between noise covariance and Hessian implies that detailed balance is broken in the stationary state reached by SGD, giving rise to intriguing nonequilibrium behavior such as persistent currents [12,19,20]. In the context of linear regression, effects of parameter-dependent (multiplicative) noise have been addressed in [36]. The relation of the noise covariance matrix to the Hessian and the Fisher matrix is further analyzed in [58].

(Deep) Linear nets. Single-layer linear networks often admit closed form analytic solutions, which qualitatively explain many properties observed in deep nonlinear networks (see, e.g., [7,8,13,59,60]). In the limit of large width, the neural tangent kernel provides a linear approximation of a nonlinear network around initialization [61,62]. Moreover, the steady-state dynamics of a trained nonlinear network becomes effectively linear—a fact that, e.g., is utilized in the construction of Bayesian sampling methods [10,12]. Notably, linear networks with two or more layers have a nonlinear gradient descent dynamics of the weights and feature a nonconvex loss landscape [4,63,64]. Deep linear nets have emerged as important toy models for the study of more general networks with nonlinear activation functions, enabling fundamental insights into training dynamics, generalization ability, implicit regularization, and spectral bias [4–6,9,24,29,65–70].

Inverse-variance flatness relation (IVFR). The inverse relation between the magnitude of the steady-state weight fluctuations and the flatness of the loss has been first discussed in [16]. As emphasized in [17], using the formalism of [71–73], the actual effective potential (understood as the inverse of the covariance matrix of the weight distribution) is not the loss, but the product of the inverse of the (generally anisotropic) diffusion term and the loss. This indeed follows also from the classic theory of stochastic processes [20,74,75] and has been realized in various machine-learning contexts [10,12]. In Ref. [18], IVFR has been derived within a simple dynamic model of slow and fast degrees of freedom. The stability of IVFR under dropout has been discussed in [76].

C. Outline of the present paper

We first consider learning of synthetic data via SGD in a single-layer linear network (Sec. II). Employing the continuum description of the training dynamics, we show that the spectrum of the noise covariance matrix in the stationary state deviates significantly from the Hessian, which we relate to the well-known broken detailed balance of the SGD dynamics. As a consequence, the weight fluctuations in the underparameterized regime are generally anisotropic, i.e., their amplitude depends on the directions of the principle modes.

Proceeding to a two-layer linear network (Sec. III), we determine a continuum description of the layerwise training dynamics under SGD. We obtain coupled stochastic differential equations for the weights and determine expressions for the drift and diffusion matrices in the large batch regime. These matrices turn out to be generally rank deficient due to the coupling between the network layers. As a consequence the weight fluctuations in each layer are anisotropic.

We finally apply our results to study the shape of the loss under a perturbation along a principle weight fluctuation mode. In the case of a single-layer network, the loss remains isotropic despite the anisotropy of the weight fluctuations. For a two-layer network, we confirm the recently observed [16] inverse relation between the amplitude of a weight fluctuation mode and the flatness of the loss along the direction of that mode. The coupling between the layers of the network is crucial for our result. The present paper thus provides a derivation of the inverse-variance flatness relation in an analytically tractable model of a deep neural network.

II. SINGLE-LAYER LINEAR NETWORKS

In this section, we consider the fluctuating dynamics of the weights of an underparameterized single-layer linear network. The network evolves under a mean-square loss and thus solves a linear regression problem. The present linear framework can be matched to a nonlinear network in a stationary state (see Appendix A).

A. Linear regression

We begin by recapitulating the theory of linear regression that is necessary for the subsequent parts of our study [77,78]. The training data consists of P samples $\mathbf{x}^\mu \in \mathbb{R}^N$ ($\mu \in 1, \dots, P$), the components being independent and identically distributed (i.i.d.) Gaussian random variables

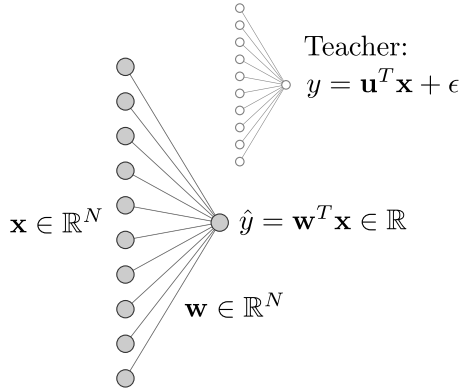


FIG. 1. Single-layer linear neural network model. The student network (left) transforms the input \mathbf{x} to the output \hat{y} using its learned weights \mathbf{w} , while the teacher network (right) generates the targets (“labels”) y using fixed random weights \mathbf{u} and noise ϵ (see text).

$x_i^\mu \sim \mathcal{N}(\mu_x, \sigma_x^2)$ with mean μ_x and variance $\sigma_x^2 = 1/N$ [79]. The targets $y^\mu \in \mathbb{R}$ are provided by a linear teacher model,

$$y^\mu = \mathbf{u}^T \mathbf{x}^\mu + \epsilon^\mu, \quad (2.1)$$

with a fixed random weight vector $\mathbf{u} \in \mathbb{R}^N$ with $u_i \sim \mathcal{N}(0, \sigma_u^2)$ (we typically take $\sigma_u^2 = 1$) and a residual error $\epsilon^\mu \sim \mathcal{N}(0, \sigma_\epsilon^2)$ modeled as a i.i.d. Gaussian noise. The error term describes the residual dependence between inputs and targets that can not be captured by a linear relationship. In a slight abuse of terminology, we will also in the present regression context refer to y and ϵ as “label” and “label noise”, respectively [20,21]. The complete set of the training data $(\mathbf{x}^\mu, y^\mu)_{\mu \in \{1, \dots, P\}}$ is called *batch*, while a randomly chosen subset $(\mathbf{x}^\mu, y^\mu)_{\mu \in \mathcal{B}}$ of size $S < P$ with $\mathcal{B} = \{\pi(1), \dots, \pi(S)\}$ (and π being a random permutation of the P elements) is called a *minibatch*. The case of statistically independent random labels y^μ , which occasionally is useful in order to analyze fluctuations, corresponds to the replacements $\mathbf{u} = 0$, $\epsilon^\mu \rightarrow y^\mu$. We define a design matrix $X \in \mathbb{R}^{N \times P}$ with the columns representing the samples \mathbf{x}^μ , and, correspondingly, a noise vector $\boldsymbol{\epsilon} = (\epsilon^1, \epsilon^2, \dots, \epsilon^P) \in \mathbb{R}^{1 \times P}$, such that Eq. (2.1) can be written as

$$Y = \mathbf{u}^T X + \boldsymbol{\epsilon} \in \mathbb{R}^{1 \times P}. \quad (2.2)$$

The linear student network is defined by $y = \mathbf{w}^T \mathbf{x}$, with the learnable weights $\mathbf{w} \in \mathbb{R}^N$ (see Fig. 1). The full-batch loss takes the form of a least-squares regression model,

$$\begin{aligned} L(\mathbf{w}) &= \frac{1}{2P} \|\mathbf{w}^T X - Y\|^2 = \frac{1}{2P} (\mathbf{w}^T X - Y)(\mathbf{w}^T X - Y)^T \\ &= \frac{1}{2} (\mathbf{w}^T \Sigma^{xx} \mathbf{w} - \mathbf{w}^T \Sigma^{xy} - \Sigma^{yx} \mathbf{w} + YY^T/P), \end{aligned} \quad (2.3)$$

with the moment matrices

$$\Sigma^{xx} = H = \frac{1}{P} X X^T, \quad \Sigma^{xy} = \frac{1}{P} X Y^T. \quad (2.4)$$

Here, $H \in \mathbb{R}^{N \times N}$ denotes the Hessian, which is identical to Σ^{xx} for this model [see Eq. (2.9) below]. Over- and underparameterized regimes correspond to $N > P$ and $N < P$, respectively. We will typically consider $\mu_x = 0$; hence $\Sigma^{xx, xy}$ represent covariance matrices and, due to the i.i.d. assumption

for \mathbf{x}^μ , in limit $P \rightarrow \infty$ one has $H_{ii} = \sigma_x^2$, while $H_{ij} = 0$ for $i \neq j$. In the following we make use of the Moore-Penrose pseudoinverse \mathcal{A}^+ of a matrix \mathcal{A} . The pseudoinverse can be expressed as $\mathcal{A}^+ = (\mathcal{A}^T \mathcal{A})^{-1} \mathcal{A}^T$ if \mathcal{A} has linearly independent columns, and as $\mathcal{A}^+ = \mathcal{A}^T (\mathcal{A} \mathcal{A}^T)^{-1}$ if \mathcal{A} has linearly independent rows.

Using the elementary derivative rule $\partial(\mathbf{x}^T \mathbf{a})/\partial \mathbf{x} = \mathbf{a} = \partial(\mathbf{a}^T \mathbf{x})/\partial \mathbf{x}$, the condition for the minimum of the loss can be written as

$$\nabla_{\mathbf{w}} L = \Sigma^{xx} \mathbf{w}^* - \Sigma^{xy} = 0. \quad (2.5)$$

In the underparameterized case, the solution readily follows as $\mathbf{w}_{\text{up}}^* = (\Sigma^{xx})^{-1} \Sigma^{yx}$. By contrast, in the overparameterized case, Σ^{xx} is not full rank, and hence infinitely many solutions can be obtained by adding an arbitrary vector of the nullspace of Σ^{xx} to any particular solution w^* . Taking advantage of this freedom, one thus typically considers the minimum norm solution obtained from the constrained minimization problem $\tilde{L} = \|\mathbf{w}\|^2 - (\mathbf{w}^T X - Y)\boldsymbol{\kappa}$, where $\boldsymbol{\kappa} \in \mathbb{R}^P$ is a vector of Lagrange multipliers [80]. The solution in the overparameterized case then follows as $\mathbf{w}_{\text{op}}^* = X(X^T X)^{-1} Y^T = (X X^T)^+ X Y^T$, where we used the fact that the so-called gram matrix $X^T X \in \mathbb{R}^{P \times P}$ now has full rank and applied the general identity $X(X^T X)^+ = (X X^T)^+ X$ of the pseudoinverse. Defining

$$J = \frac{1}{P} \sum_{\mu} \mathbf{x}^\mu \epsilon^\mu = \frac{1}{P} X \boldsymbol{\epsilon}^T \in \mathbb{R}^N, \quad (2.6)$$

the minimum-norm solution of the linear regression problem can be compactly written as

$$\begin{aligned} \mathbf{w}^* &= (\Sigma^{xx})^+ \Sigma^{yx} \\ &= \begin{cases} (X X^T)^{-1} X Y^T = \mathbf{u} + H^{-1} J, & N < P \quad (\text{underpar.}) \\ X (X^T X)^{-1} Y = H^+ H \mathbf{u} + H^+ J, & N > P. \quad (\text{overpar.}) \end{cases} \end{aligned} \quad (2.7)$$

The quantity $H^+ H$ can be understood as a projector onto the orthogonal complement of the nullspace of H .

Notably, by using only general properties of the pseudoinverse [and not Eq. (2.7)], the loss in Eq. (2.3) can be rewritten as [21]

$$\begin{aligned} L(\mathbf{w}) &= \frac{1}{2} (\mathbf{w} - \mathbf{u} - H^+ J)^T H (\mathbf{w} - \mathbf{u} - H^+ J) \\ &\quad - \frac{1}{2} J^T H^+ J + \frac{1}{2P} \boldsymbol{\epsilon} \boldsymbol{\epsilon}^T. \end{aligned} \quad (2.8)$$

Using the symmetry of H and the elementary relation $H^+ H H^+ = H^+$ of the pseudoinverse, one can readily show that Eq. (2.8) is minimized by Eq. (2.7). Accordingly, Eq. (2.8) can be reformulated as

$$L(\mathbf{w}) = \frac{1}{2} (\mathbf{w} - \mathbf{w}^*)^T H (\mathbf{w} - \mathbf{w}^*) - \frac{1}{2} J^T H^+ J + \frac{1}{2P} \boldsymbol{\epsilon} \boldsymbol{\epsilon}^T, \quad (2.9)$$

making the minimizing property of \mathbf{w}^* apparent. In the overparameterized case, the last two terms in Eq. (2.9) cancel, such that $L(\mathbf{w}^*) = 0$, showing that the model is able to perfectly fit the data generated via Eq. (2.1). This cancellation can be proven by using the SVD $X = V D^{1/2} U^T$ [$V \in \mathbb{R}^{N \times N}$, $D^{1/2} \in \mathbb{R}^{N \times P}$, $U \in \mathbb{R}^{P \times P}$, with V, U being orthogonal

and $D^{1/2}$ having $\min(P, N)$ singular values on the diagonal] and the relation $X^+ = U(D^{1/2})^+V^T$, which implies that $X^T(XX^T)^+X = \mathbb{1}_{P \times P}$. By contrast, the latter relation does not hold in the underparameterized case, since $U\mathbb{1}_{P \times P}^N U^T \neq \mathbb{1}_{P \times P}$ if $N < P$, where $\mathbb{1}_{P \times P}^N$ is the unit matrix with the last $P - N$ elements on the diagonal being zero.

B. Dynamics

Next, we consider the continuum approximation of the stochastic gradient descent, i.e., the *stochastic gradient flow* (SGF), which we will use to study the spectrum of the noise and its dependence on the number of parameters and on the label noise. We consider the stationary state at the end of training, where the weights fluctuate around the solution \mathbf{w}^* .

1. SGD and continuum approximation

With the minibatch loss ($S = |\mathcal{B}|$ is the minibatch size)

$$L_B(\mathbf{w}) = \frac{1}{S} \sum_{\mu \in \mathcal{B}} \ell^\mu(\mathbf{w}), \quad \ell^\mu(\mathbf{w}) = \frac{1}{2}(\mathbf{w}^T \mathbf{x}^\mu - y^\mu)^2, \quad (2.10)$$

and the learning rate λ , the weights evolve according to SGD as

$$\begin{aligned} \mathbf{w}(\tau + 1) &= \mathbf{w}(\tau) - \lambda \nabla_{\mathbf{w}} L_B(\mathbf{w}(\tau)) \\ &= \mathbf{w}(\tau) - \lambda \nabla L(\mathbf{w}(\tau)) - \lambda \boldsymbol{\xi}(\mathbf{w}(\tau)), \end{aligned} \quad (2.11)$$

where τ denotes the discrete time step. The minibatch noise $\boldsymbol{\xi}$ is defined as

$$\boldsymbol{\xi}(\mathbf{w}) \equiv \nabla L_B(\mathbf{w}) - \nabla L(\mathbf{w}) \quad (2.12)$$

and fulfils $\nabla L = \langle \nabla L_B \rangle_{\mathcal{B}}$, where $\langle \cdot \rangle_{\mathcal{B}}$ denotes the expectation over the full batch.

In order to proceed to the continuum limit, we follow [11,12,20,32] and assume the learning rate λ to be sufficiently small, such that Eq. (2.11) can be approximated, in the weak sense, by a SDE (Langevin equation),

$$\begin{aligned} \partial_t \mathbf{w}(t) &= -\nabla L(\mathbf{w}(t)) + \sqrt{\lambda} \boldsymbol{\xi}(t), \\ \nabla L(\mathbf{w}) &= H(\mathbf{w} - \mathbf{w}^*), \\ \langle \boldsymbol{\xi}(t) \boldsymbol{\xi}(t')^T \rangle_{\mathcal{B}} &= C(\mathbf{w}) \delta(t - t'), \end{aligned} \quad (2.13)$$

where $t = \tau/\lambda$ is a new time variable. The noise $\boldsymbol{\xi}(t)$ has weight-dependent Gaussian correlations in weight space, as characterized by the correlation matrix $C(\mathbf{w})$, which is further analyzed below. The noise is interpreted in an Ito sense, since, according to Eq. (2.11), $\mathbf{w}(\tau)$ is independent of $\mathbf{w}(\tau + 1)$ [75]. Although the SGD-SDE mapping is not strictly rigorous [38,39], we find that the results obtained from Eqs. (2.13) and (2.11) agree well even for rather large learning rates.

2. Noise covariance

According to Eq. (2.12), the covariance matrix of the noise, averaged over batches (with replacement), is given by (see Appendix C)

$$C = \frac{1}{S} \left[\frac{1}{P} \sum_{\mu=1}^P (\nabla \ell^\mu)(\nabla \ell^\mu)^T - (\nabla L)(\nabla L)^T \right]. \quad (2.14)$$

The expression in the case of batching without replacement differs by the prefactor and vanishes for $S = P$, but approaches the above when $S \ll P$. Note that for batching with replacement, which we use mainly in the present paper, gradient noise is nonzero even if $S = P$. In the case of linear regression, we have $\nabla_{\mathbf{w}} \ell^\mu = (\mathbf{w}^T \mathbf{x}^\mu - y^\mu) \mathbf{x}^\mu$ [see Eq. (2.10)]. Using Eq. (2.1), the covariance takes the form

$$\begin{aligned} C_{ab} &= \frac{1}{S} \left[\frac{1}{P} \sum_{\mu} [\bar{\mathbf{w}}^T \mathbf{x}^\mu - \epsilon^\mu]^2 x_a^\mu x_b^\mu - H \bar{\mathbf{w}} \bar{\mathbf{w}}^T H \right], \\ \bar{\mathbf{w}} &\equiv \mathbf{w} - \mathbf{u}. \end{aligned} \quad (2.15)$$

This result shows that, beside the external source of randomness stemming from the minibatching of input data \mathbf{x}^μ and the label noise ϵ^μ , also the fluctuations of the weights \mathbf{w} contribute to the SGD noise, rendering it state dependent [21,57,81]. Estimating the typical magnitudes of the various terms in Eq. (2.15) (see Appendix D) indicates that the w -dependent terms dominate if $\|\bar{\mathbf{w}}\|^2 \|\sigma_x^2\| \gg \sigma_\epsilon^2$. In the absence of label noise ($\sigma_\epsilon^2 = 0$), this condition is fulfilled at all times. By contrast, if $\sigma_\epsilon^2 > 0$, we write $\mathbf{w} = \mathbf{w}^* + \delta \mathbf{w}$, where $\delta \mathbf{w}$ denotes the weight fluctuations around the stationary solution \mathbf{w}^* [Eq. (2.7)]. Since the typical stationary variance $\langle \delta w_i^2 \rangle \sim (\lambda/S) \sigma_\epsilon^2$ [cf. Eq. (2.34) below and Appendix D] and $\|H^{-1}J\|^2 \sim \sigma_\epsilon^2/\sigma_x^2$, multiplicative noise terms will dominate in the stationary state only if $(\lambda/S) \sigma_x^2 \gg 1$. In the case of nonvanishing label noise ($\sigma_\epsilon^2 > 0$), the additive noise approximation is therefore warranted for typical choices of model parameters and normalized data. Exceptions can arise at early stages of the training, where $\|\bar{\mathbf{w}}\|$ can be large. In the following, we will study in more detail the various approximations to the SGD noise.

Additive noise approximation for finite sample number ($P < \infty$). The additive noise approximation consists of evaluating C [Eq. (2.15)] on the stationary solution $\mathbf{w} = \mathbf{w}^*$ [Eq. (2.7)], assuming low-noise magnitude ($S \gg 1$). For an overparameterized network, each individual loss gradient vanishes for $\mathbf{w} = \mathbf{w}^*$, i.e., $\nabla_{\mathbf{w}} \ell^\mu(\mathbf{w}^*) = 0$ [82]. Consequently, in this case, the noise and the weight fluctuations vanish in the steady state [see also the discussion after Eq. (2.9)],

$$C = 0 \quad \text{if } N > P \text{ (overparameterized)}. \quad (2.16)$$

By contrast, in the underparameterized regime, only the sum over the total batch is zero, $\sum_{\mu} \nabla_{\mathbf{w}} \ell^\mu = 0$. We remark that the analysis here also applies for output dimensions larger than one, i.e., $\mathbf{y}^\mu \in \mathbb{R}^{d_o}$ with $d_o > 1$, since the different output channels are not coupled by the loss.

Since generally $J \neq 0$ [see Eq. (2.6)] for finite sample number, Eq. (2.15) takes the form

$$C_{ab} = \frac{1}{S} \left[\frac{1}{P} \sum_{\mu} (J^T H^+ \mathbf{x}^\mu - \epsilon^\mu)^2 x_a^\mu x_b^\mu \right] = \frac{1}{S} (X K X^T)_{ab}, \quad (2.17)$$




FIG. 2. Deviations from the Hessian approximation of the noise covariance of SGD of a linear network [Eq. (2.17)]. (a) Variance of the matrix elements of $\mathcal{P} = X^+X$, which enters in Eq. (2.18), as a function of the sample number P rescaled by the input dimension N . Curves for various N and P superimpose when expressed in appropriately scaled coordinates. (b) Entries, ordered by magnitude, of the diagonal random matrix K [Eq. (2.18)], which determines the covariance of the SGD noise in Eq. (2.17). The various solid curves correspond to sample numbers ranging between $P = 70$ and $P = 10\,000$ (from bottom to top) with fixed input dimension $N = 50$, averaged over 100 realizations of K . The dotted-black curve represents the approximation $K^{\mu\mu} = (\epsilon^\mu)^2$, where ϵ^μ is a vector of independent Gaussian random numbers of zero mean and variance σ_ϵ^2 . The scaling of $\text{var}(\mathcal{P}) \propto N^{-1}s^{-1.8}$ observed in (a) ensures that the elements of the diagonal matrix K , which characterizes the deviation of the noise from the Hessian [see Eq. (2.17)], are Gaussian i.i.d. for $s \rightarrow \infty$.

with a diagonal matrix in sample space,

$$\begin{aligned} K^{\mu\nu} &= \frac{1}{P} \left[\sum_{j,k}^N \sum_{\gamma}^P \frac{1}{P} \epsilon^\gamma x_k^\gamma H_{kj}^+ x_j^\mu - \epsilon^\mu \right]^2 \delta^{\mu\nu} \\ &= \frac{1}{P} \left[\left(\frac{1}{P} \boldsymbol{\epsilon} X^T H^+ X - \boldsymbol{\epsilon} \right)^\mu \right]^2 \delta^{\mu\nu}. \end{aligned} \quad (2.18)$$

Note that these expressions are general and do not rely on specific assumptions for the statistics of the data or the label noise. In the case of fully random labels y^μ , they apply with the replacement $\epsilon^\mu \rightarrow y^\mu$. Similar forms of C in various limits have also been obtained in [36,83].

In the *overparameterized* case, one readily shows, using identities for the pseudoinverse, that $X^T (X X^T)^+ X = \mathbb{1}$, hence $K = 0$. In the *underparameterized* case, instead, one has $X^T (X X^T)^+ X = X^+ X =: \mathcal{P}$, which is a projector on the row space of X and is independent of the variance of X . In the large sample regime $P \gg N$, a numerical analysis [see Fig. 2(a)] reveals that the variance of the elements of this matrix behaves as $\text{var}(\mathcal{P}) \propto P^{-1.8} N^{-1}$ (the variance is slightly larger on the diagonal), while the mean value along the diagonal $\langle \mathcal{P}_{ii} \rangle \simeq N/P$. For the purpose of estimating the variance of $\boldsymbol{\epsilon} \mathcal{P}$, we write $\mathcal{P} \simeq (N/P) \mathbb{1}_P + P^{-0.9} N^{-0.5} \tilde{\mathcal{P}}$, where $\tilde{\mathcal{P}}$ is a random matrix with i.i.d. entries of zero mean and unit variance. Accordingly, we infer $\text{var}(\boldsymbol{\epsilon} X^+ X)_{ij} \sim P \sigma_\epsilon^2 \text{var}(\mathcal{P}) \rightarrow 0$ as $P \rightarrow \infty$, which implies $K^{\mu\nu} \simeq (\boldsymbol{\epsilon}^\mu)^2 \delta^{\mu\nu}$, i.e., in the relevant oversampled regime and for $P \sigma_\epsilon^2 \gg 1$, K essentially becomes a diagonal matrix of squared Gaussian random numbers. This is illustrated in Fig. 2(b), where the diagonal entries of K (ordered by magnitude) are displayed for various sample sizes P .

In Fig. 3(a), the spectrum of the noise covariance matrix obtained from numerical experiments via minibatch sampling in the stationary state is shown and compared to the theoretical prediction of Eq. (2.17) averaged over the distributions of the data x and the label noise ϵ (which we take here to be

Gaussian i.i.d., see Sec. II A). The good agreement of the two spectra supports the validity of the additive noise approximation in this regime. As shown in Fig. 3(b), the noise spectrum deviates in general from the Hessian approximation of Eq. (2.20), but approaches the latter in the strongly underparameterized limit $P/N \rightarrow \infty$. This is also confirmed by the calculations leading to Eq. (2.20) below. The deviation is caused by the intermixing between the data eigenmodes due to the label noise [see Eq. (2.17)], as illustrated in Fig. 3(c). This becomes also clear by noting that, from the SVD of the design matrix $X = V D^{1/2} U^T$ (with $V \in \mathbb{R}^{N \times N}$, $U \in \mathbb{R}^{P \times N}$, and a positive-semidefinite diagonal matrix $D^{1/2} \in \mathbb{R}^{N \times N}$), the noise covariance in the Hessian eigenbasis takes the form $V^T C V = \frac{1}{S} D^{1/2} U^T K U D^{1/2}$, which differs from D/S since $U^T K U \neq \mathbb{1}$.

Large sample number limit ($P \rightarrow \infty$). In this regime, using the Gaussian i.i.d. assumption of the data and the label noise, the matrix C can be evaluated as [21,84]

$$C = \frac{1}{S} \left[H \bar{\mathbf{w}} \bar{\mathbf{w}}^T H + \text{tr}(H \bar{\mathbf{w}} \bar{\mathbf{w}}^T) H + \sigma_\epsilon^2 H \right], \quad P \rightarrow \infty, \quad (2.19)$$

which explicitly reveals the multiplicative character of the gradient noise. Using the fact that $J = 0$ [see Eq. (2.6)] for $P \rightarrow \infty$ (since x^μ and ϵ^μ are uncorrelated), we have $\mathbf{w} \simeq \mathbf{w}^* = \mathbf{u}$ in the stationary state [see Eq. (2.7)]. Setting thus $\bar{\mathbf{w}} = 0$ in Eq. (2.19) renders the additive noise approximation

$$C \simeq \frac{\sigma_\epsilon^2}{S} H, \quad P \rightarrow \infty. \quad (2.20)$$

This direct proportionality between the gradient covariance and the Hessian is often assumed in the literature [11,20,46].

Multiplicative noise process for $\sigma_\epsilon^2 = 0$. Returning to the general case $\mathbf{w} \neq \mathbf{w}^*$, further insight into the nature of the stochastic process associated with SGD can be gained in the case of vanishing label noise $\epsilon = 0$. In this situation,

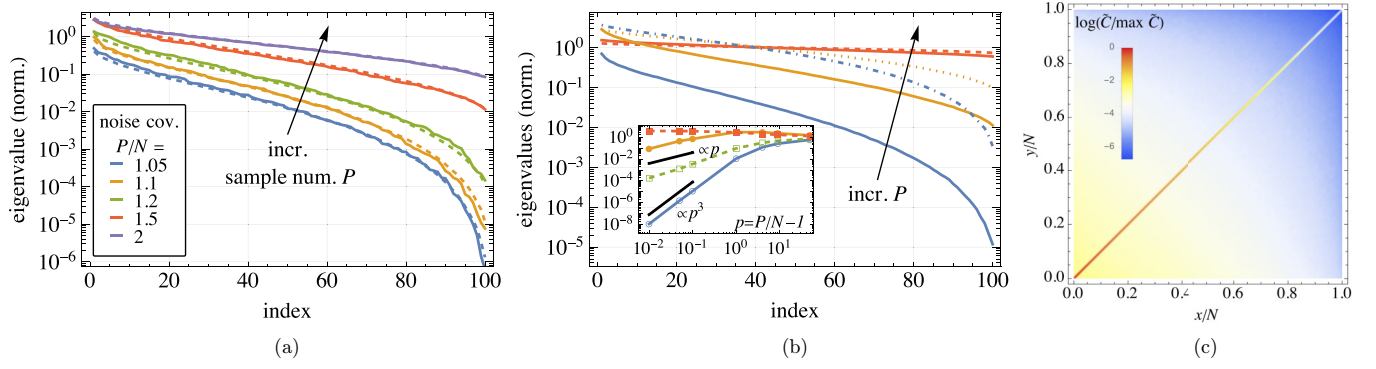


FIG. 3. Gradient noise of SGD for a linear network [see Eq. (2.15)]: anisotropy, dependence on sample number, and deviations from the Hessian approximation. (a) Eigenvalues of the gradient noise covariance matrix C , ordered by their magnitude, for varying number of samples P (increasing from bottom to top) in the underparameterized regime and at input dimension $N = 100$. Numerical results obtained from SGD (single data realization with learning rate $\lambda = 0.1$, minibatch size $S = 10$, label noise $\sigma_\epsilon^2 = 10^{-2}$; solid lines) are compared to the theoretical predictions obtained from Eq. (2.17) (averaged over multiple data realizations). (b) Validity of Hessian approximation of the noise for $P/N \gg 1$: Eigenvalues of C given by Eq. (2.17) (solid lines) compared to those of Eq. (2.20) (broken lines, essentially corresponding to the eigenvalues of the Hessian H), for $P/N = 1.1, 2, 50$ (from bottom to top right) at $N = 100$. The inset shows the maximum and minimum eigenvalues of C from Eq. (2.17) (solid connecting lines) and Eq. (2.20) (dashed connecting lines) for varying P . Eigenvalues in (a) and (b) are normalized by the factor $\sigma_x^2 \sigma_\epsilon^2 / S$, where $\sigma_x^2 = 1/N$ is the variance of the samples [see also Eq. (2.20)]. (c) Visualization of the noise covariance (for $N/P \simeq 0.67$) in the basis of the Hessian H [Eq. (2.4)], i.e., $\tilde{C} = \langle (V^T C V) \rangle_{x,\epsilon}$, where the average is over several data realizations and the matrix V consists of the eigenvectors of H as columns. The color range is scaled logarithmically and normalized to the maximum entry of \tilde{C} . The magnitude of the off-diagonal entries decrease upon increasing P .

Eq. (2.14) becomes

$$\begin{aligned}
 C_{ab} &= \frac{1}{S} \left[\frac{1}{P} \sum_{\mu} \bar{w}_i x_i^{\mu} x_a^{\mu} x_b^{\mu} x_j^{\mu} \bar{w}_j \right. \\
 &\quad \left. - \left(\frac{1}{P} \bar{w}_i \sum_{\mu} x_i^{\mu} x_a^{\mu} \right) \left(\frac{1}{P} \sum_{\mu} x_b^{\mu} x_j^{\mu} \bar{w}_j \right) \right] \\
 &= \frac{1}{S} \text{cov}_x[(\bar{\mathbf{w}} \cdot \mathbf{x})\mathbf{x}]_{ab}. \tag{2.21}
 \end{aligned}$$

Accordingly, Eq. (2.13) can be formulated as

$$\partial_t \bar{\mathbf{w}} = -H \bar{\mathbf{w}} + \sqrt{\frac{\lambda}{S}} \bar{\mathbf{w}}^T Z, \tag{2.22}$$

with a matrix-valued Gaussian white noise Z_{ij} with covariance $\langle \Delta x_i x_j \rangle \equiv x_i x_j - \langle x_i \rangle \langle x_j \rangle$,

$$\begin{aligned}
 \langle Z_{ij}(t) Z_{kl}(t') \rangle &= \langle (\Delta x_i x_j) (\Delta x_k x_l) \rangle \delta(t - t') \\
 &= (\langle x_i x_k \rangle \langle x_j x_l \rangle + \langle x_i x_l \rangle \langle x_j x_k \rangle) \delta(t - t'). \tag{2.23}
 \end{aligned}$$

In the infinite-sample limit ($P \rightarrow \infty$), the i.i.d. assumption (as well as $\mu_x = 0$) for the x_i implies $\langle Z_{ij}(t) Z_{kl}(t) \rangle = (\delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}) \langle x^2 \rangle \delta(t - t')$ or, equivalently, $C_{ab} = S^{-1} (\bar{w}_a \bar{w}_b + \bar{\mathbf{w}}^2 \delta_{ab}) \langle x^2 \rangle^2$. The noise thus generally couples the different weights \bar{w}_i during the evolution of Eq. (2.22), even for whitened data (where $H = \mathbb{1}_N$). In the case of $N = 1$ dimensions (but arbitrary finite P), the noise covariance simplifies to $\langle (\Delta x^2)^2 \rangle = \langle x^4 \rangle - \langle x^2 \rangle^2 = 2 \langle x^2 \rangle^2$ and Eq. (2.22) reduces to the standard geometric Brownian motion process [36]

$$\partial_t \bar{w} = -\langle x^2 \rangle \bar{w} + \sqrt{\frac{2\lambda}{S}} \langle x^2 \rangle \bar{w} \zeta, \quad \langle \zeta(t) \zeta(t') \rangle = \delta(t - t'), \tag{2.24}$$

where ζ is a Gaussian white noise. The mean and variance of \bar{w} for this process are given by $\langle \bar{w}(t) \rangle = \bar{w}(0) e^{-\langle x^2 \rangle t}$ and $\langle \bar{w}(t)^2 \rangle = \bar{w}(0)^2 [e^{2\langle x^2 \rangle (\frac{\lambda}{S} \langle x^2 \rangle - 1)t} - e^{-2\langle x^2 \rangle t}]$, showing that both converge to zero provided $\langle x^2 \rangle < S/\lambda$ (which is fulfilled for typical parameter choices). If $\langle x^2 \rangle = S/\lambda$ and $\bar{w}(0)^2 = \lambda \langle x^2 \rangle / S = 1$, mean and variance become identical to those of an Ornstein-Uhlenbeck process.

3. Relaxational dynamics

In order to study the dynamics induced by the SGD, we project the weights onto the eigenmodes of the data. Assuming the SVD of the design matrix to be $X = P^{1/2} V D^{1/2} U^T$, with $V \in \mathbb{R}^{N \times N}$, $U \in \mathbb{R}^{P \times N}$, and a positive-semidefinite diagonal matrix $D^{1/2} \in \mathbb{R}^{N \times N}$ (the sample number P is introduced in order for the entries of D to coincide with the eigenvalues of H), the Langevin Eq. (2.13), within the additive noise approximation, can be transformed to

$$\dot{\mathbf{z}} = -D(\mathbf{z} - \mathbf{z}^*) + \sqrt{\lambda} \boldsymbol{\eta}, \quad \langle \boldsymbol{\eta} \boldsymbol{\eta}^T \rangle = \frac{P}{S} D^{1/2} U^T K U D^{1/2}, \tag{2.25}$$

where $\mathbf{z} \equiv V^T \mathbf{w} \in \mathbb{R}^N$ is the new dynamical variable, $\mathbf{z}^* \equiv V^T \mathbf{w}^*$ are the transformed solution weights [see Eq. (2.7)], and $\boldsymbol{\eta} = V^T \boldsymbol{\xi}$ is a new noise (Gaussian distributed with zero mean). Analogously, the loss [Eq. (2.9)] can be expressed as

$$L = \frac{1}{2} (\mathbf{z} - \mathbf{z}^*)^T D (\mathbf{z} - \mathbf{z}^*) + L^* = \frac{1}{2} \sum_k^N (z_k - z_k^*)^2 \omega_k + L^*, \tag{2.26}$$

with $\omega_k \equiv D_{kk}$ and $L^* \equiv L(\mathbf{z}^*) = -\frac{1}{2} J^T H^+ J + \frac{1}{2P} \mathbf{e} \mathbf{e}^T = \frac{1}{2P} \mathbf{e} (\mathbb{1}_P - X^T (X X^T)^{-1} X) \mathbf{e}^T$ denoting the loss of the fully trained model. Equation 2.25 describes a multivariate Ornstein-Uhlenbeck process driven by correlated Gaussian




FIG. 4. Exemplary relaxation behavior of the loss $L(\mathbf{z})$ (dash-dotted curve) and the weight modes $z_k(t) = V_{kj}^T w_j$ [for $k = 1, 5, 25$, see Eq. (2.25)] towards the solution z_k^* (single run). Modes obtained from SGD experiments (solid lines) are compared to theoretical predictions $\langle z_k(t) \rangle$ [Eq. (2.27), averaged over the noise; dashed], for a slightly underparameterized linear net with $P/N \simeq 1.04$. The inset shows the loss and the mode $|z_0(t)|$ in double-logarithmic representation. The gray line in the inset represents the theoretical prediction L^* for the final loss. Parameters used for SGD experiments: learning rate $\lambda = 0.1$, label noise $\sigma_\epsilon^2 = 10^{-2}$, input dimension $N = 50$, sample number $P = 52$, batch size $S = 10$.

noise. Owing to the presence of label noise, encoded in the matrix K [see Eq. (2.18)], the covariance of the noise $\boldsymbol{\eta}$ is not diagonal but instead induces correlations between the weight modes z_i . The solution of Eq. (2.25) is given by

$$z_k(t) = z_k^* + e^{-\omega_k(t-t_0)}[z_k(t_0) - z_k^*] + \int_{t_0}^t ds e^{\omega_k(s-t)} \eta_k(s), \quad (2.27)$$

where, henceforth, we assume the relaxation rates to be ordered as $\omega_0 \leq \omega_1 \leq \dots$. For times larger than the smallest relaxation time ($t \gg 1/\omega_0$), the influence of the initial value term $\mathbf{z}(t_0) - \mathbf{z}^*$ has decayed and \mathbf{z} fluctuates around the equilibrium solution \mathbf{z}^* . Within the additive noise approximation considered here, the gradient descent (GD) solution follows by averaging $z_k(t)$ over the noise using $\langle \eta_k \rangle = 0$. Near the interpolation threshold $N = P$, the eigenvalues of the Hessian H become arbitrarily small, i.e., $\omega_0 \rightarrow 0^+$, implying a diverging relation time [7,85,86]. Specifically, for Gaussian data we consider here [see Sec. II A], the Marchenko-Pastur law of random matrix theory implies $\omega_0 = \sigma_x^2(\sqrt{P/N} - 1)^2$ asymptotically for $P, N \rightarrow \infty$ with $P/N = \text{const}$ [7,87,88].

The typical relaxation behavior of the modes $z_k(t)$ under SGD near the interpolation threshold is exemplified in Fig. 4. We find that the relaxation behavior is overall well described by Eq. (2.27), with minor deviations from the purely exponential relaxation for $P \simeq N$. An interesting observation (see inset to Fig. 4) is that the loss settles to the constant L^* already for $t \ll \omega_0$, although the slowest modes z_k have not yet all reached their final value. This is a consequence of the presence of the factor ω_k in the loss [Eq. (2.26)], due to which the contribution of modes for small ω_k can be hidden below the noise floor L^* , especially for large label noise σ_ϵ^2 .

C. Steady-state properties

In the following, we analyze the stationary covariance matrix of the weights induced by the SGF of a single-layer linear network. We focus, in particular, on the spectral properties of the weight fluctuations and their dependence on the sample number and the label noise. We then use these results to show that the IVFR is generally not fulfilled for a single-layer linear net.

1. Steady-state solution

The Fokker-Planck equation for the distribution $p(t, \mathbf{w})$ induced by Eq. (2.13) takes the form [74,75]

$$\begin{aligned} \partial_t p(t, \mathbf{w}) &= - \sum_i \frac{\partial}{\partial w_i} \mathcal{J}_i, \\ \mathcal{J}_i &= - \left[\frac{\partial}{\partial w_i} L(\mathbf{w}) \right] p(t, \mathbf{w}) \\ &\quad - \sum_j \frac{\partial}{\partial w_j} \left[\frac{\lambda}{2} C_{ij}(\mathbf{w}) p(t, \mathbf{w}) \right], \end{aligned} \quad (2.28)$$

with $L(\mathbf{w}) = H(\mathbf{w} - \mathbf{w}^*)$. We focus in the following on the w -independent noise defined by C in Eq. (2.17), which, as discussed above, is expected to apply in the stationary state. The steady-state covariance matrix of the weights $M \equiv \langle \delta \mathbf{w} \delta \mathbf{w}^T \rangle = M^T$ (with $\delta \mathbf{w} \equiv \mathbf{w} - \mathbf{w}^*$) is determined by the Lyapunov equation [10,20,74,89]

$$HM + MH^T = \lambda C. \quad (2.29)$$

The most general solution of this equation fulfills $\nabla \cdot \mathcal{J} = 0$, but not necessarily $\mathcal{J} = 0$, i.e., while the probability density remains stationary [see Eq. (2.28)], the current \mathcal{J} can still be nonvanishing (broken detailed balance). This situation is indeed realized for SGD and one obtains [12,20]

$$\begin{aligned} M &\equiv \langle \delta \mathbf{w} \delta \mathbf{w}^T \rangle = \frac{\lambda}{2} H^{-1} (C + Q), \\ Q &= V \tilde{Q} V^T, \quad \tilde{Q}_{ij} = \frac{\omega_i - \omega_j}{\omega_i + \omega_j} (V^T C V)_{ij}, \quad \text{where} \\ V \text{diag}(\omega_i) V^T &= H, \end{aligned} \quad (2.30)$$

i.e., ω_i are the eigenvalues of H and the columns of V contain the eigenvectors. The stationary probability distribution and current follow as

$$\begin{aligned} p_{\text{ss}}(\mathbf{w}) &= \frac{1}{Z} \exp \left[-\frac{1}{2} (\mathbf{w} - \mathbf{w}^*)^T M^{-1} (\mathbf{w} - \mathbf{w}^*) \right], \\ \mathcal{J}_{\text{ss}} &= -Q M^{-1} (\mathbf{w} - \mathbf{w}^*) p_{\text{ss}}(\mathbf{w}), \end{aligned} \quad (2.32)$$

with a normalization factor Z . The inverse covariance matrix M^{-1} can be interpreted as an effective potential $U(\mathbf{w}) \equiv (1/2) \delta \mathbf{w}^T M^{-1} \delta \mathbf{w}$, which, due to the presence of the kinetic term $(C + Q)^{-1}$, differs in general from the loss [12,17]. The matrix Q is antisymmetric ($Q = -Q^T$) and determines a stationary current \mathcal{J}_{ss} , which flows in planes of constant force, i.e., $\mathcal{J}_{\text{ss}} \cdot (\nabla U) = 0$ [20,72]. In order to aid comparison between various conventions used, e.g., in numerical codes, we remark that, if the loss [Eq. (2.10)] is scaled by a factor α , i.e., $\ell^\mu \rightarrow \alpha \ell^\mu$, the noise covariance [Eq. (2.14)] scales




FIG. 5. Statistics of weight fluctuations for a linear network in the stationary state. (a) Behavior of the weight variance as a function of input dimension for $P = 200$ samples. Simulation results (connected symbols) obtained for a single run of SGD with learning rate $\lambda = 0.1$, label noise $\sigma_\epsilon^2 = 0.25$, and minibatch size $S = 10$ are compared to the theoretical predictions of Eqs. (2.34) and (2.30) for $\langle (\delta w_i(t)^2)_i \rangle = N^{-1} \sum_i M_{ii}$ (dotted and dash-dotted lines), which represents the temporal variance of the weights in the stationary state [where $\mathbf{w}(t \rightarrow \infty) = \mathbf{w}^*$, see Eq. (2.7)] averaged over all weights. The quantity $\langle (\delta w_i(t \rightarrow \infty)^2)_i \rangle$ (filled circles and dashed line) represents the variance across all weights at an arbitrary but fixed time in the stationary state. The train and test losses are also shown for comparison. (b) Eigenvalues of the covariance matrix of the weights $M = \langle \delta \mathbf{w} \delta \mathbf{w}^T \rangle$, ordered by their magnitude and normalized by $\lambda \sigma_\epsilon^2 / (2S)$ [see Eq. (2.34)], for varying number of samples ($P/N = 1.05, 1.1, 1.2, 1.5, 2.0, 4.0, 10$, from bottom to top) in the underparameterized regime (for $\lambda = 0.1$, $S = 10$, and input dimension $N = 100$). Numerical results obtained from SGD (solid lines) are compared to the theoretical predictions obtained from Eq. (2.30) (dashed lines; averaged over the multiple data realizations). The straight (dashed-dotted) black line represents the prediction of Eq. (2.34), while the dotted lines represent the theoretical eigenvalues under the detailed balance assumption [Eq. (2.33)]. (c) Smallest and largest eigenvalues of M as given by the exact result Eq. (2.30) (solid lines) and the detailed balance approximation Eq. (2.33) (dotted lines), normalized by $\lambda \sigma_\epsilon^2 / (2S)$ and for varying sample numbers P . The dashed lines represent the smallest and largest absolute values of the eigenvalues of Q [Eq. (2.30)] in the same normalization. We note that the impact of Q diminishes in the oversampled regime.

as $C \rightarrow \alpha^2 C$ (analogously $Q \rightarrow \alpha^2 Q$) and the Hessian $H \rightarrow \alpha H$, which implies $M \rightarrow \alpha M$.

A solution fulfilling detailed balance can be obtained from straightforward integration of the condition $\mathcal{J} = 0$ in Eq. (2.28), which yields $\frac{1}{2} \lambda C_{ij} \cdot \partial \ln p / \partial w_j = -\partial L(\mathbf{w}) / \partial w_j$ and eventually

$$M = \langle \delta \mathbf{w} \delta \mathbf{w}^T \rangle = \frac{\lambda}{2} H^{-1} C, \quad Q = 0. \quad (\text{detailed balance}) \quad (2.33)$$

While, generally, detailed balance is broken for SGD, we find that Eq. (2.33) constitutes a reasonable approximation sufficiently deep in the underparameterized regime (specifically, for $P/N \gtrsim 2$, see discussion of Fig. 5 below). In fact, asymptotically for $P/N \rightarrow \infty$, one has $C \propto H$ [see Eq. (2.20)], which implies $V^T C V \propto \delta_{ij}$ and thus $Q = 0$ in Eq. (2.30). In this limit detailed balance is thus exactly fulfilled and one obtains isotropic weight fluctuations [20,21],

$$M = \frac{\lambda \sigma_\epsilon^2}{2S} \mathbb{1}, \quad P/N \rightarrow \infty. \quad (2.34)$$

In Fig. 5, the behavior of the weight fluctuations of a fully trained linear network is illustrated. As shown in panel (a), the steady-state variance $\langle (\delta w_i(t)^2)_i \rangle$ obtained from SGD experiments (by computing the temporal variance of each weight w_i and the averaging over all weights $i = 1, \dots, N$) agrees with the theoretical prediction given by the diagonal entries of M in Eq. (2.30). Notably, the variance is nonvanishing only for $N < P$, since an overparameterized network can fit the data in Eq. (2.1) exactly [see also the comment after Eq. (2.15)]. In the limit $N/P \rightarrow 0$, the variance approaches the isotropic result in Eq. (2.34) (dotted-horizontal line). For comparison, the train and test loss are included in the plot, where the

latter shows a “double descent” peak at the interpolation threshold ($P = N$), characteristic for learning data with noisy labels [7,8,59,90–92]. Beside the temporal variance, we show also the spatial variance of the weights $\langle (\delta w_i(t \rightarrow \infty)^2)_i \rangle = \langle (\delta w_i^*)^2 \rangle$, which peaks at $P = N$ [83]. We recall that a stationary state is only reached for nonvanishing label noise. By contrast, for $\sigma_\epsilon = 0$, the network can fit the teacher weights exactly such that the loss exponentially approaches zero (up to machine precision).

Figure 5(b) illustrates eigenvalues of the weight covariance matrix $M = \langle \delta \mathbf{w} \delta \mathbf{w}^T \rangle$, ordered by their magnitude, for various sample numbers P in the underparameterized regime. The numerical results from SGD are well captured by Eq. (2.30) (dashed lines), while significant deviations from Eq. (2.34) are noticeable, indicating broken detailed balance, except for large P .

In Fig. 5(c), the smallest and largest eigenvalues of the weight covariance matrix M [Eqs. (2.30) and (2.33), solid and dotted lines, respectively] are shown as functions of the reduced sample number $P/N - 1$. We infer that the detailed balance approximation for M given in Eq. (2.33) is applicable for $P \gtrsim 2N$. This is also corroborated by the behavior of the eigenvalues of Q [Eq. (2.31), dashed lines], which peak around $P \approx N$, and decay for larger sample numbers. (Note that the eigenvalues of the antisymmetric matrix Q are purely imaginary and hence we plot their absolute value.) From the gap between the minimum and maximum eigenvalues we finally conclude that the weight fluctuations become essentially isotropic [see Eq. (2.34)] for $P \gtrsim 100N$, consistent with the behavior of the noise eigenvalues in Fig. 3(b).

Before proceeding, we summarize some important steady-state properties for linear regression in the following Table I.

TABLE I. Properties of a linear network in the stationary (late-time) regime, trained on a teacher-student task [see Eq. (2.1)]. Note that these properties are independent of the sampling method (i.e., sampling with or without replacement). The variance of the target (label) noise is denoted by σ_ϵ^2 .

| | Student weights [Eq. (2.7)] | Final loss [Eq. (2.9)] | Temporal weight fluctuations |
|--------------------|--|--|---|
| Underparameterized | approach teacher weights as $\epsilon \rightarrow 0$ | nonzero in general; vanishes $\propto \sigma_\epsilon^2$ | nonzero in general; vanish $\propto \sigma_\epsilon^2$ [Eq. (2.34)] |
| Overparameterized | differ from teacher weights | vanishes (also if $\sigma_\epsilon^2 > 0$) | vanish [Eqs. (2.16) and (2.20)] |

2. Loss perturbation

We now investigate whether the IVFR [16] applies in the context of a single-layer linear network. Perturbing the linear regression loss in Eq. (2.9) around the solution \mathbf{w}^* [Eq. (2.7)] renders

$$\delta L(\theta \mathbf{v}) = L(\mathbf{w}^* + \theta \mathbf{v}) - L(\mathbf{w}^*) = \frac{\theta^2}{2} \mathbf{v}^T H \mathbf{v}, \quad (2.35)$$

where \mathbf{v} is assumed to be an eigenvector of the weight covariance matrix M [Eq. (2.30)]. Owing to the different structural forms of M and H , the eigenvectors of these matrices can be considered as essentially unrelated here and we thus do not expect $\delta L(\theta \mathbf{v})$ to depend in any significant way on the covariance eigenvalues, even if the latter strongly vary in magnitude. This is confirmed by SGD simulations for Gaussian independent and Gaussian correlated data (see Fig. 6), as well as for purely random labels y^μ .

Analytic insight can be obtained in the large (infinite) sample number regime, where Eq. (2.34) implies the eigenvectors $\mathbf{v} \in \{\mathbf{e}_i\}_{i=1, \dots, N}$ with a common eigenvalue $\gamma = \lambda \sigma^2 / S$. This yields $\delta L(\theta \mathbf{e}_i) = \frac{1}{2} \theta^2 H_{ii} = \frac{1}{2} \theta^2 \langle x_i^2 \rangle$, i.e., the loss perturbation reveals the variance of the i th feature. Accordingly, unless one engineers a particular variance dependence on the feature index, the loss perturbation behaves in an essentially random and isotropic way. We thus conclude that the IVFR is not fulfilled for a typical single-layer linear network.

III. TWO-LAYER LINEAR NETWORK

We now transfer the Langevin formalism presented in the preceding section to a two-layer linear network, which will

enable us to analyze the fluctuations of the weights in the individual layers. Since the dynamics of the weight matrices is now nonlinear, we develop a linearization around the deterministic solution. We will finally use these results in order to derive the IVFR. In order to facilitate analytical calculations, we will focus, where necessary, on the regime of large batch sizes, small learning rate, and large sample numbers.

A. Model

We consider training data consisting of P samples $\mathbf{x}^\mu \in \mathbb{R}^{N_i}$, $\mathbf{y}^\mu \in \mathbb{R}^{N_o}$ ($\mu = 1, \dots, P$), with i.i.d. Gaussian $x_i^\mu \sim \mathcal{N}(0, \sigma_x^2)$ and $y_i^\mu \sim \mathcal{N}(0, \sigma_y^2)$, which, for simplicity, have zero mean. Within the teacher-student framework used in Eq. (2.1), the targets (“labels”) \mathbf{y}^μ used here can be interpreted as pure noise, while the teacher weights identically vanish. This setup (see Fig. 7) is sufficient to investigate the stationary fluctuating weight dynamics, which is driven by label noise [cf. Eq. (2.17)]. The data is conveniently organized in terms of the design matrices $X \in \mathbb{R}^{N_i \times P}$, $Y \in \mathbb{R}^{N_o \times P}$. The (full-batch) loss of a linear two-layer network is given by

$$\begin{aligned} L(W_1, W_2) &= \frac{1}{2P} \sum_{\mu=1}^P \|\mathbf{y}^\mu - W_2 W_1 \mathbf{x}^\mu\|^2 \\ &= \frac{1}{2P} \text{tr}[(Y - W_2 W_1 X)(Y - W_2 W_1 X)^T] \\ &= \frac{1}{2} \text{tr}(\Sigma_0^{xx} W^T W - W^T \Sigma_0^{yx} - \Sigma_0^{yx, T} W + Y Y^T / P), \end{aligned} \quad (3.1)$$

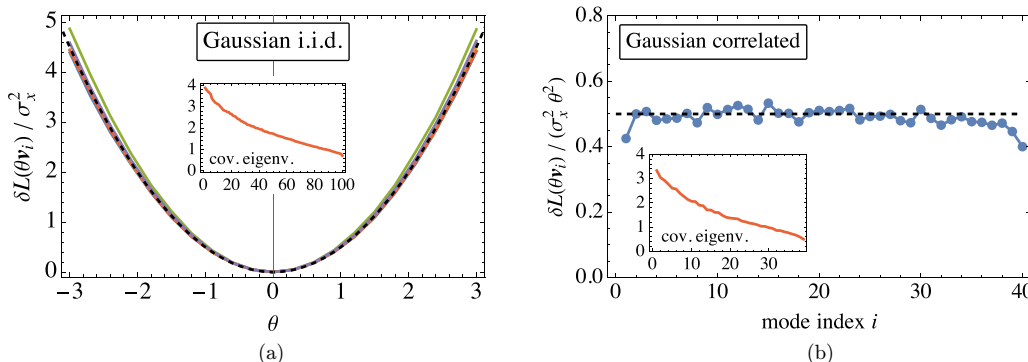


FIG. 6. Perturbing the linear regression loss [Eq. (2.9)] along eigenvectors of the empirical weight covariance matrix $M = \langle \delta \mathbf{w} \delta \mathbf{w}^T \rangle$ obtained from SGD (learning rate $\lambda = 0.1$, minibatch size $S = 10$) of a linear network. The insets show the eigenvalues of the weight covariance matrix (ordered by their magnitude vs mode index i). (a) Perturbed loss curves $\delta L(\theta \mathbf{v}_i)$ for various eigenvectors \mathbf{v}_i for $N = 100$ dimensional input and $P = 2000$ Gaussian i.i.d. samples (single run). The simulation results superimpose on the theoretical prediction $\delta L = \frac{1}{2} \theta^2 \sigma_x^2$ (dashed curve) for large sample numbers. (b) Perturbed loss $\delta L(\theta \mathbf{v}_i) / \theta^2$ at fixed (arbitrary) θ as a function of eigenmode index i for $P = 200$ correlated samples in $N = 40$ dimensions obtained from a multivariate Gaussian defined by a random covariance matrix Σ of Wishart type [i.e., $\Sigma \simeq M M^T$, with $M_{ij} \sim \mathcal{N}(0, 1)$]. The black-dashed line represents the theoretical prediction for i.i.d. data, while the connected symbols are obtained by averaging over several experimental runs.




FIG. 7. Two-layer linear neural network model with input \mathbf{x} , output $\hat{\mathbf{y}}$, and weight matrices W_1 , W_2 . The targets are given as i.i.d. noise \mathbf{y} (see text).

where $W_1 \in \mathbb{R}^{N_h \times N_i}$, $W_2 \in \mathbb{R}^{N_o \times N_h}$ are the weight matrices, $w \equiv W_2 W_1 \in \mathbb{R}^{N_o \times N_i}$ denotes the product matrix, and

$$\Sigma_0^{xx} = \frac{1}{P} \mathbf{X} \mathbf{X}^T, \quad \Sigma_0^{yx} = \frac{1}{P} \mathbf{Y} \mathbf{X}^T \quad (3.2)$$

are the full-batch input-input and output-input covariance matrices. The minimum of the loss is determined by the linear regression solution

$$w^* \equiv (W_2 W_1)^* = \Sigma_0^{yx} (\Sigma_0^{xx})^+, \quad (3.3)$$

which effectively fixes only $N_{\text{eff}} = N_o N_i$ degrees of freedom out of the total number of parameters $N_{\text{tot}} = N_h(N_i + N_o)$. As was the case for a single-layer net [see Eq. (2.16)], a deep linear network of the form considered here has nonvanishing weight fluctuations in the stationary state provided $N_i < P$, irrespective of the output dimension N_o . We will thus restrict ourselves to this regime in the following, in which case the pseudoinverse in Eq. (3.3) accordingly becomes the usual inverse. In certain cases, where indicated, we will also assume $N_i \geq N_h \geq N_o$.

Introducing the minibatch covariance matrices as

$$\Sigma^{xx} = \frac{1}{S} \sum_{\mu \in \mathcal{B}} \mathbf{x}^\mu \mathbf{x}^{\mu,T}, \quad \Sigma^{yx} = \frac{1}{S} \sum_{\mu \in \mathcal{B}} \mathbf{y}^\mu \mathbf{x}^{\mu,T}, \quad (3.4)$$

where \mathcal{B} denotes the set of sample indices in the current minibatch (of size $S = |\mathcal{B}|$), the SGD equations take the form

$$\frac{1}{\lambda} [W_1(\tau + 1) - W_1(\tau)] = - \frac{\partial L(w, \tau)}{\partial W_1} = W_2^T (\Sigma^{yx} - W_2 W_1 \Sigma^{xx}), \quad (3.5a)$$

$$\frac{1}{\lambda} [W_2(\tau + 1) - W_2(\tau)] = - \frac{\partial L(w, \tau)}{\partial W_2} = (\Sigma^{yx} - W_2 W_1 \Sigma^{xx}) W_1^T, \quad (3.5b)$$

with the learning rate λ .

The full-batch dynamics of Eq. (3.5) in the limit of small λ (*gradient flow*) has been analyzed previously [4,5,9,93]. We will derive the continuum limit of Eq. (3.5) in the next section, but already note here that, in the case of gradient flow with an explicitly time-dependent loss L , one can show that the following quantity is conserved for a linear network ($t = \tau/\lambda$),

$$W_2^T(t) W_2(t) - W_1(t) W_1^T(t) = \text{const.} \quad (3.6)$$

The proof of this relation follows the same steps as in [65], where it was demonstrated for linear networks with losses that do not explicitly depend on time. When the weights are initialized such that the constant matrix on the r.h.s. of Eq. (3.6) vanishes, the weights are conventionally called *zero balanced*.

Notably, for whitened inputs $\Sigma_0^{xx} \propto \mathbb{1}$ and zero-balanced weights, an explicit solution of this gradient flow can be constructed [9,94]: Denoting the compact singular value decomposition of Σ_0^{yx} by

$$\Sigma_0^{yx} = \tilde{U} \tilde{S} \tilde{V}^T, \quad \text{with } \tilde{U} \in \mathbb{R}^{N_o \times N_m}, \tilde{S} \in \mathbb{R}^{N_m \times N_m}, \tilde{V} \in \mathbb{R}^{N_i \times N_m}, \quad (3.7)$$

where \tilde{S} is a diagonal matrix with N_m nonzero eigenvalues on its diagonal, the late-time solution can be expressed as

$$W_1^{0,T} W_1^0 = \tilde{V} \tilde{S} \tilde{V}^T, \quad W_2^0 W_2^{0,T} = \tilde{U} \tilde{S} \tilde{U}^T. \quad (\text{zero-balanced}) \quad (3.8)$$

However, explicit expressions for the individual full-batch W_i^0 are not available. Unless otherwise mentioned, we will in the following not assume any special initial conditions.

B. Continuum description of SGD

We now derive the continuum limit of the SGD dynamics in each layer, which renders two coupled nonlinear SDEs. A crucial step is to determine the covariances of the effective noise terms induced by the minibatching. In order to analyze the nonlinear SDEs in the stationary state, we focus on the large batch-regime, where fluctuations around the deterministic gradient flow solution are small.

1. Noise sources and stochastic gradient flow

Since the only source of randomness enters via the minibatching of $\Sigma^{xx,yx}$ [95], we have

$$\Sigma^{xx} = \Sigma_0^{xx} + \delta \Sigma^{xx}, \quad \Sigma^{yx} = \Sigma_0^{yx} + \delta \Sigma^{yx}, \quad (3.9)$$

where we note that $\Sigma_0^{xx} = \langle \Sigma^{xx} \rangle_{\mathcal{B}}$ and $\Sigma_0^{yx} = \langle \Sigma^{yx} \rangle_{\mathcal{B}}$. Accordingly, the noise terms are defined as

$$\delta \Sigma^{xx} = \Sigma^{xx} - \Sigma_0^{xx}, \quad \delta \Sigma^{yx} = \Sigma^{yx} - \Sigma_0^{yx}, \quad (3.10)$$

and their covariances can be expressed (using Appendix C) by a four-dimensional tensor,

$$\begin{aligned} \text{cov}(\delta \Sigma^{xx})_{ij,kl} &\equiv \langle \delta \Sigma_{ij}^{xx} \delta \Sigma_{kl}^{xx} \rangle_{\mathcal{B}} \\ &= \left\langle \left(\frac{1}{S} \sum_{\mu} x_i^{\mu} x_j^{\mu} - \Sigma_{0,ij}^{xx} \right) \left(\frac{1}{S} \sum_{\nu} x_k^{\nu} x_l^{\nu} - \Sigma_{0,kl}^{xx} \right) \right\rangle_{\mathcal{B}} \\ &= \frac{1}{S} \left[\frac{1}{P} \sum_{\mu=1}^P x_i^{\mu} x_j^{\mu} x_k^{\mu} x_l^{\mu} - \Sigma_{0,ij}^{xx} \Sigma_{0,kl}^{xx} \right]. \end{aligned} \quad (3.11)$$

In general, a four-dimensional tensor \mathcal{M}_{ijkl} can be reshaped into a conventional two-dimensional matrix m by defining $m_{ab} = [\mathcal{M}_{ij,kl}]_{a=(j-1)N_i+i, b=(l-1)N_k+k}$ where N_i and N_k are the sizes of dimensions indexed by i and k , respectively. This definition is consistent with the usual vectorization

$m = \text{vec}(\mathcal{A})\text{vec}(\mathcal{B})^T$ in the case where $\mathcal{M}_{ijkl} = \mathcal{A}_{ij}\mathcal{B}_{kl}$, which applies to the covariance tensor. The vectorization operation $\text{vec}(\mathcal{A})$ is defined by stacking the columns of the matrix \mathcal{A} on top of each other to form a column vector.

In order to arrive at tractable expressions, we evaluate the noise covariance in the large sample number limit ($P \rightarrow \infty$), which is analogous to the Hessian approximation given by Eq. (2.20). Practically, the Hessian approximation for the noise applies already for $P \gtrsim 100N$, as discussed in Sec. II C 1. In this limit, the first term in the square brackets in Eq. (3.11) converges to a batch average, such that Wick's rule for products of Gaussian variables can be applied: $\frac{1}{P} \sum_{\mu=1}^P x_i^\mu x_j^\mu x_k^\mu x_l^\mu \simeq \langle x_i x_j x_k x_l \rangle_{\mathcal{B}} = \langle x_i x_j \rangle_{\mathcal{B}} \langle x_k x_l \rangle_{\mathcal{B}} + \langle x_i x_k \rangle_{\mathcal{B}} \langle x_j x_l \rangle_{\mathcal{B}} + \langle x_i x_l \rangle_{\mathcal{B}} \langle x_j x_k \rangle_{\mathcal{B}}$. Hence, altogether one obtains

$$\text{cov}(\delta \Sigma^{xx})_{ij,kl} = \frac{1}{S} [\Sigma_{0,ik}^{xx} \Sigma_{0,jl}^{xx} + \Sigma_{0,il}^{xx} \Sigma_{0,jk}^{xx}], \quad (3.12a)$$

$$\text{cov}(\delta \Sigma^{yx})_{ij,kl} = \frac{1}{S} [\langle y_i^2 \rangle \delta_{ik} \Sigma_{0,jl}^{yx} + \Sigma_{0,il}^{yx} \Sigma_{0,kj}^{yx}], \quad (3.12b)$$

$$\text{cov}(\delta \Sigma_{ij}^{xx}, \delta \Sigma_{kl}^{yx}) = \frac{1}{S} [\Sigma_{0,ki}^{yx} \Sigma_{0,jl}^{xx} + \Sigma_{0,il}^{xx} \Sigma_{0,kj}^{yx}], \quad (3.12c)$$

where we used the factor that the y_i are uncorrelated zero mean random variables. In the case of whitened inputs, i.e.,

$$\Sigma_0^{xx} = \sigma_x^2 \mathbb{1}, \quad (3.13)$$

the above covariances simplify to

$$\begin{aligned} \text{cov}(\delta \Sigma^{xx})_{ij,kl} &= \frac{1}{S} [\sigma_x^4 \delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}], \\ \text{cov}(\delta \Sigma^{yx})_{ij,kl} &= \frac{1}{S} [\langle y_i^2 \rangle \sigma_x^2 \delta_{ik} \delta_{jl} + \Sigma_{0,il}^{yx} \Sigma_{0,jk}^{yx}]. \end{aligned} \quad (3.14)$$

Upon inserting these expansions of the data covariance matrices into Eq. (3.5) and proceeding to the continuum limit analogously to the single-layer case [Eq. (2.13)], we obtain the following *stochastic gradient flow* (SGF) equations:

$$\partial_t W_1 = W_2^T (\Sigma_0^{yx} - W_2 W_1 \Sigma_0^{xx}) + \mathcal{R}_1, \quad (3.15a)$$

$$\partial_t W_2 = (\Sigma_0^{yx} - W_2 W_1 \Sigma_0^{xx}) W_1^T + \mathcal{R}_2, \quad (3.15b)$$

with the multiplicative noise matrices $\mathcal{R}_1 \in \mathbb{R}^{N_h \times N_i}$, $\mathcal{R}_2 \in \mathbb{R}^{N_o \times N_h}$,

$$\mathcal{R}_1 = \sqrt{\lambda} W_2^T (\delta \Sigma^{yx} - W_2 W_1 \delta \Sigma^{xx}), \quad (3.16a)$$

$$\mathcal{R}_2 = \sqrt{\lambda} (\delta \Sigma^{yx} - W_2 W_1 \delta \Sigma^{xx}) W_1^T. \quad (3.16b)$$

2. Perturbation expansion

In order to analyze the stationary fluctuation dynamics, we introduce a suitable linearization into Eq. (3.15). Accordingly, we decompose the weight matrices as $W_{1,2} = W_{1,2}^0 + \delta W_{1,2}$, where $W_{1,2}^0$ denotes the solution of the full-batch GD and $\delta W_{1,2}$ describe small fluctuations around this solution. Upon expanding Eq. (3.15) up to linear order in the fluctuations of W and Σ , and collecting terms of the same order, we obtain the deterministic gradient flow equations

$$\partial_t W_1^0 = W_2^{0,T} (\Sigma_0^{yx} - W_2^0 W_1^0 \Sigma_0^{xx}), \quad (3.17a)$$

$$\partial_t W_2^0 = (\Sigma_0^{yx} - W_2^0 W_1^0 \Sigma_0^{xx}) W_1^{0,T}, \quad (3.17b)$$

and the dynamics of the first-order perturbations,

$$\begin{aligned} \partial_t \delta W_1 &= (\delta W_2^T \Sigma_0^{yx} - \delta W_2^T w^* \Sigma_0^{xx} - W_2^{0,T} \delta W_2 W_1^0 \Sigma_0^{xx} \\ &\quad - W_2^{0,T} W_2^0 \delta W_1 \Sigma_0^{xx}) + \mathcal{R}_1, \end{aligned} \quad (3.18a)$$

$$\begin{aligned} \partial_t \delta W_2 &= (\Sigma_0^{yx} \delta W_1^T - \delta W_2 W_1^0 \Sigma_0^{xx} W_1^{0,T} - W_2^0 \delta W_1 \Sigma_0^{xx} W_1^{0,T} \\ &\quad - w^* \Sigma_0^{xx} \delta W_1^T) + \mathcal{R}_2. \end{aligned} \quad (3.18b)$$

in which the noises [Eq. (3.16)] now become additive (and are assumed to be quantities of the same order as $\delta W_{1,2}$),

$$\begin{aligned} \mathcal{R}_1 &= \sqrt{\lambda} W_2^{0,T} (\delta \Sigma^{yx} - W_2^0 W_1^0 \delta \Sigma^{xx}) \\ &= \sqrt{\lambda} W_2^{0,T} (\delta \Sigma^{yx} - w^* \delta \Sigma^{xx}), \end{aligned} \quad (3.19a)$$

$$\begin{aligned} \mathcal{R}_2 &= \sqrt{\lambda} (\delta \Sigma^{yx} - W_2^0 W_1^0 \delta \Sigma^{xx}) W_1^{0,T} \\ &= \sqrt{\lambda} (\delta \Sigma^{yx} - w^* \delta \Sigma^{xx}) W_1^{0,T}. \end{aligned} \quad (3.19b)$$

Crucially, we focus henceforth on a (quasi)stationary state of training, where the $W_{1,2}^0$ are essentially time independent or at least change on a significantly longer time scale than the $\delta W_{1,2}$, allowing us to analyze Eqs. (3.18) and (3.19) with $W_{1,2}^0 \simeq \text{const}$. Since explicit expressions for W_i^0 in terms of the data and initial conditions are not available, we will consider the $W_{1,2}^0$ as given outcomes of the gradient flow [Eq. (3.17)], obtained, e.g., from a numerical solution [96].

For whitened inputs [see Eq. (3.13)], we have $w^* = \Sigma_0^{yx} / \sigma_x^2$ [see Eq. (2.7)] as well as

$$\begin{aligned} (\Sigma_0^{yx} \Sigma_0^{yx,T})_{ij} &= \frac{1}{P^2} \sum_{\mu, \nu} \sum_k y_i^\mu x_k^\mu x_k^\nu y_j^\nu \simeq \frac{1}{P^2} \sum_{\mu} (y^\mu)^2 \sigma_x^2 \delta_{ij} \\ &= \frac{1}{P} \sigma_x^2 \langle y^2 \rangle \delta_{ij}, \end{aligned} \quad (3.20)$$

where we used $\mathbf{x}^\mu \cdot \mathbf{x}^\nu \simeq \sigma_x^2 \delta^{\mu\nu}$ and the fact that y^μ are i.i.d. variates. With these simplifications, the noise covariance takes the following form (see Appendix E):

$$\begin{aligned} \text{cov}(\mathcal{R}_1)_{ij,kl} &= \sigma_{\mathcal{R}}^2 \delta_{jl} (1 - 1/P) (W_2^{0,T} W_2^0)_{ik}, \\ &\quad j, l = 1, \dots, N_i, \end{aligned} \quad (3.21a)$$

$$\begin{aligned} \text{cov}(\mathcal{R}_2)_{ij,kl} &= \sigma_{\mathcal{R}}^2 \delta_{ik} (1 - 1/P) (W_1^0 W_1^{0,T})_{jl}, \\ &\quad i, k = 1, \dots, N_o, \end{aligned} \quad (3.21b)$$

$$\begin{aligned} \text{cov}(\mathcal{R}_1, \mathcal{R}_2)_{ij,kl} &= \sigma_{\mathcal{R}}^2 (1 - 1/P) W_2^{0,T} W_{2,ik}^{0,T} W_{1,jl}^{0,T} \quad \text{with} \\ \sigma_{\mathcal{R}}^2 &\equiv \frac{1}{S} \lambda \sigma_x^2 \langle y^2 \rangle, \end{aligned} \quad (3.21c)$$

while Eq. (3.18) becomes

$$\begin{aligned} \partial_t \delta W_1 &= -\sigma_x^2 [W_2^{0,T} W_2^0 \delta W_1 + W_2^{0,T} \delta W_2 W_1^0] + \mathcal{R}_1 \\ &= -\frac{\partial \tilde{\mathcal{L}}}{\partial \delta W_1} + \mathcal{R}_1, \end{aligned} \quad (3.22a)$$

$$\begin{aligned} \partial_t \delta W_2 &= -\sigma_x^2 [\delta W_2 W_1^0 W_1^{0,T} + W_2^0 \delta W_1 W_1^{0,T}] + \mathcal{R}_2 \\ &= -\frac{\partial \tilde{\mathcal{L}}}{\partial \delta W_2} + \mathcal{R}_2, \end{aligned} \quad (3.22b)$$

with the reduced loss $\tilde{\mathcal{L}}(\delta W_1, \delta W_2) \equiv \frac{1}{2} \sigma_x^2 \text{tr}[W_2^{0,T} W_2^0 \delta W_1 \delta W_1^T + \delta W_2^{0,T} W_2^0 \delta W_1 W_1^{0,T} + \delta W_2^T \delta W_2 W_1^0 W_1^{0,T}]$ quadratic in the fluctuations.

Recalling the matrix rank equation $\text{rank}(AB) = \min(\text{rank}A, \text{rank}B)$, we note that, for one-dimensional outputs ($N_o = 1$), one has $\text{rank}(W_2^0) = 1 = \text{rank}(W_2^{0,T} W_2^0)$ and thus $\text{rank}([\text{cov}(\mathcal{R}_1)]) = N_i$ according to Eq. (3.21a). For zero-balanced initial conditions [i.e., Eq. (3.6) with vanishing right-hand side], one has $\text{rank}([\text{cov}(\mathcal{R}_2)]) = \text{rank}(W_1^0 W_1^{0,T}) = \text{rank}(W_2^{0,T} W_2^0) = 1$, while for random initial conditions, typically $\text{rank}([\text{cov}(\mathcal{R}_2)]) = \text{rank}(W_1^0 W_1^{0,T}) = N_h$.

3. Transformations

In order to further analyze Eq. (3.22), we expand $\delta W_{1,2}$ in the eigenbasis of $W_{1,2}^0$. To this end, we introduce the thin SVD $W_1^0 = AS_1B^T$ [where A , S_1 and B are (N_h, N_h) , (N_h, N_i) and (N_i, N_i) matrices], and the standard SVD $W_2^0 = US_2V^T$ [where U , S_2 and V are (N_o, N_o) , (N_o, N_o) and (N_h, N_o) matrices], which yield

$$W_1^0 W_1^{0,T} = AD_1A^T \quad \text{with} \quad D_1 \equiv S_1 S_1^T \in \mathbb{R}^{N_h \times N_h}, \quad (3.23a)$$

$$W_2^{0,T} W_2^0 = VD_2V^T \quad \text{with} \quad D_2 \equiv S_2^T S_2 \in \mathbb{R}^{N_o \times N_o}. \quad (3.23b)$$

The eigenvalues of $W_1^0 W_1^{0,T}$ and $W_2^{0,T} W_2^0$ encoded in the diagonal positive-semidefinite matrices $D_{1,2}$ will play a crucial role in the following. Introducing new dynamical variables $z_1 \equiv V^T \delta W_1 \in \mathbb{R}^{N_o \times N_i}$, $z_2 \equiv \delta W_2 A \in \mathbb{R}^{N_o \times N_h}$, and noises $\eta_1 \equiv V^T \mathcal{R}_1 \in \mathbb{R}^{N_o \times N_i}$, $\eta_2 \equiv \mathcal{R}_2 A \in \mathbb{R}^{N_o \times N_h}$, reduces Eq. (3.22) to

$$\dot{z}_1 = -D_2 z_1 - (US_2)^T z_2 S_1 B^T + \eta_1, \quad (3.24a)$$

$$\dot{z}_2 = -z_2 D_1 - US_2 z_1 B S_1^T + \eta_2. \quad (3.24b)$$

The covariance of the new noises $\eta_{1,2}$ follows readily by using the SVDs in Eq. (3.21),

$$\begin{aligned} \text{cov}(\eta_1)_{mj,nl} &= \sigma_{\mathcal{R}}^2 \delta_{jl} D_{2,mn}, \\ &(j, l = 1, \dots, N_i, \quad m, n = 1, \dots, N_o), \end{aligned} \quad (3.25a)$$

$$\begin{aligned} \text{cov}(\eta_2)_{im,kn} &= \sigma_{\mathcal{R}}^2 \delta_{ik} D_{1,mn}, \\ &(m, n = 1, \dots, N_h, \quad i, k = 1, \dots, N_o), \end{aligned} \quad (3.25b)$$

$$\begin{aligned} \text{cov}(\eta_1, \eta_2)_{mj,kn} &= \text{cov}(\eta_2, \eta_1)_{nk,jm} = \sigma_{\mathcal{R}}^2 (US_2)^T_{mk} (B S_1^T)_{jn}, \\ &(j = 1, \dots, N_i, \quad n = 1, \dots, N_h, \\ &m, k = 1, \dots, N_o). \end{aligned} \quad (3.25c)$$

Since the D_i are diagonal, these results represent a significant simplification over Eqs. (3.21) and (3.22).

We now introduce the joint vectors $\mathbf{Z} \equiv \{\text{vec}(z_1), \text{vec}(z_2)\} \in \mathbb{R}^{N_o(N_i+N_h)}$, $\boldsymbol{\eta} \equiv \{\text{vec}(\eta_1), \text{vec}(\eta_2)\}$, which allows us to express Eq. (3.24) in the compact form

$$\dot{\mathbf{Z}} = -\Gamma \mathbf{Z} + \boldsymbol{\eta}. \quad (3.26)$$

The drift matrix Γ and the diffusion matrix $\Delta \equiv \langle \boldsymbol{\eta} \boldsymbol{\eta}^T \rangle$ turn out to be identical symmetric square matrices of size $[N_o(N_i + N_h)]^2$,

$$\Gamma / \sigma_x^2 = \begin{pmatrix} \mathbb{1}_{N_i} \otimes D_2 & B S_1^T \otimes (US_2)^T \\ S_1 B^T \otimes US_2 & D_1 \otimes \mathbb{1}_{N_o} \end{pmatrix} = \Delta / \sigma_{\mathcal{R}}^2. \quad (3.27)$$

Here, \otimes denotes the Kronecker product and in the derivation, we used $\text{vec}(ABC) = (C^T \otimes A) \text{vec}(B)$ as well as $\text{vec}(AB) = (\mathbb{1}_{\dim B} \otimes A) \text{vec}(B) = (B^T \otimes \mathbb{1}_{\dim A}) \text{vec}(A)$ where $\dim_n A$ denotes the size of the n th dimension of A [97]. Next, we determine the eigenvalues and the rank of Γ / σ_x^2 and $\Delta / \sigma_{\mathcal{R}}^2$. After some algebra, using the determinant formula for block matrices [97] and assuming $N_i \geq N_h \geq N_o$, we find the eigenvalues [98]

$$\begin{aligned} \text{multiplicity } 1 : \quad & D_{1,kk} + D_{2,ll}, \quad \text{with } k = 1, \dots, N_h, \\ & l = 1, \dots, N_o, \end{aligned} \quad (3.28a)$$

$$\text{multiplicity } N_i - N_h : \quad D_{2,ll}, \quad \text{with } l = 1, \dots, N_o, \quad (3.28b)$$

$$\text{multiplicity } N_h N_o : \quad 0. \quad (3.28c)$$

Accordingly, Γ and Δ are positive semidefinite matrices with $\text{rank}(\Gamma) = \text{rank}(\Delta) = N_o N_i$. The rank deficiency of Γ and Δ turns out to be a crucial property in the following.

4. General properties

From the results of the preceding subsection one can already infer a few general properties of the fluctuations $\delta W_{1,2}$, which become particularly simple in the case $N_o = 1$: First, from the definition $\delta W_1(t) = V z_{1,1}(t)$, with the column vector V , we infer that all N_h elements in the i th column of δW_1 fluctuate in the same way, i.e., their dynamics is controlled by the scalar $z_{1,i}(t)$, up to time-independent prefactors contained in V . This applies independently from the initialization. Next, we recall $\delta W_2(t) = z_2(t) A^T = (z_2(t) \cdot \bar{a}_1, z_2(t) \cdot \bar{a}_2, \dots, z_2(t) \cdot \bar{a}_{N_h})$, where \bar{a}_i denotes the i th column of A^T . In the case of zero-balanced initial conditions [described by Eq. (3.6) with vanishing right-hand side], one has $1 = \text{rank}(W_2^0) = \text{rank}(W_2^{0,T} W_2^0) = \text{rank}(W_1^0 W_1^{0,T})$ and hence $D_1 = \text{diag}(r, 0, \dots, 0)$ [see Eq. (3.23)]. Accordingly, in this case, the scalars $z_{2,i}(t) = 0$ for $i > 1$, such that $\delta W_2(t) = (\bar{a}_{1,1}, \bar{a}_{2,1}, \dots, \bar{a}_{N_h,1}) z_{2,1}(t)$, i.e., all elements of δW_2 fluctuate identically up to individual scaling factors determining their variance. By contrast, in the case of random initialization, D_1 is typically of maximum rank, implying that all elements of δW_2 fluctuate independently. Numerical results show that the above statements apply also for nonwhite Σ^{xx} .

C. Steady-state fluctuations of weights

Having obtained the linear Langevin equation [Eq. (3.26)] for the weights in each layer, we proceed to determine the resulting steady-state statistics of the weight fluctuations. We find that, as a genuine consequence of the structure of a deep linear network, weight fluctuations are generally anisotropic. We finally provide a derivation of the IVFR.

1. Covariances

The covariance matrix $M \equiv \text{cov}(\mathbf{Z}) = \langle \mathbf{Z} \mathbf{Z}^T \rangle$ describing the fluctuations of the z_i is determined by a Lyapunov equation [10,20,74,89],

$$\Gamma M + M \Gamma^T = \Delta. \quad (3.29)$$




FIG. 8. Covariance of the weights of a two-layer linear network in the quasistationary state. (a) Visualization of the covariance matrix M [Eq. (3.30)] of the transformed weights z_i for a typical simulation with $N_i = 30$, $N_h = 20$, $N_o = 1$. In the special basis defined by the SVD of W_1^0 and W_2^0 [see Eq. (3.23)], the covariance of the second-layer weights is fully diagonal. [(b),(c)] Eigenvalues of M obtained from SGD experiments (symbols), compared to the theoretical predictions obtained from M_{11} and M_{22} in Eq. (3.30) (solid lines), for a single realization of the data. In the simulations, we use nonwhite Gaussian i.i.d. samples $x_i^t \sim \mathcal{N}(0, 1/N_i)$, label noise variance of $\sigma_y^2 = 0.1$, learning rate $\lambda = 0.1$, minibatch size $S = 10$, and an initial variance of the weights of (b) $\sigma_{W_1}^2 = 10^{-4}/N_i$, $\sigma_{W_2}^2 = 10^{-4}/N_h$, and (c) $\sigma_{W_1}^2 = 1/N_i$, $\sigma_{W_2}^2 = 1/N_h$. For comparison, we show the eigenvalues of M obtained by assuming Gaussian i.i.d. distributed elements for W_1^0 and W_2^0 as dashed and dot-dashed lines, respectively. Eigenvalues are normalized by $\lambda\sigma_y^2/(2S)$ [cf. Eq. (3.34)].

As one readily checks (using the fact that $\Gamma \propto \Delta$ and the identity $\mathcal{A}^+ \mathcal{A} = \mathcal{A}$ for a matrix \mathcal{A}), Eq. (3.29) is solved by

$$M \equiv \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} = \frac{1}{2} \Gamma^+ \Delta, \quad (3.30)$$

which is symmetric of size $N_o(N_i + N_h) \times N_o(N_i + N_h)$, while $M_{11} \in \mathbb{R}^{N_o N_i \times N_o N_i}$, $M_{12}, M_{21}^T \in \mathbb{R}^{N_o N_i \times N_o N_h}$, $M_{22} \in \mathbb{R}^{N_o N_h \times N_o N_h}$. This solution assumes detailed balance, which is approximately realized for SGD in the case of sufficient underparameterization [see Eq. (2.33) and the related discussion]. The covariances of the weight fluctuations δW_1 and δW_2 in the original space follow as

$$\text{cov}(\text{vec}(\delta W_1)) = (\mathbb{1}_{N_i} \otimes V) M_{11} (\mathbb{1}_{N_i} \otimes V^T), \quad (3.31a)$$

$$\text{cov}(\text{vec}(\delta W_2)) = (A \otimes \mathbb{1}_{N_o}) M_{22} (A^T \otimes \mathbb{1}_{N_o}), \quad (3.31b)$$

with analogous relations for the cross-covariances. Since $(A \otimes B)^T = A^T \otimes B^T$, the eigenvalues of M_{22} coincide with the eigenvalues of $\text{cov}(\text{vec}(\delta W_2))$, while the eigenvalues of M_{11} coincide with the nonzero eigenvalues of $\text{cov}(\text{vec}(\delta W_1))$, noting that the latter has at least $N_i(N_h - N_o)$ vanishing eigenvalues. A typical covariance matrix M is visualized in Fig. 8(a). Upon averaging over the data distribution, all off-diagonal elements vanish, leaving only the diagonal ones as “universal” imprints of the data statistics (see also the discussion below). However, since we are interested in properties pertaining to a specific realization of the data, we do not perform such averaging here.

Due to the rank deficiency of Γ , the matrix M is, in general, nondiagonal. This can also be inferred from the eigen-decompositions $\Gamma/\sigma_x^2 = \Delta/\sigma_R^2 = R\Omega R^T$, where Ω is a diagonal matrix of rank $N_o N_i$ [see Eq. (3.28)], which implies $M = (R^T)^{-1} \Omega^+ \Omega R^T$ with a diagonal matrix of eigenvalues

$$\Omega^+ \Omega = \frac{\lambda\sigma_y^2}{2S} \mathbb{1}_{N_o(N_i+N_h)}^{N_o N_i}. \quad (3.32)$$

Accordingly, M is of rank $N_o N_i$ and has $N_o N_h$ vanishing eigenvalues.

In order to limit complexity of the analysis, we focus in the following on a one-dimensional output $N_o = 1$. In this case, z_i and η_i become row vectors, while U , S_2 , D_2 become scalars (in fact, $U = \pm 1$), and

$$\begin{aligned} \Gamma/\sigma_x^2 = \Delta/\sigma_R^2 &= \begin{pmatrix} D_2 \mathbb{1}_{N_i} & (US_2)BS_1^T \\ (US_2)S_1 B^T & D_1 \end{pmatrix} \\ &= \begin{pmatrix} V^T W_2^{0,T} W_2^0 V & (W_2^0 V)^T W_1^{0,T} A \\ (W_2^0 V) A^T W_1^0 & A^T W_1^0 W_1^{0,T} A \end{pmatrix}. \end{aligned} \quad (3.33)$$

Using expressions for the pseudoinverse of a block matrix [99], we find, after some tedious but straightforward algebra, that M_{22} [Eq. (3.30)] has a simple diagonal form

$$M_{22} = \text{cov}(z_2) = \frac{\lambda\sigma_y^2}{2S} \text{diag} \left\{ \frac{D_{1,ii}}{D_{1,ii} + D_2} \right\}_{i=1, \dots, N_h}. \quad (3.34)$$

Accordingly, the components of z_2 [see Eq. (3.24)] fluctuate independently and $M_{22,ii}$ in fact represent the eigenvalues of $\text{cov}(\delta W_2)$ [see Eq. (3.31b)]. By contrast, M_{11} is, in general, not diagonal [see Fig. 8(a)], although its off-diagonal elements are typically small, implying that the components of z_1 are weakly cross correlated.

2. Covariances: discussion and numerical results

We conclude this analysis with several remarks: (1) Without the cross coupling between δW_1 and δW_2 [see Eq. (3.22)], Γ and Δ would be diagonal and full rank, which would imply $M \propto \mathbb{1}_{N_i+N_h}$ and thus $[\text{cov}(\delta W_1)] \propto \mathbb{1}_{N_i+N_h}$ and $[\text{cov}(\delta W_2)] \propto \mathbb{1}_{N_o \times N_h}$, i.e., isotropic fluctuations. (2) In the case of Gaussian i.i.d. data considered here, the analysis starting from Eq. (3.23) essentially relied only on the eigenstructure of the weight product matrices $W_1^0 W_1^{0,T}$ and $W_2^{0,T} W_2^0$, which enter both the drift and the diffusion matrix [Eq. (3.27)]. The data is assumed to be structureless [with whitened inputs, see Eq. (3.13)] and thus appears here only as scaling factors σ_x^2 , σ_y^2 . Numerical results confirm that similar fluctuation spectra are also obtained for nonwhite data. (3) The

solution in Eq. (3.3) only determines the $N_{\text{eff}} < N_{\text{tot}}$ degrees of freedom of $W_2^0 W_1^0$ uniquely in terms of the data, implying that the individual weight matrices W_1^0 and W_2^0 (and hence also $D_{1,2}$) are additionally affected by the initialization. Specifically, in the limit $N_{\text{tot}} \gg N_{\text{eff}}$, the weights hardly change from their initial values during training (lazy training regime) [100,101], allowing one to approximate a nonlinear neural network by its linearization around initialization [61,62]. Assuming $W_i^0(t=0) \sim \mathcal{N}(0, \sigma_{W_i})$, $W_1^0 W_1^{0,T}$ and $W_2^{0,T} W_2^0$ in this case can be approximated as Wishart matrices, the eigenvalues of which asymptotically follow a Marchenko-Pastur distribution with characteristic magnitude $D_i \propto \sigma_{W_i}^2$. When initializing the weights conventionally with [102]

$$\sigma_{W_1}^2 \simeq 1/N_i, \quad \sigma_{W_2}^2 \simeq 1/N_h, \quad (3.35)$$

and assuming $N_h \lesssim N_i$, one accordingly expects $D_{1,ii} \lesssim D_2$ [see Eq. (3.23)], such that Eq. (3.34) can be approximated as

$$M_{22} \simeq \frac{\lambda \sigma_y^2}{2SD_2} D_1, \quad (3.36)$$

except for the few largest eigenvalues in D_1 . We conclude that for deep linear nets initialized in this way, weight fluctuations are generally *anisotropic* and their spectrum has a certain universal character encoded in the limiting eigenvalue distribution. A notable exception occurs for zero-balanced initial conditions [see Eq. (3.6)], where, owing to $D_{1,ii} \propto \delta_{i,0}$ (see Sec. III B 4), the fluctuation spectrum is isotropic. (4) We recall that the derivation leading to Eqs. (3.34) and (3.36), demonstrating the anisotropy of the fluctuation spectrum, is obtained within the large-sample number limit [$P \rightarrow \infty$; see Eq. (3.12) and the related discussion], which, in the case of a single-layer net, would imply isotropic weight fluctuations [see Eq. (2.34)].

Figures 8(b) and 8(c) shows simulation results obtained from SGD training of a two-layer linear network on (non-white) Gaussian i.i.d. data for Gaussian weight initializations [Eq. (3.35)] of (b) small and (c) large variance. We find good agreement with the covariance spectrum of the weights predicted by Eq. (3.30), which is a linearization of the full stochastic gradient flow [see Eq. (3.18)]. In particular, the assumption of whitened data in the theory does not appear to qualitatively affect the results. Confirming Eq. (3.34), the spectrum of the fluctuations of W_2 is essentially determined by the eigenvalues of $W_1^0 W_1^{0,T}$ [see Eq. (3.23)]. Eigenvalues of $\text{cov}(\delta W_1)$ with index larger than N_i vanish according to the linear theory [Eqs. (3.30) and (3.31)], while in our simulations, they are nonzero but decrease in magnitude with increasing index. Comparing the two panels, we observe that, for large values of the initial weight variance, the spectrum remains closer to the shape implied for Gaussian i.i.d. elements of $W_{1,2}^0$. For weights close to being zero balanced, the fluctuations of W_2 are essentially governed by a single independent mode (see discussion in Sec. III B 4), which is reflected by the prominent peak of the lowest mode in Fig. 8(b). The eigenvalue peak of the covariance of W_1 seen in (c) can be associated to an overall drift of the mean of W_1 , which occurs on a large time scale. In fact, although the loss is stationary, we observe in our SGD experiments a transient dynamics of

the individual weights, resulting in a slow but steady drop of the overall magnitude of the eigenvalues (except for the zero mode). This dynamics, which is most pronounced for small batch sizes, is not captured by the present analysis, since the latter only provides a relation between the covariances of $W_{1,2}$ and the products $W_1^0 W_1^{0,T}$, $W_2^{0,T} W_2^0$, assuming quasistationary means W_1^0 and W_2^0 . A complete description requires a solution of the GF equations (3.17) and is reserved for a future study.

3. Inverse-variance flatness relation

We now study how the fluctuation variance relates to the loss, thereby providing an analytical derivation of the IVFR [16]. The reduced quadratic loss [see Eq. (3.22)] is given by $\tilde{L}(z_1, z_2) = \frac{1}{2} \sigma_x^2 \text{tr}[D_2 z_1 z_1^T + z_2^T z_2 D_1 + z_2^T (US_2) z_1 (BS_1^T)]$, which indicates that the perturbation of the loss due to fluctuations of z_1 and z_2 is controlled by the spectrum of $W_2^{0,T} W_2^0$ and $W_1^0 W_1^{0,T}$, respectively [see Eq. (3.23)]. In order to capture the effect of the fluctuations in the original weight space, we start from Eq. (3.1) and recall that the principal modes $\mathbf{a}_{i=1, \dots, N_h}$ of the fluctuations of W_2 (i.e., the columns of A) fulfill $W_1^0 W_1^{0,T} \mathbf{a}_i = D_{1,ii} \mathbf{a}_i$ with $D_{1,ii} \simeq 2SD_2 M_{22,ii} / (\lambda \sigma_y^2)$ [see Eqs. (3.31) and (3.36)]. Using the basic relation $\text{tr}(AB) = \text{tr}(BA)$ as well as the fact that $W_2^0 W_1^0 = \sigma_x^{-2} \Sigma_0^{yx}$ for whitened inputs, the perturbation of the loss [Eq. (3.1)] along \mathbf{a}_i follows, after a bit of algebra, as [cf. Eq. (2.35)]

$$\begin{aligned} \delta L(\theta \mathbf{a}_i) &= L(W_1^0, W_2^0 + \theta \mathbf{a}_i^T) - L(W_1^0, W_2^0) \\ &= \frac{1}{2} \sigma_x^2 \theta^2 D_{1,ii} \simeq \frac{SD_2 \sigma_x^2}{\lambda \sigma_y^2} \theta^2 M_{22,ii}. \end{aligned} \quad (3.37)$$

As $W_2^0 \in \mathbb{R}^{N_o \times N_h}$ and $N_o = 1$, we reshaped the vector \mathbf{a}_i into a compatible form via its transpose. Numerical experiments with SGD in the large batch regime as well as with SGF agree well with this prediction [see Fig. 9(a)]. Since the ‘‘flatness’’ F_i of the loss $\propto (\text{curvature } \kappa_i)^{-1/2}$ and the curvature κ_i of the parabolas defined by Eq. (3.37) is given by their quadratic coefficient, i.e., $\kappa_i = 2S\sigma_x^2 D_2 M_{22,ii} / (\lambda \sigma_y^2)$, we recover the inverse-variance flatness relation in the form $M_{22,ii} \propto F_i^{-\psi}$ with $\psi = 2$. In Ref. [16], empirically $\psi \approx 4$ was found in the case of cross-entropy loss, but it was emphasized that the value of the exponent can depend on the specific training setting. We remark that, due to the transient weight dynamics persisting under SGD (see Sec. III C 2), we observe that δL slowly diminishes with time, while remaining consistent with the IVFR.

By contrast, for a perturbation of the loss along eigendirections \mathbf{v}_i of $\text{cov}(\text{vec } \delta W_1)$ [Eq. (3.31a)], we do not observe any dependence on the mode index i , i.e., $\delta L(\theta \mathbf{v}_i) = L(W_1^0 + \theta[\mathbf{v}_i], W_2^0) - L(W_1^0, W_2^0) \simeq \text{const.}$, even though the weight covariance is anisotropic [see Fig. 9(b)]. Here, $[\mathbf{v}_i]$ represents the vector \mathbf{v}_i reshaped into the form of W_1^0 . Note, however, that δL determined within the linearized fluctuation theory [i.e., using Eq. (3.30)] does not capture the bimodal structure of δL observed in numerical solutions of SGD or SGF [Eq. (3.15)], but instead remains flat. We thus find that, for a two-layer linear network, the IVFR is not fulfilled for perturbations of the first-layer weights W_1 .

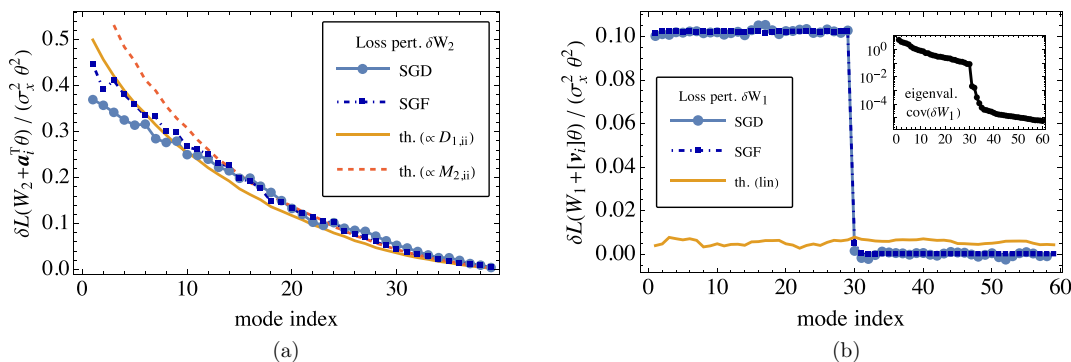


FIG. 9. Perturbing the loss L of a two-layer linear network [Eq. (3.1)] by the eigenvectors of the covariance matrix of the fluctuations of (a) δW_2 [see Eq. (3.37)] and (b) δW_1 . Results of training on Gaussian i.i.d. data ($\sigma_x^2 = 1/N_i$, $\sigma_y^2 = 1$) with SGD (connected circles) are compared to the theoretical predictions (solid lines), given by Eq. (3.37) in (a), and by the numerically evaluated $\delta L(\theta \mathbf{v}_i)$ using the eigenvectors \mathbf{v}_i of $\text{cov}(\text{vec}(\delta W_1))$ [Eq. (3.31a)] in (b). The dashed-dotted line represents the prediction based on the nonlinear stochastic gradient flow. In (a), beside the exact result $\delta L = \sigma_x^2 \theta^2 D_{1,ii}/2$ (solid line), also the approximation [see Eq. (3.37)] in terms of the eigenvalues of the covariance M_{22} is plotted (dashed line). The inset in (b) shows the eigenvalues of $\text{cov}(\text{vec}(\delta W_1))$ ordered by magnitude vs the mode index i . For $i > N_i$, the eigenvalues abruptly drop to a smaller level, while they are predicted to vanish according to the linear model (see discussion in Sec. III C 2). The corresponding drop of δL at $i = N_i$ is captured by the nonlinear stochastic gradient flow simulations, while the linear theory [solid line in (b)] instead predicts a flat profile for δL . We used $N_o = 1$, $N_h = 40$, $N_i = 50$ in (a) and $N_o = 1$, $N_h = 20$, $N_i = 30$ in (b). The behaviors for whitened and nonwhitened data are similar.

D. Relaxation dynamics

We illustrate here how the linearized SGF derived in the previous section can be used to analyze the layerwise relaxation dynamics of the weights. Focusing on $N_o = 1$, we diagonalize the (symmetric) drift and diffusion matrices as $\Gamma/\sigma_x^2 = \Delta/\sigma_R^2 = R\Omega R^T$ [see Eq. (3.33)], with a diagonal $\Omega = \text{diag}(\omega_k)$ ($k = 1, \dots, N_h + N_i$) and an orthogonal matrix R with eigenvectors in the columns. According to Eq. (3.28), the eigenvalues are given by

$$\begin{aligned} \omega_{k=1, \dots, N_h} &= D_{1,kk} + D_2, & \omega_{k=N_h+1, \dots, N_i} &= D_2, \\ \omega_{k=N_i+1, \dots, N_i+N_h} &= 0, \end{aligned} \quad (3.38)$$

with $D_{1,2}$ defined in Eq. (3.23). Note that all $\omega_k \geq 0$. Introducing new dynamical variables $\chi = R^T \mathbf{Z}$ and noises $\zeta = R^T \eta$, we transform Eq. (3.26) to

$$\begin{aligned} \dot{\chi}_k(t) &= \sigma_x^2 \omega_k \chi_k(t) + \zeta_k(t), \\ \langle \zeta_k(t) \zeta_{k'}(t') \rangle &= \sigma_R^2 \omega_k \delta_{kk'} \delta(t - t'). \end{aligned} \quad (3.39)$$

The general solution is given by $\chi_k(t) = e^{-\omega_k(t-t_0)} \chi_k(t_0) + \int_{t_0}^t ds e^{\omega_k(s-t)} \zeta_k(s)$ with some initial time t_0 . Accordingly, in the stationary state, one obtains the correlation functions

$$\langle \chi_k(t) \chi_{k'}(0) \rangle = \langle \chi_k(0) \chi_{k'}(0) \rangle e^{-\sigma_x^2 \omega_k t}, \quad (3.40)$$

$$\begin{aligned} \langle \chi_k(t) \chi_{k'}(t) \rangle &= e^{-\sigma_x^2 (\omega_k + \omega_{k'}) t} \langle \chi_k(0) \chi_{k'}(0) \rangle \\ &+ \delta_{kk'} \frac{\sigma_R^2}{2\sigma_x^2} (1 - e^{-2\sigma_x^2 \omega_k t}), \end{aligned} \quad (3.41)$$

where we assumed time-translation invariance, which allows us to chose the arbitrary time origin as $t_0 = 0$. Modes χ_k belonging to eigenvalues $\omega_k = 0$ remain fixed at their initialization $\chi_k(t) = \chi_k(0)$. By contrast, modes with $\omega_k > 0$ relax

towards zero [Eq. (3.40)], while their variance saturates at late times at the equilibrium value [Eq. (3.41)],

$$\langle \chi_k \chi_{k'} \rangle_{\text{eq}} = \begin{cases} \langle \chi_k(0) \chi_{k'}(0) \rangle, & \omega_k = 0 = \omega_{k'} \\ \delta_{kk'} \frac{\sigma_R^2}{2\sigma_x^2}, & \text{otherwise.} \end{cases} \quad (3.42)$$

Notably, as a consequence of the equality of Γ and Δ (up to an overall scale), the fluctuations of the modes χ are isotropic (independent of k) at late times and only determined by the relative strength of the noise (σ_R^2) and the loss (σ_x^2). Since the initial values of $W_{1,2}$ are uncorrelated, we may assume $\langle \chi_k(0) \chi_{k'}(0) \rangle \propto \delta_{kk'}$. In order to transform back into the original weight space, one uses $\delta W_1 = V[\mathcal{P}_1 R \chi]^T$ and $\delta W_2 = [\mathcal{P}_2 R \chi]^T A^T$ [see Eq. (3.23) and notes that \mathbf{Z} is a column vector], with the projection matrices $\mathcal{P}_1 = \begin{pmatrix} \mathbb{1}_{N_i} & 0 \\ 0 & 0 \end{pmatrix}$ and $\mathcal{P}_2 = \begin{pmatrix} 0 & 0 \\ 0 & \mathbb{1}_{N_h} \end{pmatrix}$, both of size $(N_i + N_h)^2$. From these considerations we infer that the relaxation dynamics of the weights is in general a superposition of exponentials, dominated by the smallest eigenvalue ω_k . A more detailed investigation of the noise-driven time evolution is reserved for a future work.

IV. SUMMARY AND OUTLOOK

In the present study, we have analytically and numerically investigated the fluctuating dynamics of linear networks trained with mean-square loss. The considered single-layer linear network trained with SGD solves an ordinary least-squares regression problem. Besides being of theoretical interest, such linear models approximate the stationary training regime of deep nonlinear networks (see Appendix A) and provide a basis for the construction of SGD-based Bayesian-sampling methods (see Appendix B). Two-layer linear networks share important properties with deep nonlinear nets, such as nonlinear training dynamics, a

nonconvex loss landscape, and a sensitivity to initial conditions [4,6,63].

In our analytical calculations, we have considered uncorrelated Gaussian i.i.d. training data with targets generated by linearly transforming the inputs and adding noise (including the limit of fully random targets). Correlated or non-Gaussian input data will in general affect the gradient noise covariance [see Eqs. (2.19) and (2.20)] and thereby the weight statistics. However, for sufficiently large minibatch sizes and small learning rates, the additive Gaussian approximation for the noise is usually warranted (see discussion in Sec. II B 2 and Refs. [10,11,51,52]). Crucially, these assumptions enable a linearization of the nonlinear SGD dynamics of two-layer networks and thereby facilitate analytical calculations. Our theory applies to neural networks in a (quasi)stationary regime, characterized by a loss fluctuating around a nonzero mean value. Typically, such a state is reached at the end of training of an underparameterized network. Overparameterized networks, by contrast, can memorize their training data [103] and approach a global optimum in which gradient noise vanishes [37,104–108]. In this case, the loss does not become strictly stationary, but decays on a timescale much larger than the weight fluctuations. Indeed, the inverse-variance flatness relation (IVFR) has been considered in such a situation in Ref. [16]. Our perturbative approach to a two-layer network approximates such a quasistationary state by assuming that the mean of the weights evolve sufficiently slowly.

The main motivation of our present paper was to study the simplest possible model that gives rise to the IVFR and to gain further insights into the nature of the SGD noise and the induced weight fluctuations. We have obtained the following main results:

(1) We have shown that one source of anisotropy of the weight fluctuations is the broken detailed balance of SGD [12,20], which emerges here in the presence of label noise for mildly underparameterized nets. In the infinite-sample limit, detailed balance is restored and the noise covariance becomes equal to the Hessian of the loss.

(2) For two-layer linear nets, another source of fluctuation anisotropy stems from the rank deficiency of the drift and diffusion matrices that describe the stochastic gradient flow of the weight fluctuations in each layer. We expect this general consequence of the network structure to also hold for deeper linear networks.

(3) Based on these insights, we have provided an analytical derivation of the IVFR for a two-layer linear network in the underparameterized regime. Our approach is based on a perturbative treatment of weight fluctuations around a quasistationary state and thereby adds a further perspective to existing explanations of the IVFR [16,18,19].

(4) As a byproduct of our investigation, we have described the layerwise continuum dynamics of SGD in two-layer linear nets, extending previous works that considered the deterministic gradient flow [4,9,24,69].

Understanding weight fluctuations in neural networks is not only interesting from a theoretical perspective [13–15], but has a number of practical implications. Noise-induced transition events, for instance, can be relevant for optimization in a rough loss landscape [42,43,46,49]. Recently, the

fluctuation-dissipation ratio, which has its origin in statistical physics [109,110], has been introduced as a novel training measure [38]. In the context of differential privacy [111], random noise is added to the gradient updates. In Ref. [16], the IVFR has been observed empirically for various deep nonlinear network architectures trained on image data under cross-entropy loss and explained within a phenomenological model of a random loss landscape. Our derivation is based on layerwise perturbation of the SGD dynamics, which should also be applicable to deep nonlinear network models. It will be interesting in the future to connect these two approaches and further clarify the relevance of the effects identified here (interlayer coupling and rank deficiency) for fluctuation and loss anisotropy. We have noted that, in the case of a two-layer linear network, the weight spectrum is significantly determined by the initialization and the IVFR emerges while the SGD dynamics is still transient despite a stationary loss. These aspects deserve a further investigation, which could consider also effects of non-Gaussianity and higher-order correlations stemming from the data or from nonlinearities in the dynamics [54,112].

APPENDIX A: NONLINEAR NETWORKS

Nonlinear networks can be linearized in the stationary state of training, provided deviations from the reference state are sufficiently small, which requires large batch sizes and small learning rates (see the discussion in Sec. II B 2). Notably, for large widths, nonlinear networks reduce to linear models [61,62] and the associated limiting dynamics of SGD is a topic of recent interest [37,107,113–115]. For the linear models considered in the present study, steady-state weight fluctuations vanish in the overparameterized regime [see Eq. (2.16)].

We denote the output function of the nonlinear network as $f(\mathbf{x}, \mathbf{w})$, where $\mathbf{w} \in \mathbb{R}^{N_{\text{tot}}}$ collectively represents the network parameters. Taylor expansion around a reference point \mathbf{w}^* renders

$$\begin{aligned} f(\mathbf{x}, \mathbf{w}) &\simeq f(\mathbf{x}, \mathbf{w}^*) + (\mathbf{w} - \mathbf{w}^*) \cdot \nabla_{\mathbf{w}} f(\mathbf{x}, \mathbf{w}^*) \\ &= b(\mathbf{x}) + \mathbf{w} \cdot \nabla_{\mathbf{w}} f(\mathbf{x}, \mathbf{w}^*), \end{aligned} \quad (\text{A1})$$

with a bias term $b(\mathbf{x}) = f(\mathbf{x}, \mathbf{w}^*) - \mathbf{w}^* \cdot \nabla_{\mathbf{w}} f(\mathbf{x}, \mathbf{w}^*)$ that is independent of \mathbf{w} . Defining a Jacobian matrix

$$\Phi = \begin{pmatrix} \nabla_{\mathbf{w}}^T f(\mathbf{x}_1, \mathbf{w}^*) \\ \vdots \\ \nabla_{\mathbf{w}}^T f(\mathbf{x}^P, \mathbf{w}^*) \end{pmatrix} \in \mathbb{R}^{P \times N}, \quad (\text{A2})$$

we can write the MSE loss as

$$\begin{aligned} L(\mathbf{w}) &= \frac{1}{2P} \sum_{\mu} [f(\mathbf{x}^{\mu}, \mathbf{w}) - y^{\mu}]^2 \\ &= \frac{1}{2P} \|\mathbf{w}^T \Phi - Y\|^2 + \text{const.} \end{aligned} \quad (\text{A3})$$

up to a constant independent of \mathbf{w} . Accordingly, the linear framework of Sec. II can be adopted by replacing the design matrix X by the Jacobian matrix Φ [see Eq. (2.3)]. We remark

that the result resembles a random feature model, although the features $\partial_{w_i} f(\mathbf{x}, \mathbf{w}^*)$ are not independent [116–119].

APPENDIX B: CONNECTION TO BAYESIAN METHODS

Since Bayesian neural networks involve the weight distribution as a central ingredient, it is informative to briefly review this approach and its connection to SGD-based sampling here [120–122]. The prediction y^* for a new data point x^* is computed from the likelihood $p(y^*|x^*, w)$ weighted by the posterior distribution $p(w|\mathcal{D})$: $p(y^*|x^*, \mathcal{D}) = \int_w p(y^*|x^*, w)p(w|\mathcal{D})dw$, where $\mathcal{D} = \{x^\mu, y^\mu\}_\mu$ denotes the training data. According to Bayes' rule, the posterior is given by $p(w|\mathcal{D}) \propto p(y|x, w)p(w)$, where $p(w)$ is the prior distribution of the weights and $p(y|x, w) \propto \exp(-\mathcal{L}(x, y, w))$ is the likelihood function represented here in terms of a loss \mathcal{L} . A regularized loss can accordingly be introduced by writing $p(w|\mathcal{D}) \propto \exp(-\mathcal{L}_{\text{reg}}(x, y, w))$ with $\mathcal{L}_{\text{reg}} = \mathcal{L} + \chi \|w\|_2^2$, where the last equation holds for a Gaussian prior $p(w)$, which gives rise to an L_2 regularization. Classical (non-stochastic) training amounts to approximating the posterior as $p(w|\mathcal{D}) = \delta(w - w_{\text{map}})$ using the point estimate $w_{\text{map}} = \text{argmin}_w \mathbb{E}_{\mathcal{D}} \mathcal{L}_{\text{reg}}(x, y, w)$, where $\mathbb{E}_{\mathcal{D}}$ denotes the average over the training data.

In the stationary state of SGD, weight samples are drawn from Eqs. (2.30) and (2.32), which represents a different distribution from the true posterior. Thus, in order to use SGD as a Bayesian sampler, several methods have been proposed in the literature, such as preconditioning [10,123] or injecting Gaussian white noise into the SGD dynamics [19,124,125]. This isotropic white noise is constructed such that it dominates over the minibatch-sampling noise in the limit of small learning rates, implying $C + Q \rightarrow \mathbb{1}$ in Eq. (2.30).

APPENDIX C: COVARIANCE OF MINI-BATCHES

Consider (possibly vectorial) quantities

$$G_B = \frac{1}{S} \sum_{\mu \in \mathcal{B}} g^\mu, \quad G = \langle G_B \rangle_{\mathcal{B}} = \frac{1}{P} \sum_{\mu=1}^P g^\mu, \quad (C1)$$

where G_B is defined as an average over minibatches \mathcal{B} of size $S = |\mathcal{B}|$, while G is given by the average of G_B over the full batch of size $P > S$. The covariance of G_B and minibatch sampling with replacement can be calculated, following similar steps as in Appendix E of [21]:

$$\begin{aligned} \text{cov}(G_B) &\equiv \langle (G_B - \langle G_B \rangle_{\mathcal{B}})(G_B - \langle G_B \rangle_{\mathcal{B}})^T \rangle_{\mathcal{B}} \\ &= \langle G_B G_B^T \rangle_{\mathcal{B}} - G G^T \\ &= \frac{1}{S} \left[\frac{1}{P} \sum_{\mu=1}^P g^\mu (g^\mu)^T - G G^T \right]. \end{aligned} \quad (C2)$$

The corresponding expression for sampling without replacement takes the same form but with the prefactor $(P - S)/(S(P - 1))$ instead of $1/S$ [21]. The two cases coincide if $S = 1$ or $P \gg S$.

APPENDIX D: ADDITIVE NOISE APPROXIMATION

We justify in the following the use of the additive noise approximation in the stationary training regime of a linear network. We start from the Langevin Eq. (2.13), which we write

$$\dot{\mathbf{w}} = -H(\mathbf{w} - \mathbf{w}^*) + \sqrt{\eta} \boldsymbol{\zeta}, \quad \eta \equiv \frac{\lambda}{S} \quad (D1)$$

with \mathbf{w}^* given by Eq. (2.7). The correlations of the noise $\boldsymbol{\zeta}$ are described by the large-sample number limit [Eq. (2.19)],

$$\begin{aligned} \langle \boldsymbol{\zeta}(t) \boldsymbol{\zeta}(t')^T \rangle &= SC(\mathbf{w}) \delta(t - t') \\ &= [H \bar{\mathbf{w}} \bar{\mathbf{w}}^T H + \text{tr}(H \bar{\mathbf{w}} \bar{\mathbf{w}}^T) H + \sigma_\epsilon^2 H] \delta(t - t'), \end{aligned} \quad (D2)$$

which is sufficient for the purpose of estimating the importance of the various contributions. We consider η as a smallness parameter and assume an expansion of the form

$$\mathbf{w} = \mathbf{w}_0 + \sqrt{\eta} \mathbf{w}_1 + \mathcal{O}(\eta). \quad (D3)$$

Analogously, we expand the noise as $\boldsymbol{\zeta} = \boldsymbol{\zeta}_0 + \boldsymbol{\zeta}_1 + \dots$, where the correlations of the new noises $\boldsymbol{\zeta}_i$ follow successively by evaluating Eq. (D2) with Eq. (D3).

Inserting Eq. (D3) into Eq. (D1) readily yields

$$\mathbf{w}_0 = \mathbf{w}^*. \quad (D4)$$

Next, using $\mathbf{w}^* = \mathbf{u} + H^{-1}J$ [Eq. (2.7)] renders

$$\begin{aligned} \langle \boldsymbol{\zeta}_0(t) \boldsymbol{\zeta}_0(t')^T \rangle &= C_0 \delta(t - t'), \\ C_0 &= [J J^T + \text{tr}(J J^T H^{-1}) H + \sigma_\epsilon^2 H] \\ &\sim \sigma_x^2 \sigma_\epsilon^2, \end{aligned} \quad (D5)$$

and the Langevin equation at the order $\sqrt{\eta}$,

$$\dot{\mathbf{w}}_1 = -H \mathbf{w}_1 + \boldsymbol{\zeta}_0. \quad (D6)$$

The variance of the resulting stationary solution is given by [cf. Eq. (2.33)]

$$\langle \mathbf{w}_1 \mathbf{w}_1^T \rangle = \frac{1}{2} H^{-1} C_0 \sim \sigma_\epsilon^2. \quad (D7)$$

Using this result in the expansion of Eq. (D2) provides the $\mathcal{O}(\eta)$ term of the noise correlation,

$$\begin{aligned} \langle \boldsymbol{\zeta}_1(t) \boldsymbol{\zeta}_1(t')^T \rangle &= C_1 \delta(t - t'), \\ C_1 &= \frac{\eta}{2} [H \langle w_1 w_1^T \rangle H + \text{tr}(H \langle w_1 w_1^T \rangle) H] \\ &\sim \eta \sigma_x^4 \sigma_\epsilon^2 \end{aligned} \quad (D8)$$

and a corresponding $\mathcal{O}(\eta^{1/2})$ term with correlation $C_{1/2} = \sqrt{\eta/2} J \mathbf{w}_1^T H + \dots \sim \sqrt{\eta} \sigma_x^2 \sigma_\epsilon^2$ stemming from the cross terms between \mathbf{w}^* and \mathbf{w}_1 . A comparison of the relative magnitude of these contributions and the one in Eq. (D5) indicates that the multiplicative noise $\boldsymbol{\zeta}_1$ remains subdominant in the stationary state if $\lambda/S \ll 1$ and $\lambda \sigma_x^2/S \ll 1$ are fulfilled, which is the case for typical setups. A more rigorous analysis based on adiabatic elimination techniques [75,126] for the Fokker-Planck equation [Eq. (2.28)] is reserved for a future work.

APPENDIX E: NOISE COVARIANCE

Assuming whitened inputs, Eqs. (3.19) and (3.20) imply

$$\begin{aligned}
\lambda^{-1} \text{cov}(\mathcal{R}_1)_{ij,kl} &= (W_2^{0,T} \Sigma_0^{yx})_{il} (\Sigma_0^{yx,T} W_2^0)_{jk} + \sigma_x^2 \langle y^2 \rangle \delta_{jl} (W_2^{0,T} W_2^0)_{ik} - \sigma_x^2 (W_2^{0,T} \Sigma_0^{yx} w^{*,T} W_2^0)_{ik} \delta_{jl} \\
&\quad - \sigma_x^2 (W_2^{0,T} w^* \Sigma_0^{yx,T} W_2^0)_{ik} \delta_{jl} - \sigma_x^2 (W_2^{0,T} \Sigma_0^{yx})_{il} (w^{*,T} W_2^0)_{jk} - \sigma_x^2 (W_2^{0,T} w^*)_{il} (\Sigma_0^{yx,T} W_2^0)_{jk} \\
&\quad + \sigma_x^4 (W_2^{0,T} w^* w^{*,T} W_2^0)_{ik} \delta_{jl} + \sigma_x^4 (W_2^{0,T} w^*)_{il} (w^{*,T} W_2^0)_{jk} \\
&= \sigma_x^2 \langle y^2 \rangle \delta_{jl} (W_2^{0,T} W_2^0)_{ik} - (W_2^{0,T} \Sigma_0^{yx} \Sigma_0^{yx,T} W_2^0)_{ik} \delta_{jl} \\
&= \sigma_x^2 \langle y^2 \rangle \delta_{jl} (1 - 1/P) (W_2^{0,T} W_2^0)_{ik}, \tag{E1}
\end{aligned}$$

$$\begin{aligned}
\lambda^{-1} \text{cov}(\mathcal{R}_2)_{ij,kl} &= (\Sigma_0^{yx} W_1^{0,T})_{il} (W_1^0 \Sigma_0^{yx,T})_{jk} + \sigma_x^2 \langle y^2 \rangle \delta_{ik} (W_1^0 W_1^{0,T})_{jl} - \sigma_x^2 (\Sigma_0^{yx} W_1^{0,T})_{il} (W_1^0 w^{*,T})_{jk} \\
&\quad - \sigma_x^2 (\Sigma_0^{yx} w^{*,T})_{ik} (W_1^0 W_1^{0,T})_{jl} - \sigma_x^2 (w^* W_1^{0,T})_{il} (W_1^0 \Sigma_0^{yx,T})_{jk} - \sigma_x^2 (w^* \Sigma_0^{yx,T})_{ik} (W_1^0 W_1^{0,T})_{jl} \\
&\quad + \sigma_x^4 (w^* W_1^{0,T})_{il} (W_1^0 w^{*,T})_{jk} + \sigma_x^4 (w^* w^{*,T})_{ik} (W_1^0 W_1^{0,T})_{jl} \\
&= \sigma_x^2 \langle y^2 \rangle \delta_{ik} (1 - 1/P) (W_1^0 W_1^{0,T})_{jl}, \tag{E2}
\end{aligned}$$

$$\begin{aligned}
\lambda^{-1} \text{cov}(\mathcal{R}_1, \mathcal{R}_2)_{ij,kl} &= (W_2^{0,T} \Sigma_0^{yx} W_1^{0,T})_{il} \Sigma_{0,jk}^{yx,T} + \sigma_x^2 \langle y^2 \rangle W_{2,ik}^{0,T} W_{1,jl}^{0,T} - \sigma_x^2 (W_2^{0,T} \Sigma_0^{yx} w^{*,T})_{ik} W_{1,jl}^{0,T} \\
&\quad - \sigma_x^2 (W_2^{0,T} w^* \Sigma_0^{yx,T})_{ik} W_{1,jl}^{0,T} - \sigma_x^2 (W_2^{0,T} \Sigma_0^{yx} W_1^{0,T})_{il} w_{jk}^{*,T} - \sigma_x^2 (W_2^{0,T} w^* W_1^{0,T})_{il} \Sigma_{0,jk}^{yx,T} \\
&\quad + \sigma_x^4 (W_2^{0,T} w^* w^{*,T})_{ik} W_{1,jl}^{0,T} + \sigma_x^4 (W_2^{0,T} w^* W_1^{0,T})_{il} w_{jk}^{*,T} \\
&= \sigma_x^2 \langle y^2 \rangle W_{2,ik}^{0,T} W_{1,jl}^{0,T} - (W_2^{0,T} \Sigma_0^{yx} \Sigma_0^{yx,T})_{ik} W_{1,jl}^{0,T} \\
&= \sigma_x^2 \langle y^2 \rangle (1 - 1/P) W_{2,ik}^{0,T} W_{1,jl}^{0,T}, \tag{E3}
\end{aligned}$$

which leads to Eq. (3.21).

-
- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [2] C. C. Aggarwal, *Neural Networks and Deep Learning: A Textbook* (Springer International Publishing, Cham, 2018).
- [3] K. P. Murphy, *Probabilistic Machine Learning: An Introduction* (The MIT Press, Cambridge, MA, 2022).
- [4] A. M. Saxe, J. L. McClelland, and S. Ganguli, Exact solutions to the nonlinear dynamics of learning in deep linear neural networks, [arXiv:1312.6120](https://arxiv.org/abs/1312.6120).
- [5] 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 (2018).
- [6] S. Arora, N. Cohen, W. Hu, and Y. Luo, Implicit regularization in deep matrix factorization, in *Advances in Neural Information Processing Systems, Vancouver, Canada*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), Vol. 32.
- [7] M. Advani, A. M. Saxe, and H. Sompolinsky, High-dimensional dynamics of generalization error in neural networks. *Neural Netw.* **132**, 428 (2020).
- [8] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, Surprises in high-dimensional ridgeless least squares interpolation, *Ann. Stat.* **50**, 949 (2022).
- [9] L. Braun, C. Dominé, J. Fitzgerald, and A. M. Saxe, Exact learning dynamics of deep linear networks with prior knowledge, in *Advances in Neural Information Processing Systems*, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Curran Associates, Inc., 2022), Vol. 35, pp. 6615–6629.
- [10] S. Mandt, M. D. Hoffman, and D. M. Blei, Stochastic gradient descent as approximate Bayesian inference, *J. Mach. Learn. Res.* **18**, 4873 (2017).
- [11] S. Jastrzebski, Z. Kenton, D. Arpit, N. Ballas, A. Fischer, A. Storkey, and Y. Bengio, Three factors influencing minima in SGD, [arXiv:1711.04623](https://arxiv.org/abs/1711.04623).
- [12] P. Chaudhari and S. Soatto, Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks, [arXiv:1710.11029](https://arxiv.org/abs/1710.11029).
- [13] A. Engel and C. Van den Broeck, *Statistical Mechanics of Learning* (Cambridge University Press, Cambridge, 2001).
- [14] Y. Bahri, J. Kadmon, J. Pennington, S. S. Schoenholz, J. Sohl-Dickstein, and S. Ganguli, Statistical mechanics of deep learning, *Annu. Rev. Condens. Matter Phys.* **11**, 501 (2020).
- [15] H. Huang, *Statistical Mechanics of Neural Networks* (Springer Nature, Singapore, 2021).
- [16] Y. Feng and Y. Tu, The inverse variance–flatness relation in stochastic gradient descent is critical for finding flat minima, *Proc. Natl. Acad. Sci. USA* **118**, e2015617118 (2021).
- [17] X. Xiong, Y.-C. Chen, C. Shi, and P. Ao, On the stochastic gradient descent and inverse variance–flatness relation in artificial neural networks, *Chinese Phys. Lett.* **40**, 080202 (2023).
- [18] N. Yang, C. Tang, and Y. Tu, Stochastic Gradient descent introduces an effective landscape-dependent regularization favoring flat solutions, *Phys. Rev. Lett.* **130**, 237101 (2023).

- [19] S. Adhikari, A. Kabakçioğlu, A. Strang, D. Yuret, and M. Hinczewski, Machine learning in and out of equilibrium, [arXiv:2306.03521](#).
- [20] D. Kunin, J. Sagastuy-Brena, L. Gillespie, E. Margalit, H. Tanaka, S. Ganguli, and D. L. K. Yamins, Limiting dynamics of SGD: Modified loss, phase space Oscillations, and anomalous diffusion, *Neural Computation* **36**, 151 (2024).
- [21] L. Ziyin, K. Liu, T. Mori, and M. Ueda, Strength of minibatch noise in SGD, ICLR (2022), [arXiv:2102.05375](#).
- [22] S. Hochreiter and J. Schmidhuber, Flat minima, *Neural Comput.* **9**, 1 (1997).
- [23] D. Soudry, E. Hoffer, M. S. Nacson, S. Gunasekar, and N. Srebro, The implicit bias of gradient descent on separable data, *J. Mach. Learn. Res.* **19**, 7822 (2018).
- [24] A. K. Lampinen and S. Ganguli, An analytic theory of generalization dynamics and transfer learning in deep linear networks, ICLR (2019), [arXiv:1809.10374](#).
- [25] D. Saad and S. A. Solla, On-line learning in soft committee machines, *Phys. Rev. E* **52**, 4225 (1995).
- [26] P. Riegler and M. Biehl, On-line backpropagation in two-layered neural networks, *J. Phys. A: Math. Gen.* **28**, L507 (1995).
- [27] S. Goldt, M. Advani, A. M. Saxe, F. Krzakala, and L. Zdeborová, Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup, *J. Stat. Mech.* (2020) 124010.
- [28] L. Arnaboldi, L. Stephan, F. Krzakala, and B. Loureiro, From high-dimensional & mean-field dynamics to dimensionless ODEs: A unifying approach to SGD in two-layers networks, [arXiv:2302.05882](#).
- [29] F. Chen, D. Kunin, A. Yamamura, and S. Ganguli, Stochastic collapse: How gradient noise attracts SGD dynamics towards simpler subnetworks, [arXiv:2306.04251](#).
- [30] S. Mandt, M. D. Hoffman, and D. M. Blei, Continuous-time limit of stochastic gradient descent revisited, *Opt. Mach. Learn.* (2015).
- [31] S. Mandt, M. D. Hoffman, and D. M. Blei, A variational analysis of stochastic gradient algorithms, in *Proceedings of the 33rd International Conference on International Conference on Machine Learning, New York, USA* (JMLR.org, 2016), Vol. 48, pp. 354–363.
- [32] Q. Li, C. Tai, and W. E, Stochastic modified equations and adaptive stochastic gradient algorithms, [arXiv:1511.06251](#).
- [33] W. Hu, C. J. Li, L. Li, and J.-G. Liu, On the diffusion approximation of nonconvex stochastic gradient descent, [arXiv:1705.07562](#).
- [34] S. Wojtowytsch, Stochastic gradient descent with noise of machine learning type. Part I: Discrete time analysis, [arXiv:2105.01650](#).
- [35] S. Wojtowytsch, Stochastic gradient descent with noise of machine learning type. Part II: Continuous time analysis, [arXiv:2106.02588](#).
- [36] A. Ali, E. Dobriban, and R. J. Tibshirani, The implicit regularization of stochastic gradient flow for least squares, in *Proceedings of the 37th International Conference on Machine Learning* (JMLR.org, 2020), pp. 233–244.
- [37] Z. Li, T. Wang, and S. Arora, What happens after SGD reaches zero loss? – A mathematical framework, ICLR (2022), [arXiv:2110.06914](#).
- [38] S. Yaida, Fluctuation-dissipation relations for stochastic gradient descent, ICLR (2019), [arXiv:1810.00004](#).
- [39] Z. Li, S. Malladi, and S. Arora, On the validity of modeling SGD with stochastic differential equations (SDEs), NeurIPS (2021), [arXiv:2102.12470](#).
- [40] Y. Zhang, A. M. Saxe, M. Advani, and A. A. Lee, Energy-entropy competition and the effectiveness of stochastic gradient descent in machine learning, *Mol. Phys.* **116**, 3214 (2018).
- [41] S. L. Smith and Q. V. Le, A Bayesian perspective on generalization and stochastic gradient descent, ICLR (2018), [arXiv:1710.06451](#).
- [42] Z. Zhu, J. Wu, B. Yu, L. Wu, and J. Ma, The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects, ICML (2019), [arXiv:1803.00195](#).
- [43] T. H. Nguyen, U. Simsekli, M. Gurbuzbalaban, and G. Richard, First exit time analysis of stochastic gradient descent under heavy-tailed gradient noise, NeurIPS (2019), [arXiv:1906.09069](#).
- [44] Z. Xie, I. Sato, and M. Sugiyama, A diffusion theory for deep learning dynamics: Stochastic gradient descent exponentially favors flat minima, ICLR (2021), [arXiv:2002.03495](#).
- [45] G. Chen, C. K. Qu, and P. Gong, Anomalous diffusion dynamics of learning in deep neural networks. *Neural Netw.* **149**, 18 (2022).
- [46] T. Mori, L. Ziyin, K. Liu, and M. Ueda, Power-law escape rate of SGD, [arXiv:2105.09557](#).
- [47] L. Wu, M. Wang, and W. J. Su, The alignment property of SGD noise and how it helps select flat minima: A stability analysis, NeurIPS (2022), [arXiv:2207.02628](#).
- [48] A. Sclocchi, M. Geiger, and M. Wyart, Dissecting the effects of SGD noise in distinct regimes of deep learning, [arXiv:2301.13703](#).
- [49] H. Daneshmand, J. Kohler, A. Lucchi, and T. Hofmann, Escaping saddles with stochastic gradients, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause (PMLR, 2018), pp. 1155–1164.
- [50] M. Wang and L. Wu, A theoretical analysis of noise geometry in stochastic gradient descent, [arXiv:2310.00692](#).
- [51] Y. Wu, R. Luo, C. Zhang, J. Wang, and Y. Yang, Revisiting the characteristics of stochastic gradient noise and dynamics, [arXiv:2109.09833](#).
- [52] J. Wu, W. Hu, H. Xiong, J. Huan, V. Braverman, and Z. Zhu, On the noisy gradient descent that generalizes as SGD, in *Proceedings of the 37th International Conference on Machine Learning* (JMLR.org, 2020), pp. 10367–10376.
- [53] U. Şimşekli, M. Gürbuzbalaban, T. H. Nguyen, G. Richard, and L. Sagun, On the heavy-tailed theory of stochastic gradient descent for deep neural networks, [arXiv:1912.00018](#).
- [54] U. Simsekli, L. Sagun, and M. Gurbuzbalaban, A tail-index analysis of stochastic gradient noise in deep neural networks, ICML (2019), [arXiv:1901.06053](#).
- [55] Z. Xie, Q.-Y. Tang, Z. He, M. Sun, and P. Li, Rethinking the structure of stochastic gradients: Empirical and statistical evidence, [arXiv:2212.02083](#).
- [56] M. Kühn and B. Rosenow, Correlated noise in epoch-based stochastic gradient descent: Implications for weight variances, [arXiv:2306.05300](#).

- [57] K. Liu, L. Ziyin, and M. Ueda, Noise and fluctuation of finite learning rate stochastic gradient descent, ICML (2021), [arXiv:2012.03636](#).
- [58] V. Thomas, F. Pedregosa, B. van Merriënboer, P.-A. Mangazol, Y. Bengio, and N. L. Roux, On the interplay between noise and curvature and its effect on optimization and generalization, in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, edited by S. Chiappa and R. Calandra (PMLR, 2020), pp. 3503–3513.
- [59] A. Krogh and J. A. Hertz, Generalization in a linear perceptron in the presence of noise, *J. Phys. A: Math. Gen.* **25**, 1135 (1992).
- [60] S. Mei and A. Montanari, The generalization error of random features regression: Precise asymptotics and double descent curve, [arXiv:1908.05355](#).
- [61] J. Lee, L. Xiao, S. S. Schoenholz, Y. Bahri, J. Sohl-Dickstein, R. Novak, and J. Pennington, Wide neural networks of any depth evolve as linear models under gradient descent, NeurIPS (2019), *J. Stat. Mech.* (2020) 124002.
- [62] A. Jacot, F. Gabriel, and C. Hongler, Neural tangent kernel: Convergence and generalization in neural networks, in *Proceedings of the 32nd International Conference on Neural Information Processing Systems, Montréal, Canada* (Curran Associates Inc., New York, 2018), pp. 8580–8589.
- [63] T. Laurent and J. H. von Brecht, Deep linear networks with arbitrary loss: All local minima are global, ICML (2018), [arXiv:1712.01473](#).
- [64] E. M. Achour, F. Malgouyres, and S. Gerchinovitz, The loss landscape of deep linear neural networks: A second-order analysis, [arXiv:2107.13289](#).
- [65] S. Arora, N. Cohen, and E. Hazan, On the optimization of deep networks: Implicit acceleration by overparameterization, ICML (2018), [arXiv:1802.06509](#).
- [66] Z. Ji and M. Telgarsky, Gradient descent aligns the layers of deep linear networks, ICLR (2019), [arXiv:1810.02032](#).
- [67] S. Du and W. Hu, Width provably matters in optimization for deep linear neural networks, in *Proceedings of the 36th International Conference on Machine Learning*, edited by K. Chaudhuri and R. Salakhutdinov, Proceedings of Machine Learning Research, Vol. 97 (PMLR, 2019), pp. 1655–1664.
- [68] G. Gidel, F. Bach, and S. Lacoste-Julien, Implicit regularization of discrete gradient dynamics in linear neural networks, NeurIPS (2019), [arXiv:1904.13262](#).
- [69] S. Tarmoun, G. França, B. D. Haeffele, and R. Vidal, Understanding the dynamics of gradient flow in overparameterized linear models, in *Proceedings of the 38th International Conference on Machine Learning*, edited by M. Meila and T. Zhang, Proceedings of Machine Learning Research Vol. 139 (PMLR, 2021), pp. 10153–10161.
- [70] L. Chizat, M. Colombo, X. Fernández-Real, and A. Figalli, Infinite-width limit of deep linear neural networks, [arXiv:2211.16980](#).
- [71] P. Ao, Potential in stochastic differential equations: Novel construction, *J. Phys. A: Math. Gen.* **37**, L25 (2004).
- [72] C. Kwon, P. Ao, and D. J. Thouless, Structure of stochastic dynamics near fixed points. *Proc. Nat. Acad. Sci. USA* **102**, 13029 (2005).
- [73] P. Ao and L. Yin, Stochastic dynamical structure (SDS) of nonequilibrium processes in the absence of detailed balance. III: Potential function in local stochastic dynamics and in steady state of Boltzmann-Gibbs type distribution function, *J. Phys. A* **39**, 8593 (2006).
- [74] H. Risken, *The Fokker-Planck Equation*, 2nd ed. (Springer, New York, 1989).
- [75] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, 4th ed. (Springer, Berlin, 2009).
- [76] Z. Zhang, H. Zhou, and Z. Xu, Dropout in training neural networks: Flatness of solution and noise structure, [arXiv:2111.0102](#).
- [77] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, Berlin, Heidelberg, 2006).
- [78] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*, 2nd ed. (Springer, New York, 2009).
- [79] This variance is chosen in order to obtain quantities of $\mathcal{O}(1)$ after applying random teacher weights as given in Eq. (2.1).
- [80] R. Schaeffer, M. Khona, Z. Robertson, A. Boopathy, K. Pistunova, J. W. Rocks, I. R. Fiete, and O. Koyejo, Double descent demystified: Identifying, interpreting & ablating the sources of a deep learning puzzle, [arXiv:2303.14151](#).
- [81] Q. Meng, S. Gong, W. Chen, Z.-M. Ma, and T.-Y. Liu, Dynamic of stochastic gradient descent with state-dependent noise, [arXiv:2006.13719](#).
- [82] This relation, which is exact in our setup, is also assumed to hold in general overparameterized deep networks [107].
- [83] X. Chen, J. D. Lee, X. T. Tong, and Y. Zhang, Statistical inference for model parameters in stochastic gradient descent, [arXiv:1610.08637](#).
- [84] It is useful to note that $(\nabla L)(\nabla L)^T = H\bar{w}\bar{w}^T H$, which cancels one such term stemming from the first term in Eq. (2.14).
- [85] J. W. Rocks and P. Mehta, Memorizing without overfitting: Bias, variance, and interpolation in over-parameterized models. *Phys. Rev. Res.* **4**, 013201 (2022).
- [86] S. P. Singh, A. Lucchi, T. Hofmann, and B. Schölkopf, Phenomenology of double descent in finite-width neural networks, [arXiv:2203.07337](#).
- [87] J. Bun, J.-P. Bouchaud, and M. Potters, Cleaning large correlation matrices: Tools from random matrix theory, *Phys. Rep.* **666**, 1 (2017).
- [88] In the overparameterized case ($P < N$), additionally a number of exactly zero eigenvalues appears.
- [89] A. J. Laub, *Matrix Analysis for Scientists and Engineers* (SIAM, Philadelphia, 2005).
- [90] M. Opper, W. Kinzel, J. Kleinz, and R. Nehl, On the ability of the optimal perceptron to generalise, *J. Phys. A: Math. Gen.* **23**, L581 (1990).
- [91] M. Belkin, D. Hsu, S. Ma, and S. Mandal, Reconciling modern machine-learning practice and the classical bias–variance trade-off, *Proc. Natl. Acad. Sci. USA* **116**, 15849 (2019).
- [92] P. Nakkiran, G. Kaplun, Y. Bansal, T. Yang, B. Barak, and I. Sutskever, Deep double descent: Where bigger models and more data hurt, [arXiv:1912.02292](#).
- [93] S. Gunasekar, B. Woodworth, S. Bhojanapalli, B. Neyshabur, and N. Srebro, Implicit regularization in matrix factorization, in *Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, California, USA* (Curran Associates Inc., New York, 2017), pp. 6152–6160.

- [94] The proof presented in Appendix C of [9] should apply for arbitrary N_i and N_o .
- [95] M. Han, J. Park, T. Lee, and J. H. Han, Fluctuation-dissipation-type theorem in stochastic linear learning, *Phys. Rev. E* **104**, 034126 (2021).
- [96] Note that the zero-balanced solution in Eq. (3.8) involves different matrix products of the W_i from the ones in Eq. (3.23) and thus cannot be directly used here.
- [97] K. B. Petersen and M. S. Pedersen, The Matrix Cookbook, available at <http://matrixcookbook.com>.
- [98] This notation means that, e.g., in Eq. (3.28a), all possible combinations of k and l are to be taken, resulting in $N_h N_o$ distinct eigenvalues, while in Eq. (3.28b), each $D_{2,il}$ occurs with multiplicity $N_i - N_h$.
- [99] J.-M. Miao, General expressions for the Moore-Penrose inverse of a 2×2 block matrix, *Linear Algebra Appl.* **151**, 1 (1991).
- [100] L. Chizat, E. Oyallon, and F. Bach, On lazy training in differentiable programming, in *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (Curran Associates Inc., New York, 2019), pp. 2937–2947.
- [101] M. Geiger, L. Petrini, and M. Wyart, Landscape and training regimes in deep learning, *Phys. Rep.* **924**, 1 (2021).
- [102] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, edited by Y. W. Teh and M. Titterton, Proceedings of Machine Learning Research Vol. 9 (PMLR, 2010), pp. 249–256.
- [103] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, Understanding deep learning requires rethinking generalization, [arXiv:1611.03530](https://arxiv.org/abs/1611.03530).
- [104] K. A. Sankararaman, S. De, Z. Xu, W. R. Huang, and T. Goldstein, The impact of neural network overparameterization on gradient confusion and stochastic gradient descent, in *Proceedings of the 37th International Conference on Machine Learning* (JMLR.org, 2020), pp. 8469–8479.
- [105] S. S. Du, X. Zhai, B. Póczos, and A. Singh, Gradient descent provably optimizes over-parameterized neural networks, ICLR (2019), [arXiv:1810.02054](https://arxiv.org/abs/1810.02054).
- [106] S. S. Du, J. D. Lee, H. Li, L. Wang, and X. Zhai, Gradient descent finds global minima of deep neural networks, [arXiv:1811.03804](https://arxiv.org/abs/1811.03804).
- [107] S. Ma, R. Bassily, and M. Belkin, The power of interpolation: Understanding the effectiveness of SGD in modern over-parametrized learning, [arXiv:1712.06559](https://arxiv.org/abs/1712.06559).
- [108] F. Mignacco and P. Urbani, The effective noise of stochastic gradient descent, *J. Stat. Mech.* (2022) 083405.
- [109] L. F. Cugliandolo, The effective temperature, *J. Phys. A: Math. Theor.* **44**, 483001 (2011).
- [110] U. M. B. Marconi, A. Puglisi, L. Rondoni, and A. Vulpiani, Fluctuation-dissipation: Response theory in statistical physics, *Phys. Rep.* **461**, 111 (2008).
- [111] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, Deep learning with differential privacy, in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria* (Association for Computing Machinery, New York, 2016), pp. 308–318.
- [112] M. Refinetti, A. Inghrosso, and S. Goldt, Neural networks trained with SGD learn distributions of increasing complexity, [arXiv:2211.11567](https://arxiv.org/abs/2211.11567).
- [113] J. A. Sirignano and K. Spiliopoulos, Mean field analysis of neural networks: A central limit theorem, *Stochastic Processes Appl.* **130**, 1820 (2020).
- [114] V. D. Bortoli, A. Durmus, X. Fontaine, and U. Simsekli, Quantitative propagation of chaos for SGD in wide neural networks, [arXiv:2007.06352](https://arxiv.org/abs/2007.06352).
- [115] B. Gess, S. Kassing, and V. Konarovskyi, Stochastic modified flows, mean-field limits and dynamics of stochastic gradient descent, [arXiv:2302.07125](https://arxiv.org/abs/2302.07125).
- [116] Y. Cao and Q. Gu, Generalization bounds of stochastic gradient descent for wide and deep neural networks, in *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (Curran Associates Inc., New York, 2019), pp. 10836–10846.
- [117] A. Wei, W. Hu, and J. Steinhardt, More than a toy: Random matrix models predict how real-world neural representations generalize, [arXiv:2203.06176](https://arxiv.org/abs/2203.06176).
- [118] S. Mei, T. Misiakiewicz, and A. Montanari, Generalization error of random features and kernel methods: Hypercontractivity and kernel matrix concentration, [arXiv:2101.10588](https://arxiv.org/abs/2101.10588).
- [119] A. Montanari and Y. Zhong, The interpolation phase transition in neural networks: Memorization and generalization under lazy training, [arXiv:2007.12826](https://arxiv.org/abs/2007.12826).
- [120] A. G. Wilson and P. Izmailov, Bayesian deep learning and a probabilistic perspective of generalization, in *Proceedings of the 34th International Conference on Neural Information Processing Systems, Vancouver, BC, Canada* (Curran Associates Inc., New York, 2020), pp. 4697–4708.
- [121] L. V. Jospin, W. Buntine, F. Boussaid, H. Laga, and M. Bennamoun, Hands-on bayesian neural networks – a tutorial for deep learning users, *IEEE Comput. Intell. Mag.* **17**, 29 (2022).
- [122] D. A. Roberts, S. Yaida, and B. Hanin, *The Principles of Deep Learning Theory: An Effective Theory Approach to Understanding Neural Networks*, new ed. (Cambridge University Press, Cambridge, 2022).
- [123] C. Li, C. Chen, D. Carlson, and L. Carin, Preconditioned stochastic gradient Langevin dynamics for deep neural networks, [arXiv:1512.07666](https://arxiv.org/abs/1512.07666).
- [124] M. Welling and Y. W. Teh, Bayesian learning via stochastic gradient Langevin dynamics, in *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML'11 (Omnipress, Madison, WI, USA, 2011), pp. 681–688.
- [125] Y.-A. Ma, T. Chen, and E. Fox, A complete recipe for stochastic gradient MCMC, in *Advances in Neural Information Processing Systems*, Vol. 28 (Curran Associates, Inc., 2015).
- [126] D. Venturelli and M. Gross, Tracer particle in a confined correlated medium: An adiabatic elimination method, *J. Stat. Mech.* (2022) 123210.