

Contrasting random and learned features in deep Bayesian linear regressionJacob A. Zavatore-Veth ^{1,2,*} William L. Tong ^{3,†} and Cengiz Pehlevan ^{3,2,‡}¹*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*²*Center for Brain Science, Harvard University, Cambridge, Massachusetts 02138, USA*³*John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*

(Received 13 March 2022; revised 11 May 2022; accepted 26 May 2022; published 16 June 2022)

Understanding how feature learning affects generalization is among the foremost goals of modern deep learning theory. Here, we study how the ability to learn representations affects the generalization performance of a simple class of models: deep Bayesian linear neural networks trained on unstructured Gaussian data. By comparing deep random feature models to deep networks in which all layers are trained, we provide a detailed characterization of the interplay between width, depth, data density, and prior mismatch. We show that both models display samplewise double-descent behavior in the presence of label noise. Random feature models can also display modelwise double descent if there are narrow bottleneck layers, while deep networks do not show these divergences. Random feature models can have particular widths that are optimal for generalization at a given data density, while making neural networks as wide or as narrow as possible is always optimal. Moreover, we show that the leading-order correction to the kernel-limit learning curve cannot distinguish between random feature models and deep networks in which all layers are trained. Taken together, our findings begin to elucidate how architectural details affect generalization performance in this simple class of deep regression models.

DOI: [10.1103/PhysRevE.105.064118](https://doi.org/10.1103/PhysRevE.105.064118)**I. INTRODUCTION**

Deep neural networks (NNs) display a rich and often-perplexing spectrum of generalization behaviors. Highly overparametrized NNs may possess the expressivity to fit random noise, yet in practice can still generalize well to unseen data [1,2]. The ability of NNs to flexibly learn features from data is widely believed to be a critical contributor to their practical success [1–4], but the precise contributions of feature learning to their generalization behavior remain incompletely understood [1–10].

In recent years, intensive theoretical work has begun to elucidate the properties of deep networks in the limit of infinite hidden layer width. In this limit, a dramatic simplification occurs, and inference in deep networks is equivalent to kernel regression or classification [6,11–16]. This correspondence has enabled detailed characterizations of inference at infinite width in both maximum-likelihood and fully Bayesian settings, providing new insights into the inductive biases that allow deep networks to overfit benignly [17–27]. Yet, understanding inference in the kernel limit is not sufficient, because kernel descriptions cannot capture feature learning [3,7–9,28].

As a result, a growing number of recent works have aimed to study the behavior of networks near the kernel limit, with the hope that leading-order corrections to the large-width behavior might elucidate how width and depth affect infer-

ence [4,29–41]. Some of these works focus on the properties of the function-space prior distribution [29–34], some consider maximum-likelihood inference with gradient descent [32,40,41], and some consider properties of the full Bayes posterior [4,30,32,35–39]. This body of research has resulted in several conjectural conditions under which narrower and deeper networks might perform better than their infinitely wide cousins in the Bayesian setting, as measured by generalization for fixed data [35,37,38] or by some alternative criterion based on entropic considerations [32].

However, previous studies of Bayesian neural network (BNN) generalization near the kernel limit have not clearly differentiated the effect of width on feature learning from its other potential effects on inference. Concretely, it is not clear whether potential improvements in generalization afforded by the leading finite-width correction reflect the benefits of feature learning, or if a similar gain would be observed in random feature (RF) models, where only the readout layer is trained. Here, we explore how random and learned features affect generalization in the simplest class of Bayesian NNs—deep linear models—when trained on unstructured, noisy data. By developing a detailed understanding of this simple setting, one might hope to gain intuition that may prove useful in studying more complex networks [31,42–49].

In this paper, we study the asymptotic generalization performance of deep linear Bayesian regression for data generated with an isotropic Gaussian covariate model. Using the replica trick [50,51], we compute learning curves for simple linear regression (LR), deep linear Gaussian RF models, and deep linear NNs. Our results are obtained using an isotropic Gaussian likelihood in the limit of small likelihood

*jzavatoreveth@g.harvard.edu

†wtong@g.harvard.edu

‡cpehlevan@seas.harvard.edu

variance, which renders this analysis analytically tractable [35,36]. Using alternative replica-free methods and numerical simulation, we show that the predictions obtained under a replica-symmetric (RS) *Ansatz* are accurate for all three model classes. In particular, the RS result for learning curves of NNs with hidden layers of equal widths is consistent with results obtained by Li and Sompolinsky [37] using a different approximation method.

In the presence of label noise, both RF and NN models display samplewise nonmonotonicity in their learning curves. As we work in a high-dimensional limit, this nonmonotonicity is of a particularly extreme form: the generalization error diverges at a particular data density. In keeping with modern deep learning parlance, we refer to this behavior as “double descent,” though this monotonicity can arise from distinct effects in different settings [2,5,10,17,20,22–25,44,45,52]. If one introduces a bottleneck layer that is narrower than the input dimension, an RF model will display modelwise double-descent behavior at fixed data density—or equivalently samplewise double descent at fixed width—even in the absence of label noise, while an NN model will not show this divergence. This distinct small-width behavior shows one advantage afforded by the flexibility to learn features. For both models, we analyze how optimal network architecture depends on data density and prior mismatch. We show that, at a given data density, RF models have a particular optimal width for fixed depth and optimal depth for fixed width that minimizes the generalization error. In contrast, it is always optimal to take an NN to be as wide or as narrow as possible, depending on the regime.

We further analyze models of arbitrary depth perturbatively in the limit in which the network depth and dataset size are small relative to the hidden layer widths, connecting these results to those of previous work on fixed-dataset perturbation theory [35]. We find that the leading-order correction to the large-width behavior of RF and NN models is identical, hence first-order perturbation theory for the generalization error cannot distinguish between random and learned features. To distinguish between training only the readout layer and training all layers, one must go to second order in perturbation theory. Therefore, at large widths, the ability to perform representation learning provides only a small advantage in generalization performance in these simple models relative to random features, which is invisible in first-order perturbation theory. In total, our results provide insight into how the generalization behavior of deep Bayesian linear regression in high dimensions depends on architectural details. Moreover, they shed light onto which qualitative features of generalization behavior can or cannot be captured by low-order perturbative corrections [31].

II. PROBLEM SETTING

In this section, we introduce the three classes of regression models we consider in this paper, as well as our generative data model. Our notation throughout is standard; we use $\|\cdot\|$ to denote the Euclidean norm, \mathbf{I}_d to denote the $d \times d$ identity matrix, and $\mathbf{1}$ to denote the vector with all elements equal to 1.

A. Regression models and parameter priors

In this paper, we consider three classes of scalar Bayesian linear regression models for a scalar-valued function of d -dimensional inputs. All three of these model classes are of the form

$$g_{\mathbf{w}}(\mathbf{x}) = \frac{1}{\sqrt{d}} \mathbf{w}^\top \mathbf{x}, \quad (1)$$

and differ in the parametrization of the “end-to-end” weight vector $\mathbf{w} \in \mathbb{R}^d$. We will choose parameter priors such that $\mathbb{E} \|\mathbf{w}\|^2 = \sigma^2 d$ for each model, where $\sigma > 0$ is a hyperparameter which sets the prior variance of the network outputs. We remark that each model is positive-homogeneous in its parameters, hence this choice is made without loss of generality. In all cases, the parameter priors are isotropic and Gaussian, as is standard in Bayesian deep learning [11–14,31,47–49].

Below, we list the three classes of models we consider, and introduce a two-letter abbreviation for each.

(1) LR: Simple Bayesian linear regression. For this model, the end-to-end weight vector is directly parametrized as

$$\mathbf{w}_{\text{LR}} = \sigma \mathbf{v} \quad (2)$$

for a trainable parameter vector $\mathbf{v} \in \mathbb{R}^d$ with isotropic Gaussian prior distribution

$$\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d). \quad (3)$$

Previous works have extensively studied this model in both maximum-likelihood and fully Bayesian settings [20,21,44,45,52,53], hence we include it as a baseline against which we will compare our results for more complicated models.

(2) RF: Deep Bayesian random feature models. For these models, the weight vector is parametrized as

$$\mathbf{w}_{\text{RF}} = \frac{\sigma}{\sqrt{n_1 \cdots n_\ell}} \mathbf{U}_1 \cdots \mathbf{U}_\ell \mathbf{v} \quad (4)$$

for matrices $\mathbf{U}_1 \in \mathbb{R}^{d \times n_1}$, $\mathbf{U}_2 \in \mathbb{R}^{n_1 \times n_2}, \dots$, and $\mathbf{U}_\ell \in \mathbb{R}^{n_{\ell-1} \times n_\ell}$ and a vector $\mathbf{v} \in \mathbb{R}^{n_\ell}$. Here, $\ell \in \mathbb{N}_{>0}$ is the network depth, while $n_1, \dots, n_\ell \in \mathbb{N}_{>0}$ are the hidden layer widths. For the RF model, only the readout weight vector \mathbf{v} is trainable, while the hidden layer weights \mathbf{U}_l are fixed and random. We choose an isotropic Gaussian prior for the readout weights,

$$\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{n_\ell}), \quad (5)$$

while the hidden layer weights are drawn from a fixed isotropic Gaussian distribution:

$$(\mathbf{U}_l)_{ij} \sim \mathcal{N}(0, 1) \quad (l = 1, \dots, \ell). \quad (6)$$

(3) NN: Deep Bayesian linear neural networks. For these models, the weight vector is parametrized as

$$\mathbf{w}_{\text{NN}} = \frac{\sigma}{\sqrt{n_1 \cdots n_\ell}} \mathbf{U}_1 \cdots \mathbf{U}_\ell \mathbf{v}. \quad (7)$$

Though NNs are parametrized identically to the RF models above, they differ in that all of the weights are trainable, not only the readout. We again choose isotropic Gaussian prior distributions:

$$(\mathbf{U}_l)_{ij} \sim \mathcal{N}(0, 1) \quad (l = 1, \dots, \ell), \quad (8)$$

$$\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{n_\ell}). \quad (9)$$

From a physical perspective, the hidden layer weights in the RF model are “quenched” disorder, whereas they are “annealed” disorder in NNs [50,51].

For all models, we denote expectation with respect to the prior distribution of the trainable parameters by $\mathbb{E}_{\mathcal{W}}$.

B. Data model and the Bayes posterior

We train all models on a dataset $\{(\mathbf{x}_\mu, y_\mu)\}_{\mu=1}^p$ of p examples, generated according to a standard isotropic Gaussian covariate model [21,44–46,52,53]. In this model, the example inputs are independent and identically distributed samples from a standard Gaussian distribution,

$$\mathbf{x}_\mu \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d), \quad (10)$$

while the labels are generated by a ground truth linear model, possibly corrupted by additive Gaussian noise:

$$y_\mu = \frac{1}{\sqrt{d}} \mathbf{w}_*^\top \mathbf{x}_\mu + \eta \xi_\mu, \quad (11)$$

where $\eta \geq 0$ sets the noise variance. The noise variables are independent and identically distributed as

$$\xi_\mu \sim \mathcal{N}(0, 1), \quad (12)$$

and are independent of the inputs. We take the “teacher” weight vector \mathbf{w}_* to have fixed norm $\|\mathbf{w}_*\|^2 = d$. In some places, we will average over teacher weights distributed uniformly on the sphere, i.e., $\mathbf{w}_* \sim \mathcal{U}[S^{d-1}(\sqrt{d})]$, though our main results will hold pointwise for any \mathbf{w}_* on the sphere. We will collect the training inputs and outputs into a matrix $(\mathbf{X})_{\mu,j} = (\mathbf{x}_\mu)_j$ and a vector $(\mathbf{y})_\mu = y_\mu$, respectively.

For a dataset thusly generated, we introduce an isotropic Gaussian likelihood of variance $1/\beta$:

$$p[\{(\mathbf{x}_\mu, y_\mu)\}_{\mu=1}^p | \mathcal{W}] \propto \exp\left(-\frac{\beta}{2} \sum_{\mu=1}^p [g_{\mathbf{w}}(\mathbf{x}_\mu) - y_\mu]^2\right), \quad (13)$$

where \mathcal{W} denotes the set of trainable parameters for a given model, and the normalization constant is implied. We will refer to β as the “inverse temperature” by standard analogy with statistical mechanics [35,51,52,54–56]. Then, the partition function of the resulting Bayes posterior is given as

$$Z = \mathbb{E}_{\mathcal{W}} \exp\left(-\frac{\beta}{2} \sum_{\mu=1}^p [g_{\mathbf{w}}(\mathbf{x}_\mu) - y_\mu]^2\right). \quad (14)$$

We denote expectations with respect to this Bayes posterior by $\langle \cdot \rangle$.

C. Generalization error in the thermodynamic limit

With the initial setup of the previous sections, we can now introduce our concrete objective. We consider a proportional asymptotic limit in which the input dimension d , the dataset size p , and (for NN and RF models) the hidden layer widths n_1, \dots, n_ℓ tend to infinity for fixed depth ℓ and fixed ratios

$$\alpha \equiv p/d = \mathcal{O}(1), \quad (15)$$

$$\gamma_l \equiv n_l/d = \mathcal{O}(1) \quad (l = 1, \dots, \ell). \quad (16)$$

Moreover, we focus on the zero-temperature limit $\beta \rightarrow \infty$, in which the likelihood tends to a constraint that the network interpolates the training set with probability 1. In the noise-free case, this limiting likelihood is matched to the true generative model of the data, but it is clearly mismatched in the presence of label noise. This limit has been considered in several recent studies of deep linear Bayesian neural networks [4,32,35–37].

Our goal is to study the average-case generalization error ϵ of the resulting model, as measured by the deviation of its end-to-end weight vector \mathbf{w} from the true teacher weight vector \mathbf{w}_* :

$$\epsilon = \lim_{\beta \rightarrow \infty} \lim_{d, p, n_1, \dots, n_\ell \rightarrow \infty} \mathbb{E}_{\mathcal{D}} \left\langle \frac{1}{d} \|\mathbf{w} - \mathbf{w}_*\|^2 \right\rangle. \quad (17)$$

Here, $\mathbb{E}_{\mathcal{D}}$ denotes expectation with respect to all quenched disorder for a given model. For all models, this includes the training inputs and label noise; for the RF model, it also includes the hidden layer weights. We emphasize that this choice means that we do not ensemble RF model predictions over realizations of the features; rather, we consider the average-case generalization error of individual realizations [17,22–25].

We remark that Eq. (17) is the average-case error of the Gibbs estimator, i.e., a single sample from the posterior; one could instead consider the error of the mean estimator $\langle \mathbf{w} \rangle$. For the LR and RF models, this corresponds to studying Bayesian minimum mean-squared error (MMSE) inference [21,53]. As one has the thermal bias-variance decomposition

$$\langle \|\mathbf{w} - \mathbf{w}_*\|^2 \rangle = \|\langle \mathbf{w} \rangle - \mathbf{w}_*\|^2 + \text{tr}(\langle \mathbf{w} \mathbf{w}^\top \rangle - \langle \mathbf{w} \rangle \langle \mathbf{w} \rangle^\top), \quad (18)$$

our results include an additional contribution to the generalization error from the posterior covariance $\langle \mathbf{w} \mathbf{w}^\top \rangle - \langle \mathbf{w} \rangle \langle \mathbf{w} \rangle^\top$ of the end-to-end weight vector, which is not identically zero. If one considered an alternative low-temperature limit in which the prior variance is proportional to $1/\beta$, then this additional contribution would vanish in the low-temperature limit. Our choice of scaling is motivated by the considerations described in our previous work [35], and is the one classically used in studies of the statistical mechanics of Bayesian inference [37,54–56]. This choice is important as it affects the relationship of our results to those in the setting of ridge regression. As discussed in Appendices C and D, in our previous work [35], and in previous works of Advani and Ganguli [53] and Barbier *et al.* [21], the zero-temperature limit of the MMSE estimator would in this case coincide with the ridge regression estimator.

We compute the limiting average generalization error using the replica method, a nonrigorous but powerful heuristic that has seen broad use in statistical mechanical studies of inference [20,25,27,50,51]. As our main results can be understood independently of calculation through which they were obtained, we relegate the details to Appendices A and B. We note the important caveat that our main results are obtained under a replica-symmetric *Ansatz*. We expect this assumption to hold exactly for the LR and RF models by virtue of the concavity of their log posteriors, but replica symmetry may be

broken in deep linear NNs [50,57]. We will not address this possibility analytically by considering *Ansätze* with broken replica symmetry [50], but will instead simply compare the RS predictions against results obtained through a combination of alternative analytical methods and numerics.

III. LEARNING CURVES FOR THE LR MODEL

We begin by briefly describing the learning curve of the simple LR model. Our result extends the classic result of Krogh and Hertz [52] for ridge regression in the ridgeless limit to the Bayesian setting:

$$\epsilon_{\text{LR}} = \begin{cases} (1 + \sigma^2)(1 - \alpha) + \frac{\alpha}{1-\alpha}\eta^2, & \text{if } \alpha < 1 \\ \frac{1}{\alpha-1}\eta^2, & \text{if } \alpha > 1. \end{cases} \quad (19)$$

For this simple model, the learning curve can also be computed directly by first evaluating the posterior average defining ϵ_{LR} for a fixed realization of the disorder, and then averaging the result over the disorder in the zero-temperature limit (see Appendix C for details). The result of [52] can be recovered from Eq. (19) by setting $\sigma = 0$. We provide further discussion of the relationship between the Bayesian LR model in the zero-temperature limit and ridge regression in the ridgeless limit in Appendix C.

$$\epsilon_{\text{RF}} = \begin{cases} (1 - \alpha)(1 + \sigma^2 \prod_{l=1}^{\ell} \frac{\gamma_l - \alpha}{\gamma_l} + \sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l - \alpha}) + (\frac{\alpha}{1-\alpha} + \sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l - \alpha})\eta^2, & \text{if } \alpha < \min\{1, \gamma_{\min}\} \\ \alpha \frac{1-\gamma_{\min}}{\alpha-\gamma_{\min}} + \frac{\gamma_{\min}}{\alpha-\gamma_{\min}}\eta^2, & \text{if } \alpha > \gamma_{\min} \text{ and } \gamma_{\min} < 1 \\ \frac{1}{\alpha-1}\eta^2, & \text{if } \alpha > 1 \text{ and } \gamma_{\min} > 1. \end{cases} \quad (20)$$

We validate the accuracy of this RS result by comparing it against the result of an alternative semianalytical approach. As shown in Appendix C, the zero-temperature posterior average in Eq. (17) can be computed for a fixed realization of the disorder. Even without explicitly evaluating the disorder average, this shows that the RF model should display the three phases indicated by the RS result Eq. (20), and confirms the prediction for in which of the phases the learning curve should depend on the prior variance σ^2 (see Appendix C). Importantly, the RF model learning curve Eq. (17) does not depend on the ordering of the hidden layer widths, which follows from the fact that the random Gaussian hidden layer weight matrices weakly commute [58,59]. For conceptual clarity, we therefore refer to the cases in which different hidden layers are the narrowest as a single phase. To quantitatively test the accuracy of the RS result, the disorder average can be evaluated numerically using sampling (see Appendix G). As shown in Figs. 1 and 2, we observe excellent agreement over a broad range of parameter values. These results are consistent with our expectation that the RS *Ansatz* should yield accurate results for the RF models [20,46,50,57].

While the LR model only exhibits double-descent behavior in the presence of label noise Eq. (19), the RF model can also exhibit double-descent behavior in the absence of label noise if any one of the hidden layers is narrower than the input dimension, i.e., $\gamma_{\min} < 1$. This phenomenon occurs in

Therefore, as [52] found in the ridge regression setting, the LR model exhibits samplewise double-descent behavior, i.e., nonmonotonicity in ϵ_{LR} as a function of α [2,5], in the presence of label noise. In the thermodynamic limit, the double-descent behavior is particularly striking: ϵ_{LR} diverges as $\alpha \rightarrow 1$. In the absence of noise, ϵ_{LR} decreases monotonically from $1 + \sigma^2$ to zero as $\alpha \uparrow 1$, and then remains at zero for all $\alpha > 1$. We remark that, for this and subsequent models, we will not conduct a detailed analysis of what happens precisely at exceptional points, e.g., $\alpha = 1$. In the ridge regression setting, the phase transition at $\alpha = 1$ has recently been analyzed in detail by Canatar *et al.* [20]. We also direct the interested reader to an expository note by Nakkiran [44] for further intuitions on double descent in ridge regression, and to work by Hastie *et al.* [45] for a detailed rigorous analysis. We will take the modelwise double-descent behavior of the LR model as a benchmark for our subsequent analyses of the more complex RF and NN models.

IV. LEARNING CURVES FOR THE RF MODEL

A. Learning curve and double-descent behavior

For RF models, we obtain a closed-form expression for the learning curve at any depth. Let $\gamma_{\min} = \min\{\gamma_1, \dots, \gamma_{\ell}\}$ be the minimum hidden layer width. Then, we find that

a modelwise fashion at fixed data density: if one considers a decreasing sequence of widths γ_{\min} at fixed α , ϵ_{RF} will diverge as $\gamma_{\min} \downarrow \alpha$ [Figs. 1(a), 1(c) and 1(e)]. Equivalently, this divergence can be observed in a samplewise fashion at fixed width, with $\epsilon_{\text{RF}} \rightarrow \infty$ as $\alpha \rightarrow \gamma_{\min}$. Moreover, as illustrated in Fig. 2, it is determined by the width of the narrowest hidden layer. If one adds more bottleneck layers, then the expression for the generalization error in the regime $\alpha < \gamma_{\min}$ will formally include more poles Eq. (20), but these poles will not be visible as one varies the size of the training set or the width of the narrowest bottleneck.

Like the LR model, the RF model exhibits samplewise double-descent behavior in the presence of label noise (Fig. 1). However, if there is a bottleneck layer with width $\gamma_{\min} < 1$, then the addition of label noise does not introduce additional divergences in ϵ_{RF} beyond that at $\alpha \rightarrow \gamma_{\min}$; the pole at $\alpha = 1$ is visible only if $\gamma_{\min} > 1$ [Figs. 1(b), 1(d) and 1(e)]. This is clearly illustrated by comparing learning curves of two-layer ($\ell = 1$) RF models with $\gamma_{\min} = 1/2$ in the absence [Fig. 1(c)] and presence [Fig. 1(d)] of label noise: ϵ_{RF} increases with the addition of noise, but the only visible divergence is at $\alpha \rightarrow \gamma_{\min}$. The presence of only a single divergence in ϵ_{RF} for two-layer models is consistent with work by d'Ascoli *et al.* on the phenomenology of double descent in two-layer RF models trained with ridge regression [22].

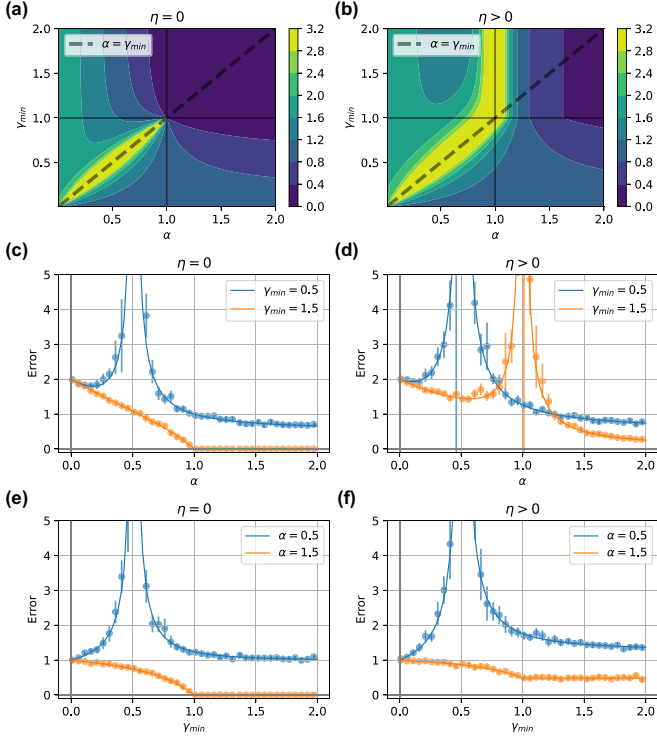


FIG. 1. Sample and modelwise double descent in deep Bayesian random feature models. (a) Contour plot in (α, γ) space of the theoretical error surface ϵ_{RF} Eq. (20) for a single-hidden-layer RF model in the absence of label noise ($\eta = 0$). For all panels, we set the input dimensionality $d = 100$ and prior variance $\sigma^2 = 1$. For details of our numerical methods, see Appendix G. (b) As in panel (a), but in the presence of label noise ($\eta = 0.5$). (c) Horizontal cross sections of panel (a). Theory curves are overlaid with experiment points, plotted with ± 2 SE bars. (d) Horizontal cross sections of panel (b). (e) Vertical cross sections of panel (a). (f) Vertical cross sections of panel (b).

B. Large-width behavior

We now analyze the behavior of RF models at large widths. As $\gamma_1, \dots, \gamma_\ell \rightarrow \infty$, $\epsilon_{\text{RF}} \rightarrow \epsilon_{\text{LR}}$ for any fixed α , σ , and η . We will refer to this simplification—the reduction of the linear curve of a deep linear model to that of simple linear regression—as the *kernel limit* [11,13–16]. To obtain a more precise understanding of the behavior of the RF model near the kernel limit, we expand Eq. (20) in the regime $\gamma_1, \dots, \gamma_\ell \gg 1$. If $\alpha > 1$, then we have $\epsilon_{\text{RF}} = \epsilon_{\text{LR}}$ in this regime. If $\alpha < 1$, we have

$$\epsilon_{\text{RF}} = \epsilon_{\text{LR}} + [(1 - \alpha)(1 - \sigma^2) + \eta^2] \sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l} + \mathcal{O}\left(\frac{\alpha^2}{\gamma^2}\right), \quad (21)$$

where $\mathcal{O}(\alpha^2/\gamma^2)$ denotes terms that include two or more factors of any combination of the layer widths.

For an RF model of equal hidden layer widths $\gamma_1 = \dots = \gamma_\ell = \gamma$, the leading correction scales as $\ell\alpha/\gamma$. For this simple architecture, we can also study the scaling of higher-order corrections relatively easily. In the regime $\alpha < \min\{1, \gamma\}$, the

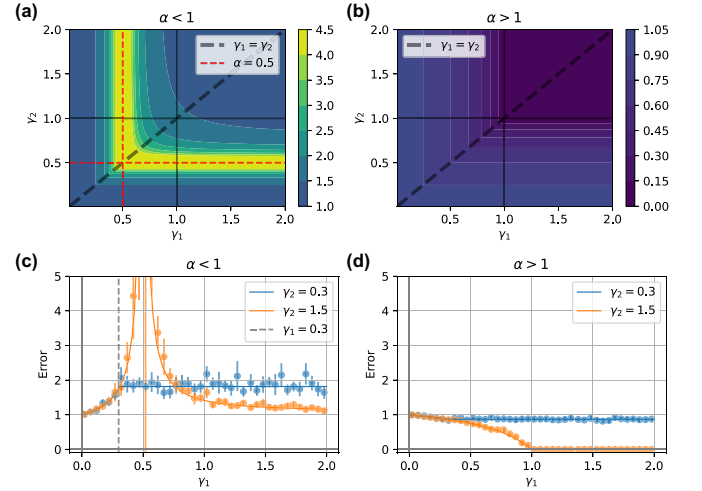


FIG. 2. Double descent in deep random feature models depends on the narrowest hidden layer. (a) Contour plot in (γ_1, γ_2) space of the theoretical error surface ϵ_{RF} Eq. (20) for a deep RF model with two hidden layers and $\alpha = 0.5$. For all panels, we set the input dimensionality $d = 100$, prior variance $\sigma^2 = 1$, and no label noise ($\eta = 0$). For details of our numerical methods, see Appendix G. (b) As in panel (a), but with $\alpha = 1.5$. (c) Horizontal cross sections of panel (a). Theory curves are overlaid with experiment points, plotted with ± 2 SE bars. (d) Horizontal cross sections of panel (b).

learning curve Eq. (20) can be written compactly as

$$\frac{\epsilon_{\text{RF}} - \epsilon_{\text{LR}}}{1 - \alpha + \eta^2} = \tilde{\sigma}^2 \left[\left(\frac{\gamma - \alpha}{\gamma} \right)^\ell - 1 \right] + \ell \frac{\alpha}{\gamma - \alpha}, \quad (22)$$

where we have defined the rescaled prior variance:

$$\tilde{\sigma}^2 \equiv \frac{\sigma^2}{1 + \eta^2/(1 - \alpha)}. \quad (23)$$

Then, for $\alpha/\gamma < 1$, we can read off the full series expansion using the binomial theorem and the geometric series:

$$\frac{\epsilon_{\text{RF}} - \epsilon_{\text{LR}}}{1 - \alpha + \eta^2} = \sum_{j=1}^{\infty} \left[(-1)^j \tilde{\sigma}^2 \binom{\ell}{j} + \ell \right] \frac{\alpha^j}{\gamma^j}, \quad (24)$$

where we note that $\binom{\ell}{j} = 0$ if $j > \ell$. Noting that $\binom{\ell}{j} = \mathcal{O}(\ell^j)$ for $\ell \gg j$, we can see that the dominant term for large ℓ at each order in α/γ will scale with $\ell\alpha/\gamma$, up to around $\mathcal{O}(\alpha^\ell/\gamma^\ell)$. Therefore, depth will have an important effect on how quickly the kernel limit is approached with varying width. At small ℓ , the j -th-order term will simply scale as $\ell(\alpha/\gamma)^j$ for all $j > \ell$, hence the effect of depth on the approach to the kernel limit can be neglected in this regime.

C. Optimal width and depth

With the formula Eq. (20) for the generalization error in hand, we can determine the optimal hidden layer width for fixed depth, noise variance, and prior variance. We focus on the regime $\alpha < \min\{1, \gamma_{\text{min}}\}$, in which the generalization error always depends nontrivially on width. In Appendix F1, we show that the optimal architecture for an RF model depends on the rescaled prior variance $\tilde{\sigma}^2$ defined in Eq. (23). If $\tilde{\sigma} \leq 1$, then $\partial\epsilon_{\text{RF}}/\partial\gamma_l < 0$ for all l and all widths in this regime, hence

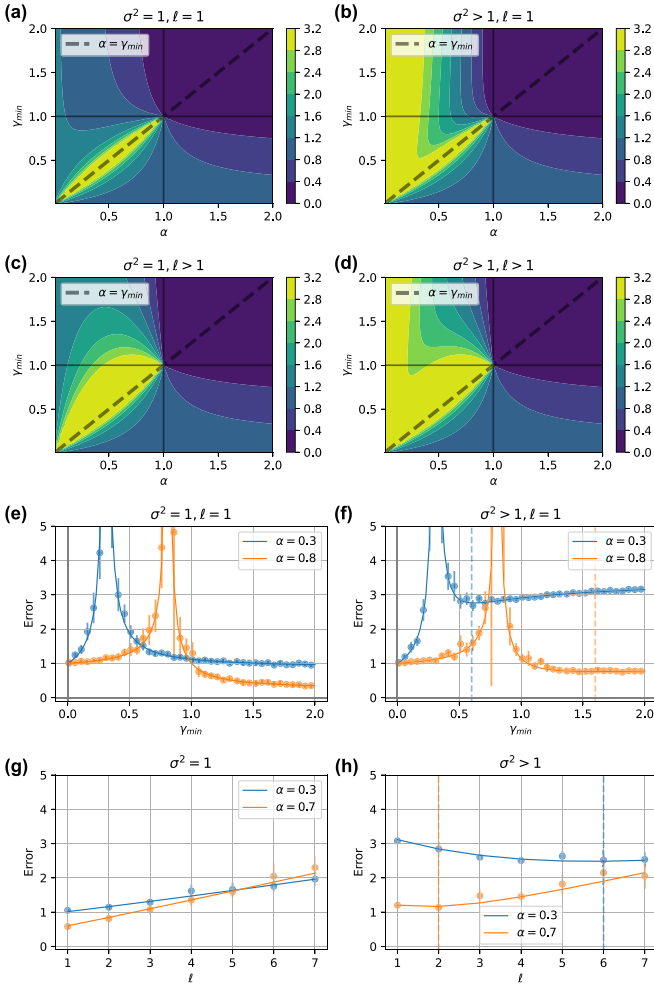


FIG. 3. Optimal RF model architecture depends on target-prior mismatch. (a) Contour plot in (α, γ) space of the theoretical error surface ϵ_{RF} Eq. (20) for a single-hidden-layer RF model with prior variance $\sigma^2 = 1$. For all panels, we have no label noise ($\eta = 0$) and set the input dimensionality $d = 100$. For details of our numerical methods, see Appendix G. (b) As in panel (a), but for a single-hidden-layer RF model with higher prior variance ($\sigma^2 = 4$). (c) As in panel (a), but for a deep RF model ($\ell = 5$) and prior variance $\sigma^2 = 1$. (d) As in panel (a), but for a deep RF model ($\ell = 5$) and with higher prior variance ($\sigma^2 = 4$). (e) Vertical cross sections of panel (a). Theory curves are overlaid with experiment points, plotted with ± 2 SE bars. (f) Vertical cross sections of panel (b). Optimal widths computed from Eq. (25) are marked with dashed vertical lines for each respective setting of α . (g) Error across different depths for prior variance $\sigma^2 = 1$ and fixed width $\gamma = 1.5$ (h) Error across different depths for prior variance $\sigma^2 = 4$ and fixed width $\gamma = 1.5$. Optimal depths computed from Eq. (26) are marked with dashed vertical lines for each respective setting of α .

increasing width always improves generalization. Thus, in this regime, the best RF model is one that behaves identically to an LR model (Fig. 3). If $\tilde{\sigma}^2 > 1$, then ϵ_{RF} is minimized by taking all $\gamma_1 = \dots = \gamma_\ell = \gamma_*$ for

$$\gamma_* = \frac{\tilde{\sigma}^{2/(\ell+1)}}{\tilde{\sigma}^{2/(\ell+1)} - 1}. \quad (25)$$

We note that the leading term in the perturbative expansion Eq. (21) would predict that generalization performance improves with increasing width if $\tilde{\sigma} < 1$, is invariant under changes of width if $\tilde{\sigma} = 1$, and degrades with increasing width if $\tilde{\sigma} > 1$. Thus, in this case the leading-order perturbative correction captures some, but not all, of the effect of width.

In the absence of noise, this yields a simple qualitative picture in which the optimal width is related to the mismatch between the scale of the prior and the target weight vector: if $\mathbb{E}_{\mathcal{W}} \|\mathbf{w}\|^2 = \sigma^2 d \leq d = \|\mathbf{w}_*\|^2$, then wider networks are always better, while otherwise one can obtain improved generalization performance by using an RF model rather than an LR model. This occurs because of the tradeoff between the terms with linear and exponential dependence on depth in Eq. (20). Label noise has the effect of increasing the effective scale of the target in a way that depends on the data density α : as $\alpha \uparrow 1$, wider models are always better in the presence of label noise. This behavior is illustrated in Fig. 3.

Similarly, one can also optimize the depth for fixed width, noise variance, and prior variance. To do so, it is convenient to assume that all layers are of the same width $\gamma_1 = \dots = \gamma_\ell = \gamma$, which allows us to analytically continue ϵ_{RF} as a function of the depth ℓ . Again specializing to the regime $\alpha < \min\{1, \gamma\}$, the generalization error is given by Eq. (22). It is then easy to see that the LR model learning curve Eq. (19) is recovered upon $\ell \downarrow 0$. In Appendix F2, we show that, if $\tilde{\sigma} \leq 1$, ϵ_{RF} is a monotonically increasing function of ℓ , hence shallower RF models always generalize better. This is consistent with our result above for optimal width, because taking $\gamma_l \rightarrow \infty$ for some l effectively reduces the depth of the RF model by 1, by eliminating that layer's contribution to Eq. (22). If $\tilde{\sigma} > 1$, then the optimal depth is given by

$$\ell_* = \begin{cases} j \text{ or } j - 1, & \text{if } \frac{\log(\tilde{\sigma}^2)}{\log[\gamma/(\gamma-\alpha)]} = j \in \mathbb{N}_{>0} \\ \lfloor \frac{\log(\tilde{\sigma}^2)}{\log[\gamma/(\gamma-\alpha)]} \rfloor, & \text{otherwise.} \end{cases} \quad (26)$$

In the former condition, taking $\ell = j$ or $j - 1$ will yield identical generalization error. Moreover, for the condition $\log(\tilde{\sigma}^2)/\log[\gamma/(\gamma-\alpha)] \in \mathbb{N}_{>0}$ to hold, the network width must be of the form

$$\gamma = \frac{\tilde{\sigma}^{2/j}}{\tilde{\sigma}^{2/j} - 1} \alpha, \quad (27)$$

which is consistent with the result for optimal width at fixed depth given in Eq. (25). Therefore, much like we found in our analysis of optimal width, the optimal depth of an RF model is related to the match between the scale of the prior and of the target. This behavior is illustrated in Fig. 3.

V. LEARNING CURVES FOR THE NN MODEL

A. Learning curve and double-descent behavior

For the NN model, we do not obtain a simple closed-form solution for the RS learning curve at general depth. As shown in Appendix B3, we find that the solution is of the form

$$\epsilon_{\text{NN}} = \epsilon_{\text{LR}} + \begin{cases} z - \sigma^2(1 - \alpha), & \text{if } \alpha < 1 \\ 0, & \text{if } \alpha > 1, \end{cases} \quad (28)$$

where $z = z(\alpha, \sigma^2, \eta^2, \gamma_1, \dots, \gamma_\ell)$ is a non-negative real root of the polynomial

$$z^{\ell+1} = \sigma^2(1 - \alpha) \prod_{l=1}^{\ell} \left[\frac{\gamma_l - \alpha}{\gamma_l} z + \frac{\alpha(1 - \alpha + \eta^2)}{\gamma_l} \right]. \quad (29)$$

We defer more detailed discussion of which root should be selected to Appendix B3, where we show that one required condition on the solution is that

$$(\gamma_l - \alpha)z + \alpha(1 - \alpha + \eta^2) > 0 \quad (30)$$

for all l . For a network with a single hidden layer ($\ell = 1$), Eq. (29) is quadratic, and we can easily obtain

$$\frac{z}{1 - \alpha + \eta^2} = \frac{\tilde{\sigma}^2(\gamma_1 - \alpha) + \sqrt{\tilde{\sigma}^4(\gamma_1 - \alpha)^2 + 4\alpha\gamma_1\tilde{\sigma}^2}}{2\gamma_1}, \quad (31)$$

where $\tilde{\sigma}^2$ is defined as in Eq. (23).

The special case of Eq. (28) for networks with hidden layers of equal widths follows from results obtained through a rather different approach in a recent study by Li and Sompolinsky [37]. Concretely, they use an iterative saddle-point argument to approximate the posterior expectation in Eq. (17) for fixed data, and then apply that result to a random Gaussian covariate model under what amounts to the assumption that the quantity $\mathbf{y}^\top(\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{y}$ concentrates rapidly. In Appendix D, we provide a detailed discussion of the mapping between the polynomial condition in terms of which their result is expressed and the RS condition Eq. (29). In Appendix D, we also use a finite-size fixed-data approach derived from our previous work [36] to show that the learning curve should be of the form Eq. (28). Concretely, this approach gives an expression for z as the thermodynamic limit of a dataset average of a ratio of prior averages, with the remaining components of the learning curve exactly matching the RS prediction. Taken together, these results suggest that the RS prediction for the learning curve correctly captures at least the coarse behavior of generalization in NNs.

To further probe whether the RS prediction is quantitatively accurate, we evaluate the finite-size data average numerically. As shown in Figs. 4 and 5, and in supplemental figures provided in Appendix G, we observe good agreement for two-layer networks. To probe the accuracy of the RS prediction for deeper networks, we solve the polynomial Eq. (29) numerically. As shown in Figs. 4 and 5, we again observe good agreement. Therefore, both alternative heuristic analytical approaches and numerical results are consistent with the RS learning curve, suggesting that it provides a reasonably accurate picture of generalization in deep NNs.

Like the previously studied models, we see that label noise can induce samplewise double descent, with $\epsilon_{\text{NN}} \rightarrow \infty$ as $\alpha \rightarrow 1$ (Fig. 4). However, unlike for the RF model, having relatively narrow hidden layers does not introduce the possibility of divergences other than at $\alpha = 1$, as z should remain bounded. This is illustrated in Fig. 5, where we repeat the analysis of Fig. 2, but do not observe similar modelwise divergences. Moreover, this means that the NN model does not display samplewise divergences in the ab-

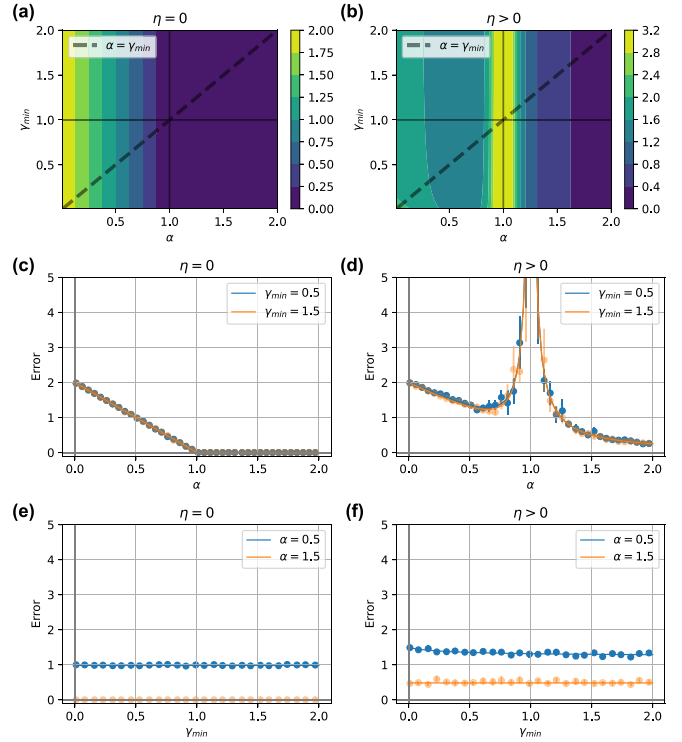


FIG. 4. Samplewise double descent in deep Bayesian neural networks. (a) Contour plot in (α, γ) space of the theoretical error surface ϵ_{NN} Eq. (28) for a two-layer NN model in the absence of label noise ($\eta = 0$). For all panels, we set the input dimensionality $d = 100$ and prior variance $\sigma^2 = 1$. For details of our numerical methods, see Appendix G. (b) As in panel (a), but in the presence of label noise ($\eta = 0.5$). (c) Horizontal cross sections of panel (a). Theory curves are overlaid with experiment points, plotted with ± 2 SE bars. (d) Horizontal cross sections of panel (b). (e) Vertical cross sections of panel (a). (f) Vertical cross sections of panel (b).

sence of label noise. Therefore, training the hidden layers affords the advantage of avoiding the possible modelwise and samplewise divergences that can arise in RF models with narrow bottlenecks. This sharp contrast makes sense, since in the RF model the presence of layers with width $\gamma_l < 1$ introduces a true bottleneck, while in the NN model one could in principle find a solution where, in all layers except the first, exactly one weight is nonzero, and the model essentially reduces to shallow linear regression. The existence of this solution reflects the fact that, from the standpoint of expressivity, NN models should be able to perform as well as LR models, and differences in performance reflect the behavior of the inference algorithm [42]. Indeed, if $\sigma = 1$ and $\eta = 0$, we have the solution $z = 1 - \alpha$, and $\epsilon_{\text{NN}} = \epsilon_{\text{LR}}$. Therefore, in this special case, the RS result predicts that depth has no effect on generalization performance. This behavior is clearly illustrated by Fig. 5, where the generalization error of a three-layer NN remains constant as the widths of the two hidden layers are varied. Even at nonzero noise levels, Fig. 4 illustrates that width has a relatively minimal effect on generalization performance when $\sigma = 1$.

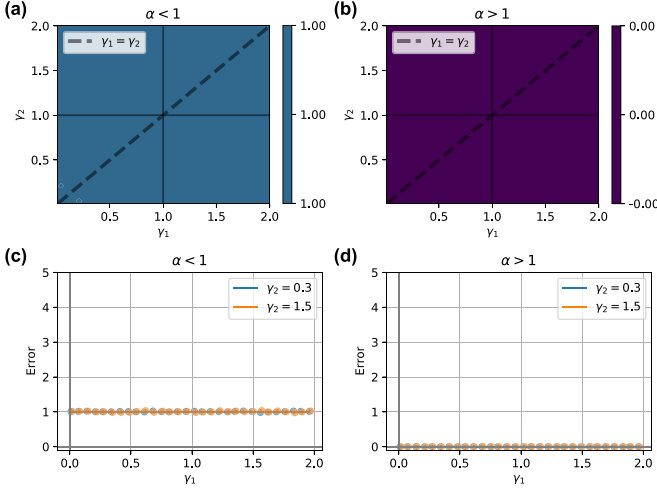


FIG. 5. Bottleneck layers do not induce modelwise double descent in deep NNs. (a) Contour plot in (γ_1, γ_2) space of the theoretical error surface ϵ_{NN} Eq. (28) for a deep Bayesian NN model with two hidden layers and $\alpha = 0.5$. For all panels, we set the input dimensionality $d = 100$, prior variance $\sigma^2 = 1$, and no label noise ($\eta = 0$). For details of our numerical methods, see Appendix G. (b) As in panel (a), but with $\alpha = 1.5$. (c) Horizontal cross sections of panel (a). Theory curves are overlaid with experiment points, plotted with ± 2 SE bars. (d) Horizontal cross sections of panel (b).

B. Large-width behavior

Beyond the special cases mentioned above, we observe that, in the limit $\gamma_1, \dots, \gamma_\ell \rightarrow \infty$, we have the solution $z = \sigma^2(1 - \alpha)$ for any fixed α, σ , and η . Therefore, as we found for the RF model, the NN model's generalization performance reduces to that of the shallow LR model in this large-width limit: $\epsilon_{\text{NN}} \rightarrow \epsilon_{\text{LR}}$. In the regime $\alpha < 1, \gamma_1, \dots, \gamma_\ell \gg 1$, we can obtain a perturbative solution for the learning curve (see Appendix B), which is given as

$$\epsilon_{\text{NN}} = \epsilon_{\text{LR}} + [(1 - \alpha)(1 - \sigma^2) + \eta^2] \sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l} + \mathcal{O}\left(\frac{\alpha^2}{\gamma^2}\right) \quad (32)$$

to leading order. This result can be compared to the leading-order perturbative computation of the zero-temperature learning curve for fixed data in our previous work [35]. As shown in Appendix E, averaging the result of [35] over data recovers the $\mathcal{O}(\alpha/\gamma)$ term resulting from the replica method computation. This suggests that the RS prediction for the NN model learning curve is accurate at large widths. Heuristically, this makes sense because the concavity of the log posterior is restored in the limit $\gamma_1, \dots, \gamma_\ell \rightarrow \infty$.

This limiting result has several interesting features. First, paralleling our analysis of the RF model at large widths, the closeness of the NN model's learning curve to that of simple linear regression is determined by a combination of depth, dataset size, and width. Second, not only do the RS learning curves for NN and RF models agree at infinite width, but the leading-order corrections agree, i.e., the term that is linear in α/γ_l [see Eq. (21)]. Thus, if one tracked only the generalization error, one could not differentiate between training

only the readout layer and training all of the layers simply by considering the leading-order perturbative correction. One could of course distinguish between these two models by considering leading-order corrections to observables that explicitly measure task-relevant feature learning in early hidden layers, such as the kernels considered in our previous work [35].

C. Generalization gap between RF and NN models

To distinguish between RF and NN models based on generalization performance, one must therefore go to higher order in perturbation theory. For convenience and clarity, we specialize to the case of networks with equal hidden layer widths $\gamma_1 = \gamma_2 = \dots = \gamma_\ell = \gamma$. Then, we find that

$$\begin{aligned} \frac{\epsilon_{\text{NN}} - \epsilon_{\text{LR}}}{1 - \alpha + \eta^2} &= (1 - \tilde{\sigma}^2) \frac{\ell \alpha}{\gamma} \\ &+ \left(\frac{\ell(\ell - 1)\tilde{\sigma}^2}{2} - \frac{\ell(\ell + 1)}{2\tilde{\sigma}^2} + \ell \right) \frac{\alpha^2}{\gamma^2} \\ &+ \mathcal{O}\left(\frac{\alpha^3}{\gamma^3}\right), \end{aligned} \quad (33)$$

where $\tilde{\sigma}$ is defined as in Eq. (23). In contrast, by truncating Eq. (24) to this order, we can see the corresponding RF model has generalization error:

$$\begin{aligned} \frac{\epsilon_{\text{RF}} - \epsilon_{\text{LR}}}{1 - \alpha + \eta^2} &= (1 - \tilde{\sigma}^2) \frac{\ell \alpha}{\gamma} + \left(\frac{\ell(\ell - 1)\tilde{\sigma}^2}{2} + \ell \right) \frac{\alpha^2}{\gamma^2} \\ &+ \mathcal{O}\left(\frac{\alpha^3}{\gamma^3}\right). \end{aligned} \quad (34)$$

Therefore, the next-to-leading order correction can distinguish between RF and NN models. Moreover, the gap in the generalization performance of the two models is, to the given order,

$$\frac{\epsilon_{\text{RF}} - \epsilon_{\text{NN}}}{1 - \alpha + \eta^2} = \frac{\ell(\ell + 1)\alpha^2}{2\tilde{\sigma}^2 \gamma^2} + \mathcal{O}\left(\frac{\alpha^3}{\gamma^3}\right). \quad (35)$$

The coefficient of the leading term is always positive, hence at very large widths training both layers should produce a small benefit relative to simply training the readout. In the two-layer case, one can use the closed-form solution for the RS generalization error to show that the generalization gap $\epsilon_{\text{RF}} - \epsilon_{\text{NN}}$ is strictly positive, except at vanishing load or in the limit $\gamma_1 \rightarrow \infty$ (see Appendix F4). These results suggest that training all layers of a deep linear network can yield improved generalization relative to training only the last layer, even if the widths are large enough such that the RF model does not display double descent in the absence of noise. See Fig. 6 for an illustration of this behavior.

D. Optimal width and depth

The leading correction term in Eq. (32) predicts that generalization error always decreases with increasing width (respectively increases with increasing depth) if $\tilde{\sigma} < 1$, is constant if $\tilde{\sigma} = 1$, and always increases (respectively decreases) if $\tilde{\sigma} > 1$. As shown in Appendix E, this condition is the dataset-averaged version of the fixed-data condition noted by Li and Sompolinsky [37] and in our previous perturbative work [35].




FIG. 6. The generalization gap between RF and NN models approaches zero with increasing width. The difference $\epsilon_{RF} - \epsilon_{NN}$ remains positive or consistent with zero within standard error throughout, highlighting the advantage in training all layers. See Appendix G for details of our numerical methods.

In Appendix F3, we show in detail that this condition captures the behavior of the full RS generalization error of an NN with one hidden layer; for other depths it follows from an argument based on implicit differentiation given by Li and Sompolinsky [37]. Therefore, like in our study of the RF model, the optimal width of an NN depends on the match between the scales of the prior and of the target. However, unlike for an RF model, the RS result suggests that the optimal generalization performance for an NN is obtained either by taking $\gamma_l \rightarrow \infty$ or by taking $\gamma_l \downarrow 0$, behavior which is fully predicted by the leading perturbative correction. This behavior is illustrated in Fig. 7.

VI. DISCUSSION AND CONCLUSIONS

A. Summary of results

In this paper, we studied the statistical mechanics of inference in deep Bayesian linear models. We characterized the learning curves of deep linear random feature models and deep linear neural networks for isotropic Gaussian covariates, using a combination of the replica trick and replica-free methods. Our primary results for how deep Bayesian linear models with random and learned features differ or resemble may be summarized as follows.

(1) In the presence of label noise, both RF and NN models display samplewise double descent (Figs. 1 and 4). For RF models, the presence of a bottleneck layer with width less than the input dimension induces modelwise double descent at fixed dataset size and samplewise double descent at fixed width (Figs. 1 and 2), while bottlenecks do not affect the double-descent behavior of NN models (Figs. 4 and 5). In particular, NN models do not display modelwise double descent, and do not display samplewise double descent in the absence of label noise.

(2) For both RF and NN models, the effect of width on generalization depends on the match between the prior variance and the true scale of the targets, with wider networks yielding better generalization when the prior variance is less than the

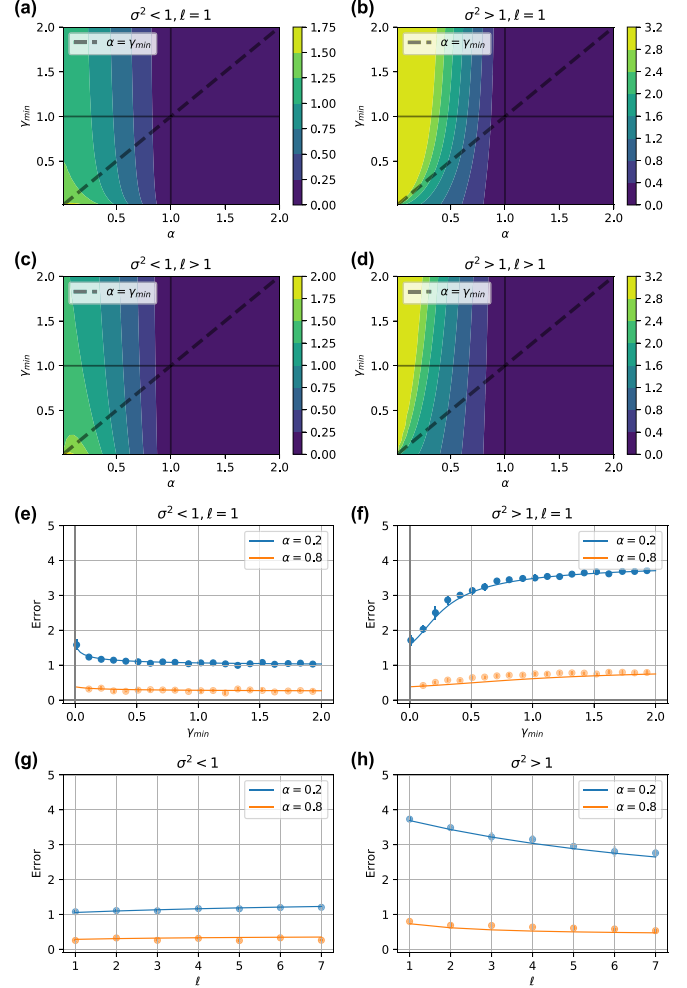


FIG. 7. Optimal NN model architecture depends on target-prior mismatch. (a) Contour plot in (α, γ) space of the theoretical error surface ϵ_{NN} Eq. (28) for a single-hidden-layer NN model with prior variance $\sigma^2 = 1/4$. For all panels, we have no label noise ($\eta = 0$) and set the input dimensionality $d = 100$. For details of our numerical methods, see Appendix G. (b) As in panel (a), but for a single-hidden-layer NN model with higher prior variance ($\sigma^2 = 4$). (c) As in panel (a), but for a deep NN model ($\ell = 5$) and prior variance $\sigma^2 = 1/4$. (d) As in panel (a), but for a deep NN model ($\ell = 5$) and with higher prior variance ($\sigma^2 = 4$). (e) Vertical cross sections of panel (a). Theory curves are overlaid with experiment points, plotted with ± 2 SE bars. (f) Vertical cross sections of panel (b). (g) Error across different depths for prior variance $\sigma^2 = 1/4$ and fixed width $\gamma = 1.5$. (h) Error across different depths for prior variance $\sigma^2 = 4$ and fixed width $\gamma = 1.5$.

average target scale (Figs. 3 and 7). For NN models, taking the network to be as wide or as narrow as possible is always optimal. In contrast, when the prior variance is greater than the average target scale, there is a particular width that yields optimal generalization in RF models.

(3) Similarly, the optimal depth for both models depends on prior-target mismatch. Paralleling the case of optimal width, deeper models always perform worse when the prior variance is less than the average target scale (Figs. 3 and 7). When prior variance is greater than the average target scale,

shallower models perform better. In this regime, as in the case of optimal width, there is a particular depth that yields optimal RF model generalization for fixed width, prior variance, and data density.

(4) Both RF and NN models display kernel-limit behavior, i.e., their learning curves reduce to those of shallow linear regression, when the depth and dataset size are small relative to the hidden layer width. Moreover, for both classes of deep models, the $\mathcal{O}(\ell\alpha/\gamma)$ perturbative correction captures much of the gross qualitative behavior of the learning curve as a function of prior variance, width, and depth.

(5) The learning curves of wide RF and NN models not only coincide in the limit $\ell\alpha/\gamma \downarrow 0$, but have identical leading-order corrections in $\ell\alpha/\gamma$. Training all layers improves generalization relative to training only the readout, but this gap is an $\mathcal{O}(\ell^2\alpha^2/\gamma^2)$ effect (Fig. 6).

B. Prior work

As noted above, our results for deep linear neural networks partially overlap with those obtained previously by Li and Sompolinsky [37]. Specifically, the RS learning curve for networks of equal hidden layer width agrees with the result they obtained through an alternative heuristic, and, as a result, their criteria for when generalization improves or degrades with width and depth coincide with those obtained here. However, they did not analyze in detail how the kernel limit is approached as $\ell\alpha/\gamma \downarrow 0$, and did not consider random feature models. Our results therefore complement their study by providing a more granular picture of how generalization performance for random datasets depends on model architecture. Moreover, the agreement between their approximations, our RS results, and our numerical simulations is consistent with the conjecture that the RS learning curve is reasonably accurate.

Double-descent phenomena have recently garnered significant interest in deep learning [2,5,10,17,20,22,24,25,44,52]. In high-dimensional random feature models like those considered in Sec. IV, divergences in the generalization error can arise through interactions between randomness in the features and randomness in the training data [17,22–25]. Moreover, as noted in our discussion of simple linear regression models in Sec. III, divergences can arise in models without additional feature randomness—including kernel regressors with deterministic nonlinear features—due to overfitting of noisy labels [20,22–25,44,45,52]. Disentangling the causes of nonmonotonic generalization performance observed in experimental settings for realistic data models remains an interesting subject for further study [2,5,10,17,20,22,24,25,44,52].

The statistical mechanics of inference in shallow linear models with more general priors and likelihoods was investigated in detail by Advani and Ganguli [53], who showed a correspondence between the performance of Bayesian MMSE inference and a class of algorithms known as M-estimators. The effect of prior mismatch on the performance of the shallow MMSE estimator has also been considered in recent rigorous work by Barbier *et al.* [21]. However, neither of these works considered the effect of depth on inference.

Here, we considered a proportional asymptotic limit in which the input dimension d , dataset size p , and hidden layer

widths n_ℓ tend jointly to infinity with fixed limiting ratios $\alpha = p/d = \mathcal{O}(1)$ and $\gamma_\ell = n_\ell/d = \mathcal{O}(1)$ and fixed depth ℓ . Moreover, we have only considered networks with scalar output. In this setting, kernel-machine behavior, i.e., approximate reduction of the learning curves of deep models to those of simple linear regression [11–14,16], emerges when the ratio $\ell\alpha/\gamma_\ell$ is vanishingly small. This is consistent with our observations in prior work [35,36], and with those of works that considered large depths but fixed dataset size [32] or large dataset size but fixed depth [37,39]. Given the large scale of contemporary regression and classification datasets, e.g., [61], careful consideration of limits in which the output dimension and dataset size are not vanishingly small relative to hidden layer width warrants further study.

This regime has thus far proven challenging to access perturbatively, as large deviations from the kernel limit may emerge [4,35–37,39]. Existing fixed-data approaches to regimes in which either the dataset size [36,37,39] or the output dimension [4,36] is not negligible relative to hidden layer width rely on saddle-point approximations that may break down when both of these parameters are large. New approaches will therefore be required to study networks in this limit nonperturbatively. With such results in hand, it will be interesting to test whether existing perturbative predictions do in fact capture qualitative features of how generalization depends on network architecture and other hyperparameters. For the simple models considered here, we found that small-sample-size perturbation theory does in fact yield largely correct predictions for when wider networks generalize better, even at large sample size.

C. Outlook

We conclude by noting that our paper has several important limitations, which will be interesting to address in future work. First, our approach is highly specialized to deep linear networks, and would not extend easily to nonlinear models. Though the utility of linear networks as a model system for studying the effect of depth on inference has been clearly established [35,37,42,43,46], rigorous characterization of the effect of nonlinearity on inference in deep Bayesian neural networks remains a largely open problem [4,31,32,35,37–39,62]. Second, we have assumed that the covariates are drawn from an isotropic Gaussian distribution. Though this is a standard generative model in theoretical studies of inference [21,44–46,52,53], it is undoubtedly not reflective of real-world data. Extending results of this form to more realistic generative models will be an interesting objective for future work [20,27,63]. We remark that some of the fixed-data results of Appendices C and D would extend immediately to anisotropic and non-Gaussian data provided that the requisite invertibility conditions hold. While BNNs are finding practical applications in physics and elsewhere [64], another important direction for future work will be to develop a rigorous theoretical understanding of how results on the generalization performance of BNNs, like those obtained here, relate to the generalization performance of networks trained with stochastic gradient-based algorithms, a link that remains incompletely understood [47,49,65,66]. Finally, our replica theory approach is of course nonrigorous. For the RF

model, we do not expect replica symmetry to be broken, and conjecture that our results might be rigorously justifiable [20,50,51,57–59]. Moreover, our replica-free analytical approaches and numerical experiments suggest that our RS results for NNs are at the very least a reasonable approximation for their true generalization performance. With that in mind, careful exploration of the possibility of replica symmetry breaking will be an interesting topic for further investigation.

Note added. Recently, results on the behavior of the ridge regression estimator—which in this case would coincide with the limiting MMSE estimator—for a model with a single layer of Gaussian linear random features were announced by Rocks and Mehta [60].

ACKNOWLEDGMENTS

We thank A. Atanasov, B. Bordon, and B. S. Ruben for helpful comments on our paper. This work was supported by a Google Faculty Research Award and NSF Grant No. 2134157. A subset of the computations in this paper was performed using the Harvard University Faculty of Arts and Sciences Division of Science Research Computing Group’s Cannon HPC cluster.

APPENDIX A: REPLICA THEORY FRAMEWORK

In this Appendix, we introduce the replica theory framework we use to compute learning curves. We direct the interested reader to [50] for more details on replica theory. Our starting point is the partition function of the Bayes posterior:

$$Z = \mathbb{E}_{\mathcal{W}} \exp \left(-\frac{\beta}{2} \sum_{\mu=1}^p [g_{\mathbf{w}}(\mathbf{x}_{\mu}) - y_{\mu}]^2 \right). \quad (\text{A1})$$

In the limit of interest, we expect the quenched free energy

$$f = - \lim_{d,p,n_1,\dots,n_{\ell} \rightarrow \infty} \frac{1}{d} \log Z \quad (\text{A2})$$

to be self-averaging, i.e., $f = \mathbb{E}_{\mathcal{D}} f$ with probability 1 [50,51]. To compute the limiting quenched average, we will resort to the replica trick, which proceeds via the identity $\mathbb{E}_{\mathcal{D}} \log Z = \lim_{m \rightarrow 0} \log(\mathbb{E}_{\mathcal{D}} Z^m)/m$, yielding

$$f = - \lim_{m \rightarrow 0} \lim_{d,p,n_1,\dots,n_{\ell} \rightarrow \infty} \frac{1}{dm} \log \mathbb{E}_{\mathcal{D}} Z^m \quad (\text{A3})$$

after a nonrigorous interchange of the limits $m \rightarrow 0$ and $d, p, n_1, \dots, n_{\ell} \rightarrow \infty$. We evaluate the moments $\mathbb{E}_{\mathcal{D}} Z^m$ for positive integer m , and assume that they can be analytically continued to $m \rightarrow 0$.

In this Appendix, we show that we can write the disorder-averaged replicated partition function $\mathbb{E}_{\mathcal{D}} Z^m$ for each model as an integral over some set of order parameter matrices. We then introduce the replica-symmetric *Ansatz* under which we will solve the resulting saddle-point equations in Appendix B.

1. Integrating out the data

We first integrate out the data. Introducing replicas indexed by $a = 1, \dots, m$, the object of interest is the disorder-

averaged replicated partition function:

$$\mathbb{E}_{\mathcal{D}} Z^m = \mathbb{E}_{\{\mathcal{W}^a\}} \mathbb{E}_{\mathcal{D}} \exp \left(-\frac{\beta}{2} \sum_{a=1}^m \sum_{\mu=1}^p [g_{\mathbf{w}^a}(\mathbf{x}_{\mu}) - y_{\mu}]^2 \right) \quad (\text{A4})$$

where \mathbf{w}^a denotes the end-to-end weight vector with appropriate replica indices for a given model. Using the fact that the training examples are independent and identically distributed, we have

$$\mathbb{E}_{\mathcal{D}} Z^m = \mathbb{E}_{\{\mathcal{W}^a\}} \mathbb{E}_{\mathcal{D} \setminus \mathcal{X}} \left[\mathbb{E}_{\mathbf{x}, \xi} \exp \left(-\frac{\beta}{2} \sum_{a=1}^m [g_{\mathbf{w}^a}(\mathbf{x}) - y]^2 \right) \right]^p, \quad (\text{A5})$$

where we use $\mathbb{E}_{\mathcal{D} \setminus \mathcal{X}}$ as shorthand for expectation with respect to all quenched disorder other than the training inputs and label noise. The expectations over \mathbf{x} and ξ are Gaussian integrals, hence we can evaluate them explicitly. After simplifying the resulting determinant with the aid of the matrix determinant lemma, the Weinstein-Aronzjan identity, and the push-through identity [67], this yields

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}, \xi} \exp \left(-\frac{\beta}{2} \sum_{a=1}^m [g_{\mathbf{w}^a}(\mathbf{x}) - y]^2 \right) \\ &= \det(\mathbf{I}_m + \beta \mathbf{Q})^{-1/2} [1 + \beta \eta^2 \mathbf{1}^{\top} (\mathbf{I}_m + \beta \mathbf{Q})^{-1} \mathbf{1}]^{-1/2}, \end{aligned} \quad (\text{A6})$$

where we have defined the $m \times m$ overlap matrix

$$Q^{ab} \equiv \frac{1}{d} (\mathbf{w}^a - \mathbf{w}_*) \cdot (\mathbf{w}^b - \mathbf{w}_*). \quad (\text{A7})$$

Enforcing the definition of the order parameter matrix \mathbf{Q} by introducing corresponding Lagrange multipliers $\hat{\mathbf{Q}}$ [50], we therefore have

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} Z^m &= \int \frac{d\mathbf{Q} d\hat{\mathbf{Q}}}{(4\pi i/d)^{m(m+1)/2}} \\ &\times \exp \left(-\frac{d}{2} [\text{tr}(\mathbf{Q}\hat{\mathbf{Q}}) + \alpha m G_1(\mathbf{Q})] \right) S(\hat{\mathbf{Q}}) \end{aligned} \quad (\text{A8})$$

where we have defined

$$\begin{aligned} G_1(\mathbf{Q}) &\equiv \frac{1}{m} \log \det(\mathbf{I}_m + \beta \mathbf{Q}) \\ &+ \frac{1}{m} \log [1 + \beta \eta^2 \mathbf{1}^{\top} (\beta^{-1} \mathbf{I}_m + \mathbf{Q})^{-1} \mathbf{1}] \end{aligned} \quad (\text{A9})$$

and

$$S(\hat{\mathbf{Q}}) \equiv \mathbb{E}_{\{\mathcal{W}^a\}} \mathbb{E}_{\mathcal{D} \setminus \mathcal{X}} \exp \left(\frac{1}{2} \sum_{a,b} \hat{Q}^{ab} (\mathbf{w}^a - \mathbf{w}_*) \cdot (\mathbf{w}^b - \mathbf{w}_*) \right). \quad (\text{A10})$$

Here, the integrals over \mathbf{Q} and $\hat{\mathbf{Q}}$ are taken over the spaces of real and imaginary $m \times m$ symmetric matrices, respectively. Our remaining task is to integrate out the weights, which we will do for each of the three models of interest in the following sections.

2. Integrating out the weights for the LR model

For simple linear regression, there is no quenched disorder other than the training inputs and we have $\mathbb{E}_{\{\mathcal{Y}^a\}} = \mathbb{E}_{\{\mathbf{w}^a\}}$. In this case, all integrals are Gaussian, and we can easily compute

$$S(\hat{\mathbf{Q}}) = \det(\mathbf{I}_m - \sigma^2 \hat{\mathbf{Q}})^{-d/2} \times \exp \left[\frac{1}{2} \mathbf{1}^\top \hat{\mathbf{Q}} (\mathbf{I}_m - \sigma^2 \hat{\mathbf{Q}})^{-1} \mathbf{1} \|\mathbf{w}_*\|_2^2 \right]. \quad (\text{A11})$$

Using the assumption that $\|\mathbf{w}_*\|_2^2 = d$ and defining

$$m G_2(\mathbf{Q}, \hat{\mathbf{Q}}) \equiv \text{tr}(\mathbf{Q} \hat{\mathbf{Q}}) + \det(\mathbf{I}_m - \sigma^2 \hat{\mathbf{Q}}) - \mathbf{1}^\top \hat{\mathbf{Q}} (\mathbf{I}_m - \sigma^2 \hat{\mathbf{Q}})^{-1} \mathbf{1} \quad (\text{A12})$$

we can write the averaged replicated partition function of a single-layer network as

$$\mathbb{E}_{\mathcal{D}} Z_{\text{LR}}^m = \int \frac{d\mathbf{Q} d\hat{\mathbf{Q}}}{(4\pi i/d)^{m(m+1)/2}} \times \exp \left(-\frac{1}{2} dm [\alpha G_1(\mathbf{Q}) + G_2(\mathbf{Q}, \hat{\mathbf{Q}})] \right). \quad (\text{A13})$$

3. Integrating out the weights for the RF model

We now consider deep random feature models, for which we have $\mathbb{E}_{\mathcal{D}, \mathcal{X}} = \mathbb{E}_{\mathbf{U}_1, \dots, \mathbf{U}_\ell}$ and $\mathbb{E}_{\{\mathcal{Y}^a\}} = \mathbb{E}_{\{\mathbf{v}^a\}}$. We of course have

$$S = \exp \left(\frac{1}{2} \mathbf{1}^\top \hat{\mathbf{Q}} \mathbf{1} \|\mathbf{w}_*\|_2^2 \right) \mathbb{E}_{\mathbf{U}_1, \dots, \mathbf{U}_\ell, \{\mathbf{v}^a\}} \times \exp \left(\frac{1}{2} \sum_{a,b} \hat{Q}^{ab} \mathbf{w}^a \cdot \mathbf{w}^b - \sum_{a,b} \hat{Q}^{ab} \mathbf{w}_* \cdot \mathbf{w}^a \right). \quad (\text{A14})$$

By introducing order parameters

$$C_1^{ab} \equiv \frac{1}{n_1 \dots n_\ell} (\mathbf{U}_2 \dots \mathbf{U}_\ell \mathbf{v}^a) \cdot (\mathbf{U}_2 \dots \mathbf{U}_\ell \mathbf{v}^b) \quad (\text{A15})$$

via Fourier representations of the Dirac distribution with corresponding Lagrange multipliers $\hat{\mathbf{C}}_1$, we can integrate out \mathbf{U}_1 , yielding

$$S = \int \frac{d\mathbf{C}_1 d\hat{\mathbf{C}}_1}{(4\pi i/n_1)^{m(m+1)/2}} \exp \left(-\frac{n_1}{2} \text{tr}(\mathbf{C}_1 \hat{\mathbf{C}}_1) - \frac{d}{2} \log \det(\mathbf{I}_m - \sigma^2 \mathbf{C}_1 \hat{\mathbf{Q}}) + \frac{\|\mathbf{w}_*\|_2^2}{2} \mathbf{1}^\top \hat{\mathbf{Q}} (\mathbf{I}_m - \sigma^2 \mathbf{C}_1 \hat{\mathbf{Q}})^{-1} \mathbf{1} \right) \times \mathbb{E}_{\mathbf{U}_2, \dots, \mathbf{U}_\ell, \{\mathbf{v}^a\}} \exp \left(\frac{1}{2n_2 \dots n_\ell} \sum_{a,b} \hat{C}_1^{ab} (\mathbf{U}_2 \dots \mathbf{U}_\ell \mathbf{v}^a)^\top \mathbf{U}_2 \dots \mathbf{U}_\ell \mathbf{v}^b \right). \quad (\text{A16})$$

It is easy to see that we can iterate this procedure forward through the network, introducing order parameters

$$C_l^{ab} \equiv \frac{1}{n_l \dots n_\ell} (\mathbf{U}_{l+1} \dots \mathbf{U}_\ell \mathbf{v}^a) \cdot (\mathbf{U}_{l+1} \dots \mathbf{U}_\ell \mathbf{v}^b) \quad (\text{A17})$$

for $l = 1, \dots, \ell - 1$ and

$$C_\ell^{ab} \equiv \frac{1}{n_\ell} \mathbf{v}^a \cdot \mathbf{v}^b \quad (\text{A18})$$

along with corresponding Lagrange multipliers, yielding

$$S(\hat{\mathbf{Q}}) = \int \frac{d\mathbf{C}_1 d\hat{\mathbf{C}}_1}{(4\pi i/n_1)^{m(m+1)/2}} \dots \int \frac{d\mathbf{C}_\ell d\hat{\mathbf{C}}_\ell}{(4\pi i/n_\ell)^{m(m+1)/2}} \exp \left(-\frac{d}{2} \log \det(\mathbf{I}_m - \sigma^2 \mathbf{C}_1 \hat{\mathbf{Q}}) + \frac{\|\mathbf{w}_*\|_2^2}{2} \mathbf{1}^\top \hat{\mathbf{Q}} (\mathbf{I}_m - \sigma^2 \mathbf{C}_1 \hat{\mathbf{Q}})^{-1} \mathbf{1} \right) \times \exp \left(-\frac{1}{2} \sum_{l=1}^{\ell-1} n_l [\text{tr}(\mathbf{C}_l \hat{\mathbf{C}}_l) + \log \det(\mathbf{I}_m - \mathbf{C}_{l+1} \hat{\mathbf{C}}_l)] \right) \exp \left(-\frac{1}{2} n_\ell [\text{tr}(\mathbf{C}_\ell \hat{\mathbf{C}}_\ell) + \log \det(\mathbf{I}_m - \hat{\mathbf{C}}_\ell)] \right). \quad (\text{A19})$$

Then, using the assumption that $\|\mathbf{w}_*\|_2^2 = d$ and defining

$$m G_2(\mathbf{Q}, \hat{\mathbf{Q}}, \mathbf{C}_1) \equiv \text{tr}(\mathbf{Q} \hat{\mathbf{Q}}) + \log \det(\mathbf{I}_m - \sigma^2 \mathbf{C}_1 \hat{\mathbf{Q}}) - \mathbf{1}^\top \hat{\mathbf{Q}} (\mathbf{I}_m - \sigma^2 \mathbf{C}_1 \hat{\mathbf{Q}})^{-1} \mathbf{1} \quad (\text{A20})$$

and

$$m G_3(\mathbf{C}_1, \hat{\mathbf{C}}_1, \dots, \mathbf{C}_\ell, \hat{\mathbf{C}}_\ell) \equiv \sum_{l=1}^{\ell-1} \gamma_l \left[\text{tr}(\mathbf{C}_l \hat{\mathbf{C}}_l) + \log \det(\mathbf{I}_m - \mathbf{C}_{l+1} \hat{\mathbf{C}}_l) \right] + \gamma_\ell \left[\text{tr}(\mathbf{C}_\ell \hat{\mathbf{C}}_\ell) + \log \det(\mathbf{I}_m - \hat{\mathbf{C}}_\ell) \right], \quad (\text{A21})$$

we can write the averaged replicated partition function as

$$\mathbb{E}_{\mathcal{D}} Z_{\text{RF}}^m = \int \frac{d\mathbf{Q} d\hat{\mathbf{Q}}}{(4\pi i/d)^{m(m+1)/2}} \int \frac{d\mathbf{C}_1 d\hat{\mathbf{C}}_1}{(4\pi i/n_1)^{m(m+1)/2}} \dots \int \frac{d\mathbf{C}_\ell d\hat{\mathbf{C}}_\ell}{(4\pi i/n_\ell)^{m(m+1)/2}} \times \exp \left(-\frac{1}{2} dm F(\mathbf{Q}, \hat{\mathbf{Q}}, \mathbf{C}_1, \hat{\mathbf{C}}_1, \dots, \mathbf{C}_\ell, \hat{\mathbf{C}}_\ell) \right), \quad (\text{A22})$$

where we have defined

$$F(\mathbf{Q}, \hat{\mathbf{Q}}, \mathbf{C}_1, \hat{\mathbf{C}}_1, \dots, \mathbf{C}_\ell, \hat{\mathbf{C}}_\ell) \equiv \alpha G_1(\mathbf{Q}) + G_2(\mathbf{Q}, \hat{\mathbf{Q}}, \mathbf{C}_1) + G_3(\mathbf{C}_1, \hat{\mathbf{C}}_1, \dots, \mathbf{C}_\ell, \hat{\mathbf{C}}_\ell). \quad (\text{A23})$$

We note that we have intentionally split the entropic contribution to the replica free energy into two pieces. The first, $G_2(\mathbf{Q}, \hat{\mathbf{Q}}, \mathbf{C}_1)$, reduces to the entropic contribution for simple linear regression upon fixing $\hat{\mathbf{C}}_1 = \mathbf{I}_m$. The second, G_3 , captures the effect of depth.

4. Integrating out the weights for the NN model

We now consider deep networks, for which there is no quenched disorder other than the training inputs, and $\mathbb{E}_{\{\mathcal{V}^a\}} = \mathbb{E}_{\{\mathbf{U}_1^a, \dots, \mathbf{U}_\ell^a, \mathbf{v}^a\}}$. In this case, we have

$$S = \exp\left(\frac{1}{2} \mathbf{1}^\top \hat{\mathbf{Q}} \mathbf{1} \|\mathbf{w}_*\|_2^2\right) \mathbb{E}_{\{\mathbf{U}_1^a, \dots, \mathbf{U}_\ell^a, \mathbf{v}^a\}} \exp\left(\frac{1}{2} \sum_{a,b} \hat{Q}^{ab} \mathbf{w}^a \cdot \mathbf{w}^b - \sum_{a,b} \hat{Q}^{ab} \mathbf{w}_* \cdot \mathbf{w}^a\right). \quad (\text{A24})$$

As we did for the RF model, we start by integrating out \mathbf{U}_1^a . We introduce analogous order parameters

$$C_1^{ab} \equiv \frac{1}{n_1 \dots n_\ell} (\mathbf{U}_2^a \dots \mathbf{U}_\ell^a \mathbf{v}^a) \cdot (\mathbf{U}_2^b \dots \mathbf{U}_\ell^b \mathbf{v}^b). \quad (\text{A25})$$

However, importantly, as the weights \mathbf{U}_j^a are annealed rather than quenched, the covariance of h_j^a is replica diagonal. For clarity of notation, we define the diagonal matrix

$$D_1^{ab} = \delta_{ab} C_1^{ab}. \quad (\text{A26})$$

Then, introducing a corresponding diagonal matrix of Lagrange multipliers $\hat{\mathbf{D}}_1$, we have

$$S = \int \frac{d\mathbf{D}_1 d\hat{\mathbf{D}}_1}{(4\pi i/n_1)^{m/2}} \exp\left(-\frac{n_1}{2} \text{tr}(\mathbf{D}_1 \hat{\mathbf{D}}_1) - \frac{d}{2} \log \det(\mathbf{I}_m - \sigma^2 \mathbf{D}_1 \hat{\mathbf{Q}}) + \frac{\|\mathbf{w}_*\|_2^2}{2} \mathbf{1}^\top \hat{\mathbf{Q}} (\mathbf{I}_m - \sigma^2 \mathbf{D}_1 \hat{\mathbf{Q}})^{-1} \mathbf{1}\right) \\ \times \mathbb{E}_{\{\mathbf{U}_2^a, \dots, \mathbf{U}_\ell^a, \mathbf{v}^a\}} \exp\left(\frac{1}{2n_2 \dots n_\ell} \sum_a \hat{D}_1^{aa} (\mathbf{U}_2^a \dots \mathbf{U}_\ell^a \mathbf{v}^a)^\top \mathbf{U}_2^a \dots \mathbf{U}_\ell^a \mathbf{v}^a\right) \quad (\text{A27})$$

upon integrating out \mathbf{U}_1^a . We can see that we can follow much the same procedure to integrate out the remaining weights as we did for the random feature model, except for the fact that we only consider the replica-diagonal component of the overlaps, which are redefined to include the replica indices of the hidden layer weights, i.e.,

$$C_l^{ab} \equiv \frac{1}{n_l \dots n_\ell} (\mathbf{U}_{l+1}^a \dots \mathbf{U}_\ell^a \mathbf{v}^a) \cdot (\mathbf{U}_{l+1}^b \dots \mathbf{U}_\ell^b \mathbf{v}^b) \quad (\text{A28})$$

for $l = 1, \dots, \ell - 1$ and

$$C_\ell^{ab} \equiv \frac{1}{n_\ell} \mathbf{v}^a \cdot \mathbf{v}^b. \quad (\text{A29})$$

We therefore obtain

$$\mathbb{E}_{\mathcal{D}} Z_{\text{NN}}^m = \int \frac{d\mathbf{Q} d\hat{\mathbf{Q}}}{(4\pi i/d)^{m(m+1)/2}} \int \frac{d\mathbf{D}_1 d\hat{\mathbf{D}}_1}{(4\pi i/n_1)^{m/2}} \dots \int \frac{d\mathbf{D}_\ell d\hat{\mathbf{D}}_\ell}{(4\pi i/n_\ell)^{m/2}} \exp\left(-\frac{1}{2} dm F(\mathbf{Q}, \hat{\mathbf{Q}}, \mathbf{D}_1, \hat{\mathbf{D}}_1, \dots, \mathbf{D}_\ell, \hat{\mathbf{D}}_\ell)\right) \quad (\text{A30})$$

where F is the same as for the random feature model and the matrices \mathbf{D}_l and $\hat{\mathbf{D}}_l$ are constrained to be replica diagonal. This difference reflects the fact that the hidden layer weights of the NN model are annealed, rather than being quenched as in the RF model.

5. The replica-symmetric Ansatz

In the thermodynamic limit, we evaluate the integral over the appropriate order parameters and Lagrange multipliers for each model via the method of steepest descent. Importantly, we note that the diagonal components of the order parameters Q^{aa} give the posterior-averaged generalization errors of the replicas, as in the thermodynamic limit the mean value of

these parameters is given by the saddle-point equations. Our eventual objective is therefore simply to evaluate the saddle-point values of \mathbf{Q} in the zero-temperature limit, and we will not consider the resulting values of the free energy.

As usual in the replica method, we seek extrema in the limit $m \rightarrow 0$ of a constrained form, known as the RS Ansatz [50]. For all three models, the RS Ansatz for the variables \mathbf{Q} and $\hat{\mathbf{Q}}$ is simply

$$\mathbf{Q}_{\text{RS}} = (Q - q) \mathbf{I}_m + q \mathbf{1} \mathbf{1}^\top, \quad (\text{A31})$$

$$\hat{\mathbf{Q}}_{\text{RS}} = (\hat{Q} - \hat{q}) \mathbf{I}_m + \hat{q} \mathbf{1} \mathbf{1}^\top. \quad (\text{A32})$$

For a deep random feature model, the RS *Ansatz* for the remaining order parameters is

$$\mathbf{C}_{l,\text{RS}} = (C_l - c_l)\mathbf{I}_m + c_l\mathbf{1}\mathbf{1}^\top \quad (l = 1, \dots, \ell), \quad (\text{A33})$$

$$\hat{\mathbf{C}}_{l,\text{RS}} = (\hat{C}_l - \hat{c}_l)\mathbf{I}_m + \hat{c}_l\mathbf{1}\mathbf{1}^\top \quad (l = 1, \dots, \ell), \quad (\text{A34})$$

while, for a deep network, the RS *Ansatz* for the remaining order parameters is

$$\mathbf{D}_{l,\text{RS}} = C_l\mathbf{I}_m \quad (l = 1, \dots, \ell), \quad (\text{A35})$$

$$\hat{\mathbf{D}}_{l,\text{RS}} = \hat{C}_l\mathbf{I}_m \quad (l = 1, \dots, \ell), \quad (\text{A36})$$

as we consider only the replica-diagonal components of the overlaps \mathbf{C}_l . With this *Ansatz*, one can simplify the expressions for the free energy and the saddle-point equations in the limit $m \rightarrow 0$. These manipulations are standard exercises using the properties of RS matrices [50], hence we will only report the results (in Appendix B).

We remark briefly on the conditions under which the RS order parameters make physical sense given their definitions. We must have $Q \geq 0$ and $C_l \geq 0$ for all l , as these quantities are the squares of norms of vectors. If $C_l = 0$ for any l , the norm of the end-to-end weight vector tends to zero, and we must have a trivial solution with $Q = 1$. We must also have $Q - q \geq 0$ and $C_l - c_l \geq 0$ for all l . Moreover, if $Q - q = 0$ (respectively, $C_l - c_l = 0$ for some l), then we must have $q \geq 0$ (respectively, $c_l \geq 0$), to obtain a nontrivial physical solution.

APPENDIX B: SOLUTION OF THE REPLICASYMMETRIC SADDLE-POINT EQUATIONS

In this Appendix, we analyze the replica-symmetric saddle-point equations in the zero-temperature limit.

1. LR model

For simple linear regression, the RS saddle point is given by a four-dimensional system of equations, which decouples into a two-dimensional nonlinear system for the replica-nonuniform components $z \equiv Q - q$ and $\hat{z} \equiv \hat{Q} - \hat{q}$,

$$\hat{z} = -\frac{\alpha}{\beta^{-1} + z}, \quad (\text{B1})$$

$$z = \frac{\sigma^2}{1 - \sigma^2\hat{z}}, \quad (\text{B2})$$

and a linear system for the replica-uniform components q and \hat{q} :

$$\hat{q} = \frac{\alpha(q + \eta^2)}{(\beta^{-1} + z)^2}, \quad (\text{B3})$$

$$q = \frac{1 + \sigma^4\hat{q}}{(1 - \sigma^2\hat{z})^2}. \quad (\text{B4})$$

Using the expression for \hat{z} as a function of z , we obtain the quadratic condition

$$z^2 - [\sigma^2(1 - \alpha) - \beta^{-1}]z - \sigma^2\beta^{-1} = 0. \quad (\text{B5})$$

The two solutions z_{\pm} to this quadratic equation have zero-temperature limits

$$\lim_{\beta \rightarrow \infty} z_{\pm} = \frac{1}{2}\sigma^2(1 - \alpha \pm |\alpha - 1|). \quad (\text{B6})$$

For $\alpha > 1$, z_- is negative, and is therefore unphysical. If $0 < \alpha < 1$,

$$z_+ = \sigma^2(1 - \alpha) + \frac{\alpha}{1 - \alpha} \frac{1}{\beta} + \mathcal{O}(\beta^{-2}), \quad (\text{B7})$$

while if $\alpha > 1$,

$$z_+ = \frac{1}{\alpha - 1} \frac{1}{\beta} + \mathcal{O}(\beta^{-2}). \quad (\text{B8})$$

These are the two low-temperature scalings we would expect to be self-consistent given the saddle-point equations; we could alternatively derive the above solutions by assuming these scalings for z .

Considering the replica-uniform components, we use the expressions for $1 - \sigma^2\hat{z}$ and \hat{q} as functions of z and q to write

$$q = \frac{z^2}{\sigma^4}(1 + \sigma^4\hat{q}) = \frac{z^2}{\sigma^4} + \frac{\alpha z^2}{(\beta^{-1} + z)^2}(q + \eta^2). \quad (\text{B9})$$

For the solution with $z \sim \mathcal{O}(1)$, we then have

$$(1 - \alpha)q = \frac{z^2}{\sigma^4} + \alpha\eta^2 \quad (\text{B10})$$

and hence, recalling that this scaling yields $z = \sigma^2(1 - \alpha)$ and is valid for $0 < \alpha < 1$,

$$q = 1 - \alpha + \frac{\alpha}{1 - \alpha}\eta^2. \quad (\text{B11})$$

For the solution with $z \sim r/\beta$ with $r \sim \mathcal{O}(1)$, we have

$$\left[1 - \alpha\left(\frac{r}{1+r}\right)^2\right]q = \frac{r^2}{\beta^2\sigma^4} + \alpha\left(\frac{r}{1+r}\right)^2\eta^2, \quad (\text{B12})$$

and hence, recalling that this scaling yields $r = 1/(\alpha - 1)$ and is valid for $\alpha > 1$,

$$\frac{\alpha - 1}{\alpha}q = \frac{1}{\beta^2\sigma^4(\alpha - 1)^2} + \frac{1}{\alpha}\eta^2 \quad (\text{B13})$$

or

$$q = \frac{1}{\alpha - 1}\eta^2. \quad (\text{B14})$$

Combining these results, we obtain a zero-temperature solution which gives the result for $\epsilon = Q$ reported in the main text.

2. RF model

For a deep random feature model, the RS saddle point is given by a $4(\ell + 1)$ -dimensional system of equations. As in the single-layer case, this system decouples into two sets of equations. Defining

$$z \equiv Q - q, \quad (\text{B15})$$

$$\hat{z} \equiv \hat{Q} - \hat{q}, \quad (\text{B16})$$

$$w_l \equiv C_l - c_l, \quad (\text{B17})$$

$$\hat{w}_l \equiv \hat{C}_l - \hat{c}_l, \quad (\text{B18})$$

the deviations from uniformity across replicas are determined by the $2(\ell + 1)$ -dimensional closed system

$$\hat{z} = -\frac{\alpha}{\beta^{-1} + z}, \quad (\text{B19})$$

$$z = \frac{\sigma^2 w_1}{1 - \sigma^2 w_1 \hat{z}}, \quad (\text{B20})$$

$$\hat{w}_1 = \frac{\sigma^2}{\gamma_1} \frac{\hat{z}}{1 - \sigma^2 w_1 \hat{z}}, \quad (\text{B21})$$

$$\hat{w}_l = \frac{\gamma_{l-1}}{\gamma_l} \frac{\hat{w}_{l-1}}{1 - \hat{w}_{l-1} w_l} \quad (l = 2, \dots, \ell), \quad (\text{B22})$$

$$w_l = \frac{w_{l+1}}{1 - w_{l+1} \hat{w}_l} \quad (l = 1, \dots, \ell - 1), \quad (\text{B23})$$

$$w_\ell = \frac{1}{1 - \hat{w}_\ell}. \quad (\text{B24})$$

We can then solve the remaining equations for the replica-uniform components,

$$\hat{q} = \frac{\alpha(q + \eta^2)}{(\beta^{-1} + Q - q)^2}, \quad (\text{B25})$$

$$q = \frac{1 + \sigma^2 c_1 + \sigma^4 w_1^2 \hat{q}}{(1 - \sigma^2 w_1 \hat{z})^2}, \quad (\text{B26})$$

$$\hat{c}_1 = \frac{\sigma^2 \hat{q} + \hat{z}^2 (1 + \sigma^2 c_1)}{\gamma_1 (1 - \sigma^2 w_1 \hat{z})^2}, \quad (\text{B27})$$

$$\hat{c}_l = \frac{\gamma_{l-1}}{\gamma_l} \frac{\hat{c}_{l-1} + \hat{w}_{l-1}^2 c_l}{(1 - \hat{w}_{l-1} w_l)^2} \quad (l = 2, \dots, \ell), \quad (\text{B28})$$

$$c_l = \frac{c_{l+1} + w_{l+1}^2 \hat{c}_l}{(1 - w_{l+1} \hat{w}_l)^2} \quad (l = 1, \dots, \ell - 1), \quad (\text{B29})$$

$$c_\ell = \frac{\hat{c}_\ell}{(1 - \hat{w}_\ell)^2}, \quad (\text{B30})$$

for fixed values of these parameters. Importantly, we note that this is a set of $2(\ell + 1)$ linear equations for $2(\ell + 1)$ variables.

a. Solving for the replica-nonuniform components

We first consider the replica-nonuniform components. We start by noting that the equations for z and \hat{z} yield

$$w_1 = \frac{1}{\sigma^2} \frac{z}{1 + z\hat{z}}, \quad (\text{B31})$$

and hence the equation for \hat{w}_1 yields

$$\hat{w}_1 = \frac{\sigma^2}{\gamma_1} \hat{z} (1 + z\hat{z}). \quad (\text{B32})$$

If $\ell = 1$, then we are nearly done. The condition $w_1 = 1/(1 - \hat{w}_1)$ gives

$$\hat{w}_1 = 1 - \frac{1}{w_1} = \frac{z - \sigma^2(1 + z\hat{z})}{z} \quad (\text{B33})$$

whence

$$\frac{\sigma^2}{\gamma_1} \hat{z} (1 + z\hat{z}) = \frac{z - \sigma^2(1 + z\hat{z})}{z}, \quad (\text{B34})$$

and therefore

$$z = \sigma^2 \frac{(1 + z\hat{z})(\gamma_1 + z\hat{z})}{\gamma_1}. \quad (\text{B35})$$

If $\ell > 1$, we observe that a solution with any $w_l = 0$ must have all $w_l = 0$ and $z = 0$. Similarly, a solution with one $\hat{w}_l = 0$ must have all $\hat{w}_l = 0$ and $\hat{z} = 0$. As $w_\ell = 1/(1 - \hat{w}_\ell)$, these situations cannot coexist. Moreover, neither is self-consistent unless $\alpha = 0$ or β is strictly infinite or zero.

With this observation in mind, we will eliminate the Lagrange multipliers \hat{w}_l . Formally defining $w_{l+1} \equiv 1$ for convenience, we have

$$\hat{w}_l = \frac{w_l - w_{l+1}}{w_l w_{l+1}} \quad (\text{B36})$$

for $l = 1, \dots, \ell$. Then, for $l = 2, \dots, \ell$, the equation

$$\hat{w}_l = \frac{\gamma_{l-1}}{\gamma_l} \frac{\hat{w}_{l-1}}{1 - \hat{w}_{l-1} w_l} \quad (\text{B37})$$

yields

$$\frac{w_l - w_{l+1}}{w_l w_{l+1}} = \frac{\gamma_{l-1}}{\gamma_l} \frac{w_{l-1}}{w_l} \frac{w_{l-1} - w_l}{w_{l-1} w_l}. \quad (\text{B38})$$

We now consider the equation

$$\hat{w}_1 = \frac{\sigma^2}{\gamma_1} \hat{z} (1 + z\hat{z}). \quad (\text{B39})$$

As

$$1 + z\hat{z} = \frac{1}{\sigma^2} \frac{z}{w_1}, \quad (\text{B40})$$

we can rewrite this as

$$\hat{w}_1 = \frac{z\hat{z}}{\gamma_1 w_1}, \quad (\text{B41})$$

which in turn implies that

$$\frac{w_1 - w_2}{w_1 w_2} = \frac{z\hat{z}}{\gamma_1 w_1}. \quad (\text{B42})$$

Then,

$$\frac{w_2 - w_3}{w_2 w_3} = \frac{\gamma_1}{\gamma_2} \frac{w_1}{w_2} \frac{w_1 - w_2}{w_1 w_2} = \frac{z\hat{z}}{\gamma_2 w_2} \quad (\text{B43})$$

and hence we can see that

$$\hat{w}_l = \frac{w_l - w_{l+1}}{w_l w_{l+1}} = \frac{z\hat{z}}{\gamma_l w_l} \quad (\text{B44})$$

for $l = 1, \dots, \ell$. This yields the backward recurrence

$$w_l = \frac{\gamma_l + z\hat{z}}{\gamma_l} w_{l+1}, \quad (\text{B45})$$

which can be solved using the end-point condition $w_{\ell+1} \equiv 1$, yielding

$$w_l = \frac{(\gamma_l + z\hat{z})(\gamma_{l+1} + z\hat{z}) \cdots (\gamma_\ell + z\hat{z})}{\gamma_l \gamma_{l+1} \cdots \gamma_\ell}. \quad (\text{B46})$$

In particular,

$$w_1 = \frac{(\gamma_1 + z\hat{z})(\gamma_2 + z\hat{z}) \cdots (\gamma_\ell + z\hat{z})}{\gamma_1 \gamma_2 \cdots \gamma_\ell}, \quad (\text{B47})$$

which yields a self-consistent equation for z

$$z = \sigma^2 \frac{(1 + z\hat{z})(\gamma_1 + z\hat{z})(\gamma_2 + z\hat{z}) \cdots (\gamma_\ell + z\hat{z})}{\gamma_1 \gamma_2 \cdots \gamma_\ell} \quad (\text{B48})$$

using the condition

$$\hat{z} = -\frac{\alpha}{\beta^{-1} + z}. \quad (\text{B49})$$

This coincides with the result we obtained earlier for $\ell = 1$.

As in the single-layer case, it can be seen that the self-consistent scalings for z in the zero-temperature limit are $z \sim \mathcal{O}(1)$ and $\mathcal{O}(1/\beta)$. If we take $\beta \rightarrow \infty$ with $z \sim \mathcal{O}(1)$, we simply have $z\hat{z} \rightarrow -\alpha$, which gives

$$z = \sigma^2 \frac{(1-\alpha)(\gamma_1-\alpha)(\gamma_2-\alpha)\cdots(\gamma_\ell-\alpha)}{\gamma_1\gamma_2\cdots\gamma_\ell}. \quad (\text{B50})$$

This scaling gives

$$w_l \rightarrow \frac{(\gamma_l-\alpha)(\gamma_{l+1}-\alpha)\cdots(\gamma_\ell-\alpha)}{\gamma_l\gamma_{l+1}\cdots\gamma_\ell} \sim \mathcal{O}(1) \quad (\text{B51})$$

for all l . As physical solutions have $z \geq 0$ and all $w_l \geq 0$, this solution is sensible in the regime $\alpha < \min\{1, \gamma_1, \gamma_2, \dots, \gamma_\ell\}$.

If we take $z \sim r/\beta$ for $r \sim \mathcal{O}(1)$, we have

$$z\hat{z} \rightarrow -\frac{\alpha r}{1+r} \quad (\text{B52})$$

and the limiting equation

$$0 = \frac{\sigma^2}{\gamma_1\gamma_2\cdots\gamma_\ell} \left(1 - \frac{\alpha r}{1+r}\right) \prod_{l=1}^{\ell} \left(\gamma_l - \frac{\alpha r}{1+r}\right). \quad (\text{B53})$$

This yields $\ell + 1$ solutions

$$r_0 = \frac{1}{\alpha - 1}, \quad (\text{B54})$$

$$r_{l_*} = \frac{\gamma_{l_*}}{\alpha - \gamma_{l_*}} \quad (l_* = 1, \dots, \ell). \quad (\text{B55})$$

For the zeroth solution with $r_0 = \frac{1}{\alpha-1}$, we have $z\hat{z} \rightarrow -1$, and thus

$$w_l \rightarrow \frac{(\gamma_l-1)(\gamma_{l+1}-1)\cdots(\gamma_\ell-1)}{\gamma_l\gamma_{l+1}\cdots\gamma_\ell} \sim \mathcal{O}(1). \quad (\text{B56})$$

This solution is therefore physical for $\alpha > 1$ and all $\gamma_l > 1$. For the l_* th such solution, we have $z\hat{z} \rightarrow -\gamma_{l_*}$, hence

$$w_l \rightarrow \frac{(\gamma_l - \gamma_{l_*})(\gamma_{l+1} - \gamma_{l_*})\cdots(\gamma_\ell - \gamma_{l_*})}{\gamma_l\gamma_{l+1}\cdots\gamma_\ell}. \quad (\text{B57})$$

Thus, we have $w_l \rightarrow 0$ for all $l \leq l_*$. For $l > l_*$, $w_l \sim \mathcal{O}(1)$, and we must have $\gamma_l \geq \gamma_{l_*}$ for all $l > l_*$ such that $w_l \geq 0$. Moreover, we must have $\alpha > \gamma_{l_*}$, such that $r_{l_*} > 0$. We will obtain further conditions on the validity of these solutions from solving for the replica-uniform components.

b. Solving for the replica-uniform components

We now consider the linear system of Eqs. Eq. (B25) that determines the replica-uniform components in terms of the

nonuniform components. We start by noting that we can let \hat{c}_1 express a function of q alone, eliminating the dependence on c_1 using the equation for q :

$$\hat{c}_1 = \frac{\sigma^2}{\gamma_1} \hat{z}^2 \left(\frac{1 + 2z\hat{z}}{\alpha} (q + \eta^2) + q \right) \quad (\text{B58})$$

where we have used the fact that $\hat{q} = \hat{z}^2(q + \eta^2)/\alpha$.

If $\ell = 1$, we can use the equation $w_1 = 1/(1 - \hat{w}_1)$ to obtain

$$c_1 = \frac{\hat{c}_1}{(1 - \hat{w}_1)^2} = w_1^2 \hat{c}_1, \quad (\text{B59})$$

which will give a closed equation for q . If $\ell > 1$, our task is somewhat more complex. We eliminate the Lagrange multipliers via

$$\hat{c}_l = \left(\frac{1 - w_{l+1}\hat{w}_l}{w_{l+1}} \right)^2 c_l - \frac{1}{w_{l+1}^2} c_{l+1} \quad (l = 1, \dots, \ell) \quad (\text{B60})$$

where we have defined $w_{\ell+1} \equiv 1$ and $c_{\ell+1} \equiv 0$ for convenience. Then, for $l = 2, \dots, \ell$, the equation

$$\hat{c}_l = \frac{\gamma_{l-1}}{\gamma_l} \frac{\hat{c}_{l-1} + \hat{w}_{l-1}^2 c_l}{(1 - \hat{w}_{l-1} w_l)^2} \quad (\text{B61})$$

yields the three-term recurrence

$$\frac{\gamma_{l-1}}{\gamma_l w_l^2} c_{l-1} = \left[\frac{(1 - w_{l+1}\hat{w}_l)^2}{w_{l+1}^2} + \frac{\gamma_{l-1}}{\gamma_l} \frac{1 - w_l^2 \hat{w}_{l-1}^2}{w_l^2 (1 - \hat{w}_{l-1} w_l)^2} \right] c_l - \frac{1}{w_{l+1}^2} c_{l+1} \quad (\text{B62})$$

with initial difference condition

$$\hat{c}_1(q) = \left(\frac{1 - w_2 \hat{w}_1}{w_2} \right)^2 c_1 - \frac{1}{w_2^2} c_2 \quad (\text{B63})$$

and end-point condition $c_{\ell+1} = 0$. Substituting in $\hat{w}_l = \frac{z\hat{z}}{\gamma_l w_l}$ and using the backward recurrence $w_l = \frac{\gamma_l + z\hat{z}}{\gamma_l} w_{l+1}$, we have

$$\frac{\gamma_{l-1}}{\gamma_l} c_{l-1} = \left[1 + \frac{\gamma_{l-1}}{\gamma_l} \frac{(\gamma_{l-1} + z\hat{z})^2 - (z\hat{z})^2}{\gamma_{l-1}^2} \right] c_l - \frac{(\gamma_l + z\hat{z})^2}{\gamma_l^2} c_{l+1}. \quad (\text{B64})$$

Similarly, we can simplify the initial difference condition to

$$\hat{c}_1(q) = \frac{1}{w_1^2} c_1 - \frac{1}{w_2^2} c_2, \quad (\text{B65})$$

and hence, substituting in $w_1 = \frac{1}{\sigma^2} \frac{z}{1+z\hat{z}}$, we have

$$w_1^2 \hat{c}_1(q) = c_1 - \left(\frac{\gamma_1 + z\hat{z}}{\gamma_1} \right)^2 c_2. \quad (\text{B66})$$

To simplify our remaining task, we define new variables u_l such that

$$c_l = \gamma_l w_1^2 \hat{c}_1(q) u_l. \quad (\text{B67})$$

If $\ell = 1$, we simply have $u_1 = 1/\gamma_1$. For $\ell > 1$, these variables are determined by the recurrence

$$\frac{\gamma_{l-1}}{\gamma_l} u_{l-1} = \left[1 + \frac{\gamma_{l-1}}{\gamma_l} \frac{(\gamma_{l-1} + z\hat{z})^2 - (z\hat{z})^2}{\gamma_{l-1}^2} \right] u_l - \frac{(\gamma_l + z\hat{z})^2}{\gamma_l^2} u_{l+1} \quad (\text{B68})$$

with initial difference condition

$$\frac{1}{\gamma_1} = u_1 - \left(\frac{\gamma_1 + z\hat{z}}{\gamma_1} \right)^2 u_2 \quad (\text{B69})$$

and end-point condition $u_{\ell+1} = 0$. We note that the initial difference and end-point conditions give the consistent result $u_1 = 1/\gamma_1$ when $\ell = 1$.

Given a solution to the recurrence for the variables u_l , we then have a closed equation for q :

$$q = (1 + z\hat{z})^2 + (z\hat{z})^2 \left(\frac{(1 + \alpha + 2z\hat{z})u_1 + 1}{\alpha} (q + \eta^2) - \eta^2 \right). \quad (\text{B70})$$

With this solution in hand, we can then obtain c_l via the relation $c_l = \gamma_l w_1^2 \hat{c}_1(q) u_l$.

We now consider the zero-temperature limits of interest. With $z \sim \mathcal{O}(1)$, we have $z\hat{z} \rightarrow -\alpha$. The limiting equation for q is then

$$q = (1 - \alpha)^2 + \alpha[(1 - \alpha)u_1 + 1](q + \eta^2) - \eta^2 \alpha^2 u_1, \quad (\text{B71})$$

which yields

$$q = (1 - \alpha) \left[1 + \frac{\alpha u_1}{1 - \alpha u_1} \right] + \left[\frac{\alpha}{1 - \alpha} + \frac{\alpha u_1}{1 - \alpha u_1} \right] \eta^2. \quad (\text{B72})$$

Considering the recurrence for u_l , we have

$$\frac{\gamma_{l-1}}{\gamma_l} u_{l-1} = \frac{\gamma_l + \gamma_{l-1} - 2\alpha}{\gamma_l} u_l - \frac{(\gamma_l - \alpha)^2}{\gamma_l^2} u_{l+1} \quad (\text{B73})$$

for $l = 2, \dots, \ell$, and the initial difference condition

$$\frac{1}{\gamma_1} = u_1 - \left(\frac{\gamma_1 - \alpha}{\gamma_1} \right)^2 u_2. \quad (\text{B74})$$

We can reexpress this recurrence as

$$\frac{\gamma_l - \alpha}{\gamma_l} u_{l+1} - u_l = \frac{\gamma_{l-1}}{\gamma_l} \frac{\gamma_l}{\gamma_l - \alpha} \left(\frac{\gamma_{l-1} - \alpha}{\gamma_{l-1}} u_l - u_{l-1} \right), \quad (\text{B75})$$

which can easily be iterated backward, yielding

$$u_l = \frac{\gamma_l - \alpha}{\gamma_l} u_{l+1} + \frac{1}{\gamma_l} \frac{\gamma_l}{\gamma_l - \alpha} \frac{\gamma_{l-1}}{\gamma_{l-1} - \alpha} \dots \frac{\gamma_2}{\gamma_2 - \alpha} [\gamma_1 u_1 - (\gamma_1 - \alpha) u_2] \quad (\text{B76})$$

for $l = 2, \dots, \ell$. Then, the termination condition $u_{\ell+1} = 0$ implies that

$$u_\ell = \frac{1}{\gamma_\ell} \frac{\gamma_\ell}{\gamma_\ell - \alpha} \frac{\gamma_{\ell-1}}{\gamma_{\ell-1} - \alpha} \dots \frac{\gamma_2}{\gamma_2 - \alpha} [\gamma_1 u_1 - (\gamma_1 - \alpha) u_2], \quad (\text{B77})$$

and hence, iterating one step backwards, we find that

$$u_{\ell-1} = \left(\frac{1}{\gamma_\ell - \alpha} + \frac{1}{\gamma_{\ell-1} - \alpha} \right) \frac{\gamma_{\ell-2}}{\gamma_{\ell-2} - \alpha} \dots \frac{\gamma_2}{\gamma_2 - \alpha} \times [\gamma_1 u_1 - (\gamma_1 - \alpha) u_2]. \quad (\text{B78})$$

It is now easy to see that we can iterate this process backwards, yielding

$$u_l = \left(\sum_{j=l}^{\ell} \frac{1}{\gamma_j - \alpha} \right) \frac{\gamma_{l-1}}{\gamma_{l-1} - \alpha} \dots \frac{\gamma_2}{\gamma_2 - \alpha} \times [\gamma_1 u_1 - (\gamma_1 - \alpha) u_2] \quad (\text{B79})$$

for $l = 2, \dots, \ell$, with

$$u_2 = [\gamma_1 u_1 - (\gamma_1 - \alpha) u_2] \sum_{j=2}^{\ell} \frac{1}{\gamma_j - \alpha} \quad (\text{B80})$$

in particular. Using the initial difference condition to express u_2 in terms of u_1 , we then obtain a closed equation for u_1 :

$$u_1 - \frac{1}{\gamma_1} = \frac{\gamma_1 - \alpha}{\gamma_1} (1 - \alpha u_1) \sum_{j=2}^{\ell} \frac{1}{\gamma_j - \alpha}. \quad (\text{B81})$$

As this equation is linear, it is easy to solve, yielding

$$1 - \alpha u_1 = \frac{\gamma_1 - \alpha}{\gamma_1 + \alpha(\gamma_1 - \alpha) \sum_{j=2}^{\ell} \frac{1}{\gamma_j - \alpha}} \quad (\text{B82})$$

$$= \frac{1}{1 + \sum_{j=1}^{\ell} \frac{\alpha}{\gamma_j - \alpha}} \quad (\text{B83})$$

under the assumption that $\alpha \neq \gamma_l$ for all l . Then, we have

$$\frac{\alpha u_1}{1 - \alpha u_1} = \sum_{j=1}^{\ell} \frac{\alpha}{\gamma_j - \alpha}, \quad (\text{B84})$$

which yields

$$q = (1 - \alpha) \left(1 + \sum_{j=1}^{\ell} \frac{\alpha}{\gamma_j - \alpha} \right) + \left(\frac{\alpha}{1 - \alpha} + \sum_{j=1}^{\ell} \frac{\alpha}{\gamma_j - \alpha} \right) \eta^2. \quad (\text{B85})$$

Moreover, for $l = 2, \dots, \ell$, we have the solution

$$u_l = \left(\sum_{j=l}^{\ell} \frac{1}{\gamma_j - \alpha} \right) \frac{\gamma_{l-1}}{\gamma_{l-1} - \alpha} \dots \frac{\gamma_1}{\gamma_1 - \alpha} (1 - \alpha u_1) \quad (\text{B86})$$

in terms of the solution for $1 - \alpha u_1$. Then, in terms of these solutions for u_l , we have

$$c_l = \frac{\alpha}{\sigma^2} \frac{1 - \alpha + \eta^2}{1 - \alpha} \left(1 + \sum_{j=1}^{\ell} \frac{\alpha}{\gamma_j - \alpha} \right) u_l. \quad (\text{B87})$$

We now consider the solutions with $z \sim r/\beta$ for $r \sim \mathcal{O}(1)$. We first consider the solution with

$$r_0 = \frac{1}{\alpha - 1}. \quad (\text{B88})$$

For this solution, $z\hat{z} \rightarrow -1$, and the limiting self-consistent equation for q reduces to

$$q = \frac{(\alpha - 1)u_1 + 1}{\alpha}(q + \eta^2) - \eta^2 u_1, \quad (\text{B89})$$

which yields

$$q = \frac{1}{\alpha - 1}\eta^2 \quad (\text{B90})$$

for any u_1 . This is non-negative throughout the expected region of physical validity ($\alpha > 1$), and therefore the overall solution makes sense given that $z \rightarrow 0$ in this regime. The recurrence for u_l simplifies to

$$\frac{\gamma_{l-1}}{\gamma_l} u_{l-1} = \left[1 + \frac{\gamma_{l-1}(\gamma_{l-1} - 1)^2 - 1}{\gamma_l \gamma_{l-1}^2} \right] u_l - \frac{(\gamma_l - 1)^2}{\gamma_l^2} u_{l+1} \quad (\text{B91})$$

with initial difference condition

$$\frac{1}{\gamma_1} = u_1 - \left(\frac{\gamma_1 - 1}{\gamma_1} \right)^2 u_2 \quad (\text{B92})$$

and end-point condition $u_{\ell+1} = 0$. This is exactly analogous to the recurrence we obtained when considering the solution with $z \sim \mathcal{O}(1)$ with α set to 1, hence we conclude immediately that u_1 is given by

$$1 - u_1 = \frac{1}{1 + \sum_{j=1}^{\ell} \frac{1}{\gamma_{j-1}}}, \quad (\text{B93})$$

while

$$u_l = \left(\sum_{j=1}^{\ell} \frac{1}{\gamma_{j-1}} \right) \frac{\gamma_{l-1}}{\gamma_{l-1} - 1} \cdots \frac{\gamma_1}{\gamma_1 - 1} (1 - u_1) \quad (\text{B94})$$

for $l = 2, \dots, \ell$. These results are positive throughout the region of interest, i.e., $\gamma_l > 1$ for all l . For this solution, we must be somewhat careful in simplifying the equation for c_l . We have the limiting equation

$$c_l = \frac{1}{\sigma^2} \left[1 + \frac{1}{\alpha} q + \sigma^2 c_1 \right] u_l. \quad (\text{B95})$$

This yields a self-consistent equation for c_1 , which gives

$$c_1 = \frac{1}{\sigma^2} \frac{q + \alpha}{\alpha} \frac{u_1}{1 - u_1} = \frac{1}{\sigma^2} \frac{q + \alpha}{\alpha} \sum_{l=1}^{\ell} \frac{1}{\gamma_{l-1}}. \quad (\text{B96})$$

Then, we have

$$c_l = \frac{1}{\sigma^2} \frac{q + \alpha}{\alpha} \frac{u_l}{1 - u_1} = \frac{1}{\sigma^2} \left(1 + \frac{\eta^2}{\alpha(\alpha - 1)} \right) \frac{u_l}{1 - u_1}, \quad (\text{B97})$$

which gives

$$c_l = \frac{1}{\sigma^2} \left(1 + \frac{\eta^2}{\alpha(\alpha - 1)} \right) \frac{\gamma_{l-1}}{\gamma_{l-1} - 1} \cdots \frac{\gamma_1}{\gamma_1 - 1} \sum_{j=1}^{\ell} \frac{1}{\gamma_{j-1}} \quad (\text{B98})$$

for all l , where the empty product is interpreted as unity. These results are positive throughout the region of physical validity

we expect from our analysis of the replica-nonuniform components; recalling that $w_l > 0$ for these solutions, no further conditions are imposed.

We now consider the solutions with

$$r_{l_*} = \frac{\gamma_{l_*}}{\alpha - \gamma_{l_*}} \quad (\text{B99})$$

for some $l_* = 1, \dots, \ell$. For these solutions, we have $z\hat{z} \rightarrow -\gamma_{l_*}$ and $w_l \rightarrow 0$ for all $l \leq l_*$. From our previous analysis, we have the condition $\gamma_l \geq \gamma_{l_*}$ for all $l > l_*$. If $l_* = 1$, the initial difference condition reduces to

$$u_1 = \frac{1}{\gamma_1} \quad (\text{B100})$$

and

$$\begin{aligned} \frac{\gamma_{l-1}}{\gamma_l} u_{l-1} = & \left[1 + \frac{\gamma_{l-1}(\gamma_{l-1} - \gamma_l)^2 - \gamma_l^2}{\gamma_l \gamma_{l-1}^2} \right] u_l \\ & - \frac{(\gamma_l - \gamma_l)^2}{\gamma_l^2} u_{l+1}. \end{aligned} \quad (\text{B101})$$

If $l_* > 1$, we have the recurrence

$$\begin{aligned} \frac{\gamma_{l-1}}{\gamma_l} u_{l-1} = & \left[1 + \frac{\gamma_{l-1}(\gamma_{l-1} - \gamma_{l_*})^2 - \gamma_{l_*}^2}{\gamma_l \gamma_{l-1}^2} \right] u_l \\ & - \frac{(\gamma_l - \gamma_{l_*})^2}{\gamma_l^2} u_{l+1} \end{aligned} \quad (\text{B102})$$

with initial difference condition

$$\frac{1}{\gamma_1} = u_1 - \left(\frac{\gamma_1 - \gamma_{l_*}}{\gamma_1} \right)^2 u_2 \quad (\text{B103})$$

and end-point condition $u_{\ell+1} = 0$. Precisely at l_* , we have the simplification

$$u_{l_*-1} = \frac{\gamma_{l_*-1} - \gamma_{l_*}}{\gamma_{l_*-1}} u_{l_*}. \quad (\text{B104})$$

Iterating one step backward, we find that

$$u_{l_*-2} = \frac{\gamma_{l_*-2} - \gamma_{l_*}}{\gamma_{l_*-2}} u_{l_*-1}. \quad (\text{B105})$$

It is then easy to see that we can iterate further back to obtain, for $l < l_*$,

$$u_l = \frac{\gamma_l - \gamma_{l_*}}{\gamma_l} u_{l+1}. \quad (\text{B106})$$

In particular, we have

$$u_1 = \frac{\gamma_1 - \gamma_{l_*}}{\gamma_1} u_2 \quad (\text{B107})$$

and hence the initial difference equation yields

$$u_1 = \frac{1}{\gamma_{l_*}}. \quad (\text{B108})$$

For $2 \leq l \leq l_*$, this yields the solution

$$u_l = \frac{1}{\gamma_{l_*}} \frac{\gamma_l - \gamma_{l_*}}{\gamma_l} \cdots \frac{\gamma_{l-1} - \gamma_{l_*}}{\gamma_{l-1}}. \quad (\text{B109})$$

The limiting equation for q is then

$$q = (1 - \gamma_{l_*})^2 + \gamma_{l_*} \frac{1 + \alpha - \gamma_{l_*}}{\alpha} (q + \eta^2) - \eta^2 \gamma_{l_*}, \quad (\text{B110})$$

which yields

$$q = \alpha \frac{1 - \gamma_{l_*}}{\alpha - \gamma_{l_*}} + \frac{\gamma_{l_*}}{\alpha - \gamma_{l_*}} \eta^2. \quad (\text{B111})$$

By the same reasoning as in our analysis of the case $r = 1/(\alpha - 1)$, we have the limiting closed set of equations

$$c_l = \frac{1}{\sigma^2} \gamma_{l_*}^2 \left[1 + \frac{1}{\alpha} q + \sigma^2 c_1 \right] u_l. \quad (\text{B112})$$

Using the fact that $u_1 = 1/\gamma_{l_*}$, we have

$$c_1 = \frac{1}{\sigma^2} \gamma_{l_*} \left[1 + \frac{1}{\alpha} q + \sigma^2 c_1 \right], \quad (\text{B113})$$

which yields

$$c_1 = \frac{1}{\sigma^2} \frac{q + \alpha}{\alpha} \frac{\gamma_{l_*}}{1 - \gamma_{l_*}}, \quad (\text{B114})$$

and thus, for $l > 1$,

$$c_l = \frac{1}{\sigma^2} \gamma_{l_*}^2 \left[1 + \frac{1}{\alpha} q + \sigma^2 c_1 \right] u_l = \frac{1}{\sigma^2} \frac{q + \alpha}{\alpha} \frac{\gamma_{l_*}^2}{1 - \gamma_{l_*}} u_l. \quad (\text{B115})$$

Then, for $2 \leq l \leq l_*$, we have

$$c_l = \frac{1}{\sigma^2} \frac{q + \alpha}{\alpha} \frac{\gamma_{l_*}}{1 - \gamma_{l_*}} \frac{\gamma_l - \gamma_{l_*}}{\gamma_l} \dots \frac{\gamma_{l-1} - \gamma_{l_*}}{\gamma_{l-1}} \quad (\text{B116})$$

using the solution for u_l obtained above. As $w_l = 0$ for $l \leq l_*$, we must have $c_l \geq 0$ for $l \leq l_*$ in order for these solutions to be physical, hence we conclude that we must have $\gamma_l \geq \gamma_{l_*}$ for all $l < l_*$. As $w_l \geq 0$ for $l > l_*$, we will not obtain further conditions on the physical validity of these solutions by solving for c_l for $l > l_*$, hence we will not attempt to do so.

3. NN model

For a deep network, the RS saddle point is determined by the $2(\ell + 2)$ -dimensional system of equations

$$\hat{z} = -\frac{\alpha}{\beta^{-1} + z}, \quad (\text{B117})$$

$$\hat{q} = \frac{\alpha(q + \eta^2)}{(\beta^{-1} + Q - q)^2}, \quad (\text{B118})$$

$$z = \frac{\sigma^2 C_1}{1 - \sigma^2 C_1 \hat{z}}, \quad (\text{B119})$$

$$q = \frac{1 + \sigma^4 C_1^2 \hat{q}}{(1 - \sigma^2 C_1 \hat{z})^2}, \quad (\text{B120})$$

$$\hat{C}_1 = \frac{\sigma^2}{\gamma_1} \left[\frac{\hat{q} + \hat{z}(1 + \hat{z} - \sigma^2 C_1 \hat{z})}{(1 - \sigma^2 C_1 \hat{z})^2} \right], \quad (\text{B121})$$

$$\hat{C}_l = \frac{\gamma_{l-1}}{\gamma_l} \frac{\hat{C}_{l-1}}{1 - \hat{C}_{l-1} C_l} \quad (l = 2, \dots, \ell), \quad (\text{B122})$$

$$C_l = \frac{C_{l+1}}{1 - C_{l+1} \hat{C}_l} \quad (l = 1, \dots, \ell - 1), \quad (\text{B123})$$

$$C_\ell = \frac{1}{1 - \hat{C}_\ell} \quad (\text{B124})$$

where, as before, we have defined $z \equiv Q - q$ and $\hat{z} \equiv \hat{Q} - \hat{q}$ for brevity. Unlike for the RF model, in this case the replica-uniform and replica-nonuniform components do not decouple

nically. However, we have fewer equations to solve. Moreover, we can exclude solutions with $C_l \sim \mathcal{O}(1/\beta)$, as they will be trivial.

From the condition

$$z = \frac{\sigma^2 C_1}{1 - \sigma^2 C_1 \hat{z}}, \quad (\text{B125})$$

we have

$$\sigma^2 C_1 \hat{z} = \frac{z \hat{z}}{1 + z \hat{z}} \quad (\text{B126})$$

and hence

$$q = \frac{1 + \sigma^4 C_1^2 \hat{q}}{(1 - \sigma^2 C_1 \hat{z})^2} = (1 + z \hat{z})^2 + z^2 \hat{q} \quad (\text{B127})$$

$$= (1 + z \hat{z})^2 + \frac{1}{\alpha} (z \hat{z})^2 (q + \eta^2), \quad (\text{B128})$$

where we have noted that $\hat{q} = \hat{z}^2 (q + \eta^2) / \alpha$.

We now seek to eliminate the Lagrange multipliers \hat{C}_l and all of the order parameters C_l except for C_1 . To do so, we will follow our earlier analysis of the RF model. We define

$$A \equiv z \frac{\hat{q} + \hat{z}(1 + \hat{z} - \sigma^2 C_1 \hat{z})}{(1 - \sigma^2 C_1 \hat{z})} \quad (\text{B129})$$

such that

$$\hat{C}_1 = \frac{A}{\gamma_1 C_1}. \quad (\text{B130})$$

If $\ell = 1$, then we can solve the equation

$$C_1 = \frac{1}{1 - \hat{C}_1} = \frac{1}{1 - \frac{A}{\gamma_1 C_1}} \quad (\text{B131})$$

yielding

$$C_1 = \frac{\gamma_1 + A}{\gamma_1}, \quad (\text{B132})$$

which will allow us to close the equations.

We now consider deeper networks ($\ell > 1$), following the solution techniques we used for the RF model. We observe that a solution with any $C_l = 0$ must have all $C_l = 0$ and $z = 0$. Similarly, a solution with one $\hat{C}_l = 0$ must have all $\hat{C}_l = 0$ and $\hat{z} = 0$. As $C_\ell = 1/(1 - \hat{C}_\ell)$, these situations cannot coexist. Moreover, neither is self-consistent unless $\alpha = 0$ or β is strictly infinite or zero. With this observation in mind, we will eliminate the Lagrange multipliers \hat{C}_l using the same method as we did for the replica-nonuniform components of the Lagrange multipliers in the RF case. Formally defining $C_{l+1} \equiv 1$ for convenience, we have

$$\hat{C}_l = \frac{C_l - C_{l+1}}{C_l C_{l+1}} \quad (\text{B133})$$

for $l = 1, \dots, \ell$. Then, for $l = 2, \dots, \ell$, the equation

$$\hat{C}_l = \frac{\gamma_{l-1}}{\gamma_l} \frac{\hat{C}_{l-1}}{1 - \hat{C}_{l-1} w_l} \quad (\text{B134})$$

yields

$$\frac{C_l - C_{l+1}}{C_l C_{l+1}} = \frac{\gamma_{l-1}}{\gamma_l} \frac{C_{l-1}}{C_l} \frac{C_{l-1} - C_l}{C_{l-1} C_l}. \quad (\text{B135})$$

Using the above-mentioned fact that we can write

$$\hat{C}_1 = \frac{A}{\gamma_1 C_1}, \quad (\text{B136})$$

we can see that this set of equations is analogous to what we obtained for the deviations from uniformity w_l and \hat{w}_l in the RF case (with, in that case, $A = z\hat{z}$). Thus, using the results of our previous calculation, we conclude that

$$C_l = \frac{(\gamma_l + A)(\gamma_{l+1} + A) \cdots (\gamma_\ell + A)}{\gamma_l \gamma_{l+1} \cdots \gamma_\ell}, \quad (\text{B137})$$

which gives us a closed set of equations for z , \hat{z} , q , \hat{q} , and C_1 .

With the scaling $z \sim \mathcal{O}(1)$, we have $z\hat{z} \rightarrow -\alpha$, and the condition on q becomes

$$q = (1 - \alpha)^2 + \alpha(q + \eta^2), \quad (\text{B138})$$

which yields

$$q = 1 - \alpha + \frac{\alpha}{1 - \alpha} \eta^2. \quad (\text{B139})$$

To determine the limiting condition on z , we note that

$$A = \frac{1}{z} \alpha (1 - \alpha + \eta^2) - \alpha \quad (\text{B140})$$

$$z_{\pm} = (1 - \alpha) \frac{\sigma^2(\gamma_1 - \alpha) \pm \sqrt{\sigma^4(\gamma_1 - \alpha)^2 + 4\alpha\gamma_1\sigma^2(1 - \alpha + \eta^2)/(1 - \alpha)}}{2\gamma_1}. \quad (\text{B145})$$

For any $\gamma_1, \sigma > 0$, $z_+ \geq 0$ if $0 < \alpha < 1$, and $z_- \geq 0$ if $\alpha > 1$. However, noting that

$$C_1 = \frac{z}{\sigma^2(1 - \alpha)}, \quad (\text{B146})$$

the z_+ solution yields a non-negative value for C_1 , and is therefore physical for $0 < \alpha < 1$, while the z_- solution yields a nonpositive value, and is therefore unphysical.

For a network with more than a single hidden layer, the polynomial condition cannot be analytically solved in a useful way. However, we can gain some insight by solving it perturbatively in the large-width regime $\gamma_1, \dots, \gamma_\ell \gg \alpha$. Concretely, we introduce a formal expansion parameter λ , and solve the equation

$$z^{\ell+1} = \sigma^2(1 - \alpha) \prod_{l=1}^{\ell} \left[\left(1 - \lambda \frac{\alpha}{\gamma_l}\right) z + \lambda \frac{\alpha}{\gamma_l} (1 - \alpha + \eta^2) \right] \quad (\text{B147})$$

order by order in λ with the *Ansatz*

$$z = \sum_{j=0}^{\infty} z_j \lambda^j. \quad (\text{B148})$$

It is easy to see that the zeroth-order condition yields

$$z_0 = \sigma^2(1 - \alpha), \quad (\text{B149})$$

and that the first-order term yields

$$z_1 = [(1 - \sigma^2)(1 - \alpha) + \eta^2] \sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l}. \quad (\text{B150})$$

in this limit, hence

$$(\gamma_l + A)z = (\gamma_l - \alpha)z + \alpha(1 - \alpha + \eta^2). \quad (\text{B141})$$

Therefore, we have the polynomial condition

$$z^{\ell+1} = \sigma^2(1 - \alpha) \prod_{l=1}^{\ell} \left[\frac{(\gamma_l - \alpha)z + \alpha(1 - \alpha + \eta^2)}{\gamma_l} \right]. \quad (\text{B142})$$

Given a candidate positive solution to this degree- $(\ell + 1)$ polynomial, we must then verify that it yields a positive value for all

$$C_l = \frac{1}{z^l} \prod_{l'=l}^{\ell} \frac{(\gamma_{l'} - \alpha)z + \alpha(1 - \alpha + \eta^2)}{\gamma_{l'}} \quad (\text{B143})$$

in order for it to be a nontrivial physical solution. This implies that we must have

$$\frac{(\gamma_l - \alpha)z + \alpha(1 - \alpha + \eta^2)}{\gamma_l z} > 0 \quad (\text{B144})$$

for all l .

For a network with a single hidden layer ($\ell = 1$), the condition is just a quadratic, with solutions

To go to higher order, it is convenient to specialize to the case of equal hidden layer widths $\gamma_1 = \gamma_2 = \dots = \gamma_\ell = \gamma$, both to simplify the calculations and to make the results easier to interpret. In this case, we can directly define the expansion parameter as $\lambda \equiv \alpha/\gamma$, hence the equation we want to solve becomes

$$z^{\ell+1} = \sigma^2(1 - \alpha) [(1 - \lambda)z + \lambda(1 - \alpha + \eta^2)]^{\ell}. \quad (\text{B151})$$

In this simplified setting, it is relatively straightforward to work out by hand or with the aid of MATHEMATICA that

$$z_2 = (1 - \alpha + \eta^2) \left(\frac{\ell(\ell - 1)\tilde{\sigma}^2}{2} - \frac{\ell(\ell + 1)}{2\tilde{\sigma}^2} + \ell \right) \quad (\text{B152})$$

with $\tilde{\sigma}$ as in Eq. (23), which yields the result reported in the main text.

For solutions with $z \sim \mathcal{O}(\beta^{-1})$ and $C_1 \sim \mathcal{O}(1)$, we have the limiting equation

$$q = \frac{1}{\alpha}(q + \eta^2), \quad (\text{B153})$$

which implies that we should have

$$q = \frac{\eta^2}{\alpha - 1}. \quad (\text{B154})$$

As $z \rightarrow 0$, we must have $q \geq 0$, hence this solution makes sense for all $\alpha > 1$. To solve for C_l for these solutions, it is most convenient to express A in terms of $C_1 \sim \mathcal{O}(1)$. Noting that

$$\hat{C}_1 \rightarrow \frac{1}{\gamma_1 C_1} \frac{1 - \sigma^2 C_1 + \eta^2/(\alpha - 1)}{\sigma^2 C_1}, \quad (\text{B155})$$

we have

$$A \rightarrow \frac{1 - \sigma^2 C_1 + \eta^2 / (\alpha - 1)}{\sigma^2 C_1}, \quad (\text{B156})$$

hence C_1 is determined by degree- $(\ell + 1)$ polynomial

$$\sigma^{2\ell} C_1^{\ell+1} = \prod_{l=1}^{\ell} \frac{1 + \sigma^2 (\gamma_l - 1) C_1 + \eta^2 / (\alpha - 1)}{\gamma_l}. \quad (\text{B157})$$

Given a candidate positive solution for C_1 , we can then determine C_l for all $l > 1$ via

$$C_l = \prod_{l'=1}^{\ell} \frac{\gamma_{l'} + A}{\gamma_{l'}} = \prod_{l'=1}^{\ell} \frac{1 + \sigma^2 (\gamma_{l'} - 1) C_1 + \eta^2 / (\alpha - 1)}{\gamma_{l'} \sigma^2 C_1}. \quad (\text{B158})$$

This implies that we must have

$$C_l > \frac{1}{\sigma^2 (1 - \gamma_l)} \left(1 + \frac{\eta^2}{\alpha - 1} \right) \quad (\text{B159})$$

for all l (including $l = 1$) in order for the candidate solution to be physical and nontrivial. As we are interested only in learning curves, we will not analyze this equation further.

APPENDIX C: DIRECT COMPUTATION OF POSTERIOR EXPECTATIONS FOR LR AND RF MODELS

For the LR and RF models, we can evaluate the zero-temperature posterior expectation in the definition of ϵ analytically. In particular, writing

$$\mathbf{F} \equiv \frac{\sigma}{\sqrt{dn_1 \cdots n_\ell}} \mathbf{U}_1 \cdots \mathbf{U}_\ell \quad (\text{C1})$$

for brevity, we have the posterior moment generating function for \mathbf{v} :

$$\mathcal{Z}(\mathbf{j}) \propto \int d\mathbf{v} \exp \left(-\frac{\beta}{2} \|\mathbf{X}\mathbf{F}\mathbf{v} - \mathbf{y}\|^2 - \frac{1}{2} \|\mathbf{v}\|^2 + \mathbf{j}^\top \mathbf{v} \right) \quad (\text{C2})$$

$$\begin{aligned} &\propto \exp \left(\beta \mathbf{y}^\top \mathbf{X}\mathbf{F}(\mathbf{I}_{n_\ell} + \beta \mathbf{F}^\top \mathbf{X}^\top \mathbf{X}\mathbf{F})^{-1} \mathbf{j} \right. \\ &\quad \left. + \frac{1}{2} \mathbf{j}^\top (\mathbf{I}_{n_\ell} + \beta \mathbf{F}^\top \mathbf{X}^\top \mathbf{X}\mathbf{F})^{-1} \mathbf{j} \right), \end{aligned} \quad (\text{C3})$$

where the implied constants of proportionality are independent of the source \mathbf{j} . Then, as $\mathbf{w} = \sqrt{d}\mathbf{F}\mathbf{v}$, the posterior mean and covariance of the end-to-end weight vector are given as

$$\langle \mathbf{w} \rangle = \sqrt{d} \mathbf{F} (\beta^{-1} \mathbf{I}_{n_\ell} + \mathbf{F}^\top \mathbf{X}^\top \mathbf{X}\mathbf{F})^{-1} \mathbf{F}^\top \mathbf{X}^\top \mathbf{y} \quad (\text{C4})$$

and

$$\langle \mathbf{w}\mathbf{w}^\top \rangle - \langle \mathbf{w} \rangle \langle \mathbf{w} \rangle^\top = d \mathbf{F} (\mathbf{I}_{n_\ell} + \beta \mathbf{F}^\top \mathbf{X}^\top \mathbf{X}\mathbf{F})^{-1} \mathbf{F}^\top, \quad (\text{C5})$$

respectively. We note that $\langle \mathbf{w} \rangle$ is simply the RF ridge regression estimator with ridge parameter $1/\beta$, as

$$\langle \mathbf{v} \rangle = \arg \min_{\mathbf{v}} \left(\|\mathbf{X}\mathbf{F}\mathbf{v} - \mathbf{y}\|^2 + \frac{1}{\beta} \|\mathbf{v}\|^2 \right). \quad (\text{C6})$$

The thermal bias-variance decomposition of the zero-temperature generalization error is then given as

$$\epsilon_b \equiv \lim_{\beta \rightarrow \infty} \frac{1}{d} \|\langle \mathbf{w} \rangle - \mathbf{w}_*\|^2 \quad (\text{C7})$$

$$= \lim_{\beta \rightarrow \infty} \left\| \mathbf{F} (\beta^{-1} \mathbf{I}_{n_\ell} + \mathbf{F}^\top \mathbf{X}^\top \mathbf{X}\mathbf{F})^{-1} \mathbf{F}^\top \mathbf{X}^\top \mathbf{y} - \frac{\mathbf{w}_*}{\sqrt{d}} \right\|^2, \quad (\text{C8})$$

$$\epsilon_v \equiv \lim_{\beta \rightarrow \infty} \frac{1}{d} \text{tr}[\langle \mathbf{w}\mathbf{w}^\top \rangle - \langle \mathbf{w} \rangle \langle \mathbf{w} \rangle^\top] \quad (\text{C9})$$

$$= \lim_{\beta \rightarrow \infty} \text{tr}[\mathbf{F} (\mathbf{I}_{n_\ell} + \beta \mathbf{F}^\top \mathbf{X}^\top \mathbf{X}\mathbf{F})^{-1} \mathbf{F}^\top]. \quad (\text{C10})$$

In terms of these quantities, we have

$$\epsilon = \lim_{d, p, n_1, \dots, n_\ell \rightarrow \infty} \mathbb{E}_{\mathcal{D}} (\epsilon_b + \epsilon_v). \quad (\text{C11})$$

With our data model, \mathbf{X} has rank $\min\{p, d\}$ with probability 1, while \mathbf{F} has rank $\min\{d, n_1, \dots, n_\ell\}$ with probability 1 [68]. Moreover, $\mathbf{X}\mathbf{F}$ has rank $\min\{p, d, n_1, \dots, n_\ell\}$ with probability 1.

If all $n_1, \dots, n_\ell > d$ and $p > d$, both $\mathbf{F}\mathbf{F}^\top$ and $\mathbf{X}^\top \mathbf{X}$ are invertible with probability 1, and, applying the push-through identity [67], we have

$$\epsilon_b = \frac{1}{d} \|\sqrt{d} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} - \mathbf{w}_*\|^2, \quad (\text{C12})$$

$$\epsilon_v = 0. \quad (\text{C13})$$

If $p < \min\{d, n_1, \dots, n_\ell\}$, then the matrix $\mathbf{F}\mathbf{F}^\top \mathbf{X}^\top \mathbf{X}$ will not be invertible, but the matrix $\mathbf{X}\mathbf{F}\mathbf{F}^\top \mathbf{X}^\top$ will be invertible with probability 1, even if $\mathbf{F}\mathbf{F}^\top$ is not. Then, with another application of the push-through identity and the aid of the Woodbury identity [67], we have

$$\epsilon_b = \frac{1}{d} \|\sqrt{d} \mathbf{F}\mathbf{F}^\top \mathbf{X}^\top (\mathbf{X}\mathbf{F}\mathbf{F}^\top \mathbf{X}^\top)^{-1} \mathbf{y} - \mathbf{w}_*\|^2, \quad (\text{C14})$$

$$\epsilon_v = \text{tr}(\mathbf{F}\mathbf{F}^\top) - \text{tr}[\mathbf{F}\mathbf{F}^\top \mathbf{X}^\top (\mathbf{X}\mathbf{F}\mathbf{F}^\top \mathbf{X}^\top)^{-1} \mathbf{X}\mathbf{F}\mathbf{F}^\top]. \quad (\text{C15})$$

Finally, if $p > \min\{n_1, \dots, n_\ell\}$ but $\min\{n_1, \dots, n_\ell\} < d$, the situation is somewhat more complicated. Let $l_{\min} = \arg \min n_l$ be the index of the narrowest layer. Then, let

$$\mathbf{A} = \frac{\sigma}{\sqrt{dn_1 \cdots n_{l_{\min}}}} \mathbf{U}_1 \cdots \mathbf{U}_{l_{\min}} \in \mathbb{R}^{d \times n_{\min}}, \quad (\text{C16})$$

$$\mathbf{B} = \frac{1}{\sqrt{n_{l_{\min}+1} \cdots n_\ell}} \mathbf{U}_{l_{\min}+1} \cdots \mathbf{U}_\ell \in \mathbb{R}^{n_{\min} \times n_\ell} \quad (\text{C17})$$

such that

$$\mathbf{F} = \mathbf{A}\mathbf{B}. \quad (\text{C18})$$

Under the stated assumptions, the matrices $\mathbf{A}^\top \mathbf{X}^\top \mathbf{X}\mathbf{A}$ and $\mathbf{B}\mathbf{B}^\top$ are invertible with probability 1, as is their product. Then, we have

$$\epsilon_b = \frac{1}{d} \|\sqrt{d} \mathbf{A} (\mathbf{A}^\top \mathbf{X}^\top \mathbf{X}\mathbf{A})^{-1} \mathbf{A}^\top \mathbf{X}^\top \mathbf{y} - \mathbf{w}_*\|^2, \quad (\text{C19})$$

$$\epsilon_v = 0. \quad (\text{C20})$$

We observe that, under the rescaling of the feature map $\mathbf{F} \mapsto \sigma \mathbf{F}$ for any $\sigma > 0$, ϵ_b is always constant, while ϵ_v is either identically zero or degree-2 homogeneous in σ . This suggests that we should be able to read off the ridgeless results from our Bayesian replica results. We also note that we have recovered the three-region phase diagram indicated by our replica calculation.

For completeness, we also remark that we can use these results to directly compute ϵ_{LR} without the use of the replica

trick. For the LR model, we have

$$\mathbf{F} = \frac{\sigma}{\sqrt{d}} \mathbf{I}_d \quad (\text{C21})$$

and two phases: $p < d$ and $p > d$. We first consider the regime $p < d$. Using the fact that $\mathbb{E}_{\mathbf{w}_*} \mathbf{w}_* \mathbf{w}_*^\top = \mathbf{I}_d$ and the formula for the expectation of an inverse Wishart matrix with identity scale matrix [68],

$$\mathbb{E}_{\mathbf{X}} \text{tr}[(\mathbf{X}\mathbf{X}^\top)^{-1}] = \frac{p}{d - p - 1}, \quad (\text{C22})$$

a short computation yields

$$\lim_{n, p \rightarrow \infty} \mathbb{E}_{\mathbf{X}, \xi, \mathbf{w}_*} \varepsilon_b = 1 - \alpha + \frac{\alpha}{1 - \alpha} \eta^2, \quad (\text{C23})$$

$$\lim_{n, p \rightarrow \infty} \mathbb{E}_{\mathbf{X}, \xi, \mathbf{w}_*} \varepsilon_v = \sigma^2 (1 - \alpha). \quad (\text{C24})$$

In the regime $p > d$, we have

$$\lim_{n, p \rightarrow \infty} \mathbb{E}_{\mathbf{X}, \xi, \mathbf{w}_*} \varepsilon_b = \frac{1}{\alpha - 1} \eta^2, \quad (\text{C25})$$

$$\lim_{n, p \rightarrow \infty} \mathbb{E}_{\mathbf{X}, \xi, \mathbf{w}_*} \varepsilon_v = 0, \quad (\text{C26})$$

using the fact that

$$\mathbb{E}_{\mathbf{X}} \text{tr}[(\mathbf{X}^\top \mathbf{X})^{-1}] = \frac{d}{p - d - 1} \quad (\text{C27})$$

in this regime. This recovers the result of our replica computation.

We remark that a similar, albeit more complex, procedure would likely allow one to derive the learning curve for a deep RF model rigorously using properties of products of large Gaussian random matrices [58]. However, from a physical perspective, the nonrigorous replica theory approach used here has the advantages of being more transparent and of allowing a relatively unified treatment of NN models.

APPENDIX D: DIRECT COMPUTATION OF POSTERIOR EXPECTATIONS FOR NN MODELS

In this Appendix, we show that the zero-temperature posterior expectation in the definition of ϵ can be evaluated semianalytically for NNs. Our approach mirrors that of our previous work in [36]: we will integrate out the weights of the first hidden layer (\mathbf{U}_1) exactly, yielding expressions for the posterior mean and variance of the end-to-end weight vector in terms of expectations over the remaining weights. These results follow by applying the results of [36] to a test dataset of d examples with trivial data matrix $\hat{\mathbf{X}} = \sqrt{d} \mathbf{I}_d$ and then passing to the zero-temperature limit, but we will provide a detailed derivation for completeness.

Writing

$$\mathbf{f} \equiv \frac{\sigma}{\sqrt{dn_1 \cdots n_\ell}} \mathbf{U}_2 \cdots \mathbf{U}_\ell \mathbf{v} \quad (\text{D1})$$

for brevity, such that $\mathbf{w} = \sqrt{d} \mathbf{U}_1 \mathbf{f}$, we can write the posterior moment generating function of \mathbf{w} as

$$\mathcal{Z}(\mathbf{j}) \propto \mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} \int d\mathbf{U}_1 \exp \left(-\frac{\beta}{2} \|\mathbf{X}\mathbf{U}_1 \mathbf{f} - \mathbf{y}\|^2 - \frac{1}{2} \|\mathbf{U}_1\|^2 + \sqrt{d} \mathbf{j}^\top \mathbf{U}_1 \mathbf{f} \right), \quad (\text{D2})$$

where we discard irrelevant constants of proportionality. This matrix Gaussian integral can be conveniently evaluated through vectorization [69]. Using standard properties of the Kronecker product, we find that

$$\mathcal{Z}(\mathbf{j}) \propto \mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} \left\{ \rho(\|\mathbf{f}\|^2) \exp \left[\beta \sqrt{d} \|\mathbf{f}\|^2 \mathbf{y}^\top (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X} \mathbf{j} + \frac{1}{2} d \|\mathbf{f}\|^2 \mathbf{j}^\top [\mathbf{I}_d - \beta \|\mathbf{f}\|^2 \mathbf{X}^\top (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X}] \mathbf{j} \right] \right\}, \quad (\text{D3})$$

where

$$\rho(\|\mathbf{f}\|^2) \equiv \det(\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1/2} \exp \left(-\frac{1}{2} \beta \mathbf{y}^\top (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} \right). \quad (\text{D4})$$

By varying this result with respect to the source, we thus obtain

$$\langle \mathbf{w} \rangle = \frac{\mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} [\rho \mathbf{z}]}{\mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} \rho} \quad (\text{D5})$$

and

$$\langle \mathbf{w} \mathbf{w}^\top \rangle - \langle \mathbf{w} \rangle \langle \mathbf{w} \rangle^\top = \frac{\mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} \{ \rho d \|\mathbf{f}\|^2 [\mathbf{I}_d - \beta \|\mathbf{f}\|^2 \mathbf{X}^\top (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X}] \}}{\mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} \rho} + \frac{\mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} [\rho \mathbf{z} \mathbf{z}^\top]}{\mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1} \rho} - \langle \mathbf{w} \rangle \langle \mathbf{w} \rangle^\top, \quad (\text{D6})$$

where we have defined

$$\mathbf{z} \equiv \beta \sqrt{d} \|\mathbf{f}\|^2 \mathbf{X}^\top (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} \quad (\text{D7})$$

for brevity. This matches the result of applying [36]'s expressions to a trivial dataset with $\hat{\mathbf{X}} = \sqrt{d} \mathbf{I}_d$.

As for the RF model, we introduce a thermal bias-variance decomposition

$$\varepsilon_b \equiv \lim_{\beta \rightarrow \infty} \frac{1}{d} \|\langle \mathbf{w} \rangle - \mathbf{w}_*\|^2, \quad (\text{D8})$$

$$\varepsilon_v \equiv \lim_{\beta \rightarrow \infty} \frac{1}{d} \text{tr}[\langle \mathbf{w}\mathbf{w}^\top \rangle - \langle \mathbf{w} \rangle \langle \mathbf{w} \rangle^\top]. \quad (\text{D9})$$

For any set of hidden layer widths, $\|\mathbf{f}\|^2$ is almost surely positive. Therefore, as the only matrix inverses present in these expressions are of the form $(\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1}$, the NN model should have two phases: $p < d$ and $p > d$.

If $p < d$, then the matrix $\mathbf{X}\mathbf{X}^\top$ is invertible with probability 1. Then, up to (divergent) multiplicative constants which will cancel in the ratios of expectations, we have the almost-sure pointwise limit

$$\lim_{\beta \rightarrow \infty} \rho \propto \|\mathbf{f}\|^{-p} \exp\left(-\frac{\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y}}{2\|\mathbf{f}\|^2}\right). \quad (\text{D10})$$

Similarly, we have the almost-sure limits

$$\lim_{\beta \rightarrow \infty} \mathbf{z} = \sqrt{d} \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} \quad (\text{D11})$$

and

$$\begin{aligned} \lim_{\beta \rightarrow \infty} \|\mathbf{f}\|^2 \text{tr}[\mathbf{I}_d - \beta \|\mathbf{f}\|^2 \mathbf{X}^\top (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X}] \\ = (d - p) \|\mathbf{f}\|^2. \end{aligned} \quad (\text{D12})$$

Therefore, noting that $\lim_{\beta \rightarrow \infty} \mathbf{z}$ is almost surely a constant function of \mathbf{f} , we have

$$\varepsilon_b = \frac{1}{d} \|\sqrt{d} \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} - \mathbf{w}_*\|^2, \quad (\text{D13})$$

$$\varepsilon_v = (1 - \alpha) d \frac{\mathbb{E}_{\mathcal{W} \setminus \mathcal{U}_1} \{ \|\mathbf{f}\|^{2-p} \exp[-\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} / 2 \|\mathbf{f}\|^2] \}}{\mathbb{E}_{\mathcal{W} \setminus \mathcal{U}_1} \{ \|\mathbf{f}\|^{-p} \exp[-\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} / 2 \|\mathbf{f}\|^2] \}}. \quad (\text{D14})$$

If $p > d$, then the matrix $\mathbf{X}\mathbf{X}^\top$ is invertible with probability zero, but the matrix $\mathbf{X}^\top \mathbf{X}$ is invertible with probability 1. By the Weinstein-Aronzjan identity,

$$\det(\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top) = \det(\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}^\top \mathbf{X}), \quad (\text{D15})$$

hence the determinant factors in ρ will yield a factor of $\|\mathbf{f}\|^{-d}$ in the zero-temperature limit. We must be more careful in considering the exponential term in ρ . Letting the orthonormal eigendecomposition of $\mathbf{X}\mathbf{X}^\top$ be

$$\mathbf{X}\mathbf{X}^\top = \sum_{j=1}^p \chi_j \mathbf{m}_j \mathbf{m}_j^\top, \quad (\text{D16})$$

$$z = \lim_{d, p, n_1, \dots, n_\ell \rightarrow \infty} \mathbb{E}_{\mathcal{D}} \varepsilon_v \quad (\text{D23})$$

$$= (1 - \alpha) \lim_{d, p, n_1, \dots, n_\ell \rightarrow \infty} d \mathbb{E}_{\mathcal{D}} \left\{ \frac{\mathbb{E}_{\mathcal{W} \setminus \mathcal{U}_1} \{ \|\mathbf{f}\|^{2-p} \exp[-\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} / 2 \|\mathbf{f}\|^2] \}}{\mathbb{E}_{\mathcal{W} \setminus \mathcal{U}_1} \{ \|\mathbf{f}\|^{-p} \exp[-\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} / 2 \|\mathbf{f}\|^2] \}} \right\}. \quad (\text{D24})$$

We remark that evaluation of the outer dataset average without resorting to the replica trick seems likely to be challenging.

For a network with a single hidden layer, we can evaluate the average over $\mathcal{W} \setminus \mathcal{U}_1 = \mathbf{v}$ analytically. In this case, we have

$$\|\mathbf{f}\|^2 = \frac{\sigma^2}{n_1 d} \|\mathbf{v}\|^2, \quad (\text{D25})$$

hence

$$d \frac{\mathbb{E}_{\mathbf{v}} \{ \|\mathbf{f}\|^{2-p} \exp[-\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} / 2 \|\mathbf{f}\|^2] \}}{\mathbb{E}_{\mathbf{v}} \{ \|\mathbf{f}\|^{-p} \exp[-\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} / 2 \|\mathbf{f}\|^2] \}} = \frac{\sigma^2 \mathbb{E}_{\mathbf{v}} [\|\mathbf{v}\|^{2-p} \exp(-q^2/2 \|\mathbf{v}\|^2)]}{n_1 \mathbb{E}_{\mathbf{v}} [\|\mathbf{v}\|^{-p} \exp(-q^2/2 \|\mathbf{v}\|^2)]} \quad (\text{D26})$$

we have the low-temperature Neumann series [67]

$$\begin{aligned} \beta (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \\ = \beta \sum_{\{j: \chi_j=0\}} \mathbf{m}_j \mathbf{m}_j^\top + \frac{1}{\|\mathbf{f}\|^2} \sum_{\{j: \chi_j>0\}} \frac{1}{\chi_j} \mathbf{m}_j \mathbf{m}_j^\top + \mathcal{O}(\beta^{-1}), \end{aligned} \quad (\text{D17})$$

hence the divergent null-space projector term $\beta \sum_{\{j: \chi_j=0\}} \mathbf{m}_j \mathbf{m}_j^\top$ does not depend on \mathbf{f} and will therefore cancel in the ratio of expectations. Thus, we have

$$\lim_{\beta \rightarrow \infty} \rho \propto \|\mathbf{f}\|^{-d} \exp\left(-\frac{1}{2\|\mathbf{f}\|^2} \sum_{\{j: \chi_j>0\}} \frac{1}{\chi_j} (\mathbf{m}_j^\top \mathbf{y})^2\right). \quad (\text{D18})$$

By a simple application of the push-through identity, we have the almost-sure pointwise limits

$$\lim_{\beta \rightarrow \infty} \mathbf{z} = \sqrt{d} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \quad (\text{D19})$$

and

$$\lim_{\beta \rightarrow \infty} \|\mathbf{f}\|^2 \text{tr}[\mathbf{I}_d - \beta \|\mathbf{f}\|^2 \mathbf{X}^\top (\mathbf{I}_p + \beta \|\mathbf{f}\|^2 \mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X}] = 0. \quad (\text{D20})$$

Therefore, noting that $\lim_{\beta \rightarrow \infty} \mathbf{z}$ is once again almost surely a constant function of \mathbf{f} , we have

$$\varepsilon_b = \frac{1}{d} \|\sqrt{d} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} - \mathbf{w}_*\|^2, \quad (\text{D21})$$

$$\varepsilon_v = 0. \quad (\text{D22})$$

Comparing this result to the discussion of the LR model in Appendix C, we can see that the bias terms in each phase are identical to those for the LR model, hence we can apply the results given there for their dataset averages. This shows that the learning curve for the NN model is of the form Eq. (28), with

where we have defined

$$q^2 \equiv \frac{n_1 d}{\sigma^2} \mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} \quad (\text{D27})$$

for brevity. Using the fact that $\|\mathbf{v}\|^2 \sim \chi^2(n_1)$ under the prior, we have

$$\frac{\mathbb{E}_{\mathbf{v}}[\|\mathbf{v}\|^{2-p} \exp(-q^2/2\|\mathbf{v}\|^2)]}{\mathbb{E}_{\mathcal{W} \setminus \mathbf{U}_1}[\|\mathbf{v}\|^{-p} \exp(-q^2/2\|\mathbf{v}\|^2)]} = \frac{\int_0^\infty dt t^{(n_1-p)/2-1+1} \exp(-t/2 - q^2/2t)}{\int_0^\infty dt t^{(n_1-p)/2-1} \exp(-t/2 - q^2/2t)} = q \frac{K_{(n_1-p)/2+1}(q)}{K_{(n_1-p)/2}(q)}, \quad (\text{D28})$$

where $K_\nu(z)$ is a modified Bessel function of the second kind [70]. Thus, for a NN with a single hidden layer, we have

$$z = \sigma^2(1 - \alpha) \lim_{d, p, n_1, \dots, n_\ell \rightarrow \infty} \mathbb{E}_{\mathcal{D}} \left\{ \frac{1}{n_1} q \frac{K_{(n_1-p)/2+1}(q)}{K_{(n_1-p)/2}(q)} \right\}, \quad (\text{D29})$$

with q defined as above.

In general, it is not immediately clear how to evaluate the limit of the nested averages in Eq. (D24). However, as argued by Li and Sompolinsky [37] in the case $\gamma_1 = \dots = \gamma_\ell = \gamma$, one can make progress under the assumption that the quantity

$$r \equiv \frac{1}{\sigma^2 \alpha} \mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} \quad (\text{D30})$$

rapidly concentrates about its limiting mean value, which is

$$\lim_{d, p \rightarrow \infty} \mathbb{E}_{\mathbf{X}, \mathbf{w}_*, \xi} r = \frac{1}{\bar{\sigma}^2} \quad (\text{D31})$$

by the results of Appendix C. Then, assuming that the outer dataset average in Eq. (D24) can be evaluated by replacing all occurrences of r with $\bar{\sigma}^{-2}$, let us make a layer-by-layer saddle-point approximation of the integrals over $\mathcal{W} \setminus \mathbf{U}_1$. Li and Sompolinsky's [37] general fixed-data analysis is not set up by first integrating out \mathbf{U}_1 as above, but, as discussed in our previous work [36] for the case $\ell = 1$, can be related to this approach. Moreover, their analysis of the Gaussian covariate model (presented in Appendix D of the Supplemental Material of [37]) is framed in terms of an approximation in which correlations are neglected, but amounts to the concentration assumption stated above.

Then, by Eqs. (26) and (27) of the main text of [37], or by Eqs. (24), (26), and (30) of the Supplemental Material of [37], the result of their approximation is that

$$z = \sigma^2(1 - \alpha)u^\ell \quad (\text{D32})$$

for u a solution of

$$1 - u = \lambda(1 - u^{-\ell} \bar{\sigma}^{-2}), \quad (\text{D33})$$

with $\lambda \equiv \alpha/\gamma$. We would like to show that this is consistent with the result of our RS calculation, which implies that z should be a non-negative root of

$$z^{\ell+1} = \sigma^2(1 - \alpha)[(1 - \lambda)z + \lambda(1 - \alpha + \eta^2)]^\ell. \quad (\text{D34})$$

Substituting in the Ansatz $z = \sigma^2(1 - \alpha)u^\ell$, we have

$$u^{\ell(\ell+1)} = [(1 - \lambda)u^\ell + \lambda \bar{\sigma}^{-2}]^\ell. \quad (\text{D35})$$

Taking the ℓ th root of both sides, we obtain

$$u^{\ell+1} = (1 - \lambda)u^\ell + \lambda \bar{\sigma}^{-2} \quad (\text{D36})$$

and hence, dividing by u^ℓ under the assumption that it is positive and rearranging terms, we obtain

$$1 - u = \lambda(1 - u^{-\ell} \bar{\sigma}^{-2}), \quad (\text{D37})$$

which recovers Li and Sompolinsky's result.

APPENDIX E: COMPARISON TO LARGE-WIDTH PERTURBATIVE CALCULATIONS WITH FIXED DATA

In this Appendix, we compare the results of the replica theory calculation of the generalization error of a deep linear network in the present paper to our previous perturbative results in [35]. Concretely, Appendix G of [35] computes the leading finite-width correction to the posterior-averaged error on a test set of \hat{p} examples for fixed data in the regime $\alpha < 1$. As noted there and in [36], this can be roughly interpreted as the leading-order correction in $\ell p/n$, as the correction that is parametrically $\mathcal{O}(1/n)$ in n scales with ℓ and p . To make contact with the present paper, we evaluate their result for a test dataset of d examples, with trivial data matrix $\hat{\mathbf{X}} = \sqrt{d}\mathbf{I}_d$. Then, the perturbative approximation for the thermal bias-variance decomposition given in [35] is

$$\varepsilon_b = \frac{1}{d} \|\sqrt{d}\mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} - \mathbf{w}_*\|^2, \quad (\text{E1})$$

$$\varepsilon_v = \frac{1}{d} \text{tr}[\mathbf{I}_d - \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X}] \left[\sigma^2 + \left(\sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l} \right) \times \left(\frac{d}{p} \mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} - \sigma^2 \right) \right] + \mathcal{O}(n^{-2}). \quad (\text{E2})$$

We remark that the results of [35,36] show that it should be safe to interchange the limit $\beta \rightarrow \infty$ with the high-dimensional limit and the expectation over data. Using the fact that $\mathbf{X}\mathbf{X}^\top$ is invertible with probability 1 in this regime, the disorder average of the thermal bias term yields

$$\lim_{n, p \rightarrow \infty} \mathbb{E}_{\mathbf{X}, \xi, \mathbf{w}_*} \varepsilon_b = 1 - \alpha + \frac{\alpha}{1 - \alpha} \eta^2, \quad (\text{E3})$$

where we have again used the formula for the expectation of an inverse Wishart matrix with identity scale matrix [68]. Similarly, the disorder average of the thermal variance term yields

$$\lim_{n, p \rightarrow \infty} \mathbb{E}_{\mathbf{X}, \xi, \mathbf{w}_*} \varepsilon_v = (1 - \alpha)\sigma^2 + [(1 - \alpha)(1 - \sigma^2) + \eta^2] \sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l} + \mathcal{O}(n^{-2}). \quad (\text{E4})$$

Therefore, the perturbative result of [35] implies a large-width disorder-averaged generalization error of

$$\epsilon_{\text{NN}} = \epsilon_{\text{LR}} + [(1 - \sigma^2)(1 - \alpha) + \eta^2] \sum_{l=1}^{\ell} \frac{\alpha}{\gamma_l} + \mathcal{O}(n^{-2}), \quad (\text{E5})$$

which agrees with the leading-order large-width solution of the RS result reported here. This makes sense, as we intuitively expect possible replica symmetry breaking effects to emerge at smaller width. Moreover, we remark that we have an exact correspondence between $Q - q$ and q and the averages of the thermal variance and bias terms, respectively. Finally, we note that the coefficient of the $\mathcal{O}(\alpha/\gamma)$ correction, which is the dataset average of $d\mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y}/p - \sigma^2$, gives the dataset average of the condition for when increasing width helps generalization noted by [35,37].

APPENDIX F: DETAILED ANALYSIS OF OPTIMAL NETWORK ARCHITECTURE

In this Appendix, we provide a detailed analysis of how width and depth affect generalization in RF and NN models.

1. Optimal width for RF models

We first consider optimizing the width of a deep random feature model. In the regime $\alpha < \min\{1, \gamma_{\min}\}$, we have

$$\frac{\partial \epsilon_{\text{RF}}}{\partial \gamma_l} = \frac{\alpha(1 - \alpha + \eta^2)}{\gamma_l^2} \left[\tilde{\sigma}^2 \prod_{l' \neq l} \frac{\gamma_{l'} - \alpha}{\gamma_{l'}} - \left(\frac{\gamma_l}{\gamma_l - \alpha} \right)^2 \right] \quad (\text{F1})$$

where we have defined

$$\tilde{\sigma}^2 \equiv \frac{1 - \alpha}{1 - \alpha + \eta^2} \sigma^2 \quad (\text{F2})$$

as in Eq. (23). If $\ell = 1$, this is simply

$$\frac{\partial \epsilon_{\text{RF}}}{\partial \gamma_1} = \frac{\alpha(1 - \alpha + \eta^2)}{\gamma_1^2} \left[\tilde{\sigma}^2 - \left(\frac{\gamma_1}{\gamma_1 - \alpha} \right)^2 \right]. \quad (\text{F3})$$

For $\tilde{\sigma} \leq 1$, $\partial \epsilon_{\text{RF}}/\partial \gamma_l < 0$ for all parameter values in this regime, and ϵ_{RF} is minimized by taking $\gamma_l \rightarrow \infty$. If $\tilde{\sigma} > 1$, there is a valid, i.e., $\gamma_l > \alpha$, stationary point $\partial \epsilon_{\text{RF}}/\partial \gamma_l = 0$ at

$$\gamma_\star = \frac{\tilde{\sigma}}{\tilde{\sigma} - 1} \alpha. \quad (\text{F4})$$

For $\gamma_l < \gamma_\star$, $\partial \epsilon_{\text{RF}}/\partial \gamma_l < 0$, while for $\gamma_l > \gamma_\star$, $\partial \epsilon_{\text{RF}}/\partial \gamma_l > 0$. This point is therefore a minimum of ϵ_{RF} .

We now consider $\ell > 1$. In this parameter regime, we have

$$\prod_{l' \neq l} \frac{\gamma_{l'} - \alpha}{\gamma_{l'}} \leq 1 \quad \text{and} \quad \left(\frac{\gamma_l}{\gamma_l - \alpha} \right)^2 \geq 1, \quad (\text{F5})$$

where both inequalities are strict if all $\gamma_l < \infty$. Thus, if $\tilde{\sigma} \leq 1$, $\partial \epsilon_{\text{RF}}/\partial \gamma_l$ is always negative for all l , and ϵ_{RF} is minimized by taking all $\gamma_l \rightarrow \infty$. If $\tilde{\sigma} > 1$, we have a nontrivial stationary point with all

$$\gamma_l = \gamma_\star \quad (l = 1, \dots, \ell) \quad (\text{F6})$$

for

$$\frac{\gamma_\star - \alpha}{\gamma_\star} = \tilde{\sigma}^{-2/(\ell+1)}, \quad (\text{F7})$$

which gives

$$\gamma_\star = \frac{\tilde{\sigma}^{2/(\ell+1)}}{\tilde{\sigma}^{2/(\ell+1)} - 1} \alpha. \quad (\text{F8})$$

To check whether this is indeed a local minimum, we compute the Hessian of ϵ_{RF} at the stationary point, which is given by

$$\frac{\partial^2 \epsilon_{\text{RF}}}{\partial \gamma_l \partial \gamma_{l'}} \Big|_{\gamma_l = \gamma_\star} = \frac{\alpha^2(1 - \alpha + \eta^2)}{\gamma_\star^4} \tilde{\sigma}^{6/(\ell+1)} [\delta_{ll'} + 1]. \quad (\text{F9})$$

Diagonalizing this matrix is trivial, yielding eigenvalue

$$\frac{\alpha^2(1 - \alpha + \eta^2)}{\gamma_\star^4} \tilde{\sigma}^{6/(\ell+1)} \quad (\text{F10})$$

with multiplicity $\ell - 1$ and eigenvalue

$$\frac{\alpha^2(1 - \alpha + \eta^2)}{\gamma_\star^4} \tilde{\sigma}^{6/(\ell+1)} (\ell + 1) \quad (\text{F11})$$

with multiplicity 1. Both of these eigenvalues are positive throughout the parameter region of interest, confirming that the Hessian is positive definite at the stationary point. Moreover, substituting γ_\star into the generalization error ϵ_{RF} , we have

$$\epsilon_{\text{RF}} \Big|_{\gamma_1 = \dots = \gamma_\ell = \gamma_\star} = \epsilon_{\text{LR}} + (1 - \alpha) \sigma^2 (\tilde{\sigma}^{-2\ell/(\ell+1)} - 1) + \ell(1 - \alpha + \eta^2) (\tilde{\sigma}^{2/(\ell+1)} - 1) \quad (\text{F12})$$

for any fixed $\alpha < \min\{1, \gamma_\star\}$. As $\tilde{\sigma} > 1$, both correction terms are negative, and this result is a monotonically decreasing function of the network depth ℓ . It can therefore be seen that this result yields better generalization than that obtained by taking any subset of the hidden layers to infinity while keeping the remainder fixed at γ_\star . This result has the interesting feature that, for any fixed $\tilde{\sigma} > 1$ and α, γ_\star is a monotonically increasing function of ℓ .

For RF models in the regime $\alpha > \gamma_{\min}$ for $\gamma_{\min} < 1$, it is easy to see that ϵ_{RF} is a monotonically decreasing function of $\gamma_{\min} \in [0, 1)$ if $\alpha > 1$, while if $\alpha < 1$ it is a monotonically increasing function of $\gamma_{\min} \in [0, \alpha)$. However, in this regime, it is important to keep track of crossings in the ordering of different layer widths.

2. Optimal depth for RF models

We now consider optimizing the depth of a RF model. For simplicity, we specialize to the case in which all hidden layers have the same width, i.e., $\gamma_1 = \dots = \gamma_\ell = \gamma$. Our starting point is therefore the expression Eq. (22), which can sensibly be analytically continued in width. For brevity, we let

$$\psi \equiv \frac{\gamma - \alpha}{\gamma}, \quad (\text{F13})$$

which is bounded as $0 < \psi < 1$ in the regime of interest. This gives

$$\frac{\epsilon_{\text{RF}} - \epsilon_{\text{LR}}}{1 - \alpha + \eta^2} = \tilde{\sigma}^2 (\psi^\ell - 1) + \ell \left(\frac{1}{\psi} - 1 \right), \quad (\text{F14})$$

where we have defined $\tilde{\sigma}$ as in Eq. (23). Treating ℓ as a continuous parameter, we then have

$$\frac{\partial \epsilon_{\text{RF}}}{\partial \ell} = (1 - \alpha + \eta^2) \left[\tilde{\sigma}^2 \psi^\ell \log(\psi) + \frac{1}{\psi} - 1 \right]. \quad (\text{F15})$$

For all $\tilde{\sigma} > 0$ and $0 < \psi < 1$, we can see that $\partial \epsilon_{\text{RF}} / \partial \ell$ is a monotonically increasing function of ℓ , as can be confirmed by inspecting

$$\frac{\partial^2 \epsilon_{\text{RF}}}{\partial \ell^2} = (1 - \alpha + \eta^2) \tilde{\sigma}^2 \psi^\ell \log(\psi)^2 > 0. \quad (\text{F16})$$

Using the lower bound $\log(\psi) > 1 - 1/\psi$, which is strict for all $0 < \psi < 1$, we have

$$\frac{\partial \epsilon_{\text{RF}}}{\partial \ell} > (1 - \alpha + \eta^2) \left(1 - \frac{1}{\psi} \right) (1 - \tilde{\sigma}^2 \psi^\ell). \quad (\text{F17})$$

For $\tilde{\sigma} \leq 1$, we have $\tilde{\sigma}^2 \psi^\ell < 1$ for all $0 < \psi < 1$ and all $\ell \geq 1$, hence the above bound shows that

$$\frac{\partial \epsilon_{\text{RF}}}{\partial \ell} > 0, \quad (\text{F18})$$

implying that ϵ_{RF} is a monotonically increasing function of ℓ if $\tilde{\sigma} \leq 1$. Therefore, shallow random feature models are optimal in this regime. This is consistent with our earlier observations regarding optimal width, as taking hidden layer widths to infinity reduces the effective depth.

If $\tilde{\sigma} > 1$, then the above lower bound shows that $\partial \epsilon_{\text{RF}} / \partial \ell > 0$ if $\psi \leq \tilde{\sigma}^{-2/\ell}$, i.e., if $\ell \geq -\log(\tilde{\sigma}^2) / \log(\psi)$. However, in this case, ϵ_{RF} is not always an increasing function of ℓ . In particular, using the upper bound $\log(\psi) < \psi(1 - 1/\psi)$, which is again strict for $0 < \psi < 1$, we have the upper bound

$$\frac{\partial \epsilon_{\text{RF}}}{\partial \ell} < (1 - \alpha + \eta^2) \left(\frac{1}{\psi} - 1 \right) (1 - \tilde{\sigma}^2 \psi^{\ell+1}), \quad (\text{F19})$$

which shows that $\partial \epsilon_{\text{RF}} / \partial \ell < 0$ for all $\psi \geq \tilde{\sigma}^{-2/(\ell+1)}$, i.e., if $\ell \leq -\log(\tilde{\sigma}^2) / \log(\psi) - 1$. We remark that the optimized value γ_* found above in our study of optimal width is covered by this crude bound, as it is defined by the condition $\psi|_{\gamma=\gamma_*} = \tilde{\sigma}^{-2/(\ell+1)}$.

In terms of ℓ , the intermediate region not covered by the bounds above is the open interval

$$-\frac{\log(\tilde{\sigma}^2)}{\log(\psi)} - 1 < \ell < -\frac{\log(\tilde{\sigma}^2)}{\log(\psi)}. \quad (\text{F20})$$

To the left of this interval, we know that $\partial \epsilon_{\text{RF}} / \partial \ell < 0$, while to the right of this interval we know that $\partial \epsilon_{\text{RF}} / \partial \ell > 0$. Therefore, for any ℓ outside this open interval, ϵ_{RF} is strictly greater than its values anywhere within the interval.

If $-\log(\tilde{\sigma}^2) / \log(\psi)$ is not an integer, this open interval will always include exactly one integer value of ℓ ,

$$\ell_* = \left\lfloor -\frac{\log(\tilde{\sigma}^2)}{\log(\psi)} \right\rfloor, \quad (\text{F21})$$

which will thus minimize the generalization error for fixed $\tilde{\sigma} > 1$ and $0 < \psi < 1$. Restoring the definition of ψ in terms of γ and α , this optimal value is

$$\ell_* = \left\lfloor \frac{\log(\tilde{\sigma}^2)}{\log[\gamma/(\gamma - \alpha)]} \right\rfloor. \quad (\text{F22})$$

If $-\log(\tilde{\sigma}^2) / \log(\psi)$ is an integer, then no integers are contained within the open interval that is not covered by the bounds on $\partial \epsilon_{\text{RF}} / \partial \ell$ derived above. In that case, the bounds on $\partial \epsilon_{\text{RF}} / \partial \ell$ imply that the optimal depth is then given by one of the boundary points, i.e.,

$$\ell_* \in \{j, j - 1\}, \quad (\text{F23})$$

where we have defined the positive integer [noting that $-\log(\tilde{\sigma}^2) / \log(\psi) > 0$ for all $\tilde{\sigma} > 1$ and $0 < \psi < 1$]

$$j \equiv -\frac{\log(\tilde{\sigma}^2)}{\log(\psi)} \in \mathbb{N}_{>0}. \quad (\text{F24})$$

For candidate depths of this form, we can use the fact that $\psi^j = \psi^{-\log(\tilde{\sigma}^2) / \log(\psi)} = \tilde{\sigma}^{-2}$ to obtain

$$\frac{\epsilon_{\text{RF}}|_{\ell=j+k} - \epsilon_{\text{RF}}|_{\ell=j}}{1 - \alpha + \eta^2} = k \left(\frac{1}{\psi} - 1 \right) + \psi^k - 1 \quad (\text{F25})$$

for any $k \in \mathbb{Z}$ such that $j + k \geq 0$. For any $0 < \psi < 1$, this generalization gap is positive for all $k < -1$, vanishes when $k = -1$, is negative for $-1 < k < 0$, vanishes for $k = 0$, and is positive for all $k > 0$. To show this, we first observe that the gap is a smooth function of ψ for any k (including noninteger real values), with

$$\left[k \left(\frac{1}{\psi} - 1 \right) + \psi^k - 1 \right] \Big|_{\psi=1} = 0. \quad (\text{F26})$$

Then, we consider

$$\frac{\partial}{\partial \psi} \left[k \left(\frac{1}{\psi} - 1 \right) + \psi^k - 1 \right] = k \psi^{-2} (\psi^{k+1} - 1), \quad (\text{F27})$$

which is strictly negative for all $0 < \psi < 1$ and vanishes from below as $\psi \uparrow 1$ if $k < -1$ or if $k > 0$, and is strictly positive for all $0 < \psi < 1$ and vanishes from above as $\psi \uparrow 1$ if $-1 < k < 0$. This shows that the desired claim should hold. Therefore, in this case the two candidate depths will yield identical generalization error, and we can take either $\ell_* = j$ or $j - 1$.

We note that the condition $-\log(\tilde{\sigma}^2) / \log(\psi) = j \in \mathbb{N}_{>0}$ implies that

$$\psi = \tilde{\sigma}^{-2/j}, \quad (\text{F28})$$

which gives a condition on the width as a function of $\tilde{\sigma}^2$ and α :

$$\gamma = \frac{\tilde{\sigma}^{2/j}}{\tilde{\sigma}^{2/j} - 1} \alpha. \quad (\text{F29})$$

In our earlier study of the optimal width for fixed depth, we found solutions of precisely this form, with $j = \ell + 1$. This result for the optimal depth at fixed width is therefore internally consistent with our earlier result for the optimal width at fixed depth.

3. Optimal width for NN models

For a two-layer NN in the regime $\alpha < 1$, we have

$$\frac{\partial \epsilon_{\text{NN}}}{\partial \gamma_1} = \frac{\alpha(1 - \alpha)\sigma^2}{2\gamma_1^2} \left[1 + \frac{(\gamma_1 - \alpha) - 2\gamma_1\tilde{\sigma}^{-2}}{\sqrt{(\gamma_1 - \alpha)^2 + 4\alpha\gamma_1\tilde{\sigma}^{-2}}} \right] \quad (\text{F30})$$

where we have defined $\tilde{\sigma}$ as in Eq. (23). If $\tilde{\sigma} = 1$, we have

$$\left. \frac{(\gamma_1 - \alpha) - 2\gamma_1\tilde{\sigma}^{-2}}{\sqrt{(\gamma_1 - \alpha)^2 + 4\alpha\gamma_1\tilde{\sigma}^{-2}}} \right|_{\tilde{\sigma}=1} = -\frac{\gamma_1 + \alpha}{\sqrt{(\gamma_1 + \alpha)^2}} = -1, \quad (\text{F31})$$

hence $\partial\epsilon_{\text{NN}}/\partial\gamma_1|_{\tilde{\sigma}=1} = 0$. Exactly at $\alpha = \gamma_1$, we have

$$\left. \frac{(\gamma_1 - \alpha) - 2\gamma_1\tilde{\sigma}^{-2}}{\sqrt{(\gamma_1 - \alpha)^2 + 4\alpha\gamma_1\tilde{\sigma}^{-2}}} \right|_{\alpha=\gamma_1} = -\frac{1}{\tilde{\sigma}}. \quad (\text{F32})$$

If $\alpha \neq \gamma_1$, we have

$$\begin{aligned} \frac{(\gamma_1 - \alpha) - 2\gamma_1\tilde{\sigma}^{-2}}{\sqrt{(\gamma_1 - \alpha)^2 + 4\alpha\gamma_1\tilde{\sigma}^{-2}}} &> -\frac{\gamma_1 + \alpha}{\sqrt{(\gamma_1 - \alpha)^2 + 4\alpha\gamma_1\tilde{\sigma}^{-2}}} \\ &> -\frac{\gamma_1 + \alpha}{\sqrt{(\gamma_1 + \alpha)^2}} = -1 \end{aligned} \quad (\text{F33})$$

if $\tilde{\sigma} > 1$; the reversed chain of strict inequalities holds if $\tilde{\sigma} < 1$. Therefore, we conclude that, for any α and γ_1 ,

$$\frac{\partial\epsilon_{\text{NN}}}{\partial\gamma_1} = \begin{cases} < 0, & 0 < \tilde{\sigma} < 1 \\ = 0, & \tilde{\sigma} = 1, \\ > 0, & \tilde{\sigma} > 1, \end{cases} \quad (\text{F34})$$

as reported in the main text.

4. Difference in generalization in two-layer NN and RF models

For two-layer networks, we have the RF-NN generalization gap

$$\frac{\epsilon_{\text{RF}} - \epsilon_{\text{NN}}}{1 - \alpha + \eta^2} = 1 - \frac{1}{\psi} + \tilde{\sigma}^2 \frac{\psi - \sqrt{\psi^2 + 4\alpha/\gamma_1\tilde{\sigma}^2}}{2} \quad (\text{F35})$$

in the regime $\alpha < \min\{1, \gamma_1\}$, where $\tilde{\sigma}$ is defined in Eq. (23) and $\psi \equiv (\gamma_1 - \alpha)/\gamma_1$. For any finite γ_1 , we have $0 < \psi <$

1, hence it is easy to see that both $1 - 1/\psi$ and $\psi - \sqrt{\psi^2 + 4\alpha/\gamma_1\tilde{\sigma}^2}$ are negative. If $\gamma_1 \rightarrow \infty$, then $\psi \rightarrow 1$, and $\epsilon_{\text{RF}} - \epsilon_{\text{NN}} \downarrow 0$. This yields the conclusion reported in the main text.

APPENDIX G: NUMERICAL METHODS

Below, we elaborate on the numerical methods used to validate the replica-symmetric learning curves. In all of the numerical simulations, we set the input dimensionality $d = 100$, and sample with a resolution of around 50–100 estimates per dimension. To produce the standard error bars, we sample ten values per estimate. Numerical procedures were written with NUMPY [71] and SCIPY [72]. Plots were generated with MATPLOTLIB [73].

1. Random feature model

Theoretical predictions for the generalization error in Bayesian random feature models can be computed directly from Eq. (20). Additionally, after sampling a set of initial weights, we use the results from Appendix C to directly compute the posterior expectations for the RF model. By then averaging the resulting error across multiple samples of weights, we numerically verify our theoretical curves.

2. Neural network model

Theoretical predictions for the generalization error in a Bayesian neural network can be computed directly from Eq. (28). We then use the results from Appendix D to directly compute the error for particular instantiations of weights, numerically verifying our theoretical results.

-
- [1] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, Understanding deep learning (still) requires rethinking generalization, *Commun. ACM* **64**, 107 (2021).
- [2] 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).
- [3] G. Yang and E. J. Hu, Tensor Programs IV: Feature learning in infinite-width neural networks, in *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by M. Meila and T. Zhang (PMLR, 2021), Vol. 139, pp. 11727–11737.
- [4] L. Aitchison, Why bigger is not always better: On finite and infinite neural networks, in *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by H. D. III and A. Singh (PMLR, 2020), Vol. 119, pp. 156–164.
- [5] P. Nakkiran, G. Kaplun, Y. Bansal, T. Yang, B. Barak, and I. Sutskever, Deep double descent: Where bigger models and more data hurt, *J. Stat. Mech.* (2021) 124003.
- [6] G. Yang, Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation, [arXiv:1902.04760](https://arxiv.org/abs/1902.04760).
- [7] M. Refinetti, S. Goldt, F. Krzakala, and L. Zdeborova, Classifying high-dimensional Gaussian mixtures: Where kernel methods fail and neural networks succeed, in *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by M. Meila and T. Zhang (PMLR, 2021), Vol. 139, pp. 8936–8947.
- [8] B. Woodworth, S. Gunasekar, J. D. Lee, E. Moroshko, P. Savarese, I. Golan, D. Soudry, and N. Srebro, Kernel and rich regimes in overparametrized models, in *Conference on Learning Theory* (PMLR, 2020), pp. 3635–3673.
- [9] M. Geiger, S. Spigler, A. Jacot, and M. Wyart, Disentangling feature and lazy training in deep neural networks, *J. Stat. Mech.* (2020) 113301.
- [10] M. Geiger, A. Jacot, S. Spigler, F. Gabriel, L. Sagun, S. d’Ascoli, G. Biroli, C. Hongler, and M. Wyart, Scaling description of generalization with number of parameters in deep learning, *J. Stat. Mech.* (2020) 023401.
- [11] R. M. Neal, Priors for infinite networks, in *Bayesian Learning for Neural Networks* (Springer, New York, 1996), pp. 29–53.
- [12] C. K. Williams, Computing with infinite networks, *Advances in Neural Information Processing Systems* **9**, 295 (1997).
- [13] J. Lee, J. Sohl-Dickstein, J. Pennington, R. Novak, S. Schoenholz, and Y. Bahri, Deep neural networks as

- Gaussian processes, in Proceedings of the International Conference on Learning Representations, 2018 (unpublished), [arXiv:1711.00165](#).
- [14] A. G. D. G. Matthews, J. Hron, M. Rowland, R. E. Turner, and Z. Ghahramani, Gaussian process behaviour in wide deep neural networks, in Proceedings of the International Conference on Learning Representations, 2018 (unpublished), [arXiv:1804.11271](#).
- [15] A. Jacot, F. Gabriel, and C. Hongler, Neural tangent kernel: Convergence and generalization in neural networks, in *Advances in Neural Information Processing Systems* (2018), Vol. 31, [arXiv:1806.07572](#).
- [16] J. Hron, Y. Bahri, R. Novak, J. Pennington, and J. Sohl-Dickstein, Exact posterior distributions of wide Bayesian neural networks, [arXiv:2006.10541](#) (2020).
- [17] S. Mei and A. Montanari, The generalization error of random features regression: Precise asymptotics and the double descent curve, *Commun. Pure Appl. Math.* **75**, 667 (2019).
- [18] H. Hu and Y. M. Lu, Universality laws for high-dimensional learning with random features, [arXiv:2009.07669](#) (2020).
- [19] S. Spigler, M. Geiger, and M. Wyart, Asymptotic learning curves of kernel methods: Empirical data versus teacher-student paradigm, *J. Stat. Mech.* (2020) 124001.
- [20] A. Canatar, B. Bordelon, and C. Pehlevan, Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks, *Nat. Commun.* **12**, 2914 (2021).
- [21] J. Barbier, W.-K. Chen, D. Panchenko, and M. Sáenz, Performance of Bayesian linear regression in a model with mismatch, [arXiv:2107.06936](#) (2021).
- [22] S. d'Ascoli, L. Sagun, and G. Biroli, Triple descent and the two kinds of overfitting: where and why do they appear?, *J. Stat. Mech.* (2021) 124002.
- [23] B. Adlam and J. Pennington, The neural tangent kernel in high dimensions: Triple descent and a multi-scale theory of generalization, in *International Conference on Machine Learning* (PMLR, 2020), pp. 74–84.
- [24] B. Adlam and J. Pennington, Understanding double descent requires a fine-grained bias-variance decomposition, in *Advances in Neural Information Processing Systems* (2020), Vol. 33, pp. 11022–11032.
- [25] S. d'Ascoli, M. Refinetti, G. Biroli, and F. Krzakala, Double trouble in double descent: Bias and variance(s) in the lazy regime, in *International Conference on Machine Learning* (PMLR, 2020), pp. 2280–2290.
- [26] H. Jin, P. K. Banerjee, and G. Montufar, Learning curves for Gaussian process regression with power-law priors and targets, in Proceedings of the International Conference on Learning Representations, 2022 (unpublished), [arXiv:2110.12231](#).
- [27] B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mezard, and L. Zdeborova, Learning curves of generic features maps for realistic datasets with a teacher-student model, in *Advances in Neural Information Processing Systems*, edited by A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan (2021), Vol. 34.
- [28] J. Lee, S. Schoenholz, J. Pennington, B. Adlam, L. Xiao, R. Novak, and J. Sohl-Dickstein, Finite versus infinite neural networks: an empirical study, in *Advances in Neural Information Processing Systems* (2020), Vol. 33, pp. 15156–15172.
- [29] J. M. Antognini, Finite size corrections for neural network Gaussian processes, [arXiv:1908.10030](#) (2019).
- [30] S. Yaida, Non-Gaussian processes and neural networks at finite widths, in *Proceedings of The First Mathematical and Scientific Machine Learning Conference*, Proceedings of Machine Learning Research, edited by J. Lu and R. Ward (PMLR, 2020), Vol. 107, pp. 165–192.
- [31] J. A. Zavatone-Veth and C. Pehlevan, Exact marginal prior distributions of finite Bayesian neural networks, in *Advances in Neural Information Processing Systems*, edited by A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan (2021), Vol. 34.
- [32] D. A. Roberts, S. Yaida, and B. Hanin, *The Principles of Deep Learning Theory: An Effective Theory Approach to Understanding Neural Networks* (Cambridge University, New York, 2022).
- [33] K. T. Grosvenor and R. Jefferson, The edge of chaos: quantum field theory and deep neural networks, [arXiv:2109.13247](#) (2021).
- [34] J. Halverson, A. Maiti, and K. Stoner, Neural networks and quantum field theory, *Mach. Learn.: Sci. Technol.* **2**, 035002 (2021).
- [35] J. A. Zavatone-Veth, A. Canatar, B. S. Ruben, and C. Pehlevan, Asymptotics of representation learning in finite Bayesian neural networks, in *Advances in Neural Information Processing Systems*, edited by A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan (2021), Vol. 34.
- [36] J. A. Zavatone-Veth and C. Pehlevan, Depth induces scale-averaging in overparameterized linear Bayesian neural networks, in *Asilomar Conference on Signals, Systems, and Computers* (2021), Vol. 55.
- [37] Q. Li and H. Sompolinsky, Statistical Mechanics of Deep Linear Neural Networks: The Backpropagating Kernel Renormalization, *Phys. Rev. X* **11**, 031059 (2021).
- [38] G. Naveh, O. Ben David, H. Sompolinsky, and Z. Ringel, Predicting the outputs of finite deep neural networks trained with noisy gradients, *Phys. Rev. E* **104**, 064301 (2021).
- [39] G. Naveh and Z. Ringel, A self consistent theory of Gaussian processes captures feature learning effects in finite CNNs, in *Advances in Neural Information Processing Systems*, edited by A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan (2021), Vol. 34.
- [40] E. Dyer and G. Gur-Ari, Asymptotics of wide networks from Feynman diagrams, in Proceedings of the International Conference on Learning Representations, 2020 (unpublished), [arXiv:1909.11304](#).
- [41] K. Aitken and G. Gur-Ari, On the asymptotics of wide networks with polynomial activations, [arXiv:2006.06687](#) (2020).
- [42] 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](#) (2013).
- [43] K. Fukumizu, Effect of batch learning in multilayer neural networks, in Proceedings of the Fifth International Conference on Neural Information Processing, 1998, pp. 67–70 (unpublished).
- [44] P. Nakkiran, More data can hurt for linear regression: Sample-wise double descent, [arXiv:1912.07242](#) (2019).
- [45] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, Surprises in high-dimensional ridgeless least squares interpolation, [arXiv:1903.08560](#) (2019).

- [46] M. S. Advani, A. M. Saxe, and H. Sompolinsky, High-dimensional dynamics of generalization error in neural networks, *Neural Networks* **132**, 428 (2020).
- [47] A. G. Wilson and P. Izmailov, Bayesian deep learning and a probabilistic perspective of generalization, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Inc., Red Hook, NY, 2020), Vol. 33, pp. 4697–4708.
- [48] F. Wenzel, K. Roth, B. Veeling, J. Swiatkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin, How good is the Bayes posterior in deep neural networks really?, in *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by H. D. III and A. Singh (PMLR, 2020), Vol. 119, pp. 10248–10259.
- [49] P. Izmailov, S. Vikram, M. D. Hoffman, and A. G. G. Wilson, What are Bayesian neural network posteriors really like?, in *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by M. Meila and T. Zhang (PMLR, 2021), Vol. 139, pp. 4629–4640.
- [50] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* (World Scientific, Singapore, 1987).
- [51] A. Engel and C. van den Broeck, *Statistical Mechanics of Learning* (Cambridge University, New York, 2001).
- [52] A. Krogh and J. A. Hertz, Generalization in a linear perceptron in the presence of noise, *J. Phys. A: Math. Gen.* **25**, 1135 (1992).
- [53] M. Advani and S. Ganguli, Statistical Mechanics of Optimal Convex Inference in High Dimensions, *Phys. Rev. X* **6**, 031034 (2016).
- [54] M. Biehl, E. Schlösser, and M. Ahr, Phase transitions in soft-committee machines, *Europhys. Lett.* **44**, 261 (1998).
- [55] S. A. Solla and E. Levin, Learning in linear neural networks: The validity of the annealed approximation, *Phys. Rev. A* **46**, 2124 (1992).
- [56] E. Levin, N. Tishby, and S. A. Solla, A statistical approach to learning and generalization in layered neural networks, *Proc. IEEE* **78**, 1568 (1990).
- [57] J. Barbier, D. Panchenko, and M. Sáenz, Strong replica symmetry for high-dimensional disordered log-concave Gibbs measures, *Information and Inference: A Journal of the IMA*, [arXiv:2009.12939](https://arxiv.org/abs/2009.12939).
- [58] G. Akemann, J. R. Ipsen, and M. Kieburg, Products of rectangular random matrices: singular values and progressive scattering, *Phys. Rev. E* **88**, 052118 (2013).
- [59] J. R. Ipsen and M. Kieburg, Weak commutation relations and eigenvalue statistics for products of rectangular random matrices, *Phys. Rev. E* **89**, 032106 (2014).
- [60] J. W. Rocks and P. Mehta, Bias-variance decomposition of overparameterized regression with random linear features, [arXiv:2203.05443](https://arxiv.org/abs/2203.05443) (2022).
- [61] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, ImageNet large scale visual recognition challenge, *Int. J. Comput. Vision* **115**, 211 (2015).
- [62] J. A. Zavatone-Veth and C. Pehlevan, Activation function dependence of the storage capacity of treelike neural networks, *Phys. Rev. E* **103**, L020301 (2021).
- [63] 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).
- [64] M. Cranmer, D. Tamayo, H. Rein, P. Battaglia, S. Hadden, P. J. Armitage, S. Ho, and D. N. Spergel, A Bayesian neural network predicts the dissolution of compact planetary systems, *Proc. Natl. Acad. Sci. USA* **118**, e2026053118 (2021).
- [65] C. Mingard, G. Valle-Pérez, J. Skalse, and A. A. Louis, Is SGD a Bayesian sampler?, *J. Mach. Learn. Research* **22**, 1 (2021).
- [66] B. He, B. Lakshminarayanan, and Y. W. Teh, Bayesian deep ensembles via the neural tangent kernel, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin (Curran Associates, Inc., Red Hook, NY, 2020), Vol. 33, pp. 1010–1022.
- [67] R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge University, New York, 2012).
- [68] R. J. Muirhead, *Aspects of Multivariate Statistical Theory* (Wiley, New York, 2009).
- [69] J. R. Magnus and H. Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics* (Wiley, New York, 2019).
- [70] DLMF, *NIST Digital Library of Mathematical Functions*, edited by F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, <http://dlmf.nist.gov/>, Release 1.1.1 of 2021-03-15 (2021).
- [71] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. Fernández del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [72] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey *et al.*, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nat. Methods* **17**, 261 (2020).
- [73] J. D. Hunter, Matplotlib: A 2d graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).