

Frequency Principle: Fourier Analysis Sheds Light on Deep Neural Networks

Zhi-Qin John Xu^{1,*}, Yaoyu Zhang², Tao Luo¹, Yanyang Xiao³ and Zheng Ma¹

¹ School of Mathematical Sciences and Institute of Natural Sciences, Shanghai Jiao Tong University, Shanghai 200240, China.

² School of Mathematics, Institute for Advanced Study, Princeton, NJ 08540, USA.

³ Brain Cognition and Brain Disease Institutes of Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen 518055, China.

Received 9 May 2020; Accepted (in revised version) 16 July 2020

Abstract. We study the training process of Deep Neural Networks (DNNs) from the Fourier analysis perspective. We demonstrate a very universal Frequency Principle (F-Principle) — DNNs often fit target functions from low to high frequencies — on high-dimensional benchmark datasets such as MNIST/CIFAR10 and deep neural networks such as VGG16. This F-Principle of DNNs is opposite to the behavior of Jacobi method, a conventional iterative numerical scheme, which exhibits faster convergence for higher frequencies for various scientific computing problems. With theories under an idealized setting, we illustrate that this F-Principle results from the smoothness/regularity of the commonly used activation functions. The F-Principle implies an implicit bias that DNNs tend to fit training data by a low-frequency function. This understanding provides an explanation of good generalization of DNNs on most real datasets and bad generalization of DNNs on parity function or a randomized dataset.

AMS subject classifications: 68Q32, 65N06, 68T01

Key words: Deep learning, training behavior, generalization, Jacobi iteration, Fourier analysis.

1 Introduction

Understanding the training process of Deep Neural Networks (DNNs) is a fundamental problem in the area of deep learning. We find a common behavior of the gradient-based training process of DNNs, that is, a Frequency Principle (F-Principle):

*Corresponding author. Email addresses: xuzhiqin@sjtu.edu.cn (Z.-Q. J. Xu), yaoyu@ias.edu (Y. Zhang), luotao41@sjtu.edu.cn (T. Luo), yy.xiao@siat.ac.cn (Y. Xiao), Wudy_MZ@sjtu.edu.cn (Z. Ma)

DNNs often fit target functions from low to high frequencies during the training process.

In another word, at the early stage of training, the low-frequencies are fitted and as iteration steps of training increase, the high-frequencies are fitted. For example, when a DNN is trained to fit $y = \sin(x) + \sin(2x)$, its output would be close to $\sin(x)$ at early stage and as training goes on, its output would be close to $\sin(x) + \sin(2x)$. Along with our previous works in [34], this paper is one of works that first discovery the F-Principle. In the same time, another group[†] independently found the F-Principle (or spectral bias) [26]. F-Principle was verified empirically in synthetic low-dimensional data with MSE loss during DNN training [26, 34]. However, in deep learning, empirical phenomena could vary from one network structure to another, from one dataset to another and could exhibit significant difference between synthetic data and high-dimensional real data. Therefore, the universality of the F-Principle remains an important problem for further study. Especially for high-dimensional real problems, because the computational cost of high-dimensional Fourier transform is prohibitive in practice, it is of great challenge to demonstrate the F-Principle. On the other hand, the mechanism underlying the F-Principle and its implication to the application of DNNs, e.g., design of DNN-based PDE solver, as well as their generalization ability are also important open problems to be addressed.

In this work, we design two methods, i.e., projection and filtering methods, to show that the F-Principle exists in the training process of DNNs for high-dimensional benchmarks, i.e., MNIST [22], CIFAR10 [21]. The settings we have considered are i) different DNN architectures, e.g., fully-connected network, convolutional neural network (CNN), and VGG16 [28]; ii) different activation functions, e.g., tanh and rectified linear unit (ReLU); iii) different loss functions, e.g., cross entropy, mean squared error (MSE), and loss energy functional in variational problems. These results demonstrate the universality of the F-Principle.

To facilitate the designs and applications of DNN-based schemes, we characterize a stark difference between DNNs and the Jacobi method, a conventional numerical scheme exhibiting the opposite convergence behavior — faster convergence for higher frequencies. Numerical methods [5,7,31], such as well-known multigrid method [5,31], are developed to accelerate the convergence for low frequency. As the DNN-based schemes have potential to solve high-dimensional problems [8, 13–15, 17, 19, 30, 32], the low-frequency bias of DNN can be adopted to accelerate the convergence of low frequencies for computational problems.

We also intuitively explain with theories under an idealized setting how the smoothness/regularity of commonly used activation functions contributes to the F-Principle. Note that this mechanism is rigorously demonstrated for DNNs of general settings in a subsequent work [23]. Finally, we discuss that the F-Principle provides an understanding of good generalization of DNNs in many real datasets [37] and poor generalization in learning the parity function [25,27], that is, the F-Principle which implies that DNNs prefer low frequencies, is consistent with the property of low frequencies dominance in

[†]We acknowledge each other after communication with authors in [26]

many real datasets, e.g., MNIST/CIFAR10, but is different from the parity function whose spectrum concentrates on high frequencies. Compared with previous studies, our main contributions are as follows:

1. By designing both the projection and filtering methods, we consistently demonstrate the F-Principle for MNIST/CIFAR10 over various architectures such as VGG16 and various loss functions.
2. For the application of solving differential equations, we show that (i) conventional numerical schemes learn higher frequencies faster whereas DNNs learn lower frequencies faster by the F-Principle, (ii) convergence of low frequencies can be greatly accelerated with DNN-based schemes.
3. We present theories under an idealized setting to illustrate how smoothness/regularity of activation function contributes to the F-Principle.
4. We discuss in detail the implication of the F-Principle to the generalization of DNNs that DNNs are implicitly biased towards a low frequency function and provide an explanation of good and poor generalization of DNNs for low and high frequency dominant target functions, respectively.

2 Frequency principle

The concept of “frequency” is central to the understanding of F-Principle. In this paper, the “frequency” means *response frequency* NOT image (or input) frequency as explained in the following.

Image (or input) frequency (NOT used in the paper): Frequency of 2-d function $I: \mathbb{R}^2 \rightarrow \mathbb{R}$ representing the intensity of an image over pixels at different locations. This frequency corresponds to the rate of change of intensity *across neighbouring pixels*. For example, an image of constant intensity possesses only the zero frequency, i.e., the lowest frequency, while a sharp edge contributes to high frequencies of the image.

Response frequency (used in the paper): Frequency of a general Input-Output mapping f . For example, consider a simplified classification problem of partial MNIST data using only the data with label 0 and 1, $f(x_1, x_2, \dots, x_{784}): \mathbb{R}^{784} \rightarrow \{0, 1\}$ mapping 784-d space of pixel values to 1-d space, where x_j is the intensity of the j -th pixel. Denote the mapping’s Fourier transform as $\hat{f}(k_1, k_2, \dots, k_{784})$. The frequency in the coordinate k_j measures the rate of change of $f(x_1, x_2, \dots, x_{784})$ with respect to x_j , i.e., the intensity of the j -th pixel. If f possesses significant high frequencies for large k_j , then a small change of x_j in the image might induce a large change of the output (e.g., adversarial example). For a dataset with multiple classes, we can similarly define frequency for each output dimension. For real data, the response frequency is rigorously defined via the standard nonuniform discrete Fourier transform (NUDFT) as follows.

In all our experiments, we consistently consider the response frequency defined for the mapping function g between inputs and outputs, say $\mathbb{R}^d \rightarrow \mathbb{R}$ and any $k \in \mathbb{R}^d$ via the standard nonuniform discrete Fourier transform (NUDFT)

$$\hat{g}_k = \frac{1}{n} \sum_{i=0}^{n-1} g(x_i) e^{-i2\pi k \cdot x_i}, \tag{2.1}$$

which is a natural estimator of frequency composition of g . As $n \rightarrow \infty$, $\hat{g}_k \rightarrow \int g(x) e^{-i2\pi k \cdot x} \nu(x) dx$, where $\nu(x)$ is the data distribution.

We restrict all the evaluation of Fourier transform in our experiments to NUDFT of $\{y_i\}_{i=0}^{n-1}$ at $\{x_i\}_{i=0}^{n-1}$ for the following practical reasons.

- (i) The information of target function is only available at $\{x_i\}_{i=0}^{n-1}$ for training.
- (ii) It allows us to perform the convergence analysis. As $t \rightarrow \infty$, in general, $h(x_i, t) \rightarrow y_i$ for any i ($h(x_i, t)$ is the DNN output), leading to $\hat{h}_k \rightarrow \hat{y}_k$ for any k . Therefore, we can analyze the convergence at different k by evaluating $\Delta_F(k) = |\hat{h}_k - \hat{y}_k| / |\hat{y}_k|$ during the training. If we use a different set of data points for frequency evaluation of DNN output, then $\Delta_F(k)$ may not converge to 0 at the end of training.
- (iii) \hat{y}_k faithfully reflects the frequency structure of training data $\{x_i, y_i\}_{i=0}^{n-1}$. Intuitively, high frequencies of \hat{y}_k correspond to sharp changes of output for some nearby points in the training data. Then, by applying a Gaussian filter and evaluating still at $\{x_i\}_{i=0}^{n-1}$, we obtain the low frequency part of training data with these sharp changes (high frequencies) well suppressed.

In practice, it is impossible to evaluate and compare the convergence of all $k \in \mathbb{R}^d$ even with a proper cutoff frequency for a very large d of $\mathcal{O}(10^2)$ (MNIST) or $\mathcal{O}(10^3)$ (CIFAR10) due to curse of dimensionality. Therefore, we propose the projection approach, i.e., fixing k at a specific direction, and the filtering approach as detailed in Section 3 and 4, respectively.

Frequency Principle: *DNNs often fit target functions from low to high (response) frequencies during the training process.* The F-Principle is rigorously defined through the frequency defined by the Fourier transform (Eq. (2.1)) and the converging speed defined by the relative error. By using high-dimensional real datasets, we then experimentally demonstrate the F-Principle at the levels of both individual frequencies (projection method) and coarse-grained frequencies (filtering method).

To illustrate the phenomenon of F-Principle, we use 1-d synthetic data to show the evolution of relative training error at different frequencies during the training of DNN. We train a DNN to fit a 1-d target function $f(x) = \sin(x) + \sin(3x) + \sin(5x)$ of three frequency components. On $n = 201$ evenly spaced training samples, i.e., $\{x_i\}_{i=0}^{n-1}$ in $[-3.14, 3.14]$, the discrete Fourier transform (DFT) of $f(x)$ or the DNN output (denoted by $h(x)$) is computed by $\hat{f}_k = \frac{1}{n} \sum_{i=0}^{n-1} f(x_i) e^{-i2\pi ik/n}$ and $\hat{h}_k = \frac{1}{n} \sum_{i=0}^{n-1} h(x_i) e^{-i2\pi ik/n}$, where k is

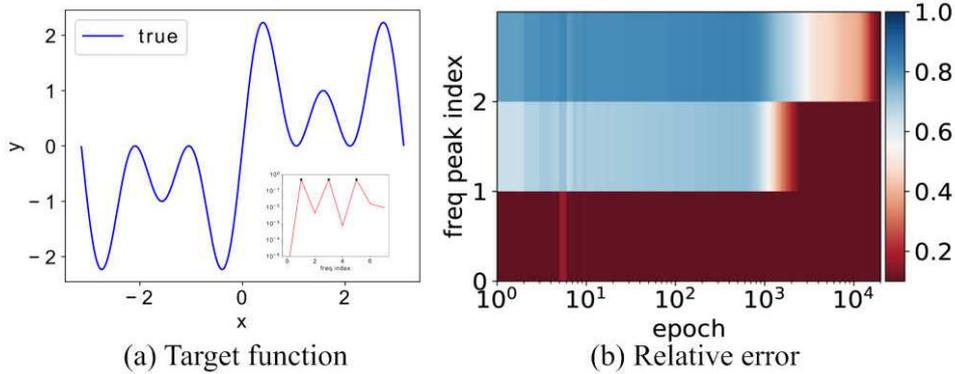


Figure 1: 1d input. (a) $f(x)$. Inset : $|\hat{f}(k)|$. (b) $\Delta_F(k)$ of three important frequencies (indicated by black dots in the inset of (a)) against different training epochs. The parameters of the DNN is initialized by a Gaussian distribution with mean 0 and standard deviation 0.1. We use a tanh-DNN with widths 1-8000-1 with full batch training. The learning rate is 0.0002. The DNN is trained by Adam optimizer [20] with the MSE loss function.

the frequency. As shown in Fig. 1(a), the target function has three important frequencies as we design (black dots at the inset in Fig. 1(a)). To examine the convergence behavior of different frequency components during the training with MSE, we compute the relative difference between the DNN output and the target function for the three important frequencies k 's at each recording step, that is, $\Delta_F(k) = |\hat{h}_k - \hat{f}_k| / |\hat{f}_k|$, where $|\cdot|$ denotes the norm of a complex number. As shown in Fig. 1(b), the DNN converges the first frequency peak very fast, while converging the second frequency peak much slower, followed by the third frequency peak.

Next, we investigate the F-Principle on real datasets with more general loss functions other than MSE which was the only loss studied in the previous works [26, 34].

3 F-Principle in MNIST/CIFAR10 through projection method

Real datasets are very different from synthetic data used in previous studies. In order to utilize the F-Principle to understand and better use DNNs in real datasets, it is important to verify whether the F-Principle also holds in high-dimensional real datasets.

In the following experiments, we examine the F-Principle in a training dataset of $\{(x_i, y_i)\}_{i=0}^{n-1}$ where n is the size of dataset. $x_i \in \mathbb{R}^d$ is a vector representing the image and $y_i \in \{0, 1\}^{10}$ is the output (a one-hot vector indicating the label for the dataset of image classification). d is the dimension of the input ($d = 784$ for MNIST and $d = 32 \times 32 \times 3$ for CIFAR10). Since the high dimensional discrete Fourier transform (DFT) requires prohibitively high computational cost, in this section, we only consider one direction in the Fourier space through a projection method for each examination.

3.1 Examination method: Projection

For a dataset $\{(x_i, y_i)\}_{i=0}^{n-1}$ we consider one entry of 10-d output, denoted by $y_i \in \mathbb{R}$. The high dimensional discrete non-uniform Fourier transform of $\{(x_i, y_i)\}_{i=0}^{n-1}$ is $\hat{y}_k = \frac{1}{n} \sum_{i=0}^{n-1} y_i \exp(-i2\pi k \cdot x_i)$. The number of all possible k grows exponentially on dimension d . For illustration, in each examination, we consider a direction of k in the Fourier space, i.e., $k = k\mathbf{p}_1$, \mathbf{p}_1 is a chosen and fixed unit vector, hence $|k| = k$. Then we have $\hat{y}_k = \frac{1}{n} \sum_{i=0}^{n-1} y_i \exp(-i2\pi(\mathbf{p}_1 \cdot x_i)k)$, which is essentially the 1-d Fourier transform of $\{(x_{\mathbf{p}_1, i}, y_i)\}_{i=0}^{n-1}$, where $x_{\mathbf{p}_1, i} = \mathbf{p}_1 \cdot x_i$ is the projection of x_i on the direction \mathbf{p}_1 [4]. For each training dataset, \mathbf{p}_1 is chosen as the first principle component of the input space. To examine the convergence behavior of different frequency components during the training, we compute the relative difference between the DNN output and the target function for selected important frequencies k 's at each recording step, that is, $\Delta_F(k) = |\hat{h}_k - \hat{y}_k| / |\hat{y}_k|$, where \hat{y}_k and \hat{h}_k are 1-d Fourier transforms of $\{y_i\}_{i=0}^{n-1}$ and the corresponding DNN output $\{h_i\}_{i=0}^{n-1}$, respectively, along \mathbf{p}_1 . Note that each response frequency component, \hat{h}_k , of DNN output evolves as the training goes.

3.2 MNIST/CIFAR10

In the following, we show empirically that the F-Principle is exhibited in the selected direction during the training process of DNNs when applied to MNIST/CIFAR10 with cross-entropy loss. The network for MNIST is a fully-connected tanh DNN (784-400-200-10) and for CIFAR10 is two ReLU convolutional layers followed by a fully-connected DNN (800-400-400-400-10). We consider one of the 10-d outputs in each case using non-uniform Fourier transform. After training, the training accuracy is 0.951 and 0.98 for MNIST and CIFAR10, respectively. The test accuracy is 0.963 and 0.72 for MNIST and CIFAR10, respectively. As shown in Fig. 2(a) and 2(c), low frequencies dominate in both real datasets. During the training, the evolution of relative errors of certain selected frequencies (marked by black squares in Fig. 2(a) and 2(c)) is shown in Fig. 2(b) and 2(d). One can easily observe that DNNs capture low frequencies first and gradually capture higher frequencies. Clearly, this behavior is consistent with the F-Principle. For other components of the output vector and other directions of \mathbf{p} , similar phenomena are also observed.

4 F-Principle in MNIST/CIFAR10 through filtering method

The projection method in the previous section enables us to visualize the F-Principle in one direction for each examination at the level of individual frequency components. However, demonstration by this method alone is insufficient because it is impossible to verify the F-Principle at all potentially informative directions for high-dimensional data. To compensate the projection method, in this section, we consider a coarse-grained

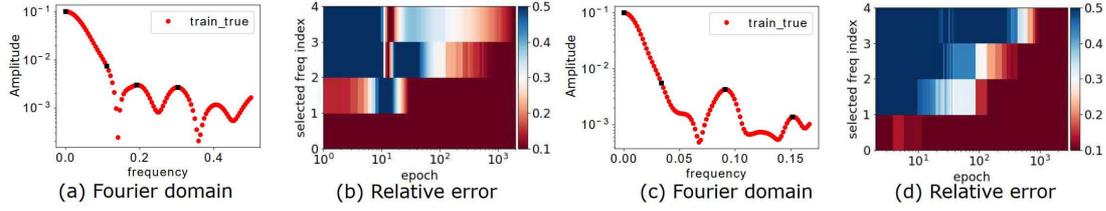


Figure 2: Projection method. (a, b) are for MNIST, (c, d) for CIFAR10. (a, c) Amplitude $|\hat{y}_k|$ vs. frequency. Selected frequencies are marked by black squares. (b, d) $\Delta_F(k)$ vs. training epochs for the selected frequencies. For MNIST dataset, we use a tanh-DNN with widths 784-400-200-10, learning rate 0.001, and batch size 10000. For CIFAR10 dataset, we use a ReLU-CNN, which consists of one convolution layer of $3 \times 3 \times 64$, a max pooling of 2×2 , one convolution layer of $3 \times 3 \times 128$, a max pooling of 2×2 , followed by a fully-connected DNN with widths 800-400-400-400-10. The learning rate is 0.003 with batch size 512. For both cases, the output layer of the network is equipped with a softmax. The DNNs are trained with cross entropy loss by Adam optimizer.

filtering method which is able to unravel whether, in the radially averaged sense, low frequencies converge faster than high frequencies.

4.1 Examination method: Filtering

The idea of the filtering method is as follows. We split the frequency domain into two parts, i.e., a low-frequency part with $|k| \leq k_0$ and a high-frequency part with $|k| > k_0$, where $|\cdot|$ is the length of a vector. The DNN is trained as usual by the original dataset $\{(x_i, y_i)\}_{i=0}^{n-1}$, such as MNIST or CIFAR10. The DNN output is denoted as \hat{h} . During the training, we can examine the convergence of relative errors of low- and high-frequency part, using the two measures below

$$e_{\text{low}} = \left(\frac{\sum_k \mathbb{1}_{|k| \leq k_0} |\hat{y}(k) - \hat{h}(k)|^2}{\sum_k \mathbb{1}_{|k| \leq k_0} |\hat{y}(k)|^2} \right)^{\frac{1}{2}}, \quad (4.1)$$

$$e_{\text{high}} = \left(\frac{\sum_k (1 - \mathbb{1}_{|k| \leq k_0}) |\hat{y}(k) - \hat{h}(k)|^2}{\sum_k (1 - \mathbb{1}_{|k| \leq k_0}) |\hat{y}(k)|^2} \right)^{\frac{1}{2}}, \quad (4.2)$$

respectively, where $\hat{\cdot}$ indicates Fourier transform, $\mathbb{1}_{|k| \leq k_0}$ is an indicator function, i.e.,

$$\mathbb{1}_{|k| \leq k_0} = \begin{cases} 1, & |k| \leq k_0, \\ 0, & |k| > k_0. \end{cases}$$

If we consistently observe $e_{\text{low}} < e_{\text{high}}$ for different k_0 's during the training, then in a mean sense, lower frequencies are first captured by the DNN, i.e., F-Principle.

However, because it is almost impossible to compute above quantities numerically due to high computational cost of high-dimensional Fourier transform, we alternatively use the Fourier transform of a Gaussian function $\hat{G}^\delta(k)$, where δ is the variance of the Gaussian function G , to approximate $\mathbb{1}_{|k| > k_0}$. This is reasonable due to the following two

reasons. First, the Fourier transform of a Gaussian is still a Gaussian, i.e., $\hat{G}^\delta(\mathbf{k})$ decays exponentially as $|\mathbf{k}|$ increases, therefore, it can approximate $\mathbb{1}_{|\mathbf{k}| \leq k_0}$ by $\hat{G}^\delta(\mathbf{k})$ with a proper $\delta(k_0)$ (referred to as δ for simplicity). Second, the computation of e_{low} and e_{high} contains the multiplication of Fourier transforms in the frequency domain, which is equivalent to the Fourier transform of a convolution in the spatial domain. We can equivalently perform the examination in the spatial domain so as to avoid the almost impossible high-dimensional Fourier transform. The low frequency part can be derived by

$$\mathbf{y}_i^{\text{low},\delta} \triangleq (\mathbf{y} * G^\delta)_i, \quad (4.3)$$

where $*$ indicates convolution operator, and the high frequency part can be derived by

$$\mathbf{y}_i^{\text{high},\delta} \triangleq \mathbf{y}_i - \mathbf{y}_i^{\text{low},\delta}. \quad (4.4)$$

Then, we can examine

$$e_{\text{low}} = \left(\frac{\sum_i |\mathbf{y}_i^{\text{low},\delta} - \mathbf{h}_i^{\text{low},\delta}|^2}{\sum_i |\mathbf{y}_i^{\text{low},\delta}|^2} \right)^{\frac{1}{2}}, \quad (4.5)$$

$$e_{\text{high}} = \left(\frac{\sum_i |\mathbf{y}_i^{\text{high},\delta} - \mathbf{h}_i^{\text{high},\delta}|^2}{\sum_i |\mathbf{y}_i^{\text{high},\delta}|^2} \right)^{\frac{1}{2}}, \quad (4.6)$$

where $\mathbf{h}^{\text{low},\delta}$ and $\mathbf{h}^{\text{high},\delta}$ are obtained from the DNN output \mathbf{h} , which evolves as a function of training epoch, through the same decomposition. If $e_{\text{low}} < e_{\text{high}}$ for different δ 's during the training, F-Principle holds; otherwise, it is falsified. Next, we introduce the experimental procedure.

Step One: Training. Train the DNN by the *original dataset* $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=0}^{n-1}$, such as MNIST or CIFAR10. \mathbf{x}_i is an image vector, \mathbf{y}_i is a one-hot vector.

Step Two: Filtering. The low frequency part can be derived by

$$\mathbf{y}_i^{\text{low},\delta} = \frac{1}{C_i} \sum_{j=0}^{n-1} \mathbf{y}_j G^\delta(\mathbf{x}_i - \mathbf{x}_j), \quad (4.7)$$

where $C_i = \sum_{j=0}^{n-1} G^\delta(\mathbf{x}_i - \mathbf{x}_j)$ is a normalization factor and

$$G^\delta(\mathbf{x}_i - \mathbf{x}_j) = \exp(-|\mathbf{x}_i - \mathbf{x}_j|^2 / (2\delta)). \quad (4.8)$$

The high frequency part can be derived by $\mathbf{y}_i^{\text{high},\delta} \triangleq \mathbf{y}_i - \mathbf{y}_i^{\text{low},\delta}$. We also compute $\mathbf{h}_i^{\text{low},\delta}$ and $\mathbf{h}_i^{\text{high},\delta}$ for each DNN output \mathbf{h}_i .

Step Three: Examination. To quantify the convergence of $\mathbf{h}^{\text{low},\delta}$ and $\mathbf{h}^{\text{high},\delta}$, we compute the relative error e_{low} and e_{high} at each training epoch through Eq. (4.6).

4.2 DNNs with various settings

With the filtering method, we show the F-Principle in the DNN training process of real datasets for commonly used large networks. For MNIST, we use a fully-connected tanh-DNN (no softmax) with MSE loss; for CIFAR10, we use cross-entropy loss and two structures, one is small ReLU-CNN network, i.e., two convolutional layers, followed by a fully-connected multi-layer neural network with a softmax; the other is VGG16 [28] equipped with a 1024 fully-connected layer. These three structures are denoted as “DNN”, “CNN” and “VGG” in Fig. 3, respectively. All are trained by SGD from *scratch*.

We scan a large range of δ for both datasets. As an example, results of each dataset for several δ 's are shown in Fig. 3, respectively. Red color indicates small relative error. In all cases, the relative error of the low-frequency part, i.e., e_{low} , decreases (turns red) much faster than that of the high-frequency part, i.e., e_{high} . Therefore, as analyzed above, the low-frequency part converges faster than the high-frequency part. We also remark that, based on the above results on cross-entropy loss, the F-Principle is not limited to MSE loss, which possesses a natural Fourier domain interpretation by the Parseval's theorem. Note that the above results holds for both SGD and GD.

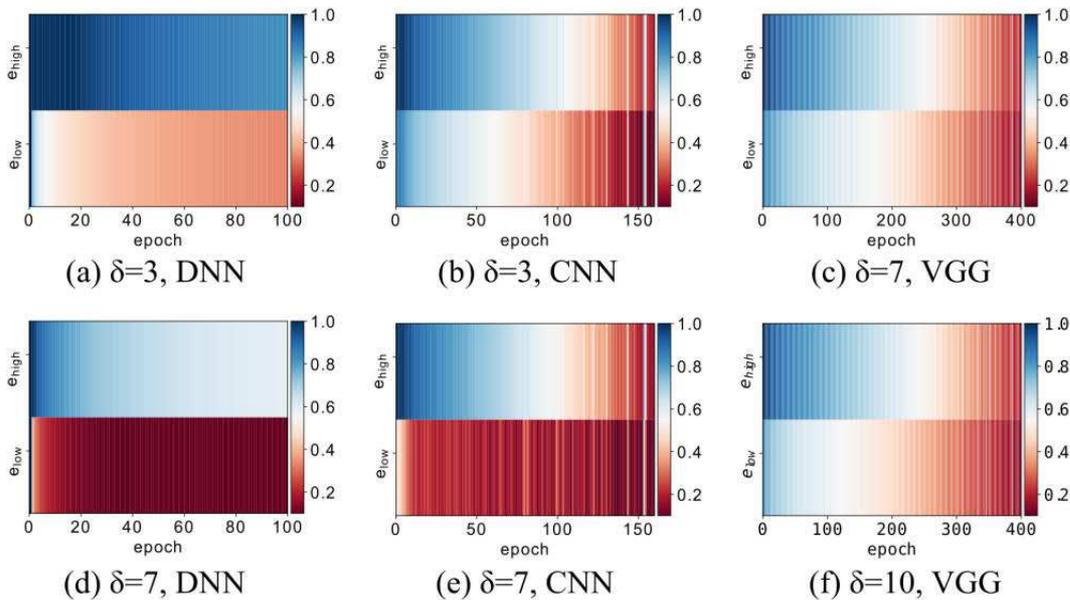


Figure 3: F-Principle in real datasets. e_{low} and e_{high} indicated by color against training epoch. For MNIST, we use a fully-connected tanh-DNN with widths 784-400-200-10 and MSE loss; for CIFAR10, we use cross-entropy loss and a ReLU-CNN, which consists of one convolution layer of $3 \times 3 \times 32$, a max pooling of 2×2 , one convolution layer of $3 \times 3 \times 64$, a max pooling of 2×2 , followed by a fully-connected DNN with widths 400-10 and the output layer of the network is equipped with a softmax. The learning rate for MNIST and CIFAR10 is 0.015 and 0.003, respectively, with batch size 10000. For VGG16, the learning rate is 10^{-5} with batch size 500. All networks are trained by Adam optimizer [20].

5 F-Principle in solving differential equation

Recently, DNN-based approaches have been actively explored for a variety of scientific computing problems, e.g., solving high-dimensional partial differential equations [8, 13, 14, 17, 19, 30, 32] and molecular dynamics (MD) simulations [15]. However, the behaviors of DNNs applied to these problems are not well-understood. To facilitate the designs and applications of DNN-based schemes, it is important to characterize the difference between DNNs and conventional numerical schemes on various scientific computing problems. In this section, focusing on solving Poisson’s equation, which has broad applications in mechanical engineering and theoretical physics [12], we highlight a stark difference between a DNN-based solver and the Jacobi method during the training/iteration, which can be explained by the F-Principle.

In this section, we consider a 1-d Poisson’s equation:

$$-\Delta u(x) = g(x), \quad x \in \Omega \triangleq (-1, 1), \tag{5.1}$$

$$u(-1) = u(1) = 0. \tag{5.2}$$

5.1 Central difference scheme and Jacobi method

Before we show the difference between a DNN-based solver and the Jacobi method, we illustrate the procedure of the Jacobi method.

$[-1, 1]$ is uniformly discretized into $n + 1$ points with grid size $h = 2/n$. The Poisson’s equation in Eq. (5.1) can be solved by the central difference scheme,

$$-\Delta u_i = -\frac{u_{i+1} - 2u_i + u_{i-1}}{(\delta x)^2} = g(x_i), \quad i = 1, 2, \dots, n, \tag{5.3}$$

resulting a linear system

$$A\mathbf{u} = \mathbf{g}, \tag{5.4}$$

where

$$A = \begin{pmatrix} 2 & -1 & 0 & 0 & \dots & 0 \\ -1 & 2 & -1 & 0 & \dots & 0 \\ 0 & -1 & 2 & -1 & \dots & 0 \\ \vdots & \vdots & \dots & & & \vdots \\ 0 & 0 & \dots & 0 & -1 & 2 \end{pmatrix}_{(n-1) \times (n-1)}, \tag{5.5}$$

$$\mathbf{u} = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ u_{n-2} \\ u_{n-1} \end{pmatrix}, \quad \mathbf{g} = (\delta x)^2 \begin{pmatrix} g_1 \\ g_2 \\ \vdots \\ g_{n-2} \\ g_{n-1} \end{pmatrix}, \quad x_i = 2\frac{i}{n}. \tag{5.6}$$

A class of methods to solve this linear system is iterative schemes, for example, the Jacobi method. Let $A = D - L - U$, where D is the diagonal of A , and L and U are the strictly lower and upper triangular parts of $-A$, respectively. Then, we obtain

$$u = D^{-1}(L + U)u + D^{-1}g. \quad (5.7)$$

At step $t \in \mathbb{N}$, the Jacobi iteration reads as

$$u^{t+1} = D^{-1}(L + U)u^t + D^{-1}g. \quad (5.8)$$

We perform the standard error analysis of the above iteration process. Denote u^* as the true value obtained by directly performing inverse of A in Eq. (5.4). The error at step $t+1$ is $e^{t+1} = u^{t+1} - u^*$. Then, $e^{t+1} = R_J e^t$, where $R_J = D^{-1}(L + U)$. The converging speed of e^t is determined by the eigenvalues of R_J , that is,

$$\lambda_k = \lambda_k(R_J) = \cos \frac{k\pi}{n}, \quad k = 1, 2, \dots, n-1, \quad (5.9)$$

and the corresponding eigenvector v_k 's entry is

$$v_{k,i} = \sin \frac{ik\pi}{n}, \quad i = 1, 2, \dots, n-1. \quad (5.10)$$

So we can write

$$e^t = \sum_{k=1}^{n-1} \alpha_k^t v_k, \quad (5.11)$$

where α_k^t can be understood as the magnitude of e^t in the direction of v_k . Then,

$$\begin{aligned} e^{t+1} &= \sum_{k=1}^{n-1} \alpha_k^t R_J v_k = \sum_{k=1}^{n-1} \alpha_k^t \lambda_k v_k, \\ \alpha_k^{t+1} &= \lambda_k \alpha_k^t. \end{aligned} \quad (5.12)$$

Therefore, the converging rate of e^t in the direction of v_k is controlled by λ_k . Since

$$\cos \frac{k\pi}{n} = -\cos \frac{(n-k)\pi}{n}, \quad (5.13)$$

the frequencies k and $(n-k)$ are closely related and converge with the same rate. Consider the frequency $k < n/2$, λ_k is larger for lower frequency. Therefore, lower frequency converges slower in the Jacobi method.

5.2 Numerical experiments

We consider the example with $g(x) = \sin(x) + 4\sin(4x) - 8\sin(8x) + 16\sin(24x)$ which has analytic solution $u_{\text{ref}}(x) = g_0(x) + c_1x + c_0$, where $g_0 = \sin(x) + \sin(4x)/4 - \sin(8x)/8 + \sin(24x)/36$, $c_1 = (g_0(-1) - g_0(1))/2$ and $c_0 = -(g_0(-1) + g_0(1))/2$. 1001 training samples $\{x_i\}_{i=0}^n$ are evenly spaced with grid size δx in $[0, 1]$. Here, we use the DNN output, $h(x; \theta)$, to fit $u_{\text{ref}}(x)$ (Fig. 4(a)). A DNN-based scheme is proposed by considering the following empirical loss function [11],

$$I_{\text{emp}} = \sum_{i=1}^{n-1} \left(\frac{1}{2} |\nabla_x h(x_i)|^2 - g(x_i)h(x_i) \right) \delta x + \beta (h(x_0)^2 + h(x_n)^2). \tag{5.14}$$

The second term in $I_{\text{emp}}(h)$ is a penalty, with constant β , arising from the Dirichlet boundary condition (5.2). After training, the DNN output well matches the analytical solution u_{ref} . Focusing on the convergence of three peaks (inset of Fig. 4(a)) in the Fourier transform of u_{ref} , as shown in Fig. 4(b), low frequencies converge faster than high frequencies as predicted by the F-Principle. For comparison, we also use the Jacobi method to solve problem (5.1). High frequencies converge faster in the Jacobi method, as shown in Fig. 4(c).

As a demonstration, we further propose that DNN can be combined with conventional numerical schemes to accelerate the convergence of low frequencies for computational problems. First, we solve the Poisson’s equation in Eq. (5.1) by DNN with M optimization steps (or epochs), which needs to be chosen carefully, to get a good initial guess in the sense that this solution has already learned the low frequencies (large eigenvalues) part. Then, we use the Jacobi method with the new initial data for the further iterations. We use $\|h - u_{\text{ref}}\|_{\infty} \triangleq \max_{x \in \Omega} |h(x) - u_{\text{ref}}(x)|$ to quantify the learning result. As shown by green stars in Fig. 4(d), $\|h - u_{\text{ref}}\|_{\infty}$ fluctuates after some running time using DNN only. Dashed lines indicate the evolution of the Jacobi method with initial data set

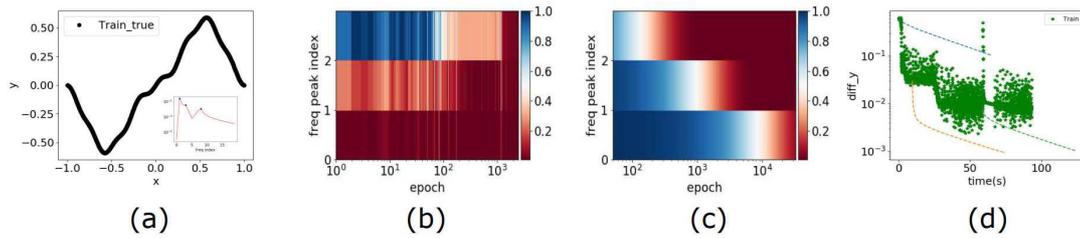


Figure 4: Poisson’s equation. (a) $u_{\text{ref}}(x)$. Inset: $|\hat{u}_{\text{ref}}(k)|$ as a function of frequency. Frequencies peaks are marked with black dots. (b,c) $\Delta_F(k)$ computed on the inputs of training data at different epochs for the selected frequencies for DNN (b) and Jacobi (c). (d) $\|h - u_{\text{ref}}\|_{\infty}$ at different running time. Green stars indicate $\|h - u_{\text{ref}}\|_{\infty}$ using DNN alone. The dashed lines indicate $\|h - u_{\text{ref}}\|_{\infty}$ for the Jacobi method with different colors indicating initialization by different timing of DNN training. We use a DNN with widths 1-4000-500-400-1 and full batch training by Adam optimizer [20]. The learning rate is 0.0005. β is 10. The parameters of the DNN are initialized following a Gaussian distribution with mean 0 and standard deviation 0.02.

to the DNN output at the corresponding steps. If M is too small (stop too early) (left dashed line), *which is equivalent to only using Jacobi*, it would take long time to converge to a small error, because low frequencies converges slowly, yet. If M is too big (stop too late) (right dashed line), *which is equivalent to using DNN only*, much time would be wasted for the slow convergence of high frequencies. A proper choice of M is indicated by the initial point of orange dashed line, in which low frequencies are quickly captured by the DNN, followed by fast convergence in high frequencies of the Jacobi method.

This example illustrates a cautionary tale that, although DNNs has clear advantage, using DNNs alone may not be the best option because of its limitation of slow convergence at high frequencies. Taking advantage of both DNNs and conventional methods to design faster schemes could be a promising direction in scientific computing problems.

6 A preliminary theoretical understanding

The key insight of the F-Principle is that the regularity of DNN converts into the decay rate of a loss function in the frequency domain. For an intuitive understanding of this key insight, we present theories under an idealized setting, which connect the smoothness/regularity of the activation function with different gradient and convergence priorities in frequency domain. Following the idea in this paper, a subsequent theoretical work [23] provides a rigorous mathematical study of the F-Principle at different frequencies for general DNNs (e.g., multiple hidden layers, different activation functions, high-dimensional inputs).

The activation function we consider is $\sigma(x) = \tanh(x)$.

$$\sigma(x) = \tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}, \quad x \in \mathbb{R}.$$

For a DNN of one hidden layer with m nodes, 1-d input x and 1-d output:

$$h(x) = \sum_{j=1}^m a_j \sigma(w_j x + b_j), \quad a_j, w_j, b_j \in \mathbb{R}, \quad (6.1)$$

where w_j , a_j , and b_j are called *parameters*, in particular, w_j and a_j are called *weights*, and b_j is also known as a *bias*. In the sequel, we will also use the notation $\theta = \{\theta_{lj}\}$ with $\theta_{1j} = a_j$, $\theta_{2j} = w_j$, and $\theta_{lj} = b_j$, $j = 1, \dots, m$. Note that $\hat{\sigma}(k) = -\frac{i\pi}{\sinh(\pi k/2)}$ where the Fourier transformation and its inverse transformation are defined as follows:

$$\hat{f}(k) = \int_{-\infty}^{+\infty} f(x) e^{-ikx} dx, \quad f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(k) e^{ikx} dk.$$

The Fourier transform of $\sigma(w_j x + b_j)$ with $w_j, b_j \in \mathbb{R}$, $j = 1, \dots, m$ reads as

$$\widehat{\sigma(w_j \cdot + b_j)}(k) = \frac{2\pi i}{|w_j|} \exp\left(\frac{ib_j k}{w_j}\right) \frac{1}{\exp\left(-\frac{\pi k}{2w_j}\right) - \exp\left(\frac{\pi k}{2w_j}\right)}. \quad (6.2)$$

Thus

$$\hat{h}(k) = \sum_{j=1}^m \frac{2\pi a_j i}{|w_j|} \exp\left(\frac{i b_j k}{w_j}\right) \frac{1}{\exp(-\frac{\pi k}{2w_j}) - \exp(\frac{\pi k}{2w_j})}. \quad (6.3)$$

We define the amplitude deviation between DNN output and the *target function* $f(x)$ at frequency k as

$$D(k) \triangleq \hat{h}(k) - \hat{f}(k).$$

Write $D(k)$ as $D(k) = A(k)e^{i\phi(k)}$, where $A(k) \in [0, +\infty)$ and $\phi(k) \in \mathbb{R}$ are the amplitude and phase of $D(k)$, respectively. The loss at frequency k is $L(k) = \frac{1}{2}|D(k)|^2$, where $|\cdot|$ denotes the norm of a complex number. The total loss function is defined as: $L = \int_{-\infty}^{+\infty} L(k) dk$. Note that according to Parseval's theorem, this loss function in the Fourier domain is equal to the commonly used loss of mean squared error, that is, $L = \int_{-\infty}^{+\infty} \frac{1}{2}(h(x) - f(x))^2 dx$. For readers' reference, we list the partial derivatives of $L(k)$ with respect to parameters

$$\frac{\partial L(k)}{\partial a_j} = \frac{2\pi}{w_j} \sin\left(\frac{b_j k}{w_j} - \phi(k)\right) E_0, \quad (6.4)$$

$$\frac{\partial L(k)}{\partial w_j} = \left[\sin\left(\frac{b_j k}{w_j} - \phi(k)\right) \left(\frac{\pi^2 a_j k}{w_j^3} E_1 - \frac{2\pi a_j}{w_j^2} \right) - \frac{2\pi a_j b_j k}{w_j^3} \cos\left(\frac{b_j k}{w_j} - \phi(k)\right) \right] E_0, \quad (6.5)$$

$$\frac{\partial L(k)}{\partial b_j} = \frac{2\pi a_j b_j k}{w_j^2} \cos\left(\frac{b_j k}{w_j} - \phi(k)\right) E_0, \quad (6.6)$$

where

$$E_0 = \frac{\text{sgn}(w_j) A(k)}{\exp(\frac{\pi k}{2w_j}) - \exp(-\frac{\pi k}{2w_j})},$$

$$E_1 = \frac{\exp(\frac{\pi k}{2w_j}) + \exp(-\frac{\pi k}{2w_j})}{\exp(\frac{\pi k}{2w_j}) - \exp(-\frac{\pi k}{2w_j})}.$$

The descent increment at any direction, say, with respect to parameter θ_{lj} , is

$$\frac{\partial L}{\partial \theta_{lj}} = \int_{-\infty}^{+\infty} \frac{\partial L(k)}{\partial \theta_{lj}} dk. \quad (6.7)$$

The absolute contribution from frequency k to this total amount at θ_{lj} is

$$\left| \frac{\partial L(k)}{\partial \theta_{lj}} \right| \approx A(k) \exp(-|\pi k / 2w_j|) F_{lj}(\theta_j, k), \quad (6.8)$$

where $\theta_j \triangleq \{w_j, b_j, a_j\}$, $\theta_{lj} \in \theta_j$, $F_{lj}(\theta_j, k)$ is a function with respect to θ_j and k , which can be found in one of Eqs. (6.4), (6.5), (6.6).

When the component at frequency k where $\hat{h}(k)$ is not close enough to $\hat{f}(k)$, $\exp(-|\pi k/2w_j|)$ would dominate $G_{lj}(\theta_j, k)$ for a small w_j . Through the above framework of analysis, we have the following theorem. Define

$$W = (w_1, w_2, \dots, w_m)^T \in \mathbb{R}^m. \tag{6.9}$$

Theorem 6.1. Consider a one hidden layer DNN with activation function $\sigma(x) = \tanh x$. For any frequencies k_1 and k_2 such that $|\hat{f}(k_1)| > 0$, $|\hat{f}(k_2)| > 0$, and $|k_2| > |k_1| > 0$, there exist positive constants c and C such that for sufficiently small δ , we have

$$\frac{\mu\left(\left\{W: \left|\frac{\partial L(k_1)}{\partial \theta_{lj}}\right| > \left|\frac{\partial L(k_2)}{\partial \theta_{lj}}\right| \text{ for all } l, j\right\} \cap B_\delta\right)}{\mu(B_\delta)} \geq 1 - C \exp(-c/\delta), \tag{6.10}$$

where $B_\delta \subset \mathbb{R}^m$ is a ball with radius δ centered at the origin and $\mu(\cdot)$ is the Lebesgue measure.

We remark that c and C depend on $k_1, k_2, |\hat{f}(k_1)|, |\hat{f}(k_2)|, \sup |a_i|, \sup |b_i|$, and m .

Proof. To prove the statement, it is sufficient to show that $\mu(S_{lj,\delta})/\mu(B_\delta) \leq C \exp(-c/\delta)$ for each l, j , where

$$S_{lj,\delta} := \left\{W \in B_\delta: \left|\frac{\partial L(k_1)}{\partial \theta_{lj}}\right| \leq \left|\frac{\partial L(k_2)}{\partial \theta_{lj}}\right|\right\}. \tag{6.11}$$

We prove this for $S_{1j,\delta}$, that is, $\theta_{1j} = a_j$. The proofs for $\theta_{lj} = w_j$ and b_j are similar. Without loss of generality, we assume that $k_1, k_2 > 0, b_j > 0$, and $w_j \neq 0, j = 1, \dots, m$. According to Eq. (6.4), the inequality $|\frac{\partial L(k_1)}{\partial a_j}| \leq |\frac{\partial L(k_2)}{\partial a_j}|$ is equivalent to

$$\begin{aligned} & \frac{A(k_2)}{A(k_1)} \left| \frac{\exp(\frac{\pi k_1}{2w_j}) - \exp(-\frac{\pi k_1}{2w_j})}{\exp(\frac{\pi k_2}{2w_j}) - \exp(-\frac{\pi k_2}{2w_j})} \right| \cdot \left| \sin\left(\frac{b_j k_2}{w_j} - \phi(k_2)\right) \right| \\ & \geq \left| \sin\left(\frac{b_j k_1}{w_j} - \phi(k_1)\right) \right|. \end{aligned} \tag{6.12}$$

Note that $|\hat{h}(k)| \leq C \sum_{j=1}^m \frac{|a_j|}{|w_j|} \exp(-\frac{\pi k}{2|w_j|})$ for $k > 0$. Thus

$$\lim_{W \rightarrow 0} \hat{h}(k) = 0 \quad \text{and} \quad \lim_{W \rightarrow 0} D(k) = -\hat{f}(k). \tag{6.13}$$

Therefore,

$$\lim_{W \rightarrow 0} A(k) = |\hat{f}(k)| \quad \text{and} \quad \lim_{W \rightarrow 0} \phi(k) = \pi + \arg(\hat{f}(k)). \tag{6.14}$$

For $W \in B_\delta$ with sufficiently small δ , $A(k_1) > \frac{1}{2}|\hat{f}(k_1)| > 0$ and $A(k_2) < 2|\hat{f}(k_2)|$. Also note that $|\sin(\frac{b_j k_2}{w_j} - \phi(k_2))| \leq 1$ and that for sufficiently small δ ,

$$\left| \frac{\exp(\frac{\pi k_1}{2w_j}) - \exp(-\frac{\pi k_1}{2w_j})}{\exp(\frac{\pi k_2}{2w_j}) - \exp(-\frac{\pi k_2}{2w_j})} \right| \leq 2 \exp\left(\frac{-\pi(k_2 - k_1)}{2|w_j|}\right). \tag{6.15}$$

Thus, inequality (6.12) implies that

$$\left| \sin\left(\frac{b_j k_1}{w_j} - \phi(k_1)\right) \right| \leq \frac{8|\hat{f}(k_2)|}{|\hat{f}(k_1)|} \exp\left(-\frac{\pi(k_2 - k_1)}{2|w_j|}\right). \tag{6.16}$$

Noticing that $\frac{2}{\pi}|x| \leq |\sin x|$ ($|x| \leq \frac{\pi}{2}$) and Eq. (6.14), we have for $W \in S_{l_j, \delta}$, for some $q \in \mathbb{Z}$,

$$\left| \frac{b_j k_1}{w_j} - \arg(\hat{f}(k_1)) - q\pi \right| \leq \frac{8\pi|\hat{f}(k_2)|}{|\hat{f}(k_1)|} \exp\left(-\frac{\pi(k_2 - k_1)}{2\delta}\right), \tag{6.17}$$

that is,

$$\begin{aligned} & -c_1 \exp(-c_2/\delta) + q\pi + \arg(\hat{f}(k_1)) \\ & \leq \frac{b_j k_1}{w_j} \leq c_1 \exp(-c_2/\delta) + q\pi + \arg(\hat{f}(k_1)), \end{aligned} \tag{6.18}$$

where $c_1 = \frac{8\pi|\hat{f}(k_2)|}{|\hat{f}(k_1)|}$ and $c_2 = \pi(k_2 - k_1)$. Define $I := I^+ \cup I^-$ where

$$I^+ := \{w_j > 0: W \in S_{1j, \delta}\}, \quad I^- := \{w_j < 0: W \in S_{1j, \delta}\}. \tag{6.19}$$

For $w_j > 0$, we have for some $q \in \mathbb{Z}$,

$$\begin{aligned} 0 & < \frac{b_j k_1}{c_1 \exp(-c_2/\delta) + q\pi + \arg(\hat{f}(k_1))} \\ & \leq w_j \leq \frac{b_j k_1}{-c_1 \exp(-c_2/\delta) + q\pi + \arg(\hat{f}(k_1))}. \end{aligned} \tag{6.20}$$

Since $W \in B_\delta$ and $c_1 \exp(-c_2/\delta) + \arg(\hat{f}(k_1)) \leq 2\pi$, we have $\frac{b_j k_1}{2\pi + q\pi} \leq w_j \leq \delta$. Then Eq. (6.20) only holds for some large q , more precisely, $q \geq q_0 := \frac{b_j k_1}{\pi\delta} - 2$. Thus we obtain the estimate for the (one-dimensional) Lebesgue measure of I^+

$$\begin{aligned} \mu(I^+) & \leq \sum_{q=q_0}^{\infty} \left| \frac{b_j k_1}{-c_1 \exp(-c_2/\delta) + q\pi + \arg(\hat{f}(k_1))} \right. \\ & \quad \left. - \frac{b_j k_1}{c_1 \exp(-c_2/\delta) + q\pi + \arg(\hat{f}(k_1))} \right| \\ & \leq 2|b_j|k_1 c_1 \exp(-c_2/\delta) \\ & \quad \cdot \sum_{q=q_0}^{\infty} \frac{1}{(q\pi + \arg(\hat{f}(k_1)))^2 - (c_1 \exp(-c_2/\delta))^2} \\ & \leq C \exp(-c/\delta). \end{aligned} \tag{6.21}$$

The similar estimate holds for $\mu(I^-)$, and hence $\mu(I) \leq C \exp(-c/\delta)$. For $W \in B_\delta$, the $(m-1)$ dimensional vector $(w_1, \dots, w_{j-1}, w_{j+1}, \dots, w_m)^T$ is in a ball with radius δ in \mathbb{R}^{m-1} . Therefore, we finally arrive at the desired estimate

$$\frac{\mu(S_{1j,\delta})}{\mu(B_\delta)} \leq \frac{\mu(I)\omega_{m-1}\delta^{m-1}}{\omega_m\delta^m} \leq C \exp(-c/\delta), \tag{6.22}$$

where ω_m is the volume of a unit ball in \mathbb{R}^m . □

Theorem 6.1 indicates that for any two non-converged frequencies, with small weights, the lower-frequency gradient exponentially dominates over the higher-frequency ones. Due to Parseval’s theorem, the MSE loss in the spatial domain is equivalent to the L2 loss in the Fourier domain. To intuitively understand the higher decay rate of a lower-frequency loss function, we consider the training in the Fourier domain with loss function of only two non-zero frequencies.

Theorem 6.2. *Considering a DNN of one hidden layer with activation function $\sigma(x) = \tanh(x)$. Suppose the target function has only two non-zero frequencies k_1 and k_2 , that is, $|\hat{f}(k_1)| > 0$, $|\hat{f}(k_2)| > 0$, $|k_2| > |k_1| > 0$, and $|\hat{f}(k)| = 0$ for $k \neq k_1, k_2$. Consider the loss function of $L = L(k_1) + L(k_2)$ with gradient descent training. Denote*

$$\mathcal{S} = \left\{ \frac{\partial L(k_1)}{\partial t} \leq 0, \frac{\partial L(k_1)}{\partial t} \leq \frac{\partial L(k_2)}{\partial t} \right\},$$

that is, $L(k_1)$ decreases faster than $L(k_2)$. There exist positive constants c and C such that for sufficiently small δ , we have

$$\frac{\mu(\{W : \mathcal{S} \text{ holds}\} \cap B_\delta)}{\mu(B_\delta)} \geq 1 - C \exp(-c/\delta),$$

where $B_\delta \subset \mathbb{R}^m$ is a ball with radius δ centered at the origin and $\mu(\cdot)$ is the Lebesgue measure.

Proof. By gradient descent algorithm, we obtain

$$\begin{aligned} \frac{\partial L(k_1)}{\partial t} &= \sum_{l,j} \frac{\partial L(k_1)}{\partial \theta_{lj}} \frac{\partial \theta_{lj}}{\partial t} \\ &= - \sum_{l,j} \frac{\partial L(k_1)}{\partial \theta_{lj}} \frac{\partial (L(k_1) + L(k_2))}{\partial \theta_{lj}} \\ &= - \sum_{l,j} \left(\frac{\partial L(k_1)}{\partial \theta_{lj}} \right)^2 - \sum_{l,j} \frac{\partial L(k_1)}{\partial \theta_{lj}} \frac{\partial L(k_2)}{\partial \theta_{lj}}, \\ \frac{\partial L(k_2)}{\partial t} &= - \sum_{l,j} \left(\frac{\partial L(k_2)}{\partial \theta_{lj}} \right)^2 - \sum_{l,j} \frac{\partial L(k_1)}{\partial \theta_{lj}} \frac{\partial L(k_2)}{\partial \theta_{lj}}, \end{aligned}$$

and

$$\frac{\partial L}{\partial t} = \frac{\partial(L(k_1) + L(k_2))}{\partial t} = - \sum_{l,j} \left(\frac{\partial L(k_1)}{\partial \theta_{lj}} + \frac{\partial L(k_2)}{\partial \theta_{lj}} \right)^2 \leq 0. \quad (6.23)$$

To obtain

$$0 > \frac{\partial L(k_1)}{\partial t} - \frac{\partial L(k_2)}{\partial t} = - \sum_{l,j} \left[\left(\frac{\partial L(k_1)}{\partial \theta_{lj}} \right)^2 - \left(\frac{\partial L(k_2)}{\partial \theta_{lj}} \right)^2 \right], \quad (6.24)$$

it is sufficient to have

$$\left| \frac{\partial L(k_1)}{\partial \theta_{lj}} \right| > \left| \frac{\partial L(k_2)}{\partial \theta_{lj}} \right|. \quad (6.25)$$

Eqs. (6.23), (6.24) also yield to

$$\frac{\partial L(k_1)}{\partial t} < 0.$$

Therefore, Eq. (6.25) is a sufficient condition for \mathcal{S} . Based on Theorem 6.1, we have proved Theorem 6.2. \square

7 Discussions

7.1 Generalization

DNNs often generalize well for real problems [37] but poorly for problems like fitting a parity function [25, 27] despite excellent training accuracy for all problems. Understanding the differences between above two types of problems, i.e., good and bad generalization performance of DNN, is critical. In the following, we show a qualitative difference between these two types of problems through *Fourier analysis* and use the *F-Principle* to provide an explanation different generalization performances of DNNs.

For MNIST/CIFAR10, we examine

$$\hat{y}_{\text{total},k} = \frac{1}{n_{\text{total}}} \sum_{i=0}^{n_{\text{total}}-1} y_i \exp(-i2\pi k \cdot x_i), \quad (7.1)$$

where $\{(x_i, y_i)\}_{i=0}^{n_{\text{total}}-1}$ consists of both the training and test datasets with certain selected output component, at different directions of k in the Fourier space. We find that $\hat{y}_{\text{total},k}$ concentrates on the low frequencies along those examined directions. For illustration, $\hat{y}_{\text{total},k}$'s along the first principle component are shown by green lines in Fig. 5(a, b) for MNIST/CIFAR10, respectively. When only the training dataset is used, $\hat{y}_{\text{train},k}$ well overlaps with $\hat{y}_{\text{total},k}$ at the dominant low frequencies.

For the parity function $f(x) = \prod_{j=1}^d x_j$ defined on $\Omega = \{-1, 1\}^d$, its Fourier transform is $\hat{f}(k) = \frac{1}{2^d} \sum_{x \in \Omega} \prod_{j=1}^d x_j e^{-i2\pi k \cdot x} = (-i)^d \prod_{j=1}^d \sin 2\pi k_j$. Clearly, for $k \in [-\frac{1}{4}, \frac{1}{4}]^d$, the power of the parity function concentrates at $k \in \{-\frac{1}{4}, \frac{1}{4}\}^d$ and vanishes as $k \rightarrow \mathbf{0}$, as illustrated in Fig. 5(c) for the direction of $\mathbf{1}_d$. Given a randomly sampled training

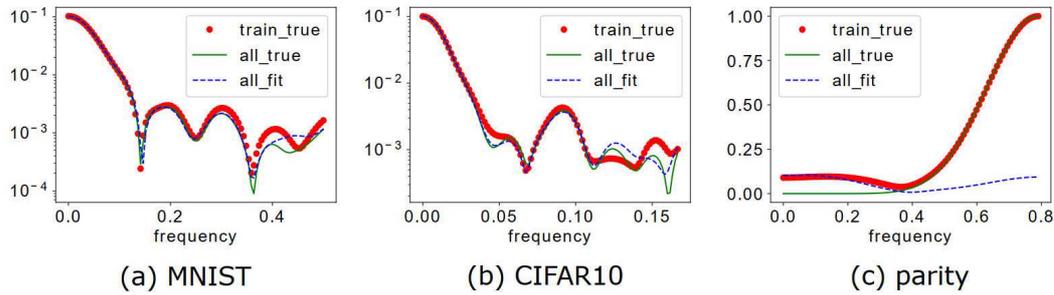


Figure 5: Fourier analysis for different generalization ability. The plot is the amplitude of the Fourier coefficient against frequency k . The red dots are for the training dataset, the green line is for the whole dataset, and the blue dashed line is for an output of well-trained DNN on the input of the whole dataset. For (c), $d = 10$. The training data is 200 randomly selected points. The settings of (a) and (b) are the same as the ones in Fig. 2. For (c), we use a tanh-DNN with widths 10-500-100-1, learning rate 0.0005 under full batch-size training by Adam optimizer. The parameters of the DNN are initialized by a Gaussian distribution with mean 0 and standard deviation 0.05.

dataset $S \subset \Omega$ with s points, the nonuniform Fourier transform on S is computed as $\hat{f}_S(\mathbf{k}) = \frac{1}{s} \sum_{x \in S} \prod_{j=1}^d x_j e^{-i2\pi \mathbf{k} \cdot x}$. As shown in Fig. 5(c), $\hat{f}(\mathbf{k})$ and $\hat{f}_S(\mathbf{k})$ significantly differ at low frequencies.

By experiments, the generalization ability of DNNs can be well reflected by the Fourier analysis. For the MNIST/CIFAR10, we observed the Fourier transform of the output of a well-trained DNN on $\{\mathbf{x}_i\}_{i=0}^{n_{\text{total}}-1}$ faithfully recovers the dominant low frequencies, as illustrated in Fig. 5(a) and 5(b), respectively, indicating a good generalization performance as observed in experiments. However, for the parity function, we observed that the Fourier transform of the output of a well-trained DNN on $\{\mathbf{x}_i\}_{i \in S}$ significantly deviates from $\hat{f}(\mathbf{k})$ at almost all frequencies, as illustrated in Fig. 5(c), indicating a bad generalization performance as observed in experiments.

The F-Principle implicates that among all the functions that can fit the training data, a DNN is implicitly biased during the training towards a function with more power at low frequencies. If the target function has significant high-frequency components, insufficient training samples will lead to artificial low frequencies in training dataset (see red line in Fig. 5(c)), which is the well-known *aliasing* effect. Based on the F-Principle, as demonstrated in Fig. 5(c), these artificial low frequency components will be first captured to explain the training samples, whereas the high frequency components will be compromised by DNN. For MNIST/CIFAR10, since the power of high frequencies is much smaller than that of low frequencies, artificial low frequencies caused by aliasing can be neglected. To conclude, the distribution of power in Fourier domain of above two types of problems exhibits significant differences, which result in different generalization performances of DNNs according to the F-Principle.

7.2 Related work

There are different approaches attempting to explain why DNNs often generalize well. For example, generalization error is related to various complexity measures [2, 10, 24], local properties (sharpness/flatness) of loss functions at minima [18, 33], stability of optimization algorithms [16], and implicit bias of the training process [1, 29, 34]. On the other hand, several works focus on the failure of DNNs [25, 27], e.g., fitting the parity function, in which a well-trained DNN possesses no generalization ability. We propose that the Fourier analysis can provide insights into both success and failure of DNNs.

F-Principle was verified in [26, 34] through simple synthetic data and not very deep networks. In the revised version, [26] examines the F-Principle in the MNIST dataset. However, they add noise to MNIST, which contaminates the labels and damages the structure of real data. They only examine not very deep (6-layer) fully connected ReLU network with MSE loss, while cross-entropy loss is widely used. This paper verified that F-Principle holds in the training process of MNIST and CIFAR10, both CNN and fully connected networks, very deep networks (VGG16) and various loss functions, e.g., MSE Loss, cross-entropy loss and variational loss function. In the aspect of theoretical study, based on the key mechanism found by the theoretical study in this paper, [23] shows a rigorous proof of the F-Principle for general DNNs. The theoretical study of the gradient of $\tanh(x)$ in the Fourier domain is adopted by [26], in which they generalize the analysis to ReLU and show similar results. A subsequent work [9] demonstrates the F-Principle through a continuous viewpoint, that is, the dynamics of the neural network is governed by an integral equation and illustrate the F-Principle by a 1d pseudo-spectrum method.

In the aspect of application, based on the study of the F-Principle in this paper, several algorithms are developed. For example, Blind et al. [3] explicitly impose higher priority for high frequencies in the loss function to significantly accelerate the simulation of fluid dynamics through DNN approach; Cai et al. [6] design DNN-based algorithms to solve high-frequency problems by shifting high frequencies to lower ones.

Note that this paper includes two arxiv papers [35, 36].

Acknowledgments

We thank Prof. Weinan E (Princeton) and Prof. David W. McLaughlin (NYU) for helpful discussions. This work is sponsored by National Key R&D Program of China (2019YFA0709503) (Z. X.), Shanghai Sailing Program (Z. X.), Natural Science Foundation of Shanghai (20ZR1429000) (Z. X.), NSF Grant No.DMS-1638352 (Y. Z.) and the Ky Fan and Yu-Fen Fan Membership Fund (Y. Z.).

References

- [1] Devansh Arpit, Stanislaw Jastrzbski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al.

- A closer look at memorization in deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 233–242, 2017.
- [2] Peter L Bartlett, Vitaly Maiorov, and Ron Meir. Almost linear vc dimension bounds for piecewise polynomial networks. In *Advances in Neural Information Processing Systems*, pages 190–196, 1999.
 - [3] Simon Biland, Vinicius C Azevedo, Byungsoo Kim, and Barbara Solenthaler. Frequency-aware reconstruction of fluid simulations with generative networks. *arXiv preprint arXiv:1912.08776*, 2019.
 - [4] Ronald Newbold Bracewell and Ronald N Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.
 - [5] James H Bramble, Joseph E Pasciak, Jun Ping Wang, and Jinchao Xu. Convergence estimates for multigrid algorithms without regularity assumptions. *Mathematics of Computation*, 57(195):23–45, 1991.
 - [6] Wei Cai, Xiaoguang Li, and Lizuo Liu. A phase shift deep neural network for high frequency wave equations in inhomogeneous media. *Arxiv preprint, arXiv:1909.11759*, 2019.
 - [7] Jingrun Chen and Carlos J García-Cervera. An efficient multigrid strategy for large-scale molecular mechanics optimization. *Journal of Computational Physics*, 342:29–42, 2017.
 - [8] Weinan E, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
 - [9] Weinan E, Chao Ma, and Lei Wu. Machine learning from a continuous viewpoint. *arXiv preprint arXiv:1912.12777*, 2019.
 - [10] Weinan E, Chao Ma, and Lei Wu. A priori estimates of the population risk for two-layer neural networks. *Communications in Mathematical Sciences*, 17:1407–1425, 2019.
 - [11] Weinan E and Bing Yu. The deep ritz method: A deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
 - [12] Lawrence C Evans. *Partial differential equations*. 2010.
 - [13] Yuwei Fan, Lin Lin, Lexing Ying, and Leonardo Zepeda-Núñez. A multiscale neural network based on hierarchical matrices. *Multiscale Modeling & Simulation*, 17(4):1189–1213, 2019.
 - [14] Jiequn Han, Chao Ma, Zheng Ma, and E Weinan. Uniformly accurate machine learning-based hydrodynamic models for kinetic equations. *Proceedings of the National Academy of Sciences*, 116(44):21983–21991, 2019.
 - [15] Jiequn Han, Linfeng Zhang, Roberto Car, et al. Deep potential: A general representation of a many-body potential energy surface. *Communications in Computational Physics*, 23(3), 2018.
 - [16] Moritz Hardt, Benjamin Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. *arXiv preprint arXiv:1509.01240*, 2015.
 - [17] Juncai He, Lin Li, Jinchao Xu, and Chunyue Zheng. Relu deep neural networks and linear finite elements. *Journal of Computational Mathematics*, 2019.
 - [18] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
 - [19] Yuehaw Khoo and Lexing Ying. Switchnet: a neural network model for forward and inverse scattering problems. *SIAM Journal on Scientific Computing*, 41(5):A3182–A3201, 2019.
 - [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
 - [21] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for ad-

- vanced research). URL <http://www.cs.toronto.edu/kriz/cifar.html>, 2010.
- [22] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
 - [23] Tao Luo, Zheng Ma, Zhi-Qin John Xu, and Yaoyu Zhang. Theory of the frequency principle for general deep neural networks. *arXiv preprint arXiv:1906.09235*, 2019.
 - [24] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 5947–5956, 2017.
 - [25] Maxwell Nye and Andrew Saxe. Are efficient deep representations learnable? 2018.
 - [26] Nasim Rahaman, Devansh Arpit, Aristide Baratin, Felix Draxler, Min Lin, Fred A Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of deep neural networks. *International Conference on Machine Learning*, 2019.
 - [27] Shai Shalev-Shwartz, Ohad Shamir, and Shaked Shammah. Failures of gradient-based deep learning. *arXiv preprint arXiv:1703.07950*, 2017.
 - [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
 - [29] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70), 2018.
 - [30] Carlos Michelen Strofer, Jin-Long Wu, Heng Xiao, and Eric Paterson. Data-driven, physics-based feature extraction from fluid flow fields using convolutional neural networks. *Communications in Computational Physics*, 25(3):625–650, 2019.
 - [31] Ulrich Trottenberg, Cornelius W Oosterlee, and Anton Schuller. *Multigrid*. Elsevier, 2000.
 - [32] E Weinan, Chao Ma, and Jianchun Wang. Model reduction with memory and the machine learning of dynamical systems. *Communications in Computational Physics*, 25(4):947–962, 2018.
 - [33] Lei Wu, Zhanxing Zhu, and Weinan E. Towards understanding generalization of deep learning: Perspective of loss landscapes. *arXiv preprint arXiv:1706.10239*, 2017.
 - [34] Zhi-Qin J Xu, Yaoyu Zhang, and Yanyang Xiao. Training behavior of deep neural network in frequency domain. *International Conference on Neural Information Processing*, pages 264–274, 2019.
 - [35] Zhi-Qin John Xu. Frequency principle in deep learning with general loss functions and its potential application. *arXiv preprint arXiv:1811.10146*, 2018.
 - [36] Zhiqin John Xu. Understanding training and generalization in deep learning by fourier analysis. *arXiv preprint arXiv:1808.04295*, 2018.
 - [37] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *The International Conference on Learning Representations*, 2017.