

Analytic theory of dropout regularization

Francesco Mori 

Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford OX1 3PU, United Kingdom

Francesca Mignacco

Graduate Center, City University of New York, New York, New York 10016, USA

and Joseph Henry Laboratories of Physics, Princeton University, Princeton, New Jersey 08544, USA



(Received 30 April 2025; accepted 8 September 2025; published 1 October 2025)

Dropout is a regularization technique widely used in training artificial neural networks to mitigate overfitting. It consists of dynamically deactivating subsets of the network during training to promote more robust representations. Despite its widespread adoption, dropout rates are often selected heuristically, and theoretical explanations of its success remain sparse. Here we analytically study dropout in two-layer neural networks trained with online stochastic gradient descent. In the high-dimensional limit, we derive a set of ordinary differential equations that fully characterize the evolution of the network during training and capture the effects of dropout. We obtain a number of exact results describing the generalization error and the optimal dropout probability at short, intermediate, and long training times. Our analysis shows that dropout reduces detrimental correlations between hidden nodes and mitigates the impact of label noise and that the optimal dropout probability increases with the level of noise in the data. Our results are validated by extensive numerical simulations.

DOI: [10.1103/jmdx-x3gr](https://doi.org/10.1103/jmdx-x3gr)

I. INTRODUCTION

Dropout is a regularization technique introduced in Refs. [1,2] to mitigate overfitting in artificial neural networks. The key idea is to randomly deactivate a subset of neurons at each training step, effectively training an ensemble of subnetworks with shared weights. Since each neuron can be either active or inactive, the total number of subnetworks grows exponentially with the number of hidden neurons. At testing time, all neurons are used, with their weights appropriately rescaled, leading to an implicit averaging of these subnetworks. This averaging reduces overfitting, as the different nodes are more independent due to dropout, and improves generalization across tasks. Additionally, because not all nodes are active simultaneously, dropout is believed to reduce detrimental coadaptations and encourage independent representations [1].

More than a decade after its introduction, dropout remains a widely used technique in training state-of-the-art models [3], with several generalizations proposed. One such variant is DropConnect [4], which applies dropout at the level of individual weights rather than entire nodes. While early studies [1,2] suggested an optimal activation rate between 0.5 and 0.8, constant over training, more recent work has advocated for adaptive schedules [5,6].

Understanding the effect of dropout on learning has motivated a series of theoretical results. Several papers focused on the worst-case scenario and derived generalization bounds [4, 7–9]. The early work [10] formalized the interpretation of dropout as a stochastic regularizer that approximates model averaging in both linear and nonlinear networks. Implicit regularization induced by training with dropout has been studied in generalized linear models [11], two-layer linear networks [12], and two-layer nonlinear networks [13]. In random deep neural networks, Ref. [14] has shown that dropout destroys the order-to-chaos critical point, resulting in an upper bound on trainable neural depth.

In contrast to these prior approaches, in this paper we focus on the typical-case scenario and derive an exact analytic characterization of online feature-learning dynamics in the presence of dropout regularization. To this end, we leverage tools from statistical physics to obtain closed-form equations for the dropout dynamics in two-layer networks [15–18] within a prototypical teacher-student framework [19–21]. This theoretical framework allows us to address several open questions about the mechanisms driving the performance improvement induced by dropout. In particular, we clarify the benefits of dropout in the presence of dataset noise, identify optimal dropout rates, and analyze its impact on the different phases of training.

The remainder of the paper is organized as follows. In Sec. II we introduce the model and provide a summary of our main results. In Sec. III we present a detailed explanation of our derivations and findings. In particular, Sec. III A offers a numerical analysis of the effects of dropout on the learning dynamics. In Sec. III B we provide an analytical characterization of the early-time dynamics. Section III C focuses on the intermediate-time regime, where the network nodes

*Contact author: francesco.mori@physics.ox.ac.uk




FIG. 1. (a) Illustration of the model presented in Sec. II A. (b) Sketch of the generalization error dynamics highlighting the three phases of learning. At early times, the weights remain close to their small-norm initialization. During the plateau (unspecialized phase), the hidden units are symmetric being equally aligned with the teacher units. After specialization, this symmetry is broken.

approximately learn the same features. In Sec. III D we analyze the long-time regime, where the nodes specialize to represent different aspects of the task, and provide an analytical prediction for how dropout affects this specialization transition. Section III E offers further details on the long-time behavior, based on numerical analysis. We conclude in Sec. IV with a summary and perspectives for future work. Some technical details are relegated to the Appendix and to supplementary *Mathematica* notebooks available from [22].

II. MODEL AND SUMMARY OF THE MAIN RESULTS

In this paper we consider supervised learning within the teacher-student framework [19–21] for two-layer neural networks. The teacher-student setup is standard in the statistical physics of neural networks and consists of a student network that is trained on data with labels generated by passing random inputs through a fixed teacher network. In this section we first introduce our model and then we provide a summary of our main results.

A. Model definition

We assume that the input data $x \in \mathbb{R}^N$ are drawn independently as standard Gaussian variables. The corresponding label y is generated via a teacher network $\phi^*(x)$ as

$$y = \phi^*(x) + \sigma \zeta, \quad (1)$$

with label noise $\zeta \sim \mathcal{N}(0, 1)$ while $\sigma > 0$ controls the noise level, and

$$\phi^*(x) = \sum_{m=1}^M v_m^* g^*\left(\frac{w_m^* x}{\sqrt{N}}\right), \quad (2)$$

where $v^* \in \mathbb{R}^M$ and $w_m^* \in \mathbb{R}^N$ denote the fixed teacher weights and g^* denotes the activation function. To learn the dataset, we employ a student network, which after μ training steps takes the form

$$\phi^\mu(x) = \sum_{i=1}^K \sigma_i^\mu v_i^\mu g\left(\frac{w_i^\mu x}{\sqrt{N}}\right), \quad (3)$$

where $v^\mu \in \mathbb{R}^K$ and $w_i^\mu \in \mathbb{R}^N$ are trainable weights. We have introduced the selection variables

$$\sigma_i^\mu = \begin{cases} 1 & \text{if node } i \text{ is active at step } \mu \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

We assume that these variables are independently drawn from a Bernoulli distribution with activation probability $0 < r \leq 1$, i.e., $\sigma_i^\mu \sim \mathcal{B}(r)$. These binary variables model the effect of dropout: At each training step μ , each hidden node is active with probability r or dropped out with probability $1 - r$. A pictorial representation of this model is shown in Fig. 1(a).

We consider training via online stochastic gradient descent (SGD) on the square loss, where each training step involves a single fresh example from the dataset. The update rules are given by

$$w_k^{\mu+1} = w_k^\mu - \frac{\eta}{\sqrt{N}} \sigma_k^\mu v_k^\mu g'(\lambda_k^\mu) \Delta^\mu x^\mu, \quad (5)$$

$$v_k^{\mu+1} = v_k^\mu - \frac{\eta_v}{N} \sigma_k^\mu g(\lambda_k^\mu) \Delta^\mu, \quad (6)$$

where η and η_v are the learning rates for the first and second layers, respectively. Here the updates are evaluated on an independent data point $\{x^\mu, y^\mu\}$ freshly drawn at each step. Note that this setup differs from offline (or batch) learning, where training examples are typically reused multiple times. We have defined the error term $\Delta^\mu = \hat{y}^\mu - y^\mu$ and the pre-activations $\lambda_k^\mu = w_k x^\mu / \sqrt{N}$ for each time step μ . This type of dynamics was previously investigated in Refs. [15, 16, 18], where closed equations for low-dimensional order parameters were derived. Here we would like to extend these results to the case of dropout.

As customary when using dropout, at testing time the full network $\phi_f(x)$ is used, after a rescaling of the rates by a factor $0 < r_f \leq 1$:

$$\phi_f(x) = \sum_{i=1}^K r_f v_i^\mu g\left(\frac{w_i^\mu x}{\sqrt{N}}\right). \quad (7)$$

This rescaling compensates for the reduced activation during training, ensuring consistent output magnitudes between




FIG. 2. Comparison of learning dynamics with and without dropout. (a) Generalization error vs training time $\alpha = \mu/N$ with ($r = 0.68$) and without ($r = 1$) dropout. (b) Observable $\Delta = (Q_{12} - R_{11}R_{21})/\sqrt{Q_{11}Q_{22}}$, which quantifies detrimental correlations between the student's hidden nodes (see the text for details). (c) Cosine similarity between the student's first hidden node weight vector w_1 and the teacher weight vector w_1^* . The parameters are $K = 2$, $M = 1$, $\eta = 1$, $\eta_v = 0$, $\alpha = 5$, and $\sigma = 0.3$. The teacher weights w^* are drawn independently and identically distributed from $\mathcal{N}(0, 1)$; the student weights are initialized independently and identically distributed from $\mathcal{N}(0, 0.1)$, with $N = 50\,000$. The second-layer weights are fixed to one. Solid lines show theoretical predictions from the ODEs (10); crosses mark numerical simulations on a single training trajectory.

training and testing. For simplicity, in the rest of the paper we consider $r_f = r$. We define the generalization error as

$$\epsilon_g = \langle \frac{1}{2} [\phi^*(x) - \phi_f(x)]^2 \rangle, \quad (8)$$

where the angular brackets denote an average over the distribution of the input x .

A remarkable feature of this teacher-student model is that it allows for an exact analytical characterization in the thermodynamic (high-dimensional) limit, where the input dimension N and the number of training steps μ jointly tend to infinity at a fixed ratio $\alpha = \mu/N$. Here α serves as a continuous version of the training step and is interpreted as the training time.

Following Ref. [18], we derive a set of ordinary differential equations (ODEs) that describe the training dynamics. To proceed, we introduce a set of low-dimensional quantities, known as order parameters in statistical physics, that fully characterize the network's performance in the thermodynamic limit. For our model, the order parameters are

$$Q_{ik}^\mu \equiv \frac{w_i^\mu \cdot w_k^\mu}{N}, \quad R_{in}^\mu \equiv \frac{w_i^\mu \cdot w_n^*}{N}, \quad T_{nm} \equiv \frac{w_n^* \cdot w_m^*}{N}, \quad (9)$$

where the centered dot denotes the scalar product and we use the convention that the indices i, j, k , and ℓ refer to the student network (running from 1 to K), while m and n refer to the teacher (running from 1 to M). Unless stated otherwise, we will assume $T_{nm} = \delta_{nm}$. The order parameters provide physical intuition into the learning dynamics. For instance, R_{in} , also known as the magnetization, quantifies the alignment between the i th student node and the n th teacher node.

In the thermodynamic limit $N, \mu \rightarrow \infty$ with $\alpha = \mu/N$ fixed, they satisfy the ODEs

$$\begin{aligned} \frac{dR_{in}}{d\alpha} &= f_{R_{in}}(Q, R, v; \eta, r), & \frac{dQ_{ik}}{d\alpha} &= f_{Q_{ik}}(Q, R, v; \eta, r), \\ \frac{dv_i}{d\alpha} &= f_{v_i}(Q, R, v; \eta, r), \end{aligned} \quad (10)$$

with the definitions and derivations of f_R , f_Q , and f_v presented in the Appendix.

For simplicity, in what follows we take $g^*(\cdot) = g(\cdot) = \text{erf}(\frac{\cdot}{\sqrt{2}})$. For this choice the expressions for the ODEs sim-

plify substantially. In this case, the generalization error can be written as

$$\begin{aligned} \epsilon_g &= \frac{r_f^2}{\pi} \sum_{i,k} v_i v_k \arcsin \left(\frac{Q_{ik}}{\sqrt{1 + Q_{ii}} \sqrt{1 + Q_{kk}}} \right) \\ &+ \frac{1}{\pi} \sum_{n,m} v_n^* v_m^* \arcsin \left(\frac{T_{nm}}{\sqrt{1 + T_{nn}} \sqrt{1 + T_{mm}}} \right) \\ &- \frac{2r_f}{\pi} \sum_{i,n} v_i v_n^* \arcsin \left(\frac{R_{in}}{\sqrt{1 + Q_{ii}} \sqrt{1 + T_{nn}}} \right). \end{aligned} \quad (11)$$

Note that by setting $r = r_f = 1$, thereby removing dropout, we recover the equations of the teacher-student model studied in Ref. [18]. Our first result is thus the extension of these equations to incorporate the effects of dropout.

Our main focus is on the overparametrized regime $K > M$, in which the student network has sufficient capacity to reproduce the teacher network exactly and also represent more complex functions. For simplicity, we do not train the second layer, by fixing $v_i = v_n^* = 1$ and $\eta_v = 0$. In this setting, the teacher and student networks are known as soft committee machines [17].

The results of numerically integrating the ODEs (10) are presented in Fig. 2 for different values of the activation probability r . The solid lines, obtained from the ODEs, show excellent agreement with the numerical simulations (crosses) performed on a single training trajectory with $N = 50\,000$. This confirms that our theoretical predictions, derived in the thermodynamic limit $N \rightarrow \infty$, remain accurate even at finite system size. Interestingly, in Fig. 2(a) we observe that dropout leads to a substantial performance improvement with respect to training without dropout.

B. Summary of the main results

To aid the reader, we provide a brief summary of the main results. Full derivations are presented in the following sections.

(1) *Analytic theory of dropout.* We develop an analytic framework describing the online SGD learning dynamics with




FIG. 3. Optimal activation rate r^* as a function of the variance σ^2 of the label noise. The parameters are $K = 2$, $M = 1$, $\eta = 1$, $\eta_v = 0$, and $\alpha = 0.1$. The teacher weight vector w^* is taken to be of unit norm; the student weights w are initialized to zero. The second-layer weights are fixed to one. The solid line represents the optimal rate obtained by numerical optimization by using the full equations of motion. The crosses correspond to the analytical result in Eq. (24), obtained from a short-time expansion.

$$r^* = \min \left\{ 1; \frac{12M}{3M\eta + 9\eta\sigma^2 + \sqrt{3\eta(M + 3\sigma^2)\{M[32(K - 1) + 3\eta] + 9\eta\sigma^2\}}} \right\} \quad (12)$$

and matches well with the numerical result shown in Fig. 3. Interestingly, r^* decreases with increasing noise variance σ^2 , consistent with the intuition that dropout prevents noise fitting. It also decreases with an increasing number of hidden nodes K , reflecting the role of dropout in mitigating negative coadaptations between hidden units. Furthermore, we obtain the phase diagram in Fig. 4, which identifies the regimes in which dropout improves generalization as a function of the problem parameters.

(4) *Unspecialized phase (analyzed in the low- η limit).* We find that the generalization error quickly reaches a plateau, corresponding to an unspecialized phase in which each student node overlaps equally with all teacher nodes. This unspecialized phase, observed numerically across different values of the learning rate, is known to be unstable in the absence of dropout [17]. The plateau corresponds to a fixed point of the ODEs, which cannot be computed analytically in general. However, in the limit of small learning rate η , it is possible to derive an analytic expression for this fixed point. In this regime, we obtain the exact expression for the generalization error (see also Fig. 5)

$$\epsilon_g^0 = \frac{1}{\pi} \left[\frac{M\pi}{6} + K^2 r^2 \arcsin \left(\frac{MR_0^2}{1 + MR_0^2} \right) - 2KMr \arcsin \left(\frac{R_0}{\sqrt{2}\sqrt{1 + MR_0^2}} \right) \right], \quad (13)$$

where

$$\frac{1}{R_0} = \sqrt{[1 + (K - 1)r]^2 - M + \sqrt{M^2 + (K - 1)r[1 + (K - 1)r]^2[2 + (K - 1)r]}}. \quad (14)$$

The optimal activation rate r_0^* is obtained by minimizing the expression for the generalization error and is shown, together with the corresponding generalization error, in Fig. 6. We observe that the activation probability decreases with increasing K . Note, however, that in the ODEs (10), the dependence on the label noise variance σ^2 disappears at leading order in η . This motivates us to compute the first-order correction in η to determine how the optimal dropout rate depends on σ^2 .

(5) *Unspecialized phase (introducing small- η corrections).* To investigate the role of label noise in determining the optimal activation rate, we derive perturbative corrections to the fixed point to leading order in η . In particular, we show analytically that the optimal activation probability takes the

form

dropout in a prototypical teacher-student model. The dynamics is captured by a set of ODEs, given in Eq. (10), which are exact in the high-dimensional limit $N \rightarrow \infty$. While these ODEs cannot be solved analytically in general, we first examine their numerical integration and then explore specific limits where analytical progress can be made. An illustration of the different learning phases under consideration is provided in Fig. 1(b).

(2) *Numerical analysis of dropout.* As a first step, we integrate Eq. (10) numerically to investigate the effects of dropout on the high-dimensional learning dynamics (see Fig. 2). We observe that, although dropout decreases the cosine similarity between student and teacher nodes, it leads to more independent student representations, resulting in a substantial improvement in performance. In particular, we find that dropout mitigates the negative effects of label noise. Indeed, the optimal activation rate r^* decreases with the variance σ^2 of the label noise (see Fig. 3).

(3) *Early-time dynamics.* At the beginning of training, we assume that the student weights are initialized with a small norm and expand the dynamics to leading order in that norm. The evolution equations can be solved exactly for arbitrary M , K , and r . As a result, we compute the optimal activation rate that minimizes the generalization error

form

$$r^* \approx r_0^* + \eta[a(K, M) + b(K, M)\sigma^2], \quad (15)$$

where $a(K, M)$ and $b(K, M)$ are coefficients we compute explicitly (see Fig. 7). Notably, we find that $b(K, M) < 0$ for




FIG. 4. Early-time phase diagram. Above the lines, obtained for different values of the number K of hidden nodes in the student, it is beneficial to use dropout, with the optimal activation probability given in Eq. (24). Below the lines, it is optimal to avoid dropout, corresponding to activation probability $r = 1$.

all values of K and M considered, confirming that the optimal activation probability decreases with increasing label noise.

(6) *Specialization transition.* As anticipated, in the absence of dropout ($r = 1$), the unspecialized phase is known to be linearly unstable [17], leading to specialization at long times, i.e., each student node aligns preferentially with one of the teacher nodes. Here we investigate how dropout modifies this specialization transition. By performing a linear stability analysis around the plateau fixed point, we derive the phase diagram shown in Fig. 8, valid for $K = M$ and small η . Interestingly, we find that specialization occurs only for sufficiently large values of r (i.e., low dropout). Below a critical threshold in r , the plateau phase becomes stable and the system does not transition to specialization. Finally, for $K = M$ and small η , we find that the optimal activation probability is given by

$$r^* \approx \min\{1; 1 + \eta[a(M) + \sigma^2 b(M)]\}, \quad (16)$$

where $a(M)$ and $b(M)$ can be computed explicitly (see Fig. 9). We find $b(M) < 0$ for all values of M considered, confirming that dropout helps to mitigate label noise even in the specialized phase.

III. RESULTS

In this section we begin by discussing general insights obtained from the numerical integration of the ODEs (10). We then proceed with an analytical study of these equations, deriving explicit results for the optimal activation rate and the corresponding generalization performance across the different phases of the learning dynamics.

A. Numerical analysis of the dropout dynamics

To understand the performance improvement introduced by dropout, it is helpful to examine the evolution of the order parameters over training time (see Fig. 2). For simplicity, we

analyze the case $K = 2$ and $M = 1$, though the following reasoning extends to scenarios where $K > M$.

During training, the weight vectors w_1 and w_2 associated with the two student hidden nodes progressively align with the teacher weight vector w_1^* . However, due to noise (arising from the randomness in the input x , label noise, and random initial conditions), w_1 and w_2 also acquire an additional random component, perpendicular to w_1^* . At any point during training, the weights can be expressed as

$$w_i = R_{i,1} w_1^* + \tilde{w}_i, \quad (17)$$

where $\tilde{w}_i$ represents the noisy component, with $\tilde{w}_i \cdot w_1^* = 0$. Ideally, we seek a large $R_{i,1}$ relative to the norm of $\tilde{w}_i$ and prefer $\tilde{w}_1$ and $\tilde{w}_2$ to be uncorrelated. Indeed, when these noisy components are uncorrelated, their adverse effects on network performance tend to average out. Conversely, if they are correlated, their detrimental effect is compounded. These detrimental correlations can be expressed in terms of the order parameters as

$$\tilde{\Delta} \equiv \frac{\tilde{w}_1 \cdot \tilde{w}_2}{\sqrt{w_1 \cdot w_1 w_2 \cdot w_2}} = \frac{Q_{12} - R_{11}R_{21}}{\sqrt{Q_{11}Q_{22}}}. \quad (18)$$

This quantity is plotted in Fig. 2(b). Interestingly, we observe that dropout reduces these correlations, mitigating the interference between student hidden nodes.

We can therefore identify two competing objectives: (i) learning the signal w_1^* effectively (favored by a low dropout rate) and (ii) avoiding fitting the noise by promoting independence between $\tilde{w}_1$ and $\tilde{w}_2$ (favored by a higher dropout rate).

Figure 2(b) illustrates that, without dropout, the quantity $\tilde{\Delta}$ is significantly larger compared to the case with dropout. This indicates that the two nodes learn similar noisy components ($\tilde{w}_1 \approx \tilde{w}_2$), as they are simultaneously trained on the same data points, exacerbating the fitting of noise. By contrast, with dropout, the nodes are trained jointly only for a fraction of the time and individually otherwise, promoting more independent representations. The readout layer then averages these representations, partially canceling the noise and leading to improved performance.

Interestingly, dropout reduces the alignment of the weights with the signal w_1^* [Fig. 2(c)]. This is expected, as dropout effectively reduces the number of training updates of each individual node. Despite the weaker signal representation, the benefit of increased independence between nodes outweighs this drawback, leading to superior generalization.

A key parameter in our setting is the variance σ^2 of the label noise. When training without dropout, the network is expected to be more prone to fitting this noise, which can degrade performance. In practical applications, dropout is commonly observed to mitigate noise fitting and promote a more robust representation of the data.

First, in Fig. 3 we examine the optimal dropout rate as a function of the label noise variance σ^2 . Here the crosses are obtained by numerical optimization of the full dynamical equations, while the solid line corresponds to the asymptotic result derived analytically in Sec. III B. In the absence of label noise ($\sigma = 0$), the optimal activation rate r^* remains close to unity, indicating minimal dropout. Conversely, as the label noise increases, the optimal strategy shifts toward a higher




FIG. 5. (a) Generalization error at the plateau vs the student's width K at fixed activation rate $r = 0.8$. (b) Generalization error at the plateau vs the activation rate r at fixed student's width $K = 3$. In both panels, the symbols are obtained by numerical integration of Eq. (10) at different values of the learning rate η and with $\sigma = 0$. The black line marks the theoretical prediction in the small- η regime from Eq. (30). The teacher's width is fixed to $M = 2$.




FIG. 6. (a) Optimal activation probability r_0^* in the plateau phase as a function of $K - M$ for different values of M . The curves are obtained via numerical minimization of the generalization error in Eq. (30), derived in the small- η limit. For large K , all curves converge to the universal value $r_0^* = 11/12$. (b) Corresponding value of the shifted generalization error $\epsilon_g - \frac{(\pi-3)M}{6\pi}$ as a function of $K - M$.




FIG. 7. Coefficients (a) $a(K, M)$ and (b) $b(K, M)$ for the plateau phase as a function of $K - M$ for different values of M .




FIG. 8. Specialization phase diagram in the (M, r) plane. In the nonspecialized phase, the plateau fixed point is linearly stable. As r increases, the fixed point becomes unstable, signaling a transition to the specialized phase where student units align with individual teacher units. The solid black line is the theoretical transition line obtained from linear stability analysis in the $\eta \rightarrow 0$ limit. The symbols are obtained by numerical integration of the full ODEs (10) up to time $\alpha = 50\,000$, with $\eta = 0.01$ and $\sigma = 0$.

dropout probability, corresponding to a lower activation probability.

As we will show in the following sections, the fact that the optimal activation rate r^* decreases with increasing label noise holds robustly across the different phases of the learning dynamics.

B. Early-time analysis

A first case in which the equations of motion can be solved exactly is in the early-time regime. In this section we derive the optimal dropout rate in this limit. We assume that the stu-

dent parameters are initialized with zero norm, i.e., $Q_{ii} = 0$, although the calculation can be readily extended to include a finite initial norm. At short times, the overlaps remain close to a symmetric configuration, where each student node is equally aligned with each of the teacher nodes, i.e.,

$$w_i = R \sum_n w_n^* + \tilde{w}_i, \quad (19)$$

where, with a slight abuse of notation, R denotes a positive scalar and $\tilde{w}_i$ is an additional node-dependent random vector, perpendicular to w_n^* for all n . This motivates the ansatz

$$Q_{ik} = Q\delta_{ik} + C(1 - \delta_{ik}), \quad R_{in} = R, \quad (20)$$

where again Q and R denote positive scalar quantities.

At the initial time $\alpha = 0$, the order parameters are $Q = C = R = 0$ and the dynamical equations become, to leading order,

$$\begin{aligned} \frac{dR}{d\alpha} &\approx \frac{\sqrt{2}}{\pi} \eta r, & \frac{dQ}{d\alpha} &\approx \frac{2}{\pi} \eta^2 r \left(\frac{M}{3} + \sigma^2 \right), \\ \frac{dC}{d\alpha} &\approx \frac{2}{\pi} \eta^2 r^2 \left(\frac{M}{3} + \sigma^2 \right). \end{aligned} \quad (21)$$

Therefore, at short times, the order parameters evolve independently as

$$\begin{aligned} R(\alpha) &\approx \frac{\sqrt{2}}{\pi} \eta r \alpha, & Q(\alpha) &\approx \frac{2}{\pi} \eta^2 r \left(\frac{M}{3} + \sigma^2 \right) \alpha, \\ C(\alpha) &\approx \frac{2}{\pi} \eta^2 r^2 \left(\frac{M}{3} + \sigma^2 \right) \alpha. \end{aligned} \quad (22)$$

Plugging these expressions into Eq. (11) and expanding for small α , we find

$$\epsilon_g \approx \frac{2}{\pi^2} \eta^2 K (M/3 + \sigma^2) [(K-1)r^4 + r^3] \alpha - \frac{2}{\pi^2} K M \eta r^2 \alpha. \quad (23)$$

Minimizing the right-hand side with respect to the activation probability r , we find the unique positive solution

$$r^* = \frac{12M}{3M\eta + 9\eta\sigma^2 + \sqrt{3\eta(M + 3\sigma^2)\{M[32(K-1) + 3\eta] + 9\eta\sigma^2\}}}. \quad (24)$$

In Fig. 3 we compare this early-time result (solid line) with the exact result obtained by numerically optimizing the full dynamics with $\alpha = 0.1$ (crosses), finding good agreement. As expected, the activation rate decreases with increasing label variance σ^2 . For large σ^2 , i.e., for $\sigma^2 \gg M \max[1/3, 32(K-1)/9\eta]$, we find $r^* \approx 2M/3\eta\sigma^2$.

Imposing $r^* \leq 1$, we find the additional condition

$$\sigma^2 \geq M \left(\frac{2}{(4K-1)\eta} - \frac{1}{3} \right). \quad (25)$$


Therefore, the optimal activation probability is $r^* = 1$ (no dropout) below a threshold value of the label noise, and it is given by the expression in Eq. (24) otherwise. Note that the

right-hand side of Eq. (25) is positive only if

$$\eta < \frac{6}{4K-1}. \quad (26)$$

These observations lead to the phase diagram in Fig. 4. We observe that dropout is beneficial at high values of the learning rate or the label noise. Interestingly, the region where dropout is useful increases for increasing number of student's hidden nodes, in agreement with the intuition that dropout prevents negative coadaptations.

Note that if we jointly optimize the generalization error in Eq. (23) with respect to both the learning rate and the activation probability, we find $r^* = 1$ and $\eta = M/2K(\sigma^2 + 1/3)$. This implies that, at short times, dropout is not necessary provided the learning rate is properly tuned.

FIG. 9. Coefficients (a) $a(M)$ and (b) $b(M)$ as a function of M for the specialized phase with $K = M$.

C. Intermediate-time analysis: Unspecialized phase

By numerically integrating the ODEs (10), we observe that the generalization error exhibits a plateau at intermediate times. This plateau corresponds to the unspecialized, or symmetric, phase, previously described in Refs. [15–18,23]. This phase is characterized by a symmetry of the student weight vectors w_i , which are equally aligned with all teacher weights. This is in contrast to the specialized phase that emerges at later times, where each student node w_i specializes by aligning predominantly with a specific teacher weight w_n^* or becoming inactive. Note that if the second layer is also trained, more complex specialization patterns can emerge, as the network gains additional flexibility in how it distributes representational roles across the hidden units.

In this section we analytically investigate the role of dropout during the intermediate phase of training. Despite the complexity of the ODEs, exact results can be obtained in the limit of small learning rate $\eta \rightarrow 0$, allowing us to derive expressions for both the optimal dropout rate and the corresponding generalization error. We then compute perturbative corrections around this limit to account for finite learning rates. Finally, we examine how dropout influences the transition to specialization. As we will show, this transition only takes place when the activation probability r exceeds a critical threshold r_c .

1. Small- η regime

In this section we investigate how dropout affects the performance of the network in the unspecialized phase. To do so, we identify fixed-point solutions of the ODEs (A2) and (A3). To make analytic progress, we first consider the small- η limit and neglect terms of order η^2 in the ODEs. Notably, this approximation eliminates the dependence on the label noise variance σ^2 . We will later perform a perturbative expansion in η to recover this dependence to leading order. By setting $f_{R_{in}} = 0$ and $f_{Q_{ik}} = 0$ to leading order in η , we find Eqs. (A16) and (A17). Since we are interested in an unspecialized solution, we assume the student weights take the form

$$w_i = R_0 \sum_n w_n^* \quad (27)$$

for some scalar R_0 to be determined. In other words, for small learning rate η , we expect the student weight vectors to lie entirely within the subspace spanned by the teacher vectors $\{w_n^*\}$, without acquiring an additional component in the orthogonal direction, as in Eq. (19). In other words, in the small- η limit the student does not overfit the noise. This can be verified by checking that the corresponding ansatz

$$Q_{ik} = R_0^2 M, \quad R_{in} = R_0 \quad (28)$$

is a fixed point of the ODEs. By inserting this ansatz into Eqs. (A16) and (A17), we find that it indeed satisfies the fixed-point equations, yielding

$$\frac{1}{R_0} = \sqrt{[1 + (K-1)r]^2 - M} + \sqrt{M^2 + (K-1)r[1 + (K-1)r]^2[2 + (K-1)r]}. \quad (29)$$

Correspondingly, the generalization error, i.e., the height of the plateau, reads

$$\epsilon_g^0 = \frac{1}{\pi} \left[\frac{M\pi}{6} + K^2 r^2 \arcsin\left(\frac{MR_0^2}{1 + MR_0^2}\right) - 2KM r \arcsin\left(\frac{R_0}{\sqrt{2}\sqrt{1 + MR_0^2}}\right) \right]. \quad (30)$$

This expression is shown in Fig. 5 as a function of K and r and is in good agreement with the results of numerically

integrating the full ODEs (10) at small but finite η . We observe that at fixed r and M , the generalization error at the plateau

shows a minimum at a finite value of K . Similarly, at fixed $M = 2$ and $K = 3$, the error is minimized at an intermediate value of the dropout rate, in this case $r^* \approx 0.95$.

In the special case $M = K$ and $r = 1$, we recover the result of Ref. [17],

$$\epsilon_g^0 = \frac{K}{\pi} \left[\frac{\pi}{6} - K \arcsin \left(\frac{1}{2K} \right) \right], \quad (31)$$

which is here generalized to arbitrary M , K , and r . It is also interesting to consider the overparametrized limit $K \gg 1$ at fixed M and r , leading to

$$\epsilon_g^0 \approx \frac{\pi - 3}{6\pi} M + \frac{11 + 12r(r - 2)}{24\pi r^2} \frac{M}{K^2} + O\left(\frac{1}{K^3}\right), \quad (32)$$

which is minimized by the optimal activation probability $r_0^* = 11/12 \approx 0.92$, independently of K and M .

The optimal activation probability r_0^* , obtained by minimizing Eq. (30), is shown in Fig. 6. Interestingly, we observe that r_0^* decreases with increasing K and increases with M . This behavior can be intuitively understood: When the student has more hidden units (larger K), fewer of them need to be active at any given time to capture the relevant features; conversely, a more expressive teacher (larger M) demands a higher number of active units in the student to effectively represent the target function.

It is worth noting that the ansatz in Eq. (27) constrains the student to lie in the M -dimensional subspace spanned by the teacher vectors w_n^* for $n = 1, \dots, M$. As a result, the student's hidden units do not overfit the label noise, which is consistent with the behavior expected in the small- η regime. For finite η , however, each vector w_i acquires an additional component $\tilde{w}_i$, orthogonal to the teacher subspace, as described by Eq. (19), leading to noise fitting.

Although the small- η analysis does not capture this noise-fitting component, we find, somewhat surprisingly, that dropout can still improve performance in this regime, as shown in Fig. 6. However, the optimal activation probability remains close to one ($r^* > 11/12$), and the gain in generalization error is modest. We expect the effect of dropout to be more pronounced in the finite- η regime, where noise-fitting becomes significant and can degrade performance. This observation motivates the perturbative expansion presented in the next section, aimed at computing leading-order corrections in η to the results discussed here.

As a final remark, we note that while choosing a small value of η mitigates noise fitting, it also slows down convergence, resulting in a tradeoff between training speed and precision that we do not explore in detail in this work.

2. Perturbative corrections to the small- η regime

In this section we compute the first-order correction to the small- η results. As discussed above, for finite η the student weight vectors are expected to acquire additional random components due to noise fitting. This motivates the refined ansatz

$$w_i = R \sum_{n=1}^M w_n^* + \tilde{w}_i, \quad (33)$$

where $\tilde{w}_i$ is a random vector orthogonal to the teacher subspace and R is a scalar. The effect of dropout in this regime is to reduce correlations among the noise components $\tilde{w}_i$ for $i = 1, \dots, K$, promoting error cancellation and thereby improving generalization.

The corresponding ansatz for the order parameters is

$$Q_{ik} = Q \delta_{ik} + C (1 - \delta_{ik}), \quad R_{\text{in}} = R. \quad (34)$$

To carry out the perturbative expansion, we expand the order parameters to linear order in η ,

$$Q = Q_0 + \eta Q_1, \quad C = Q_0 + \eta C_1, \quad R = R_0 + \eta R_1, \quad (35)$$

with $Q_0 = R_0^2 M$ and R_0 defined in Eq. (29).

Using *Mathematica*, we insert this ansatz into the full ODEs given in Eqs. (A2) and (A3) and expand to leading order in η . This yields analytical expressions for the first-order corrections Q_1 , C_1 , and R_1 . These expressions are rather lengthy and not particularly illuminating, so we do not report them here. However, we provide a detailed *Mathematica* notebook in [22] that reproduces the derivation and enables the user to compute the explicit forms of the corrections for arbitrary values of K , M , and r .

To leading order in η , the generalization error takes the form

$$\epsilon_g \approx \epsilon_g^0(r) + \eta \epsilon_g^1(r, \sigma), \quad (36)$$

where $\epsilon_g^0(r)$ is given in Eq. (30), and an explicit expression for $\epsilon_g^1(r, \sigma)$ is provided in the notebook in [22]. Recalling that r_0^* denotes the minimizer of $\epsilon_g^0(r)$, we find that the optimal activation probability for small η can be written as

$$r^* \approx r_0^* + r_1^*(\sigma) \eta, \quad (37)$$

where the linear correction is given by

$$r_1^*(\sigma) = -\frac{\partial_r \epsilon_g^1(r_0^*, \sigma)}{\partial_r^2 \epsilon_g^0(r_0^*)} = a(K, M) + b(K, M) \sigma^2, \quad (38)$$

with the coefficients $a(K, M)$ and $b(K, M)$ available in [22]. For instance, in the case $K = 2$ and $M = 1$, we find $r_0^* = 0.931 \dots$, $a(2, 1) = 0.023 \dots$, and $b(2, 1) = -6.445 \dots$. In Fig. 7 we plot the coefficients $a(K, M)$ and $b(K, M)$ for various values of M and K . We observe that both coefficients are negative for most values of K and M , although $a(K, M)$ becomes positive when both K and M are small. This suggests that, in general, increasing the label noise variance σ^2 leads to a decrease in the optimal activation probability r .

D. Onset of specialization

In this section we investigate how the specialization transition is affected by the presence of dropout. Recall that in the absence of dropout ($r = 1$) this transition has been characterized in Refs. [17,24,25], where it was shown that the symmetric unspecialized phase is linearly unstable. Small perturbations, arising either from noise or from the initial conditions, grow over time, ultimately driving the system toward a specialized phase in which each student node aligns predominantly with a distinct teacher node or is deactivated. For example, in the case $K > M$, a typical specialized configuration takes the form $w_i \approx w_i^*$ for $i \leq M$ and $w_i = 0$ for $i > M$.

The precise dependence of this specialization transition on the initialization of the online dynamics has been investigated in [26]. Moreover, changing the scaling of the student readout can affect the rate of convergence toward specialization [27]. Note that the hidden-unit specialization transition has also been studied within the asymptotic Bayes-optimal framework, as discussed, for example, in Refs. [28,29].

Here we analyze how the introduction of dropout modifies this transition to specialization. For simplicity, we focus on the case $K = M$ and the small- η regime. To proceed, we consider the ansatz

$$Q_{ij} = Q\delta_{ij} + C(1 - \delta_{ij}), \quad R_{in} = R\delta_{in} + S(1 - \delta_{in}), \quad (39)$$

where the structure of the R_{in} matrix now allows for specialization. Using this ansatz, the evolution equations for the order parameters then can be written in a simplified form, given in Eqs. (A18)–(A21).

As shown in Sec. III C 1, the plateau phase corresponds to the fixed point

$$Q = C = R_0^2 M, \quad R = S = R_0, \quad (40)$$

where R_0 is given in Eq. (29). To study the stability of this fixed point, we perform a linear stability analysis by introducing small perturbations

$$\begin{aligned} Q &= R_0^2 M + Q_1, & C &= R_0^2 M + C_1, \\ R &= R_0 + R_1, & S &= R_0 + S_1, \end{aligned} \quad (41)$$

where Q_1 , C_1 , R_1 , and S_1 denote small deviations from the fixed-point values.

Using *Mathematica*, we expand the ODEs to linear order in the perturbations and obtain the system of equations

$$\frac{d}{d\alpha} \begin{pmatrix} R_1 \\ S_1 \\ Q_1 \\ C_1 \end{pmatrix} = A(M, r) \begin{pmatrix} R_1 \\ S_1 \\ Q_1 \\ C_1 \end{pmatrix}, \quad (42)$$

where $A(M, r) \in \mathbb{R}^{4 \times 4}$ is a matrix whose entries depend on M and r . The explicit form of $A(M, r)$ is lengthy and provided in the accompanying *Mathematica* notebook in [22].

By inspecting the eigenvalues of $A(M, r)$, we find the phase diagram in Fig. 8, where we identify two phases: a nonspecialized phase where the plateau fixed point is linearly stable and a specialized phase which is reached if the fixed point becomes unstable. The theoretical line corresponds to the critical activation probability r_c at which the real part of the largest eigenvalue of $A(M, r)$ vanishes. It shows excellent agreement with the results obtained by numerically integrating the full ODEs in Eq. (10). Let us also stress that in the specialized phase the student vectors will not perfectly align with the teacher weights; however, they will show a preferential overlap with one of the teacher weights.

E. Long-time behavior

1. Perturbative expansion for small η ($K = M$)

To investigate the long-time behavior, we first focus on the case $K = M$ and the small- η regime. In the limit $\eta \rightarrow 0$, perfect specialization, corresponding to a student perfectly reproducing the teacher with vanishing generalization error,

is obtained if $r = 1$. In this limit, the specialized order parameters are $Q_{ij} = R_{ij} = \delta_{ij}$.

We compute the first-order corrections in η . The ansatz in this case is

$$\begin{aligned} Q_{ij} &= \delta_{ij} + \eta\delta_{ij}Q_1 + \eta(1 - \delta_{ij})C_1, \\ R_{ij} &= \delta_{ij} + \eta\delta_{ij}R_1 + \eta(1 - \delta_{ij})S_1. \end{aligned} \quad (43)$$

One can then perform a perturbative calculation identical to that in Sec. III C 2. We find that in the small- η limit the optimal activation probability is given by

$$r^* \approx \min\{1; 1 + \eta[a(M) + \sigma^2 b(M)]\}. \quad (44)$$

The coefficients $a(M)$ and $b(M)$ are shown in Fig. 9 and can be computed for arbitrary M via the corresponding *Mathematica* notebook in [22]. Once again, we observe that the coefficient $b(M)$ is negative, implying that larger noise σ requires setting lower values of the activation probability.

2. Numerical analysis

As shown in Sec. III D, the long-time behavior depends strongly on the activation probability r . For $r < r_c$, the plateau phase remains stable and the network stays trapped in this phase even at long times. Conversely, specialization occurs at long times when $r > r_c$. As a result, the optimal activation probability r^* depends on both the total training time α and the label noise variance σ^2 . This behavior is summarized in Fig. 10.

For small label noise [Fig. 10(a)], a strategy with little or no dropout ($r \approx 1$) is optimal both before and after the specialization transition. In the opposite case of large noise [Fig. 10(c)], even though a large value of r (e.g., $r = 0.96$) induces specialization, the resulting error is larger than that of an unspecialized configuration achieved with a smaller r (e.g., $r = 0.7$), reflecting the beneficial effects of dropout at higher noise levels.

Interestingly, at intermediate noise levels [Fig. 10(b)], the optimal strategy depends on the total training time α . At intermediate times (before the specialization transition), a lower value of r is preferable. However, for larger α , it is preferable to choose a higher r to allow specialization and achieve a lower final error.

IV. CONCLUSION

In this paper we have studied the effects of dropout regularization on the training dynamics of two-layer neural networks in the teacher-student setting. We have derived exact closed-form equations that describe the training dynamics under online SGD in the high-dimensional limit. By analyzing these equations, we have shown that dropout helps reduce detrimental correlations between the nodes of the network. Furthermore, we have studied the optimal dropout rate and the corresponding generalization error, deriving several analytical results at short, intermediate, and long times. Our findings demonstrate that dropout can significantly mitigate the negative effects of label noise and that the optimal activation probability decreases as the label noise increases. Finally, we have shown analytically that high




FIG. 10. Generalization error as a function of time α for $K = 2$, $M = 1$, $\eta = 0.1$, and different values of node-activation rate $r = 0.7, 0.8, 0.9, 0.96, 1$. Each panel represents a different value of label noise (a) $\sigma = 0.1$, (b) $\sigma = 0.7$, and (c) $\sigma = 1$.

dropout rates prevent the network from reaching a specialized configuration.

Our results rely on the assumption of online training, where a new sample is used at each iteration. An interesting direction for future work is to consider offline training, where the same samples are revisited multiple times. This would enable the study of how dropout affects memorization and overfitting. This analysis could be carried out using dynamical mean-field theory [30,31].

While here we have focused on constant dropout strategies, an interesting open question is how to find optimal dropout schedules that vary the dropout probability over the course of training. Such schedules have been investigated empirically in Refs. [5,6], where they were shown to improve generaliza-

tion performance across various settings. Identifying optimal dropout schedules could be possible using the optimal control framework introduced in Ref. [32]. Such adaptive strategies will be analyzed elsewhere [33].

ACKNOWLEDGMENTS

This work was supported by a Leverhulme Trust International Professorship grant (Award No. LIP-2020-014) and by the Simons Foundation (Award No. 1141576).

DATA AVAILABILITY

The data that support the findings of this article are openly available [22].

APPENDIX: EQUATIONS FOR THE TRAINING DYNAMICS

In this Appendix we present the dynamical equations discussed in the main text. The derivation of these equations follows closely that of Ref. [18], with the difference that an additional binary variable is present to describe dropout. To this end, it is useful to define the notation

$$\mathcal{N}[r, \{i, j, k, \dots, l\}] = r^n, \quad (\text{A1})$$

where $n = |\{i, j, k, \dots, l\}|$ denotes the cardinality of the set $\{i, j, k, \dots, l\}$, i.e., the number of distinct indices in the set. Following [18], we show that, for $N, \mu \rightarrow \infty$ with $\alpha = \mu/N$ fixed,

$$\frac{dR_{\text{in}}}{d\alpha} = f_{R_{\text{in}}} \equiv \eta v_i \left(\sum_{m=1}^M \mathcal{N}[r, \{i\}] v_m^* I_3(i, n, m) - \sum_{j=1}^K \mathcal{N}[r, \{i, j\}] v_j I_3(i, n, j) \right), \quad (\text{A2})$$

$$\begin{aligned} \frac{dQ_{ik}}{d\alpha} = f_{Q_{ik}} \equiv & \eta v_i \left(\sum_{m=1}^M \mathcal{N}[r, \{i\}] v_m^* I_3(i, k, m) - \sum_{j=1}^K \mathcal{N}[r, \{i, j\}] v_j I_3(i, k, j) \right) \\ & + \eta v_k \left(\sum_{m=1}^M \mathcal{N}[r, \{k\}] v_m^* I_3(k, i, m) - \sum_{j=1}^K \mathcal{N}[r, \{k, j\}] v_j I_3(k, i, j) \right) \\ & + \eta^2 v_i v_k \left(\sum_{n=1}^M \sum_{m=1}^M \mathcal{N}[r, \{i, k\}] v_n^* v_m^* I_4(i, k, n, m) - 2 \sum_{j=1}^K \sum_{n=1}^M \mathcal{N}[r, \{i, k, j\}] v_j v_n^* I_4(i, k, j, n) \right. \\ & \left. + \sum_{j=1}^K \sum_{l=1}^K \mathcal{N}[r, \{i, j, k, l\}] v_j v_l I_4(i, k, j, l) + \mathcal{N}[r, \{i, k\}] \sigma^2 J_2(i, k) \right), \end{aligned} \quad (\text{A3})$$

$$\frac{dv_i}{d\alpha} = \eta v_i \left(\sum_{n=1}^M \mathcal{N}[r, \{i\}] v_n^* I_2(i, n) - \sum_{j=1}^K \mathcal{N}[r, \{i, j\}] v_j I_2(i, j) \right), \quad (\text{A4})$$

where

$$J_2 \equiv \frac{2}{\pi} (1 + c_{11} + c_{22} + c_{11}c_{22} - c_{12}^2)^{-1/2}, \quad (\text{A5})$$

$$I_2 \equiv \frac{1}{\pi} \arcsin \left(\frac{c_{12}}{\sqrt{1 + c_{11}} \sqrt{1 + c_{12}}} \right), \quad (\text{A6})$$

$$I_3 \equiv \frac{2}{\pi} \frac{1}{\sqrt{\Lambda_3}} \frac{c_{23}(1 + c_{11}) - c_{12}c_{13}}{1 + c_{11}}, \quad (\text{A7})$$

$$I_4 \equiv \frac{4}{\pi^2} \frac{1}{\sqrt{\Lambda_4}} \arcsin \left(\frac{\Lambda_0}{\sqrt{\Lambda_1 \Lambda_2}} \right) \quad (\text{A8})$$

and

$$\Lambda_4 = (1 + c_{11})(1 + c_{22}) - c_{12}^2, \quad (\text{A9})$$

$$\Lambda_3 = (1 + c_{11}) * (1 + c_{33}) - c_{13}^2, \quad (\text{A10})$$

$$\Lambda_0 = \Lambda_4 c_{34} - c_{23}c_{24}(1 + c_{11}) - c_{13}c_{14}(1 + c_{22}) + c_{12}c_{13}c_{24} + c_{12}c_{14}c_{23}, \quad (\text{A11})$$

$$\Lambda_1 = \Lambda_4(1 + c_{33}) - c_{23}^2(1 + c_{11}) - c_{13}^2(1 + c_{22}) + 2c_{12}c_{13}c_{23}, \quad (\text{A12})$$

$$\Lambda_2 = \Lambda_4(1 + c_{44}) - c_{24}^2(1 + c_{11}) - c_{14}^2(1 + c_{22}) + 2c_{12}c_{14}c_{24}. \quad (\text{A13})$$

Note that we are using the indices i, j, k , and l to indicate the student's nodes, while n and m indicate the teacher's nodes. When using the functions I_2, I_3 , and I_4 , we adopt the compact notation used in Ref. [18]. For instance, $I(i, n)$ takes as input the correlation matrix of the preactivations corresponding to the indices i and n , i.e., $\lambda_i = w_i x / \sqrt{N}$ and $\lambda_n^* = w_n^* x / \sqrt{N}$. For this example, the correlation matrix would be

$$C = \begin{pmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{pmatrix} = \begin{pmatrix} \langle \lambda_i \lambda_i \rangle & \langle \lambda_i \lambda_n^* \rangle \\ \langle \lambda_n^* \lambda_i \rangle & \langle \lambda_n^* \lambda_n^* \rangle \end{pmatrix} = \begin{pmatrix} Q_{ii} & R_{in} \\ R_{in} & T_{nn} \end{pmatrix}. \quad (\text{A14})$$

Finally, the generalization error can be written as

$$\begin{aligned} \epsilon_g &= \frac{r_f^2}{\pi} \sum_{i,k} v_i v_k \arcsin \left(\frac{Q_{ik}}{\sqrt{1 + Q_{ii}} \sqrt{1 + Q_{kk}}} \right) + \frac{1}{\pi} \sum_{n,m} v_n^* v_m^* \arcsin \left(\frac{T_{nm}}{\sqrt{1 + T_{nn}} \sqrt{1 + T_{mm}}} \right) \\ &\quad - \frac{2r_f}{\pi} \sum_{i,n} v_i v_n^* \arcsin \left(\frac{R_{in}}{\sqrt{1 + Q_{ii}} \sqrt{1 + T_{nn}}} \right). \end{aligned} \quad (\text{A15})$$

We report here the expression of the fixed-point equations $f_R = 0$ and $f_Q = 0$ in the small- η limit. These can be obtained by neglecting the terms in η^2 in Eqs. (A2) and (A3), yielding

$$\sum_{m=1}^M \mathcal{N}[r, \{i\}] I_3(i, n, m) - \sum_{j=1}^K \mathcal{N}[r, \{i, j\}] I_3(i, n, j) = 0, \quad (\text{A16})$$

$$\begin{aligned} &\sum_{m=1}^M \mathcal{N}[r, \{i\}] I_3(i, k, m) - \sum_{j=1}^K \mathcal{N}[r, \{i, j\}] I_3(i, k, j) \\ &\quad + \sum_{m=1}^M \mathcal{N}[r, \{k\}] I_3(k, i, m) - \sum_{j=1}^K \mathcal{N}[r, \{k, j\}] I_3(k, i, j) = 0. \end{aligned} \quad (\text{A17})$$

Additionally, assuming the specialization ansatz in Eq. (39), the ODEs can be written, in the case $K = M \geq 2$ and in the small- η limit, as

$$\frac{dR}{d\alpha} = \eta r I_3 \begin{pmatrix} Q & R & R \\ R & 1 & 1 \\ R & 1 & 1 \end{pmatrix} + \eta r (M - 1) I_3 \begin{pmatrix} Q & R & S \\ R & 1 & 0 \\ S & 0 & 1 \end{pmatrix} - \eta r I_3 \begin{pmatrix} Q & R & Q \\ R & 1 & R \\ Q & R & Q \end{pmatrix} - \eta r^2 (M - 1) I_3 \begin{pmatrix} Q & R & C \\ R & 1 & S \\ C & S & Q \end{pmatrix}, \quad (\text{A18})$$

$$\begin{aligned} \frac{dS}{d\alpha} = & \eta r I_3 \begin{pmatrix} Q & S & R \\ S & 1 & 0 \\ R & 0 & 1 \end{pmatrix} + \eta r I_3 \begin{pmatrix} Q & S & S \\ S & 1 & 1 \\ S & 1 & 1 \end{pmatrix} + \eta(M-2)r I_3 \begin{pmatrix} Q & S & S \\ S & 1 & 0 \\ S & 0 & 1 \end{pmatrix} \\ & - \eta r I_3 \begin{pmatrix} Q & S & Q \\ S & 1 & S \\ Q & S & Q \end{pmatrix} - \eta r^2 I_3 \begin{pmatrix} Q & S & C \\ S & 1 & R \\ C & R & Q \end{pmatrix} - \eta(M-2)r^2 I_3 \begin{pmatrix} Q & S & C \\ S & 1 & S \\ C & S & Q \end{pmatrix}, \end{aligned} \quad (A19)$$

$$\begin{aligned} \frac{dQ}{d\alpha} = & 2\eta r I_3 \begin{pmatrix} Q & Q & R \\ Q & Q & R \\ R & R & 1 \end{pmatrix} + 2\eta r(M-1) I_3 \begin{pmatrix} Q & Q & S \\ Q & Q & S \\ S & S & 1 \end{pmatrix} \\ & - 2\eta r I_3 \begin{pmatrix} Q & Q & Q \\ Q & Q & Q \\ Q & Q & Q \end{pmatrix} - 2\eta r^2(M-1) I_3 \begin{pmatrix} Q & Q & C \\ Q & Q & C \\ C & C & Q \end{pmatrix}, \end{aligned} \quad (A20)$$

$$\begin{aligned} \frac{dC}{d\alpha} = & 2\eta r I_3 \begin{pmatrix} Q & C & R \\ C & Q & S \\ R & S & 1 \end{pmatrix} + 2\eta r I_3 \begin{pmatrix} Q & C & S \\ C & Q & R \\ S & R & 1 \end{pmatrix} + 2\eta(M-2)r I_3 \begin{pmatrix} Q & C & S \\ C & Q & S \\ S & S & 1 \end{pmatrix} \\ & - \eta r(r+1) I_3 \begin{pmatrix} Q & C & Q \\ C & Q & C \\ Q & C & Q \end{pmatrix} - \eta r(r+1) I_3 \begin{pmatrix} Q & C & C \\ C & Q & Q \\ C & Q & Q \end{pmatrix} - 2\eta(M-2)r^2 I_3 \begin{pmatrix} Q & C & C \\ C & Q & C \\ C & C & Q \end{pmatrix}. \end{aligned} \quad (A21)$$

-
- [1] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov, Improving neural networks by preventing co-adaptation of feature detectors, [arXiv:1207.0580](https://arxiv.org/abs/1207.0580).
- [2] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [3] I. Salehin and D.-K. Kang, A review on dropout regularization approaches for deep neural networks within the scholarly domain, *Electronics* **12**, 3106 (2023).
- [4] L. Wan, M. Zeiler, S. Zhang, Y. Le Cun, and R. Fergus, Regularization of neural networks using dropconnect, in *Proceedings of the 30th International Conference on Machine Learning* (PMLR, 2013), pp. 1058–1066.
- [5] P. Morerio, J. Cavazza, R. Volpi, R. Vidal, and V. Murino, in *Proceedings of the IEEE International Conference on Computer Vision* (IEEE, Piscataway, 2017), pp. 3544–3552.
- [6] Z. Liu, Z. Xu, J. Jin, Z. Shen, and T. Darrell, in *Proceedings of the 40th International Conference on Machine Learning* (PMLR, Cambridge, 2023), pp. 22233–22248.
- [7] W. Mou, Y. Zhou, J. Gao, and L. Wang, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause (PMLR, Cambridge, 2018), Vol. 80, pp. 3645–3653.
- [8] R. Arora, P. Bartlett, P. Mianjy, and N. Srebro, in *Proceedings of the 38th International Conference on Machine Learning* (PMLR, Cambridge, 2021), pp. 351–361.
- [9] K. Zhai and H. Wang, in *Proceedings of the Sixth International Conference on Learning Representations, Vancouver, 2018* (ICLR, La Jolla, 2018).
- [10] P. Baldi and P. J. Sadowski, in *Advances in Neural Information Processing Systems 26*, edited by C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (Curran, Red Hook, 2013).
- [11] S. Wager, S. Wang, and P. Liang, Dropout training as adaptive regularization, in *Advances in Neural Information Processing Systems 26* (Curran Associates, Red Hook, 2013), pp. 351–359.
- [12] P. Mianjy, R. Arora, and R. Vidal, On the implicit bias of dropout, in *Proceedings of the 35th International Conference on Machine Learning* (PMLR, 2018), Vol. 80, pp. 3540–3548.
- [13] Z. Zhang and Z.-Q. J. Xu, Implicit regularization of dropout, *IEEE Trans. Pattern Anal. Mach. Intell.* **46**, 4206 (2024).
- [14] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein, in *Proceedings of the Fifth International Conference on Learning Representations, Toulon, 2017* (ICLR, La Jolla, 2017).
- [15] M. Biehl and H. Schwarze, Learning by on-line gradient descent, *J. Phys. A: Math. Gen.* **28**, 643 (1995).
- [16] D. Saad and S. A. Solla, Exact solution for on-line learning in multilayer neural networks, *Phys. Rev. Lett.* **74**, 4337 (1995).
- [17] D. Saad and S. A. Solla, On-line learning in soft committee machines, *Phys. Rev. E* **52**, 4225 (1995).
- [18] S. Goldt, M. Advani, A. M. Saxe, F. Krzakala, and L. Zdeborová, in *Advances in Neural Information Processing Systems 32, Vancouver, 2019*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran, Red Hook, 2019).
- [19] E. Gardner and B. Derrida, Three unfinished works on the optimal storage capacity of networks, *J. Phys. A: Math. Gen.* **22**, 1983 (1989).
- [20] H. S. Seung, H. Sompolinsky, and N. Tishby, Statistical mechanics of learning from examples, *Phys. Rev. A* **45**, 6056 (1992).
- [21] A. Engel, *Statistical Mechanics of Learning* (Cambridge University Press, Cambridge, 2001).
- [22] For full details on the calculations see the *Mathematica* notebooks at https://github.com/francescomori/analytic_dropout.git (2025).

- [23] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová, Modeling the influence of data structure on learning in neural networks: The hidden manifold model, *Phys. Rev. X* **10**, 041044 (2020).
- [24] M. Biehl, P. Riegler, and C. Wöhler, Transient dynamics of on-line learning in two-layered neural networks, *J. Phys. A: Math. Gen.* **29**, 4769 (1996).
- [25] M. Straat and M. Biehl, On-line learning dynamics of ReLu neural networks using statistical physics techniques, [arXiv:1903.07378](https://arxiv.org/abs/1903.07378).
- [26] D. Jarvis, S. Lee, C. C. J. Dominé, A. M. Saxe, and S. S. Mannelli, in *The 13th International Conference on Learning Representations, Singapore, 2025* (ICLR, La Jolla, 2025).
- [27] F. Richert, R. Worschech, and B. Rosenow, Soft mode in the dynamics of over-realizable online learning for soft committee machines, *Phys. Rev. E* **105**, L052302 (2022).
- [28] E. Oostwal, M. Straat, and M. Biehl, Hidden unit specialization in layered neural networks: ReLu vs. sigmoidal activation, *Physica A* **564**, 125517 (2021).
- [29] B. Aubin, A. Maillard, J. Barbier, F. Krzakala, N. Macris, L. Zdeborová, in *Advances in Neural Information Processing Systems 31, Montreal, 2018*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran, Red Hook, 2018).
- [30] E. Agoritsas, G. Biroli, P. Urbani, and F. Zamponi, Out-of-equilibrium dynamical mean-field equations for the perceptron model, *J. Phys. A: Math. Theor.* **51**, 085002 (2018).
- [31] F. Mignacco, F. Krzakala, P. Urbani, and L. Zdeborová, in *Advances in Neural Information Processing Systems 33, 2020*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M.-F. Balcan, and H.-T. Lin (Curran, Red Hook, 2020), p. 9540.
- [32] F. Mori, S. S. Mannelli, and F. Mignacco, Optimal protocols for continual learning via statistical physics and control theory, *J. Stat. Mech.* (2025) 084004.
- [33] F. Mignacco and F. Mori, A statistical physics framework for optimal learning, [arXiv:2507.07907](https://arxiv.org/abs/2507.07907).