A Solvable Model of Neural Scaling Laws

Alexander Maloney,[a][*] Daniel A. Roberts,[bc][*] and James Sully[de][*]

[a] *Department of Physics, McGill University,*
Montréal, Quebec H3A 2T8, Canada

[b] *Center for Theoretical Physics and*
Department of Physics, Massachusetts Institute of Technology
Cambridge, Massachusetts 02139, USA

[c] *Salesforce, Cambridge, Massachusetts 02139, USA*

[d] *Department of Physics and Astronomy, University of British Columbia,*
Vancouver, BC V6T 1Z1, Canada

[e] *Anthropic, San Francisco, California 94960, USA*

`alex.maloney@mcgill.ca, drob@mit.edu, jsully@anthropic.com`

Abstract

Large language models with a huge number of parameters, when trained on near internet-sized number of tokens, have been empirically shown to obey *neural scaling laws*: specifically, their performance behaves predictably as a power law in either parameters or dataset size until bottlenecked by the other resource. To understand this better, we first identify the necessary properties allowing such scaling laws to arise and then propose a statistical model – a joint generative data model and random feature model – that captures this neural scaling phenomenology. By solving this model in the dual limit of large training set size and large number of parameters, we gain insight into *(i)* the statistical structure of datasets and tasks that lead to scaling laws, *(ii)* the way nonlinear feature maps, such as those provided by neural networks, enable scaling laws when trained on these datasets, *(iii)* the optimality of the *equiparameterization* scaling of training sets and parameters, and *(iv)* whether such scaling laws can break down and how they behave when they do. Key findings are the manner in which the power laws that occur in the statistics of natural datasets are extended by *nonlinear* random feature maps and then translated into power-law scalings of the test loss and how the finite extent of the data's spectral power law causes the model's performance to plateau.

---

[*] Equal contribution.

Contents

1   Introduction

Large language models (LLMs) such as GPT-3 [1], LaMDA [2], and Palm [3] have made fantastic advances in the generation of language, so much so that they can convincingly write text that fools humans into thinking it's written by other humans. Built from the transformer architecture [4], these and similar dense LLMs [5–8] are "large" as in *size*, with Palm topping out at 540 billion parameters, and also "large" as in (big) *data*, with Chinchilla [8] trained on 1.4 trillion tokens. These regime that these models operate in – jointly large parameter and large data – differs from both the regime covered by classical statistical approaches to machine learning (see, e.g., [9]) – typically an *underparameterized*

setting of large datasets and a fixed number of parameters and characterized by a bias-variance tradeoff – and the regime typically studied by modern theoretical approaches to deep learning [10–16] – an *overparameterized* setting of fixed datasets and a large number of parameters and characterized by *interpolation* [17] in which models memorize their training sets.

Inspired by the performance gains of the successive scaling up of LLMs, Ref. [18] comprehensively studied the test loss of such autoregressive transformer models trained on language model tasks across a large variety of model and dataset sizes. Impressively, they found that the overall performance can behave as a *power law* in any of parameters, dataset size, and compute, so long as the model isn't bottlenecked by any of the other two. (See, e.g., Fig. 1.) Moreover, by mapping the bottleneck and then jointly scaling parameters, data, and compute, practitioners can learn how to most efficiently apply their finite resources towards engineering bigger models, gathering more data, or burning their FLOPS. Thus, given the breadth of this empirical investigation over a number of orders of magnitude, the existence of these *neural scaling laws*, as they've been dubbed, have led many to believe a *scaling hypothesis* [19]: performance on language modeling tasks can be made *predictably* good simply by taking current transformer models and continuing to scale up parameters, data, and compute.

After such a study, a number of follow ups appeared showing even more general applicability and more detailed understanding [20–23] – and even improved performance scaling with data size from a power law to an exponential falloff with clever pruning [24].[1] At the same time, autoregressive generative modeling with transformers has continued to be applied to broader AI tasks such as coding [27], quantitative reasoning [28], and even on the suite of computer vision tasks, with the advent of the Vision Transformer (ViT) [29] family of models.

Given the ever growing breadth of tasks that these models can accomplish [30] and given their continuing gains in performance as we engineer ever bigger models and scrape ever bigger datasets, it is increasingly important to understand the origin of these neural scaling laws. The set of important questions include:

- What are the properties of datasets and tasks that lead to scaling laws?

- Which classes of models support scaling laws when trained on these datasets?

- How do scaling laws arise, or what mechanism leads to such predictive behavior?

- Can this predictive behavior break down, and what happens in such regimes?

Addressing these questions can not only help us improve our AI systems practically, but also help us better understand the structure of AI tasks – such as language modeling –

---

[1]Earlier works with similar ideas include [25], which predicts the test loss for different deep learning scenarios and identifies a power law scaling with training set size, and [26], which models the scaling of performance with both data and model size.

that seem to require gigantic amounts of data to reach (approximate) human-level performance.[2]

In this paper, we will provide some initial answers to these questions by jointly studying a *generative data model* and *random feature model* that together exhibit neural scaling laws analogous to those found in [18]. This provides a theoretical framework for deriving the observed phenomenology in a class of large-parameter and big-data models, just like the "microscopic" framework of *statistical mechanics* can be used to derive the "macroscopic" laws of *thermodynamics* in physics. We will explain how our model captures the essential statistical aspects of natural datasets and of the feature representation of non-linear networks, and then we will systematically solve this joint model to compute its test loss as a function of *both* dataset size and number of parameters.[3] We will thus show how our model matches the empirically observed behavior of LLMs, and then we will use this setting to better understand how scaling laws arise and break down.

One of our main results is a lack of universality of scaling laws across differently structured data generation processes: datasets that lead to scaling laws have a particular power-law structure in their spectral statistics, which ultimately leads to a power-law scaling of the test loss when there are no resource bottlenecks present. Moreover, we find that an essential role of nonlinear feature maps is extending the power law in the spectrum of the representation as a function of the number of features. This ability to extend the power law differentiates the performance of different deep neural network (DNN) models and, although we don't investigate it here, is presumably an important reason why – from the perspective of this analysis – transformers enable neural scaling law phenomenology. Finally, for generalized linear models – i.e., linear regressions of potentially nonlinear feature maps – we learn that exact *equiparameterization* – scaling the number of features identically with the size of the training set – is optimal when some kind of regularization is applied.[4] Intuitively, for the sort of data that leads to scaling laws, each additional sample can be used to learn about an additional feature in the latent feature space, and the model should have an additional parameter in order to represent the information from this new latent feature.

An important insight that emerges from our analysis is the role of a new scale that determines when the empirical behavior found by [18] breaks down. This scale can be understood as the size of the *latent space* from which the data is generated and must be much larger than *both* the size of the training set and the number of parameters of the

---

[2]How do we understand the contrast between Chinchilla [8], trained on 1.4 trillion tokens, and a human, for which the size of the training set is perhaps only of order 10 million words [31, 32]? By another estimate, LLMs may receive 1000x the linguistic data that a typical ten-year-old child might have received [33].

[3]As our analysis makes use of an exact optimization solution, it's effectively in the regime of "infinite" compute: thus, our framework can only teach us about tradeoffs between data and parameter resources.

[4]This is consistent with the finding of [8], though is slightly counter to the initial empirical results in [18]. However, both of those references concern empirical investigations of LLMs, while our analysis concerns generalized linear models and may not apply in the same way for nonlinear models that learn representations. (See §5 under the subheading *Representation Learning?* for further discussion.)

model in order to observe the power-law scaling and bottleneck behavior of Ref. [18].[5] If either of these two resource scales exceed the size of the latent space, our analysis shows a new regime of different behaviors for the test loss that has not yet been seen in the LLM experiments. Since we have a generative model of the data we control this scale directly in our analysis, but it would be extremely interesting to understand this scale in natural data, such as images or text.

As the main set of tools we use to solve our joint data and feature model is random matrix theory (RMT), one of the technical contributions in this paper is the development of a diagrammatic approach – borrowed from theoretical physics – for computing random matrix expectations in machine learning. Our techniques are particularly well suited to the regime of jointly large dataset and number of parameters, in which a restriction to diagrams with a *planar* topology captures the dominant contributions to RMT expectations. While such techniques will be familiar to physicists, we were able to apply them here to vastly simplify a number of previous RMT derivations in machine learning; for example, we avoid the lengthy concentration-of-measure proofs of [34], and we also sidestep the need to use replica methods as in [35, 36].

While most of the previous work on scaling laws has been empirical, there have been a few papers that have sought theoretical explanations or models of neural scaling laws [37–40], and a few others that have found a power-law scaling of the test loss with dataset size [41, 35]. The most directly relevant of these, [38], considers a student-teacher model with the right phenomenology though studies it only in certain limits: in particular, the authors assume a power law of *infinite* extent for the data and then import a result from [36] to find a power-law scaling of the test loss; our work leaves the relative ratios of the size of the latent space, the size of the training set, and the number of features all finite. Additionally, the calculation of [36] gives the averaged test loss of a non-generalized linear regression without any random feature map to extend the power law. Our work is more closely related to [34], which was not focused on scaling laws but finds a general expression for the test loss averaged over nonlinear random feature maps with fixed training data and labels; we extend this further to student-teacher models, finding a very simple expression for the test error and then are able to average it over our random data model as well. In the process, we also find much simpler derivations of the central results in [34] and [36] using our diagrammatic techniques.

$$\blacklozenge\blacklozenge\blacklozenge\blacklozenge$$

The plan of this paper is as follows:

In §2, we provide a non-technical overview of the data and feature map settings that lead to neural scaling laws: to begin, we review the phenomenology of the joint parameter and dataset test loss behavior discovered by [18]; then, in §2.1 and §2.2, we use examples from natural datasets to explain the specific spectral properties of the input data and

---

[5]This is perhaps surprising given a general expectation that natural data should live on manifold of smaller intrinsic dimension than its embedding dimension; see §4.3 for further discussion.

4

nonlinear feature maps, respectively, needed to have both power law scaling and bottleneck behavior.

In §3, we present (§3.1) and then solve (§3.2) our statistical model of neural scaling laws, consisting jointly of a generative model for the data and random feature map, deriving a formula for the test loss as a function of the size of the dataset and number of features of the model and then showing that it precisely matches experiment. We also outline (§3.3) how we could use our same RMT tools to model spectral power-law extension in nonlinear feature maps and comment (§3.4) on the relationship of our work to other RMT results in the machine learning literature.

In §4, we interpret our calculations from the previous section and expand on our results. Most importantly, in §4.1, we characterize the *breakdown* of neural scaling law behavior in our model by considering our result from §3 in the limit where the size of the latent space becomes smaller than either the size of the training set or the number of features in the model. We also confirm the validity of our calculation in this limit by comparing against numerical simulations in the same regime. Then, in §4.2 we explain the optimality of the *equiparameterized* regime for neural scaling, contrasting with the overparameterized regime and discussing the double descent phenomenon, while in §4.3 we further consider our new scale that controls the size of the latent space and the breakdown of scaling laws in the context of traditional notions of dimensionality reduction. We close in §4.4 by discussing some limitations of our minimal power-law spectral data model that could be improved in future analyses.

Finally, in §5 we conclude and give an outlook towards a future research direction. In particular, we provide a guide on how one could use the tools from [16] to move beyond our random-feature linear regression and incorporate the type of representation learning present in nonlinear models, such as those used in realistic deep learning scenarios.

To make the paper tractable, a few additional analyses and technical details have been consigned to appendices. In Appendix A, we present and solve a progression of simpler data and feature models with increasing complexity: first, (§A.1) we show that the simplest possible model, a linear model where the input data has independent and identically distributed Gaussian components – i.e. data with Marchenko-Pastur spectral statistics – does not have scaling laws; then, (§A.2) we explain why linear regression on data sampled from a more realistic data model – but without any feature mapping – also does not exhibit the right behavior. Finally, in Appendix B we explain how to find analytical formulae for the trace of the resolvent with the covariance matrix, the quantity that controls the test loss of our model.

2   Prerequisites for Neural Scaling

An exciting empirical observation of [18] was that the test loss for large-scale transformer models [4] can be *predicted* by fitting by a **phenomenological model** of an extremely

simple form:

$$\mathcal{L}(N, T) = \left[\left(\frac{N_c}{N} \right)^{\frac{\alpha_N}{\alpha_T}} + \frac{T_c}{T} \right]^{\alpha_T}. \tag{1}$$

Here, N is the number of (non-embedding) parameters characterizing the size of the model, T is the number of datapoints in the training set characterizing how many examples the model can learn from, and N_c, T_c, α_N, and α_T are all fit constants. A cartoon plot of (1) as a function of the training set size, T, and for a variety of different model sizes, N, is shown in Fig. 1.
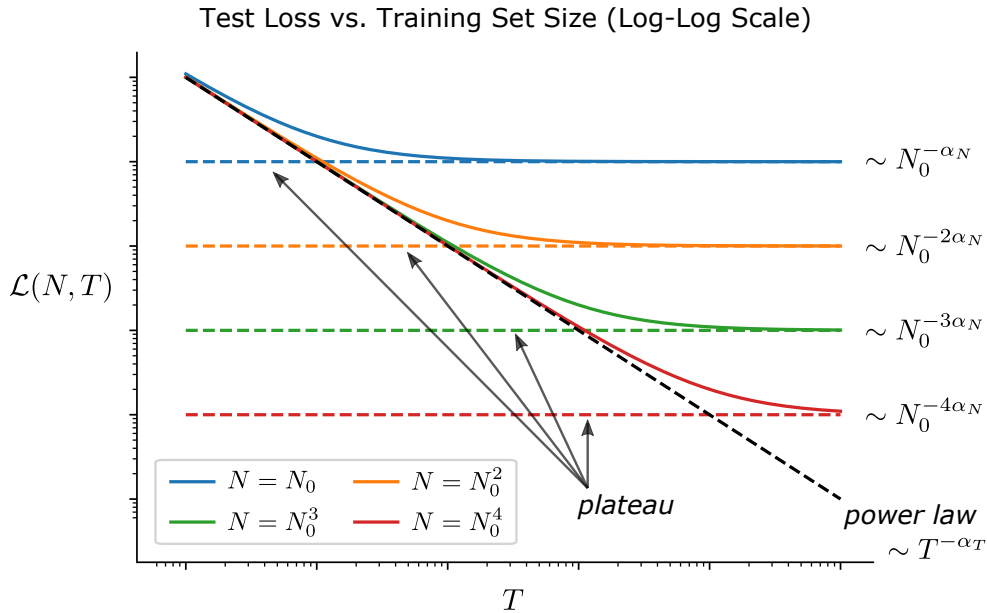

Figure 1: Cartoon plot of the empirical scaling laws discovered by Ref. [18] demonstrating that the test loss of LLMs trained with early stopping are predictably described by a simple phenomenological model, (1), plotted as a function of dataset size, T, for different model sizes, $N = \{N_0, N_0^2, N_0^3, N_0^4\}$: if the model isn't bottlenecked by the number of parameters ($N \to \infty$), the test loss behaves as a *power law* in the training set size, $\mathcal{L}(N, T) \sim T^{-\alpha_T}$; otherwise, if the number of parameters is too small for a given training set, then the test loss stalls at a *plateau* at a value that depends predictably on the parameters, $\mathcal{L}(N, T) \sim N^{-\alpha_N}$. Similar statements hold reversing the role of the training set and parameter resources, and scaling both training set and parameters jointly with relative ratio $N \sim T^{\alpha_T/\alpha_N}$ ensures the overall best performance.

This formula is quite interesting for a number of reasons:

(i) On the one hand, taking the training set, T, or the model size, N, to be large the

model improves as a power law in the scaled parameter:

$$\mathcal{L}(N) \equiv \mathcal{L}(N, \infty) \sim N^{-\alpha_N} \,, \tag{2}$$

$$\mathcal{L}(T) \equiv \mathcal{L}(\infty, T) \sim T^{-\alpha_T} \,, \tag{3}$$

where we now see that α_N is the **scaling exponent** characterizing the behavior of the loss as the number of parameters is increased, and α_T is the scaling exponent characterizing the its behavior as the size of the training set is increased. The first **scaling law**, (2), is particularly interesting in practice as often LLMs are effectively trained in infinite data regime; in this case, we achieve predictable and continuing performance gains for increases in the size of our models.[6] If there is truly infinite data, this means that for tasks and models that have scaling law (2), we can become arbitrarily good at those tasks simply by engineering bigger and bigger models. More realistically, we get a scaling law like (2) so long as the size of the model is much smaller than training set size, $N \ll T$, and we get a scaling law like (3) so long as the size of the training set is much smaller than the size of the model, $T \ll N$.[7]

(ii) On the other hand, once the scaled parameter exceeds the fixed parameter, the test loss asymptotes to a **plateau** that depends on the fixed parameter. For instance, studying the test loss as a function of the parameters for $N \gg T$ just gives a constant,

$$\mathcal{L}_{\text{plateau}}(N) \equiv \lim_{N \gg T} \mathcal{L}(N, T) = \left(\frac{T_c}{T} \right)^{\alpha_T} \,, \tag{4}$$

while analogously studying the loss as a function of the training set size for $T \gg N$ gives a similar constant,

$$\mathcal{L}_{\text{plateau}}(T) \equiv \lim_{T \gg N} \mathcal{L}(N, T) = \left(\frac{N_c}{N} \right)^{\alpha_N} \,. \tag{5}$$

This means that the model performance can be inhibited by a **bottleneck** when either the size of the training set or the size of the model is limited. This has practical consequences as well: when the performance is bottlenecked by the training data, no matter how good our engineering talent is in training larger and larger models, the loss will not be able to improve further. Of course, in this case, if instead of building a larger model we collect more training data and reinterpret (4) as a function of the training set size, e.g. as (3), the loss will again improve as a power law, though it will be in the training set size, T.

---

[6]The Chinchilla model and related investigation [8] suggest that if we train long enough, data might actually be a bottleneck, or soon will be in the future.

[7]More precisely, the relative scaling to get a power law (2) should be stated as $N \ll T^{\alpha_T/\alpha_N}$, while the relative scaling for power law (3) should be stated as $T \ll N^{\alpha_N/\alpha_T}$; analogously, we should more precisely have the reverse of these relations for accessing the plateau regimes in (4) and (5).

(iii) On our other other hand, we can instead interpret the test loss scaling with either resource as a power law when not bottlenecked by the other in the following way: if we jointly scale both the number of parameters and the size of the training set in a particular way, then we can always achieve power-law gains in our performance and avoid the plateau behavior. For instance, if we parameterize the size of our model in terms of the amount of training data we collected and make a power law ansatz, $N(T) \sim T^p$, we can then make both terms inside the square brackets of (1) contribute equally, ensuring the loss overall decreases as a power law, by scaling our model size as

$$N(T) = N_c \left(\frac{T}{T_c} \right)^{\frac{\alpha_T}{\alpha_N}} ,\qquad(6)$$

where the power $p \equiv \alpha_T/\alpha_N$ controls how the size of the model scales as we grow the training set.

In principle this phenomenological model of the test loss, (1), predicts that these LLMs – and any related models that can be fit by this equation – can become arbitrarily proficient at their underlying tasks so long as we continue to jointly scale both training data and model size as (6).

Given that (1) is an empirical observation over some range of training set sizes and models sizes, and for a particular set of AI tasks and deep learning architectures, it's natural to wonder how general it actually is, and whether the behavior will continue for especially large scales N and T. In fact, there are a number of details leading to the behavior (1) that are so far implicit in this discussion and should be made explicit. For instance, the fit of the exponents, and the relative scaling, (6), can depend on the details of the learning algorithm;[8] and to find the fit (1), the authors of [18] needed to regularize their models, e.g., by using early stopping. Most importantly, in [18] the authors trained on a specific natural dataset: an extended WebText dataset built from human language [42]. Thus, by identifying the mechanism that leads to such scaling laws, we will see that they arise for far more general dataset–model combinations.

In the rest of this section we will discover the properties of the data (§2.1) and model (§2.2) that must go into a minimal modeling scenario that contains a version of (1), including both the *scaling law* limit and the *plateau* limit. In particular, we'll explain how the data distribution must have special statistical properties, which we'll identify in natural datasets, and how a machine learning model must transform those statistical properties in a special manner, which we'll see is generically present in nonlinear DNNs.

---

[8]In the original scaling laws paper, [18], the power-law exponents for the training set and parameters were measured to be different, with their ratio positive, $\alpha_T/\alpha_N > 1$; in contrast, Ref. [8] found equal exponents, $\alpha_T = \alpha_N$, by training their models longer and on more training data.

2.1 Data Properties

AI tasks in different domains have very different underlying data – the input token features of textual data used to train LLMs for natural language processing (NLP) is a priori very different than the input pixel features of image data used for computer vision (CV) applications – yet, as we will see, both domains can exhibit the full neural scaling law phenomenology of (1). However, purely random data without any structure does not.[9] Thus, to understand such behavior, we should try to identify the structure universal to these natural datasets that exhibit scaling-law behavior.

Consider a generic raw data point, x_α, in a dataset of T samples, for $\alpha = \{1, \dots, T\}$. For our data model, we will think of each x_α as being sampled independently from some distribution $p(x)$ and refer to a particular one as a **sample**. We will denote the individual components of a sample as

$$x_{i;\alpha}, \quad \text{with} \quad i = 1, \dots, N_{\text{in}}, \tag{7}$$

where i indexes the N_{in} different **input features**, and which is supposed to represent, e.g., a particular pixel or token.[10] Accordingly, the statistics of the dataset endow the data with the structure that allows for the power law and plateau in the test loss.

Intuitively, the correlations between the different input features, $x_{i;\alpha}$, should characterize the dataset. For instance, if the $x_{i;\alpha}$ are pixels of an image, we may expect that different pixels will vary similarly across images that are similar. In contrast, the mean value of an input feature is uninformative, and so we will assume our data is centered in a preprocessing stage. Thus, a object of interest for us will be the **empirical feature-feature covariance matrix** of the dataset:

$$\frac{1}{T} \sum_{\alpha=1}^{T} x_{i;\alpha} x_{j;\alpha}. \tag{8}$$

As this covariance matrix will typically contain nonzero off-diagonal components, instead it will be simpler to consider its eigenvalues. Let us denote a particular eigenvalue as λ_i, and we'll refer to all of the eigenvalues collectively as the **spectrum** of the data. Note that since the covariance, (8), is a N_{in}-by-N_{in}-dimensional matrix, there are N_{in} eigenvalues in the spectrum. Moreover, if the size of the dataset is smaller than the number of input features, $T < N_{\text{in}}$, then the rank of the covariance is at most T, and at least $N_{\text{in}} - T$ of those eigenvalues will be zero.

The spectrum is a simple summary quantity that we can use to characterize a dataset. To see why, let's look at some spectra from real natural datasets in different domains: in Fig. 2 we plot example spectra for different dataset sizes, T, from a computer vision image

---

[9]See Appendix A.1 where we analyze data with Marchenko-Pastur statistics.

[10]Technically for NLP, we want to first pass our input tokens through a fixed embedding so that x_i represents a component of the embedded token.

dataset (left panel) and a tokenized and embedded natural language dataset (right panel). Right away, we see that both these representative natural datasets have interesting and in fact, very special structure to their spectra:

(1) For a few orders of magnitude beginning from around $i \approx 10$, the spectra are well fit by a power law,

$$\lambda_i \sim \frac{1}{i^{1+\alpha}} \, , \tag{9}$$

where here we've decomposed the exponent of decay as $1 + \alpha$ with some foresight.[11]

(2) For each fixed size T, the power law terminates in very rapid decline in the value of the eigenvalues, $\lambda_i \to 0$, as the index approaches the size of the dataset, $i \to T$, for datasets smaller than the number of input features, $T < N_{\text{in}}$, or as the index approaches the number of input features, $i \to N_{\text{in}}$, for larger datasets, $T > N_{\text{in}}$. This characterizes the **tail** of the spectrum.[12]

(3) Varying the number of samples, T, we also vary the *extent* of the power-law fit, (9).

As we will explain, these properties will translate to the power law scaling and plateau features of the test loss (1).[13]

To give more intuition for the spectral property *(1)*, let's compare this situation to what we might have naively expected. A familiar data analysis setting in which we analyze spectra is *principal component analysis* (PCA) [48]: PCA is a dimensionality-reduction tool in which the covariance matrix of a dataset is diagonalized to find linear combinations of the input features, x_i, that account for the majority of the variance of the data. Typically when PCA is useful the spectrum has a **gap**, an index, $i = M$, for $M \ll N_{\text{in}}$, such that these few M large eigenvalues account for the majority of the total variance. In that case, projecting the data onto the subspace spanned by the top M eigenvectors is a way of reducing the naive N_{in}-dimensional *input feature* space to a much smaller M-dimensional **latent feature** space. When this works, we might think of the bulk of the spectrum, $\lambda_{M+1}, \ldots, \lambda_{N_{\text{in}}}$, as uninformative noise, and that the true *generative* process describing the distribution $p(x)$ lives on this smaller dimensional latent feature space. However, the spectra of our natural datasets in Fig. 2 essentially have a *continuous* spectra without any gap: this implies that the data was generated from a space without any natural cutoff for

---

[11] This point was also emphasized by [38]. Note that when the test loss has scaling law phenomenology, the exponent α should be a positive real number, $0 < \alpha < \infty$.

[12] To see the tail for larger datasets, $T > N_{\text{in}}$, see Fig. 3.

[13] Moreover, whether the underlying process that generates the spectrum actually comes from a power law distribution – versus, e.g., a log-normal distribution – doesn't matter; we actually only need for the spectrum to be *approximately* described by a power law, as in *(1)*, *(2)*, *(3)*, in order for the scaling law phenomenology of the test loss to arise. (For further discussion of processes that give power law vs. log-normal distributions, see [43], and to learn more about the difficulty of identifying true power laws in nature, see [44].)
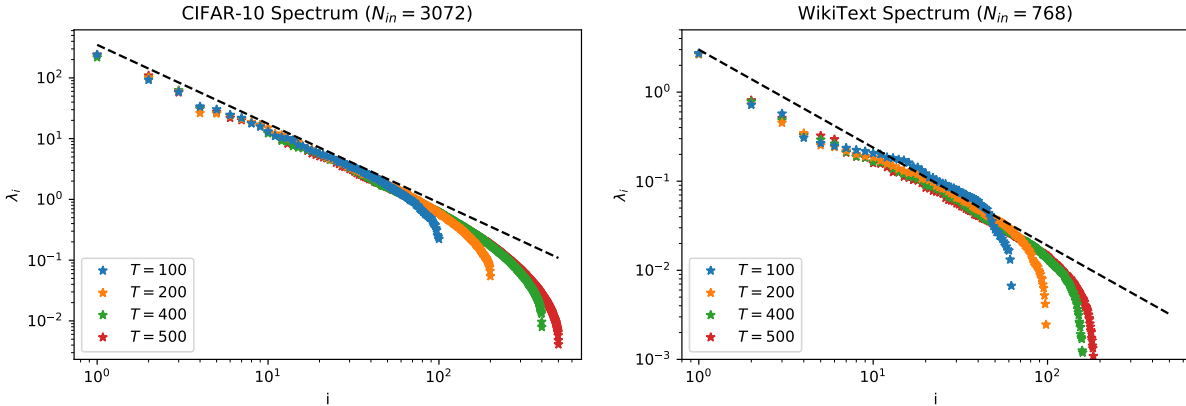
Figure 2: Log-log plot of example spectra for different dataset sizes, T, from different data domains. Increasing the dataset size, T, increases the extent of the approximate power-law fit (dashed line) so long as $T < N_{\text{in}}$. **Left:** CIFAR-10 [45], a CV dataset of 32×32-pixel natural color images. The 3 color channels bring the total number of input features per image is $N_{\text{in}} = 3 \times 32 \times 32 = 3072$. **Right:** WikiText, an NLP dataset taken from the verified Good and Featured articles on Wikipedia [46]. The input data was tokenized and then embedded using Hugging Face's implementation [47] of GPT-2 [42], and the embedding we use has dimension $N_{\text{in}} = 768$.

separating uninformative and informative features.[14] As such, especially in the portion of the spectrum that can be modeled by a power law, you can always do fractionally better at capturing the variance of the data by including more eigen-features.[15]

Considering item *(2)* above, while the eigenvalues in the part of the spectrum modeled by a power law are important in capturing the variance, the *tail* of the spectrum – in which the eigenvalues rapidly approach zero – is not. As per item *(3)*, if we increase the size of the dataset up to the number of input features, $T \to N_{\text{in}}$, we can increase the *useful* portion of the spectrum that participates in the power law. This suggests a related question: if we instead fixed a dataset size T, and subsampled the input features, N_{in}, do we still get a power law and is it now limited by N_{in}? To answer this question, in Fig. 3 we plot example spectra from the same datasets as before, but this time for a fixed dataset size, T, and with the input features subsampled from the total number N_{in}; we see that increasing the number of subsampled input features *does* extend (and slightly rescale) the power law, preserving this structure in the spectrum.

Thus, we see that the inclusion of *either* additional samples *or* additional input features can be used to extend the spectral power law, which we expect might be useful given our discussion of continuous spectra and PCA above. Unfortunately, the extent of the

---

[14]See [49] for a renormalization group perspective on this lack of cutoff for continuous spectra.

[15]This point, along with an extended discussion of latent space dimensionality, is explored more in §4.3.
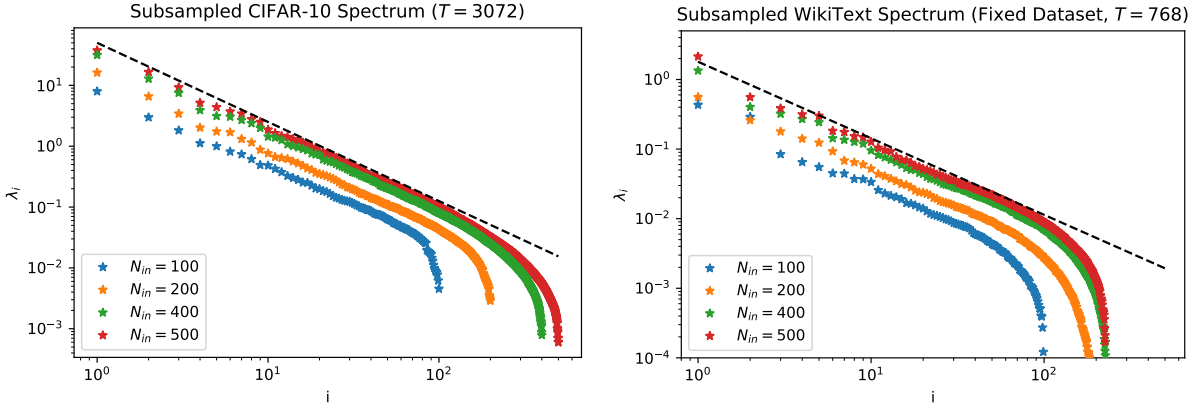
Figure 3: Example spectra from different data domains for a fixed dataset size and sub-sampled input features. (For a more detailed description of the datasets, see the caption of Fig. 2.) Increasing the number of input features in the subsample extends the length of the approximate power-law fit for the bulk (dashed line). **Left:** CIFAR-10, with pixels sub-sampled from the total 3072 input features and a dataset size of $T = 3072$. **Right:** Wiki-Text, with the components of the embedding subsampled from the total 768-dimensional embedding vector for each token and a dataset size of $T = 768$.

power law ultimately seems to be limited by the number of input features N_{in}. However, presumably if we *increased* the number of input features, for instance if we acquired high-resolution versions of our images, we'd find an even longer power law? Relatedly, CIFAR-10 contains 50,000 images in its training set despite having only 3072 input features: are those extra samples beyond the first 3072 informative?

2.2   Feature Map Properties

To answer these questions, let's try mapping the input data to a **feature space**, N, that's *larger* than the input space, N_{in}. We define a collection of feature functions as

$$\varphi_j(x), \quad \text{with} \quad j = 1, \dots, N, \tag{10}$$

where j indexes the N different **features** of the *representation* of the input x. At this point, the $\varphi_j(x)$ could be the features of a deep neural network or they could be a simpler random feature model. We are interested in studying the spectrum of this representation, which we can find by forming the *empirical feature-feature covariance matrix* of features,

$$\frac{1}{T} \sum_{\alpha=1}^{T} \varphi_{i;\alpha} \varphi_{j;\alpha}, \tag{11}$$

and then computing its eigenvalues, λ_j. In particular, we would like to understand how the spectrum of the feature representation compares to the spectrum of the input representation for different types of feature maps.

As a naive first feature map, let's pass our input dataset through a linear transformation:

$$\varphi_{j;\alpha} \equiv \sum_{k=1}^{N_{\text{in}}} u_{jk} x_{k;\alpha} \,, \tag{12}$$

where u_{jk} is a N-by-N_{in}-dimensional weight matrix, which we assume to be full rank. Concretely, we can assume each component of the weight matrix is sampled independently from a zero-mean Gaussian distribution with variance given by unity over fan-in, $1/N_{\text{in}}$. In the left panel of Fig. 4 we've plotted the spectrum of this linear feature map applied to a fixed-sized image dataset. After inspecting this figure, we remember a basic fact about linear algebra that our linear map to the larger space can only create linearly-*dependent* columns, and thus can only add zero eigenvalues to our spectrum. Thus, to meaningfully extend our spectrum, we will need to do something *nonlinear*.

As a simple example of a nonlinear feature map, let's apply a nonlinear activation function after the linear transformation (12):

$$\varphi_{j;\alpha} \equiv \sigma\left(\sum_{k=1}^{N_{\text{in}}} u_{jk} x_{k;\alpha}\right) \,, \tag{13}$$

where σ is a scalar function that acts on each individual component x_k of an input data point. We can think of this nonlinear feature map, (13), as representing the activations of a single hidden-layer neural network. As a concrete example, let's set the activation as the ReLU,

$$\sigma(z) = \begin{cases} z\,, & z > 0\,, \\ 0\,, & z \leq 0\,, \end{cases} \tag{14}$$

and again we will take each element of the weight matrix, u_{jk}, to be independent and initialized identically according to a zero-mean Gaussian with variance $2/N_{\text{in}}$. With these choices, (13) is a type of nonlinear **random feature model**. In the right panel of Fig. 4 we've plotted the spectrum of this nonlinear feature map applied to the same fixed-sized image dataset as before. Importantly, compared to the spectrum of the bare input data (blue stars), we see that increasing the number features in the feature map *extends* the portion of the spectrum that's approximately fit by a power law. In this way, we see that by applying a nonlinear transformation to our data we can build additional useful features when we have more samples than input features, $T > N_{\text{in}}$.

Together, this means that both the size of the model, N, and the size of the dataset, T, control the length of the power law in the spectrum of features: on the one hand, when the model is feature limited ($N < T$) we can increase the power-law bulk by increasing the size of the model; on the other hand, when the model is data limited ($T < N$) the
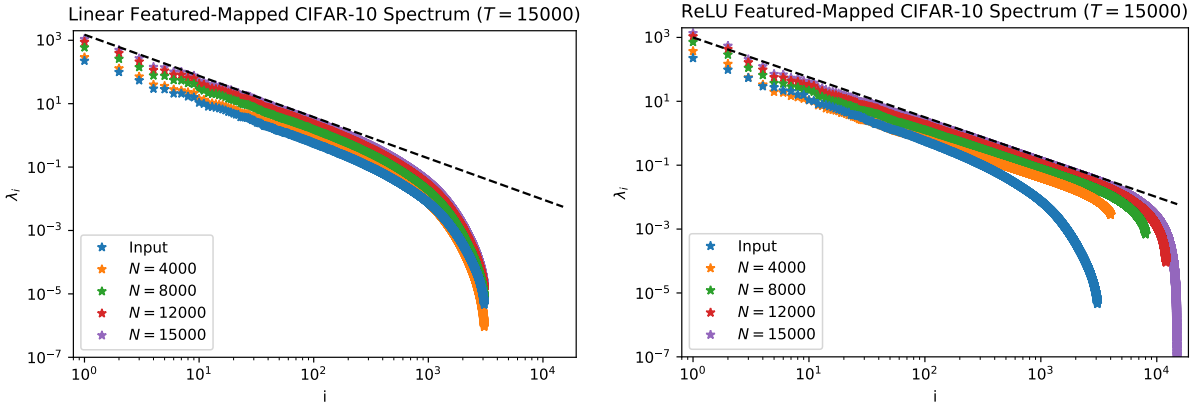
Figure 4: Spectra of the feature representation from CIFAR-10 ($N_{\text{in}} = 3072$) of a fixed dataset ($T = 15000$), with an approximate power-law fit for the bulk (dashed line). **Left:** A linear map, (12), does not extend the length of the approximate power-law fit. **Right:** For a nonlinear map, a ReLU activation applied after a linear map, (13), increasing the number of features, N, increases the extent of the approximate power-law fit. This extension is limited by the dataset size, T.

feature-feature covariance matrix, (11), is rank limited by the size of the dataset, and so the extra capacity is unnecessary as the extent of the power law is similarly limited by T. Thus, so long as they are greater than the number of input features, $N, T > N_{\text{in}}$, the minimum of these two resource scales will control how many useful features there are.[16] In the next section, we will construct a joint statistical model of datasets and feature maps that has this precise property, and by solving this model we will see how this is translated into the power-law scaling and performance plateau in the test loss of a trained model.

Aside: **Random Feature Maps vs. DNNs**

Before moving on to discuss our solvable model, let's just discuss more general nonlinear feature maps, namely deep neural networks. Even though they are both nonlinear models, a single-hidden-layer ReLU network is a very different model than the 540 billion parameter Palm based on the transformer architecture: in particular, DNNs have specially designed components, such as the multi-headed self-attention mechanism that powers transformers; moreover, they are not random feature models – at least at finite width, see, e.g., [16] – *learning* nonrandom representations of inputs; and finally, the scaling laws of [18] concern the number of *parameters*, but here we instead focused on the number of *features*. Let's address these concerns one by one in reverse order.

---

[16]Note that this extension effect is special for natural datasets with the properties enumerated in §2.1 and will not be true in general.

Firstly, there is often some ambiguity in the feature map of a DNN; e.g. with LLMs like BERT [50], practitioners sometimes use the activations of the last few layers of the model as features for a downstream task. However, the distinction between parameters is sharp even for our single-layer ReLU random feature map (13): as we are using it, this model has $N \times N_{\text{in}}$ parameters, but only N features. The resolution is that the proper way to think about the features of a network is in terms of the NTK [13, 14]: the NTK is a type of *data-data covariance matrix* of features,

$$\widehat{H}_{\alpha_1 \alpha_2} \equiv \sum_{j=1}^{N} \varphi_{j;\alpha_1} \varphi_{j;\alpha_2} , \tag{15}$$

and it's easy to see that this will have the same spectral properties as the feature-feature covariance matrix that we've been considering in (11) up to an overall rescaling. For network architectures that have an infinite-width limit [10–12], DNNs trained by gradient descent are generalized linear models with the NTK identified as the kernel. Moreover, if $z(x_\alpha; \theta)$ is the (scalar) output of the network when evaluated on a sample x_α, and θ_j is an N dimensional vector that indexes *all* the parameters, then the definition of the NTK [13, 14] tells us to identify the feature map with the derivative of the network output:

$$\varphi_{j;\alpha} \equiv \frac{dz(x_\alpha; \theta)}{d\theta_j} . \tag{16}$$

Thus, where this correspondence between DNN and linear models holds, then there's precise correspondence giving a feature for every parameter, and increasing the number of parameters increases the effective number of features accessible to the linear model.[17] To this end, if we were to plot the spectrum of the features derived from the NTK, we would see a similar phenomenology to what we observed in §2.2 for our simple nonlinear feature map (13).

Away from the infinite-width limit, at least perturbatively [16], the model output still depends on the NTK with a *parameter*-number of features, but rather than a random feature model, the features of a finite-width network learn nontrivial representations of inputs from the data. The fact that our model in the next section exhibits the neural scaling phenomenology of power law and plateau suggests that probably feature learning isn't an essential part of scaling laws; we will address more concrete means of understanding this relationship between representation learning and scaling laws in the last part of §5.

Finally, what of the broader and essential differences between a one-hidden-layer ReLU network to an LLM? A standard principle of computer science is GIGO: *garbage in, garbage out*. Perhaps the biggest takeaway lesson in our setting is PIPO: *power-law in, power-law out*. To that end, we conjecture that better DNN architectures are better able to preserve power law structure when transforming the spectra of input datasets and leave it to future

---

[17] For a more detailed discussion of this point and the correspondence, see §10.4 of [16].

work to understand how to translate these performance of better models into statements about the spectra of features.[18]

3   A Statistical Model

We want to construct a generative data model and random feature model that captures the broad empirical properties of real datasets composed with nonlinear feature maps such that the resulting statistical model's test loss exhibits the scaling law phenomenology illustrated in Fig. 1. Recall from the previous section, our key observation is an approximate power law in the spectrum of the feature representation, with the extent of the power-law portion controlled by the minimum of the number of features, N, and the size of the training set, T. After finding features with these properties in a simplified model, we can then use them in a generalized linear regression problem such that the test loss exhibits our desired behavior.

  Our goal will be to compute that averaged test loss analytically. We will accomplish this using tools from random matrix theory, using some simple diagrammatic techniques that can quickly and easily extract the properties of these models when N and T are sufficiently large. We will begin our journey in §3.1 by setting up and defining our model as well as verifying its properties with numerical simulations. The bulk of the section will be spent in §3.2, where we will explain how to average over our generative data model and random features in order to derive a formula for the model's test loss. Then, in §3.3 we outline how we could use our RMT tools to model spectral extension in nonlinear feature maps such as neural networks. Finally, in §3.4 we compare our methods and results to other related RMT machine-learning calculation.

3.1   Setup and Verifying

Return here often as you explore the other subsections in this section.

Generative Data Model

We will start by defining a generative model for the dataset.

---

[18]For instance, a more careful investigation of the raw input (blue stars) in the right panel of Fig. 4 would show that the exponent α that characterizes the spectrum in (9) actually decreases slightly after the ReLU layer. As we will explain in the next section, α ultimately will become the exponent in the power-law portion of the test loss of our model; thus, even though the power-law gets extended to give the scaling law, the slight decrease in its value ultimately leads to worse performance than if it were otherwise preserved.

  A second issue worth considering when comparing a single ReLU layer to an LLM is that we haven't modeled the eigenvectors, which may need to be considered in a more detailed model. (For a discussion of scaling laws that does consider eigenvectors, see [40].)

Rather than generating data in the raw input space, we will generate data in a **latent space**. Consider a latent data point, x, whose components are denoted

$$x_I, \quad \text{with} \quad I = 1, \dots, M, \tag{17}$$

where I indexes the M different *latent features*. To distinguish latent features from the features following a random feature map, we will use capital roman indices from the middle of the alphabet $(I, J, K, \dots)$ for the former and lower-cased roman indices $(i, j, k, \dots)$ for the latter. Importantly, to get the right behavior we will need the dimension of the latent space to be larger than any other scales in the problem:

$$M \gg N, T. \tag{18}$$

For each data point, we will sample components from a zero-mean Gaussian distribution with a covariance matrix Λ:

$$\langle x_I \rangle = 0, \qquad \langle x_I x_J \rangle = \Lambda_{IJ}, \tag{19}$$

where we denote expectations over random variables with the notation

$$\langle f(u) \rangle \equiv \int du\, p(u)\, f(u), \tag{20}$$

where u includes *all* random variables in the expression $f(u)$. If we instead want to take expectations over only some of the random variables, we will use a subscript notation on the bracket as

$$\langle f(u, v) \rangle_u \equiv \int du\, p(u|v)\, f(u, v). \tag{21}$$

When possible we will keep our derivation generic and not make assumptions about the covariance Λ, other than that it is full rank.

However, for the goal of understanding scaling laws we will be motivated to consider a class of models where the covariance spectrum has the form of a power law. In particular, we will assume the eigenvalues of Λ are well-approximated by a smooth number density of eigenvalues,

$$n(\lambda)\, d\lambda = M(\beta - 1)\lambda_-^{\beta-1}\lambda^{-\beta}\theta(\lambda - \lambda_-)\, d\lambda, \tag{22}$$

where λ_- is the minimum eigenvalue, β is an exponent that characterizes the tail of the distribution, $\theta(\lambda)$ is the Heaviside step function, and the constants are chosen such that the density integrates to M. Alternatively, we can write the spectrum as a function of index I as

$$\lambda_I = \lambda_+ \left(\frac{1}{I}\right)^{1+\alpha}. \tag{23}$$

In this form it is convenient to (hyper-)parameterize the spectrum in terms of a maximal eigenvalue,

$$\lambda_+ \equiv \lambda_- M^{1+\alpha}. \tag{24}$$

The exponent α in (23) is related to the exponent β appearing in (22) by

$$\alpha \equiv \frac{2 - \beta}{\beta - 1} \,, \tag{25}$$

as can be checked by integrating the density of states (23). Since α will ultimately be power-law exponent of the test loss, we must have $0 < \alpha < \infty$, which in turn implies $1 < \beta < 2$.[19] This is actually a rather compact range and implies that natural datasets have surprisingly heavy tails. Finally, for large M it generally does not matter whether the eigenvalues are drawn at random from a distribution of the form (22) or taken from a fixed spectrum of the form (23), and we will generally be agnostic about this choice in our theory.[20]

In Fig. 5, we numerically sample some datasets from our generative data model, (19) and (23), for different (hyper)-parameters T, M, and α to check that we have a good model of the natural datasets discussed in §2.1. We see that for $T < M$ the extent of the power law increases with the size of the dataset, and for $T \geq M$ increasing the size of the dataset sharpens the rapid decline towards zero but leaves the extent of the power law fixed. This confirms that we have captured the broad spectral properties of the inputs from the natural datasets we have analyzed.

Finally, for every latent datapoint x_I, we will also generate a C-dimensional label

$$y_i = \sum_{I=1}^{M} w_{iI} x_I, \quad \text{with} \quad i = 1, \ldots, C \,, \tag{26}$$

using a C-by-M-dimensional weight matrix, $w \equiv w_{iI}$, whose elements we will take to be independent and drawn from a zero-mean Gaussian distribution, so that

$$\langle w_{iI} \rangle = 0 \,, \qquad \langle w_{i_1 I_1} w_{i_2 I_2} \rangle = \frac{\sigma_w^2}{M} \delta_{i_1 i_2} \delta_{I_1 I_2} \,. \tag{27}$$

It is important that each label is allowed to depend on all M latent features of an input to ensure that the difficulty of the problem scales with M. Such a scaling is needed in order to approximate the self-supervised generative modeling tasks that LLMs perform.

Random Feature Model

Now let's define a random feature model that we will use to map our latent data to a feature representation. Our goal is to find a representation where the spectrum contains an approximate power-law fit that is controlled by the number of feature functions, N, in the model.

---

[19]For power-law probability distributions of the form (22), the distribution is normalizable only for $\beta > 1$ and has finite mean only for $\beta > 2$. However, if instead we fix a maximal eigenvalue λ_+, then the mean (and higher moments) will exist, but the distribution will no longer be normalizable.

[20]In our simulations, we find it convenient to use (23) and characterize the spectrum by M, λ_+, and α.
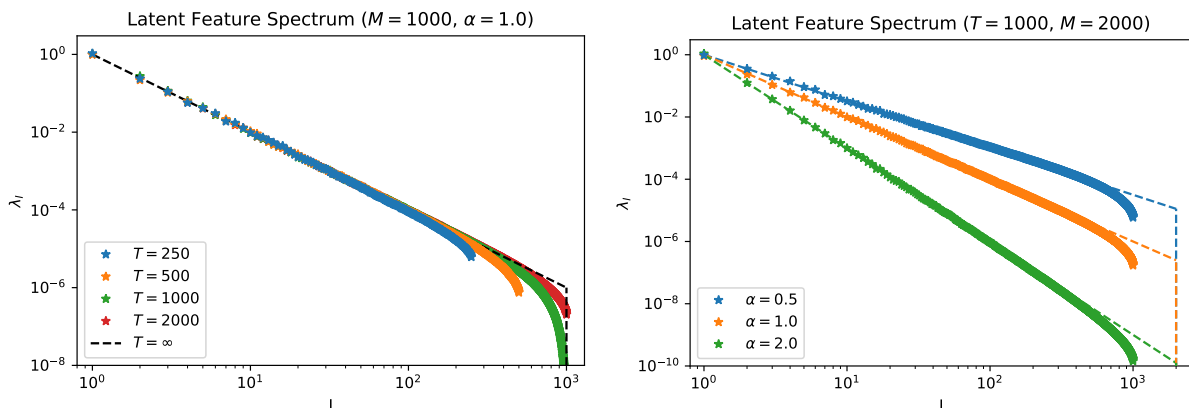
Figure 5: Spectrum λ_I from numerical simulations (stars) of our latent data generative model, (19) and (23), with the maximum eigenvalue fixed ($\lambda_+ = 1$). **Left:** The size of the dataset, T, is varied while the size of the latent space and the power-law exponent are fixed ($M = 1000$, $\alpha = 1$). These spectra follow a pattern similar to the ones displayed in Fig. 2 for natural data: for dataset size smaller than the size of the latent space ($T < M$, blue and orange) the spectrum has a bulk power law portion that terminates in a very rapid decline ($\lambda_I \to 0$) as the index approaches the size of the dataset ($I \to T$), and the extent of the power law increases with increasing dataset; for dataset size equal to and greater than the size of the latent space ($T \geq M$, green and red), the power law terminates at the size of the latent space, but the rapid decline becomes sharper and sharper as the size of the dataset increases, forming a kink in the limit of infinite data ($T \to \infty$, dashed black line). **Right:** The power-law exponent, α, is varied as the sizes of the dataset and latent space are held fixed ($T = 1000$, $M = 2000$), and the spectrum for infinite data is plotted for comparison (dashed lines). As all three simulations have the same size datasets, their power laws all terminate at the same point ($T = 1000$).

The main advantage of generating our data in a large latent space ($M > N$) rather than a smaller input space ($N_{\text{in}} < N$) is that we can use a simpler *linear* map from the larger latent space to the smaller feature space rather than having to analyze a *nonlinear* map from the smaller input space to the larger feature space. We will define our collection of feature functions by

$$\varphi_j(x) \equiv \sum_{I=1}^{M} u_{jI} x_I \, , \tag{28}$$

where j indexes the N different features of the representation of the latent input x, and $u \equiv u_{jI}$ is a $N \times M$ matrix of random feature weights drawn from a zero-mean Gaussian:

$$\langle u_{jI} \rangle = 0 \, , \qquad \langle u_{j_1 I_1} u_{j_2 I_2} \rangle = \frac{\sigma_u^2}{M} \delta_{j_1 j_2} \delta_{I_1 I_2} \, . \tag{29}$$

In Fig. 6, we take datasets sampled from our generative data model, (19) and (23),

and map them through our random feature model, (28), for a fixed set of features weights, u, in order to verify that their spectra has the properties discussed in §2.2. We see that the power-law portion of each spectrum is controlled by the minimum of the number of features, N, and the size of the dataset, T. These are precisely the properties we sought to find in our simplified joint data and feature model.
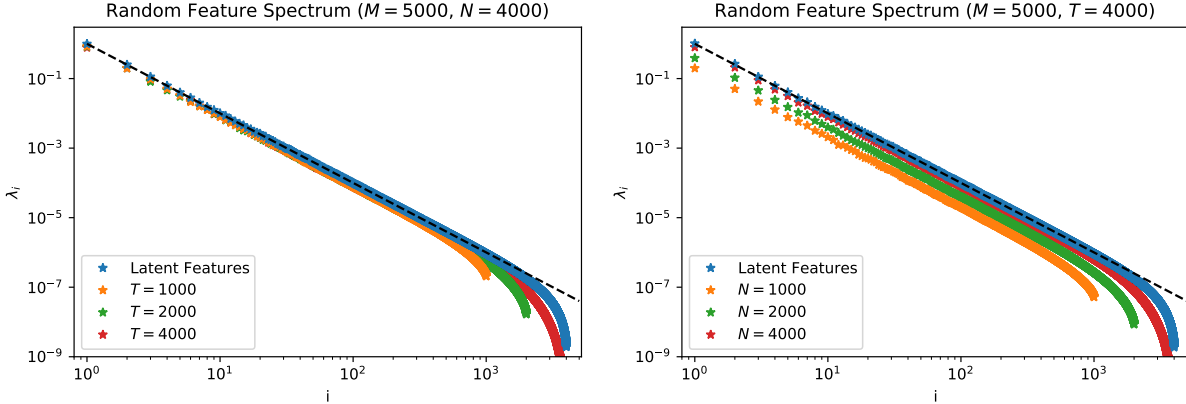
Figure 6: Spectrum, λ_i, from numerical simulations of our random feature model, (28), mapping sampled data from our generative data model, with the size of the latent space, the maximum eigenvalue, and power-law exponent fixed ($M = 5000$, $\lambda_+ = 1$, $\alpha = 1$). These spectra show that the random features of our joint model follow closely to what we observed for CIFAR-10 in Fig. 4: the approximate power-law fit is controlled by the minimum of the number of features and the size of the dataset, $\min(N, T)$. The latent feature spectrum is also plotted for comparison (blue) as is the power-law fit (dashed line). **Left:** The size of the dataset, T, is varied while the number of random features is held fixed ($N = 4000$). **Right:** The number of random features, N, is varied while the size of the dataset is held fixed ($T = 4000$).

(Generalized) Linear Regression

Now that we have features, we will "train" a (generalized) linear model to reproduce the labels, y_i, generated from the underlying latent features, (26), by learning a linear transformation of the random features $\varphi_j(x)$:

$$z_i(x; \theta) \equiv \sum_{j=1}^{N} \theta_{ij} \varphi_j(x), \tag{30}$$

where $\theta \equiv \theta_{ij}$ is a set of learnable parameters.

To fix these parameters, we will use our generative data model (19) to draw a collection of T pairs of samples $\{x_I, y_i\}$ to form our training set $\mathcal{A}$:

$$\{x, y\} \quad \Longleftrightarrow \quad \{x_{I;\alpha}, y_{i;\alpha}\} . \tag{31}$$

We will typically denote our training set of latent data and labels using the matrix notation (left-hand side), although when clarity dictates we may also use the index notation (right-hand side), with $\alpha = \{1, \ldots, T\}$ used to index particular samples in $\mathcal{A}$. Accordingly, we can use our feature functions, (28), to construct a corresponding matrix of random features derived from the training set:

$$\varphi \equiv \varphi(x) \quad \Longleftrightarrow \quad \varphi_{j;\alpha} \equiv \varphi_j(x_\alpha) . \tag{32}$$

We can then use the training set to fit the parameters by optimizing a standard MSE loss function with a ridge regression term:

$$\mathcal{L}_{\mathcal{A}}(\theta) \equiv \frac{1}{2} \sum_{i=1}^{C} \sum_{\alpha=1}^{T} (z_{i;\alpha} - y_{i;\alpha} - \epsilon_{i;\alpha})^2 + \frac{\gamma}{2} \sum_{i=1}^{C} \sum_{j=1}^{N} \theta_{ij}^2 \tag{33}$$

$$= \frac{1}{2} \|\theta\varphi - y - \epsilon\|^2 + \frac{\gamma}{2} \|\theta\|^2 ,$$

where γ is the ridge parameter. Here, we've also introduced the ability to corrupt our labels with a matrix of random noise, ϵ, which has a separate entry for each training sample and label component, and each of which entry is drawn from another zero-mean Gaussian with statistics

$$\langle \epsilon_{i;\alpha} \rangle = 0 , \qquad \langle \epsilon_{i_1;\alpha_1} \epsilon_{i_2;\alpha_2} \rangle = \sigma_\epsilon^2 \delta_{i_1 i_2} \delta_{\alpha_1 \alpha_2} . \tag{34}$$

(Note that there is no normalization factor in the variance as there was in our other variances, e.g., (27) and (29).)

Finally, optimizing the loss (33) with respect to θ has a well-known solution:

$$\theta_{ij}^\star \equiv \sum_{k=1}^{N} \sum_{\alpha_2=1}^{T} \left(\gamma \delta_{jk} + \sum_{\alpha_1=1}^{T} \varphi_{j;\alpha_1} \varphi_{k;\alpha_1} \right)^{-1} \varphi_{k;\alpha_2} (y_{i;\alpha_2} + \epsilon_{i;\alpha_2})$$

$$= (y + \epsilon)\varphi^T q , \tag{35}$$

where on the second line we switched to matrix notation and also introduced the *feature-feature resolvent* matrix

$$q(\gamma) \equiv \frac{1}{\gamma I_N + \varphi\varphi^T} \quad \Longleftrightarrow \quad q_{jk}(\gamma) \equiv \left(\gamma \delta_{jk} + \sum_{\alpha=1}^{T} \varphi_{j;\alpha} \varphi_{k;\alpha} \right)^{-1} , \tag{36}$$

where $I_N \equiv \delta_{ij}$ represents the identity matrix on feature space.

Computing Performance

To evaluate our model, we will again use our generative model (19) to draw a collection of $\widehat{T}$ pairs of samples to form our training set $\mathcal{B}$:

$$\{\widehat{x}, \widehat{y}\} \quad \Longleftrightarrow \quad \{\widehat{x}_{I;\beta}, \widehat{y}_{i;\beta}\}, \tag{37}$$

where we will generally use a *hat* to emphasize test-set quantities, and when using indices we will use $\beta = \{1, \ldots, \widehat{T}\}$ to index particular samples in $\mathcal{B}$. Accordingly, we can use our feature functions, (28), to construct a matrix of random features derived from the test set,

$$\widehat{\varphi} \equiv \varphi(\widehat{x}) \quad \Longleftrightarrow \quad \widehat{\varphi}_{j;\beta} \equiv \varphi_j(\widehat{x}_\beta), \tag{38}$$

and then use our solution, (35), for inference on these test examples:

$$\widehat{z}^\star \equiv \theta^\star \widehat{\varphi} \quad \Longleftrightarrow \quad \widehat{z}^\star_{i;\beta} = \sum_{j=1}^{N} \theta^\star_{ij} \widehat{\varphi}_{j;\beta}. \tag{39}$$

The model's performance may then be measured by a test loss

$$
\begin{aligned}
\mathcal{L}_{\mathcal{B}}(\theta^\star) &\equiv \frac{1}{2\widehat{T}} \left\| \widehat{z}^\star - \widehat{y} \right\|^2 \\
&= \frac{1}{2\widehat{T}} \left\| (y+\epsilon)\varphi^T q\widehat{\varphi} - \widehat{y} \right\|^2 \\
&= \frac{1}{2\widehat{T}} \left\| (wx+\epsilon)\varphi^T q\widehat{\varphi} - w\widehat{x} \right\|^2,
\end{aligned} \tag{40}
$$

where on the second line we substituted in for the test predictions, $\widehat{z}^\star$, using (39), and then the optimal parameters, $\theta^\star$, using (35), and on the final line we substituted in for the labels using (26). Note that this MSE loss has a different normalization than the training loss, (33), so that it represents a *per sample* loss if averaged and has a nice large-test-set limit. Furthermore, note that by using this analytical form of the linear regression solution, we are effectively in the limit of infinite training, in which the model has been allowed to converge. This means that *(a)* we will not have to worry about the way that the performance can depend on the details of the learning algorithm (see, e.g. [8]), but also that *(b)* our statistical model will not capture capture the compute-limited scaling laws studied by Ref. [18].

One advantage of a joint model of data *and* features is that we are able to numerically simulate it for different (hyper)-parameters to confirm that it has the right behavior. In Fig. 7, we plot the test loss (40) for a variety of model sizes, N, as a function of the training set size, T: first, we generate latent training and test sets by sampling from our generative data model, (19) and (23); then, we map both sets through random feature models of different sizes, (28); next, we use the linear regression solution (35) to compute test-set predictions using (39) for different values of the ridge parameter, γ, and evaluate the test

loss (40) as a function of γ; finally, we optimize the ridge parameter ($\gamma = \gamma^\star$) and plot the test loss that gives the best performance.[21] The figure illustrates the way our statistical model exhibits the same *power-scaling law* and *plateau* regions as the early-stopped LLMs studied in Ref. [18], cf. Fig. 1: in particular, our numerical simulations are predicted by a phenomenological model of an extremely simple form,

$$\mathcal{L}(N, T) = L_0 \left(\frac{1}{N} + \frac{1}{T} \right)^\alpha , \tag{41}$$

where α is the power-law exponent that parameterizes the spectrum defined in (25), and L_0 is a constant that we will compute explicitly in the following subsection but for now can be thought of as a constant to be fit. This equation, (41), follows from the original phenomenological model of the test loss, (1), by setting

$$\alpha \equiv \alpha_N = \alpha_T , \qquad L_0 \equiv N_0^\alpha = T_0^\alpha , \tag{42}$$

and seems to be the appropriate simplification for an optimally-regularized random feature model used for linear regression.

Average Goals

Having confirmed that our statistical model has the right properties, our goal now is to analytically compute the expected value of the test loss, $\langle \mathcal{L}_\mathcal{B}(\theta^\star) \rangle$, averaged according to the statistics of our generative data model and random feature maps, in order to understand the full phenomenology of neural scaling laws. This will involve: averaging over different realizations of the latent training inputs, x, and latent test inputs, $\widehat{x}$; averaging over different label weights, w, used to compute training labels, y, and test labels, $\widehat{y}$; averaging over realizations of the noise, ϵ, added to the training labels; and averaging over different feature weights u, that determine the training features, φ, and the test features, $\widehat{\varphi}$, given the latent inputs, x and $\widehat{x}$, respectively. Since these random variables are all matrices, the computation of the expected test loss is a problem in random matrix theory.

Some of these averages are very easy to perform and can be evaluated immediately:

- The test loss, (40), is quadratic in the random label noise, ϵ. Expanding in ϵ and using its statistics, (34), we find:

$$\langle \mathcal{L}_\mathcal{B}(\theta^\star) \rangle_\epsilon = \frac{1}{2\widehat{T}} \left\| w(x\varphi^T q\widehat{\varphi} - \widehat{x}) \right\|^2 + \frac{C\sigma_\epsilon^2}{2\widehat{T}} \left\| \varphi^T q\widehat{\varphi} \right\|^2 . \tag{43}$$

The first term is independent of the noise, and the second term is independent of the labels but depends on the random features $\varphi, \widehat{\varphi}$. Therefore, we will refer to these two terms as the *label term* and the *noise term*, respectively.

---

[21]For the most stable results, we use the form (35) of the linear regression solution when the model is underparameterized ($N < T$), and we use (58) when the model is overparameterized ($N > T$).
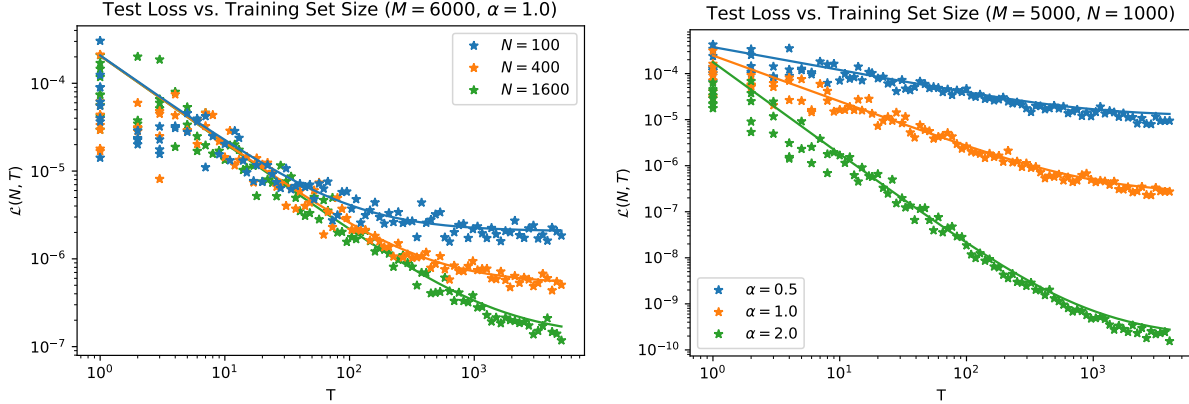
Figure 7: Test loss from numerical simulations (stars) of our optimally regularized ($\gamma = \gamma^\star$) joint statistical model ($\sigma_u^2 = 1$, $\lambda_+ = 1$, $\sigma_w^2 = 1$, $\sigma_\epsilon^2 = 0$) as a function of training set size demonstrating the same rich behavior as LLMs and a simple fit (solid lines) given by (41): analogous to Fig. 1, if the model isn't bottlenecked by the number of parameters, $N \to \infty$, the test loss behaves as a *power law* in the training set size, $\mathcal{L}(N, T) \sim T^{-\alpha}$; otherwise, if the number of parameters is too small for a given training set, then the test loss stalls at a *plateau* at a value that depends predictably on the parameters, $\mathcal{L}(N, T) \sim N^{-\alpha}$. Similar statements would hold if we plotted the test loss as a function of the number of features. (For smaller values of T, the variance of any particular realization is large, and so we've plotted multiple simulations for $T \lesssim 10$.) **Left:** The size of the training set, T, is varied for a few different sized models, N, while the size of the latent space and the power-law exponent is held fixed ($M = 6000$, $\alpha = 1.0$). **Right:** The size of the training set, T, is varied for a few different power-law exponents, α, while the size of the latent space and the size of the model is held fixed ($M = 6000$, $N = 1000$).

- The label term now involves a simple square of the label weights, w, which can be averaged over using its statistics, (27), to find:

$$\langle \mathcal{L}_\mathcal{B}(\theta^\star)\rangle_{\epsilon, w} = \frac{C\sigma_w^2}{2\widehat{T}M} \left\| x\varphi^T q\widehat{\varphi} - \widehat{x}\right\|^2 + \frac{C\sigma_\epsilon^2}{2\widehat{T}} \left\|\varphi^T q\widehat{\varphi}\right\|^2 . \tag{44}$$

Since the output dimension C just gives an overall scaling of the test loss, we will simply set $C = 1$ for the rest of the paper without loss of generality.

Two of the three remaining averages, over the latent input training set, x, and over the random feature weights, u, will be more challenging to compute.[22] They will be carried out in the remainder of the section, with some of the less conceptual and more mechanical

---

[22]The average over the latent test inputs, $\widehat{x}$, is relatively easy but we will find it convenient to defer this computation until later.

details relegated to Appendix B. Before moving on to the details, let us meditate on the mechanics of these averages.

The feature functions are defined as the product of two zero-mean Gaussian variables, x and u, cf. (28). Holding either x or u fixed and averaging over the other is straightforward given their statistics, (19) and (29). From this it follows that the features are centered,

$$\langle \varphi_{j;\alpha} \rangle = 0 \,, \qquad \langle \widehat{\varphi}_{j;\beta} \rangle = 0 \,, \tag{45}$$

but their covariances are non-trivial.

It will be convenient to decompose these covariances in terms of matrices that have either sample indices or feature indices, but not both. For instance, when averaging over the random features, using (29) we find

$$\langle \varphi_{j_1;\alpha_1} \varphi_{j_2;\alpha_2} \rangle_u = \Sigma_{\alpha_1 \alpha_2} \, \delta_{j_1 j_2} \,, \quad \langle \varphi_{j_1;\alpha} \widehat{\varphi}_{j_2;\beta} \rangle_u = \widetilde{\Sigma}_{\alpha \beta} \, \delta_{j_1 j_2} \,, \quad \langle \widehat{\varphi}_{j_1;\beta_1} \widehat{\varphi}_{j_2;\beta_2} \rangle_u = \widehat{\Sigma}_{\beta_1 \beta_2} \, \delta_{j_1 j_2} \,, \tag{46}$$

where we have defined the matrices

$$\Sigma \equiv \frac{\sigma_u^2}{M} x^T x \,, \qquad \widetilde{\Sigma} \equiv \frac{\sigma_u^2}{M} x^T \widehat{x} \,, \qquad \widehat{\Sigma} \equiv \frac{\sigma_u^2}{M} \widehat{x}^T \widehat{x} \,, \tag{47}$$

which have *sample indices* only. Note that, as they depend on x and $\widehat{x}$, these matrices are themselves random variables. Relatedly, when averaging either over the training inputs or over the test inputs, using (19) we find

$$\langle \varphi_{j_1;\alpha_1} \varphi_{j_2;\alpha_2} \rangle_x = \Omega_{j_1 j_2} \delta_{\alpha_1 \alpha_2} \,, \qquad \langle \widehat{\varphi}_{j_1;\beta_1} \widehat{\varphi}_{j_2;\alpha_2} \rangle_{\widehat{x}} = \Omega_{j_1 j_2} \delta_{\beta_1 \beta_2} \,, \tag{48}$$

where we have defined the random matrix

$$\Omega \equiv u \Lambda u^T \qquad \Longleftrightarrow \qquad \Omega_{j_1 j_2} \equiv \sum_{I_1,I_2=1}^{M} u_{j_1 I_1} u_{j_2 I_2} \Lambda_{I_1 I_2} \,, \tag{49}$$

which has *random feature indices* only and is essentially a projection of the latent-space covariance matrix, Λ, onto our model's random feature space. Lastly, again using (19) to average over training inputs or test inputs, there's a nontrivial cross-correlation between the latent inputs and the training or test features,

$$\langle x_{I;\alpha_1} \varphi_{j;\alpha_2} \rangle_x = \widetilde{\Omega}_{Ij} \, \delta_{\alpha_1 \alpha_2} \,, \qquad \langle \widehat{x}_{I;\beta_1} \widehat{\varphi}_{j;\beta_2} \rangle_{\widehat{x}} = \widetilde{\Omega}_{Ij} \, \delta_{\beta_1 \beta_2} \,, \tag{50}$$

where we have defined a final random matrix,

$$\widetilde{\Omega} \equiv \Lambda u^T \qquad \Longleftrightarrow \qquad \widetilde{\Omega}_{Ij} \equiv \sum_{J=1}^{M} u_{jJ} \Lambda_{IJ} \,. \tag{51}$$

Note that this matrix has a feature index and a latent index but does not depend on samples, and, as a partial projection, the relation

$$\Omega = u \widetilde{\Omega} \tag{52}$$

follows from the above definitions.

Finally, we note that the following training-set-averaged covariances vanish,

$$\langle x_{I;\alpha}\widehat{\varphi}_{j;\beta}\rangle_x = \langle \widehat{x}_{I;\beta}\varphi_{j;\alpha}\rangle_x = \langle \varphi_{j_1;\alpha}\widehat{\varphi}_{j_2;\beta}\rangle_x = 0\,, \tag{53}$$

and the analogous set of test-set-averaged covariances vanish,

$$\langle x_{I;\alpha}\widehat{\varphi}_{j;\beta}\rangle_{\widehat{x}} = \langle \widehat{x}_{I;\beta}\varphi_{j;\alpha}\rangle_{\widehat{x}} = \langle \varphi_{j_1;\alpha}\widehat{\varphi}_{j_2;\beta}\rangle_{\widehat{x}} = 0\,, \tag{54}$$

together indicating no cross-correlation between train and test latent or random features; this is as expected, since each sample is drawn independently. We also note that any mixed covariance between random features and latent will vanish when averaged over the random feature weights,

$$\langle x_{I;\alpha_1}\varphi_{j;\alpha_2}\rangle_u = \langle x_{I;\alpha}\widehat{\varphi}_{j;\beta}\rangle_u = \langle \widehat{x}_{I;\beta}\varphi_{j;\alpha}\rangle_u = \langle \widehat{x}_{I;\beta_1}\widehat{\varphi}_{j;\beta_2}\rangle_u = 0\,, \tag{55}$$

since these expressions are linear u.

Final Definitions

The main difficulty of our analysis is that the resolvent, $q(\gamma)$, involves inverses of the feature functions, (36). Roughly speaking, our approach to compute the averaged loss involves expanding it around $\gamma \to \infty$, thus the expanding factors of $q(\gamma)$ it contains, and then using the data-averaged covariances, (48) and (50), to evaluate the resulting infinite sum of Gaussian expectations. This leads to an implicit equation for a quantity that ultimately determines the test loss, which can be solved in certain limits as well as averaged over the random features.[23]

One wrinkle is that the above only works well in the underparameterized regime with $N < T$. To analyze the overparameterized regime, $N > T$, it will be useful to rewrite the linear regression solution, (35), in terms of a *data-data resolvent* matrix, defined as

$$Q(\gamma) \equiv \frac{1}{\gamma I_T + \varphi^T\varphi} \quad \Longleftrightarrow \quad Q_{\alpha_1\alpha_2}(\gamma) \equiv \left(\gamma\delta_{\alpha_1\alpha_2} + \sum_{j=1}^{N}\varphi_{j;\alpha_1}\varphi_{j;\alpha_2}\right)^{-1}\,, \tag{56}$$

where here $I_T \equiv \delta_{\alpha_1\alpha_2}$ represents the identity matrix on sample space. To see how to

---

[23]This is an important difference from Ref. [38], where the authors performed averages over random training data *only* using the results of a replica calculation from Refs. [35, 36]. In Appendix A, we discuss such (non-generalized) linear models using the simpler techniques of this paper. In particular, in §A.2 we explain how models with the right generative data model, but without any random feature maps, behave qualitatively differently than the LLMs observed in Ref. [18].

rewrite the linear regression solution, note that

$$q\varphi = \left[\sum_{s=0}^{\infty} \frac{1}{\gamma} \left(-\frac{\varphi\varphi^T}{\gamma} \right)^s \right] \varphi \tag{57}$$

$$= \frac{\varphi}{\gamma} \left[I_T - \frac{\varphi^T\varphi}{\gamma} + \left(\frac{\varphi^T\varphi}{\gamma} \right)^2 + \dots \right]$$

$$= \varphi Q \,,$$

where on the first line we used the definition of q, (36), to expand the resolvent around $\gamma \to \infty$, on the second line we pulled out a φ from the sum to the left and put the φ from the right into the sum, and on the third line we resummed the geometric series and used the definition (56). Using the transpose of this commutation relation, $\varphi^T q = Q\varphi^T$, we can rewrite the linear regression solution, (35), as

$$\theta^\star = (y + \epsilon)Q\varphi^T \,. \tag{58}$$

Now, let us state two simple identities that we can use to drastically simplify our calculations. First, from the definitions of the two resolvents, (36) and (56) note that

$$q(\gamma)^2 = -\frac{\partial}{\partial\gamma}q(\gamma) \,, \qquad Q(\gamma)^2 = -\frac{\partial}{\partial\gamma}Q(\gamma) \,. \tag{59}$$

With these, we can simplify the averaging by eliminating powers of q and Q from various expressions. Second, we also see from these definitions that

$$(\gamma I_N + \varphi\varphi^T)q = I_N \,, \qquad (\gamma I_T + \varphi^T\varphi)Q = I_T \,, \tag{60}$$

which will similarly be used to eliminate factors of $\varphi\varphi^T$ and $\varphi^T\varphi$.

Finally, we will find it convenient to adopt the following notation:

$$\overline{q} \equiv \langle q \rangle_x \qquad \overline{Q} \equiv \langle Q \rangle_x \,, \tag{61}$$

where the *overline* notation represents a training set average of the resolvent.

3.2   Data and Feature Averaging

We now begin with the more challenging part of the calculation, the dataset and random feature averages. As we will explain, the expectation of the test loss cannot be computed for all values of M, T, N: instead we will have to settle for expressions that are valid in the limit where $M, T, N \gg 1$, though their ratios, M/T, M/N, and T/N, may still take any value. This is not a problem: the neural scaling laws that LLMs exhibit in practice arise for very large data and models sizes, where our solutions are extremely accurate.[24] And, although we will not need to do so, the techniques that we describe below can be used to systematically compute the subleading corrections to the loss, which are suppressed by inverse powers of M, T, and N.

---

[24]From our numerics, cf. Fig. 8, we will see that subleading corrections are only important for very small training set sizes and numbers of features, $T, N \lesssim 10$.

3.2.1 The Noise Term

We begin with the *noise term* in (44). As the calculation that we perform here can be repurposed and significantly generalized when we analyze the label term, this section will also serve as a gentle introduction to our techniques.

Considering the noise term in the partially-averaged test loss, (44), and taking expectations over both datasets, x, $\widehat{x}$, and the random feature weights, u, we get:

$$
\begin{aligned}
\frac{\sigma_\epsilon^2}{2\widehat{T}} \left\langle \left\| \varphi^T q \widehat{\varphi} \right\|^2 \right\rangle_{\widehat{x},x,u} &= \frac{\sigma_\epsilon^2}{2\widehat{T}} \left\langle \mathrm{tr}\{ \varphi\varphi^T q \widehat{\varphi}\widehat{\varphi}^T q \} \right\rangle_{\widehat{x},x,u} . \\
&= \frac{\sigma_\epsilon^2}{2} \left\langle \mathrm{tr}\{ \varphi\varphi^T q \Omega q \} \right\rangle_{x,u} \\
&= \frac{\sigma_\epsilon^2}{2} \left\langle \left(1 + \gamma \frac{\partial}{\partial\gamma} \right) \mathrm{tr}\{ \Omega \bar{q} \} \right\rangle_u ,
\end{aligned}
\tag{62}
$$

Here, in the first line we expanded the square and expressed it as a trace, and in the second line we used our expression for the test-set covariance, (48), to perform the average over the test set. Finally, in the third line, we first used our first identity, (60), to eliminate the $\varphi\varphi^T$; we second used our second identity, (59), to exchange the q^2 term for a derivative; and we third used the definition, (61), to replace q with $\bar{q}$. In this way, we've reduced the computation to three steps: *(i)* compute the quantity

$$
\Delta^F \equiv \mathrm{tr}\{ \Omega \bar{q} \} ,
\tag{63}
$$

which involves evaluating the training-set average of the resolvent, $\bar{q}$, and then *(ii)* apply the differential operator, and finally *(iii)* evaluate its random feature average.

The only nontrivial step will be *(i)*, which we will now describe: we need to compute the training-set-averaged resolvent, (36),

$$
\bar{q} \equiv \langle q \rangle_x , \qquad q \equiv \frac{1}{\gamma I_N + \varphi\varphi^T} ,
\tag{64}
$$

using the fact that, for fixed u, the training-set average of φ is given by

$$
\langle \varphi_{j;\alpha} \rangle_x = 0 , \qquad \langle \varphi_{j_1;\alpha_1} \varphi_{j_2;\alpha_2} \rangle_x = \Omega_{j_1 j_2} \delta_{\alpha_1 \alpha_2} , \qquad \Omega \equiv u \Lambda u^T ,
\tag{65}
$$

cf. (45) and (48). To evaluate $\bar{q}$, we will take the limit of large training set and large number of random features, $T, N \to \infty$, with their ratio fixed.[25] This computation is a classic result of random matrix theory that goes back to the original work of Marchenko and Pastur [51] and was further studied in later works (see e.g. [52]).[26] Here we give a

---

[25]We are also implicitly taking the size of the latent space to be large, $M \to \infty$, though it may also have fixed ratios with T and N.

[26]In particular, the quantity $\mathrm{tr}\{\bar{q}(\gamma)\}$ is the Stieltjes transform of the eigenvalue distribution of the matrix $\varphi\varphi^T$: $\mathrm{tr}\{q(\gamma)\}$ is a meromorphic function with poles given by the (negative of the) eigenvalues of $\varphi\varphi^T$; after averaging over x to get $\mathrm{tr}\{\bar{q}(\gamma)\}$ these poles condense to a branch cut, and so the discontinuity across this branch cut determines the eigenvalue density. In the case where Ω is proportional to the identity, this reproduces the famous Marchenko-Pastur distribution.

simple derivation using *Feynman diagram* techniques that can be easily generalized to the other averages we will consider later.[27]

To begin, note from its definition, (64), that we can expand the resolvent in a power series in the ridge parameter:

$$q(\gamma) = \gamma^{-1} \sum_{s=0}^{\infty} (-\gamma)^{-s} \left(\varphi\varphi^T\right)^s. \tag{66}$$

Each term in this expansion is a power of the elements of the matrix φ. For fixed u, the elements of the matrix φ are Gaussian random variables with statistics (65). The higher-order moments of a zero-mean Gaussian distribution are determined entirely by the covariance matrix, $\Omega_{j_1 j_2}\delta_{\alpha_1 \alpha_2}$, and can be determined by a repeated application of (65). Let us enumerate the first few terms, which will make our overall strategy clear:

- The $s = 0$ term is trivial and simply given by γ^{-1}.

- The $s = 1$ term is also trivial and can be evaluated directly using the middle equation in (65), giving
$$- \gamma^{-2} \left\langle \varphi\varphi^T \right\rangle_x = -\gamma^{-2} T\Omega, \tag{67}$$
where the factor of T comes from the sum over the training set.

- The $s = 2$ term involves the average of the quartic matrix $\varphi\varphi^T\varphi\varphi^T$. Generally, the expectation value of a product of four arbitrary φ components is a sum of the three different ways that these features can be "paired up" together, with each pairing weighted by the appropriate covariance:[28]
$$\left\langle \varphi_{j_1;\alpha_1}\varphi_{j_2;\alpha_2}\varphi_{j_3;\alpha_3}\varphi_{j_4;\alpha_4} \right\rangle_x = \left\langle \varphi_{j_1;\alpha_1}\varphi_{j_2;\alpha_2} \right\rangle_x \left\langle \varphi_{j_3;\alpha_3}\varphi_{j_4;\alpha_4} \right\rangle_x + \tag{68}$$
$$\left\langle \varphi_{j_1;\alpha_1}\varphi_{j_3;\alpha_3} \right\rangle_x \left\langle \varphi_{j_2;\alpha_2}\varphi_{j_4;\alpha_4} \right\rangle_x +$$
$$\left\langle \varphi_{j_1;\alpha_1}\varphi_{j_4;\alpha_4} \right\rangle_x \left\langle \varphi_{j_2;\alpha_2}\varphi_{j_3;\alpha_3} \right\rangle_x,$$
which can be evaluated using (65) to give
$$\Omega_{j_1 j_2}\Omega_{j_3 j_4}\delta_{\alpha_1\alpha_2}\delta_{\alpha_3\alpha_4} + \Omega_{j_1 j_3}\Omega_{j_2 j_4}\delta_{\alpha_1\alpha_3}\delta_{\alpha_2\alpha_4} + \Omega_{j_1 j_4}\Omega_{j_2 j_3}\delta_{\alpha_1\alpha_4}\delta_{\alpha_2\alpha_3}. \tag{69}$$
We can use this to evaluate $\left\langle \varphi\varphi^T\varphi\varphi^T \right\rangle_x$ by setting $\alpha_1 = \alpha_2$, $\alpha_3 = \alpha_4$, $j_2 = j_3$ and summing, which altogether gives for the $s = 2$ term:
$$\gamma^{-3} \left\langle \left(\varphi\varphi^T\right)^2 \right\rangle_x = \gamma^{-3}\left(T^2\Omega^2 + T\Omega^2 + T\Omega\operatorname{tr}\{\Omega\}\right). \tag{70}$$
The three terms in this expression correspond directly to the three terms in (68).

---

[27]A similar diagrammatic derivation of $\bar{q}$ can be found in [53]. For other applications of Feynman diagrams in machine learning, see also [54].

[28]This an example of a general fact about Gaussian distributions: the expectation value of a product of Gaussian random variables is always a sum over the different ways that the variables can be paired up. This is known as *Isserlis' theorem* in probability theory and *Wick's theorem* in the physics literature. Note that if the statistics of φ were non-Gaussian, there could be an additional contribution to the right-hand side of (68) related to the fourth cumulant of the distribution.

In principle, it is straightforward – though tedious – to proceed order-by-order in this way, evaluating the average of each term. The main challenge is that the super-exponential proliferation of pairings that arise for larger s makes it difficult to keep track of how the indices in the resulting expressions are summed over: as is already evident in (70), many different types of terms can arise, with different powers of T, Ω, and $\mathrm{tr}\{\Omega\}$.

The method of **Feynman diagrams** is simply a diagrammatic technique that we can use to keep track of these various terms. For example, we represent the expectation value of the power series for the averaged resolvent, (66), as the following diagrammatic expression:

$$\overline{q} \;\equiv\; \text{———}\,\boxed{\overline{q}}\,\text{———} \tag{71}$$

$$=\; \gamma^{-1}\,\text{————————}$$

$$-\,\gamma^{-2}\,\text{——}\overset{\frown}{\underset{\ \ }{\text{-----}}}\text{——}$$

$$+\,\gamma^{-3}\,\text{——}\overset{\frown}{\text{---}}\text{——}\overset{\frown}{\text{---}}\text{——} \;+\;\gamma^{-3}\,\text{——}\overset{\frown\frown}{\text{---}}\text{——} \;+\;\gamma^{-3}\,\text{——}\overset{\frown\frown}{\text{---}}\text{——}$$

$$+\,\ldots$$

In this expression, each diagram on the right-hand side represents a term in the expansion and can be interpreted using the following rules:

- Each *single horizontal solid line* represents a *random-feature j-type index* as

$$\delta_{j_1 j_2} \;\equiv\; j_1\,\text{————————}\,j_2 \quad . \tag{72}$$

- Each *single horizontal dashed line* represents a *training-sample α-type index* as

$$\delta_{\alpha_1 \alpha_2} \;\equiv\; \alpha_1\,\text{- - - - - - - - -}\,\alpha_2 \quad . \tag{73}$$

- Then, we can represent $\varphi \equiv \varphi_{j;\alpha}$ as the *vertex* of a horizontal solid line and a horizontal dashed line as

$$\varphi_{j;\alpha} \;\equiv\; j\,\text{————}\,\text{- - - - -}\,\alpha \quad , \tag{74}$$

with φ^T drawn as the mirror image, with the dashed line preceding the solid line.

- Each *set of curved double lines* represents the *random-feature covariance matrix* as

$$\Omega_{j_1 j_2}\delta_{\alpha_1 \alpha_2} \;\equiv\; {}_{j_1;\alpha_1}\overset{\frown}{\frown}{}_{j_2;\alpha_2} \quad . \tag{75}$$

Note that it doesn't matter whether the dashed line is above or below the solid line, and in some diagrams they may even twist around each other.

30

- If two lines are connected, the corresponding indices are set equal and summed over.

- If a closed loop appears, the corresponding index is summed over.

Note also that for each of the diagrams in the expansion, we have written explicitly the factors of γ which accompany each of the terms in the sum. Please convince yourself that with these rules the second, third, and fourth lines in the diagrammatic expression (71) correspond to the $s = 0$, $s = 1$, and $s = 2$ terms, respectively, in the expansion computed above.

For example, consider the $s = 1$ term in the third line: first, we write $\Omega_{j_1 j_2} \delta_{\alpha_1 \alpha_2}$ for the curved double lines; then, we set $\alpha_1 = \alpha_2$ for the horizontal dashed line connecting to the curved dashed line; lastly, we sum over its index to give a factor of T. Including the correct factors of the ridge parameter and the overall sign, we find $-\gamma^{-2} T \Omega$, matching the $s = 1$ term, (67). Similarly, one can check that the three diagrams on the final line correspond to each of the three terms in the $s = 2$ contribution, (70).[29]

This Feynman diagram representation of the expansion is useful because it helps us select the terms that dominate in the large-T and large-N limit. Each closed loop in a Feynman diagram represents a trace: a loop of a dashed line will give a factor of T coming from the trace of the identity matrix I_T, while a loop of a solid line will give a factor of $\mathrm{tr}\{\Omega^p\}$ for some integer p, which will overall scale linearly with N in the large-N limit.[30]

With these rules, it turns out that a simple way to evaluate the overall scaling of a diagram is to consider its *topology*: the diagrams that dominate at large T and N are *planar*, i.e. can be drawn on the plane without any lines crossing one another. For example, inspecting (71), the only non-planar diagram is the middle diagram on the final line; this reflects the fact that the middle term in (70) scales only *linearly* with T, while the first and last terms scale *quadratically*, as T^2 and $T\mathrm{tr}\{\Omega\} \sim TN$, respectively. Indeed, when we take the limit of large T and N with the fixed ratio T/N, then there is a simple way to organize our calculation as an expansion in $1/N$: diagrams that can only be drawn on a surface of a more complicated topology are suppressed relative to the leading contributions by powers of N.[31] As we are only interested in the leading contribution, we will only need

---

[29]Note that you can also read off the form of the term before any averaging by considering the horizontal straight lines along the bottom of a diagram: the matrix $\varphi \equiv \varphi_{j;\alpha}$ is represented by a solid line followed by a dashed line, and its transpose, $\varphi^T \equiv (\varphi_{j;\alpha})^T$, is represented a dashed line followed by a solid line. The curved lines emerging from these vertices then determine the different pairings that result when averaging over x.

[30]This latter statement depends on the assumption that the elements $\Omega_{j_1 j_2}$ remain of order one as N is increased. This is explicitly true for the data models considered in our paper.

[31]In physics, this is known as the 't Hooft large-N expansion [55], and the limit of large T and N is often referred to as the *planar limit*. In mathematics, this observation is the starting point for the study of *free probability* (see, e.g., [56]): in free probability theory, we replace the usual cumulants, which are sums over partitions, by sums only over *non-crossing* partitions. The non-crossing criterion is, in the language of Feynman diagrams, just the statement that we include only planar diagrams. For example, in free probability we would label the three terms in (68) by the three possible partitions of the index

to consider planar diagrams in our calculation. By focusing on the leading contribution, we can thus make a significant reduction in the number of terms that we have to consider.

Even so, there are *still* an infinite number of diagrams that contribute $\bar{q}$. To simplify things further, we will organize the expansion of the resolvent, (66), in a slightly different way. First, let us write the *averaged* resolvent as

$$\bar{q}(\gamma) \equiv \frac{1}{\gamma I_N + S^F(\gamma)}, \tag{76}$$

which we take as the definition of some $N \times N$ matrix, S^F. Note importantly that S^F is independent of x, as x has already been averaged over in this formula. Comparing the original definition q, (64), with the definition of $\bar{q}$, (76), we see that the entire effect of averaging has been packaged into a substitution of the random matrix $\varphi\varphi^T$ for the fixed matrix S^F.[32] Next, expanding the definition (76) for large γ, we can express $\bar{q}(\gamma)$ as series,

$$q(\gamma) = \gamma^{-1} \sum_{s=0}^{\infty} (-\gamma)^{-s} \left(S^F\right)^s, \tag{77}$$

and use our Feynman rules to represent the terms in that series diagrammatically as

where the circle labeled with "S^F" represents an insertion of the matrix $S^F \equiv S^F_{j_1 j_2}$. Of course, this is only a reorganization of our original expansion (71): all that we've done is assumed that $\bar{q}$ can be resummed in a particular way and then repackaged our ignorance into the as-yet unknown matrix S^F.

To make this useful, let's develop a Feynman diagram expansion for S^F. This expansion must have the property that when inserted into the $\bar{q}$ expansion, (78), it reproduces our original expansion, (71). In the planar limit, it turns out that the correct expansion is:

set $\{1, 2, 3, 4\}$ into pairs, as $\{1, 2\} \cup \{3, 4\}$, $\{1, 3\} \cup \{2, 4\}$ and $\{1, 4\} \cup \{2, 3\}$. The first and third are non-crossing partitions, while the second "crossing" partition is excluded in free probability. Comparing to (71), this is just the planar limit.

[32]In the physics literature, quantities like S^F that appear in the denominators of averaged resolvents or propagators of interacting systems are often referred to as *self-energies*. As an N-by-N-dimensional matrix that in principle depends on the random feature weights, u, one might therefore call S^F the *self-energy of the features*. We only mention this name as an explanation for our choice of symbol, S^F.

Note that *(a)*, that rather than using solid horizontal lines instead we use subdiagrams labeled with "$\bar{q}$", which represents an insertion of the *averaged* $\bar{q}$ matrix, and *(b)* that the signs alternate between diagrams.[33]

To see why this is the correct expansion, note that diagrams in this expansion have the property that they cannot be split into two parts by cutting a single or double line in one particular place.[34] Contrast that with the first diagram on the fourth line in (71): there, if we cut the middle solid line connecting the two rising suns, then the diagram will be split into two separate pieces. This underscores the reason why organizing the computation in terms of S^F is helpful and suggests the strategy for its computation: first, write down an expansion for *all* the diagrams that have the property that they cannot be split by cutting a solid line according to our conditions *(a)–(b)*; then, by packaging them back into the expansion (78), we will generate *all* the needed diagrams, including those that fall into two disconnected pieces when a line is cut. To check that this claim is correct, you can insert this definition, (79), for S^F into the expansion (78); this will give exactly the same expansion as our original one, (71), provided that every time either $\bar{q}$ or S^F appears you also recursively insert expansions (78) or (79), respectively. The utility of this is that drawing all planar irreducible-by-cutting diagrams is much simpler than trying to draw all diagrams overall, and so by restricting to this set of diagrams – or equivalently, the terms in the expansion represented by these diagrams – we can make the computation of $\bar{q}$ not only tractable, but simple. Note that for this to work, the large-N limit restriction to planar diagrams was absolutely crucial: at finite N, simple recursive representations like this do not exist, although in many cases it is possible to systematically compute subleading $1/N$ corrections.

All that remains is to actually evaluate S^F. Evaluating the diagrams in the expansion (79) according to our rules for Feynman diagrams, we see that the pattern is

$$S^F = T\Omega - T\Omega \operatorname{tr}\{\Omega\bar{q}\} + T\Omega \operatorname{tr}\{\Omega\bar{q}\}^2 + \cdots + T\Omega \left(-\operatorname{tr}\{\Omega\bar{q}\}\right)^s + \ldots . \tag{80}$$

Recognizing yet another geometric series, we can sum it up to get

$$S^F(\gamma) = \frac{T\Omega}{1 + \Delta^F(\gamma)} , \tag{81}$$

where we've recalled our earlier definition, (63).[35] Together (76) and (81) form a system of two coupled matrix equations that can be solved to determine the x-averaged resolvent $\bar{q}$, and the quantity that we ultimately need, Δ^F.[36]

---

[33]In the physics literature, $\bar{q}$ is sometimes referred to as a *dressed* propagator, while, by contrast, the matrix corresponding to the solid horizontal line, I_N, is known as a *bare* propagator.

[34]In the physics literature, such diagrams are known as one-particle irreducible (1PI).

[35]Analogous to S^F, you might call Δ^F the *Delta of the features*.

[36]In the physics literature, (78) and (79) are sometimes known as Schwinger-Dyson equations [57, 58]. Note also that there is no simple closed form solution of these equations, except in certain special cases.

To finally finish item *(i)* in our three-step plan for calculating the noise term, let us rewrite our system of equations as an implicit equation for Δ^F as a function of Ω:

$$\Delta^F(\gamma) = \text{tr}\left\{\frac{\Omega}{\gamma I_N + \frac{T\Omega}{1+\Delta^F(\gamma)}}\right\}. \tag{82}$$

Here, we have used (81) to eliminate S^F from (76), multiplied $\bar{q}$ by Ω, and then took the trace of both sides. This equation can be solved numerically as a function of Ω, T, and γ, or can be analyzed analytically in the ridgeless limit ($\gamma \to 0$). In this paper, our analysis will mostly be concerned with the ridgeless limit. Now, let's continue with our plan with step *(ii)*, applying the differential operator, $\left(1 + \gamma\frac{\partial}{\partial\gamma}\right)$, and step *(iii)*, averaging over the random features, u.

To apply step *(ii)* in the ridgeless limit, we should expand Δ^F near $\gamma \to 0$:[37]

$$\Delta^F(\gamma) \equiv \frac{\Delta^F_{-1}}{\gamma} + \Delta^F_0 + \Delta^F_1\gamma + \dots. \tag{83}$$

Applying our differential operator to this expansion eliminates the leading term as

$$\lim_{\gamma \to 0}\left(1 + \gamma\frac{\partial}{\partial\gamma}\right)\Delta^F(\gamma) = \Delta^F_0, \tag{84}$$

which means that in this limit the noise term, (62), has a very simple expression:

$$\frac{\sigma^2_\epsilon}{2\widehat{T}}\left\langle\left\|\varphi^T q\widehat{\varphi}\right\|^2\right\rangle_{\widehat{x},x,u} = \frac{\sigma^2_\epsilon}{2}\left\langle\Delta^F_0\right\rangle_u, \tag{85}$$

where Δ^F_0 is understood to be determined by solving the self-consistent equation (82). It turn out the solution for Δ^F_0 depends on whether the model is underparameterized ($T > N$) or overparameterized ($T < N$) and also on whether the model is in the regime valid for neural scaling behavior ($N < M$) or in a regime where that behavior breaks down ($M > N$). The computation of Δ^F_0 in these four cases is somewhat technical and not particularly illuminating, so is relegated to Appendix B.2 (under the final subheading labeled Δ_0 *for the Truncated Power-Law Spectrum*).

Finally, for step *(iii)* we note that the solution for Δ^F_0 given in (267) does not explicitly depend on the fixed random feature weights, u. Thus, the average over u is trivial, and,

---

One scenario in which it is possible to jointly solve these equations is when the covariance proportional to the identity matrix: $\Omega = \sigma^2 I_N$. In this case, it's easy to check that the system, (78) and (79), reduces to a simple quadratic equation for $\text{tr}\{\bar{q}\}$. We discuss this case explicitly in §A.1 of the appendix.

[37]The justification for such an expansion is provided in Appendix B, but intuitively note the possibility of a $1/\gamma$ divergence as $\gamma \to 0$ by considering the self-consistent equation (82) as Δ^F becomes large.

substituting (267) for Δ_0^F in (85), we find a final answer for the noise term:

$$
\mathcal{L}_\epsilon \equiv \frac{\sigma_\epsilon^2}{2\widehat{T}} \left\langle \|\varphi^T q \widehat{\varphi}\|^2 \right\rangle_{\widehat{x},x,u} = \frac{\sigma_\epsilon^2}{2} \begin{cases} \alpha + \frac{1}{N/T-1} \,, & T < N \,, \ N < M \,, \\ \frac{1}{T/N-1} \,, & T > N \,, \ N < M \,, \\ \alpha + \frac{1}{M/T-1} \,, & T < M \,, \ N > M \,, \\ \frac{1}{T/M-1} \,, & T > M \,, \ N > M \,. \end{cases}
\tag{86}
$$

3.2.2 The Label Term

We now move on to the more complicated *label term* in (44). Starting from the partially-averaged expression and taking expectations over both datasets, x, $\widehat{x}$, and the random feature weights, u, we expand the square to express the term as a sum of traces:

$$
\begin{aligned}
\mathcal{L}_w &\equiv \frac{\sigma_w^2}{2\widehat{T}M} \left\langle \|x\varphi^T q \widehat{\varphi} - \widehat{x}\|^2 \right\rangle_{\widehat{x},x,u} \\
&= \frac{\sigma_w^2}{2\widehat{T}M} \left\langle \left(\mathrm{tr}\{\widehat{x}\widehat{x}^T\} - 2\mathrm{tr}\{\widehat{x}\widehat{\varphi}^T q\varphi x^T\} + \mathrm{tr}\{\widehat{\varphi}\widehat{\varphi}^T q\varphi x^T x\varphi^T q\} \right) \right\rangle_{\widehat{x},x,u} \\
&= \frac{\sigma_w^2}{2M} \left(\mathrm{tr}\{\Lambda\} - 2\left\langle \mathrm{tr}\{\widetilde{\Omega}q\varphi x^T\} \right\rangle_{x,u} + \left\langle \mathrm{tr}\{\Omega q\varphi x^T x\varphi^T q\} \right\rangle_{x,u} \right) \,.
\end{aligned}
\tag{87}
$$

In the final line, we used the random feature covariance, (48), and the mixed latent-feature-random-feature covariance, (50), to perform the averages over $\widehat{x}$. Thus, we have two expectations to evaluate.

To calculate these x and u expectations, we will need to separately consider the cases of underparameterization ($N < T$) and overparameterization ($N > T$). These computations are somewhat technical and follow very similar techniques to those discussed in the previous subsection; if you feel you've already got the hang of these techniques and are now ready for the final result please feel free to skip ahead to §3.2.3.

The Underparameterized Regime

We start with the case of $N < T$, proceeding in two steps: *(i)* first, we'll compute the average over the training set, x, and *(ii)* then, we'll perform the average over the random feature weights, u.

(i) Training Set Average

To perform the training set average, we will apply the same Feynman diagram techniques considered earlier to more involved expectations. The essential difference is that we now must consider both the random and latent features, φ and x, as Gaussian random matrices, with covariances given by

$$
\langle x_{I_1;\alpha_1} x_{I_2;\alpha_2}\rangle_x = \Lambda_{I_1 I_2}\delta_{\alpha_1\alpha_2}\,, \qquad \langle \varphi_{j_1;\alpha_1}\varphi_{j_2;\alpha_2}\rangle_x = \Omega_{j_1 j_2}\delta_{\alpha_1\alpha_2}\,, \qquad \langle x_{I;\alpha_1}\varphi_{j;\alpha_2}\rangle_x = \widetilde{\Omega}_{Ij}\delta_{\alpha_1\alpha_2}\,,
\tag{88}
$$

cf. (19), (48), and (50) . To account for these statistics, we will have to augment our Feynman rules from before:

- In addition to the horizontal solid line, which carried a *random-feature j-type index*, (72), and the horizontal dashed line, which carried a *training sample α-type index*, (73), we will need to introduce the horizontal *wavy line*, which will carry a *latent-feature I-type index* as

$$\delta_{I_1 I_2} \quad \equiv \quad {}_{I_1} \wwbar {}_{I_2} \quad . \tag{89}$$

- Then, we can represent $x \equiv x_{I;\alpha}$ as the *vertex* of a horizontal wavy line and a horizontal dashed line as

$$x_{I;\alpha} \quad \equiv \quad {}_I \wwbar \text{-----}{}_\alpha \quad . \tag{90}$$

x^T is drawn as the mirror image, with the dashed line preceding the wavy line.

- Similarly, in addition to the curved dashed-and-solid double lines in (75), which represents the covariance $\langle \varphi_{j_1;\alpha_1} \varphi_{j_2;\alpha_2} \rangle_x$, we must introduce two additional double lines which represent the other covariances in (88). The *curved dashed-and-wavy double line* will represent the covariance of the latent features, $\langle x_{I_1;\alpha_1} x_{I_2;\alpha_2} \rangle_x$, as

$$\Lambda_{I_1 I_2} \delta_{\alpha_1 \alpha_2} \quad \equiv \quad {}_{I_1;\alpha_1} \overset{\frown}{} {}_{I_2;\alpha_2} \quad , \tag{91}$$

and the *curved dashed-and-solid-turns-into-wavy* double lines represents the mixed covariance of random and latent features, $\langle \varphi_{j;\alpha_1} x_{I;\alpha_2} \rangle_x$, as

$$\widetilde{\Omega}_{Ij} \delta_{\alpha_1 \alpha_2} \quad \equiv \quad {}_{I;\alpha_1} \overset{\frown}{} {}_{j;\alpha_2} \quad . \tag{92}$$

Here, the dot on the *solid-and-wavy-line* reminds us that $\widetilde{\Omega} \equiv \widetilde{\Omega}_{Ij}$ transforms a latent feature index, I, into a random feature index, j.

As before, we will work in the limit where all the scales in the problem are all large $(M, N, T \to \infty)$ but with fixed ratios. Thus, we can continue to restrict our attention to only planar diagrams.

With these rules defined, let's now look at the expectations appearing in (87). Consider the first,

$$\left\langle \operatorname{tr}\left\{ \widetilde{\Omega} \, q\varphi x^T \right\} \right\rangle_{x,u} = \left\langle \operatorname{tr}\left\{ \widetilde{\Omega} \left\langle q\varphi x^T \right\rangle_x \right\} \right\rangle_u . \tag{93}$$

Here we have used the fact that $\widetilde{\Omega}$ is independent of the training set to move it outside the x average. The evaluation of $\left\langle q\varphi x^T \right\rangle_x$ is similar to that of $\langle q \rangle_x$, except that now the

36

factors of $(\varphi\varphi^T)^s$ appearing in the expansion of q, (66), can have correlations with φ *and* x, cf. (88).

To proceed, we expand q order-by-order in γ as in (66). Representing the terms in the expansion of $\left\langle q\varphi x^T\right\rangle_x$ using our Feynman rules, we find:

$$\left\langle q\varphi x^T\right\rangle_x \quad = \quad \gamma^{-1} \; \rule{0pt}{0pt}$$

$$- \quad \gamma^{-2} \; \rule{0pt}{0pt} \quad - \quad \gamma^{-2} \; \rule{0pt}{0pt}$$

$$+ \quad \dots \quad , \tag{94}$$

Since this is more involved than our previous expansions, it is worth summarizing what each of these diagrams represents. Each line represents a different order in the expansion (66) of q:

- The first line represents $\left\langle \gamma^{-1}\varphi x^T\right\rangle_x$, the term resulting from the $s = 0$ part of the expansion of $q(\gamma)$. To interpret the diagram, we first read along the bottom horizontal lines: the feature φ is represented by the vertex that connects a solid line and a dashed line, and then the latent feature x^T is represented by the the vertex that connects a dashed line and a wavy line. The average of these two matrices is given by (88),

$$\left\langle \varphi x^T\right\rangle_x = T\widetilde{\Omega}^T \, ; \tag{95}$$

  the double curved lines represent this covariance as per our rule (92). Note that our rules also tell us to include the factor of T due to the closed dashed-line loop.

- The second line represents $\left\langle -\gamma^{-2}\varphi\varphi^T\varphi x^T\right\rangle_x$, which arises from the $s = 1$ part of the expansion of $q(\gamma)$. Reading along the bottom of either diagram, we see a pattern of solid, dashed, and wavy lines that represents the quantity to be averaged, $\varphi\varphi^T\varphi x^T$. There are three ways to pair up these four matrices, but one of these would result in a non-planar diagram that is subleading. The two planar diagrams are depicted in the second line, and represent the pairings

$$\left\langle \varphi\varphi^T\right\rangle_x \left\langle \varphi x^T\right\rangle_x = T^2\Omega\widetilde{\Omega}^T \, , \qquad \left\langle \varphi x^T\right\rangle_x \left\langle \varphi^T\varphi\right\rangle_x = T\widetilde{\Omega}^T\mathrm{tr}\{\Omega\} \, , \tag{96}$$

  respectively. These diagrams both contribute at the same order in our large-(N, T) counting, since the first diagram scales like T^2 and the second diagram scales like $T\mathrm{tr}\{\Omega\} \sim TN$.

We could continue on expanding in this manner, but there is a better way. In fact, this far is enough to let us notice a pattern that allows us to immediately write down the answer in terms of quantities we've already computed.

To explain this, we will need to define one more object. Consider the data-data resolvent, $Q(\gamma)$, defined in (56). We can compute its data average, $\overline{Q} \equiv \langle Q \rangle_x$, via a diagrammatic expansion just as we did for $\overline{q}$:

$$\overline{Q} \quad \equiv \quad \cdots\!\!-\!\!\boxed{\overline{Q}}\!\!-\!\!\cdots \tag{97}$$

$$= \quad \gamma^{-1}\; \text{-----------} \quad - \quad \gamma^{-2}\; \text{---}\frown\text{---}$$

$$+ \gamma^{-3}\; \text{---}\frown\frown\text{---} \quad + \quad \gamma^{-3}\; \text{---}\frown\frown\text{---} \quad - \quad \cdots$$

$$= \quad \gamma^{-1}\; \text{-----------} \; - \gamma^{-1}\; \text{---}\boxed{\overline{q}}\frown\text{---} \; + \gamma^{-1}\; \text{---}\boxed{\overline{q}}\frown\boxed{\overline{q}}\frown\text{---} \; - \; \cdots .$$

Here, to reach the final line, we recognized a pattern: the diagrams appearing in $\overline{Q}$ contain as subdiagrams the same diagrams that appeared in the expansion for $\overline{q}$.[38] Note also that all but one of the factors of γ^{-1} that appear in the expansion of $\overline{Q}$ are accounted for by the factors of $\overline{q}$ appearing in the final line. Evaluating these diagrams on the final line, we can sum the resulting geometric series:

$$\overline{Q} = \gamma^{-1}I_T - \gamma^{-1}\text{tr}\{\Omega\overline{q}\}\, I_T + \gamma^{-1}\left(\text{tr}\{\Omega\overline{q}\}\right)^2 I_T - \ldots \tag{98}$$
$$= \frac{I_T}{\gamma(1 + \Delta^F)}\,,$$

where to get the resulting expression we used the definition $\Delta^F \equiv \text{tr}\{\Omega\overline{q}\}$. As a sanity check, note that $\overline{Q}$ must be proportional to the identity because none of our covariance matrices have nontrivial training-set indices after averaging over the data, cf. (88).

Returning now to (94), we can similarly see a pattern in the expansion:

$$\langle q\varphi x^T \rangle_x \quad = \quad \gamma\; \text{-}\boxed{\overline{q}}\text{-}\frown\text{-}\boxed{\overline{Q}}\text{-}\sim \quad , \tag{99}$$

which can be easily checked by recursively substituting in the diagrams for $\overline{q}$ and $\overline{Q}$. Note that there is an explicit factor of γ in this expression to account for the factors of γ^{-1} that

---

[38]It is easiest to see this by using the (78) form of the $\overline{q}$ expansion and plugging in for S^F with its diagrams, (79). One could also define a self-energy matrix like S^F as an intermediate step in the computation of $\overline{Q}$; the diagrammatic expansion for that object makes the result (98) almost immediately apparent. In fact, from (98), we see that the analog of S^F for $\overline{Q}$ is just the matrix $\gamma\Delta^F I_T$.

are included in *both* $\bar{q}$ and $\overline{Q}$. Using our rules to evaluate this diagram, we find the first x expectation in the label term, (93) is:

$$\left\langle \mathrm{tr}\left\{\widetilde{\Omega}\, q\varphi x^T\right\}\right\rangle_{x,} = \left(\frac{T}{1+\Delta^F}\right) \mathrm{tr}\left\{\widetilde{\Omega}^T\widetilde{\Omega}\,\bar{q}\right\}, \tag{100}$$

where, to get this expression, we used here the cyclicity of the trace and substituted for the trace of (98),

$$\gamma\,\mathrm{tr}\left\{\overline{Q}\right\} = \frac{T}{1+\Delta^F}\,. \tag{101}$$

Now let's consider the final expectation in the label term (87):

$$\left\langle \mathrm{tr}\left\{\Omega q\varphi x^T x\varphi^T q\right\}\right\rangle_{x,u} = \left\langle \mathrm{tr}\left\{\Omega\left\langle q\varphi x^T x\varphi^T q\right\rangle_x\right\}\right\rangle_u \tag{102}$$

where we've used the fact that Ω is independent of the training set to take it outside the x average. This will be the most complicated expectation to evaluate, since terms like $(\varphi\varphi^T)^s$ resulting from the expansion of *either* q in the expression can pair with any of the other four matrices: $\varphi, x^T, x, \varphi^T$.

To proceed, we will use a similar diagrammatic strategy as before: we will consider the various diagrams that result from the dual expansion of the resolvents in $\left\langle q\varphi x^T x\varphi^T q\right\rangle_x$ and try to recognize patterns in the subdiagrams. Before we begin, in this case note that the different terms arising in the expectation must be determined by the different ways in which the x and x^T are paired. In particular, there are four different possibilities to consider: *(1)* the x and x^T are paired together; *(2)* the x is paired with a φ^T, and the x^T is paired with a φ to their left; *(3)* the x is paired with a φ^T, and the x^T is paired with a φ to their right; and *(4)* the x^T is paired with a φ to the left and the x^T is paired with a φ^T to the right. As you can check, any other pairings will involve crossings, leading to non-planar diagrams, and therefore subleading contributions in the large-N limit.

Next, when expanding both q's and making these pairings, you'll notice that for the first three pairings, *(1)–(3)*, the expectation will factorize into a part with a common factor, while for the fourth pairing, *(4)*, there will be two terms, one that factorizes the same way and one that does not. Altogether, the structure will be

$$\left\langle q\varphi x^T x\varphi^T q\right\rangle_x = \left\langle q\varphi\left(\widetilde{P}^{(1)} + \widetilde{P}^{(2)} + \widetilde{P}^{(3)} + \widetilde{P}^{(4)}\right)\varphi^T q\right\rangle_x + \widetilde{P}^{(4,\mathrm{c})}, \tag{103}$$

where the various $\widetilde{P}^{(a)} \equiv \widetilde{P}^{(a)}_{\alpha_1\alpha_2}$ are $(T \times T)$-dimensional *x-independent* matrices that represent the four factorized contributions from the four different pairings, and $\widetilde{P}^{(4,\mathrm{c})} \equiv \widetilde{P}^{(4,\mathrm{c})}_{j_1 j_2}$ is an $(N \times N)$-dimensional *x-independent* matrix representing the so-called *connected* contribution that does not factorize the same way as the other. This decomposition, (103), can be shown by expanding in both q's and recognizing that the resulting diagrammatic expansion can be reorganized into sets of *disconnected* diagrams and *connected* diagrams. Importantly, in doing so you will see that the overall factor for the disconnected diagrams is the same for each pairing.

Moreover, as we already explained above after (98) when evaluating $\overline{Q}$, the only matrix with training-data indices that can appear after x-averaging is the identity matrix I_T. Thus, each of these factorized pairings can be written as $\widetilde{P}^{(a)} \equiv P^{(a)} I_T$, for scalar $P^{(a)}$, which lets us further simplify (103) as

$$\left\langle q\varphi x^T x \varphi^T q \right\rangle_x = \left\langle q\varphi\varphi^T q \right\rangle_x \left(P^{(1)} + P^{(2)} + P^{(3)} + P^{(4)} \right) + \widetilde{P}^{(4,c)} . \qquad (104)$$

Finally, the factor, $\left\langle q\varphi\varphi^T q \right\rangle_x$, can be expressed in terms of already known quantities as

$$\left\langle q\varphi\varphi^T q \right\rangle_x = \left(1 + \gamma \frac{\partial}{\partial\gamma} \right) \overline{q} , \qquad (105)$$

where we've used our identities, (60) and then (59), and finally our definition $\overline{q} \equiv \langle q \rangle_x$.
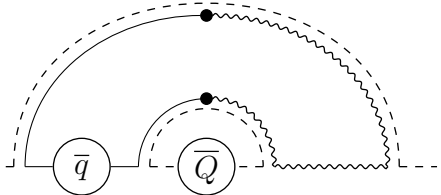
Now, let's re-enumerate the four possible pairings of x and x^T to evaluate the four $\widetilde{P}^{(a)}$ and $\widetilde{P}^{(4,c)}$:

(1) The simplest case is when x and x^T are paired together:

$$\widetilde{P}^{(1)} \equiv \left\langle x^T x \right\rangle_x \quad = \quad \text{} \quad = \quad \text{tr}\{\Lambda\}\, I_T . \qquad (106)$$

This is trivial, but it's still nice to practice our drawing.

(2) The next contribution is when both x and x^T pair with a φ^T and a φ to their left:

$$\widetilde{P}^{(2)} \equiv -\left\langle \varphi^T \left\langle q\varphi x^T \right\rangle_x x \right\rangle_x \quad = \quad -\gamma\ \text{} \quad . \qquad (107)$$

Here we have written this term as a set of nested expectation values such that x^T and x can both be contracted to the left without any crossings. To understand why the diagram is the correct way to represent these nested expectations, note that the inner expectation, $\left\langle q\varphi x^T \right\rangle_x$, was already analyzed in (99) and appears as a subdiagram here, while the outer expectation is evaluated using our Feynman rule (92). To understand why the overall factor of $-\gamma$ appears, let's go back to how this term would naturally arise in our resolvent expansion: to find this term, we would expand the first q in $\left\langle q\varphi x^T x \varphi^T q \right\rangle_x$ as $q = \gamma^{-1} \sum_s (-\gamma)^{-s} \left(\varphi\varphi^T \right)^s$ and then rewrite it in terms of two sums for a pair of resolvents q and Q; the point of this reorganization would be to extract a factor of φ to pair up with the x^T and a factor of φ^T to pair up with the x. Since each factor of $\varphi\varphi^T$ in the expansion comes with a factor of

40

$-\gamma^{-1}$, when we pull out only a factor $\varphi\varphi^T$ in this way, we must also multiply by an explicit factor of $-\gamma$ to compensate. Evaluating the diagram, we get:

$$P^{(2)} = -\gamma \operatorname{tr}\{\overline{Q}\} \operatorname{tr}\{\widetilde{\Omega}^T \widetilde{\Omega}\overline{q}\} = -\left(\frac{T}{1+\Delta^F}\right) \operatorname{tr}\{\widetilde{\Omega}^T \widetilde{\Omega}\overline{q}\} \ , \qquad (108)$$

where in the last expression we used (101) to substitute for $\gamma \operatorname{tr}\{\overline{Q}\}$.

(3) Similarly, both x and x^T can pair with a φ^T and φ to their right:

$$\widetilde{P}^{(3)} \equiv -\left\langle x^T \left\langle x\varphi^T q \right\rangle_x \varphi \right\rangle_x \qquad (109)$$

As this is just the transpose of $\widetilde{P}^{(2)}$, the scalar trace is identical:

$$P^{(3)} = P^{(2)} = -\left(\frac{T}{1+\Delta^F}\right) \operatorname{tr}\{\widetilde{\Omega}^T \widetilde{\Omega}\overline{q}\} \ . \qquad (110)$$

(The diagram representing this would be the mirror image of (107).)

(4) Finally, x^T can pair with a φ to the left, and x can pair with a φ^T to the right:

$$\widetilde{P}^{(4)} \equiv \left\langle \varphi^T \left\langle q\varphi x^T \right\rangle_x \left\langle x\varphi^T q \right\rangle_x \varphi \right\rangle_x = \gamma^2 \ \ \text{} \ , \qquad (111)$$

To see why the pairing given by the middle expression is correct, note that this is the appropriate and only way to contract the x and x^T in different directions without any crossings. Similar to before, we can see why this is the correct diagram to represent these nested expectations by using our Feynman rule, (75), for the outer part, and then by using (99) and its mirror as subdiagrams to represent the two middle expectations.[39] Evaluating this diagram, we get:

$$P^{(4)} = \left(\gamma \operatorname{tr}\{\overline{Q}\}\right)^2 \operatorname{tr}\{\overline{q}\widetilde{\Omega}^T \widetilde{\Omega}\overline{q}\Omega\} = \left(\frac{T}{1+\Delta^F}\right)^2 \operatorname{tr}\{\overline{q}\widetilde{\Omega}^T \widetilde{\Omega}\overline{q}\Omega\} \ , \qquad (112)$$

where, as is our recent *M.O.*, we used (101) to substitute for $\gamma \operatorname{tr}\{\overline{Q}\}$.

---

[39]The factor of γ^2 appears because in this case we've extracted two factors of $\varphi\varphi^T$ when reorganizing our expansion, one from each of the q's in $\left\langle q\varphi x^T x\varphi^T q \right\rangle_x$, and so we must multiply by $(-\gamma)^2$ to compensate.

(4, c) Finally prime, again x^T can pair with a φ to the left, and x can pair with a φ^T to the right:

$$\widetilde{P}^{(4,\mathrm{c})} \equiv \big\langle q\varphi x^T \big\rangle_x \big\langle x\varphi^T q \big\rangle_x \;=\; \gamma^2 \;-\!\!\!\overline{q}\!\!\!-\cdots\overline{Q}\cdots\sim\sim\cdots\overline{Q}\cdots-\!\!\!\overline{q}\!\!\!-\;,$$

(113)

and we can use (99) and (101) to evaluate this as

$$\widetilde{P}^{(4,\mathrm{c})} = \left(\frac{T}{1+\Delta^F}\right)^2 \overline{q}\widetilde{\Omega}^T\widetilde{\Omega}\overline{q}\,.$$

(114)

Assembling all five terms back into (103), and then substituting into (102), we find

$$\left\langle \mathrm{tr}\{\Omega q\varphi x^T x\varphi^T q\} \right\rangle_x = \left(\frac{T}{1+\Delta^F}\right)^2 \mathrm{tr}\{\Omega\overline{q}\widetilde{\Omega}^T\widetilde{\Omega}\overline{q}\} + \left[\left(1+\gamma\frac{\partial}{\partial\gamma}\right)\Delta^F\right] \times$$

(115)

$$\left[\mathrm{tr}\{\Lambda\} - 2\left(\frac{T}{1+\Delta^F}\right)\mathrm{tr}\{\widetilde{\Omega}^T\widetilde{\Omega}\overline{q}\} + \left(\frac{T}{1+\Delta^F}\right)^2\mathrm{tr}\{\overline{q}\widetilde{\Omega}^T\widetilde{\Omega}\overline{q}\Omega\}\right],$$

where the first term on the first line is the connected contribution, $\widetilde{P}^{(4,\mathrm{c})}$, (114), and where the second term gives the four factorized contributions *(1)–(4)*, for which we used $\Delta^F \equiv \mathrm{tr}\{\Omega\overline{q}\}$, with the $\overline{q}$ coming from the factor in (105).

Finally, we can now collect the contribution to the x average of the label term in the underparameterized regime. Starting from our expression (87) and substituting in both evaluated expectations, (100) and (115), we get

$$\mathcal{L}_w = \frac{\sigma_w^2}{2M} \left\langle \left[1 + \left(1+\gamma\frac{\partial}{\partial\gamma}\right)\Delta^F\right] \times \right.$$

(116)

$$\left. \left[\mathrm{tr}\{\Lambda\} - 2\left(\frac{T}{1+\Delta^F}\right)\mathrm{tr}\{\widetilde{\Omega}^T\widetilde{\Omega}\overline{q}\} + \left(\frac{T}{1+\Delta^F}\right)^2\mathrm{tr}\{\overline{q}\widetilde{\Omega}^T\widetilde{\Omega}\overline{q}\Omega\}\right] \right\rangle_u,$$

where all the other terms just conspired to give "1" inside the first square brackets. (Note that the terms in the second square brackets could be written as a square, if desired.) Lastly, note that from combining the Schwinger-Dyson equations, (76) and (81), we can write

$$\left(\frac{T}{1+\Delta^F}\right)\Omega\overline{q} = I_N - \gamma\overline{q}\,,$$

(117)

where to get this we first inverted (76) to get an expression for S^F, and then we it equal to the right-hand side of (81) and multiplied both sides by $\overline{q}$. Using this new identity, (117), and remembering that Ω and $\overline{q}$ commute and that the trace is cyclic, we can finally rewrite $\mathcal{L}_w$ compactly as

$$\mathcal{L}_w = \frac{\sigma_w^2}{2M} \left\langle \left[1 + \left(1+\gamma\frac{\partial}{\partial\gamma}\right)\Delta^F\right]\left(\mathrm{tr}\{\Lambda\} - \left(\frac{T}{1+\Delta^F}\right)\mathrm{tr}\{(I_N+\gamma\overline{q})\overline{q}\widetilde{\Omega}^T\widetilde{\Omega}\}\right) \right\rangle_u.$$

(118)

(ii) Random Feature Average

We now turn to the average over the random features, u. This will be simpler than the x averages we just performed: our strategy will be to massage $\mathcal{L}_w$ into a form where we can apply the same (simpler) machinery that was used in §3.2.1 for the noise term.

We begin by noting that the latent feature covariance matrix, Λ, is symmetric and positive semi-definite, by definition. This implies that it is possible to find a symmetric, positive semi-definite matrix, v, that is the "square root" of Λ, such that[40]

$$\Lambda \equiv vv^T, \qquad v = v^T. \tag{119}$$

Recalling definitions (49) and (51), we can express our other two covariances as

$$\widetilde{\Omega} = vv^T u^T, \qquad \Omega = uvv^T u^T, \tag{120}$$

and, using (76) and (81), express our x-averaged resolvent as

$$\bar{q}(\gamma) = \frac{1}{\gamma I_N + \left(\frac{T}{1+\Delta^F}\right) uvv^T u^T}. \tag{121}$$

As we will soon see, the utility of this is that as our expression for $\mathcal{L}_w$ had various products of $\bar{q}$ and $\widetilde{\Omega}$, with our new notation we can now rewrite these using

$$\bar{q}uv = uv\mathbb{Q}, \tag{122}$$

where we have defined the *latent-feature resolvent*:

$$\mathbb{Q}(\gamma) \equiv \frac{1}{\gamma I_M + \left(\frac{T}{1+\Delta^F}\right) v^T u^T uv}. \tag{123}$$

This is an $(M \times M)$ matrix $\mathbb{Q} \equiv \mathbb{Q}_{I_1 I_2}$ that operates on our latent space.[41] Contrasting carefully this new definition with the one for our (random) feature-feature resolvent q, (36), and the data-data resolvent Q, (56), we note that those resolvents are defined in terms of the matrix φ, which depend on *both* the random training set x and random feature weights u, while here $\mathbb{Q}$ depends *only* on the random feature weights u.

Much like our other resolvents, $\mathbb{Q}$ satisfies identities

$$\mathbb{Q}^2(\gamma) = -\frac{\partial}{\partial \gamma}\mathbb{Q}(\gamma), \qquad \left[\gamma I_M + \left(\frac{T}{1+\Delta^F}\right) v^T u^T uv\right]\mathbb{Q} = I_M. \tag{124}$$

---

[40]Since v is symmetric we could equally well have written this equation as $\Lambda \equiv v^2$. We prefer to use the form (119) because it makes manifest that Λ is symmetric.

[41]To derive (122) we can follow the same manipulations we used to derive the commutation relation (57), with the substitutions $q \to \bar{q}$ and $\varphi \to uv$.

Multiplying the second identity by Λ, taking the trace, and rearranging, we find

$$\text{tr}\{\Lambda\} - \left(\frac{T}{1+\Delta^F}\right)\text{tr}\left\{\bar{q}\tilde{\Omega}^T\tilde{\Omega}\right\} = \gamma\text{tr}\{\Lambda\mathbb{Q}\} \ , \tag{125}$$

where in the second term on the left-hand side we used (120), (122), and the fact that $v = v^T$. Taking $-\gamma\frac{\partial}{\partial\gamma}$ of this equation gives

$$-\gamma\left(\frac{T}{1+\Delta^F}\right)\text{tr}\left\{\bar{q}^2\tilde{\Omega}^T\tilde{\Omega}\right\} = -\gamma\text{tr}\{\mathbb{Q}\Lambda\} - \gamma^2\frac{\partial}{\partial\gamma}\text{tr}\{\Lambda\mathbb{Q}\} \ , \tag{126}$$

where we used the identity (59) to rewrite the derivative of $\bar{q}$ in terms of its square. We now note that the sum of the left hand sides of (125) and (126) is exactly the combination that appears in $\mathcal{L}_w$, allowing us to rewrite (118) as

$$\mathcal{L}_w = \frac{\sigma_w^2}{2M}\left\langle\left[1 + \left(1 + \gamma\frac{\partial}{\partial\gamma}\right)\Delta^F\right]\left(-\gamma^2\frac{\partial}{\partial\gamma}\text{tr}\{\Lambda\mathbb{Q}\}\right)\right\rangle_u \ . \tag{127}$$

This is progress, since we have written $\mathcal{L}_w$ entirely in terms of the expectation value of the new resolvent, $\mathbb{Q}$.

To proceed let's consider the ridgeless limit $\gamma \to 0$, where this expression simplifies considerably. In §B.2, we explain that in the underparameterized regime ($N < T$), we have an expansion that begins with the constant term,[42]

$$\Delta^F = \Delta_0^F + \Delta_1^F\gamma + \dots \ , \tag{128}$$

and thus in the ridgeless limit we have:

$$\lim_{\gamma\to 0}\Delta^F = \Delta_0^F \ . \tag{129}$$

Similarly, the leading term in the square brackets has the limit

$$\lim_{\gamma\to 0}\left(1 + \gamma\frac{\partial}{\partial\gamma}\right)\Delta^F(\gamma) = \Delta_0^F \ . \tag{130}$$

Moreover, as we discussed at the conclusion of §3.2.1 and in more detail in §B.2 under the final subheading labeled Δ_0 *for the Truncated Power-Law Spectrum*, Δ_0^F does not explicitly depend on the fixed random feature weights, u, and is given by (86). In the underparameterized regime for $N < M$, this is simply

$$\Delta_0^F = \frac{1}{T/N - 1} \ . \tag{131}$$

---

[42]To see this from our discussion in §B.2 under the subheading $\mu > M$, make the substitutions $\mu \to T$ and $M \to N$. This expansion holds regardless of the form of the covariance.

We are thus left with evaluating the term in round brackets in (127):

$$\lim_{\gamma \to 0} \left\langle -\gamma^2 \frac{\partial}{\partial \gamma} \text{tr}\{\Lambda \mathbb{Q}\} \right\rangle_u = \lim_{\gamma \to 0} \left(-\gamma^2 \frac{\partial}{\partial \gamma} \text{tr}\{\Lambda \langle \mathbb{Q} \rangle_u\} \right). \tag{132}$$

The quantity $\text{tr}\{\Lambda \langle \mathbb{Q} \rangle_u\}$ is quite similar to the quantity $\Delta^F \equiv \text{tr}\{\Lambda \bar{q}\}$ that was computed back in §3.2.1. To proceed, we define a new random variable

$$\omega \equiv \left(\frac{M}{\sigma_u^2} \right)^{1/2} v^T u^T \quad \Longleftrightarrow \quad \omega_{Ij} \equiv \left(\frac{M}{\sigma_u^2} \right)^{1/2} \sum_{J=1}^{M} u_{jJ} v_{JI}, \tag{133}$$

which, since v is a fixed matrix, has the statistics of zero-mean Gaussian when averaged over u:

$$\langle \omega_{Ij} \rangle_u = 0, \qquad \langle \omega_{I_1 j_1} \omega_{I_2 j_2} \rangle_u = \Lambda_{I_1 I_2} \delta_{j_1 j_2}, \tag{134}$$

where to evaluate the covariance in (134) we used statistics (29) and the definition (119). In terms of this variable ω, $\mathbb{Q}$ takes the form

$$\mathbb{Q}(\gamma) = \frac{1}{\gamma I_M + \zeta \, \omega \omega^T}. \tag{135}$$

where we have defined

$$\zeta \equiv \frac{\sigma_u^2}{M} \left(\frac{T}{1 + \Delta^F} \right). \tag{136}$$

Naively it would be a problem that the factor of Δ^F appearing in ζ itself depends on u, which would make it rather difficult to compute $\langle \mathbb{Q} \rangle_u$ in general.[43] However, as we noted above in (129), Δ^F is independent of u in the ridgeless limit. So as long as we are working in this limit we may treat ζ as a constant which is independent of u.

At this point u-averaging (135) would be form-identical to the x-averaging of q we already performed, but for the constant ζ. It's not to hard to see that if we were to follow the same method we used to solve for $\bar{q}$ in §3.2.1 – i.e. by rescaling Λ – we'd find

$$\langle \mathbb{Q} \rangle_u = \frac{1}{\gamma I_M + \frac{\zeta N \Lambda}{1 + \Delta^L}}, \tag{137}$$

where $\Delta^L \equiv \text{tr}\{\Lambda \mathbb{Q}\}$ solves the self-consistent equation[44]

$$\Delta^L = \text{tr}\left\{ \frac{\zeta \Lambda}{\gamma I_M + \frac{\zeta N \Lambda}{1 + \Delta^L}} \right\}. \tag{138}$$

As per our discussion at the beginning of §B.1, for $N < M$ we have an expansion in the $\gamma \to 0$ limit as

$$\Delta^L = \Delta^L_{-1} \gamma^{-1} + \Delta^L_0 + \dots, \tag{139}$$

---

[43]While Δ_0^F is independent of u, is not necessarily true that Δ^F is similarly independent.

[44]The "L" in Δ^L stands for *latent features*.

and plugging this expansion back into the self-consistent equation (138) tells us that Δ_{-1}^L solves the non-linear equation

$$1 = \text{tr}\left\{\frac{\zeta\Lambda}{\Delta_{-1}^L I_M + \zeta N\Lambda}\right\}, \tag{140}$$

while plugging the expansion into the averaged resolvent (137) gives

$$\langle\mathbb{Q}\rangle_u = \frac{\Delta_{-1}^L I_M}{\Delta_{-1}^L I_M + \zeta N\Lambda}\gamma^{-1} + \dots . \tag{141}$$

In fact, although it is not immediately obvious from (141), it turns out that the leading part of $\langle\mathbb{Q}\rangle_u$ as $\gamma \to 0$ written in this formula is actually independent of ζ. To see this, we can take the derivative of (140) with respect to ζ to obtain

$$0 = \text{tr}\left\{\frac{\Lambda}{(\Delta_{-1}^L I_M + \zeta N\Lambda)^2}\right\}\left(\Delta_{-1}^L - \zeta\frac{\partial\Delta_{-1}^L}{\partial\zeta}\right) \quad\Longrightarrow\quad \Delta_{-1}^L - \zeta\frac{\partial\Delta_{-1}^L}{\partial\zeta} = 0. \tag{142}$$

We can then use this to evaluate the derivative with respect to ζ of the coefficient of γ^{-1} term in the $\langle\mathbb{Q}\rangle_u$ expansion in (141):

$$\frac{d}{d\zeta}\left(\frac{\Delta_{-1}^L I_M}{\Delta_{-1}^L I_M + \zeta N\Lambda}\right) = \frac{N\Lambda}{(\Delta_{-1}^L I_M + \zeta N\Lambda)^2}\left(\zeta\frac{\partial\Delta_{-1}^L}{\partial\zeta} - \Delta_{-1}^L\right) = 0. \tag{143}$$

We thus conclude that this term is independent of ζ.

This implies that, as long so we are concerned only with the $\gamma \to 0$ limit, we may set $\zeta = 1$ in our computation of $\langle\mathbb{Q}\rangle_u$, a computation which is now manifestly the same as the one we already performed. In other words, using limits (130) and (132), substituting for the constant (131), we obtain our our final result for the label term, (127), in the under-parameterized regime $(N < T)$, and $(N < M)$:

$$\mathcal{L}_w = \frac{\sigma_w^2}{2M}\left(\frac{\Delta_{-1}^L}{1 - N/T}\right), \tag{144}$$

where Δ_{-1}^L is now determined by solving the ζ-independent non-linear equation

$$1 = \text{tr}\left\{\frac{\Lambda}{\Delta_{-1}^L I_M + N\Lambda}\right\}. \tag{145}$$

Curiously, this ζ-independence means that the variance of the random feature weights, σ_u^2, drops out of the solution (144) as well.

The Over-Parameterized Regime

We can now consider the over-parameterized regime where $N > T$. In this regime, we will need to average over the random features u *first*, and then analyze the resulting expression in ridgeless limit to determine the x average. Don't despair: luckily, there is a clever trick that we can use to find the answer with nearly no additional work.

Starting from pre-$\widehat{x}$-averaged form of the label term – the middle line in (87) – let us rewrite this term as

$$\mathcal{L}_w = \left\langle \frac{\sigma_w^2}{2\sigma_u^2 \widehat{T}} \left(\text{tr}\big\{\widehat{\Sigma}\big\} - 2\text{tr}\big\{\widetilde{\Sigma}^T Q \varphi^T \widehat{\varphi}\big\} + \text{tr}\big\{\Sigma Q \varphi^T \widehat{\varphi}\widehat{\varphi}^T \varphi Q\big\} \right) \right\rangle_{\widehat{x},x,u} \tag{146}$$

To get this expression, we've used the definitions of the matrices $\widehat{\Sigma}$, $\widetilde{\Sigma}$ and Σ, cf. (47), used the commutation relation $q\varphi = \varphi Q$, cf. (57), and used, on the middle term, the fact that the trace is invariant under transpose. Now, note that this form of the pre-averaged label term, (146), is in a very similar form to the post-$\widehat{x}$-averaged form of the label term – the final line in (87) – with the matrix substitutions:

$$\widehat{\Sigma} \to \Lambda, \quad \widetilde{\Sigma}^T \to \widetilde{\Omega}, \quad \Sigma \to \Omega, \quad Q \to q, \quad \varphi^T \to \varphi, \quad \widehat{\varphi} \to x^T, \tag{147}$$

and the overall rescaling

$$\sigma_u^2 \widehat{T} \to M. \tag{148}$$

Moreover, considering the definition of the training set features φ, cf. (28) and (32),

$$\varphi \equiv ux \quad \Longleftrightarrow \quad \varphi_{j;\alpha} \equiv \sum_{I=1}^{M} u_{jI} x_{I;\alpha}, \tag{149}$$

we see that transformation $\varphi^T \to \varphi$ actually is enacting a *training-data–random-feature duality*:

$$x^T \to u, \quad u^T \to x, \tag{150}$$

which also implies

$$T \to N, \quad N \to T. \tag{151}$$

What this means is that the u average of the label term in the form (146) is really the same calculation as the x average we performed in the underparameterized regime so long as we make the transformations (147) and (150).[45] Ultimately this means that we will find the same final answer (144) so long as we exchange $N \to T$ and $T \to N$.

Let us see how this works in slightly more detail. If you are already happy with the application of the duality, you may skip to the end of the section. To proceed carefully, let

---

[45]This remarkable simplification is due to a very general duality transformation for certain linear models. Such a duality is not completely new and has been noted before in some settings [38].

us apply the inverse of these transformations (147) and (151) to rewrite our computation of the x-averaged label term, (118):

$$\mathcal{L}_w = \frac{\sigma_w^2}{2\sigma_u^2 \widehat{T}} \left\langle \left[1 + \left(1 + \gamma \frac{\partial}{\partial \gamma} \right) \Delta^{F'} \right] \times \right. \tag{152}$$
$$\left. \left(\text{tr}\{\widehat{\Sigma}\} - \left(\frac{N}{1 + \Delta^{F'}} \right) \text{tr}\{ (I_T + \gamma \langle Q \rangle_u) \langle Q \rangle_u \widetilde{\Sigma} \widetilde{\Sigma}^T \} \right) \right\rangle_{x,\widehat{x}}.$$

Here, the u-average of Q is given by

$$\langle Q \rangle_u = \frac{1}{\gamma I_T + \left(\frac{N}{1 + \Delta^{F'}} \right) \Sigma}, \tag{153}$$

where we've define $\Delta^{F'}$, the dual of Δ^F, as this quantity satisfied by another self-consistent equation:

$$\Delta^{F'}(\gamma) \equiv \text{tr}\left\{ \frac{\Sigma}{\gamma I_T + \frac{N\Sigma}{1 + \Delta^{F'}(\gamma)}} \right\}. \tag{154}$$

Now, to compute the data averages over x and $\widehat{x}$, we follow similar steps as before: define

$$\mathbb{q}(\gamma) \equiv \frac{1}{\gamma I_M + \left(\frac{N}{1 + \Delta^{F'}} \right) \frac{\sigma_u^2}{M} x x^T}, \tag{155}$$

which satisfies $\langle Q \rangle_u x^T = x^T \mathbb{q}$, and where we have used the definition of Σ, (47). The same manipulations as before, cf. (127), gives

$$\mathcal{L}_w = \frac{\sigma_w^2}{2\sigma_u^2 \widehat{T}} \left\langle \left[1 + \left(1 + \gamma \frac{\partial}{\partial \gamma} \right) \Delta^{F'} \right] \left(-\gamma^2 \frac{\partial}{\partial \gamma} \text{tr}\{\widehat{\Sigma} \mathbb{q}\} \right) \right\rangle_{x,\widehat{x}}. \tag{156}$$

Now, note that the only test-set dependence is in $\widehat{\Sigma}$, since $\Delta^{F'}$ is independent of $\widehat{x}$. This means we can easily perform the average over the test data $\widehat{x}$:

$$\mathcal{L}_w = \frac{\sigma_w^2}{2M} \left\langle \left[1 + \left(1 + \gamma \frac{\partial}{\partial \gamma} \right) \Delta^{F'} \right] \left(-\gamma^2 \frac{\partial}{\partial \gamma} \text{tr}\{\Lambda \mathbb{q}\} \right) \right\rangle_{x}, \tag{157}$$

where to get this we used definition (47) and the latent data statistics (19).

Finally, the x average proceeds now the same as the u average did before. Starting from (157) and taking the ridgeless limit:

- The factor in the first square brackets simplifies as before using

$$\lim_{\gamma \to 0} \left(1 + \gamma \frac{\partial}{\partial \gamma} \right) \Delta^{F'}(\gamma) = \Delta_0^{F'}; \tag{158}$$

this is quantity independent of x and in the overparameterized regime is given by

$$\Delta_0^{F'} = \frac{1}{N/T - 1} , \tag{159}$$

cf. (131) and then exchange $N \leftrightarrow T$.

- The factor in the round brackets gives

$$\lim_{\gamma \to 0} \left(-\gamma^2 \frac{\partial}{\partial \gamma} \right) \Delta^{L'} = \Delta_{-1}^{L'} , \tag{160}$$

where $\Delta_{-1}^{L'}$ is defined by the solution to the non-linear equation

$$1 \equiv \mathrm{tr} \left\{ \frac{\Lambda}{\Delta_{-1}^{L'} I_M + T\Lambda} \right\} . \tag{161}$$

Thus, the final answer for the label term takes a very similar form to the underparameterized regime, with the number of random features exchanged with the training set size $(N \leftrightarrow T)$:

$$\mathcal{L}_w = \frac{\sigma_w^2}{2M} \left(\frac{\Delta_{-1}^{L'}}{1 - T/N} \right) . \tag{162}$$

3.2.3 The Result

Assembling our results from the different regimes together, (144)and (162), we find for the label term:

$$\mathcal{L}(N, T; M) \equiv \mathcal{L}_w = \frac{\sigma_w^2}{2M} \begin{cases} (1 - N/T)^{-1} \Delta_{-1}(N, M), & N < T, \\ (1 - T/N)^{-1} \Delta_{-1}(T, M), & N > T, \end{cases} \tag{163}$$

where $\Delta_{-1}(\mu, M)$, for $\mu = \{N, T\}$, is the solution of the non-linear equation

$$1 = \mathrm{tr} \left\{ \frac{\Lambda}{\Delta_{-1}(\mu, M) I_M + \mu\Lambda} \right\} . \tag{164}$$

This quantity, $\mathcal{L}(N, T; M)$, represents the test loss in the noiseless limit $(\sigma_\epsilon^2 \to 0)$ and is symmetric under the interchange of $T \leftrightarrow N$, a natural manifestation of the aforementioned training-data–random-feature duality. Moreover, this formula holds for a random feature linear regression trained on latent data with completely general covariance matrices Λ; at no point did we assume anything about the latent feature covariance, Λ, to derive this result. Finally, note that (163) is only valid in the regime of $N, T < M$; we will give a formula for the case of a small latent space when we discuss the breakdown of neural scaling laws in §4.1.

49

Importantly, the factors of $\Delta_{-1}(\mu, M)$ appearing in (163) have an intuitive explanation. In §3.2.2, we saw that $\Delta_{-1}(N, M)$ arose when we studied a resolvent $\mathbb{Q}$, whose average encodes the statistical properties of the matrix $(uv)^T uv$. This is an $M \times M$ matrix that lives in the latent space and essentially describes the "lift" of our empirical covariance $\Omega = uvv^T u^T$, cf. (120), from the N-dimensional feature space to M-dimensional latent space. Since $M > N$ this lifted matrix is not full rank and (typically) has a null space of dimension $M - N$; this null space describes the latent features which our model cannot access. Indeed, the matrix given by the ridgeless limit of $\langle \mathbb{Q} \rangle_u$, (141), can be interpreted as the projector onto this null space. The result is that $\Delta_{-1}(N, M) \equiv \text{tr}\{\Lambda \lim_{\gamma \to 0} (\gamma \langle \mathbb{Q} \rangle_u)\}$ is the projection of the latent-space covariance matrix onto the null space. *In other words: $\Delta_{-1}(N, M)$ measures how much of the latent covariance matrix is missed because we have too few features.* Similarly, $\Delta_{-1}(T, M)$ measures how much of the latent covariance matrix is missed because we have too few samples. Both effects – poor performance due to too few samples and poor performance due to too few features – can occur in our model, but which one dominates in the ridgeless limit depends on whether $N < T$ or $N > T$.

Interestingly, even when the label noise is turned off, the factors $(1 - N/T)^{-1}$ and $(1 - T/N)^{-1}$ indicate a spectrum-independent *universal* noise-like divergence at the point of equiparameterization: $N \approx T$. This divergence occurs because our random features are still noisy samples of the underlying latent features, and our model drastically overfits to this kind of random noise at $N = T$. However, as we will discuss in a few paragraphs and in more detail in §4.2, this divergence is a relic of the ridgeless limit and disappears when the model is optimally regularized.

Now, in order to learn about neural scaling phenomenology, let's specialize to the case where the latent features have a power-law spectrum, (23). In this case, the self-consistent equations for Δ_{-1} can be easily solved numerically or approximated analytically, the details of which are given in §B.1, starting from (233). The result is (252):

$$\Delta_{-1}(\mu, M) = \begin{cases} \frac{\lambda_+}{M^\alpha} \left\{ k \left[\left(\frac{M}{\mu}\right)^\alpha - 1 \right] + [2 + \alpha(1 - k)] \left(1 - \frac{\mu}{M}\right) \right\}, & \mu < M, \\ 0, & \mu > M, \end{cases} \tag{165}$$

with the α-dependent constant k defined as

$$k \equiv \left[\frac{\frac{\pi}{1+\alpha}}{\sin\left(\frac{\pi}{1+\alpha}\right)} \right]^{1+\alpha}. \tag{166}$$

In Fig. 8, we plot numerical simulations of the test loss, (40) in a variety of ways to demonstrate the fit of our RMT result, (163) with the power-law spectral solution (165). The simulations were performed analogously to those shown in Fig 7: we generate latent training and test sets by sampling from our generative data model, (19) and (23); then, we map both sets through random feature models of different sizes, (28); finally, we use the linear regression solution (35) to compute test-set predictions using (39). The top left, top right, and bottom left panels of the figure demonstrate the close fit of our RMT result,

(163), in the limit of no regularization ($\gamma = 0$), and the top panels illustrate that the test loss of this statistical model *is* symmetric under the exchange $T \leftrightarrow N$, as we expected from our discussion of the duality (147).[46]

From the top two panels of the figure, we can see that our model captures the phenomenological behavior of the LLM test loss, e.g., Fig. 1. On the one hand, taking either T or N (and M) large in (165) and plugging into (163), we find a *power law scaling law*, cf. (2) and (3):

$$\mathcal{L}(N) \equiv \lim_{T,M \to \infty} \mathcal{L}(N, T; M) \sim N^{-\alpha}, \tag{167}$$

$$\mathcal{L}(T) \equiv \lim_{N,M \to \infty} \mathcal{L}(N, T; M) \sim T^{-\alpha}. \tag{168}$$

On the other hand, once the scaled resource parameter exceeds the large-but-fixed parameter, we see that the test loss asymptotes to a *plateau* that depends on that fixed parameter, cf. (4) and (5):

$$\mathcal{L}_{\text{plateau}}(N) \equiv \lim_{N \gg T} \mathcal{L}(N, T; M) = \frac{\sigma_w^2}{2M} \Delta_{-1}(T, M), \tag{169}$$

$$\mathcal{L}_{\text{plateau}}(T) \equiv \lim_{T \gg N} \mathcal{L}(N, T; M) = \frac{\sigma_w^2}{2M} \Delta_{-1}(N, M). \tag{170}$$

Finally, the bottom right panel of Fig. 8 simulates a model that is optimized over different values of the ridge parameter ($\gamma = \gamma^\star$), similar to the regularization used in Ref. [18] that led to the phenomenological model (1) and the test loss curves in Fig. 1. In this panel, we compare the version of this model appropriate to our setup ($\alpha_T = \alpha_N$, $N_0^\alpha = T_0^\alpha$), cf. (41),

$$\mathcal{L}_{\text{reg}}(N, T) = L_0 \left(\frac{1}{N} + \frac{1}{T} \right)^\alpha, \tag{171}$$

with an improved model that incorporates the size of the latent space,

$$\mathcal{L}_{\text{reg}}(N, T; M) = L_0 \left[\left(\frac{1}{N} + \frac{1}{T} \right)^\alpha - \left(\frac{1}{M} \right)^\alpha \right], \tag{172}$$

with the constant L_0 extracted from our $\gamma = 0$ result as

$$L_0 \equiv \frac{\lambda_+ \sigma_w^2 k}{2M}, \tag{173}$$

and k defined above in (166). We see that as the number of features approaches the size of the latent space ($N \to M$), the initial fit, (171), ignores the finiteness of the latent space and begins to diverge from the simulations, while the improved fit, (172), continues to fit the experiments closely. This breakdown of the empirical model from Ref. [18] suggests that for LLMs observed in practice, we still haven't seen effects resulting from finite M. As LLMs become even bigger, it would be interesting to try to fit (172) and measure the value of M for natural language data.

---

[46]However, this exact duality is broken in the presence of noise, as per our result for the noise term of the loss, (86).
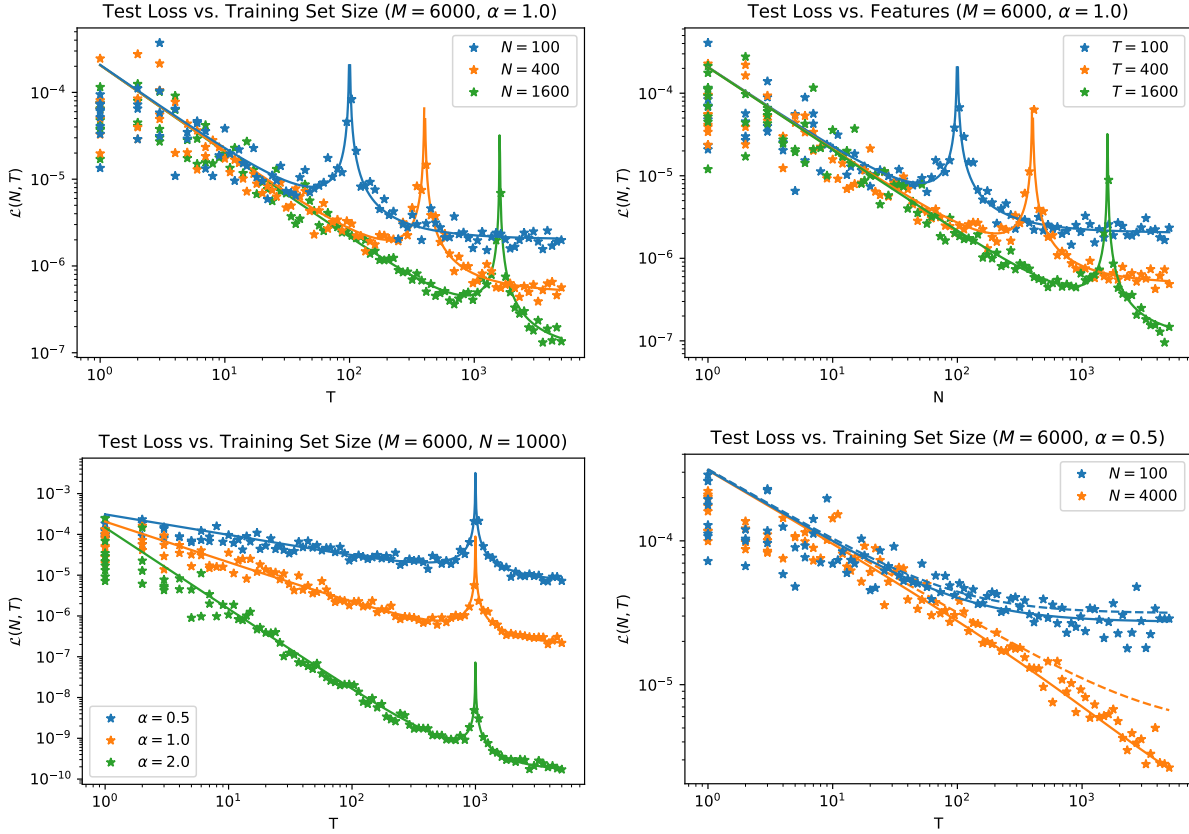
Figure 8: Test loss from numerical simulations (stars) of our joint statistical model ($\sigma_u^2 = 1$, $\lambda_+ = 1$, $\sigma_w^2 = 1$, $\sigma_\epsilon^2 = 0$) compared to the very close fit (solid lines, top and bottom left panels) of the result from our RMT calculation, (163). For the top and bottom left panels, no regularization is used ($\gamma = 0$), while for the bottom right panel the ridge parameter is optimized ($\gamma = \gamma^\star$). The top panels illustrate the symmetry of the test loss under the exchange $T \leftrightarrow N$. (As before, for smaller values of T, N, the variance of any particular realization is large, and so we've plotted multiple simulations for $T, N \lesssim 10$.) **Top Left:** The size of the training set, T, is varied for a few different sized models, N, while the size of the latent space and the power-law exponent is held fixed ($M = 6000$, $\alpha = 1.0$). **Top Right:** The number of features, N, is varied for a few different training set sizes, T, while the size of the latent space and the power-law exponent is held fixed ($M = 6000$, $\alpha = 1.0$). **Bottom Left:** The size of the training set, T, is varied for a few different power-law exponents, α, while the size of the latent space and the size of the model is held fixed ($M = 6000$, $N = 1000$). **Bottom Right:** When the ridge parameter is optimized ($\gamma = \gamma^\star$), the original phenomenological model, (171) is a poor fit (dashed lines) as the number of features approaches the size of the latent space ($N \to M$), but our improved model that incorporates the size of the latent space, (172), is a much better fit (solid lines).

52

3.3   (Towards) Modeling Spectral Extension

In our statistical model, the fact that our N random features were linearly-related to the M latent features made it possible to understand how the scaling laws worked analytically. However, our minimal statistical model is really an approximation of a more complicated phenomenon: rather than having a large latent space ($M > N$) with a *linear* transformation, in a more realistic model the random features are instead *nonlinear* functions of the smaller N_{in}-dimensional input space ($N_{\text{in}} < N$).

In §2.2, we explained that the essential property of such nonlinear feature maps – which include one-hidden-layer ReLU networks, (13), or the transformer networks that power LLMs – is the way they *extend* the portion of the spectrum that's approximately fit by a power law. From this we observed that the minimum of the data set size, T, and the number of features N, controlled the extent of the power law; the fact that the statistical model we described in §3.1 also had a power law controlled by the minimum of these two resource scales is what enabled it to capture the behavior we wanted to model. In a more realistic model, we can view the large number of useful features that we access from increasing T and N as being decoded from a non-linear transformation of the input features.

Thus, we would like to better understand how a nonlinear map for generating random features can extent the power law in the spectrum of input data. To be more concrete, let's consider a simple one-hidden-layer network with a nonlinear activation as in (13):

$$\varphi_{j;\alpha} \equiv \sigma\left(\sum_{k=1}^{N_{\text{in}}} u_{jk} x_{k;\alpha}\right) , \tag{174}$$

with σ a scalar activation function that acts on each individual component of the input data x_k, and $u \equiv u_{jk}$ is a $N \times N_{\text{in}}$ matrix of random feature weights drawn from a zero-mean Gaussian:

$$\langle u_{jk} \rangle = 0 , \qquad \langle u_{j_1 k_1} u_{j_2 k_2} \rangle = \frac{\sigma_u^2}{N_{\text{in}}} \delta_{j_1 j_2} \delta_{k_1 k_2} . \tag{175}$$

As long as the input dimension is sufficiently large, these random features will be approximately Gaussian,[47] and due to that Gaussianity this nonlinear model is fully-described by the same random matrix techniques described in the preceding section.[48] In particular,

---

[47] By adding an additional layer or by considering preactivations instead of activations we can remove the requirement of a large input dimension.

[48] The fact that one can use RMT tools to to study nonlinear random feature maps was a key point in Ref. [34]. Note also that these techniques go beyond other approaches in the literature, such as Ref. [59]: while the RMT techniques used here are "linear" in so far as they treat the features as Gaussian, they still use the full *nonlinear* kernel, Σ. If you further linearized the calculation of Σ in terms of the inputs, as was done in [59], you would not see the power-law extension we are interested in finding; this further approximation is excellent at studying the largest eigenvalues, but fails to adequately capture the spectral density for small eigenvalues.

the spectrum of a fixed dataset can be computed from the u-averaged resolvent:[49]

$$\langle Q(\gamma) \rangle_u = \frac{1}{\gamma I_T + \frac{N\Sigma}{1+\Delta^{F'}}} , \tag{176}$$

with the quantity $\Delta^{F'}$ the solution to the self-consistent equation,

$$\Delta^{F'}(\gamma) \equiv \mathrm{tr}\left\{ \frac{\Sigma}{\gamma I_T + \frac{N\Sigma}{1+\Delta^{F'}(\gamma)}} \right\} , \tag{177}$$

cf. (153) and (154), and Σ the data-data covariance matrix of the random features,

$$\Sigma \equiv \frac{1}{N}\left\langle \varphi^T \varphi \right\rangle_u \quad \Longleftrightarrow \quad \Sigma_{\alpha_1 \alpha_2} \equiv \frac{1}{N}\sum_{j=1}^{N} \left\langle \varphi_{j;\alpha_1} \varphi_{j;\alpha_2} \right\rangle_u . \tag{178}$$

Importantly, when the features are Gaussian this means that the spectrum is solely determined from knowledge of Σ.[50]

As noted in [34] and in many other places in the literature, the kernel of a single-hidden-layer network, Σ, is explicitly calculable and well-known for many choices of activation function as a function of the input data. As an explicit example, for the ReLU activation (14), one has that

$$\Sigma_{\alpha_1 \alpha_2} = \frac{\sigma_u^2}{2\pi}\|x_{\alpha_1}\|\|x_{\alpha_2}\| \left(\angle(\alpha_1,\alpha_2)\,\mathrm{acos}\left(-\angle(\alpha_1,\alpha_2)\right) + \sqrt{1-\angle(\alpha_1,\alpha_2)^2} \right) , \tag{179}$$

where the norms and inner products are defined as

$$\|x_\alpha\|^2 \equiv \sum_{k=1}^{N_{\mathrm{in}}} x_{k;\alpha}^2 , \qquad \angle(\alpha_1,\alpha_2) \equiv \frac{1}{\|x_{\alpha_1}\|\,\|x_{\alpha_2}\|}\sum_{k=1}^{N_{\mathrm{in}}} x_{k;\alpha_1} x_{k;\alpha_2} . \tag{180}$$

If we could find the spectrum of this infinite-random-feature kernel, (179), for a given data model or natural dataset, we could use that in conjunction with the resolvent, (176), to learn about the spectral extension for the finite-random-feature kernel evaluated on power-law data. Such a rigorous analytical derivation would be interesting to work out and perhaps give insights into in the interaction between datasets and activation functions.

---

[49]This was explained in footnote 26: the eigenvalue density can be determined in terms of the discontinuity across the branch cut along the real axis.

[50]In the context of the theoretical literature [10–16], this matrix is often known as the *kernel*. Sometimes it's helpful to think of this as arising from a random feature model with infinite features ($N \to \infty$).

3.4 Comparing to Other Methods

In this section, we compare our results to other RMT derivations in the machine learning literature, in particular to Refs. [34] and [35, 36]. In particular, we discuss the uses and limitations of our Feynman diagrammatic methods and contrast with other methods.

First, random matrix theory techniques used in this paper have many similarities with the approach taken in [34], which was an influential starting point for the present work. Our intermediate result, (152), is essentially the same result that appears as the random feature-averaged test loss in Conjecture 1 of their paper, up to combining some terms with the use of our averaging over the random teacher weights. While our diagrammatic derivation may then appear less general than [34] because of the random teacher assumption for the labels, our derivation nevertheless can be used in either setting.

In particular, in following our derivation in §3, one can see that the covariance matrices appearing at the beginning of each term in (146) could be substituted for arbitrary matrices without changing the structure of the following calculation. The formula from Ref. [34], without the random teacher averaging, then follows (up to trivial constants) by replacing

$$\widehat{\Sigma} \to \widehat{y}^T \widehat{y}, \qquad \widetilde{\Sigma} \to y^T \widehat{y}, \qquad \Sigma \to y^T y, \tag{181}$$

at the beginning of each term and following these substitutions through the calculation of the random feature u average. The resulting expression is then consistent with what appears in [34].

Nevertheless, there's also some notable differences between our methods and results. First, our expression is somewhat simpler than the one that appears in [34] due to the fact that we rewrote the prefactor in terms of a differential operator. Second, we also believe our diagrammatic techniques give a more transparent and extensible methodology for computing similar averaged expressions. Finally, our work generalizes the result in [34], both by deriving a dual expression for averaging over training samples instead of random features, (118), and by calculating the test loss when both random feature and data averages are taken at the same time, (163). This final *double* average is entirely novel to our work and does not otherwise appear in the literature.

Relatedly, our work also concerns similar types of averages to Refs. [35, 36]. In these papers, the authors used replica methods to compute the expected test loss of a kernel regression. One way to interpret their work is in terms of a single average over the data x. We find similar results using our methods in Appendix A where we study (non-generalized) linear models and compute a single x average over the model's features, x_j, since, in these models, there are no random features to average over. Although they are in different forms, our computation, (216), matches the main result of Ref. [36], cf. their Eq. 4.

One reason we were able to simplify the calculation of their very nice result is, by rewriting the loss in terms of differential operators, we immediately see that the test loss is determined entirely by the averaged quantity $\Delta \equiv \mathrm{tr}\{\Lambda \overline{q}\}$. The computation of this quantity for a general spectrum Λ is a classic result of random matrix theory that goes back to the original work of Marchenko and Pastur [51] and has been studied or rediscovered

many times since. In particular, this quantity Δ, or equivalently, the x averaged resolvent, $\bar{q}$, can be computed by many means, e.g., by the Feynman diagrammatic approach of §3.2.1 and, of course, also by the replica methods in Refs. [35, 36].

On the one hand, we think that the computation using Feynman diagrams has a few advantages over replica techniques: in addition to being more straightforward, one also does not need to prove that there is a unique analytic continuation in the replica index. Moreover, although we have no need for it here, the Feynman diagrammatic techniques can be easily adapted to compute subleading $1/N$ corrections to averages, if desired.

On the other hand, even for Gaussian distributions, some quantities cannot be readily computed using Feynman diagrams. This generally occurs when one computes the expectation value of a non-analytic function, such as the logarithm of a function of φ.[51] In such a case, one can use replica methods.[52]

On our third and final hand, other more advanced techniques to study matrix integrals at large N include the use of loop equations and the method of orthogonal polynomials. These techniques are often necessary to treat more complicated (non-Gaussian) probability distributions.[53] These methods allow one to effectively sum up infinite classes of Feynman diagrams but are typically overkill for Gaussian matrix expectations. To our knowledge, these techniques have not yet been used in a machine learning context.

4 Discussion of Results

Now that we have a statistical model of scaling laws that we understand for jointly large-but-finite model size N, training set size T, and latent space size M, in this section we have a discussion of what we can learn from our statistical model of scaling laws.

In §4.1, we interpret our results from §3 in the limit that the model size or training set size approaches the size of the latent space, $N, T \sim M$, and the neural scaling law phenomenology of §1 breaks down.

In §4.2, we discuss how the regime of scaling laws, of large training set size and large model size, pushes resource efficient and properly regularized models towards the equiparameterization regime, and how the phenomenon of double descent is not really relevant for such models.

In §4.3, we try to reconcile the large latent space, M, required for datasets that allow for neural scaling laws with the traditional idea that input datasets are embedded in high-dimensional spaces, N_{in}, and can be compressed to a latent space with much smaller

---

[51]This occurs in the information theory context in the study of von Neumann entropies and in the physics context in the study of quenched disorder.

[52]This requires one to instead consider n copies of the original random variable and compute the expectation value of a logarithm by analytically continuing in n using $\log z = \lim_{n \to 0} (z^n - 1)/n$. This method only works for certain classes of probability distributions (and quantities being averaged) such that they are so strongly constrained that there is a unique analytic continuation in n.

[53]For a useful review of these methods, see [60].

intrinsic dimension, d_{in}. We note that there are a number of notions of dimensionality, and the particular power-law structure of the datasets that give rise to scaling laws makes different notions meaningful for different questions.

Finally, in §4.4, we identify some limitations of our generative data model that could be improved in future analyses.

4.1   The Breakdown of Neural Scaling Laws

In this subsection, we would like to understand whether the empirically-derived phenomenological model of the test loss, (1), can break down.[54] In particular, we would like to know *(a)* when is such a formula no longer predictive of performance and *(b)* what is the behavior in the new regime?

To give insight into the answer of these questions, we can use our statistical model from §3. One hint towards an answer is that model (1) only depends on two of the scales in the problem, N and T, while the formula we derived, (163), also depends on size of the latent space of the data distribution, M. In particular, in our regularized ($\gamma > 0$) simulations of the test loss of our model (Fig. 7), to fit (1) it was important that the size of the latent space be the largest scale in the problem, $N, T \ll M$. Here we will explore what happens when this condition is violated.

To proceed, it is convenient to organize our discussion around both whether or not there's noise added to the training labels and whether or not there's a nonzero ridge parameter regularizing the regression.

No noise

First, let's consider the case without noise. In the left panel of Fig. 9 we plot a simulation of our statistical model with no regularization ($\gamma = 0$) and also plot our RMT calculation of the test loss, (163), for models for which the number of features is much larger than the size of the latent space and the number of samples in the training set ($N > T, M$); in the right panel, we use optimal regularization ($\gamma = \gamma^\star$) in the simulation and also plot a new fit that we will discuss below. In both panels, we learn that "breakdown" of scaling laws – without noise – is a lot like a *singularity*: all of the sudden at $T = M$ the test loss drops very rapidly to zero!

To understand these results, let's turn to our calculation: if we consider a model in the large-latent-space regime, $T, M \ll N \to \infty$, from (163) and (165) we see that the test loss

---

[54]In some parts of the literature, e.g. in [37], "breakdown of scaling laws" is used to mean the transition from the power law to the plateau in the test loss; as should be clear, this is not what we mean here.
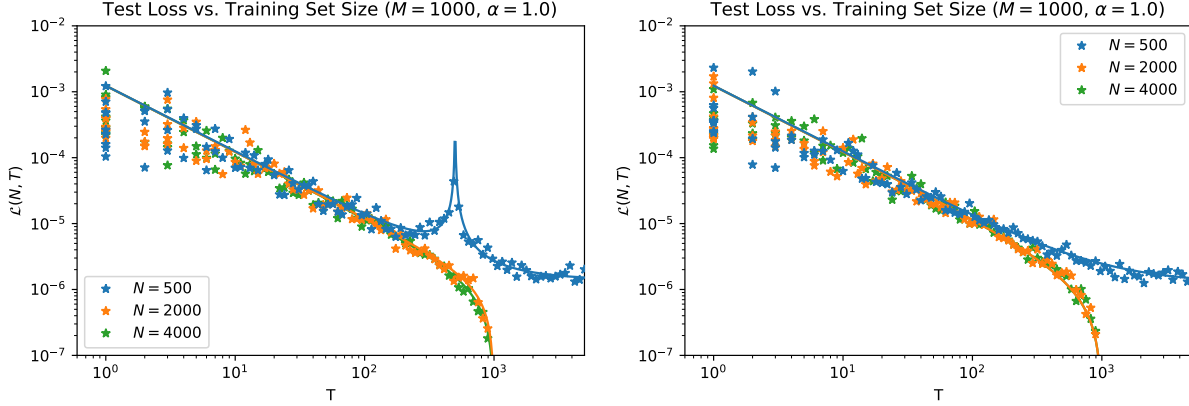
Figure 9: Simulation (stars) and fit (solid) of the scaling-law model from §3 with exponent $\alpha = 1$, latent-space size $M = 1000$, and the following entirely inconsequential choices: $\lambda_+ = \sigma_w^2 = \sigma_u^2 = 1$. The orange and green curves represent models in different large-latent-space regimes $(T, M < N)$ that demonstrate novel behavior, while the blue curves represent models in the usual regime $(T, N < M)$, including both power law and plateau behavior, and is included for purposes of comparison. As the size of the training set is increased to the size of the latent space, the test loss very rapidly approaches zero. **Left:** No regularization is used $(\gamma = 0)$, and the fit is the result of our RMT calculation (163). **Right:** The ridge parameter is optimized $(\gamma = \gamma^\star)$, which is intended to model the early stopping regularization used in Ref. [18], and the semi-heuristic fit is given by (185).

of our model simplifies as

$$\lim_{N \to \infty} \mathcal{L}(N, T; M) = \frac{\sigma_w^2}{2M} \Delta_{-1}(T, M) \tag{182}$$
$$= \begin{cases} \frac{\sigma_w^2 \lambda_+}{2M^{\alpha+1}} \left\{ k \left[\left(\frac{M}{T} \right)^\alpha - 1 \right] + [2 + \alpha(1 - k)] \left(1 - \frac{T}{M} \right) \right\}, & T < M, \\ 0, & T > M. \end{cases}$$

Recalling our analysis in §B.1, we see that the *scaling* regime controls the test loss when the training set is smaller than the size of the latent space, $T \ll M$, and we find the same power law scaling law as before:

$$\lim_{T \ll M} \lim_{N \to \infty} \mathcal{L}(N, T; M) \sim T^{-\alpha}. \tag{183}$$

However, when the size of the training set approaches the size of the latent space, $T \to M$, the test loss is instead controlled by *coincident* regime, and we have

$$\lim_{T \nearrow M} \lim_{N \to \infty} \mathcal{L}(N, T; M) = \frac{2 + \alpha}{2} \left(\frac{\sigma_w^2 \lambda_+}{M^{\alpha+1}} \right) \left(1 - \frac{T}{M} \right), \tag{184}$$

which vanishes as T approach M from below.

58

Adding and optimizing the ridge parameter doesn't really change this as regularization only really matters around the point of equiparameterization ($T \approx N$), and here we're interested in the limit of $N \to \infty$. For this reason, it's easy to construct the fit to the simulated in the right panel of Fig. 9. The blue curve is included for comparison and shows a similar regime as in the experiments plotted in bottom right panel of Fig. (8) ($N < M$), for which we are using the improved fit (172). For the new large-latent-space regime (orange and green curves) with $N > M$, we can simply use the unregularized large feature formula, (182). Altogether, we can write our no-noise optimal regularization fit to the simulations in this plot as

$$\mathcal{L}_{\text{reg}}(N, T; M) = \frac{\sigma_w^2 \lambda_+}{2M} \begin{cases} k \left[\left(\frac{1}{N} + \frac{1}{T} \right)^\alpha - \frac{1}{M^\alpha} \right] & T, N \leq M\,, \\ k \left(\frac{1}{T^\alpha} - \frac{1}{M^\alpha} \right) + \frac{1}{M^\alpha} \left[2 + \alpha(1-k) \right] \left(1 - \frac{T}{M} \right)\,, & T \leq M < N\,, \\ 0\,, & T, N > M\,. \end{cases}$$
(185)

In other words, in the right panel of Fig. 9, we use the top line for the blue curve, cf. (172), and the middle and bottom lines for the green and orange curves.

Finally, given the duality that exchanges the training set size and number of features, $N \leftrightarrow T$, similar formulas will hold in a regime where we have a super-sized training set, $T > M$: in the unregularized case as we increase the model size we will see behavior described by (182) but with the swapping of number of features with the size of the training set, $N \leftrightarrow T$, and an analogous statement will hold in the case of regularization and (185).

Noise

Now, let's turn on the label noise ($\sigma_\epsilon^2 > 0$). In addition to the label term that made up the loss immediately above, we now also must include the contribution from the noise term, (86), when $N > M$:

$$\mathcal{L}_\epsilon = \frac{\sigma_\epsilon^2}{2} \begin{cases} \alpha + \frac{1}{M/T - 1}\,, & T < M\,, \quad N > M\,, \\ \frac{1}{T/M - 1}\,, & T > M\,, \quad N > M\,. \end{cases}$$
(186)

In the left panel of Fig. 10, we see that the inclusion of this term gives an excellent fit to noisy numerical simulations.

Although we are still unable to analytically optimize over the ridge parameter, γ, we can again write a phenomenological formula for the regularization of the noise term by finding a simple solution that interpolates between the asymptotics of the above formula. Such a formula could be

$$\frac{\sigma_\epsilon^2}{2} \left[\frac{1}{M/(T + \alpha M) + T/M} \right]\,.$$
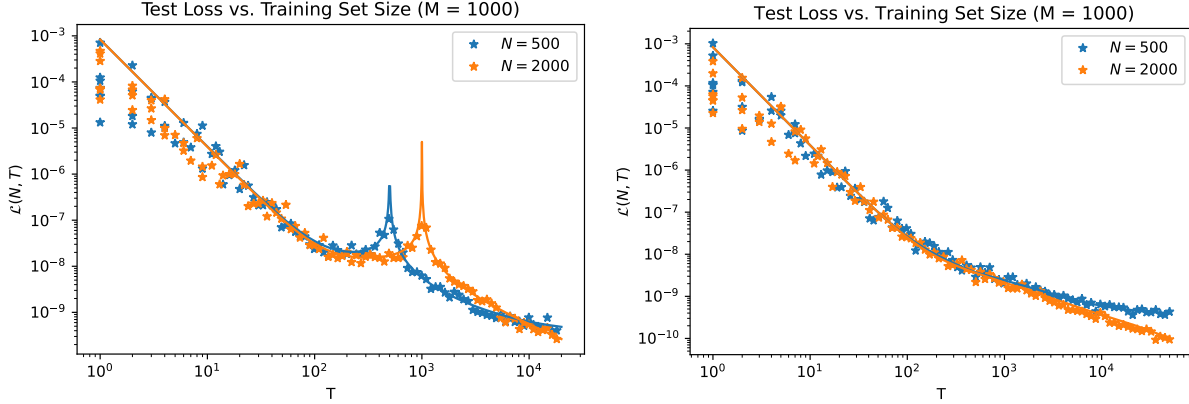(187)

Figure 10: Simulation (stars) and fit (solid) of the scaling-law model from §3 with exponent $\alpha = 2.33$, latent-space size $M = 1000$, $\sigma_\epsilon^2 = 10^{-8}$, and the following choices that are made so that there's good signal to noise in the data: $\lambda_+ = \sigma_w^2 = \sigma_u^2 = 1$. The orange curves represent models in different large-latent-space regimes ($M < N$) that demonstrate novel behavior, while the blue curves represent models in the usual regime ($N < M$), including both power law and plateau behavior, and is included for purposes of comparison. **Left:** No regularization is used ($\gamma = 0$) and the fit is our RMT calculation $\mathcal{L}(N, T; M) + \mathcal{L}_\epsilon$, cf. (182) and (186). **Right:** Optimal regularization ($\gamma = \gamma^\star$) and the fit is to our phenomenological formula $\mathcal{L}_{\text{reg}}(N, T; M) + \mathcal{L}_{\epsilon,\text{reg}}$, cf. (185) and (188), shows the transition from the power-law regime with $\mathcal{L} \sim T^{-\alpha}$ scaling to the noise-dominated $\mathcal{L} \sim T^{-1}$ scaling (orange) and to resource-bottlenecked plateau (blue) for the large-latent-space regime and the usual regime, respectively.

However, we can do slightly better when both terms in the denominator are of comparable size by normalizing with a factor of $1/2$ using a bump function:

$$\mathcal{L}_{\epsilon,\text{reg}} = \frac{\sigma_\epsilon^2}{2} \left[\frac{1}{M/(T + \alpha M) + T/M} \right] \left[\frac{1}{1 + e^{-(M/T + T/M)/c}} \right] , \qquad (188)$$

for some constant c that we can choose to fit the experiment. We see in the right panel of Fig. 10 that this formula, in conjunction with the label term (185), gives an excellent fit to the simulations, although we emphasize that the exact functional form near equiparameterization is chosen in a rather ad hoc way and should not be understood to be the definitive functional form of the optimized solution.

Most notably, for large T we see that both the unregularized case, (186), and the regularized case, (188), there's a *universal* $\sim 1/T$ falloff of the test loss when the model size and training set size jointly exceed the size of the latent space ($N, T > M$). If such a transition in powers appears in a model at an otherwise undefined scale, it could be suggestive of a breakdown associated with having reached the size of the latent space.[55]

---

[55]However, note that this universal $\sim 1/T$ behavior depends on the label noise not being correlated

Interestingly, the consequences of this for the practitioner now depend on the size of the power-law exponent α: for $\alpha > 1$, this transition would limit the model's performance gains with increasing resources N and T, while for exponents $\alpha < 1$, it would enhance such gains.

4.2 The Battle of the Parameterizations

Despite common wisdom coming from classical statistical intuition, it's been long observed in the modern deep learning setting that very large models trained to near zero training error can generalize well without *overfitting*, even with minimal to no regularization [17, 61–66].[56] This behavior is a feature of the overparameterized regime, in which the number of parameters of the model is greater than the size of the training set, $N > T$. In this regime, inference is said to succeed by *interpolating* between the learned training points, and the performance of such fully-trained models will continue to improve as model sizes increase for a fixed training set.

In contrast, in the underparameterized regime in which the size of the training set exceeds the model size, $T > N$, the **bias-variance tradeoff** of classical statistics dictates performance of under-regularized models (see, e.g., [9]). For a fixed training set, a model may *underfit* if it's too small, exhibiting a *bias* as it's not expressive enough to learn the target function; it may overfit if it's too big and unhelpfully learns the noisy fluctuations in the training set due to the *variance* between different realizations of the data: thus, there's an extremal size balancing the bias and variance at which performance peaks.

Typically, this point is only a local minimum of the test error, and – for a fixed training set size – increasing the model size into the overparameterized regime can eventually lead to models that outperform the optimal underparameterized model. This overall description of the test error as a function of parameters has been dubbed **double descent** [17] and a large amount of effort has been spent on understanding the setting for and the mechanism of this behavior.[57]

In Fig. 11 we exhibit this phenomenon (solid blue curve) based on (163), the now-familiar result of our RMT calculation of the test loss of our statistical model presented in §3.[58] The plot depicts the performance of our model for a fixed training set size ($T = T_0$) in the ridgeless limit ($\gamma = 0$): in the underparameterized region to the left of the *equiparameterization* peak at $N = T$, the test loss has a local minimum ($N = N_\star$), which

---

from sample to sample, as in our noise model (34); it's not obvious to us whether this is an appropriate assumption for self-supervised generative modeling tasks.

[56]Such regularization includes both *explicit* regularization, such as a nonzero ridge parameter γ, or *implicit* regularization, such as early stopping the gradient-based optimization of an objective function.

[57]See, e.g., [66] and additional references therein.

[58]Technically, we should distinguish between the test loss and the generalization error (the difference between the test loss and training loss) in the underparameterization regime since the model can underfit. However, this difference is not important for the qualitative point we want to make here.

optimally trades off the bias of too small a model with the variance of too large a model; in the overparameterized region to the right of the peak, the test loss continues to improve past the local minimum at $N = N_\star$ after overcoming the non-analytic peak, though this improvement is slow and asymptotic.[59]
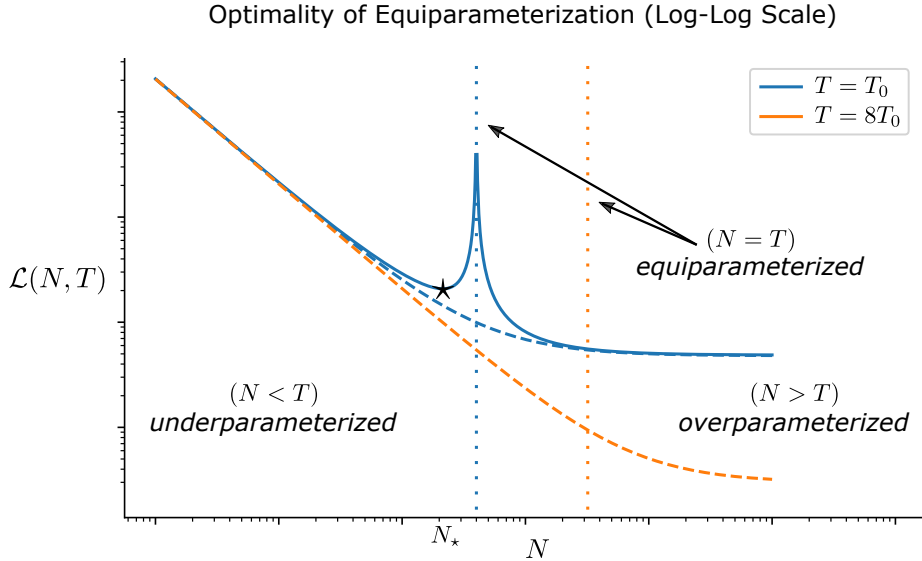

Figure 11: Sketch of test losses of our statistical model from §3 on a log-log scale for different fixed training set sizes for both unregularized or ridgeless (solid) and optimally regularized (dashed) models. The solid blue curve exhibits the double descent phenomenon, with a local minimum of performance in the underparameterized region (black star, $N = N_\star$) and with performance further improving asymptotically in the overparameterized region. The two blue curves illustrate how the double-descent peak is an artifact of the ridgeless limit ($\gamma = 0$), with performance monotonically improving through the point of equiparameterization (vertical dotted lines) when the models are properly regularized. Comparison of the dashed blue ($T = T_0$) and orange ($T = 8T_0$) curves illustrates the optimality of near-equiparameterization when using regularization properly: the best performance boost results from scaling the model jointly with the size of the training set ($N \sim T$).

However, from the LLM scaling law phenomenology of [18] and the explicit realization of the test loss (1) by our statistical model, we know that the classical "U" shape of the underparameterization region can be modified with proper regularization such that the test loss can be made analytic, monotonically decreasing as models become larger.[60] The dashed curves in Fig. 11 – generated from (172) – depict this behavior, showing that

---

[59]Of course, from our discussion in §4.1, we know that there eventually will be a qualitative change to this behavior when the model size and the training set size exceed the size of the latent space, $N, T > M$.

[60]The observation that double descent can be mitigated by optimal regularization was first made by [67].

optimally regularized models exhibit monotonically decreasing test loss with model size and do not have any special behavior around the point of equiparameterization (vertical dotted lines). Further, with this regularization it's apparent that the slow asymptotic improvement for very large models at fixed training set size is simply the plateau region of (1). Furthermore, comparison of the unregularized and regularized models – from our numerical simulations and exhibited in the figure by the blue solid and blue dashed curves, respectively – we see that regularization only seems to be important around the equiparameterization peak.[61] From this perspective, the double descent phenomenon is an artifact of not using regularization in a small region around $N \sim T$.

Thus, when using proper regularization, we can achieve the best test loss by *jointly* scaling the sizes of the model and the size of the training set. Comparison of the two regularized (dashed) curves in Fig. 11 illustrates the way that increasing model size alone eventually encounters a plateau, while increasing the size of the training set extends the power-law portion of the performance gains. As originally pointed out by [18], the reason for the optimality of this joint scaling is avoiding resource bottlenecks: each additional sample in the training set is informative about one additional eigen-feature in the power-law portion of the spectrum of data's latent spectrum (23), and we need an additional feature in our model to represent that eigen-feature.[62]

As such, given finite resources and an ability to both scale models as well as gather training points, an optimal allocation involves a kind of joint near-equiparameterization scaling: for generalized linear models such as our statistical model, there is only a single exponent, α, that controls the test loss power-law behavior in both the number of features and size of the training set, and for this model scaling the number of features of the model to equal the size of the training set, $N(T) = T$, will avoid the plateau region; for other models such as the LLMs discussed in [18], we may have a more general scaling relation, $N(T) \sim T^p$, as in (6).[63] Even in this more general case, jointly scaling both training data and model size pushes the performance away from the tails of the test loss curves and back towards the termination of the power-law region – back towards the non-analytic peak of the unregularized model – which is the region of large-data and large-parameter *equiparameterization*. Altogether, we conclude that this regime with proper regularization – and not the overparameterization regime – is the practical setting of interest for deep learning.[64]

---

[61]A related point: early stopping, which can implicitly play the same role of a ridge parameter [68, 69], only helps models in the regime of near-equiparameterization [66].

[62]Cf. our discussion of PCA in §2.1: if the latent data had a gap instead of a continuous spectrum without a natural cutoff, then the performance would not continue to improve substantially by jointly increasing model and training-set sizes.

[63]However, recall that [8] gives evidence that the scaling for LLMs should actually be linear, $N(T) \sim T$, if trained sufficiently. Note also that this doesn't necessarily mean that the number of parameters must equal the size of the training set, but only the relative ratio should always remain constant.

[64]Interestingly, a curated dataset of machine learning systems taken from highly-cited and highly-

4.3 The Battle of the Dimensions

Typically, we expect that natural data lives on a data manifold of smaller **intrinsic dimension**, d_{in}, than its embedding dimension, N_{in}. The representation that lives on the latent data manifold is supposed to be a type of encoded representation of the input, and we imagine there exists a transformation that decompresses the latent representation into the higher-dimensional embedding or input space, where, for example, the pixels of an image would live. Usually we are interested in the inverse of such a transformation, which can be used as a dimensionality-reduction technique for data.[65]

However, this seems to be in tension with the equiparameterization picture for natural data discussed in the last subsection: the power-law scaling of the test loss arises from a model learning increasingly more eigen-features from the latent-space representation of the data. For this to work, the **latent-space dimension**, M, must be the largest scale in the problem: $N_{\text{in}}, N, T \ll M$. In this picture, rather than decompressing data, the transformation from the latent representation to the input representation is a type of *projection*, and many samples are needed to reconstruct the true underlying description.

If we accept both of these pictures, then it must be the case that the intrinsic dimension, d_{in}, is a different quantity than the size of the latent space, M.

Defining and Measuring Intrinsic Dimension

There are multiple sensible ways to estimate the intrinsic dimension, d_{in}, of a dataset [72]. A nice method considers the typical (Euclidean) distance between neighboring points [73, 74]: if we define $x'(x_\alpha)$ to be the nearest neighbor (NN) of a sample, x_α, then the NN-distance, δ_α, is defined by

$$\delta_\alpha^2 \equiv \sum_{i=1}^{N} \left[x_i'(x_\alpha) - x_{i;\alpha} \right]^2 \, , \tag{189}$$

and the typical distance $\langle \delta \rangle$ in a dataset of T points is given by the sample average,

$$\langle \delta \rangle \equiv \frac{1}{T} \sum_{\alpha=1}^{T} \delta_\alpha \, . \tag{190}$$

So, if we increase the density of points, we expect the average NN-distance to decrease:

$$\frac{T}{V} \sim \langle \delta \rangle^{-d_{\text{in}}} \, , \tag{191}$$

---

influential papers from 1952 to 2021 [70] gives strong evidence that skilled practitioners have always been implicitly working in this jointly large-training-set-and-large-model-size equiparameterized regime: plotting the parameter counts vs. training set sizes of the models in this dataset on a log-log scale gives a linear fit with a slope extremely close to unity (see, [71]). We thank Ben Adlam for bringing this point to our attention.

[65] A very simple example of such a technique is PCA.

where V is the volume of the space. We will use this equation to *define* the intrinsic dimension d_{in}: intuitively, the power that relates the linear-dimensional NN-distance, $\langle \delta \rangle$, to the density of points T/V, is operationally what we mean by the dimension of a space.[66] Moreover, this definition, (191), gives a simple way to measure d_{in}: plot the scaling of the average NN-distance versus the size of the dataset and look at its slope on a log-log plot.[67]

Natural (Power-Law) Data

Now, let's apply this definition to the types of natural datasets we've considered in this paper. Recall that such datasets have power-law spectra (§2.1), and we expect that models trained on such data will have power-law scalings for their test losses when not otherwise bottlenecked by the number of parameters. Recalling this power law, (168), and rearranging our definition (191), we note curiously that both the test loss and the NN-distance are power laws in the size of the training set:

$$\langle \delta \rangle \sim T^{-1/d_{\text{in}}}, \qquad \mathcal{L}(T) \sim T^{-\alpha}. \tag{192}$$

Naively, there's no reason to relate these two quantities. However, Ref. [37] argued that the test loss should be monotonically related to $\langle \delta \rangle$, the typical linear size of a subregion occupied by a training point: the smaller the typical region, the greater that any test point will fall near a training point and inference can succeed via interpolation, the smaller the test loss.[68] With this proposal, we relate the two equations in (192) and identify

$$d_{\text{in}} = \frac{\#}{\alpha}, \tag{193}$$

where the order-one numerical factor $\#$ can depend on the details of the manifold [37].[69]

To verify this relationship, we can use the scaling of the NN-distance with the dataset size, (191), to measure d_{in} for our generative data model and see how it varies with the (hyper-)parameters of the distribution: α and M. In Fig. 12, we numerically simulate latent datasets of different sizes according to (19) and (23), varying the power-law exponent α (left panel), and the size of the latent space M (right panel). These experiments support the relationship (193) between d_{in} and α and further suggest that d_{in} is entirely unrelated to M. Thus, it must be the case that, while both dimensions d_{in} and M characterize the latent data manifold, they represent very different intrinsic properties.

---

[66]For instance, if we place points uniformly on a Euclidean lattice, $\langle \delta \rangle$ would be the fixed lattice size, d_{in} would be the usual Euclidean dimension of the lattice, and the density of points would be exactly $T/V = \langle \delta \rangle^{-d_{\text{in}}}$.

[67]There are much more sophisticated ways of measuring d_{in}, see, e.g., [73, 74, 37], but (191) is sufficient for the qualitative point we wish to make here.

[68]The connection between dataset size, the intrinsic dimension, and scaling laws was also made by [41].

[69]The numerical factor suggests that $\mathcal{L}(T) = \langle \delta \rangle^p$, for some power p.

Figure 12: Typical NN-distance from numerical simulations (stars) of our latent data generative model, with the maximum eigenvalue fixed ($\lambda_+ = 1$), fit (dashed lines) by a simple power law, $\langle \delta \rangle \sim T^{-\alpha/12}$. These experiments support the identification of the intrinsic dimension with the inverse of the power-law exponent: $d_{\text{in}} = \#/\alpha$. **Left:** The size of the dataset, T, is varied for a few different power-law exponents, α, while the dimension of the latent space is held fixed ($M = 3000$). **Right:** The size of the dataset, T, is varied for a few different latent-space dimensions, while power-law exponent is held fixed ($\alpha = 1.0$). These experiments suggest that d_{in} does not vary with M.

To understand this better, recall our discussion of PCA in §2.1: when the spectrum of the data has a gap, the few principal components above the gap are taken to span the intrinsic data manifold, the other components are thought to be random noise, and the number of principal components gives an estimate of d_{in}.[70] However, for continuous spectra arising from natural data and modeled by our latent data generative model, there's no natural cutoff for separating informative and uninformative features. Instead, the exponent characterizing the spectrum, α, provides a way of estimating d_{in}: intuitively, if the exponent is large, $\alpha \gg 1$, then the spectrum very rapidly decreases to zero, and much of the variance of the data is captured by the first few principal components. Thus, we might expect that there's a way to compress such data onto a manifold with small intrinsic dimension $d_{\text{in}} \sim 1/\alpha$.[71]

To conclude, on the one hand, for power-law data there's no natural cutoff in the spectrum, and you can always do fractionally better by modeling additional eigen-features: this is why there are neural scaling laws. On the other hand, although such data has M

---

[70]Note, PCA can overestimate the intrinsic dimension of manifolds with large curvature [72, 37].

[71]Note that in our work, α and M are independent (hyper-)parameters of our latent data model: α can be any positive real number, and so we can create datasets with any d_{in}. In contrast, the scaling exponents in LLMs [18] and images [26, 41] are small, $\alpha < 1$, but large enough such that the intrinsic dimension is smaller than the input dimension ($d_{\text{in}} < N_{\text{in}}$). It's unclear what the interpretation would be of a natural dataset if α is so small such that $d_{\text{in}} > N_{\text{in}}$.

independent components, the power-law structure is a type of constraint: the data manifold is not very typical among M-dimensional manifolds, and the intrinsic dimension, d_{in}, is determined by α, not M.

Latent Features vs. Trained Features

The argument of Ref. [37] that led to the identification (193) implicitly relied on not just the input data, but also the labels: in principle, only the features relevant for the underlying task will be useful for reducing the test loss. To account for that, the authors actually considered the intrinsic dimension of the final-layer activations of a trained network.

In general, we should expect our estimate of d_{in} to depend on whether we consider trained features or raw input data (or the latent features of our statistic model): if the task is designed so that the label depends on only a few of the input features, then the trained network will learn to ignore the rest of the features, and we might expect the trained features to have a much smaller intrinsic dimension. For example, consider input features that represent the N_{in} pixels of a black and white image: if the label is simply whether the first pixel is on or off, then we should expect a much smaller intrinsic for the trained features than if the label had depended on correlations among the N_{in} pixels of the input, as in a typical image classification task; in the former case, we might expect $d_{\text{in}} \sim 1$, while in the latter case, if there's $C \ll N_{\text{in}}$ classification categories, we might instead expect $d_{\text{in}} \sim C$.

In our statistical model, there are two reasons that make this distinction between latent features and trained features unimportant: *(1)* our labels, y, depend on *all* latent features of an input, x, cf. (26); and *(2)* our model uses *random* features – not trained features – and those random features cannot depend on any label information.[72] This first point is necessary to see the full set of observed neural scaling law phenomenology – the power-law and plateau regions – and so this suggests it's an important property of self-supervised generative modeling tasks, such as those that LLMs perform. As to the second point, it would be interesting to find a statistical model for which the features can learn nontrivial representations of the inputs based on the tasks. We discuss some ideas towards this end at the end of §5.

4.4   The Breakdown of our Data Model

The main goal of this paper has been to show how a simple generative data model combined with a simple random feature linear regression model could capture the rich neural scaling phenomenology observed in a broad class of real world large-scale models. An essential part of our data model was that the spectrum of the input data's covariance matrix had many decades of power-law eigenvalues: when we jointly increased the number of

---

[72]Additionally, we expect that the N random features would have a similar intrinsic dimension as the M latent features, given experiment in the right panel of Fig. 12. showing that the size of the latent space doesn't affect the intrinsic dimension.

random features or training samples, we were able to see a larger and larger fraction of the underlying latent eigenvalues in the empirical covariance matrix; this ultimately led to a power-law scaling of the test loss when neither features nor training data were constrained. If the data model were any simpler – e.g., that purely Gaussian data with an isotropic spectrum that we study in §A.1 – we would not find power-law scaling of the test loss. However, there are certain limitations to the simplicity of the data model, which we will discuss in the rest of this subsection.

In particular, one might wonder whether we can use our random feature and generative data model from §3 as a model for *input data*:

- If we take a dataset sampled from this random feature model (28) and interpret the size of the random feature space as the data's input dimension ($N \to N_{\text{in}}$), and if we then act on the dataset with a *nonlinear* random feature map, e.g. (174), with dimension $N_{\text{ReLU}} > N$, will the power law of the feature spectrum extend?

- If the spectrum does extend, will it also extend the power law in the test loss?

To answer these questions, we can numerically study our data model with a single-hidden-layer ReLU network and compare against a natural dataset.

In Fig. 13 we answer this first question. In these plots we show the spectrum of random features for our data model (left panel) with the input features given by an explicit power law and for a common computer vision (CV) dataset (right panel). To generate these plots, we used the explicit form of the single-layer-ReLU-network kernel, (179) and then explicitly diagonalized the kernel matrix. Recalling that such a kernel can be thought of as an infinite-dimensional feature model ($N_{\text{ReLU}} \to \infty$), we can thus study a fixed set of input features $N_{\text{in}} = 784$ and see how the spectrum extends as we vary the size of the dataset, T. As we see from the figure, the answer to our first bulleted point above is *yes*: our data model behaves identically to the natural CV dataset. In particular, since the number of features is implicitly infinite and not a bottleneck, we see that the extent of the power law is determined by the size of the limiting resource T. This property is quite universal, and we have observed it across a wide range of natural datasets and activation functions.

In both panels of this figure, we also see another interesting effect, one that was already pointed out in footnote 18 and present in the right panel of Fig. 4: when data with a power-law spectrum passes through the ReLU layer – for both our data model and natural datasets – the exponent α seems to decrease slightly. Since this exponent also characterizes the power-law behavior of the model's test loss, this decrease ultimately leads to worse performance for a given set of resources than if it were otherwise preserved. It would be interesting to use our RMT tools, e.g. as discussed in §3.3, to try to understand the origin of this effect.

Next, in Fig. 14, we answer the second bulleted question above. To generate plots, we map the N_{in} input features through a single-hidden-layer ReLU network with a much larger number of random features ($N_{\text{ReLU}} > N_{\text{in}}$), and then we train a linear regression
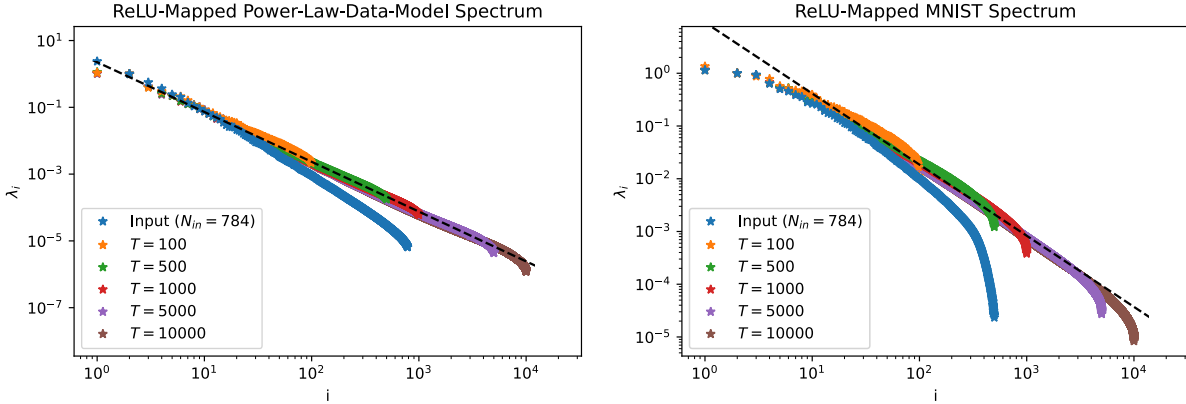
Figure 13: Spectra of the activations of a single-hidden-layer random ReLU network obtained by diagonalizing the data-data covariance matrices, Σ, defined in (179), which can be thought of as an infinite-dimensional feature model ($N_{\text{ReLU}} \to \infty$). In both panels, we see that the power-law portion of the spectrum extends beyond N_{in} as we increase the dataset size, T; we also see in both cases that passing the data through the ReLU feature map *decreases* the exponent α of the power law fit. (Note also that the spectra have been rescaled by their largest eigenvalue so that the curves lie on top of one another.) **Left:** The input dataset is obtained by sampling latent data points, $x_{I;\alpha}$, with statistics (19) and power-law eigenvalues (23) ($M = 20000, \alpha = 1, \lambda_+ = 1$) and then subsampling these latent features with the linear map (28) ($\sigma_u^2 = 1$) to obtain input features with $N_{\text{in}} = 784$. **Right:** The input dataset is obtained from samples of MNIST dataset [75], a CV dataset of natural black and white images with input dimension $N_{\text{in}} = 784$.

model on these random features and plot the test loss. In the left panel, we see the results for our random-feature-and-label model from §3.1. Note that the performance plateaus as soon as the number of training samples, T, surpasses the size of the input dimension, N_{in}, rather than at the much larger random-feature dimension, N_{ReLU}: in other words, our random feature model interpreted as a data model does not itself give a power-law scaling of the test loss when the spectral extension is due to a nonlinear network. This means that the answer to our second bulleted point above is *no*: despite the fact that the spectrum extends when passed through the ReLU (cf. Fig. 13), this extension is simply not enough to translate to extending the power law scaling of the test loss.

For the sake of comparison, in the right panel of Fig. 14 we compare to the natural CV dataset trained on a simplified classification problem. We can see clearly that, while our model has a plateau at N_{in}, the natural dataset extends the power-law scaling of the test loss past N_{in}, indicating that the spectral extension *was* translated to the test loss.[73]

---

[73]N.B. in this plot, the particular maximum value of T was chosen because there are only 12,214 images in the training set of MNIST for the "7" and "9" digits; we expect that the power-law (orange curve)
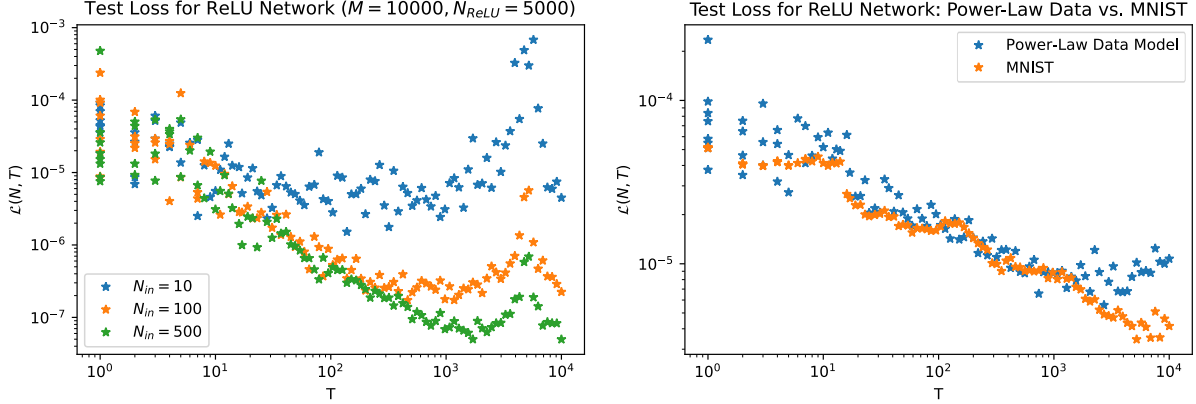
Figure 14: Model performance for our data model, with a comparison to MNIST, when the random features are the activations of a single-hidden-layer ReLU network of width $N_{\text{ReLU}} > N_{\text{in}}$. **Left:** Test loss from numerical simulations of our data model with $M = 10000$, $\alpha = 1.5$, and $\lambda_+ = 1$ as a function of training set size T: first, we generate latent data samples, $x_{I;\alpha}$, with statistics (19) and power-law eigenvalues (23) and subsample them with the linear map (28), with $\sigma_u^2 = 1$, to obtain input features with different N_{in}; then, we pass the inputs through a single-layer random ReLU network of width $N_{\text{ReLU}} = 5000$; finally, we compute the test loss via (40) with our generated labels (26). We see that the power-law scaling ends and the loss plateaus after the size of the training set, T, reaches the number of input features, N_{in}, rather than the number of ReLU features, N_{ReLU}. **Right:** We compare our data model (blue) with $M = 20000$, $\alpha = 0.4$, $\lambda_+ = 1$, $\sigma_u^2 = 1$, and input features of size $N_{\text{in}} = 784$, against a simplified MNIST classification problem (orange), where we only include 7s and 9s. In both cases, we pass the data through a single-hidden-layer ReLU network of width $N = 20000$ and compute the test loss via (40). We see that while our data model plateaus for training set sizes T greater than N_{in}, the MNIST performance continues to follow a power-law falloff up to higher scales. (Note: the MNIST loss is rescaled to lay on top of the power-law data's loss, and the exponent of our data model was chosen to match that of MNIST in order to facilitate the comparison.)

A heuristic explanation for this discrepancy is clear: when we down-sample the latent space to generate the random features in (28), we are throwing away many of the latent eigenvalues *and eigenvectors* present in the latent data x_I necessary to predicting the label y_I, cf. (26). While the nonlinear feature map (174) does extend the power-law eigenvalues (cf. Fig. 13), there is no reason to expect these nonlinear features to be aligned with the latent eigenvectors as they were determined at random. In other words, this information was fundamentally lost when our down-sampling projected out these directions. Hence, the N_{in} inputs do not generate any more useful features when passed through a ReLU network.

However, unlike our generative data model, natural datasets *do* seem to contain sufficient information in their limited input features to determine the outputs: seemingly, very little has been thrown away! Thus, it would be interesting to extend our minimal model to understand further how the nonlinear activation functions are able to decode relevant eigenvectors and extend the power-law regime in the test loss.[74]

5 Conclusion and Featured Directions

In this paper we presented a solvable model of neural scaling laws. After exploring natural datasets and their behavior under nonlinear feature maps (§2), we constructed a *minimal* statistical model that contains the full set of neural scaling behavior exhibited by large language models in practice (§3). We also explained why our model is minimal by showing that even simpler data or feature models (§A) do not have the right behavior. By solving our statistical model using techniques from random matrix theory (§3.2), we were able to better understand the mechanism driving these practically relevant neural scaling laws.

With our solution, we were then able to explore a regime of our model ($M < N, T$) that goes beyond the regimes explored experimentally for LLMs (§4.1). This is a nice benefit of theoretical analyses such as ours; in addition to trying to understand experimentally observed phenomena in contexts that we've accessed experimentally, we can also use our models to probe new contexts. Relatedly, now we can ask the experimentalists, can we measure the scale M for LLMs? Perhaps we can use the fact that the prediction of the test loss fit is better when you include M, as per our analysis, cf. (172)? Knowledge of the size of the latent space for language would have practical benefit by elucidating, in some sense, how much better LLMs can get using power-law performance gains before transitioning into a noise-dominated regime.

Finally, we hope that the large-N diagrammatic approach to RMT calculations (§3.2) and the differential operator approach to simplifying expressions (§A) will be useful theoretical tools for simplifying machine learning calculations.

<div align="center">♦♦♦♦</div>

---

would continue further if more training samples were available.

[74]Note that Ref. [40] has a toy model of scaling laws that incorporates an eigenvector dependence.

We close by featuring a future direction.

Representation Learning?

In the paper, we have studied generalized linear regression models with random feature maps, and we have seen that the ability of such maps to extend the power law present in the spectrum of input data is the mechanism that leads to the power law scaling of the test loss as the number of random features increases. A natural extension of this analysis is to move beyond linear regression with random feature maps and consider models that *learn* representations of their inputs. A natural setting for this is the **quadratic models** framework of [16], as they are minimal models of representation learning.[75]

A quadratic model extends the generalized linear model, (30), in the following way:

$$z(x;\theta) \equiv \sum_{j=1}^{N} \theta_j \varphi_j(x) + \frac{\varepsilon}{2} \sum_{j_1,j_2=1}^{N} \theta_{j_1} \theta_{j_2} \psi_{j_1 j_2}(x) \,, \tag{194}$$

where each $\psi_{j_1 j_2}(x)$ is a **meta feature function**, which is a (potentially) nonlinear transformation of an input x that is distinct from the feature functions $\varphi_j(x)$, ε is an adjustable hyperparameter that controls the importance of the quadratic deformation of the linear model, and, for convenience, we assume that the model outputs a scalar z. Due to the constraint tying the parameters in the linear term, θ_j, to the product of parameters in the quadratic term, $\theta_{j_1} \theta_{j_2}$, the quadratic model cannot be written in the form of a generalized linear model (30); it is this nonlinear-in-the-parameters property that enables it to learn nontrivial representations of inputs x over the course of training.[76]

Following §11.4 of [16], if we optimize the MSE loss, (33), with the quadratic model, (194), we can find a solution in perturbative expansion in the deformation parameter ε:

$$\theta^\star \equiv \theta^{\mathrm{L}} + \varepsilon\, \theta^{\mathrm{NL}} \,, \tag{195}$$

where the *linear* part, θ^{L}, is the linear regression solution with $\varepsilon = 0$, (35), and the *nonlinear* part is given in terms of the linear solution to leading order as

$$\theta_j^{\mathrm{NL}} = -\frac{1}{2} \sum_{j_1,j_2,j_3=0}^{N} \theta_{j_1}^{\mathrm{L}} \theta_{j_2}^{\mathrm{L}} \left[\sum_{\alpha=1}^{T} \psi_{j_1 j_2}(x_\alpha)\, \varphi_{j_3;\alpha} \right] q_{j_3 j} + O(\varepsilon) \,, \tag{196}$$

and would have a more complicated expression were we to substitute in θ^{L}. With this solution, test predictions can similarly be decomposed as

$$\widehat{z}^\star \equiv \widehat{z}^{\mathrm{L}} + \varepsilon\, \widehat{z}^{\mathrm{NL}} \,, \tag{197}$$

---

[75]As we discussed in our Aside: *Random Feature Maps vs. DNNs* in §2.1, certain DNNs in the infinite-width limit correspond to linear regression models. In contrast, nonlinear models, of which *quadratic models* are the simplest example, can instead correspond to more realistic finite-width networks [16].

[76]See §11.4 of [16] for the details of how this works.

72

with $\widehat{z}^{\mathrm{L}}$ the standard *linear* regression inference, (39), and $\widehat{z}^{\mathrm{NL}}$ the correction coming from the *nonlinear* quadratic term.

Now to the point: if we substitute these predictions into the test loss,

$$\mathcal{L}_{\mathcal{B}}(\theta^\star) = \frac{1}{2\widehat{T}} \left\| \widehat{z}^\star - \widehat{y} \right\|^2 \,, \tag{198}$$

we can try to evaluate the performance of the quadratic model given a particular data model. In particular, we can evaluate if the nonlinear contribution – and therefore if nontrivial representation learning – improves the model's performance by computing

$$\frac{d\mathcal{L}_{\mathcal{B}}(\theta^\star)}{d\varepsilon} = \frac{1}{\widehat{T}} \sum_{\beta=1}^{\widehat{T}} \widehat{z}_\beta^{\mathrm{NL}} \left(\widehat{z}_\beta^{\mathrm{L}} - \widehat{y}_\beta \right) + O(\varepsilon) \,. \tag{199}$$

This quantity captures the linear shift in the model's performance due to the nonlinear deformation; here we see that it depends on how nonlinear part of the test prediction, $\widehat{z}_\beta^{\mathrm{NL}}$, aligns with the linear part of the prediction error, $\left(\widehat{z}^{\mathrm{L}} - \widehat{y} \right)$, when averaged over the test set. If these quantities are anti-aligned for a given data model,

$$\frac{d\mathcal{L}_{\mathcal{B}}(\theta^\star)}{d\varepsilon} < 0 \,, \tag{200}$$

then we can conclude that the quadratic model improves the performance on that data, at least for small enough values of ε.[77] The qualification "for a given data model" is important, as it's easy to construct simple datasets that give either sign for $d\mathcal{L}_{\mathcal{B}}(\theta^\star)/d\varepsilon$!

To assess a realistic setting of interest, we can use our generative data model from §3, modeling the data with a power law covariance spectrum. Then, we can use our RMT tools to find the average of (199) over the data and random features.[78] Altogether this may lead to a more complex, if not even more realistic, model of neural scaling laws: the form of the quadratic model, (194), breaks the data-parameter duality in the average test loss; perhaps because of this we may find a much richer phenomenology for test loss, e.g. unequal scaling exponents α_N, α_T as was originally observed in [18].

Acknowledgments

---

[77] If we wanted to go further and find the optimal value of the hyperparameter $\varepsilon = \varepsilon^\star$, we'd need to compute $\theta^\star$ to order ε^2 in order to get a nontrivial equation when solving $d\mathcal{L}_{\mathcal{B}}(\theta^\star)/d\varepsilon = 0$.

[78] To do so, we would need to determine the statistics of ψ; as per (11.152) in [16], ψ should have a nontrivial cross correlation with φ. (N.B. the hyperparameter ε is defined differently in that reference.)

A Linear Models

In §3, we studied a joint generative data model and random feature model that minimally captured the rich neural scaling phenomenology present in the test loss of LLMs. In this appendix, we will study some even simpler models that capture some, but not all, of that behavior. These models are instructive in explaining the origin of different parts of the test-loss curve by highlighting the different properties of a statistical model that lead to their absence or presence. Although these models have been described in detail in literature, here we will attempt to give slightly different derivations in order to connect to the methods used the main body of this paper as well as to illustrate the simplicity of our approach.

In particular, in this appendix we will study (non-generalized) linear models for different generative data models. This means that we will learn a model of the form

$$z_i(x; \theta) \equiv \sum_{j=1}^{N} \theta_{ij} x_j \,, \tag{201}$$

where $\theta \equiv \theta_{ij}$ is a set of learnable parameters, and the input data x lives in an N-dimensional feature space with statistics

$$\langle x_j \rangle = 0 \,, \qquad \langle x_{j_1} x_{j_1} \rangle = \Lambda_{j_1 j_2} \,. \tag{202}$$

Our analysis will proceed with general covariance matrices Λ, and we will later evaluate our solution for different data models. Analogously, our labels are defined as

$$y_i = \sum_{j=1}^{N} w_{ij} x_j, \quad \text{with} \quad i = 1, \dots, C \,, \tag{203}$$

where $w \equiv w_{ij}$ is an C-by-N-dimensional weight matrix whose elements we will take to be independent and drawn from a zero-mean Gaussian distribution:

$$\langle w_{ij} \rangle = 0 \,, \qquad \langle w_{i_1 j_1} w_{i_2 j_2} \rangle = \frac{\sigma_w^2}{N} \delta_{i_1 i_2} \delta_{j_1 j_2} \,. \tag{204}$$

Comparing (201) and the rest of this setup with the generalized linear model, (30), in the main text, the main difference here is that there's no separate latent feature space,

$M \to N$, and no random feature model, $\varphi(x) \to x$. Importantly, unlike the labels in the main text, (26), here the complexity of the label, (203), increases as we increase the number of features in the linear model: this means that we have no way to increase the capacity of our model without also making the task harder. On the plus side, the analysis will be much simpler since we will no longer have to average over random features, and the only nontrivial average will be over the training set x.

Using the same linear regression setup and same notational conventions discussed in §3.1, we can write down a training loss with a ridge parameter,

$$\mathcal{L}_{\mathcal{A}}(\theta) \equiv \frac{1}{2} \|\theta x - y - \epsilon\|^2 + \frac{\gamma}{2} \|\theta\|^2 ,$$

where the random noise is defined the same as before (34),

$$\langle \epsilon_{i;\alpha} \rangle = 0, \qquad \langle \epsilon_{i_1;\alpha_1} \epsilon_{i_2;\alpha_2} \rangle = \sigma_\epsilon^2 \delta_{i_1 i_2} \delta_{\alpha_1 \alpha_2} , \tag{205}$$

and straight away use the well-known linear-regression solution for the parameters (35):

$$\theta^\star \equiv (y + \epsilon) x^T q , \tag{206}$$

with the feature-feature resolvent now defined as

$$q \equiv \frac{1}{\gamma I_N + x x^T} . \tag{207}$$

Finally, using this solution for inference on our test set $\widehat{x}$ as

$$\widehat{z}^\star \equiv \theta^\star \widehat{x}, \tag{208}$$

we define the test loss,

$$\mathcal{L}_{\mathcal{B}}(\theta^\star) \equiv \frac{1}{2\widehat{T}} \|\widehat{z}^\star - \widehat{y}\|^2 = \frac{1}{2\widehat{T}} \left\| (wx + \epsilon) x^T q \widehat{x} - w \widehat{x} \right\|^2 , \tag{209}$$

where on the final expression we've substituted in the test predictions, (208), the optimal parameters (206), and the labels, (203). Our goal now is to average $\mathcal{L}_{\mathcal{B}}(\theta^\star)$ over the training inputs x, the test inputs $\widehat{x}$, the label weights, w, and the random noise ϵ. This calculation is analogous to the one studied in Refs. [35, 36] via replica methods.

The first part of the calculation proceeds identically to the first steps in §3.1 under the subheading *Average Goals*: since they are both quadratic, the average over the random noise, cf. (43), and the subsequent average over the label weights, cf. (44), can be performed trivially. Thus, our nontrivial point of departure is the noise-and-label-averaged test loss expression:

$$\langle \mathcal{L}_{\mathcal{B}}(\theta^\star) \rangle_{\epsilon,w} = \frac{C \sigma_w^2}{2\widehat{T} N} \left\| x x^T q \widehat{x} - \widehat{x} \right\|^2 + \frac{C \sigma_\epsilon^2}{2\widehat{T}} \left\| x^T q \widehat{x} \right\|^2 , \tag{210}$$

75

which we again will refer to as the *label term* and *noise term*, respectively. As in the main text, from this point forward we will set $C = 1$ as it's just a trivial rescaling.

Now, we will need to recall our simple identities, (59) and (60), which we will reprint here for convenience:

$$(\gamma I_N + xx^T)q = I_N , \qquad q(\gamma)^2 = -\frac{\partial}{\partial \gamma}q(\gamma) . \tag{211}$$

Rearranging the first identity, note that we can rewrite the label term as

$$\left\| xx^T q\widehat{x} - \widehat{x} \right\|^2 = \gamma^2 \left\| q\widehat{x} \right\|^2 , \tag{212}$$

which eliminates the explicit dependence on x. Next, using (202), let's perform the average over the test set. Since both terms are quadratic in $\widehat{x}$, and remembering the test set and training set don't correlate since each sample is drawn independently, we find

$$\langle \mathcal{L}_{\mathcal{B}}(\theta^\star) \rangle_{\epsilon, w, \widehat{x}} = \frac{\sigma_w^2 \gamma^2}{2N} \mathrm{tr}\{\Lambda q^2\} + \frac{\sigma_\epsilon^2}{2}\left(\mathrm{tr}\{\Lambda q\} - \gamma \mathrm{tr}\{\Lambda q^2\} \right) , \tag{213}$$

where to simplify the noise term we again made use of a rearrangement of the first identity in (211) to eliminate xx^T. While in principle the training set average of $\mathrm{tr}\{q^2 \Lambda\}$ would be complicated given the q^2, luckily we have our second identity in (211) which lets us write the test loss as

$$\langle \mathcal{L}_{\mathcal{B}}(\theta^\star) \rangle_{\epsilon, w, \widehat{x}} = \frac{1}{2}\left[-\frac{\sigma_w^2}{N}\gamma^2 \frac{\partial}{\partial \gamma} + \sigma_\epsilon^2\left(1 + \gamma \frac{\partial}{\partial \gamma} \right) \right] \mathrm{tr}\{\Lambda q\} . \tag{214}$$

This is determined entirely by the training set average of a single quantity: $\mathrm{tr}\{\Lambda q\}$.

Now, since this expression is linear in q, we can "take" the trainset average by making the replacement $q \to \overline{q}$. Doing so, we immediately recognize the resulting quantity,

$$\mathrm{tr}\{\Lambda \overline{q}\} \equiv \Delta(T, N) = \mathrm{tr}\left\{ \frac{\Lambda}{\gamma I_N + \frac{T\Lambda}{1 + \Delta(T,N)}} \right\} , \tag{215}$$

as the nontrivial object that appears in the coupled set of equations that determine $\overline{q}$. Here, as the right-most expression, we've written the self-consistent equation satisfied by Δ, cf. (82).[79] This means that we can express our final answer as a differential operator acting on this quantity:

$$\langle \mathcal{L}_{\mathcal{B}}(\theta^\star) \rangle_{\epsilon, w, \widehat{x}, x} = \frac{1}{2}\left[-\frac{\sigma_w^2}{N}\gamma^2 \frac{\partial}{\partial \gamma} + \sigma_\epsilon^2\left(1 + \gamma \frac{\partial}{\partial \gamma} \right) \right] \Delta . \tag{216}$$

---

[79]The calculation that gives the self-consistent equation for $\Delta(T, N)$ follows identically to the calculation in §3.2.1, with the following substitutions: $\varphi \to x$ and $\Omega \to \Lambda$. These reflect the fact that *this* resolvent is built from xx^T, cf. (207), and has statistics (202), respectively.

Finally, in the ridgeless limit, $\gamma \to 0$, we can expand Δ as

$$\Delta(T, N) \equiv \frac{\Delta_{-1}}{\gamma} + \Delta_0 + \Delta_1 \gamma + \dots , \tag{217}$$

cf. (232) in Appendix B, and associate these differentials to specific terms in the expansion:

$$\lim_{\gamma \to 0} \left(-\gamma^2 \frac{\partial}{\partial \gamma} \right) \Delta(T, N) = \Delta_{-1} , \qquad \lim_{\gamma \to 0} \left(1 + \frac{\partial}{\partial \gamma} \right) \Delta(T, N) = \Delta_0 . \tag{218}$$

Using these identities, the test loss takes an extremely simple form

$$\langle \mathcal{L}_\mathcal{B}(\theta^\star) \rangle_{\epsilon, w, \widehat{x}, x} = \frac{1}{2} \left[\frac{\sigma_w^2}{N} \Delta_{-1} + \sigma_\epsilon^2 \Delta_0 \right] , \tag{219}$$

and Δ_{-1} and Δ_0 can be very simply determined given an underlying data spectrum.

On the one hand, this solution for the linear model in the ridgeless limit is similar to the random feature model in the main text: in that case, the noise term was proportional to $\Delta_0^F \equiv \Delta_0^F(T, N)$, cf. (85), and the label term was proportional to $\Delta_{-1}(N, M)$ or $\Delta_{-1}(T, M)$ depending on whether the model was under- or overparameterized, cf. (163). On the other hand, a key difference is the lack of separation of scales between the number of random features, N, and the number of latent features, M: this means that there's no duality in the exchange of training set size and number of features, $N \leftrightarrow T$, no phase transition that exchanges $\Delta_{-1}(N, M)$ with $\Delta_{-1}(T, M)$ as the model crosses $T = N$, and, consequently, the test loss is not symmetric under $N \leftrightarrow T$. Instead, the linear model has a different phase transition solely in the function $\Delta_{-1}(T, N)$ at the point of $N = T$; this phase transition is analogous only to the other phase transition in the random feature model that occurs when T or N crosses M, cf., e.g., (182), and leads to different overall behavior.

To probe these properties of the linear model further, let's explicitly evaluate the ridgeless test loss, (219), for two different data models: first we'll consider completely structureless data with Marchenko-Pastur statistics (§A.1), and then we'll make a direct comparison to the main text by considering power-law data (§A.2).

A.1 Marchenko-Pastur Data

As the simplest possible data model, let's take the covariance of the input features to be isotropic,

$$\Lambda = \sigma_x^2 I_N , \tag{220}$$

which means that each component of the input is independent and identically distributed. The distribution of empirical covariance matrices of a finite-sized dataset is the Marchenko-Pastur (MP) distribution [51]. The test loss of linear models with MP data have been studied extensively in the literature.

For MP data, the calculation of Δ_{-1} and Δ_0 are both trivial because the sum over eigenvalues doesn't depend on their index. Starting with Δ_{-1}, plugging the expansion (217) into the self-consistent equation (215), converting the sum into a trace, and substituting in the spectrum (220), we have find a simple equation that determines Δ_{-1}:

$$\Delta_{-1} = \sum_{i=1}^{N} \frac{\Delta_{-1}}{T + \Delta_{-1}\sigma_x^{-2}} = \frac{N\Delta_{-1}}{T + \Delta_{-1}\sigma_x^{-2}} \, . \tag{221}$$

Solving this trivial equation gives two solutions, one valid for when the model when the model is under(over)-parameterized:

$$\Delta_{-1}(T, N) = \begin{cases} (N - T)\sigma_x^2, & T < N \, , \\ 0, & T > N \, . \end{cases} \tag{222}$$

The reason for this assignment is as follows: for $T < N$, the empirical covariance has zeros, which means that the resolvent, q, and its average, $\bar{q}$, must blow up as γ^{-1} as $\gamma \to 0$, cf. (207), and thus $\Delta_{-1}(T, N) \neq 0$; for $T > N$, the empirical covariance is full rank, there's no pole in q as $\gamma \to 0$, and thus $\Delta_{-1}(T, N) = 0$.

For calculation of Δ_0, we have to consider $T < N$ and $T > N$ separately. For $T < N$, starting with the $O(\gamma^0)$ term in the power-series equation of the self-consistent equation (215),

$$\Delta_0 = (1 + \Delta_0) \sum_{I=1}^{M} \frac{\mu\lambda_I^2}{(\mu\lambda_I + \Delta_{-1})^2} \, , \tag{223}$$

plugging in $\Delta_{-1} = (N - T)\sigma_x^2$ and $\Lambda = \sigma_x^2 I_N$, and performing the trivial sum, we find

$$\Delta_0 = (1 + \Delta_0)\frac{T}{N} \quad \implies \quad \Delta_0(T, N) = \frac{1}{N/T - 1} \, . \tag{224}$$

For $N < T$, our ansatz (217) instead should begin at $O(\gamma^0)$: plugging that expansion into the self-consistent equation (215), we find

$$\Delta_0 = \sum_{i=1}^{N} \frac{1 + \Delta_0}{T} = (1 + \Delta_0)\frac{N}{T} \quad \implies \quad \Delta_0(T, N) = \frac{1}{T/N - 1} \, . \tag{225}$$

Altogether, we have

$$\Delta_0(T, N) = \begin{cases} \frac{1}{N/T - 1}, & T < N \, , \\ \frac{1}{T/N - 1}, & T > N \, , \end{cases} \tag{226}$$

independent of the scale of the size of the "signal" σ_x^2.[80]

---

[80]It's interesting to compare this expression, (226), with our formula for Δ_0 with power-law data, which is computed in the next appendix, (262): as the first argument of Δ_0 approaches the second, these formula agree completely; however, the MP formula, (226), is *exact*, while the coincident regime of the power-law formula, (262), is only the leading part of the ε expansion. The reason for this agreement is that as two scales become equal, the tail of the spectrum of the empirical covariance is universally determined by MP statistics for any underlying Λ.

Since we've evaluated Δ_{-1} and Δ_0 exactly, we can write an *exact* formula for the test loss in the ridgeless limit ($\gamma \to 0$): trivially substituting (222) and (226) into (219), we get

$$\mathcal{L}_{\mathrm{MP}}(N,T) \equiv \langle \mathcal{L}_{\mathcal{B}}(\theta^\star) \rangle_{\epsilon,w,\widehat{x},x} = \frac{1}{2} \begin{cases} \sigma_w^2 \sigma_x^2 \left(1 - \frac{T}{N}\right) + \frac{\sigma_\epsilon^2}{N/T - 1}, & T < N, \\ \frac{\sigma_\epsilon^2}{T/N - 1}, & T > N. \end{cases} \tag{227}$$

In Fig. 15, we plot some example test loss curves using numerical simulations of MP data for different signal strengths σ_x^2 and compare with our analytical formula (227): considering both fixed number of features while varying the training set size (left panel) and fixed training set size while varying the number of features (right panel), we see an excellent fit to the experimental data.
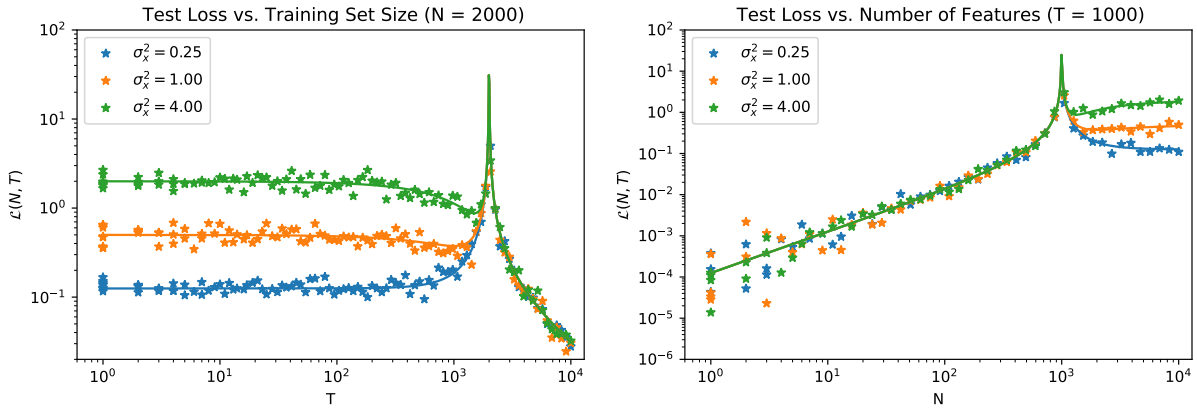
Figure 15: Test loss from numerical simulations (stars) and our analytical expression (solid lines) of a linear model trained on MP data ($\sigma_w^2 = 1$, $\sigma_\epsilon^2 = 0.25$). As the two panels show, the test loss of this model is not symmetric under the exchange $T \leftrightarrow N$. **Left:** The size of the training set, T, is varied for a few different sized signals, σ_x^2, while the number of features is held fixed ($N = 2000$); note that the linear decrease of the test loss $(1 - T/N)$ with increasing T is not easily perceptible on a log-log scale. **Right:** The number of features, N, is varied for a few different sized signals, σ_x^2, while the size of the training set is held fixed ($T = 1000$).

Overall, we can characterize the performance of the model according to its different regimes:

- When the model is *overparameterized*, $T < N$, the test loss decreases linearly towards zero as $(1 - T/N)$ with increasing T; each new training point gives the model access to an additional eigenvalue of the covariance matrix, each of which is equal size.

- At the point of *equiparameterization*, $T = N$, the performance formally diverges a universal way, overfitting to the nonzero noise. (In the noiseless limit, $\epsilon \to 0$, the test loss would be identically zero.)

- When the model is *underparameterized, $T > N$,* the only contribution to the test loss comes from the noise; this eventually falls off with increasing T as the noise is distributed independently for each training sample and ultimately averages away as $T \to \infty$.

This phenomenology is quite different than the random feature model from the main text, (163):

(i) The random feature model has a regime that exhibits power-law scaling law of the test loss. This cannot happen for MP data since there's no power-law in the spectrum to translate into the test loss; essentially, MP *data* is itself structureless noise that cannot capture the rich behavior of the LLMs.

(ii) (For small noise, σ_ϵ^2) the test loss of the random feature model in the main text was *symmetric* under the exchange of the training set size and the number of features $(T \leftrightarrow N)$, while the MP-data linear model is *not.* (This is clear from inspection of the analytical solution, (227), and illustrated by the two panels of Fig. 15.) The problem with the linearly model is that as we increase the number of features, N, we simultaneously increase the number of features used to determine the label, cf. (203). In other words, the learning problem of training a small number of features on a large training set is fundamentally easier than training many features on a small training set.

To fix the first problem, we need a more realistic generative data model (§A.2). To fix the second problem, we need to decouple the size of the model used in the regression from the complexity of the learning problem (§3).

A.2   Power-Law Data

To address that first problem, let's now use the same underlying data model as in §3: we'll assume the spectrum of Λ has a number density that's well-approximated by

$$ n(\lambda)d\lambda = N(\beta - 1)\lambda_-^{\beta-1}\lambda^{-\beta}\theta(\lambda - \lambda_-)d\lambda \,, \tag{228} $$

where here λ_- is the minimum eigenvalue, β is an exponent that characterizes the tail of the distribution, and the constants are chosen such that the density integrates to N. Recall from our earlier discussion that we can instead characterize the spectrum by a maximum eigenvalue, λ_+ ,and exponent, α, which are related to λ_- and β by, cf. (24) and (25),

$$ \lambda_+ \equiv \lambda_- M^{1+\alpha} \,, \qquad \alpha \equiv \frac{2 - \beta}{\beta - 1} \,, \tag{229} $$

with the range $1 < \beta < 2$ translated to $0 < \alpha < \infty$.

For power-law data, the calculations of $\Delta_{-1}(T, N)$ and $\Delta_0(T, N)$ are performed in Appendix B. Importing the results of those calculations, (252) and (262), and substituting

into our expression for the test loss, (219), we can write a formula for the test loss of the linear model:

$$\mathcal{L}_{\mathrm{PL}}(N,T) \equiv \langle \mathcal{L}_{\mathcal{B}}(\theta^\star) \rangle_{\epsilon,w,\widehat{x},x} \qquad (230)$$

$$= \frac{1}{2} \begin{cases} \frac{\lambda_+ \sigma_w^2}{N^{\alpha+1}} \left\{ k \left[\left(\frac{N}{T} \right)^\alpha - 1 \right] + [2 + \alpha(1-k)] \left(1 - \frac{T}{N} \right) \right\} + \sigma_\epsilon^2 \left(\alpha + \frac{1}{N/T-1} \right), & T < N, \\ \frac{\sigma_\epsilon^2}{T/N-1}, & T > N. \end{cases}$$

Comparison with the MP formula (227) highlights the way that near the coincident regime, $T \to N$, the loss behaves as if the data were structureless MP data. In Fig. 16, we plot some example test loss curves using numerical simulations of power-law data and compare with our analytical formula (230). This shows an excellent fit of this formula to experiment.

On the one hand, from the expression of the test loss, (230), and the figure, we can see that with power-law data there is a scaling regime, $T \ll N$, where the test loss has power-law behavior, $\mathcal{L}_{\mathrm{PL}}(N,T) \sim T^{-\alpha}$, analogous to the random feature model in the main text. On the other hand, as discussed in the previous subsection, there cannot be a dual regime where the model behaves as a power-law in the number of features. In other words, from the perspective of increasing T, there's no plateau behavior as exhibited by (1); instead, for $T > N$, we only get the universal noise contribution.

Thus, we conclude that a simple linear model does not have the same rich phenomenology exhibited by the LLMs of [18]. To make a sufficient minimal model of neural scaling phenomenology, we need to distinguish the latent features of the data from the model's features. We do this in the main body of the paper using the random feature model presented in §3.

B Explicit Solutions for Δ

We would like to self-consistently solve the equation

$$\Delta(\mu, M) = \mathrm{tr} \left\{ \frac{\Lambda}{\gamma I_M + \frac{\mu\Lambda}{1+\Delta(\mu,M)}} \right\}, \qquad (231)$$

in the limit where the ridge parameter is vanishing, $\gamma \to 0$. Here, μ is a generic scale that can either represent the number of features, N, or the size of the training set, T. We will proceed to solve this equation order-by-order.

B.1 Explicit Solution for Δ_{-1}

First, let us consider the parameter regime where $\mu < M$. Then, consideration of the above expression as $\Delta(\mu, M)$ becomes large suggests that it should diverge as $1/\gamma$ as $\gamma \to 0$. Thus, to find the leading behavior, we can insert an ansatz expansion

$$\Delta(\mu, M) \equiv \frac{\Delta_{-1}}{\gamma} + \Delta_0 + \Delta_1 \gamma + \dots, \qquad (232)$$
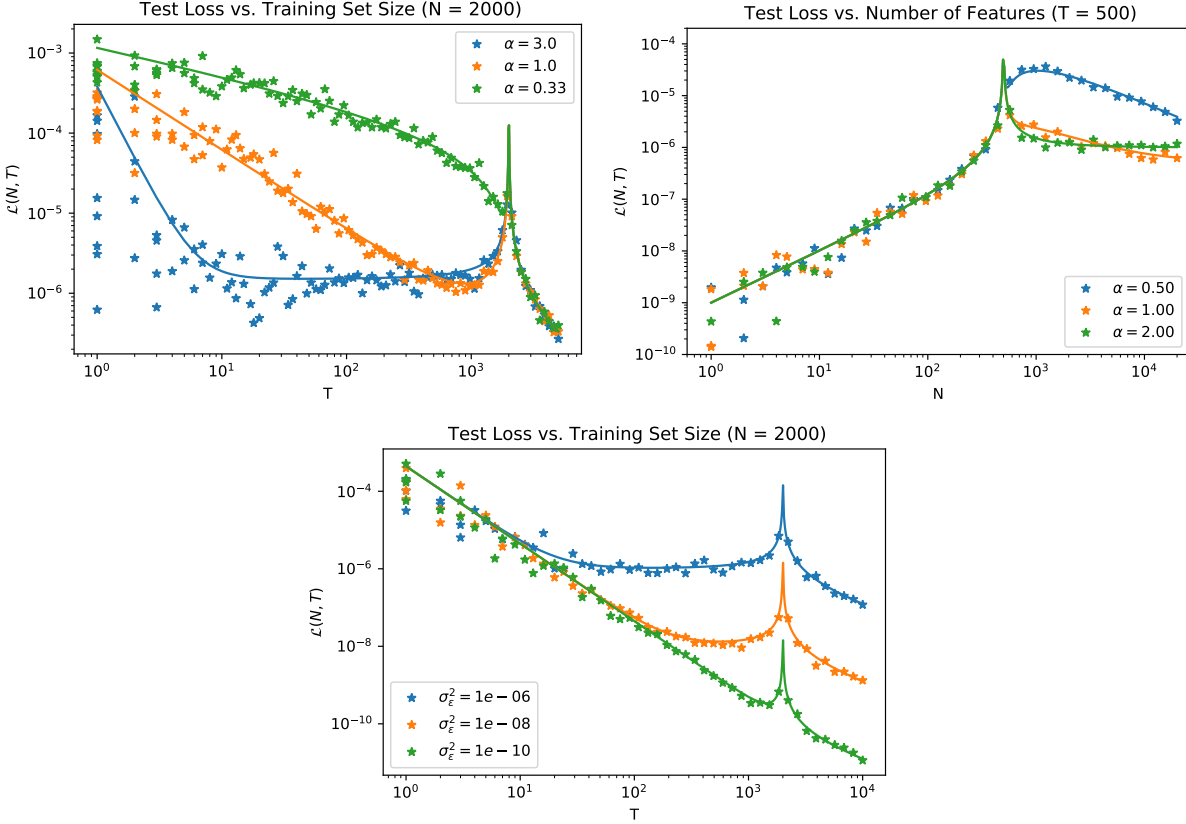
Figure 16: Test loss from numerical simulations (stars) and our analytical expression (solid lines) of a linear model trained on power-law data ($\lambda_+ = 1$, $\sigma_w^2 = 1$). The top panels illustrate that the test loss of this model is not symmetric under the exchange $T \leftrightarrow N$. **Top Left:** The size of the training set, T, is varied for a few different power-law exponents, α, while the number of features and the noise variance are held fixed ($N = 2000$, $\sigma_\epsilon^2 = 10^{-6}$); for $T \ll N$, the test loss behaves as a power law in the training set size, $\mathcal{L}_{\text{PL}}(N, T) \sim T^{-\alpha}$, but for $T > N$ there's no plateau behavior. **Top Right:** The number of features, N, is varied for a few different power-law exponents, α, while the size of the training set and the noise variance are held fixed ($T = 1000$, $\sigma_\epsilon^2 = 10^{-6}$); for $N \ll T$, the test loss actually *grows* with increasing features as $\mathcal{L}_{\text{PL}}(N, T) \sim N$. **Bottom:** The number of features, N is varied for several different values of the noise variance, σ_ϵ^2, while the power-law exponent and training set size are held fixed ($\alpha = 1.0$, $N = 2000$).

and rearrange to find:

$$1 = \text{tr}\left\{\frac{\Lambda}{\Delta_{-1} I_M + \mu\Lambda}\right\} + O(\gamma) \ . \tag{233}$$

Covering the trace into a sum over eigenvalues, we get

$$1 = \sum_{I=1}^{M} \frac{1}{\mu + \Delta_{-1}\lambda_I^{-1}} + O(\gamma) \ , \tag{234}$$

where λ_I is the spectrum of Λ. Our goal is to find an explicit expression for $\Delta_{-1}(\mu, M)$. We will proceed by considering two asymptotic regimes and then by constructing a solution that naturally interpolates between them.

The first regime concerns the limit where the scale, μ, approaches the size of the latent space, M. It's easy to see that in the exact limit of $\mu \to M$, the equation (234) is satisfied by $\Delta_{-1} = 0$. This suggests that in this *coincident* regime, we might find a solution in terms of an expansion

$$\Delta_{-1}(\varepsilon) \equiv a\varepsilon + a'\varepsilon^2 + \dots , \tag{235}$$

with the closeness parameter defined as

$$\varepsilon \equiv 1 - \frac{\mu}{M} \ . \tag{236}$$

Inverting this definition, (236), for μ and substituting both that and the expansion (235) into the summation equation, (234), we find:

$$1 = \frac{1}{M}\sum_{I=1}^{M}\left[1 + \left(1 - \frac{a}{M\lambda_I}\right)\varepsilon + O(\varepsilon^2)\right] \ . \tag{237}$$

From this, we see that zeroth order term is satisfied by construction and that the constant a must satisfy

$$a = \frac{M^2}{\sum_{I=1}^{M}\lambda_I^{-1}} \ . \tag{238}$$

If desired, this can be systematically continued to find the higher-order coefficients in the expansion (235), but for our needs the first order term is sufficient.

Up to this point this expression, (238), is completely universal given a spectrum λ_I. To simplify further, let us insert our latent space power law (23):

$$\lambda_I = \lambda_+ \left(\frac{1}{I}\right)^{1+\alpha} \ , \tag{239}$$

where recall λ_+ represents the largest eigenvalue, and $\alpha > 0$ controls the exponent of the power-law limit of the scaling law.[81] Inserting this spectrum into (238), then approximating

---

[81]Alternatively, we could instead average over a smooth density of eigenvalues, (22): regardless, at large M the smooth density will behave similarly to the fixed spectrum.

the sum as an integral and taking the large-latent-space limit, $M \to \infty$, we find:

$$a = \frac{(2+\alpha)\lambda_+}{M^\alpha} + O(M^{-1-\alpha}) \ . \tag{240}$$

The second regime is a *scaling* regime: to find another solution we essentially reverse the order of limits we just took; for the coincident regime we first expanded around $\mu \to M$ and then took $M \to \infty$, but in the scaling regime we will first take $M \to \infty$ and then look for a power-law scaling solution in μ as $\mu \to \infty$. In this limit, we can safely approximate the sum as an integral from the start,

$$\sum_{I=1}^{M} \frac{1}{\mu + \Delta_{-1}\lambda_I^{-1}} \to \int_1^\infty \frac{dI}{\mu + \Delta_{-1}I^{1+\alpha}/\lambda_+} \ , \tag{241}$$

where importantly we have taken the limit of the integral to infinity as $M \to \infty$, and we have also substituted in for the power-law spectrum (239).[82] While this integral can be evaluated directly in terms of a hypergeometric function (as was done in [38]) instead let's instead substitute in a scaling ansatz:

$$\Delta_{-1}(\mu) \equiv \frac{\lambda_+ k}{\mu^\alpha} \ , \tag{242}$$

with a constant k assumed to be independent of μ, and with the function $\Delta_{-1}(\mu)$ independent of M since we've taken $M \to \infty$. In addition to that substitution, changing integration variables as

$$I = \frac{s\mu}{k^{1/(1+\alpha)}} \ , \tag{243}$$

and setting the integral equal to unity as per (234), we can find an integral expression for the constant k:

$$k = \left[\int_{k^{1/(1+\alpha)}\mu^{-1}}^\infty \frac{ds}{1+s^{1+\alpha}} \right]^{1+\alpha} \ . \tag{244}$$

Since in our scaling limit, $\mu \to \infty$, the constant k in our scaling ansatz, (242), is assumed to be independent of μ, we can now safely take the second limit of $\mu \to \infty$ and evaluate the dimensionless integral in the square brackets in terms of some gamma functions:

$$\int_0^\infty \frac{ds}{1+s^{1+\alpha}} = \Gamma\left(\frac{\alpha}{1+\alpha}\right)\Gamma\left(\frac{2+\alpha}{1+\alpha}\right) \ . \tag{245}$$

This is just some number, and using a gamma function identity, we can rewrite this number to give an expression for k:

$$k = \left[\frac{\frac{\pi}{1+\alpha}}{\sin\left(\frac{\pi}{1+\alpha}\right)} \right]^{1+\alpha} \ . \tag{246}$$

---

[82]N.B. while the coincident limit has the universal form, (238), being determined by a simple function of the underlying spectrum, the scaling limit is specific to the power law spectrum.

Lastly, having solved (234) in two non-commuting limits, as a final step we need to find a solution that interpolates between these two regimes. Consider a final ansatz

$$\Delta_{-1}(\mu, M) \equiv \lambda_+ k \left(\mu^{-\alpha} - M^{-\alpha} \right) + b\varepsilon \,, \tag{247}$$

for some constant b that we assume to be independent of μ. To make sure that we match the correct behavior in the coincident regime, we expand (247) as $\mu \to M$,

$$\Delta_{-1}(\mu, M) = (b + \alpha \lambda_+ k) \, \varepsilon + O\!\left(\varepsilon^2\right) \,, \tag{248}$$

and set it equal to (235), giving

$$b = [2 + \alpha(1 - k)] \, \frac{\lambda_+}{M^\alpha} \,. \tag{249}$$

Then, substituting this back in our ansatz (247), we see that in the scaling regime we will match onto our solution (242) as we take $M \to \infty$:

$$\Delta_{-1}(\mu, \infty) \to \Delta_{-1}(\mu) = \frac{\lambda_+ k}{\mu^\alpha} \,. \tag{250}$$

Overall, we can write our explicit interpolating solution as

$$\Delta_{-1}(\mu, M) = \frac{\lambda_+}{M^\alpha} \left\{ k \left[\left(\frac{M}{\mu} \right)^\alpha - 1 \right] + [2 + \alpha(1 - k)] \left(1 - \frac{\mu}{M} \right) \right\} \,, \tag{251}$$

with the dimensionless constant k defined in (246). To test our solution, we can also solve for $\Delta_{-1}(\mu, M)$ numerically using (234). In Fig. 17 we plot our numerical solution against our analytical solution (251) as a function of μ, for different α and M. Overall we see excellent match everywhere except for small $\mu \lesssim 10$, where our scaling solution would receive corrections and where we expect our random matrix theory analysis of §3 also to be subject to corrections.

Finally, let us note that for $\mu > M$ the empirical covariance is full rank, so there's no pole in the resolvent as $\gamma \to 0$: this means that there's no $\gamma \to 0$ pole in Δ either. Thus, overall we have:

$$\Delta_{-1}(\mu, M) = \begin{cases} \frac{\lambda_+}{M^\alpha} \left\{ k \left[\left(\frac{M}{\mu} \right)^\alpha - 1 \right] + [2 + \alpha(1 - k)] \left(1 - \frac{\mu}{M} \right) \right\} \,, & \mu < M \,, \\ 0 \,, & \mu > M \,. \end{cases} \tag{252}$$

B.2   Explicit Solution for Δ_0

We will need to divide the analysis into the cases of $\mu < M$ and $\mu > M$.

Figure 17: Comparison analytical and numerical solutions for $\Delta_{-1}(\mu, M)$ for different exponents α and fixed μ: the numerical solution of (234) (dots) closely matches our analytical solution (251) (solid lines) except for small small $\mu \lesssim 10$, where we also expect the random matrix theory analysis of the paper to deviate.

$\mu < M$

Let's continue with our expansion of Δ in the parameter regime where $\mu > M$ to find the NLO contribution Δ_0. The $O(\gamma^0)$ term in our power series expansion of (231) gives

$$\Delta_0 = (1 + \Delta_0) \sum_{I=1}^{M} \frac{\mu \lambda_I^2}{(\mu \lambda_I + \Delta_{-1})^2} \,, \tag{253}$$

which can be solved for Δ_0 in terms of Δ_{-1}. To begin, first note that we can rewrite (253) as

$$\kappa = \sum_{I=1}^{M} \frac{\mu}{\left(\mu + \Delta_{-1}\lambda_I^{-1}\right)^2} \,, \qquad \Delta_0 = \frac{\kappa}{1 - \kappa} \,. \tag{254}$$

Given our solution for Δ_{-1}, we can evaluate the sum in the left expression in terms of known quantities and then easily find Δ_0. As before, let us consider two asymptotic regimes for the sum and then construct an interpolating solution.

In the coincident regime ($\mu \to M$), substituting in $\Delta_{-1} = a\varepsilon + O(\varepsilon^2)$ and $\mu = M(1-\varepsilon)$, and using (238), we can evaluate the sum to linear order an find

$$\kappa = 1 - \varepsilon + O(\varepsilon^2) \,, \tag{255}$$

which gives

$$\Delta_0 = \frac{1}{M/\mu - 1} + \dots \,. \tag{256}$$

86

Interestingly, this leading order answer is completely independent of the spectrum. As mentioned before, we can systematically continue this ε expansion to find more terms if needed.

In the scaling regime, we can follow the same procedure as in the previous subsection: first, we substitute in $\Delta_{-1} = \lambda_+ k\mu^{-\alpha}$ and $\lambda_I = \lambda_+ I^{-1-\alpha}$ for the leading order scaling solution and spectrum, respectively; next, we approximate the sum in (254) as an integral; then, we change variables in (243). Finally, we safely take the $\mu \to \infty$ limit to get

$$\kappa = \frac{1}{k^{1/(1+\alpha)}} \int_0^\infty \frac{ds}{(1 + s^{1+\alpha})^2} = \frac{\alpha}{1+\alpha} \quad \Longrightarrow \quad \Delta_0 = \alpha \tag{257}$$

in the scaling regime. Putting (256) and (257) together, it's easy to see that

$$\Delta_0(\mu, M) = \alpha + \frac{1}{M/\mu - 1} \tag{258}$$

behaves as a simple interpolating solution between these two regimes.

$\mu > M$

As we discussed in previously, when we assume $M < \mu$, then Δ no longer diverges as $1/\gamma$ as $\gamma \to 0$, but instead approaches a constant. Thus, to find the new leading behavior, we should use a different ansatz,

$$\Delta(\mu, M) \equiv \Delta_0 + \Delta_1 \gamma + \dots, \tag{259}$$

for which we will only need to evaluate the first term. Inserting this expansion into the self-consistent equation, we find

$$\Delta_0 = \sum_{I=1}^M \frac{1 + \Delta_0}{\mu} . \tag{260}$$

This can be simply summed and rearranged to give

$$\Delta_0 = \frac{1}{\mu/M - 1} , \tag{261}$$

another universal answer independent of the spectrum.

Thus, altogether we have:

$$\Delta_0 = \begin{cases} \alpha + \frac{1}{M/\mu - 1} , & \mu < M , \\ \frac{1}{\mu/M - 1} , & \mu > M . \end{cases} \tag{262}$$

To verify this, in Fig. 18 we plot a comparison of this analytical approximation against a numerical evaluation of Δ_0. As before, we see deviations for very small $\mu \lesssim 10$, where our scaling solution would receive corrections. Additionally, we also note that for $\mu < M$, the fit in the crossover between the scaling and coincident regions has some very small error: this could be reduced by solving the self-consistent equations for Δ_{-1} and Δ_0 to order ε^2 in the coincident region.

Figure 18: Comparison analytical and numerical solutions for $\Delta_0(\mu, M)$ for different exponents α and fixed μ: the numerical solution (dots) closely matches our analytical solution (258) (solid lines) except for small $\mu \lesssim 10$ where we also expect the random matrix theory analysis of the paper to deviate.

Δ_0 for the Truncated Power-Law Spectrum

In §3.2.1, we needed to calculate Δ_0^F, which is determined by the self-consistent equation (82):

$$\Delta^F(T, N) = \text{tr}\left\{ \frac{\Omega}{\gamma I_N + \frac{T\Omega}{1 + \Delta^F(T,N)}} \right\} . \tag{263}$$

This is similar to the self-consistent equation (231) that we just solved, with trivial substitutions $\mu \to T$ and $N \to M$, and, less trivially, with the covariance Λ, and its strict power-law spectrum (239), replaced by a random projection of that spectrum, $\Omega = u\Lambda u^T$.

Let's first consider the typical case where our number of features is less than the dimension of the latent space, $N < M$. Then, for a fixed u, the spectrum of Ω will have a power-law bulk and then terminate in a non-power-law tail with a rapid decline to zero.

When our model is also underparameterized, $N < T$, or overparameterized but near the coincident region, $N \gtrsim T$, we know from our analysis in the previous subsection that the answer is universal, independent of the spectrum: thus, automatically our calculation from before *almost* carries over to this setting. On the other hand, we might expect a deviation when $T \ll N$, since Δ_0 was not universal in the scaling region, cf. (257). However, in this limit there are only significant deviations from a strict power-law spectrum at at very small eigenvalues; since these small eigenvalues will make a negligible contribution to the integrals determining Δ_{-1}^F and Δ_0^F, we expect this non-power-law tail to have negligible effect on the calculation of Δ_0^F. (We also expect the eigenvalues to shift by a multiplicative

88

factor, but this scaling does not appear in Δ_0^F and won't contribute.)

Thus, as analogous to the answer in previous subsection (262), even in this projected-spectrum case we expect

$$\Delta_0^F(T, N) = \begin{cases} \alpha + \frac{1}{N/T - 1}\,, & T < N\,,\ N < M\,, \\ \frac{1}{T/N - 1}\,, & T > N\,,\ N < M\,. \end{cases} \qquad (264)$$

In Fig. 19, we confirm this analysis in the overparameterized setting by comparing to a numerical evaluation of Δ_0^F. In this case, we find a reasonably good fit: we see the expected deviations for very small $T \lesssim 10$, and we also see some more deviations in the crossover region for larger α as compared to our fit in Fig. 18; in this region, we expect that the higher-order spectrum-dependent terms are less negligible. Nevertheless, these deviations are overall insignificant as the precise crossover behavior is not particularly important and could always be improved by analyzing additional terms in the ε expansion.
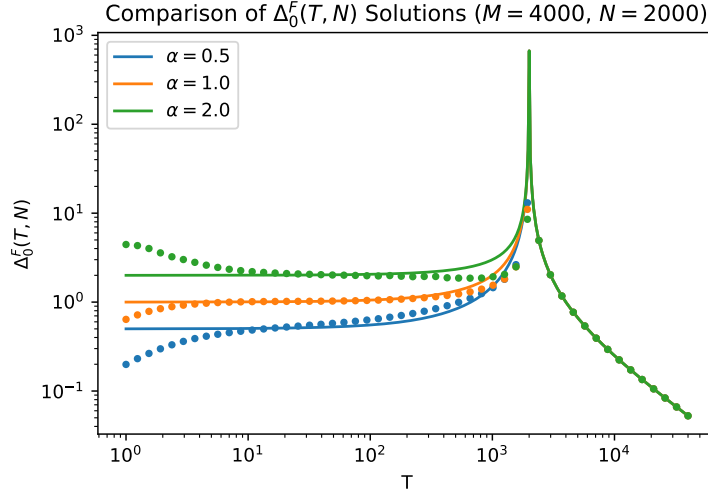

Figure 19: Comparison of analytical and numerical solutions for $\Delta_0^F(T, N)$, determined by a truncated power-law spectrum ($\lambda_+ = 2$, $\sigma_u^2 = 1$), for different exponents α and fixed number of latent and random features ($M = 4000$, $N = 2000$): the numerical solution (dots) matches our analytical solution (264) (solid lines) except for small $T \lesssim 10$, where we also expect the random matrix theory analysis of the paper to deviate. For larger α, we also see some non-negligible deviations in the crossover from the scaling regime to the coincident regime; this could be improved by analyzing the next term in the ε expansion.

We can also consider the somewhat less typical case of $N > M$, which is of interest in §4.1 where we considered the breakdown of neural scaling phenomenology. In this situation, for a typical fixed u, the spectrum of Ω will have the complete power-law of Λ plus the addition of $(N - M)$ zero eigenvalues to fill out the rest of the spectrum. Considering the

self-consistent equation for Δ^F, (263), it's clear that a similar analysis follows except that the sum now truncates at M:

$$\Delta^F(T, N) = \sum_{i=1}^{N} \frac{\lambda_i}{\gamma + \frac{T\lambda_i}{1+\Delta^F(T,N)}} = \sum_{i=1}^{M} \frac{\lambda_i}{\gamma + \frac{T\lambda_i}{1+\Delta^F(T,N)}} \,. \tag{265}$$

This means that in this regime we have

$$\Delta^F(T, N) \to \Delta(T, M) \,, \tag{266}$$

and the complete table of values for Δ^F is given by

$$\Delta_0^F(T, N) = \begin{cases} \alpha + \frac{1}{N/T-1} \,, & T < N \,, \quad N < M \,, \\ \frac{1}{T/N-1} \,, & T > N \,, \quad N < M \,, \\ \alpha + \frac{1}{M/T-1} \,, & T < M \,, \quad N > M \,, \\ \frac{1}{T/M-1} \,, & T > M \,, \quad N > M \,. \end{cases} \tag{267}$$

Interestingly, though the self-consistent equation, (263), depends explicitly on u through $\Omega = u\Lambda u^T$, we see here that $\Delta_0^F(T, N)$ is actually *independent* of the particular u, instead only caring about its size, N.

References

[1] T. B. Brown, B. P. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krüger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. J. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language Models are Few-Shot Learners," `arXiv:2005.14165 [cs.CL]`.

[2] R. Thoppilan, D. De Freitas, J. Hall, N. Shazeer, A. Kulshreshtha, H.-T. Cheng, A. Jin, T. Bos, L. Baker, Y. Du, *et al.*, "LaMDA: Language Models for Dialog Applications," `arXiv:2201.08239 [cs.CL]`.

[3] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, *et al.*, "Palm: Scaling language modeling with pathways," `arXiv:2204.02311 [cs.CL]`.

[4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Advances in Neural Information Processing Systems*, vol. 30, pp. 5998–6008. 2017. `arXiv:1706.03762 [cs.CL]`.

[5] J. W. Rae, S. Borgeaud, T. Cai, K. Millican, J. Hoffmann, F. Song, J. Aslanides, S. Henderson, R. Ring, S. Young, *et al.*, "Scaling language models: Methods, analysis & insights from training gopher," `arXiv:2112.11446 [cs.CL]`.

[6] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, *et al.*, "Opt: Open pre-trained transformer language models," `arXiv:2205.01068 [cs.CL]`.

[7] S. Smith, M. Patwary, B. Norick, P. LeGresley, S. Rajbhandari, J. Casper, Z. Liu, S. Prabhumoye, G. Zerveas, V. Korthikanti, E. Zheng, R. Child, R. Y. Aminabadi, J. Bernauer, X. Song, M. Shoeybi, Y. He, M. Houston, S. Tiwary, and B. Catanzaro, "Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model," `arXiv:2201.11990 [cs.CL]`.

[8] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. d. L. Casas, L. A. Hendricks, J. Welbl, A. Clark, *et al.*, "Training Compute-Optimal Large Language Models," `arXiv:2203.15556 [cs.CL]`.

[9] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning.* Springer Science & Business Media, 2009.

[10] R. M. Neal, "Priors for infinite networks," in *Bayesian Learning for Neural Networks*, pp. 29–53. Springer, 1996.

[11] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein, "Deep Neural Networks as Gaussian Processes," in *International Conference on Learning Representations.* 2018. `arXiv:1711.00165 [stat.ML]`.

[12] A. G. de G. Matthews, M. Rowland, J. Hron, R. E. Turner, and Z. Ghahramani, "Gaussian Process Behaviour in Wide Deep Neural Networks," in *International Conference on Learning Representations.* 2018. `arXiv:1804.11271 [stat.ML]`.

[13] A. Jacot, F. Gabriel, and C. Hongler, "Neural tangent kernel: Convergence and generalization in neural networks," in *Advances in Neural Information Processing Systems*, vol. 31, pp. 8571–8580. 2018. `arXiv:1806.07572 [cs.LG]`.

[14] J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington, "Wide Neural Networks of Any Depth Evolve as Linear Models Under Gradient Descent," in *Advances in Neural Information Processing Systems*, vol. 32, pp. 8572–8583. 2019. `arXiv:1902.06720 [stat.ML]`.

[15] S. Yaida, "Non-Gaussian Processes and Neural Networks at Finite Widths," in *Mathematical and Scientific Machine Learning Conference.* 2020. `arXiv:1910.00019 [stat.ML]`.

[16] D. A. Roberts, S. Yaida, and B. Hanin, *The Principles of Deep Learning Theory.* Cambridge University Press, 2022. `arXiv:2106.10165 [cs.LG]`. `https://deeplearningtheory.com`.

[17] M. Belkin, D. Hsu, S. Ma, and S. Mandal, "Reconciling modern machine-learning practice and the classical bias-variance trade-off," *Proceedings of the National Academy of Sciences* **116** no. 32, (2019) 15849–15854, `arXiv:1812.11118 [stat.ML]`.

[18] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," `arXiv:2001.08361 [cs.LG]`.

[19] G. Branwen, "The Scaling Hypothesis," May, 2020. `https://www.gwern.net/Scaling-hypothesis`.

[20] T. Henighan, J. Kaplan, M. Katz, M. Chen, C. Hesse, J. Jackson, H. Jun, T. B. Brown, P. Dhariwal, S. Gray, *et al.*, "Scaling laws for autoregressive generative modeling," `arXiv:2010.14701 [cs.LG]`.

[21] D. Hernandez, J. Kaplan, T. Henighan, and S. McCandlish, "Scaling laws for transfer," `arXiv:2102.01293 [cs.LG]`.

[22] M. A. Gordon, K. Duh, and J. Kaplan, "Data and Parameter Scaling Laws for Neural Machine Translation," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 5915–5922. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, Nov., 2021.

[23] D. Hernandez, T. Brown, T. Conerly, N. DasSarma, D. Drain, S. El-Showk, N. Elhage, Z. Hatfield-Dodds, T. Henighan, T. Hume, *et al.*, "Scaling Laws and Interpretability of Learning from Repeated Data," `arXiv:2205.10487 [cs.LG]`.

[24] B. Sorscher, R. Geirhos, S. Shekhar, S. Ganguli, and A. S. Morcos, "Beyond neural scaling laws: beating power law scaling via data pruning," `arXiv:2206.14486 [cs.LG]`.

[25] J. Hestness, S. Narang, N. Ardalani, G. Diamos, H. Jun, H. Kianinejad, M. Patwary, M. Ali, Y. Yang, and Y. Zhou, "Deep learning scaling is predictable, empirically," `arXiv:1712.00409 [cs.LG]`.

[26] J. S. Rosenfeld, A. Rosenfeld, Y. Belinkov, and N. Shavit, "A Constructive Prediction of the Generalization Error Across Scales," in *International Conference on Learning Representations*. 2020. `arXiv:1909.12673 [cs.LG]`. `https://openreview.net/forum?id=ryenvpEKDr`.

[27] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, *et al.*, "Evaluating large language models trained on code," `arXiv:2107.03374 [cs.CL]`.

[28] A. Lewkowycz, A. Andreassen, D. Dohan, E. Dyer, H. Michalewski, V. Ramasesh, A. Slone, C. Anil, I. Schlag, T. Gutman-Solo, *et al.*, "Solving Quantitative Reasoning Problems with Language Models," `arXiv:2206.14858 [cs.CL]`.

[29] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," `arXiv:2010.11929 [cs.CV]`.

[30] A. Srivastava, A. Rastogi, A. Rao, A. A. M. Shoeb, A. Abid, A. Fisch, A. R. Brown, A. Santoro, A. Gupta, A. Garriga-Alonso, *et al.*, "Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models," `arXiv:2206.04615 [cs.CL]`.

[31] B. C. Roy, M. C. Frank, P. DeCamp, M. Miller, and D. Roy, "Predicting the birth of a spoken word," *Proceedings of the National Academy of Sciences* **112** no. 41, (2015) 12663–12668.

[32] T. Linzen, "How Can We Accelerate Progress Towards Human-like Linguistic Generalization?," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5210–5217. 2020. `arXiv:2005.00955 [cs.CL]`.

[33] A. Warstadt and S. Bowman, "What Artificial Neural Networks Can Tell Us About Human Language Acquisition," `arXiv:2208.07998 [cs.CL]`.

[34] C. Louart, Z. Liao, and R. Couillet, "A Random Matrix Approach to Neural Networks," *The Annals of Applied Probability* **28** no. 2, (2018) 1190–1248, `arXiv:1702.05419 [math.PR]`.

[35] B. Bordelon, A. Canatar, and C. Pehlevan, "Spectrum dependent learning curves in kernel regression and wide neural networks," in *International Conference on Machine Learning*, pp. 1024–1034, PMLR. 2020. `arXiv:2002.02561 [cs.LG]`.

[36] A. Canatar, B. Bordelon, and C. Pehlevan, "Spectral Bias and Task-Model Alignment Explain Generalization in Kernel Regression and Infinitely Wide Neural Networks," *Nature Communications* **12** no. 1, (2021) 2914, `arXiv:2006.13198 [stat.ML]`.

[37] U. Sharma and J. Kaplan, "Scaling Laws from the Data Manifold Dimension," *Journal of Machine Learning Research* **23** no. 9, (2022) 1–34, `arXiv:2004.10802 [cs.LG]`. `http://jmlr.org/papers/v23/20-1111.html`.

[38] Y. Bahri, E. Dyer, J. Kaplan, J. Lee, and U. Sharma, "Explaining Neural Scaling Laws," `arXiv:2102.06701 [cs.LG]`.

[39] M. Hutter, "Learning curve theory," `arXiv:2102.04074 [cs.LG]`.

[40] A. Wei, W. Hu, and J. Steinhardt, "More Than a Toy: Random Matrix Models Predict How Real-World Neural Representations Generalize," `arXiv:2203.06176 [cs.LG]`.

[41] S. Spigler, M. Geiger, and M. Wyart, "Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm," *Journal of Statistical Mechanics: Theory and Experiment* **2020** no. 12, (2020) 124001, `arXiv:1905.10843 [stat.ML]`.

[42] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language Models are Unsupervised Multitask Learners," *OpenAI Blog* (2019) . `https://github.com/openai/gpt-2`.

[43] M. Mitzenmacher, "A brief history of generative models for power law and lognormal distributions," *Internet mathematics* **1** no. 2, (2004) 226–251.

[44] A. Clauset, C. R. Shalizi, and M. E. J. Newman, "Power-Law Distributions in Empirical Data," *SIAM Review* **51** no. 4, (2009) 661–703, `arXiv:0706.1062 [physics.data-an]`.

[45] A. Krizhevsky, "Learning multiple layers of features from tiny images," 2009. `https://www.cs.toronto.edu/~kriz/cifar.html`.

[46] S. Merity, C. Xiong, J. Bradbury, and R. Socher, "Pointer Sentinel Mixture Models," `arXiv:1609.07843 [cs.CL]`.

[47] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, *et al.*, "Huggingface's transformers: State-of-the-art natural language processing," `arXiv:1910.03771 [cs.CL]`.

[48] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philos. Mag.* **2** no. 11, (1901) 559–572.

[49] S. Bradde and W. Bialek, "PCA meets RG," *Journal of Statistical Physics* **167** no. 3, (2017) 462–475, `arXiv:1610.09733 [physics.bio-ph]`.

[50] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," `arXiv:1810.04805 [cs.CL]`.

[51] V. A. Marčenko and L. A. Pastur, "Distribution of Eigenvalues for Some Sets of Random Matrices," *Mathematics of the USSR-Sbornik* **1** no. 4, (Apr, 1967) 457–483.

[52] J. W. Silverstein and Z. D. Bai, "On the Empirical Distribution of Eigenvalues of a Class of Large Dimensional Random Matrices," *Journal of Multivariate Analysis* **54** no. 2, (1995) 175–192.

[53] Z. Burda, A. Görlich, A. Jarosz, and J. Jurkiewicz, "Signal and noise in correlation matrix," *Physica A: Statistical Mechanics and its Applications* **343** (2004) 295–310.

[54] E. Dyer and G. Gur-Ari, "Asymptotics of Wide Networks from Feynman Diagrams," in *International Conference on Learning Representations*. 2020. `arXiv:1909.11304 [cs.LG]`.

[55] G. 't Hooft, "A Planar Diagram Theory for Strong Interactions," *Nucl. Phys.* **B72** (1974) 461.

[56] J. A. Mingo and R. Speicher, *Free Probability and Random Matrices*, vol. 35. Springer, 2017.

[57] F. J. Dyson, "The S Matrix in Quantum Electrodynamics," *Phys. Rev.* **75** (Jun, 1949) 1736–1755.

[58] J. Schwinger, "On the Green's functions of quantized fields. I," *Proceedings of the National Academy of Sciences* **37** no. 7, (1951) 452–455.

[59] B. Adlam, J. Levinson, and J. Pennington, "A Random Matrix Perspective on Mixtures of Nonlinearities for Deep Learning," 2019.

[60] B. Eynard, T. Kimura, and S. Ribault, "Random matrices," `arXiv:1510.04430 [math-ph]`.

[61] M. Opper, "Statistical mechanics of learning: Generalization," *The handbook of brain theory and neural networks* (1995) 922–925.

[62] M. Opper, "Learning to generalize," *Frontiers of Life* **3** no. part 2, (2001) 763–775.

[63] M. S. Advani, A. M. Saxe, and H. Sompolinsky, "High-dimensional dynamics of generalization error in neural networks," *Neural Networks* **132** (2020) 428–446, `arXiv:1710.03667 [stat.ML]`.

[64] S. Spigler, M. Geiger, S. d'Ascoli, L. Sagun, G. Biroli, and M. Wyart, "A jamming transition from under-to over-parametrization affects loss landscape and generalization," `arXiv:1810.09665 [cs.LG]`.

[65] M. Geiger, S. Spigler, S. d'Ascoli, L. Sagun, M. Baity-Jesi, G. Biroli, and M. Wyart, "Jamming transition as a paradigm to understand the loss landscape of deep neural networks," *Phys. Rev. E* **100** (Jul, 2019) 012115, `arXiv:1809.09349 [cond-mat.dis-nn]`.

[66] P. Nakkiran, G. Kaplun, Y. Bansal, T. Yang, B. Barak, and I. Sutskever, "Deep double descent: Where bigger models and more data hurt," *Journal of Statistical Mechanics: Theory and Experiment* **2021** no. 12, (2021) 124003, `arXiv:1912.02292 [cs.LG]`.

[67] P. Nakkiran, P. Venkat, S. M. Kakade, and T. Ma, "Optimal Regularization can Mitigate Double Descent," in *International Conference on Learning Representations*. 2021. `arXiv:2003.01897 [cs.LG]`. `https://openreview.net/forum?id=7R7fAoUygoa`.

[68] J. Friedman and B. E. Popescu, "Gradient directed regularization for linear regression and classification," tech. rep., 2002. `https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.3.4934&rep=rep1&type=pdf`.

[69] G. Raskutti, M. J. Wainwright, and B. Yu, "Early stopping and non-parametric regression: an optimal data-dependent stopping rule," *The Journal of Machine Learning Research* **15** no. 1, (2014) 335–366, `arXiv:1306.3574 [stat.ML]`.

[70] J. Sevilla and P. Villalobos, "Parameter counts in Machine Learning," Jun, 2021. `https://www.alignmentforum.org/posts/GzoWcYibWYwJva8aL/parameter-counts-in-machine-learning`.

[71] B. Adlam, "Research statement." `https://web.archive.org/web/20220908153010/https://www.benadlam.com/`.

[72] F. Camastra and A. Staiano, "Intrinsic dimension estimation: Advances and open problems," *Information Sciences* **328** (2016) 26–41.

[73] E. Levina and P. J. Bickel, "Maximum Likelihood Estimation of Intrinsic Dimension," in *Proceedings of the 17th International Conference on Neural Information Processing Systems*, NIPS'04, p. 777–784. MIT Press, Cambridge, MA, USA, 2004.

[74] E. Facco, M. d'Errico, A. Rodriguez, and A. Laio, "Estimating the intrinsic dimension of datasets by a minimal neighborhood information," *Scientific Reports* **7** no. 1, (2017) 12140.

[75] Y. LeCun, C. Cortes, and C. J. Burges, "The MNIST Database of handwritten digits,". `http://yann.lecun.com/exdb/mnist`.