

# The Slow Deterioration of the Generalization Error of the Random Feature Model

**Chao Ma**

*PACM, Princeton University, Princeton, NJ 08544*

CHAOM@PRINCETON.EDU

**Lei Wu**

*PACM, Princeton University, Princeton, NJ 08544*

LEIWU@PRINCETON.EDU

**Weinan E**

*Department of Mathematics and PACM*

*Princeton University, Princeton, NJ 08544*

WEINAN@MATH.PRINCETON.EDU

## Abstract

The random feature model exhibits a kind of resonance behavior when the number of parameters is close to the training sample size. This behavior is characterized by the appearance of large generalization gap, and is due to the occurrence of very small eigenvalues for the associated Gram matrix. In this paper, we examine the dynamic behavior of the gradient descent algorithm in this regime. We show, both theoretically and experimentally, that there is a dynamic self-correction mechanism at work: The larger the eventual generalization gap, the slower it develops, both because of the small eigenvalues. This gives us ample time to stop the training process and obtain solutions with good generalization property.

**Keywords:** Random feature model, gradient descent, generalization error, early stopping, Gram matrix

## 1. Introduction

The properties of a machine learning model are largely controlled by its two most important parameters: the number of parameters  $m$  and the training sample size  $n$ . Of particular interest is when  $m = n$ , the transition point between under- and over-parametrized regimes. It has been reported that for some machine learning models, the generalization error undergoes a “double descent” transition at this point: Roughly speaking, it increases as a function of  $\gamma = m/n$  before  $\gamma = 1$  and decreases after  $\gamma = 1$  (Belkin et al., 2019a). In particular, the generalization error is abnormally large around  $\gamma = 1$ .

Figure 1(a) illustrates the result of fitting the MNIST dataset using random feature models. Gradient descent (GD) initialized at 0 is used to train these models. The blue dashed line is the test error for minimum-norm solutions, i.e. the GD solutions at  $t = \infty$ . As expected this curve shows a “double descent” behavior, with a peak around  $m = n$  (Belkin et al., 2019a). Shown in the red dashed line are the smallest eigenvalues of the associated Gram matrices as a function of the number of features  $m$ . One can see that the test error shows a very strong negative correlation with the size of the smallest eigenvalues.

Of more interest is the blue solid line in Figure 1(a), which shows the test error of “practical” GD solutions, found by running GD for one million iterations. Roughly speaking, this curve shows

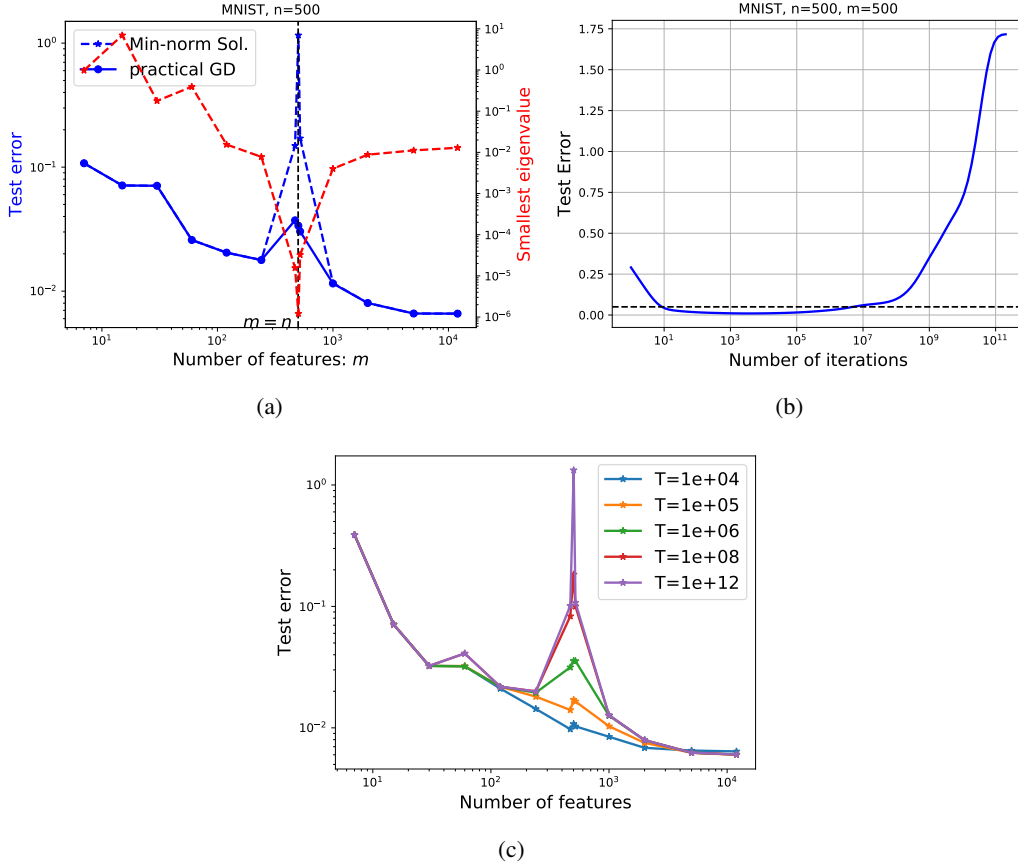


Figure 1: Fitting the MNIST dataset using the random feature model  $f(\mathbf{x}; \mathbf{a}) = \sum_{j=1}^m a_j \sigma(\mathbf{b}_j^T \mathbf{x} + c_j)$ , with  $(\mathbf{b}_j, c_j)$  independently drawn from the uniform distribution over the  $\ell_2$  unit sphere  $\mathbb{S}^d$ , and  $\sigma(t) = \max(0, t)$ . Here  $d = 784$ , and we randomly chose  $n = 500$  samples from the class 0 and 1 to form the new training set. The learning rate of GD is  $\eta = 1/\lambda_{\max}$ , where  $\lambda_{\max}$  is the largest eigenvalue of the associated Gram matrix. (a) Test errors and the smallest eigenvalues of the associated Gram matrices for different values of  $m$ . The blue dashed line shows test error of the minimum-norm solutions. The blue solid line shows the result of the GD solution found after  $10^6$  iterations. (b) Test error along the GD path for  $m = n = 500$ . The black dashed line corresponds to test error 0.05. (c) The test error of GD solutions obtained by running different number of iterations.

a decreasing behavior as a function of  $m$ , with only a very small bump detected at  $\gamma = 1$ . Note that no delicate early stopping strategy is used—we simply ran GD for one million iterations. The reason behind this can be seen from Figure 1(a) (b) which shows the test error along the GD dynamics at  $\gamma = 1$ . One can see that over a very long period of time interval, the test error remains small before it finally increases to a very large value. Another way to see this is shown in Figure 1(c): the test error curves of GD solutions behave nicely for  $T = 10^4, 10^5, 10^6$ . The “singularity” at  $\gamma = 1$  can only be seen for  $T \geq 10^8$ .

The reason behind all these is the small eigenvalues of the Gram matrix. These small eigenvalues cause the test error of the minimum-norm solution to nearly blow up at  $\gamma = 1$ . But at the same time,

they also give rise to very slow dynamics for their effect to set in, i.e. it takes a long time for the large generalization errors to develop dynamically. A practitioner will likely terminate the iteration process before these large error starts to take an effect.

Another way to say this is that the effect of the “double descent” is washed out by the dynamics, as shown in the solid blue curve in Figure 1(a).

In the following sections, we show that the **slow deterioration** feature of the GD dynamics happens in more general cases (for general  $\gamma$  and choice of kernels) and provide a theoretical analysis of this phenomenon. The analysis is based on the expansion of GD dynamics in the spectral space of the associated kernel operator. The eventual poor generalization behavior is caused by the large norm of the solution, which in turn is caused by the GD dynamics trying to finish the last mile for fitting the training data by resorting to the subspace of eigenvectors corresponding to the small eigenvalues. This is indeed a form of over-fitting. However, this over-fitting is not caused by over-parametrization, but rather by the GD dynamics trying to do its best to fit the training data.

## 2. Related work

The double descent phenomenon (also called jamming transition) was experimentally pointed out in (Advani and Saxe, 2017; Belkin et al., 2019a; Spigler et al., 2019). In (Belkin et al., 2019b), several simple models are analyzed to explain the double descent phenomenon. A more detailed analysis was carried out in (Mei and Montanari, 2019) with precise asymptotic results. Analysis of other models can be found in (Deng et al., 2019; Kini and Thrampoulidis, 2020). However, these theoretical results do not consider the generalization error of solutions along the GD trajectory.

As far as the GD dynamics is concerned, the existing literature (Yao et al., 2007; Suggala et al., 2018; E et al., 2020; Carratino et al., 2018) showed that there exist generalizable solutions along the GD path, which can be found by early stopping according to some specific rules. As a comparison, our analysis of the slow deterioration phenomenon shows that there is a large interval in which the training process can be stopped without deteriorating the test accuracy.

Also relevant to this work is the “frequency principle” investigated in (Xu, 2018; Xu et al., 2019; Zhang et al., 2019). It was observed that for random feature and neural network models, GD learns low-frequency components faster than high-frequency components. This can be viewed as the general mechanism behind the slow deterioration phenomenon.

The slow deterioration phenomenon was also observed in a simpler setting in (Advani and Saxe, 2017), but the analysis of (Advani and Saxe, 2017) still focused on optimal early stopping, for example how the optimal stopping time depends on the signal-to-noise ratio. Furthermore, their results cannot be directly applied in our setting, since the feature maps do not follow the Gaussian distribution assumption in that work.

## 3. Experimental study

Consider a random feature model with features  $\phi(\mathbf{x}; \mathbf{b})$ . Let  $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$  be the training data sampled from distribution  $\pi$ , and  $\{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_m\}$  be the random features sampled from  $\mu$ . Let  $\Phi$  be an  $n \times m$  matrix with  $\Phi_{ij} = \phi(\mathbf{x}_i; \mathbf{b}_j)$ . Then, the random feature model seeks to find a prediction function

$$\hat{f}(\mathbf{x}) = \sum_{k=1}^m a_k \phi(\mathbf{x}; \mathbf{b}_k), \quad (1)$$




Figure 2: The training error, testing error and the norm of the solution along the GD trajectory, initialized at zero. The random feature model is considered, with the feature being  $\sigma(\mathbf{b}^T \mathbf{x})$ , where  $\sigma$  is the ReLU function,  $\mathbf{b}$  and  $\mathbf{x}$  are both sampled uniformly from  $\mathbb{S}^{d-1}$ .  $d = 10$ ,  $m = n = 500$ . The left axis shows the errors and the axis on the right shows the norms. The results show that the testing error keeps being small from  $t \approx 10^1$  to  $t \approx 10^6$ , and approaches that of the minimal norm solution only after  $t > 10^9$ .

where  $\mathbf{a} = (a_1, b_2, \dots, a_m)^T$  are the parameters. To find  $\mathbf{a}$ , GD is used to optimize the following least squares objective function,

$$\min_{\mathbf{a} \in \mathbb{R}^m} \frac{1}{2n} \|\Phi \mathbf{a} - \mathbf{y}\|^2, \quad (2)$$

starting from the origin, where  $\mathbf{y}$  is a vector containing the value of the target function at  $\mathbf{x}_i$ ,  $i = 1, 2, \dots, n$ . Instead of the discrete gradient descent algorithm, we will consider the gradient flow, which is the limit of the gradient descent algorithm with learning rate tending to 0. The dynamics of  $\mathbf{a}$  is then given by

$$\frac{d}{dt} \mathbf{a}(t) = -\frac{1}{m} \frac{\partial}{\partial \mathbf{a}} \frac{1}{2n} \|\Phi \mathbf{a} - \mathbf{y}\|^2 = -\frac{1}{mn} \Phi^T (\Phi \mathbf{a} - \mathbf{y}). \quad (3)$$

Let  $\Phi = U \Sigma V^T$  be the singular value decomposition of  $\Phi$ , where  $\Sigma = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_{\min\{n,m\}}\}$  with  $\lambda_i$  being the singular values of  $\Phi$ , in descending order. Then the GD solution of (2) at time  $t \geq 0$  is

$$\mathbf{a}(t) = \sum_{i: \lambda_i > 0} \frac{1 - e^{-\lambda_i^2 t / (mn)}}{\lambda_i} (\mathbf{u}_i^T \mathbf{y}) \mathbf{v}_i. \quad (4)$$

With this solution, we can conveniently compute the training and testing error at any time  $t$ , without doing GD. The equation above shows clearly the double-sided effect of small eigenvalues. On one hand, at  $t = \infty$ , they give an  $O(1/\lambda)$  contribution to the minimum-norm solution. On the other hand, for  $t < O(mn/\lambda^2)$ , their effect can be neglected.

We first consider the random feature model with the feature vector given by  $\sigma(\mathbf{b}^T \mathbf{x})$ , where  $\sigma$  is the ReLU function,  $\mathbf{b}$  and  $\mathbf{x}$  are both sampled uniformly from  $\mathbb{S}^{d-1}$ ,  $d = 10$ . Figure 2 shows the results for the case when  $\gamma = 1$ . Results for more general values of parameters are shown in Figure 3. This figure suggests several things. The first is that the level of overfitting for the minimum solution is drastically reduced away from the transition point. In fact the minimal norm solution performs much better even for  $\gamma = 0.9$  or  $1.1$  and the degree of overfitting is much reduced.

In fact, when  $\gamma = 1/2$  or  $2$ , the test error of the minimal norm solution is only a few times larger than the best value along the GD path. If  $m$  is chosen to be  $O(\sqrt{n})$  or  $O(n^2)$ , no appreciable overfitting is observed:

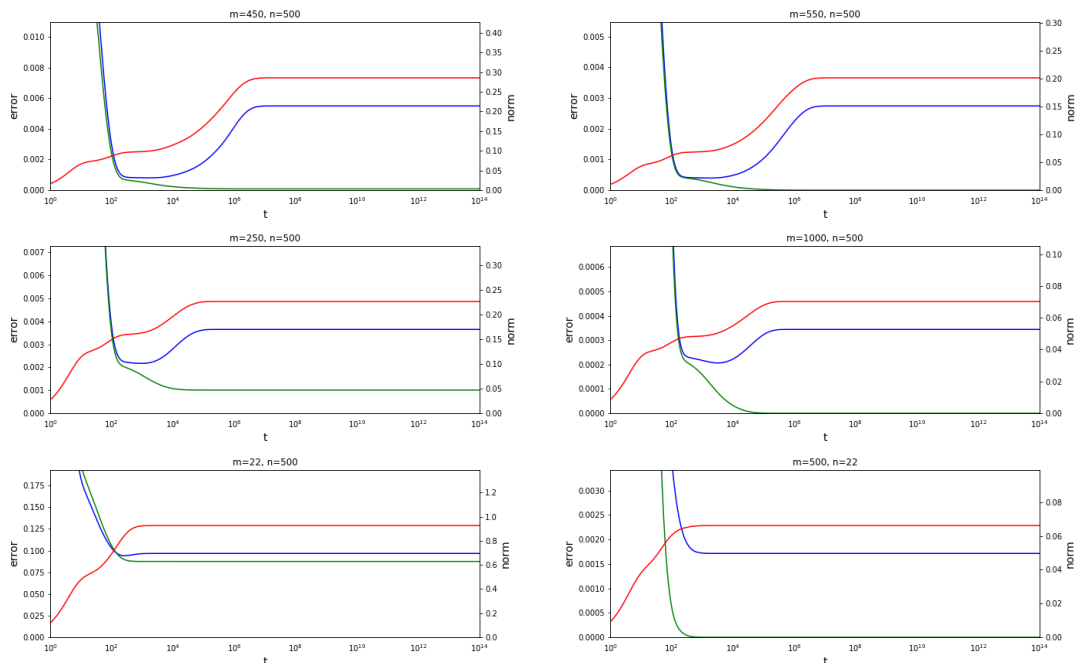


Figure 3: Time history of the training error, testing error, and solution norm along the GD path for the general case. Green lines, blue lines and red lines denote the training error, testing error, and solution norms, respectively. The minimal norm solutions perform better when  $m$  is farther away from  $n$ . No appreciable overfitting is observed for  $m = n^2$  or  $\sqrt{n}$ .

Another important observation is that the worse the overfitting, the slower it is reached. For example, in Figure 2, it takes nearly  $10^{10}$  time units for GD to converge to the minimal norm solution. On the contrary, when  $\gamma$  is far from 1, the ascending process takes shorter time and the minimal norm solutions are reached earlier, as shown in Figure 3. In Figure 4, we put together testing error curves from many different experiments, with fixed  $n$  but different  $m$ . The figure suggests that, though the testing errors of the final stable states are different, larger testing error always requires longer time to reach. Moreover, after the first period where the testing errors decrease, almost all the curves can be upper bounded by the function  $y = c\sqrt{t}$ , shown in the figure by the dashed red line, where  $c$  is a constant. This upper bound works uniformly for both under and over-parameterized models. It suggests that if we can tolerate a constant multiple of the minimal testing error, then we will have plenty of time to stop the GD algorithm before getting a solution with bad testing error.

Figure 5 shows the results for another random feature model, with feature  $\mathbf{1}_{\{b^T x > 0\}}$ . The phenomenon described above also exists for this model, e.g. the larger the testing error, the longer it takes to reach. In addition, the testing error curves can be uniformly bounded from above by a function proportional to  $\sqrt{t}$ . In this figure, we have translated the curves along the  $y$ -axis to eliminate the influence of the approximation error induced by different values of  $m$ .




Figure 4: Testing error curves for fixed  $n$  ( $n = 500$ ) and different  $m$ 's. The red dashed line is the function  $y = c\sqrt{t}$ , where  $c$  is a constant.




Figure 5: Testing error curves for fixed  $n$  ( $n = 500$ ) and different values of  $m$  for the model with features being given by  $\mathbf{1}_{\{\mathbf{b}^T \mathbf{x} > 0\}}$ , in which  $\mathbf{b}$  and  $\mathbf{x}$  are still sampled from the uniform distribution on  $\mathbb{S}^{d-1}$ . Note that the curves are translated along the  $y$ -axis so that all the curves have the same minimal value.

#### 4. Theoretical analysis

In this section we explain the slow deterioration process of GD for the random feature model from a theoretical perspective. Intuitively, since the parameters of the random feature model depend on the inverse of the Gram matrix, small eigenvalues of the Gram matrix give rises to minimum-norm solutions with large norms and consequently large testing error. At the same time, the GD dynamics proceeds slowly in the directions that correspond to these small eigenvalues. Hence, though the minimal norm solution may generalize poorly, it takes a long time for this behavior to set in along the GD path.

Let  $B = [\mathbf{b}_1, \dots, \mathbf{b}_m]$ , and with an abuse of notation let  $\phi(\mathbf{x}; B) = (\phi(\mathbf{x}; \mathbf{b}_1), \phi(\mathbf{x}; \mathbf{b}_2), \dots, \phi(\mathbf{x}; \mathbf{b}_m))^T$ , recall the trajectory of the GD (4), the prediction function is given by

$$\hat{f}_t(\mathbf{x}) = \phi(\mathbf{x}; B)^T \mathbf{a}(t) = \sum_{i:\lambda_i > 0} \frac{1 - e^{-\lambda_i^2 t / mn}}{\lambda_i} (\mathbf{u}_i^T \mathbf{y})(\mathbf{v}_i^T \phi(\mathbf{x}; B)). \quad (5)$$

Define a function  $d(\lambda, t)$  for  $\lambda \geq 0$  and  $t \geq 0$  by

$$d(\lambda, t) = \begin{cases} \frac{1}{\lambda} \left(1 - e^{-\frac{\lambda^2 t}{mn}}\right), & \text{if } \lambda > 0, \\ 0 & \text{if } \lambda = 0, \end{cases} \quad (6)$$

and let  $D(\Sigma, t) \in \mathbb{R}^{n \times m}$  be a matrix obtained by applying  $d(\cdot, t)$  elementwise to  $\Sigma$ , then (5) can be written as

$$\hat{f}_t(\mathbf{x}) = \mathbf{y}^T U D(\Sigma, t) V^T \phi(\mathbf{x}; B). \quad (7)$$

Therefore, we have

$$\begin{aligned} \|\hat{f}_t(\mathbf{x})\|_{l_2} &= \|\mathbf{y}^T U D(\Sigma, t) V^T \phi(\mathbf{x}; B)\|_{l_2} \\ &\leq \|\mathbf{y}^T U\|_2 \|D(\Sigma, t)\|_2 \|V^T \phi(\mathbf{x}; B)\|_{l_2} \\ &= \|\mathbf{y}\|_2 \|D(\Sigma, t)\|_2 \|\phi(\mathbf{x}; B)\|_{l_2} \\ &\leq \|\mathbf{y}\|_2 \|\phi(\mathbf{x}; B)\|_{l_2} \max_{\lambda \geq 0} \frac{1}{\lambda} \left(1 - e^{-\frac{\lambda^2 t}{mn}}\right), \end{aligned} \quad (8)$$

where  $\|\cdot\|_2$  denotes the 2-norm of vectors for matrices, and  $\|\cdot\|_{l_2}$  is the  $l_2$  norm of functions based on the measure  $\pi$ . For the first part of the right hand side of (8) we know it equals approximately to  $\sqrt{n} \|f^*(\mathbf{x})\|_{l_2}$ , where  $f^*(\cdot)$  is the target function. For the second part, it equals approximately to  $\sqrt{m \mathbb{E}_{\mathbf{b} \sim \mu} \|\phi(\mathbf{x}; \mathbf{b})\|_{l_2}^2}$ . For the third part, we have  $(1 - e^{-\lambda^2 t/mn})/\lambda \leq 1/\lambda$ , and  $(1 - e^{-\lambda^2 t/mn})/\lambda \leq \lambda t/mn$  because of  $1 - e^{-x} \leq x$ . Hence, we have

$$\max_{\lambda \geq 0} \frac{1}{\lambda} \left(1 - e^{-\frac{\lambda^2 t}{mn}}\right) \leq \max_{\lambda \geq 0} \min \left\{ \frac{1}{\lambda}, \frac{\lambda t}{mn} \right\} \leq \sqrt{\frac{t}{mn}}, \quad (9)$$

and the equality is achieved at  $\lambda = \sqrt{mn/t}$ . Hence, approximately we have

$$\|\hat{f}_t(\mathbf{x})\|_{l_2} \leq \|f^*(\mathbf{x})\|_{l_2} \sqrt{\mathbb{E}_{\mathbf{b} \sim \mu} \|\phi(\mathbf{x}; \mathbf{b})\|_{l_2}^2} \sqrt{t}. \quad (10)$$

More precisely, we have the following theorem.

**Theorem 1** *Assume there exists a constant  $M$  such that  $|f^*(\mathbf{x})| \leq M$  and  $|\phi(\mathbf{x}; \mathbf{b})| \leq M$  for any  $\mathbf{b}$  sampled from  $\mu$ . Then, for any  $\delta > 0$ , with probability no less than  $1 - \delta$  over the choice of  $\{\mathbf{x}_i\}$  and  $\{\mathbf{b}_k\}$ , we have*

$$\|\hat{f}_t(\mathbf{x})\|_{l_2} \leq \left( \|f^*(\mathbf{x})\|_{l_2}^2 + \sqrt{\frac{2M^2 \log 2/\delta}{n}} \right)^{1/2} \left( \mathbb{E}_{\mathbf{b} \sim \mu} \|\phi(\mathbf{x}; \mathbf{b})\|_{l_2}^2 + \sqrt{\frac{2M^2 \log 2/\delta}{m}} \right)^{1/2} \sqrt{t}. \quad (11)$$

Theorem 1 can be proved by simply applying the Hoeffding inequality to  $\|\mathbf{y}\|_2^2$  and  $\|\phi(\mathbf{x}; B)\|_{l_2}^2$  and inserting the results into (8). From Theorem 1, we obtain the following corollary on the testing error along the GD trajectory.

**Corollary 2** Consider a random feature model with feature  $\phi(\mathbf{x}; \mathbf{b})$  and target function  $f^*(\mathbf{x})$ . Assume that  $\phi(\mathbf{x}; \mathbf{b})$  and  $f^*(\mathbf{x})$  satisfy the boundedness condition in Theorem 1 with bound  $M$ . Then, for any  $0 \leq t < s$  and  $\delta > 0$ , with probability no less than  $1 - \delta$  over the choice of the training data and random features, we have

$$\begin{aligned} \left\| \hat{f}_s(\mathbf{x}) - f^*(\mathbf{x}) \right\|_{l_2} &\leq \left\| \hat{f}_t(\mathbf{x}) - f^*(\mathbf{x}) \right\|_{l_2} \\ &+ \left( \|f^*(\mathbf{x})\|_{l_2}^2 + \sqrt{\frac{2M^2 \log 2/\delta}{n}} \right)^{1/2} \left( \mathbb{E}_{\mathbf{b} \sim \mu} \|\phi(\mathbf{x}; \mathbf{b})\|_{l_2}^2 + \sqrt{\frac{2M^2 \log 2/\delta}{m}} \right)^{1/2} \\ &\cdot \sqrt{s-t}. \end{aligned} \quad (12)$$

Hence, if there exists  $t_0 > 0$  and  $\epsilon > 0$  such that  $\left\| \hat{f}_{t_0}(\mathbf{x}) - f^*(\mathbf{x}) \right\|_{l_2} \leq \epsilon$ , then for any  $t > t_0$  we have

$$\begin{aligned} \left\| \hat{f}_s(\mathbf{x}) - f^*(\mathbf{x}) \right\|_{l_2} &\leq \left( \|f^*(\mathbf{x})\|_{l_2}^2 + \sqrt{\frac{2M^2 \log 2/\delta}{n}} \right)^{1/2} \left( \mathbb{E}_{\mathbf{b} \sim \mu} \|\phi(\mathbf{x}; \mathbf{b})\|_{l_2}^2 + \sqrt{\frac{2M^2 \log 2/\delta}{m}} \right)^{1/2} \sqrt{t-t_0} \\ &+ \epsilon. \end{aligned} \quad (13)$$

Corollary 2 can be proved by using

$$\left\| \hat{f}_s(\mathbf{x}) - f^*(\mathbf{x}) \right\|_{l_2} \leq \left\| \hat{f}_t(\mathbf{x}) - f^*(\mathbf{x}) \right\|_{l_2} + \left\| \hat{f}_s(\mathbf{x}) - \hat{f}_t(\mathbf{x}) \right\|_{l_2}, \quad (14)$$

and then bounding the second term on the right hand side using similar techniques as for Theorem 1.

These results show that the growth of the testing error in the overfitting regime is controlled by a square root function. The square root function is the result of adding up the exponential contributions from all the small eigenvalues. However, these results are not optimal, in the sense that it treats all the eigenvalues and eigenvectors equally. In reality the leading eigenvalues closely approximate the corresponding eigenvalues of the kernel operator, and the associated eigenvectors are also close to that of the kernel operator evaluated at the training dataset. If we further assume that the target function concentrates on the low-frequency modes, i.e. it mainly lies in the subspace spanned by the leading eigenfunctions of the kernel space, then we will have  $\sum_{i=1}^p (\mathbf{u}_i^T \mathbf{y})^2 \approx \|\mathbf{y}\|^2$  for a small  $p$ , and  $\sum_{i=p+1}^n (\mathbf{u}_i^T \mathbf{y})^2$  is close to 0. Since the leading eigenvalues are large and do not contribute to overfitting, when estimating the overfitting effect we may limit ourselves to the eigenvalues  $\lambda_i$  with  $i \geq p+1$ . This improves the bound for the testing error.

Specifically, let  $K(\cdot, \cdot)$  be the kernel induced by the random features  $\phi(\mathbf{x}; \mathbf{b})$  defined by

$$K(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\mathbf{b} \sim \mu} \phi(\mathbf{x}; \mathbf{b}) \phi(\mathbf{x}'; \mathbf{b}), \quad (15)$$

and let  $\mathcal{K}$  be the corresponding kernel operator:

$$\mathcal{K}f(\mathbf{x}) = \int_{\Omega} K(\mathbf{x}, \mathbf{x}') f(\mathbf{x}') \pi(d\mathbf{x}'). \quad (16)$$

Then for bounded  $\phi$ ,  $\mathcal{K}$  is a trace class operator and has eigenvalues  $\{\mu_i\}_{i=1}^{\infty}$  and eigenfunctions  $\{\psi_i(\cdot)\}_{i=1}^{\infty}$ . We assume that the eigenvalues are in descending order. For convenience, we assume  $m = n$ , and let  $G = \frac{1}{n^2} \Phi \Phi^T$ . By the SVD decomposition of  $\Phi$  we have  $G = \frac{1}{n^2} U \Sigma \Sigma^T U^T$ . Hence, the  $\mathbf{u}_i$ 's are the eigenvectors of  $G$  and the  $\lambda_i^2/n^2$ 's are the eigenvalues. {For the target function, we assume  $f^*$  is the first eigenfunction of  $\mathcal{K}$ .



**Assumption 4.1** Assume

$$f^*(\mathbf{x}) = \psi_1(\mathbf{x}), \quad (17)$$

In the following assumption, we assume that the top eigen-pairs of  $G$  are close to that of  $\mathcal{K}$  (the eigenvectors of  $G$  are close to the eigenfunctions of  $\mathcal{K}$  evaluated at the training data). This is numerically verified in the Appendix.

**Assumption 4.2** Assume there exists a constant  $C$ , such that

$$\begin{aligned} 1 \quad & \left| \frac{\|\mathbf{y}\|^2}{n} - 1 \right| \leq \frac{C}{\sqrt{n}} \\ 2 \quad & \left| \frac{\mathbf{u}_1^T \mathbf{y}}{\sqrt{n}} - 1 \right| < \frac{C}{\sqrt{n}} \text{ and } \left\| \frac{\sqrt{n}}{\lambda_1} \mathbf{v}_1^T \phi(\mathbf{x}; B) - \psi_1(\mathbf{x}) \right\|_{l_2} < \frac{C}{\sqrt{n}} \\ 3 \quad & \left| \left\langle \frac{\sqrt{n}}{\lambda_i} \mathbf{v}_i^T \phi(\mathbf{x}; B), \frac{\sqrt{n}}{\lambda_j} \mathbf{v}_j^T \phi(\mathbf{x}; B) \right\rangle - \delta_{i,j} \right| < \frac{C}{\sqrt{n}}, \text{ for any } 2 \leq i, j \leq \lfloor \sqrt{n} \rfloor, \text{ where } \langle \cdot, \cdot \rangle \text{ is the } \\ & l_2 \text{ inner product of two functions.} \end{aligned}$$

The first assumption in 4.2 is just a Monte-Carlo approximation of  $\|f^*\|_{l_2}^2$  and is easy to verify. The second assumption in 4.2 characterizes the approximation of the largest eigenvalue and the corresponding eigenvector of  $G$  to that of  $\mathcal{K}$ . The third assumes that functions  $\frac{\sqrt{n}}{\lambda_i} \mathbf{v}_i^T \phi(\mathbf{x}; B)$  are nearly orthogonal for  $2 \leq i \leq \lfloor \sqrt{n} \rfloor$ . If we multiply  $\mathbf{u}_k^T$  on both side of the equation  $\Phi \Phi^T = U \Sigma V^T \Phi$ , we get

$$\lambda_k^2 \mathbf{u}_k^T = \mathbf{u}_k^T \Phi \Phi^T = \mathbf{u}_k^T U \Sigma V^T \Phi^T = \lambda_k \mathbf{v}_k^T \Phi. \quad (18)$$

Hence, for  $i = 1, 2, \dots, n$ , we have  $\mathbf{v}_k \phi(\mathbf{x}_i; B) / \lambda_k = \mathbf{u}_{k,i}$ , which gives

$$\frac{1}{n} \sum_{k=1}^n \frac{\sqrt{n}}{\lambda_i} \mathbf{v}_i^T \phi(\mathbf{x}_k; B) \cdot \frac{\sqrt{n}}{\lambda_j} \mathbf{v}_j^T \phi(\mathbf{x}_k; B) = \delta_{i,j}, \quad (19)$$

and this makes the second assumption of 4.2 quite reasonable.

Then, we have the following results on the error between  $\hat{f}_t$  and  $f^*$ :

**Theorem 3** Let  $\hat{\lambda}_i = \frac{\lambda_i}{n}$ . If Assumption 4.2 holds, then for sufficiently large  $n$  such that  $\frac{C}{\sqrt{n}} < 1$ , we have

$$\|\hat{f}_t - f^*\| \leq 3e^{-2\hat{\lambda}_1^2 t} + (5C + 1 + 2\sqrt{C}M d(t))^2 n^{-\frac{1}{2}}, \quad (20)$$

where  $M^2 = \frac{1}{n} \int \|\sigma(B^T \mathbf{x})\|^2 \pi(d\mathbf{x})$ , and

$$d(t) = \min \left\{ \sqrt{t}, \hat{\lambda}_{\lfloor \sqrt{n} \rfloor + 1} t, \hat{\lambda}_n^{-1} \right\}.$$

**Proof** First, by (5) we have

$$\begin{aligned} \|\hat{f}_t - f^*\| &= \left\| \sum_{i=1}^n \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} (\mathbf{u}_i^T \mathbf{y}) (\mathbf{v}_i^T \phi(\mathbf{x}; B)) - \psi_1(\mathbf{x}) \right\| \\ &\leq \left\| \frac{1 - e^{-\hat{\lambda}_1^2 t}}{\lambda_1} (\mathbf{u}_1^T \mathbf{y}) (\mathbf{v}_1^T \phi(\mathbf{x}; B)) - \psi_1(\mathbf{x}) \right\| + \left\| \sum_{i=2}^{\lfloor \sqrt{n} \rfloor} \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} (\mathbf{u}_i^T \mathbf{y}) (\mathbf{v}_i^T \phi(\mathbf{x}; B)) \right\| \\ &\quad + \left\| \sum_{i=\lfloor \sqrt{n} \rfloor + 1}^n \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} (\mathbf{u}_i^T \mathbf{y}) (\mathbf{v}_i^T \phi(\mathbf{x}; B)) \right\| \\ &=: I + J + K. \end{aligned} \quad (21)$$

Next, we estimate  $I$ ,  $J$  and  $K$  separately.

For  $I$ , by Assumption 4.2, we get

$$\begin{aligned}
 I &\leq \left\| e^{-\hat{\lambda}_1^2 t} \psi_1(\mathbf{x}) \right\| + \left\| (1 - e^{-\hat{\lambda}_1^2 t}) \left( \frac{1}{\lambda_1} (\mathbf{u}_1^T \mathbf{y}) (\mathbf{v}_1^T \phi(\mathbf{x}; B)) - \psi_1(\mathbf{x}) \right) \right\| \\
 &\leq e^{-\hat{\lambda}_1^2 t} + \left\| \left( \frac{1}{\lambda_1} (\mathbf{u}_1^T \mathbf{y}) (\mathbf{v}_1^T \phi(\mathbf{x}; B)) - \psi_1(\mathbf{x}) \right) \right\| \\
 &\leq e^{-\hat{\lambda}_1^2 t} + \left\| \left( \frac{\mathbf{u}_1^T \mathbf{y}}{\sqrt{n}} - 1 \right) \frac{\sqrt{n} \mathbf{v}_1^T \phi(\mathbf{x}; B)}{\lambda_1} \right\| + \left\| \frac{\sqrt{n} \mathbf{v}_1^T \phi(\mathbf{x}; B)}{\lambda_1} - \psi_1(\mathbf{x}) \right\| \\
 &\leq e^{-\hat{\lambda}_1^2 t} + \frac{C}{\sqrt{n}} \left( 1 + \frac{C}{\sqrt{n}} \right) + \frac{C}{\sqrt{n}} \\
 &\leq e^{-\hat{\lambda}_1^2 t} + 3 \frac{C}{\sqrt{n}}.
 \end{aligned} \tag{22}$$

For  $J$ , we consider  $J^2$  and have

$$\begin{aligned}
 J^2 &= \sum_{i,j=2}^{\lfloor \sqrt{n} \rfloor} (1 - e^{-\hat{\lambda}_i^2 t}) (1 - e^{-\hat{\lambda}_j^2 t}) \frac{\mathbf{u}_i^T \mathbf{y}}{\sqrt{n}} \frac{\mathbf{u}_j^T \mathbf{y}}{\sqrt{n}} \left\langle \frac{\sqrt{n}}{\lambda_i} \mathbf{v}_i^T \phi(\mathbf{x}; B), \frac{\sqrt{n}}{\lambda_j} \mathbf{v}_j^T \phi(\mathbf{x}; B) \right\rangle \\
 &\leq \sum_{i,j=2}^{\lfloor \sqrt{n} \rfloor} \left| \frac{\mathbf{u}_i^T \mathbf{y}}{\sqrt{n}} \frac{\mathbf{u}_j^T \mathbf{y}}{\sqrt{n}} \right| \left| \left\langle \frac{\sqrt{n}}{\lambda_i} \mathbf{v}_i^T \phi(\mathbf{x}; B), \frac{\sqrt{n}}{\lambda_j} \mathbf{v}_j^T \phi(\mathbf{x}; B) \right\rangle \right|.
 \end{aligned} \tag{23}$$

By Assumption 4.2 we have

$$\begin{aligned}
 J^2 &\leq \frac{C}{\sqrt{n}} \sum_{i,j=2}^{\lfloor \sqrt{n} \rfloor} \left| \frac{\mathbf{u}_i^T \mathbf{y}}{\sqrt{n}} \frac{\mathbf{u}_j^T \mathbf{y}}{\sqrt{n}} \right| + \sum_{i=2}^{\lfloor \sqrt{n} \rfloor} \frac{(\mathbf{u}_i^T \mathbf{y})^2}{n} \\
 &\leq \frac{C+1}{n} \sum_{i=2}^{\lfloor \sqrt{n} \rfloor} (\mathbf{u}_i^T \mathbf{y})^2 \\
 &\leq (C+1) \left( \frac{\|\mathbf{y}\|^2}{n} - \frac{(\mathbf{u}_1^T \mathbf{y})^2}{n} \right) \\
 &\leq (C+1) \left( \left| \frac{\|\mathbf{y}\|^2}{n} - 1 \right| + \left| \frac{(\mathbf{u}_1^T \mathbf{y})^2}{n} - 1 \right| \right) \\
 &\leq (C+1) \frac{4C}{\sqrt{n}}.
 \end{aligned} \tag{24}$$

Hence, for  $J$  we obtain  $J \leq (2C+1)n^{-1/4}$ .

Finally, we estimate  $K$ . Let  $\mathbf{u} = [\mathbf{u}_{\lfloor \sqrt{n} \rfloor + 1}, \dots, \mathbf{u}_n]$  and  $\mathbf{v} = [\mathbf{v}_{\lfloor \sqrt{n} \rfloor + 1}, \dots, \mathbf{v}_n]$ . Then, we can bound  $K$  as

$$K \leq \max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n} \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} \left\| \frac{\tilde{\mathbf{U}}^T \mathbf{y}}{\sqrt{n}} \right\| \cdot \left\| \sqrt{n} \tilde{\mathbf{V}}^T \phi(\mathbf{x}; B) \right\|. \tag{25}$$

For  $\left\| \frac{\tilde{\mathbf{U}}^T \mathbf{y}}{\sqrt{n}} \right\|$ , similar to the estimate of  $J$  we have

$$\left\| \frac{\tilde{\mathbf{U}}^T \mathbf{y}}{\sqrt{n}} \right\| \leq \sqrt{\frac{\|\mathbf{y}\|^2}{n} - \frac{(\mathbf{u}_1^T \mathbf{y})^2}{n}} \leq 2\sqrt{C} n^{-1/4}. \tag{26}$$

For  $\left\| \sqrt{n} \tilde{V}^T \phi(\mathbf{x}; B) \right\|$ , we have

$$\left\| \sqrt{n} \tilde{V}^T \phi(\mathbf{x}; B) \right\| \leq \sqrt{n} \left\| V^T \phi(\mathbf{x}; B) \right\| = \sqrt{n} \|\phi(\mathbf{x}; B)\| \leq nM. \quad (27)$$

For  $\max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n}$ , by (9) we have

$$\max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n} \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} \leq \frac{\sqrt{t}}{n}. \quad (28)$$

On the other hand, it is easy to see

$$\max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n} \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} \leq \max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n} \frac{1}{\lambda_i} \leq \frac{\hat{\lambda}_n^{-1}}{n}, \quad (29)$$

and

$$\max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n} \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} \leq \max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n} \frac{\hat{\lambda}_i t}{n}. \quad (30)$$

Combining (28)-(30), we have

$$\max_{\lfloor \sqrt{n} \rfloor + 1 \leq i \leq n} \frac{1 - e^{-\hat{\lambda}_i^2 t}}{\lambda_i} \leq \frac{1}{n} d(t), \quad (31)$$

and for  $K$  we have estimate

$$K \leq 2\sqrt{C} n^{-1/4} M d(t). \quad (32)$$

Combining the estimates for  $I$ ,  $J$  and  $K$  completes the proof.  $\blacksquare$

From the theorem above, we see that the testing error in time can be divided into three regimes. The first regime is an exponential decay regime governed by the first term on the right hand side of (20). The second regime is a period in which the testing error keeps being small, after the decaying term has been sufficiently reduced and before the effect represented by the term with  $d(t)$  really shows up. In the third regime, the last term on the right hand side of (20) begins to manifest and finally becomes very large (because  $1/\lambda_n$  is very small). The existence of the second regime is caused by the gap between the leading eigenvalue  $\lambda_1$  and the eigenvalues appearing in the third term, and this period is long when this gap is large, which occurs when  $n$  is large. The following corollary roughly characterizes the length of this period under a specific assumption about the decay rate of  $\hat{\lambda}_t$ .

**Corollary 4** *If there exist constant  $C'$  such that  $|\hat{\lambda}_k| \leq \frac{C'}{\sqrt{k}}$  holds for any  $1 \leq k \leq \lfloor \sqrt{n} \rfloor + 1$ . Then, there exist constants  $c_1$  and  $c_2$  (which may depend on  $\hat{\lambda}_1$ ) such that*

$$\|\hat{f}_t - f^*\|^2 \leq \frac{c_1}{\sqrt{n}},$$

when  $c_2 \log n \leq t \leq c_2 n^{\frac{1}{4}}$ . Hence, the length of the second regime described above is at least in the order of  $n^{1/4}$ .

**Proof** Let  $c_2 = \frac{1}{4\lambda_1^2}$ , then, for  $t$  that satisfies  $c_2 \log n \leq t \leq c_2 n^{\frac{1}{4}}$ , we have

$$\|\hat{f}_t - f^*\| \leq n^{-\frac{1}{4}} + (5C + 1 + 2\sqrt{C}M\hat{\lambda}_{\lfloor\sqrt{n}\rfloor+1}c_2n^{\frac{1}{4}})n^{-\frac{1}{4}}. \quad (33)$$

By the condition that  $|\hat{\lambda}_k| \leq \frac{C'}{\sqrt{k}}$  we have

$$\|\hat{f}_t - f^*\| \leq n^{-\frac{1}{4}}(1 + 5C + 1 + 2\sqrt{C}MC'c_2). \quad (34)$$

The proof is completed by taking  $c_1 = 2 + 5C + 2\sqrt{C}C'c_2M$ . ■

**Remark 5** *Theorem 3 shows that the estimates of the testing error can be refined if the target function is band-limited, and the leading eigen-pairs of the Gram matrix align well with that of the kernel operator. In the appendix we numerically demonstrate that the assumptions for this theorem are likely to hold in practice.*

**Remark 6** *The testing error of the final solution (the minimal norm solution) is largely controlled by the smallest eigenvalue of the Gram matrix. In the appendix we demonstrate by numerical experiments that the minimal eigenvalue of the Gram matrix is very small only when  $m \approx n$ .*

## 5. Discussions

In this paper, we studied the double-sided effect of the small eigenvalues of the Gram matrix in a random feature model. An obvious question is how small the smallest eigenvalues are, In the appendix, we provide numerical evidence that the smallest eigenvalue approximately obeys the predictions of the Marchenko-Pastur distribution in the random matrix theory when  $\gamma$  is close to 1.

The second obvious question is why should the smallest eigenvalues be so small when  $\gamma = 1$ , i.e. why should the Gram matrix be almost singular when  $\gamma = 1$ . While we do not have a simple answer to this question, we can address the related question: Why should the Gram matrix be non-singular away from  $\gamma = 1$ ? The reason is that for large values of  $\gamma$ , the spectrum of the Gram matrix converges to the top spectrum of the corresponding kernel operator, which is obviously non-singular. This can be seen from the last figure in Figure 6. The same is true for very small values of  $\gamma$  if we consider the ReLU feature  $\sigma(\mathbf{b}^T \mathbf{x})$ . In this case, the non-zero eigenvalues of  $G = \Phi\Phi^T/mn$  equal to the eigenvalues of  $\Phi^T\Phi/mn$ , and the latter can be viewed as a Gram matrix with  $\mathbf{x}$  being the features and  $\mathbf{b}$  being the data, and  $\gamma$  much larger than 1, due to the symmetry of  $\mathbf{b}$  and  $\mathbf{x}$ .

The third question is how relevant this phenomenon is for other models, such as the two-layer neural network model? We guess similar “resonance” phenomenon also exists, but extensive experimental and theoretical study is left for future work.

## References

- Madhu S Advani and Andrew M Saxe. High-dimensional dynamics of generalization error in neural networks. *arXiv preprint arXiv:1710.03667*, 2017.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019a.

- Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *arXiv preprint arXiv:1903.07571*, 2019b.
- Luigi Carratino, Alessandro Rudi, and Lorenzo Rosasco. Learning with SGD and random features. In *Advances in Neural Information Processing Systems*, pages 10192–10203, 2018.
- Zeyu Deng, Abba Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *arXiv preprint arXiv:1911.05822*, 2019.
- Weinan E, Chao Ma, and Lei Wu. A comparative analysis of optimization and generalization properties of two-layer neural network and random feature models under gradient descent dynamics. *arXiv preprint arXiv:1904.04326*, *Science China Mathematics*, pages 1–24, 2020.
- A Erdlyi, W Magnus, F Oberhettinger, and FG Tricomi. Tables of integral transforms. vol. ii. based, in part, on notes left by harry bateman, 1954.
- Christopher Frye and Costas J Efthimiou. Spherical harmonics in  $p$  dimensions. *arXiv preprint arXiv:1205.3548*, 2012.
- Ganesh Kini and Christos Thrampoulidis. Analytic study of double descent in binary classification: The impact of loss. *arXiv preprint arXiv:2001.11572*, 2020.
- Vladimir Alexandrovich Marchenko and Leonid Andreevich Pastur. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik*, 114(4):507–536, 1967.
- Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355*, 2019.
- Stefano Spigler, Mario Geiger, Stéphane d’Ascoli, Levent Sagun, Giulio Biroli, and Matthieu Wyart. A jamming transition from under-to over-parametrization affects generalization in deep learning. *Journal of Physics A: Mathematical and Theoretical*, 52(47):474001, 2019.
- Arun Suggala, Adarsh Prasad, and Pradeep K Ravikumar. Connecting optimization and regularization paths. In *Advances in Neural Information Processing Systems*, pages 10608–10619, 2018.
- Bo Xie, Yingyu Liang, and Le Song. Diverse neural network learns true target functions. In *Artificial Intelligence and Statistics*, pages 1216–1224, 2017.
- Zhi-Qin John Xu. Understanding training and generalization in deep learning by fourier analysis. *arXiv preprint arXiv:1808.04295*, 2018.
- Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019.
- Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Constructive Approximation*, 26(2):289–315, 2007.
- Yaoyu Zhang, Zhi-Qin John Xu, Tao Luo, and Zheng Ma. Explicitizing an implicit bias of the frequency principle in two-layer neural networks. *arXiv preprint arXiv:1905.10264*, 2019.

## Appendix A. Spectrum of the Gram matrix

Let  $\pi$  be the uniform distribution on the unit sphere  $\mathbb{S}^{d-1}$ , where  $d \geq 3$  is the dimension. The random feature model we consider has features  $\phi(\mathbf{x}; \mathbf{b}) = \sigma(\mathbf{b}^T \mathbf{x})$ , where  $\mathbf{x} \sim \pi$ ,  $\mathbf{b} \sim \pi$ , and  $\sigma(\cdot)$  is the ReLU function. Then, the kernel defined by (15) have a closed form

$$k(\mathbf{x}, \mathbf{x}') = \sqrt{1 - (\mathbf{x}^T \mathbf{x}')^2} + \mathbf{x}^T \mathbf{x}' (\pi - \arccos(\mathbf{x}^T \mathbf{x}')). \quad (35)$$

We study the following three quantities.

- Gram matrix  $G = (G_{i,j})$  with  $G_{i,j} = \frac{1}{nm} \sum_{s=0}^m \sigma(\mathbf{b}_s^T \mathbf{x}_i) \sigma(\mathbf{b}_s^T \mathbf{x}_j)$ .
- Kernel matrix  $K = (K_{i,j})$  with  $K_{i,j} = \frac{1}{n} k(\mathbf{x}_i, \mathbf{x}_j)$ .
- The kernel operator  $\mathcal{K}$  defined by

$$\mathcal{K}f(\mathbf{x}) = \int_{\mathbb{S}^{d-1}} k(\mathbf{x}, \mathbf{x}') f(\mathbf{x}') d\pi(\mathbf{x}').$$

In Figure 6, the spectra of Gram matrices for various value of  $m$  are displayed. As a comparison, the spectra of the corresponding kernel operator and kernel matrix are also plotted. Let  $\gamma = m/n$ . We have the following observations.

- The spectrum of the kernel operator is stage-like.
- The large eigenvalues of kernel matrix approximate the corresponding eigenvalues of kernel operator very well. The small eigenvalues are away from the smallest eigenvalues of kernel operator, but still relatively large.
- The spectrum of the Gram matrix is close to the spectrum of the kernel matrix when  $\gamma$  is large. For example, the two spectra are almost indistinguishable at  $\gamma = 8$ .
- The spectrum of the Gram matrix at  $\gamma = 1$  has a very long tail at 0, i.e. the tail is extremely small and have outliers close to zero. However, the tail of the spectrum for  $\gamma \neq 1$  does not have outliers near zero.

In the following, we provide some theoretical evidence why the spectrum of behave like this.

### A.1. The stage-like spectrum of the kernel operator

For this special case, in the following we show that the spectrum of kernel operator  $\mathcal{K}$  has explicit formula. Since  $k(\mathbf{x}, \mathbf{x}')$  only depends on the inner product of its inputs, and  $\mathbf{x}$  are uniformly distributed on the sphere, we know that the eigenfunctions of  $K$  are spherical harmonics, and the eigenvalues for spherical harmonics with the same frequency are the same. Let  $\{\psi_j(\mathbf{x})\}_{j=0}^{\infty}$  be the sequence of spherical harmonics on  $\mathbb{S}^{d-1}$ , with non-decreasing frequencies, and  $\mu_j$  be the corresponding eigenvalues, then  $k(\mathbf{x}, \mathbf{x}')$  can be represented as

$$k(\mathbf{x}, \mathbf{x}') = \sum_{j=0}^{\infty} \mu_j \psi_j(\mathbf{x}) \psi_j(\mathbf{x}'). \quad (36)$$

For  $\mathcal{K}$ , we have the following theorem about its eigenvalues and eigenfunctions.

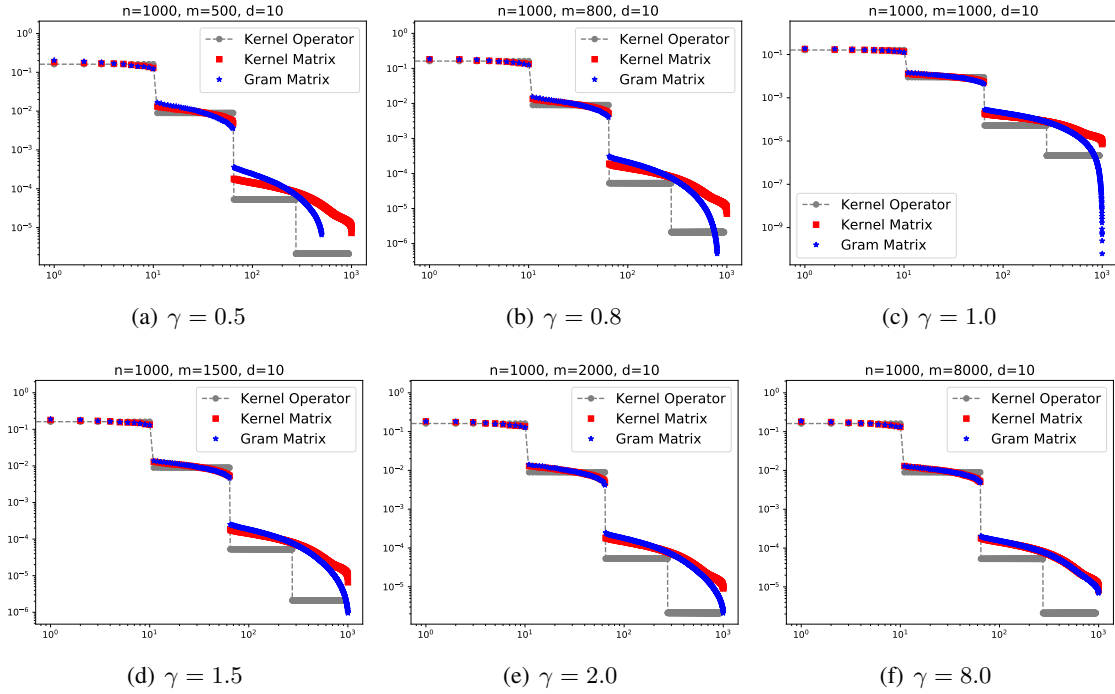


Figure 6: Comparison of the spectra of the kernel operator, kernel matrix and the corresponding Gram matrix. We see that the spectrum of Gram matrix consist of large discrete spectrum and small “continuous” spectrum. The former are very close to the top eigenvalues of the kernel operator, while the latter is rather random especially when  $\gamma \approx 1$ . Moreover, as  $\gamma \rightarrow \infty$ , the spectrum of the Gram matrix converges to the spectrum of the kernel matrix.

**Theorem 7** Let  $\{Y_{n,i} : n \geq 0, 1 \leq i \leq N(d,n)\}$  be the spherical harmonics on  $\mathbb{S}^{d-1}$ , with  $N(d,0) = 1$  and

$$N(d,n) = \frac{2n+d-2}{n} \binom{n+d-3}{n-1}, \quad (37)$$

for  $n \geq 1$ . Then,  $Y_{n,i}(\cdot)$  are eigenfunctions of  $\mathcal{K}$  for any  $(n,i)$ , and the corresponding eigenvalues only depend on  $n$ . Let the eigenvalue associated with  $Y_{n,i}$  be  $\lambda_n$ , then we have

$$\lambda_0 = \frac{2\sqrt{\pi}d\Gamma(\frac{d}{2})}{\Gamma(d)\Gamma(\frac{d-1}{2})}, \quad (38)$$

and  $\lambda_n = C(d)\Lambda(d,n)$  for  $n \geq 1$ , where

$$C(d) = 2^{\frac{d-5}{2}} \pi^{\frac{2d+3}{4}} d(d-2) \left[ \frac{\Gamma(\frac{d-2}{2})\Gamma(d-1)}{\Gamma(\frac{d-1}{2})\Gamma(\frac{d}{2})} \right]^{\frac{1}{2}}, \quad (39)$$

$$\Lambda(d,n) = \frac{2^{n-\frac{1}{2}} \Gamma(\frac{n+d-2}{2})}{(n+d-3)!(n+d-1)!\Gamma(\frac{3-n}{2})^2 \Gamma(\frac{n+d-1}{2})}. \quad (40)$$

**Proof** By the close form of the kernel (35), the kernel  $k(\mathbf{x}, \mathbf{x}')$  only depends on the inner product of  $\mathbf{x}$  and  $\mathbf{x}'$ . Hence, by (Xie et al., 2017), we know that spherical harmonics are eigenfunctions of  $\mathcal{K}$ , and the eigenvalues are the same for spherical harmonics with the same order  $n$ . Hence,  $K(\mathbf{x}, \mathbf{x}')$  has the following decomposition,

$$K(\mathbf{x}, \mathbf{x}') = \sum_{n=0}^{\infty} \lambda_n \sum_{i=1}^{N(d,n)} Y_{n,i}(\mathbf{x})Y_{n,i}(\mathbf{x}'). \quad (41)$$

A direct integral gives (38). For  $n \geq 1$ , let  $P_n(t)$  be the Legendre Polynomials defined in (Frye and Efthimiou, 2012), then by Theorem 4.11 thereof,

$$P_n(\mathbf{x}^T \mathbf{x}') = \frac{\Omega_{d-1}}{N(d, n)} \sum_{i=1}^{N(d,n)} Y_{n,i}(\mathbf{x})Y_{n,i}(\mathbf{x}'), \quad (42)$$

where  $\Omega_{d-1}$  is the surface area of  $\mathbb{S}^{d-1}$ . Hence, combining (41) and (42), we have

$$\lambda_n = \frac{1}{\Omega_{d-1}} \int \int_{\mathbf{x}, \mathbf{x}'} K(\mathbf{x}, \mathbf{x}') P_n(\mathbf{x}^T \mathbf{x}') d\Omega_{d-1} d\Omega_{d-1} \quad (43)$$

$$= \frac{1}{\Omega_{d-1}} \int_{\mathbf{x}} K(\mathbf{x}, \mathbf{x}') P_n(\mathbf{x}^T \mathbf{x}') d\Omega_{d-1} \quad (44)$$

$$= \frac{1}{\Omega_{d-1}} \int_{-1}^1 k(t) P_n(t) (1-t^2)^{\frac{d-3}{2}} dt, \quad (45)$$

where we let  $k(t) = \sqrt{1-t^2} + t(\pi - \arccos(t))$ . By (Frye and Efthimiou, 2012),  $P_n$  and  $P_m$  are orthogonal when  $m \neq n$ , and

$$\int_{-1}^1 (1-t^2)^{\frac{d-3}{2}} P_n(t)^2 dt = \frac{\Omega_{d-1}}{\Omega_{d-2} N(d, n)}. \quad (46)$$

Hence, if we let

$$C_n(t) = \left[ \frac{\Omega_{d-2} N(d, n) 2^{3-d} \pi \Gamma(n+d-2)}{\Omega_{d-1} (n + \frac{d-2}{2}) \Gamma(n+1) \Gamma(\frac{d-2}{2})^2} \right]^{\frac{1}{2}} P_n(t), \quad (47)$$

then  $C_n(t)$  are Gegenbauer polynomials defined through the generating function

$$\sum_{n=0}^{\infty} C_n(t) s^n = \frac{1}{(1-2st+s^2)^{\frac{d-2}{2}}}. \quad (48)$$

According to (48) and (Erdlyi et al., 1954), we can compute the following integrals,

$$\int_{-1}^1 (1-t^2)^{\frac{d-2}{2}} C_n(t) dt = \frac{\pi^{3/2} 2^{n-2} (d-2) \Gamma(\frac{n+d-2}{2})}{n! \Gamma(\frac{1-n}{2}) \Gamma(\frac{3-n}{2}) \Gamma(\frac{n+d+1}{2})}, \quad (49)$$

and

$$\int_{-1}^1 (1-t^2)^{\frac{d-3}{2}} t(\pi - \arccos t) C_n(t) dt = \frac{\pi^{3/2} 2^{n-3} (d-2)(n^2 + (d-2)n + 1) \Gamma(\frac{n+d-2}{2})}{(n+d-1)n! \Gamma(\frac{3-n}{2})^2 \Gamma(\frac{n+d+1}{2})}. \quad (50)$$



Combining (49) and (50), we have

$$\int_{-1}^1 (1-t^2)^{\frac{d-3}{2}} k(t) C_n(t) dt = \frac{\pi^{3/2} d(d-2) 2^{n-2} \Gamma(\frac{n+d-2}{2})}{(n+d-1)^2 n! \Gamma(\frac{3-n}{2})^2 \Gamma(\frac{n+d-1}{2})}. \quad (51)$$

Finally, combining (51) with (45) and (47) gives the results of  $\lambda_n$ . ■

**Remark 8** *From Theorem 7 we know that the spectrum of  $\mathcal{K}$  is stage-like, and the stage becomes wider as the dimension  $d$  gets larger. And when  $d$  is fixed, the gap between stages is large. Specifically, by (40) we have*

$$\lambda_{n+2} = \frac{(n-1)^2}{(n+d-1)^2(n+d+1)(n+d)} \lambda_n \leq \frac{\lambda_n}{(n+d)^2}. \quad (52)$$

### A.2. The smallest eigenvalues

To gain some insights about the small eigenvalues, we draw some inspiration from the well-known Marchenko-Pastur distribution, which characterizes the spectrum of random matrix  $Y_n = \frac{1}{n} X X^T$ , where  $X \in \mathbb{R}^{n \times m}$  with the entries being *i.i.d.* random variables with mean 0 and variance 1. Let  $\mu_n(\lambda) = \frac{1}{m} \sum_{j=1}^m \delta(\lambda - \lambda_j(Y_n))$  denote the spectrum of  $Y_n$ . Random matrix theory shows that  $\mu_n$  converges to the the following Marchenko-Pastur (MP) (Marchenko and Pastur, 1967) distribution as  $m, n \rightarrow \infty$ ,

$$\mu(\lambda) = \begin{cases} (1 - \frac{1}{\gamma})\delta(\lambda) + v_{1/\gamma}(\lambda), & \text{if } \gamma > 1 \\ v_{\gamma}(\lambda), & \text{if } \gamma \leq 1, \end{cases}$$

where

$$v_{\gamma}(\lambda) = \frac{1}{2\pi} \frac{\sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}}{\lambda} \mathbf{1}_{\lambda \in [\lambda_-, \lambda_+]},$$

and  $\lambda_{\pm} = (1 \pm \sqrt{\gamma})^2$ .

The smallest eigenvalue is given by  $\lambda_-$ . When  $\gamma = 1$ , we have  $\lambda_- = 0$  and  $\mu(\lambda) \propto \frac{1}{\sqrt{\lambda}}$  at  $\lambda \approx 0$ . This implies that the smallest eigenvalues is zero and has a long tail. This is consistent with the small part of the spectrum of the Gram matrix at  $\gamma = 1$  shown in Figure 6(c). Assuming that the MP distribution can be used to characterize the tail, the smallest eigenvalue  $\lambda_n(\gamma)$  of the Gram matrix at  $\gamma \approx 1$  should then obey

$$\lambda_n(\gamma) \approx c_{n,d} \begin{cases} (1 - \sqrt{\gamma})^2 & \text{if } \gamma \leq 1 \\ (1 - \sqrt{1/\gamma})^2 & \text{if } \gamma > 1, \end{cases} \quad (53)$$

where  $c_{n,d}$  is a constant that depends on  $n, d$ . The Figure 7 illustrates the difference between the true smallest eigenvalues and the prediction in (53). We see that the prediction is pretty accurate when  $\gamma \approx 1$ .




Figure 7: The smallest eigenvalues of the Gram matrices for different values of  $\gamma$ . In this experiment,  $n = 1000$ ,  $d = 10$ , and for each  $\gamma$ , we report the mean values of 10 independent experiments.