MULTI-SCALE FEATURE LEARNING DYNAMICS: INSIGHTS FOR DOUBLE DESCENT

Anonymous authors
Paper under double-blind review

ABSTRACT

A key challenge in building theoretical foundations for deep learning is the complex optimization dynamics of neural networks, resulting from the high-dimensional interactions between the large number of network parameters. Such non-trivial interactions lead to intriguing model behaviors such as the phenomenon of "double descent" of the generalization error. The more commonly studied aspect of this phenomenon corresponds to *model-wise* double descent where the test error exhibits a second descent with increasing model complexity, beyond the classical U-shaped error curve. In this work, we investigate the origins of the less studied *epoch-wise* double descent in which the test error undergoes two non-monotonous transitions, or descents as the training time increases. By leveraging tools from random matrix theory and statistical physics, we study a linear teacher-student setup exhibiting epoch-wise double descent similar to that in deep neural networks. In this setting, we derive closed-form analytical expressions for the evolution of generalization error over training. We find that double descent can be attributed to distinct features being learned at different scales: as fast-learning features overfit, slower-learning features start to fit, resulting in a second descent in test error. We validate our findings through numerical experiments where our theory accurately predicts empirical findings and remains consistent with observations in deep neural networks.

1 INTRODUCTION

Classical wisdom in statistical learning theory predicts a trade-off between the generalization ability of a machine learning model and its complexity, with highly complex models less likely to generalize well (Friedman et al., 2001). If the number of parameters measures complexity, deep learning models sometimes go against this prediction (Zhang et al., 2016): deep neural networks trained by stochastic gradient descent exhibit a so-called *double descent* behavior (Belkin et al., 2019b) with increasing model parameters. Specifically, with increasing complexity, the generalization error first obeys the classical U-shaped curve consistent with statistical learning theory. However, a second regime emerges as the number of parameters is further increased past a transition threshold where generalization error drops again, hence the "double descent" or more accurately *model-wise double descent* (Nakkiran et al., 2019).

Nakkiran et al. (2019) showed that the phenomenon of double descent is not limited to varying model size but is also observed as a function of training time or epochs. In this case as well, the so-called *epoch-wise double descent* is in apparent contradiction with the classical understanding of over-fitting (Vapnik, 1998), where one expects that longer training of a sufficiently large model beyond a certain threshold should result in over-fitting. This has important implications for practitioners and raises questions about one of the most widely used regularization method in deep learning (Goodfellow et al., 2016): early stopping. Indeed, while one might expect early stopping to prevent over-fitting, it might in fact prevent models from being trained at their fullest potential.

Since the 1990s, there has been much interest in understanding the origins of non-trivial generalization behaviors of neural networks (Opper, 1995; Opper & Kinzel, 1996). The authors of Krogh & Hertz (1992b) were among the first to provide theoretical explanations for (model-wise) double descent in linear models. Summarily, at intermediate levels of complexity, where the model size is equal to the number of training examples, the model is very sensitive to noise in training data and

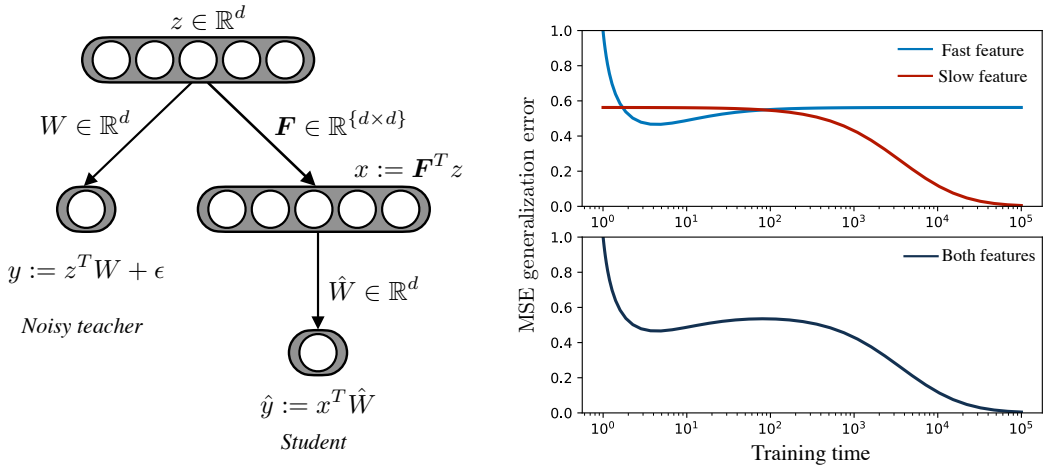

Figure 1: **Left**: The teacher-student setup for which we derive analytical results. The teacher is the data generating process that operates on isotropic Gaussian inputs $\boldsymbol{z}$. The student is trained on data generated by the teacher with its input being modulated by a matrix $\boldsymbol{F}$. As we study shortly, κ, the condition number of $\boldsymbol{F}$ determines how much faster a fraction of features are learned than the others. **Right**: The generalization error as the training time proceeds. (top): The case where only the fast-learning feature or slow-learning feature are trained. (bottom): The case where both features are trained with $\kappa = 1e2$.

hence, generalizes poorly. This sensitivity to noise reduces if the model complexity is either decreased or increased. More recently, the double descent phenomena has been also studied for more complex models such as two-layer neural networks and random feature models (Ba et al., 2019; Mei & Montanari, 2019; D'Ascoli et al., 2020).

The majority of previous work in this direction focuses on understanding the *asymptotic* behavior of model performance, i.e., where training time $t \to \infty$. In recent years, there has been an interest in studying the *non-asymptotic* (finite training time) performance, suggesting that several intriguing properties of neural networks can be attributed to different features being learned at different scales. Among limited work studying the particular epoch-wise double descent, Nakkiran et al. (2019) introduces the notion of *effective model complexity* and hypothesizes that it increases with training time and hence unifies both model-wise and epoch-wise double descent. In an empirical study Heckel & Yilmaz (2020), it is found that the time evolution dynamics show two or more overlapping, U-shape bias/variance curves with different minima times, thus leading to non-monotonic test error curves.

In this work, we build on (Bös et al., 1993; Bös, 1998; Advani & Saxe, 2017; Mei & Montanari, 2019) which analyze model-wise double descent through the lens of linear models, and probe the origins of epoch-wise double descent. Our main contributions are:

- In Section 2.2, we introduce a linear teacher-student model which, despite its simplicity, exhibits some of intriguing properties of generalization dynamics in deep neural networks.
- In the limit of high dimensions, we leverage the replica method developed in statistical physics to derive closed-form expressions for the generalization dynamics of our teacher-student setup, as a function of training time and regularization strength (Section 2.3).
- Consistent with recent findings, we provide an explanation for the existence of epoch-wise double descent through the lens of multi-scale feature learning. (Figure 1)
- We perform simulation experiments to validate our analytical predictions. We also conduct experiments with deep networks on CIFAR-10, showing that our teacher-student setup exhibits generalization behavior which is qualitatively similar to that of large nonlinear networks. (Figure 2)

2 ANALYTICAL RESULTS

2.1 PRELUDE

Before diving into the theory, we invite the reader to recall a simple equation from thermodynamics. Consider an ideal gas in a container with its large number of molecules moving around, colliding with each other, all while obeying Newton's laws. While the exact dynamics of each of such molecules is intractable, the system's macroscopic behavior can be characterized in terms of a handful of scalar quantities, namely, the pressure P, the volume V, and the temperature T. By averaging over suitable probability measures and applying the principle of free-energy minimization, one arrives at a remarkably simple relationship between these three macroscopic variables, i.e., the well-known $PV = nRT$ (n: number of moles of gas, R: gas constant) (Reif, 2009).

The same principle can be ported to neural networks. Stochastic Gradient Descent (SGD) — the de facto optimization algorithm for neural networks — exhibits complex dynamics arising from the interactions between a large number of parameters (Kunin et al., 2020). However, it is possible to describe some aspects of the high-dimensional *microscopic* dynamics of neural networks in terms of low-dimensional understandable *macroscopic* entities. In a series of seminal papers by Gardner (Gardner, 1988; Gardner & Derrida, 1988; 1989), the *replica method* of statistical physics was adopted to derive expressions describing the generalization behavior of large linear models trained using SGD. In this paper, we employ Gardner's analysis to build upon an established line of work studying linear and generalized linear models (Seung et al., 1992; Kabashima et al., 2009; Krzakala et al., 2012). While most of previous work study the asymptotic generalization behavior in the limit of large training time, we adapt these methods to study transient learning dynamics of generalization. We study a simplified teacher-student model that exhibits key characteristics of modern neural networks.

2.2 A TEACHER-STUDENT SETUP

Teacher: We study a supervised linear regression problem in which the training labels y, are generated by a noisy linear model (Figure 1),

$$y := y^* + \epsilon, \qquad y^* := \boldsymbol{z}^T W, \qquad z_i \sim \mathcal{N}(0, \frac{1}{\sqrt{d}}), \tag{1}$$

where $z \in \mathbb{R}^d$ is the teacher's input and $y^*, y \in \mathbb{R}$ are the teacher's noiseless and noisy outputs, respectively. $W \in \mathbb{R}^d$ represents the (fixed) weights of the teacher and $\epsilon \in \mathbb{R}$ is the noise. Both W and ϵ are drawn i.i.d. from Gaussian distributions with zero mean and variances of 1 and σ_ϵ^2, respectively.

Student: A student model is correspondingly chosen to be a similar shallow network with trainable weights $\hat{W} \in \mathbb{R}^d$. The student model is trained on n training pairs $\{(\boldsymbol{x}^\mu, y^\mu)\}_{\mu=1}^n$, with the labels y^μ being generated by the above teacher network, as,

$$\hat{y} := \boldsymbol{x}^T \hat{W}, \qquad s.t. \qquad \boldsymbol{x} := \boldsymbol{F}^T \boldsymbol{z}, \tag{2}$$

where the matrix $\boldsymbol{F} \in \mathbb{R}^{d \times d}$ is a predefined and fixed **modulation matrix** regulating the student's access to the true input $\boldsymbol{z}$. One can think of $\boldsymbol{z}$ as the latent factors of variation on which the teacher operates, while $\boldsymbol{x}$ can be thought as the pixels that the student learns from.

Learning paradigm: To train our student network, we use stochastic gradient descent (SGD) on the regularized mean squared loss, evaluated on the n training examples as,

$$\mathcal{L}_{\mathcal{T}} := \frac{1}{2n} \sum_{\mu=1}^{n} (y^\mu - \hat{y}^\mu)^2 + \frac{\lambda}{2} ||\hat{W}||_2^2 \tag{3}$$

where $\lambda \in [0, \infty)$ is the regularization coefficient. Optimizing Eq. 3 with stochastic gradient descent (SGD) yields the following update rule,

$$\hat{W}_t \leftarrow \hat{W}_{t-1} - \eta \nabla_{\hat{W}} \mathcal{L}_{\mathcal{T}} + \xi, \tag{4}$$

in which t denotes the training step and η is the learning rate. Additionally, $\xi \in \mathbb{R}^d$ models the stochasticity noise of the optimization. $\xi \sim \mathcal{N}(0, \frac{2}{\beta}\boldsymbol{I})$ is assumed be Gaussian with zero mean and a covariance of $\frac{2}{\beta}\boldsymbol{I}$ where β in thermodynamic terminology is referred to as the inverse temperature.

Macroscopic variables: The quantity of interest in this work, is the expected generalization error of the student, determined by averaging the student's error over all possible input-target pairs and noise realizations, as,

$$\mathcal{L}_{\mathcal{G}} := \frac{1}{2}\mathbb{E}_{\boldsymbol{x},W,\epsilon}\big[(y^* - \hat{y})^2\big]. \tag{5}$$

As shown in Stephenson & Lee (2021), if $n, d \to \infty$ with a constant ratio $\frac{n}{d}$, Eq. 5 can be written as a function of **scalar** macroscopic variables $R, Q \in \mathbb{R}$,

$$\mathcal{L}_{\mathcal{G}} = \frac{1}{2}(1 + Q - 2R + \sigma_\epsilon^2), \tag{6}$$

where σ_ϵ^2 is the variance of the teacher's output noise and,

$$R := \frac{1}{d}W^T \boldsymbol{F}\hat{W}, \qquad Q := \frac{1}{d}\hat{W}^T \boldsymbol{F}^T \boldsymbol{F}\hat{W}. \tag{7}$$

Both R and Q have clear interpretations; R is the dot-product between the teacher's weights W and the student's *modulated* weights $\boldsymbol{F}\hat{W}$, hence can be interpreted as the **alignment between the teacher and the student**. Similarly, Q can be interpreted as the **student's modulated norm**. The negative sign of R in Eq. 6 suggests that the larger R is, the smaller the generalization error gets. At the same time, Q appears with a positive sign suggesting the students with a smaller (modulated) norm generalize better.

As remark, note that both R and Q are functions of $\hat{W}$ which itself is a function of training time t and the regularization strength λ. Therefore, hereafter, we denote the above quantities as $\mathcal{L}_{\mathcal{G}}(t, \lambda)$, $R(t, \lambda)$, and $Q(t, \lambda)$.

2.3 Main Result

General case. Let $Z, X \in \mathbb{R}^{n\times d}$ denote the input matrices for the teacher and student, where $X := \boldsymbol{F}^T Z$. For a general modulation matrix $\boldsymbol{F}$, the diagonal matrix Λ contains the eigenvalues of the student's input covariance matrix,

$$C := X^T X = \boldsymbol{F}^T Z^T Z \boldsymbol{F} = P\Lambda P^T, \tag{8}$$

where in the last step we have performed a singular value decomposition (SVD). Solving the dynamics of gradient descent as in Eq. 4, we arrive at the following exact analytical expressions for $R(t, \lambda)$ and $Q(t, \lambda)$,

$$\begin{aligned} R(t,\lambda) &= \frac{n}{d}\left(1 - \frac{1}{n}\sum_{\mu=1}^{n}\Big[((1-\eta\lambda)\mathbb{I} - \eta C)^t\Big]_{\mu\mu}\right), \\ Q(t,\lambda) &= \frac{\sigma_\epsilon^2}{d}\sum_{\mu=1}^{n}\left[C^{-1}\left(\mathbb{I} - ((1-\eta\lambda)\mathbb{I} - \eta C)^t\right)^2\right]_{\mu\mu} + \frac{1}{d}\sum_{\mu=1}^{n}\left[\left(\mathbb{I} - ((1-\eta\lambda)\mathbb{I} - \eta C)^t\right)^2\right]_{\mu\mu}. \end{aligned} \tag{9}$$

Special case: Slow and Fast features. For interpretability, we now study a case where the modulation matrix $\boldsymbol{F}$ has a specific structure described in Assumption 1.

Assumption 1 *The modulation matrix $\boldsymbol{F}$ with a singular value decomposition $\boldsymbol{F} := U\Sigma V^T$ has two sets of singular values such that the first p singular values are equal to σ_1 and the remaining $d-p$ singular values are equal to σ_2. We let the condition number of $\boldsymbol{F}$ to be denoted by $\kappa := \frac{\sigma_1}{\sigma_2} > 1$.*

By employing the replica method of statistical physics (Gardner, 1988; Gardner & Derrida, 1988), we now derive approximate expressions for $R(t, \lambda)$. We first define the following auxiliary variables,

$$\alpha_1 = \frac{n}{p},\ \alpha_2 = \frac{n}{d-p}, \qquad \gamma_1 = \frac{\lambda + \frac{1}{\eta t}}{\frac{p}{d}\sigma_1^2},\ \gamma_2 = \frac{\lambda + \frac{1}{\eta t}}{\frac{d-p}{d}\sigma_2^2}, \tag{10}$$

and also let,

$$a_i = 1 + \frac{2\gamma_i}{(1-\alpha_i-\gamma_i) + \sqrt{(1-\alpha_i-\gamma_i)^2 + 4\gamma_i}}, \qquad \text{for} \qquad i \in \{1,2\}. \tag{11}$$

The closed-from scalar expression for $R(t,\lambda)$ is then given by,

$$R(t,\lambda) = R_1 + R_2, \quad \text{where,} \quad R_1 := \frac{n}{a_1 d}, \quad \text{and,} \quad R_2 := \frac{n}{a_2 d} \tag{12}$$

For $Q(t,\lambda)$, we accordingly define two more auxiliary variables,

$$b_i = \frac{\alpha_i}{a_i^2 - \alpha_i}, \quad c_i = 1 - 2R_i - \frac{n}{d}\frac{2-a_i}{a_i} \qquad \text{for} \qquad i \in \{1,2\}, \tag{13}$$

with which the closed-from scalar expression for $Q(t,\lambda)$ reads,

$$Q(t,\lambda) = Q_1 + Q_2, \quad \text{where,} \quad Q_1 := \frac{b_1 b_2 c_2 + b_1 c_1}{1 - b_1 b_2}, \quad \text{and,} \quad Q_1 := \frac{b_1 b_2 c_1 + b_2 c_2}{1 - b_1 b_2}. \tag{14}$$

By plugging Eqs. 12 and 14 into Eq. 6, one obtains a closed-form expression for $\mathcal{L}_{\mathcal{G}}(t,\lambda)$ as a function of training time and the regularization strength. Note that the speed at which R_1 and R_2 (similarly Q_1 and Q_2) evolve is a function of the condition number, $\kappa := \frac{\sigma_1}{\sigma_2}$.

2.4 Sketch of derivations

In this Section, we sketch the key steps in the derivation of our main results. For the sake of simplicity, here we only treat the noiseless teacher. Detailed proofs and further details are presented in Appendix A

Exact dynamics of SGD. Recall the gradient descent update rule in Eq. 4. For the linear model defined in Eqs. 1-2, learning is governed by the following discrete-time dynamics,

$$\hat{W}_t = \hat{W}_{t-1} + \frac{\eta}{n} X^T (XW - X\hat{W}_{t-1}) - \eta \hat{W}_{t-1} \tag{15}$$

For simplicity, we furthermore assume the student to be initialized at the origin, $\hat{W}_0 = \mathbf{0}$, where the dynamics admit the following exact closed-form solution,

$$\hat{W}_t = \left[I - \left((1-\eta\lambda)I - \eta X^T X\right)^t\right](X^T X + \lambda I)^{-1} X^T X W := \tilde{W}. \tag{16}$$

With a SVD on $X^T X$, Eqs. 9 can then be obtained by substituting $\hat{W}_t$ in Eqs. 7. As a remark, note that one can recover the results of Advani & Saxe (2017) by setting $\boldsymbol{F} \to I$ and $\eta \to 0$. In that case, the eigenvalues of $X^T X$ follow a Marchenko–Pastur distribution (Marchenko & Pastur, 1967).

Induced probability density of SGD. It is well-known (Kuhn & Bos, 1993) that probability distribution of weight configurations for network weights $\hat{W}$ trained via SGD on a loss $\mathcal{L}(\hat{W})$, tend to the Gibbs distribution such that,

$$P(\hat{W}) = \frac{1}{Z_\beta} e^{-\beta \mathcal{L}(\hat{W})}, \tag{17}$$

in which Z_β is the partition function $\left(\int \mathrm{d}\hat{W} \exp(-\beta \mathcal{L}(\hat{W}))\right)$ and β is called the *inverse temperature* of the thermodynamic system modeling our network setup. Intuitively, for small β, the distribution of $P(\hat{W})$ is almost uniform, while as $\beta \to \infty$, $P(\hat{W})$ becomes more concentrated around the minimum of the training loss.

It is important to highlight that Eq. 17 describes the equilibrium distribution of the student network's weights, i.e., at the end of training ($t \to \infty$). However, we are interested in studying the trajectory

of student's weights *during* the course of training, i.e., for finite t. To that end, the time-dependent probability density over $\hat{W}$ is given by,

$$P(\hat{W}, t) = \frac{1}{Z_\beta} e^{-\beta \tilde{\mathcal{L}}(\hat{W}, t)}, \quad \text{where} \quad \tilde{\mathcal{L}}(\hat{W}, t) := \frac{1}{2n} ||X\hat{W} - X\tilde{W}(t)||_2^2 + \frac{1}{2}\lambda ||\hat{W}||_2^2, \tag{18}$$

where $\tilde{W}(t)$ is defined in Eq. 16. To further simplify $P(\hat{W}, t)$, we note that $\tilde{\mathcal{L}}(\hat{W}, t)$ in Eq. 18 can be approximated by,

$$\tilde{\mathcal{L}}(\hat{W}, t) \approx \frac{1}{2n} ||X\hat{W} - XW||_2^2 + \big(\frac{1}{\eta t} + \lambda\big) ||\hat{W}||_2^2. \tag{19}$$

The updated expression for $\tilde{\mathcal{L}}(\hat{W}, t)$ enables us to use the replica method for the finite time learning.

Free-Energy To determine the average training and generalization error at training time t, one proceeds by first computing the free-energy of the system as,

$$f := -\left\langle \left\langle \frac{\ln Z_\beta}{\beta n} \right\rangle \right\rangle_{\boldsymbol{W}^*, \boldsymbol{x}}. \tag{20}$$

and extremize the above entity w.r.t. the macroscopic variables of interest, R, Q, to determine their average case estimates. The reason behind the evaluation of the free-energy f is that, unlike the (input and teacher averaged) partition function, its average and most probable values coincide.

Due to the logarithm inside the expectation, analytical computation of Eq. 20 is intractable. However, the replica method (Mézard et al., 1987) allows us to tackle this through the following identity,

$$\mathbb{E}\left[\ln Z_\beta\right] = \lim_{r \to 0} \frac{\mathbb{E}\left[Z_\beta\right]^r - 1}{r}. \tag{21}$$

Computation of the free-energy via this procedure and its subsequent extremization w.r.t R and Q, we arrive at Eqns. 12 and 14.

To summarize, using the replica method, we are able to cast the high-dimensional dynamics of SGD into simple scalar equations governing R and Q and consequently, the generalization error $\mathcal{L}_G$. While our analysis is limited to the specific teacher and student setup , this simple model already exhibits dynamics closely matching those observed in more complex networks, as we now illustrate.

3 Experimental Results

In this Section, we conduct numerical simulations to validate our analytical results and provide clear insights on the macroscopic dynamics of generalization. We also conduct experiments on real-world neural networks. We demonstrate that our teacher-student setup exhibits generalization behavior which is qualitatively similar to that of deep neural networks.

For simulations and evaluation of the theory, a teacher generates $n = 300$ examples with $d = 200$. The standard deviation of teacher noise is 0.1. Assumption 1 holds with $p = 100$ where we vary the condition number κ. For numerical results, we use SGD with a learning rate of 0.01 and 0.0 momentum. Numerical simulations are averaged over 100 random seeds.

For real-world experiments, we train a **ResNet18** (He et al., 2016) with large layer widths $[64, 2 \times 64, 4 \times 64, 8 \times 64]$. We follow the training setup of Nakkiran et al. (2019); Label noise with a probability 0.15 randomly assign an incorrect label to training examples. Noise is samples only once before the training starts. We train using Adam (Kingma & Ba, 2014) with learning rate of $1e-4$ for 1K epochs. Real-world experiments are averaged over 50 random seeds. To ensure reproducibility, we include the complete source code in a `GitHub repository` as well as an anonymous `Collab notebook`.

3.1 Match between theory and real-world experiments

We conduct an experiment on the classification task of CIFAR-10 (Krizhevsky et al., 2009) with varying amount of weight decay regularization strength λ. We monitor the generalization error (0-1 test error) during the course of training and visualize a heat-map of the generalization error for different λ's in Figure 2 (a).

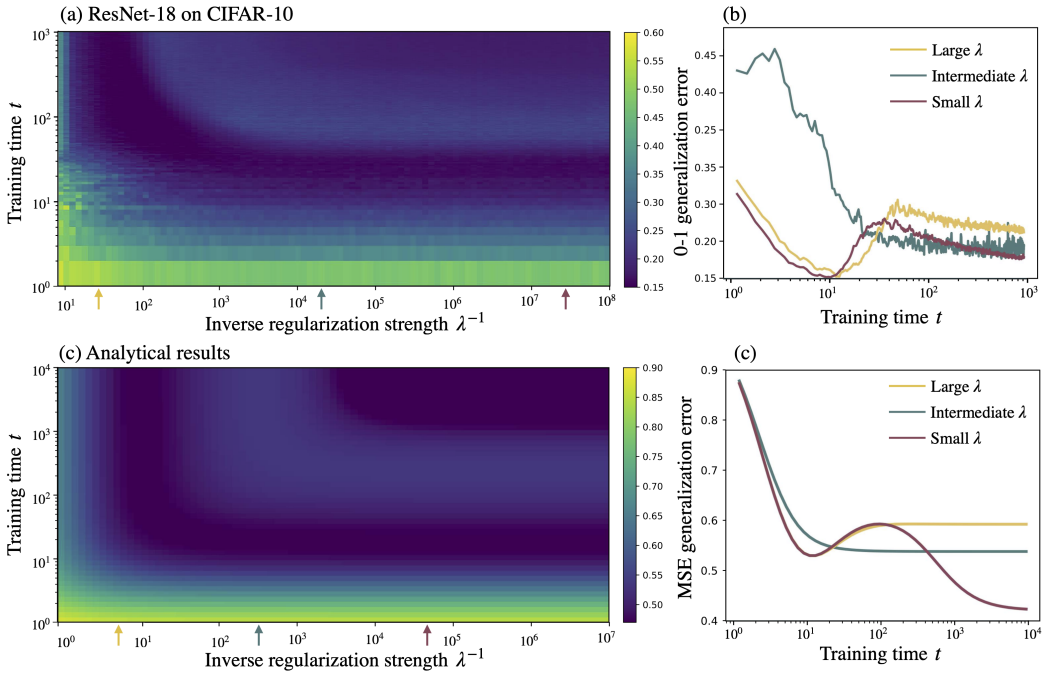

Figure 2: A qualitative comparison between a ResNet-18 and our analytical results. **(a)**: Heat-map of empirical generalization error (0-1 classification error) for the ResNet-18 trained on CIFAR-10 with 15% label noise. X-axis denotes the inverse of weight-decay regularization strength and Y-axis represents the training time. **(b)**: Heat-map of the analytical generalization error (mean squared error) for the linear teacher-student setup with $k = 100$, the condition number of the modulation matrix. **(c, d)**: Three slices of the heat-maps for large, intermediate, and small amounts of regularization. **Analysis**: As predicted by Eqs. 12 and 14, $k = 100$ implies that a fraction of features are learned 100 times faster that the rest. Intuitively, large amounts of regularization allow for the fast-learning features to be learned by not to overfit. Intermediate levels of regularization result in a classical U-shaped generalization curve but prevent slow features from learning. Small amounts of regularization allow for both fast and slow features to be learned, leading to double descent.

We also conduct a similar experiment with the teacher-student setup presented in Section 2.2. We visualize a heat-map of the generalization error which is the mean squared error (MSE) over test distribution in Figure 2 (b). Particularly, we plot Eqs. 12 and 14 with a constant $\kappa = 100$. As a remark, we note that a $\kappa = 100$ implies that some features are learned 100 times faster than other features.

Critically, it can be observed that in both experiments a model with intermediate levels of regularization displays a typical overfitting behavior where the generalization error decreases first and then overfits. This is consistent with Eq. 10 which indicates larger amounts of regularization prevent slow feature from being learned as λ and the inverse of time are summed. In our particular teacher-student setup, learning of slow features requires weights with large magnitudes which are mitigated by the weight-decay regularization. On the other hand, a model with smaller amount of regularization exhibits a larger double descent generalization curve, *i.e.*, with more training, the generalization error will eventually improve.

We also validate our derived analytical expressions by running numerical simulations which are presented in Figure 4.

3.2 The Phase Diagram

To further investigate the transition between the two phases of *classical single descent* and *double descent*, we explore the phase diagram. Recall that with Eq. 6, one can fully characterize the evo-

lution of the generalization dynamics in terms of two scalar variables instead of the d-dimensional parameter space. R and Q presented in Eq. 7 are macroscopic variables where R represents **the alignment between the teacher and the student** and Q is the **student's (modulated) norm**. Hence, a better generalization performance is achieved with larger R and smaller Q.

R and Q are not free parameters and both depend on the training dynamics through Eqs. 12 and 14. Nevertheless, it is instructive to visualize the generalization error for all pairs of $(R, Q) \in [0.0, 0.8] \times [0.0, 1.6]$. We also visualize the trajectories taken by SGD. At the time of initialization, $(R, Q) = (0, 0)$ as the models are initialized at the origin. As training time proceeds, values of R and Q follow the depicted trajectories. In Figure 3, different trajectories correspond to different values of κ, the condition number of the modulation matrix $\boldsymbol{F}$ in Eq. 2. The condition number implies that some features are learned κ times faster than the others. It is important to note that *the closer a trajectory gets to the lower-right, the smaller the generalization error is.*

The yellow curve which corresponds to the case with large $\kappa = 1e5$ meaning that some features are extremely slower than the other features that practically do not get learned. In that case, generalization error exhibits traditional over-fitting due to over-training. On the phase diagram, the yellow trajectory starts at $(0, 0)$ and moves towards Point A which has the lowest generalization error of this curve. Then as the training continues, Q increases and as $t \to \infty$ the trajectory lands at Point B which has the worse generalization error.

The curves in orange, green and blue correspond to trajectories with $\kappa = 1e3$, $\kappa = 1e2$, $\kappa = 1e1$, respectively. They follow the case of $\kappa = 1e5$ up to the vicinity of Point B, but then the trajectories slowly incline towards another fixed point, Point C signalling a second descent in the generalization error. The phase diagram along with the corresponding generalization curves in Figure 2 illustrate that features that are learned on a faster time-scale are responsible for the initial conventional U-shaped generalization curve, while the second descent can be attributed to the features that are learned at a slower time-scale.

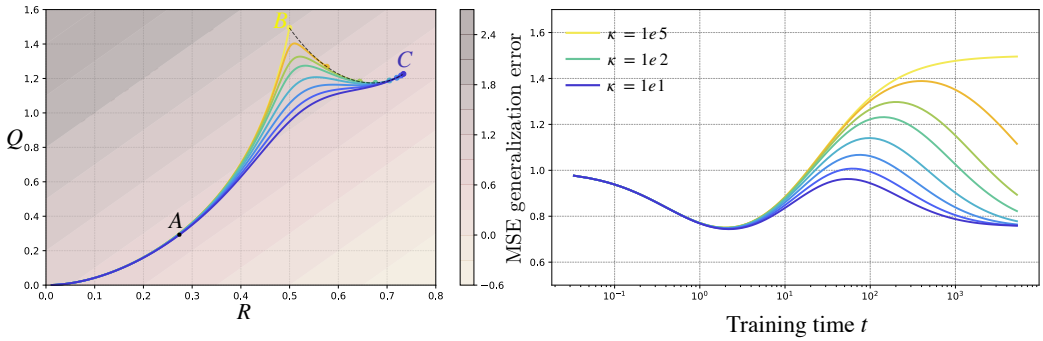

Figure 3: **Left**: Phase diagram of the generalization error as a function of $R(t)$ and $Q(t)$ (Eqs. 12 and 14). The generalization error for all pairs of $(R, Q) \in [0.0, 0.8] \times [0.0, 1.6]$ is contour-plotted in the background in shades of beige, with the best generalization performance being attained on the lower right part of the plot. The trajectories describe the evolution of $R(t)$ and $Q(t)$ as training proceeds. Each trajectory correspond to a different κ, the condition number of the modulation matrix $\boldsymbol{F}$ in Eq. 2. κ describes the ratio of the rates at which two sets of features are learned. **Right**: The corresponding generalization curves for different plotted over the training time axis. **Analysis**: The trajectory with $\kappa = 1e5$ (bright yellow) starts at the origin and advances towards point A (a descent in generalization error). Then by over-training, it converges to point B (an ascent in generalization error). For the other trajectories with smaller κ, a first descent in generalization error occurs up to the point A, then an ascent happens, but they no longer converge to point B. Instead, by further training, these trajectories converge to point C implying a second descent.

4 Related Work and Discussion

Double Descent: Although the term (model-wise) *double descent* has been introduced recently in reference to the population risk demonstrating a cusp at the *interpolation* point, a similar phe-

nomenon had already been studied in several decades-old works in the regression setting (Krogh & Hertz, 1992a; Opper, 1995; Opper & Kinzel, 1996; Bös, 1998) under a statistical physics framework. More recently, these behaviors have been investigated in the context of modern machine learning, both from an empirical (Nakkiran et al., 2019; Amari et al., 2020; Yang et al., 2020) and theoretical (Engel & Van den Broeck, 2001; Belkin et al., 2019a; Geiger et al., 2019; Advani & Saxe, 2017; Mei & Montanari, 2019; Gerace et al., 2020; d'Ascoli et al., 2020; Ba et al., 2019) point of view, to determine various limiting behavior of the training and generalization error.

Our work is closely related to that of Heckel & Yilmaz (2020), which provides *upper bounds* on the risk of single and two layer models in a regression setting during gradient-based training, on input-data with distinct feature variances. In that work, the authors demonstrate empirically that epoch-wise double descent can be viewed as a superposition of two or more bias-variance tradeoff curves due to different parts of the network being learned at different epochs. Furthermore, the authors show that epoch-wise double descent can be eliminated by appropriate selection of learning rates for individual network weights. Consistent with these findings, our work formalizes this phenomenon in terms of feature learning scales and provides closed-form predictions. Furthermore, by leveraging tools from statistical physics such as the replica method, we derive explicit equations for the generalization error as a function of regularization strength and training time in a teacher-student setup.

We believe our analysis takes a step towards studying the generalization dynamics of neural networks, by characterizing their behavior using a finite set of informative macroscopic parameters. In summary, while supporting the notion that the interaction of different feature learning speeds causes epoch-wise double descent, our work provides formal predictions of the dynamics that unfold during training.

Random matrix theory: Hastie et al. (2019); Advani et al. (2020); Belkin et al. (2020) use random matrix theory (RMT) tools to characterize the asymptotic generalization behavior of over-parameterized linear and random feature models. In an influential work, Mei & Montanari (2019) extend the same analysis to a random feature model and theoretically derive the model-wise double descent curve for a model with Tikhonov regularization. Jacot et al. (2020) also study double descent in ridge estimators and show an equivalence to kernel ridge regression. Pennington & Worah (2019) used RMT to study the curvature of single-hidden-layer neural network in an attempt to understand the efficacy of first-order optimization methods in training DNNs. In addition, Liang & Rakhlin (2020) take a similar approach to investigate implicit regularization in high dimensional ridgeless regression with nonlinear kernels.

Generalization properties of over-parameterized networks: Neyshabur et al. (2014) demonstrated that over-parameterized networks does not necessarily overfit thus suggesting the need of a new form of measure of model complexity other than network size. Subsequently, Neyshabur et al. (2018) suggest a novel complexity measure based on unit-wise capacities which correlates better with the behavior of test error with increasing network size. Chizat & Bach (2020) study the global convergence and superior generalization behavior of infinitely wide two-layer neural networks with logistic loss. Goldt et al. (2020) make use of the Gaussian Equivalence Theorem to study the generalization performance of two-layer neural networks and kernel models trained on data drawn from pre-trained generative models. Bai & Lee (2020) investigated the gap between the empirical performance of over-parameterized networks and their NTK counterparts, first proposed by Jacot et al. (2018).

In relation to the aspect of different feature learning scales, Rahaman et al. (2019) had observed that DNNs have a tendency towards learning simple target functions first that can allow for good generalization behavior of various data samples. Pezeshki et al. (2020) also identify and provide explanation for a feature learning imbalance exhibited by over-parameterized networks trained via gradient descent on cross-entropy loss, with the networks learning only a subset of the full feature spectrum over training. More recently though, Zhang & Wu (2020), show that certain DNNs models prioritize learning high-frequency components first followed by the learning of slow but informative features, leading to the second descent of the test error as observed in epoch-wise double descent.

Limitations. It should be noted that studying finer details of the dynamics would require a more precise model of the neural networks. Clearly, our proposed model is not a universal and unique way to model the dynamics of the complex, over-parameterized deep neural networks.

Social Impact. The authors do not foresee a negative social impact specifically arising from this rather theoretical work.

REFERENCES

Madhu S Advani and Andrew M Saxe. High-dimensional dynamics of generalization error in neural networks. *arXiv preprint arXiv:1710.03667*, 2017.

Madhu S Advani, Andrew M Saxe, and Haim Sompolinsky. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132:428–446, 2020.

Shun-ichi Amari, Jimmy Ba, Roger Grosse, Xuechen Li, Atsushi Nitanda, Taiji Suzuki, Denny Wu, and Ji Xu. When does preconditioning help or hurt generalization? *arXiv preprint arXiv:2006.10732*, 2020.

Jimmy Ba, Murat Erdogdu, Taiji Suzuki, Denny Wu, and Tianzong Zhang. Generalization of two-layer neural networks: An asymptotic viewpoint. In *International conference on learning representations*, 2019.

Yu Bai and Jason D. Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks, 2020.

Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019a.

Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019b.

Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.

Carl M Bender and Steven A Orszag. *Advanced mathematical methods for scientists and engineers I: Asymptotic methods and perturbation theory*. Springer Science & Business Media, 2013.

S Bös, W Kinzel, and M Opper. Generalization ability of perceptrons with continuous outputs. *Physical Review E*, 47(2):1384, 1993.

Siegfried Bös. Statistical mechanics approach to early stopping and weight decay. *Physical Review E*, 58(1):833, 1998.

Lenaic Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pp. 1305–1338. PMLR, 2020.

Stéphane D'Ascoli, Maria Refinetti, Giulio Biroli, and Florent Krzakala. Double trouble in double descent: Bias and variance(s) in the lazy regime. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 2280–2290. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/d-ascoli20a.html.

Stéphane d'Ascoli, Maria Refinetti, Giulio Biroli, and Florent Krzakala. Double trouble in double descent: Bias and variance (s) in the lazy regime. In *International Conference on Machine Learning*, pp. 2280–2290. PMLR, 2020.

Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001.

Jerome Friedman, Trevor Hastie, Robert Tibshirani, et al. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.

Elizabeth Gardner. The space of interactions in neural network models. *Journal of physics A: Mathematical and general*, 21(1):257, 1988.

Elizabeth Gardner and Bernard Derrida. Optimal storage properties of neural network models. *Journal of Physics A: Mathematical and general*, 21(1):271, 1988.

Elizabeth Gardner and Bernard Derrida. Three unfinished works on the optimal storage capacity of networks. *Journal of Physics A: Mathematical and General*, 22(12):1983, 1989.

Mario Geiger, Stefano Spigler, Stephane d'Ascoli, Levent Sagun, Marco Baity-Jesi, Giulio Biroli, and Matthieu Wyart. Jamming transition as a paradigm to understand the loss landscape of deep neural networks. *Physical Review E*, 100(1):012115, 2019.

Federica Gerace, Bruno Loureiro, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Generalisation error in learning with random features and the hidden manifold model. In *International Conference on Machine Learning*, pp. 3452–3462. PMLR, 2020.

Sebastian Goldt, Galen Reeves, Marc Mézard, Florent Krzakala, and Lenka Zdeborová. The gaussian equivalence of generative models for learning with two-layer neural networks. *arXiv e-prints*, pp. arXiv–2006, 2020.

Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*. MIT press Cambridge, 2016.

Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Reinhard Heckel and Fatih Furkan Yilmaz. Early stopping in deep networks: Double descent and how to eliminate it. *arXiv preprint arXiv:2007.10099*, 2020.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8571–8580, 2018.

Arthur Jacot, Berfin Simsek, Francesco Spadaro, Clément Hongler, and Franck Gabriel. Implicit regularization of random feature models. In *International Conference on Machine Learning*, pp. 4631–4640. PMLR, 2020.

Yoshiyuki Kabashima, Tadashi Wadayama, and Toshiyuki Tanaka. A typical reconstruction limit for compressed sensing based on lp-norm minimization. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(09):L09003, 2009.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

Anders Krogh and John A Hertz. Generalization in a linear perceptron in the presence of noise. *Journal of Physics A: Mathematical and General*, 25(5):1135, 1992a.

Anders Krogh and John A Hertz. A simple weight decay can improve generalization. In *Advances in neural information processing systems*, pp. 950–957, 1992b.

Florent Krzakala, Marc Mézard, François Sausset, YF Sun, and Lenka Zdeborová. Statistical-physics-based reconstruction in compressed sensing. *Physical Review X*, 2(2):021005, 2012.

R Kuhn and S Bos. Statistical mechanics for neural networks with continuous-time dynamics. *Journal of Physics A: Mathematical and General*, 26(4):831, 1993.

Daniel Kunin, Javier Sagastuy-Brena, Surya Ganguli, Daniel LK Yamins, and Hidenori Tanaka. Neural mechanics: Symmetry and broken conservation laws in deep learning dynamics. *arXiv preprint arXiv:2012.04728*, 2020.

Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel "ridgeless" regression can generalize. *The Annals of Statistics*, 48(3), Jun 2020. ISSN 0090-5364. doi: 10.1214/19-aos1849. URL http://dx.doi.org/10.1214/19-AOS1849.

Vladimir Alexandrovich Marchenko and Leonid Andreevich Pastur. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik*, 114(4):507–536, 1967.

Song Mei and Andrea Montanari. The generalization error of random features regression: precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355*, 2019.

Marc Mézard, Giorgio Parisi, and Miguel Virasoro. *Spin glass theory and beyond: an introduction to the Replica Method and its applications*, volume 9. World Scientific Publishing Company, 1987.

Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: where bigger models and more data hurt. *ICLR 2020, arXiv preprint arXiv:1912.02292*, 2019.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.

Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. Towards understanding the role of over-parametrization in generalization of neural networks, 2018.

Manfred Opper. Statistical mechanics of learning: Generalization. *The handbook of brain theory and neural networks*, pp. 922–925, 1995.

Manfred Opper and Wolfgang Kinzel. Statistical mechanics of generalization. In *Models of neural networks III*, pp. 151–209. Springer, 1996.

Jeffrey Pennington and Pratik Worah. Nonlinear random matrix theory for deep learning. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124005, 2019.

Mohammad Pezeshki, Sékou-Oumar Kaba, Yoshua Bengio, Aaron C. Courville, Doina Precup, and Guillaume Lajoie. Gradient starvation: A learning proclivity in neural networks. *CoRR*, abs/2011.09468, 2020. URL https://arxiv.org/abs/2011.09468.

Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pp. 5301–5310. PMLR, 2019.

Frederick Reif. *Fundamentals of statistical and thermal physics*. Waveland Press, 2009.

Hyunjune Sebastian Seung, Haim Sompolinsky, and Naftali Tishby. Statistical mechanics of learning from examples. *Physical review A*, 45(8):6056, 1992.

Cory Stephenson and Tyler Lee. When and how epochwise double descent happens. *arXiv preprint arXiv:2108.12006*, 2021.

Vladimir N. Vapnik. *The nature of statistical learning theory*. Wiley, New York, 1st edition, September 1998. ISBN 978-0-471-03003-4.

Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. In *International Conference on Machine Learning*, pp. 10767–10777. PMLR, 2020.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

Xiao Zhang and Dongrui Wu. Rethink the connections among generalization, memorization and the spectral bias of dnns. *CoRR*, abs/2004.13954, 2020. URL https://arxiv.org/abs/2004.13954.

A APPENDIX

A.1 REPLICA TRICK

In the following, we detail the mathematical arguments leading to the *replica trick* expression. For some $r \to 0$, we can write for any scalar x:

$$\begin{aligned} x^r &= \exp(r \ln x) = \lim_{r\to 0} 1 + r \ln x \\ \Rightarrow \lim_{r\to 0} r \ln x &= \lim_{r\to 0} x^r - 1 \\ \Rightarrow \ln x &= \lim_{r\to 0} \frac{x^r - 1}{r} \\ \therefore \mathbb{E}[\ln x] &= \lim_{r\to 0} \frac{\mathbb{E}[x^r] - 1}{r}, \;\; \mathbb{E} : \text{averaging} \end{aligned} \tag{22}$$

A.2 GENERALIZATION ERROR

The generalization error for our linear teacher-student setup is defined to be the average of the MSE over all possible teacher and student input configurations, $\boldsymbol{x}$, i.e.,

$$\begin{aligned} \mathcal{E}_{\mathcal{G}} &= \frac{1}{2N} \int \left(y(\boldsymbol{W}, \boldsymbol{x}) - y^*(\boldsymbol{W}^*, \boldsymbol{z})\right)^2 \rho(\boldsymbol{x}) \mathrm{d}\boldsymbol{x} \\ &= \frac{1}{2N} \int \left(\boldsymbol{W} \cdot \boldsymbol{x} - \boldsymbol{W}^* \cdot \boldsymbol{x}\right)^2 \rho(\boldsymbol{x}) \mathrm{d}\boldsymbol{x} \\ &= \frac{1}{2N} \int \left(\boldsymbol{W}^T \boldsymbol{x} \boldsymbol{x}^T \boldsymbol{W} - 2\boldsymbol{W}^T \boldsymbol{x} (\boldsymbol{x})^T \boldsymbol{W} + (\boldsymbol{W})^T \boldsymbol{x} (\boldsymbol{x})^T \boldsymbol{W}\right) \rho(\boldsymbol{x}) \mathrm{d}\boldsymbol{x} \\ &= \frac{1}{2N} \left(\boldsymbol{W}^T \Gamma \boldsymbol{W} - 2\boldsymbol{W}^T \Delta \boldsymbol{W} + |\boldsymbol{W}|^2\right) \\ &= \frac{1}{2} \left(Q_0 - 2R + 1\right) \end{aligned} \tag{23}$$

where, N is the number of input samples, $Q_0 = \boldsymbol{W}^T \Gamma \boldsymbol{W}$, $R = \boldsymbol{W}^T \Delta \boldsymbol{W}$ and $|\boldsymbol{W}|^2 = 1$. It is to be noted that in the above, we have used:

$$\langle\langle x^\mu \rangle\rangle = 0, \; \langle\langle x_i^\mu x_j^\nu \rangle\rangle = \delta^{\mu\nu} \delta_{ij} \tag{24}$$

$$\tag{25}$$

A.3 THERMAL AND SELF-AVERAGED TRAINING AND GENERALIZATION ERROR

The self-averaged free energy (per unit weight) of our student network, is given by (Engel & Van den Broeck, 2001),

$$-\beta f = \frac{1}{d} \langle\langle \ln Z \rangle\rangle_{\boldsymbol{x}, \boldsymbol{W}^*} \tag{26}$$

Here, $\beta = 1/T$ is the inverse temperature parameter corresponding to our statistical ensemble, d the (teacher) student network width, and Z the partition function of the system defined as (n: number of training examples),

$$\begin{aligned} Z &= \int_{-\infty}^{\infty} \mathrm{d}\mu\left(\mathbf{W}\right) e^{-\beta n \mathcal{L}_{\mathcal{T}}} \\ \text{with,} \;\; \mathrm{d}\mu\left(\mathbf{W}\right) &= \int_0^{\infty} \frac{\mathrm{d}Q_0}{\sqrt{2\pi}} \prod_{i=1}^{P} \frac{\mathrm{d}W_i}{\sqrt{2\pi}} \delta\left(\boldsymbol{W}^T \Gamma \boldsymbol{W} - P Q_0\right) \end{aligned} \tag{27}$$

Here, Q_0 can be perceived to be a target norm the student weights $\boldsymbol{W}$ are being constrained to. Next, by making use of the replica trick (Eq. 22), we can next write Eq. 26 as,

$$-\beta f = \frac{1}{d} \lim_{r\to 0} \frac{\langle\langle Z^n \rangle\rangle_{\boldsymbol{x}, \boldsymbol{W}} - 1}{r}, \quad r \in \mathbb{Z} \tag{28}$$

Here, Z^r corresponds to the combined average of n different replicas of our student network with the same realization of the random inputs. To evaluate $\langle\langle Z^r\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}}$, we begin by introducing the teacher and the student network functions, $y^{*\mu}$ and $y^{a,\mu}$,

$$\begin{aligned} y^{*\mu} &= \frac{1}{\sqrt{d}}\boldsymbol{W}^{*T}\boldsymbol{x}^{*\mu} \\ y^{\mu} &= \frac{1}{\sqrt{d}}\boldsymbol{W}^{T}\boldsymbol{x}^{\mu} \end{aligned} \tag{29}$$

as Gaussian variables (with $P \to \infty$), in the partition function, to obtain,

$$\begin{aligned} \langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} &= \prod_{a=1}^{n}\prod_{\mu=1}^{N}\int \mathrm{d}\mu\left(\mathbf{W}^a\right)\mathrm{d}y_a^\mu \mathrm{d}(y^*)^\mu e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)} \\ &\quad\times\left\langle\left\langle \delta\left(y^{*\mu}-\frac{1}{\sqrt{d}}\boldsymbol{W}^{*T}\boldsymbol{x}^{*\mu}\right)\delta\left(y_a^\mu-\frac{1}{\sqrt{d}}\boldsymbol{W}_a^T\boldsymbol{x}^\mu\right)\right\rangle\right\rangle_{\boldsymbol{x},\boldsymbol{W}} \\ &= \prod_{a=1}^{n}\prod_{\mu=1}^{N}\int \mathrm{d}\mu\left(\mathbf{W}^a\right)\frac{\mathrm{d}y_a^\mu \mathrm{d}\hat{y}_a^\mu}{2\pi}\frac{\mathrm{d}y^{*\mu}\mathrm{d}\hat{y}^{*\mu}}{2\pi}e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)}e^{iy^{*\mu}\hat{y}^{*\mu}+iy_a^\mu\hat{y}_a^\mu} \\ &\quad\times\left\langle\left\langle \exp\left(-\frac{i}{\sqrt{d}}\hat{y}^{*\mu}\boldsymbol{W}^{*T}\boldsymbol{x}^{*\mu}-\frac{i}{\sqrt{d}}\hat{y}_a^\mu\boldsymbol{W}_a^T\boldsymbol{x}^\mu\right)\right\rangle\right\rangle_{\boldsymbol{x},\boldsymbol{W}} \end{aligned} \tag{30}$$

where in the last line above, we have expressed the inserted δ functions using their integral representations. Now, since our teacher and student inputs satisfy Eq. 24, hence, one can write in the limit $P \to \infty$,

$$\begin{aligned} &\left\langle\left\langle \exp\left(-\frac{i}{\sqrt{d}}\hat{y}^{*\mu}\boldsymbol{W}^{*T}\boldsymbol{x}^{*\mu}-\frac{i}{\sqrt{d}}\hat{y}_a^\mu\boldsymbol{W}_a^T\boldsymbol{x}^\mu\right)\right\rangle\right\rangle_{\boldsymbol{x},\boldsymbol{W}} \\ &= \left\langle\left\langle \exp\Big(-\frac{1}{2}\sum_{\mu,a,b}\hat{y}_a^\mu\hat{y}_b^\mu\frac{1}{d}\sum_{ij}W_a^i\Gamma_{ij}W_b^j-\sum_{\mu,a}\hat{y}^{*\mu}\hat{y}_a^\mu\frac{1}{d}\sum_{ij}W_a^i\Delta_{ij}W^{*j} \right.\right. \\ &\qquad \left.\left. -\frac{1}{2}\sum_\mu(\hat{y}^{*\mu})^2\frac{1}{d}\sum_i(W^{*i})^2\Big)\right\rangle\right\rangle_{\boldsymbol{W}} \\ &= \left\langle\left\langle \exp\Big(-\frac{Q_0}{2}\sum_{\mu,a}(\hat{y}_a^\mu)^2-\frac{1}{2}\sum_{\mu,\langle a,b\rangle}\hat{y}_a^\mu\hat{y}_b^\mu\frac{1}{d}\sum_{ij}W_a^i\Gamma_{ij}W_b^j \right.\right. \\ &\qquad \left.\left. -\sum_{\mu,a}\hat{y}^{*\mu}\hat{y}_a^\mu\frac{1}{d}\sum_{ij}W_a^i\Delta_{ij}W^{*j}-\frac{1}{2}\sum_\mu(\hat{y}^{*\mu})^2\Big)\right\rangle\right\rangle_{\boldsymbol{W}} \end{aligned} \tag{31}$$

where, $\langle a,b\rangle$ stands for a sum over a and b such that $a \neq b$. Additionally, in the last line, we have used the relations, $\boldsymbol{W}^T\Gamma\boldsymbol{W} = PQ_0$ and, $\boldsymbol{W}^{*T}\boldsymbol{W} = P$. Therefore, assimilating the above, we can write,

$$\begin{aligned} \langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} &= \prod_{a=1}^{r}\prod_{\mu=1}^{n}\int \mathrm{d}\mu\left(\mathbf{W}^a\right)\frac{\mathrm{d}y_a^\mu \mathrm{d}\hat{y}_a^\mu}{2\pi}\frac{\mathrm{d}y^{*\mu}\mathrm{d}\hat{y}^{*\mu}}{2\pi}e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)}e^{iy^{*\mu}\hat{y}^{*\mu}+iy_a^\mu\hat{y}_a^\mu} \\ &\quad\times\left\langle\left\langle \exp\Big(-\frac{Q_0}{2}\sum_{\mu,a}(\hat{y}_a^\mu)^2-\frac{1}{2}\sum_{\mu,\langle a,b\rangle}\hat{y}_a^\mu\hat{y}_b^\mu\frac{1}{d}\sum_{ij}W_a^i\Gamma_{ij}W_b^j \right.\right. \\ &\qquad \left.\left. -\sum_{\mu,a}\hat{y}^{*\mu}\hat{y}_a^\mu\frac{1}{d}\sum_{ij}W_a^i\Delta_{ij}W^{*j}-\frac{1}{2}\sum_\mu(\hat{y}^{*\mu})^2\Big)\right\rangle\right\rangle_{\boldsymbol{W}} \end{aligned} \tag{32}$$

To make further progress, we introduce the auxiliary variables,

$$\sum_{ija} W_a^i\Delta_{ij}W^{*j} = dR_a, \tag{33}$$

$$\sum_{ij\langle a,b\rangle} W_a^i\Gamma_{ij}W_b^j = dQ_{ab} \tag{34}$$

via the respective δ functions, to arrive at,

$$\begin{aligned}
\langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} &= \prod_{\mu,a,b}\int \mathrm{d}\mu\left(\mathbf{W}^a\right)\frac{\mathrm{d}y_a^\mu \mathrm{d}\hat{y}_a^\mu}{2\pi}\frac{\mathrm{d}y^{*\mu}\mathrm{d}\hat{y}^{*\mu}}{2\pi}e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)}e^{iy^{*\mu}\hat{y}^{*\mu}+iy_a^\mu\hat{y}_a^\mu} \\
&\quad\times\int P\mathrm{d}Q^{ab}\int P\mathrm{d}R^a\;\delta\left(\sum_{i,j,a}W_a^i\Delta_{i,j}W^{*j}-PR_a\right)\delta\left(\sum_{ij\langle a,b\rangle}W_a^i\Gamma_{ij}W_b^j-PQ_{ab}\right) \\
&\quad\times\left\langle\left\langle \exp\Big(-\frac{Q_0}{2}\sum_{\mu,a}(\hat{y}_a^\mu)^2-\frac{1}{2}\sum_{\mu,\langle a,b\rangle}\hat{y}_a^\mu\hat{y}_b^\mu Q_{ab}-\sum_{\mu,a}\hat{y}^{*\mu}\hat{y}_a^\mu R_a-\frac{1}{2}\sum_{\mu}(\hat{y}^{*\mu})^2\Big)\right\rangle\right\rangle_{\boldsymbol{W}} \\
&= \prod_{\mu,a,b}\int \mathrm{d}\mu\left(\mathbf{W}^a\right)\frac{\mathrm{d}y_a^\mu \mathrm{d}\hat{y}_a^\mu}{2\pi}\frac{\mathrm{d}y^{*\mu}\mathrm{d}\hat{y}^{*\mu}}{2\pi}e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)}e^{iy^{*\mu}\hat{y}^{*\mu}+iy_a^\mu\hat{y}_a^\mu} \\
&\quad\times\int P\mathrm{d}Q^{ab}\int P\mathrm{d}R^a\;\delta\left(\sum_{ija}W_a^i\Delta_{ij}W^{*j}-PR_a\right)\delta\left(\sum_{ij\langle a,b\rangle}W_a^i\Gamma_{ij}W_b^j-PQ_{ab}\right) \\
&\quad\times\left\langle\left\langle \exp\Big(-\frac{Q_0}{2}\sum_{\mu,a}(\hat{y}_a^\mu)^2-\frac{1}{2}\sum_{\mu,\langle a,b\rangle}\hat{y}_a^\mu\hat{y}_b^\mu Q_{ab}-\sum_{\mu,a}\hat{y}^{*\mu}\hat{y}_a^\mu R_a-\frac{1}{2}\sum_{\mu}(\hat{y}^{*\mu})^2\Big)\right\rangle\right\rangle_{\boldsymbol{W}}
\end{aligned} \tag{35}$$

Repeating the procedure of expressing the above δ functions using their integral representations, we then get ($\alpha = n/d$),

$$\begin{aligned}
\langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{x}^*,\boldsymbol{W}} &= \int\prod_{a,b}\frac{\mathrm{d}Q_0}{\sqrt{2\pi}}\frac{\mathrm{d}\hat{Q}_{0a}}{4\pi}\frac{\mathrm{d}Q_{ab}\hat{Q}_{ab}}{2\pi/d}\frac{\mathrm{d}R_a\hat{R}_a}{2\pi/d}\exp\Big(\frac{iP}{2}\sum_a Q_0\hat{Q}_{0a}+iP\sum_{a<b}Q^{ab}\hat{Q}^{ab} \\
&\quad+iP\sum_a R^a\hat{R}^a\Big)\int\prod_{i,a}\frac{\mathrm{d}W_i^a}{\sqrt{2\pi}}\exp\Big(-\frac{i}{2}\sum_{i,j,a}\hat{Q}_{0a}W_a^i\Gamma_{ij}W_a^j \\
&\quad-i\sum_{i,j,a<b}\hat{Q}_{ab}W_a^i\Gamma_{ij}W_b^j-i\sum_{i,j,a}\hat{R}_a\Delta_{ij}W_a^j\Big)\times \\
&\quad\int\prod_{\mu,a}\frac{\mathrm{d}y_a^\mu\mathrm{d}\hat{y}_a^\mu}{2\pi}\frac{\mathrm{d}y^{*\mu}}{\sqrt{2\pi}}e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)}\exp\Big(-\frac{1}{2}\sum_\mu(y^{*\mu})^2+i\sum_{\mu,a}\hat{y}_a^\mu\hat{y}_a^\mu \\
&\quad-\frac{1}{2}\sum_{a,\mu}\left(1-R_a^2\right)(\hat{y}_a^\mu)^2-\frac{1}{2}\sum_{\mu,\langle a,b\rangle}\hat{y}_a^\mu\hat{y}_b^\mu\left(Q^{ab}-R^aR^b\right)-i\sum_{\mu,a}y^{*\mu}\hat{y}_a^\mu R^a\Big)
\end{aligned} \tag{36}$$

If we now, perform a singular value decomposition of the covariance matrix Γ as, $\Gamma = \mathrm{U}^T\mathrm{S}\mathrm{U} = \mathrm{V}^T\mathrm{V}$, where S: matrix of singular values of Γ, and we have expressed, $\mathrm{V} = \mathrm{S}^{1/2}\mathrm{U}$, then one can proceed to write,

$$\begin{aligned}
\langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} &= \frac{1}{\det|V|}\int\prod_{a,b}\frac{\mathrm{d}Q_0}{\sqrt{2\pi}}\frac{\mathrm{d}\hat{Q}_{0a}}{4\pi}\frac{\mathrm{d}Q_{ab}\hat{Q}_{ab}}{2\pi/d}\frac{\mathrm{d}R_a\hat{R}_a}{2\pi/d}\exp\Big(\frac{iP}{2}\sum_a Q_0\hat{Q}_{0a} \\
&\quad+iP\sum_{a<b}Q^{ab}\hat{Q}^{ab}+iP\sum_a R^a\hat{R}^a\Big)\int\prod_{i,a}\frac{\mathrm{d}\tilde{W}_i^a}{\sqrt{2\pi}}\exp\Big(-\frac{i}{2}\sum_{i,a}\hat{Q}_{0a}\left(\tilde{W}_a^i\right)^2 \\
&\quad-i\sum_{i,a<b}\hat{Q}_{ab}\tilde{W}_a^i\tilde{W}_b^i-i\sum_{i,j,a}\hat{R}_a\tilde{W}_a^j\Big)\times\int\prod_{\mu,a}\frac{\mathrm{d}y_a^\mu\mathrm{d}\hat{y}_a^\mu}{2\pi}\frac{\mathrm{d}y^{*\mu}}{\sqrt{2\pi}}e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)} \\
&\quad\exp\Big(-\frac{1}{2}\sum_\mu(y_\mu^*)^2+i\sum_{\mu,a}\hat{y}_a^\mu\hat{y}_a^\mu-\frac{1}{2}\sum_{a,\mu}\left(1-R_a^2\right)(\hat{y}_a^\mu)^2-i\sum_{\mu,a}y^{*\mu}\hat{y}_a^\mu R^a \\
&\quad-\frac{1}{2}\sum_{\mu,\langle a,b\rangle}\hat{y}_a^\mu\hat{y}_b^\mu\left(Q^{ab}-R^aR^b\right)\Big)
\end{aligned} \tag{37}$$

having expressed, $\tilde{\boldsymbol{W}}_a = \mathrm{V}\boldsymbol{W}_a$, and identifying $\Delta = \mathrm{S}^{1/2}\mathrm{U}$ from our definitions. Now, since in the above, the W_i^a integrals factorize in i, and similarly the y_a^μ, $\hat{y}_a^\mu$ and $\mathrm{d}y^{*\mu}$ factorize in μ, one can proceed to write:

$$\begin{aligned}\langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} = \frac{1}{\det|V|}\int\prod_{a,b}\frac{\mathrm{d}Q_0\mathrm{d}\hat{Q}_{0a}}{\sqrt{2\pi}4\pi}\frac{\mathrm{d}Q_{ab}\hat{Q}_{ab}}{2\pi/d}\frac{\mathrm{d}R_a\hat{R}_a}{2\pi/d}\exp\Big(P\Big[\frac{i}{2}\sum_a Q_0\hat{Q}_{0a} \\ + i\sum_{a<b}Q^{ab}\hat{Q}^{ab} + i\sum_a R^a\hat{R}^a + G_S(\hat{Q}_{0a},\hat{Q}^{ab},\hat{R}^a) + \alpha G_E(Q^{ab},R^a)\Big]\Big)\end{aligned} \tag{38}$$

where,

$$\begin{aligned}G_S(\hat{Q}_{0a},\hat{Q}^{ab},\hat{R}^a) &= \ln\int\prod_a\frac{\mathrm{d}\tilde{W}^a}{\sqrt{2\pi}}\exp\Big(-\frac{i}{2}\sum_a\hat{Q}_{0a}\tilde{W}_a^i\tilde{W}_a^i - i\sum_{a<b}\hat{Q}_{ab}\tilde{W}_a\tilde{W}_b - i\sum_a\hat{R}_a\tilde{W}_a\Big) \\ G_E(Q^{ab},R^a) &= \ln\int\prod_a\frac{\mathrm{d}y_a\mathrm{d}\hat{y}_a}{2\pi}\frac{\mathrm{d}y^*}{\sqrt{2\pi}}e^{-\beta N\mathcal{E}_{\mathcal{T}}(y_a,y^*)}\exp\Big(-\frac{1}{2}(y^*)^2 + i\sum_a\hat{y}_a\hat{y}_a \\ &\quad -\frac{1}{2}\sum_a\left(1-R_a^2\right)(\hat{y}_a)^2 - \frac{1}{2}\sum_{\langle a,b\rangle}\hat{y}_a\hat{y}_b\left(Q^{ab}-R^aR^b\right) - iy^{*\mu}\sum_a\hat{y}_aR^a\Big)\end{aligned} \tag{39}$$

Now, in the limit $d\to\infty$, Eq. 38 can be approximated using the saddle-point approach (Bender & Orszag, 2013),

$$\begin{aligned}\langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} \approx \textbf{extr}_{Q_0,\hat{Q}_{0a},Q^{ab},\hat{Q}^{ab},R^a,\hat{R}^a}\exp\Big(P\Big[\frac{i}{2}\sum_a Q_0\hat{Q}_{0a} + i\sum_{a<b}Q^{ab}\hat{Q}^{ab} \\ + i\sum_a R^a\hat{R}^a + G_S(\hat{Q}_{0a},\hat{Q}^{ab},\hat{R}^a) + \alpha G_E(Q^{ab},R^a)\Big]\Big)\end{aligned} \tag{40}$$

where, **extr** corresponds to extremization of $\langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}}$ over the respective order parameters. Performing this extremization over $\hat{Q}_{0a}$, $\hat{Q}^{ab}$ and $\hat{R}^a$, then generates an expression of the form,

$$\begin{aligned}\langle\langle Z^n\rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} = \textbf{extr}_{Q_0,Q,R}\exp\Bigg\{nN\Bigg(\frac{1}{2}\frac{Q-R^2}{Q_0-Q} + \frac{1}{2}\ln(Q_0-Q) - \frac{\alpha}{2}\ln\left[1+\beta(Q_0-Q)\right] \\ - \frac{\alpha\beta}{2}\frac{1-2R+Q}{1+\beta(Q_0-Q)}\Bigg)\Bigg\}\end{aligned} \tag{41}$$

where we have invoked *replica symmetry* in the form, $Q^{ab} = Q$ and $R^a = R$, and that $\mathcal{E}_{\mathcal{T}} = (y^*-y)^2/2$. Plugging this back into Eq. 28, then finally yields,

$$\begin{aligned}\beta f = -\textbf{extr}_{Q_0,Q,R}\Bigg\{\frac{1}{2}\frac{Q-R^2}{Q_0-Q} + \frac{1}{2}\ln(Q_0-Q) - \frac{\alpha}{2}\ln\left[1+\beta(Q_0-Q)\right] \\ - \frac{\alpha\beta}{2}\frac{1-2R+Q}{1+\beta(Q_0-Q)}\Bigg\}\end{aligned} \tag{42}$$

The remaining pair of order parameters generate the following set of transcendental equations on extremization (Bös, 1998):

$$\begin{aligned}R &= \frac{\alpha}{a} \\ Q &= \frac{\alpha}{a^2-\alpha}\left(1-\frac{2-a}{a}\alpha\right) \\ Q_0 &= Q + \frac{1}{\beta\,(a-1)}\end{aligned} \tag{43}$$

where, $a = \max[1,\alpha]$ for $T\to 0$.

Now, the above determined values of R, Q and Q_0 can be perceived as the *maximally likely* values of R, Q and Q_0 of our teacher-student setup, for an inverse temperature β parameterizing the system.

Generalization Error: The self-averaged generalization error can be obtained by plugging in the expressions for $R^{(*)}$ and $Q^{(*)}$ from above, into Eq. 23. This yields,

$$\langle\langle \mathcal{L}_{\mathcal{G}} \rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} = \frac{1}{2\left(a^2-\alpha\right)}\left[a^2-(2a-\alpha)\alpha\right] + \frac{1}{2\beta(a-1)} \tag{44}$$

Training Error: The self-averaged training error, on the other hand is defined to be,

$$\begin{aligned}\langle\langle \mathcal{L}_{\mathcal{T}} \rangle\rangle_{\boldsymbol{x},\boldsymbol{W}} &= \frac{1}{\alpha}\frac{\partial\left(\beta f\right)}{\partial\beta} \\ &= \frac{1}{2\left(a^2-\alpha\right)}\left[a^2-(2a-\alpha)\alpha\right]\left(\frac{a-1}{a}\right)^2 + \frac{1}{2\beta a}\end{aligned} \tag{45}$$

B Ridge regularization and Free-Energy

The addition of a ridge or L_2 regularization term to the square-loss of our teacher student setup coupled with the choice of the regularization matrix $\boldsymbol{F}$ as detailed in 1, then yields,

$$\mathcal{L}_{\mathcal{T}} = \frac{1}{2N}\left[\sum_{i=1}^{N}\left(\boldsymbol{W}^T\boldsymbol{x}_i - \boldsymbol{W}^{*T}\boldsymbol{x}_i^*\right)^2 + \lambda||\boldsymbol{W}||_2^2\right] \tag{46}$$

causes a modification of our free-energy expression of Eq. 42 to yield,

$$R(t,\lambda) = R_1 + R_2, \quad \text{where,} \quad R_1 := \frac{n}{a_1 d}, \quad \text{and,} \quad R_2 := \frac{n}{a_2 d} \tag{47}$$

$$Q(t,\lambda) = Q_1 + Q_2, \quad \text{where,} \quad Q_1 := \frac{b_1 b_2 c_2 + b_1 c_1}{1-b_1 b_2}, \quad \text{and,} \quad Q_1 := \frac{b_1 b_2 c_1 + b_2 c_2}{1-b_1 b_2}. \tag{48}$$

for a_i, b_i and c_i $(i = 1, 2)$ defined in the main text.

C Extended experiments

Figure 4 presents the analytical generalization dynamics for two values of $kappa$ and provides comparison between the theory and simulation results of the same model. We observe that the theory and simulations accurately match. Further experiments are provided in the following anonymous Collab notebook.

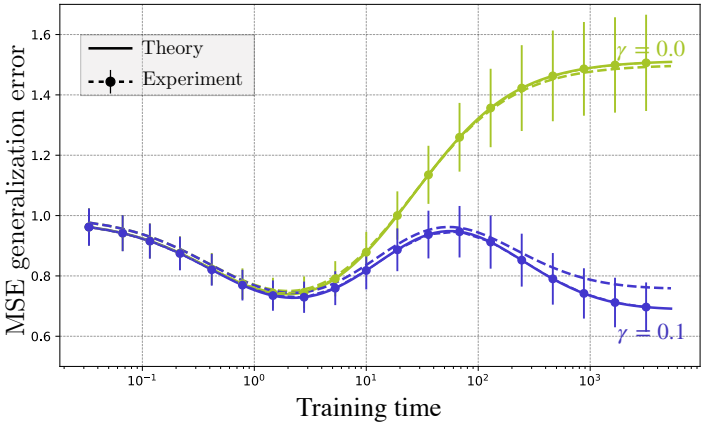

Figure 4: The teacher-student set-up in Sec. equation 2.2. We compare the analytical solutions to simulations performed on our teacher-student setup with $d = 100$, $p = 50$, $n = 150$ and we plot the error bars over 100 random seeds. The solutions and the simulations match closely and we observe double descent over the generalization error.