

FINER: Flexible spectral-bias tuning in Implicit NEural Representation by Variable-periodic Activation Functions*

Zhen Liu^{1,†}, Hao Zhu^{1,†}, Qi Zhang², Jingde Fu³, Weibing Deng³, Zhan Ma¹, Yanwen Guo⁴, Xun Cao¹

¹ School of Electronic Science and Engineering, ³ Department of Mathematics,

⁴ Department of Computer Science and Technology, Nanjing University, Nanjing 210023, China

² AI Lab, Tencent Company, Shenzhen 518054, China

[†] Equal contribution. Corresponding author: caoxun@nju.edu.cn

Abstract

Implicit Neural Representation (INR), which utilizes a neural network to map coordinate inputs to corresponding attributes, is causing a revolution in the field of signal processing. However, current INR techniques suffer from a restricted capability to tune their supported frequency set, resulting in imperfect performance when representing complex signals with multiple frequencies. We have identified that this frequency-related problem can be greatly alleviated by introducing variable-periodic activation functions, for which we propose FINER. By initializing the bias of the neural network within different ranges, sub-functions with various frequencies in the variable-periodic function are selected for activation. Consequently, the supported frequency set of FINER can be flexibly tuned, leading to improved performance in signal representation. We demonstrate the capabilities of FINER in the contexts of 2D image fitting, 3D signed distance field representation, and 5D neural radiance fields optimization, and we show that it outperforms existing INRs.

1. Introduction

The way a signal is represented is the foundation of all the following problems and determines the paradigm for solving them. Traditional representations, such as the image matrices, point cloud or volumes [35], focus on recording the elements individually and have offered significant contributions in history. However, this representation is increasingly inadequate for addressing the numerous inverse problems arising in modern times, such as neural rendering [35], inverse imaging [39] and simulations [14]. On the contrary, the implicit neural representation (INR) [32], which characterizes a signal by preserving

*The work was supported in part by the National Natural Science Foundation of China (Grants No. 62025108, 12171237).

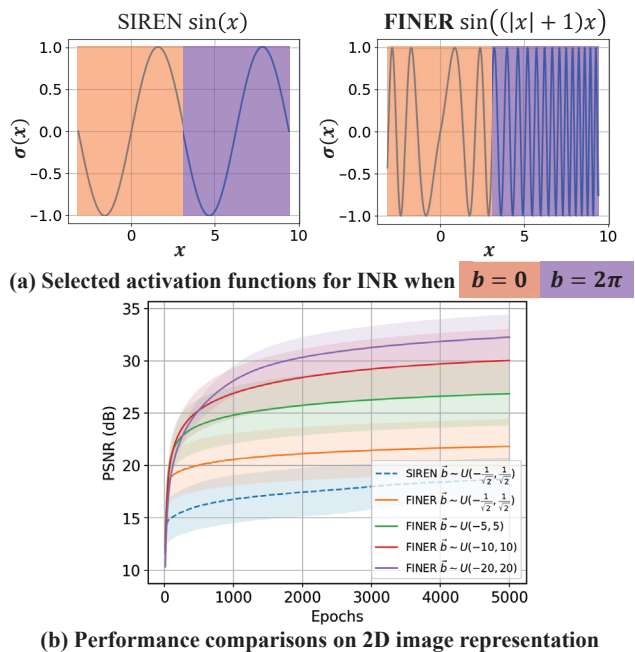


Figure 1. **Flexible spectral-bias tuning in implicit neural representation (FINER)**. We observe that the supported frequency set in classical INRs is limited due to the underutilization of activation functions’ definition domain, *i.e.*, they mainly employ the central region near the origin point. To overcome this limitation, we propose a novel variable-periodic activation function for INR. This innovation allows for tuning the supported frequency set by adjusting the initialization range of the bias vector in the neural network. (a) visualizes the selected narrow activation functions in classical periodic activation $\sin(x)$ alongside our proposed variable-periodic one $\sin((|x| + 1)x)$ with different bias settings. (b) plots the training curves of SIREN and FINER, demonstrating the impact of different initializations applied to the bias vector b (see Sec. 5.1.1 for more details).

the mapping function from the coordinates to corresponding attributes using neural networks, is gaining increasing attention thanks to the advantages of querying attributes continu-

ously and incorporating the differentiable physical process seamlessly. As a result, INR has found widespread application in solving domain-specific inverse problems [8, 14], particularly in cases where large-scale paired datasets are unavailable, and only measurements and forward physical process are provided.

However, existing INR techniques suffer the well-known spectral-bias [26, 40], *i.e.*, the low-frequency components of the signal are more easily learned. To address this bias, the positional encoding [23, 34] which aims at embedding multiple orthogonal Fourier or Wavelet bases [6] into the subsequent network is proposed. However, a significant challenge arises from the fact that the frequency distribution of a signal to be inversely solved is often unknown. This can potentially lead to a mismatch between the pre-defined bases' frequency set and the characteristics of the signal itself, resulting in an imperfect representation [38]. Apart from the pre-defined frequency set, there is a growing focus within the research community on automatic frequency tuning, achieved through the use of periodic [32] or non-periodic activation [28, 30] functions. Nevertheless, the supported frequency set is still limited and the representational accuracy could be further improved.

Such a problem is closely related to the underutilization of the used activation functions. While activation functions have an infinite domain, practical applications typically make use of only the regions centered around the origin. By 'interspersing' narrow activation functions with different frequencies along the full definition domain and then selecting the ideal one by controlling the range of input values, the supported frequency set could be significantly expanded, resulting in improved expressive power of current INRs. Following this idea, we propose the FINER, where variable-periodic functions are applied to activate the neurons within the INR.

Different from previous explorations [28, 30, 32] which focus on optimizing the weight matrix for manipulating frequency candidates with better matching degree, FINER opens a novel way to achieve frequency tuning by modulating the bias vector, or in other words, the phase of the variable-periodic activation functions. We demonstrate that both the shift-invariance and eigenvalues distribution of FINER's neural tangent kernel (NTK) can be enhanced (see Figs 1, 4) by increasing the standard deviation during bias initialization, thus the spectral bias could be flexibly tuned and the expressive power are significantly improved. To verify the performance, extensive experiments are conducted on 2D image fitting, 3D signed distance field representation and 5D neural radiance field optimization. Specifically, the main contributions of the work include,

1. We introduce a novel implicit neural representation with flexible spectral-bias tuning for representing and optimizing signals.

2. We propose a novel initialization scheme and prove its effectiveness and efficiency from both the perspectives of geometry and neural tangent kernel.
3. We substantiate that FINER surpasses prior INRs activated with other functions for 2D image fitting, 3D signed distance field representation and 5D neural radiance field optimization.

2. Related Work

2.1. Implicit Neural Representation

INRs [32, 34] serve as the foundational building blocks for neural scene representations. These representations are designed to learn continuous functions using a multi-layer perception (MLP) that maps coordinates to visual signals, such as images [5, 19, 38], videos [3, 15], and 3D scenes [23, 36]. Neural Radiance Fields or NeRF [23], a notable breakthrough in this domain, learns a 5D INR to reconstruct the radiance fields (density and view-dependent color) of a scene. With the widespread application of NeRF and its variants [2, 24] on realistic view synthesis, INRs have rapidly expanded into various fields of vision and signal processing, such as cross-model media representation/compression [7, 33], neural camera representations [10, 11], microscopy imaging [20, 41] and partial differential equations solver [14, 27]. Despite the interest and success of these implicit representations, current approaches often suffer from the well-known spectral-bias problem. As a result, the INRs may struggle to capture high-fidelity details of complex signals, leading to suboptimal performance in fitting functions and ineffectiveness in various applications.

2.2. Spectral-bias Problem

The issue of spectral-bias in deep learning [26] refers to the innate propensity of these models to selectively learn specific patterns or features from input data. In the case of INR-based methods, this problem manifests as a preference for learning low-frequency components of a signal more readily than high-frequency components. To address the spectral-bias problem, several innovative strategies have been proposed for INR-based methods. In particular, spatial encoding is applied to the input data, such as frequency or polynomial decomposition [25, 31, 34], high-pass filtering [6, 37], to emphasize high-frequency components before feeding into the model. [34] use the Fourier features mapped from spatial coordinates as the input of MLPs to improve the performance of INRs in adequately expressing high-frequency information of signals. Additionally, various architectural modifications have been integrated into INRs, including multi-scale or pyramid representations [19, 29, 42], which can aid in capturing both low-frequency and high-frequency components of a scene. [19]

implement a multi-scale network architecture with a band-limited Fourier spectrum to minimize artifacts during the downsampling or upsampling process. However, the frequency distribution of a signal requiring inverse solving is often unknown, making it difficult to design an appropriate representation or model for the signal without prior knowledge of its frequency content.

In addition to the positional encoding with pre-defined frequency bases, there has been a growing interest in the research community for automatic frequency tuning through the use of nonlinear activation functions [4, 28, 30, 32]. [32] propose the Sinusoidal Representation Network (SIREN), a method designed to represent complex signals and functions using periodic activation functions, especially sine functions. SIREN has demonstrated its effectiveness in representing intricate details and high-frequency components when compared to traditional activation functions like ReLU or sigmoid. However, it is important to note that SIREN may necessitate careful initialization and hyperparameter tuning to achieve optimal performance. The periodic nature of sine functions can lead to oscillations and slower convergence, posing challenges in the training process. Furthermore, the size of the embedding space is limited, as the frequency distribution bias remains constant for different inputs, making it difficult to represent the diversity of frequency distributions. As a result, it is crucial to develop an adaptive periodic function for activating nonlinearity and complex frequency distribution.

3. FINER: Flexible spectral-bias tuning in Implicit Neural Representation

This section first provides an overview of SIREN, where the capacity-convergence gap is summarized. Then the FINER is proposed in detail.

3.1. SIREN and the Capacity-Convergence Gap

Pipeline of SIREN. Given a signal sequence $\{\vec{x}_i, \vec{y}_i\}_{i=1}^N$, where \vec{x}_i and \vec{y}_i respectively represent the coordinate and the corresponding attributes of the i -th element, N is the number of elements in the signal, SIREN focuses on pursuing a neural network $f(\vec{x}; \theta)$ to characterize them as accurate as possible and could be formulated as follows:

$$\begin{aligned} \vec{z}^0 &= \vec{x} \\ \vec{z}^l &= \sin(\omega_0(W^l \vec{z}^{l-1} + \vec{b}^l)), \quad l = 1, 2, \dots, L-1, \\ f(\vec{x}; \theta) &= W^L \vec{z}^{L-1} + \vec{b}^L \end{aligned} \quad (1)$$

where \vec{z}^l denotes the output of layer l , $\theta = \{W^l, \vec{b}^l \mid l = 1, 2, \dots, L-1\}$ refers to the network parameters to be optimized, L is the number of layers, ω_0 is an empirical parameter for controlling the frequency.




Figure 2. Comparisons of SIREN for fitting a 2K image. There is a performance gap between the results of “Standard initialization multiply ω_0 ” and “Initialization with a larger range”.

Capacity-Convergence Gap. According to [40], the capacity of the SIREN is limited by the choice of ω_0 , *i.e.*, the functions that could be represented by $f(\vec{x}; \theta)$ should be able to be decomposed by a linear combination of certain harmonics of ω_0 , *i.e.*,

$$f(\vec{x}; \theta) \in \left\{ \sum_{\omega' \in \mathcal{F}_{\omega_0}} c_{\omega'} \sin(\omega' \vec{x} + \phi_{\omega'}) \mid c_{\omega'}, \phi_{\omega'} \in Q \right\} \quad (2)$$

where Q is the set of rational number, \mathcal{F}_{ω_0} is the supported frequency set defined by the ω_0 . Because ω_0 only plays the role of scaling neurons, it could be removed by changing the initialization range of network parameters $\{W^l, \vec{b}^l \mid l = 1, \dots, L-1\}$. As a result, it is unnecessary to introduce the parameter ω_0 in theory.

On the other hand, since the sin function is non-convex, the network parameters should be initialized carefully to guarantee the inputs of sin mainly fall into the range of $[-\pi, \pi]$ [32], otherwise SIREN will not converge well and does not have a high accuracy for representing signals. For example, W is often initialized from $\mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$, n is the number of inputs for each neuron [32]. As a result, network parameters $\{W^l, \vec{b}^l \mid l = 1, \dots, L-1\}$ could not be scaled unboundedly considering the convergence. In summary, there is a capacity-convergence gap in SIREN that,

Proposition 1 *The function set that SIREN could be represented can be enlarged by increasing the initialization range of network parameters, which violates the rule of guaranteeing the convergence, resulting in a performance gap between theory and practice.*

A simple example: fitting a high-resolution image. Fig. 2 compares the performance of SIREN with different initialization strategies. The blue line are the results follows the standard form of SIREN, *i.e.*, $W \sim \mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$ and an additional angular frequency ω_0 is applied to multiply the $W\vec{x} + \vec{b}$.

On the contrary, the green line refers to the performance by removing the angular frequency ω_0 in Eqn. 1 and initializing $W \sim \mathcal{U}(-\omega_0\sqrt{6/n}, \omega_0\sqrt{6/n})$. Although they are equivalent in the mathematical form, there is a performance gap in practice.

3.2. Variable-periodic Activation Functions

To overcome the capacity-convergence gap, SIREN introduces the parameter ω_0 to enlarge the supported frequency set. However, ω_0 is often set manually which poses challenges for novices and requires additional tuning process for inverse problems where the frequency distribution of the signal to be solved is unknown. To address this problem, we propose the FINER which introduces a variable-periodic activation function $\sin((|x| + 1)x)$, and the Eqn. 1 could be rewritten as,

$$\begin{aligned} \bar{y}^l &= W^l \bar{z}^{l-1} + \bar{b}^l \\ \bar{z}^l &= \sin(\omega_0(|\bar{y}^l| + 1)\bar{y}^l), \quad l = 1, 2, \dots, L-1 \end{aligned} \quad (3)$$

More utilization of activation function closes the capacity-convergence gap. We notice that the periodic activation function $\sin(x)$ has an infinite definition domain while only the regions centered around the origin point are used. This underutilization of definition domain is caused by the periodic behavior of \sin . As shown in Fig. 1(a), because $\sin(x)$ has a period 2π , the behavior after activation will not be changed when shifting the used regions to other areas, such as orange and purple boxes in Fig. 1(a) where $\bar{b} = 0$ and 2π , respectively. As a result, it is unnecessary to increase the utilization rate of the definition domain in SIREN where only the central region $[-2\pi, 2\pi]$ is used [32]. Different from the periodic function $\sin(x)$, FINER adopts a variable-periodic function $\sin((|x| + 1)x)$. Because the scale parameter $|x| + 1$ increases with the increase of input variable, the frequency of the used activation function will be changed when the used regions are shifted, *i.e.*, different initialization of network parameters \bar{b} . As shown in Fig. 1(a), $\sin((|x| + 1)x)$ for $\bar{b} = 2\pi$ has higher frequencies than itself when $\bar{b} = 0$. Thus the capacity-convergence gap could be closed by the proposed variable-periodic activation function.

4. Flexible Spectral-bias Initialization

In this section, we will first give the initialization scheme of FINER, then the behaviors of FINER's supported frequency set and training dynamics under different initialization are analysed from geometrical and neural tangent kernel perspectives, respectively.

4.1. Initialization scheme

As demonstrated in the above section, the supported frequency set of FINER could be manipulated by changing the bias \bar{b} . However, due to the non-convex nature of the variable-periodic activation function, there are many local minimums in $\sin((|x| + 1)x)$ and gradient-based optimizations (*e.g.*, SGD or Adam) could not guarantee moving \bar{b} to




Figure 3. Comparisons of used activation function $\sin((|x| + 1)x)$ under different bias \bar{b} . More sub-functions with high-frequency are included when b is set with a larger value.

global optimum if \bar{b} is not initialized well. Traditional initialization strategies, which apply uniform or Gaussian sampling in a narrow region centered at 0, makes the supported frequency set of FINER be limited by the frequency of the first cycle in variable-periodic function (*i.e.*, the Fig. 3(a)), resulting in a waste of other cycles (Fig. 3(b)) in variable-periodic function which have difference frequencies.

To get rid of the limited supported frequency set using traditional initialization methods, we derive a novel initialization scheme for \bar{b} for tuning the supported frequency set flexibly, meanwhile the initialization for W follows [32]. We propose to initialize \bar{b} following the uniform distribution $\mathcal{U}(-k, k)$ with a larger range k than the default one in traditional methods,

$$\bar{b} \sim \mathcal{U}(-k, k), \quad k > 0. \quad (4)$$

4.2. Geometrical Perspective

Supposing the supported frequency set of SIREN and FINER are \mathcal{F}_{ω_0} and $\mathcal{F}_{\omega_0, k}$, respectively. To analyse their relationship, let us start from the simplest case.

k is close to the origin point. Because the initialization for W follows [32], the term $W\bar{x}$ has similar distribution with the one in SIREN, that $|W\bar{x}| \leq \pi$. As a result, the term $|(|W\bar{x} + \bar{b}| + 1)(W\bar{x} + \bar{b})|$ drops into the area of $[-\pi^2 - \pi, \pi^2 + \pi]$. As shown in Fig. 3(a), the activation function $\sin((|W\bar{x} + \bar{b}| + 1)(W\bar{x} + \bar{b}))$ mainly spans two narrow sub-functions with different frequencies. For the points dropped into the first sub-function (*i.e.*, $|W\bar{x}| \leq \frac{\sqrt{4\pi+1}-1}{2}$), the orange areas in Fig. 3(a), the supported frequency set $\mathcal{F}_{\omega_0, k}^1$ is slightly larger than \mathcal{F}_{ω_0} , and could be approximated as

$$\mathcal{F}_{\omega_0, k}^1 \approx \left\{ \frac{2\pi}{\sqrt{4\pi+1}-1} \omega \mid \omega \in \mathcal{F}_{\omega_0} \right\}, \quad (5)$$

where the term $\frac{2\pi}{\sqrt{4\pi+1}-1}$ measures the frequency changes from the standard \sin function to the first function in $\sin((|x| + 1)x)$. Considering the fact that the range $[-\pi, \pi]$ of $W\bar{x}$ is continuous, *i.e.*, every value could be produced, the \mathcal{F}_{ω_0} is also a continuous set, as a result,

$$\mathcal{F}_{\omega_0} \subset \mathcal{F}_{\omega_0, k}^1. \quad (6)$$

For the points dropped into the second sub-function (i.e., $\frac{\sqrt{4\pi+1}-1}{2} \leq |W\vec{x}| \leq \frac{\sqrt{12\pi+1}-1}{2}$, the purple areas in Fig. 3(a)), the supported frequency set $\mathcal{F}_{\omega_0,k}^2$ differs from $\mathcal{F}_{\omega_0,k}^1$ since the base frequency of the used activation changes. Compared with the $\mathcal{F}_{\omega_0,k}^1$ in the first sub-function, the $\mathcal{F}_{\omega_0,k}^2$ in the second one could be approximated as,

$$\mathcal{F}_{\omega_0,k}^2 \approx \left\{ \frac{\sqrt{4\pi+1}-1}{\sqrt{12\pi+1}-\sqrt{4\pi+1}} \omega \mid \omega \in \mathcal{F}_{\omega_0,k}^1 \right\}, \quad (7)$$

where $\frac{\sqrt{4\pi+1}-1}{\sqrt{12\pi+1}-\sqrt{4\pi+1}}$ characterizes the scale of frequency changes from the first sub-function to the second one. As a result, the supported frequency set $\mathcal{F}_{\omega_0,k}$ for k is close to the origin is,

$$\mathcal{F}_{\omega_0,k} = \mathcal{F}_{\omega_0,k}^1 \cup \mathcal{F}_{\omega_0,k}^2. \quad (8)$$

k is far away from the origin point. For example, \vec{b} is initialized as 10. Because the frequencies of each sub-functions are further increased for $\vec{b} = 10$ (in Fig. 3(b), the frequency is increased from the orange box to purple box), the supported frequency set $\mathcal{F}_{\omega_0,k}$ is increased a lot compared with the one in Eqn. 8. In summary,

Proposition 2 *The supported frequency set $\mathcal{F}_{\omega_0,k}$ of FINER increases with the increase of the initialization range of \vec{b} , and the supported frequency set \mathcal{F}_{ω_0} of SIREN is a subset of $\mathcal{F}_{\omega_0,k}$.*

4.3. Neural Tangent Kernel Perspective

Neural tangent kernel (NTK) theory [12] views the training of neural network as kernel regression, where the convergence of the network on fitting signals could be derived by analysing the diagonal property of the NTK or the distribution of NTK's eigenvalues. Generally speaking, stronger diagonal property results in better shift-invariance and better convergence, larger eigenvalues leads to faster convergence for high-frequency components [1, 34].

Without loss of generality, we focus on a simple case, i.e., the signal to be learned has 1D input and 1D output and FINER has 1 hidden layer with n neurons, such a network could be written as $f(x; \theta) = \sum_{k=1}^n c_k \sigma(w_k x + b_k)$, where $\sigma(x) = \sin((|x|+1)x)$ is the activation function. According to the definition, the NTK of FINER could be calculated as¹,

$$\begin{aligned} \mathbf{K}(x_i, x_j) &= \mathbb{E}_{\theta} \langle \nabla_{\theta} f(x_i; \theta), \nabla_{\theta} f(x_j; \theta) \rangle \\ &= (x_i x_j + 1) \mathbb{E}_{\theta} \sum_{k=1}^n c_k^2 \underbrace{(2|g_k(x_i)| + 1)(2|g_k(x_j)| + 1)}_{\text{Scale term}} \\ &\quad \underbrace{\cos((|g_k(x_i)| + 1)g_k(x_i)) \cos((|g_k(x_j)| + 1)g_k(x_j))}_{\text{Sign term}} \end{aligned}$$

where $g_k(x_i) = w_k x_i + b_k$.

(9)

¹Please refer the supplemental material for details of derivation.

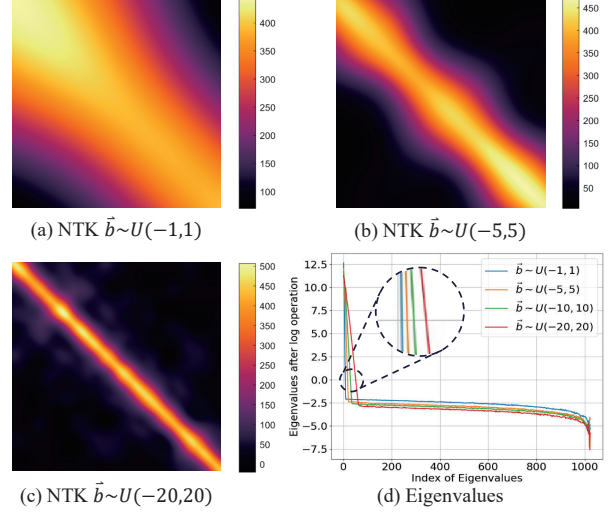


Figure 4. Visualizations of NTKs and the corresponding eigenvalues in FINER. (a)-(c) visualize the NTKs when \vec{b} is initialized following $\mathcal{U}(-1, 1)$, $\mathcal{U}(-5, 5)$ and $\mathcal{U}(-20, 20)$, respectively. (d) plots the corresponding eigenvalues. Because the max eigenvalue is much larger than the smallest one, all eigenvalues are processed by a log function for visualization.

It is observed that, the scale term is approximately proportional to the absolute of bias b_k for all nodes of the kernel, however the change rule of sign term for diagonal elements differs significantly from non-diagonal elements. Specifically, the sign term is always a positive value for diagonal elements while could be either positive or negative for non-diagonal elements.

As a result, the diagonal elements $\mathbf{K}(x_i, x_i)$ are increased with the increase of $|b|$, while the non-diagonal elements $\mathbf{K}(x_i, x_j)$ can be very small, appearing as a diagonal enhanced kernel. According to [12, 34], NTK with a stronger diagonal property provides better shift-invariance, i.e., the coordinates in the training set are little coupled with each other during the training process, thus the signal could be better learned.

Figs. 4(a)-(c) visualize the NTKs of FINER for learning a 1D signal with 1024 coordinates. It is observed that the diagonal property of NTK is enhanced with the increase of initialization range of b , verifying the analysis mentioned above. Fig. 4(d) visualizes the changes of eigenvalue distribution, it is observed that the number of eigenvalues which are larger than 10^0 is significantly increased when larger initialization range is applied to bias. As a result, FINER provides high capacity for learning high-frequency components.

5. Experiments

To verify the behaviors and performance of FINER, three experiments are conducted, including the 2D image fitting, 3D signed distance field representation and the 5D neural

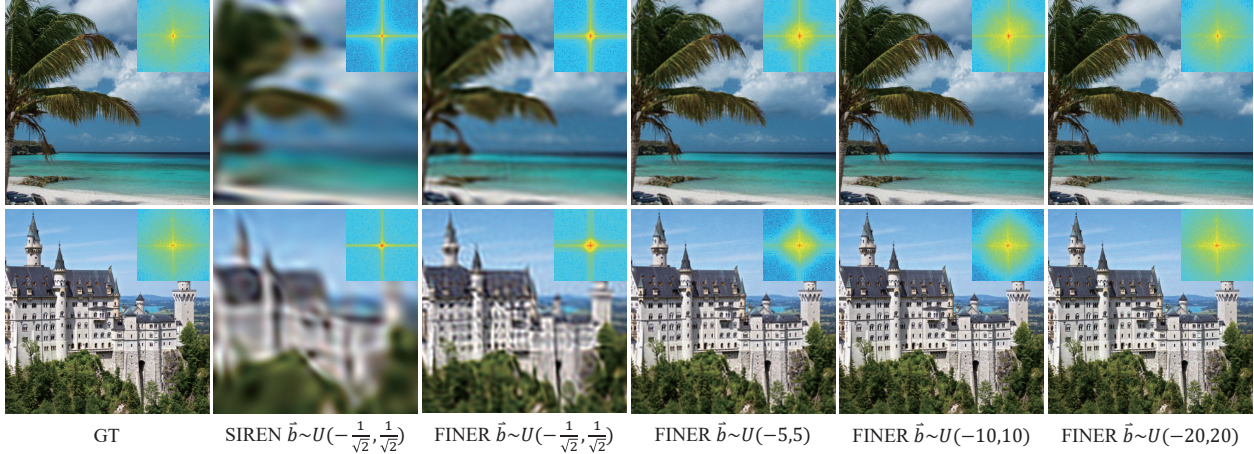


Figure 5. Comparisons of SIREN and FINER with different initializations applied to bias vector \vec{b} . For each image, the right-top box visualizes its Fourier spectrum.




Figure 6. Comparisons of the first layer outputs between SIREN and FINER. For each method, the first row demonstrates 4 neurons with smallest frequencies of the first layer, and the second row refers to 4 neurons with largest frequencies.

radiance fields optimization, where the ω_0 is equal to the one in SIREN and k is set as $\frac{1}{\sqrt{2}}$, 1 and $\frac{1}{\sqrt{3}}$, respectively.

5.1. 2D Image Fitting

For the task of 2D image representation, the INR aims at learning a 2D function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ with 2D pixel location input and 3D RGB color output, the loss function is designed as the L_2 distance between the network output and the ground truth. To evaluate the performance of INR, the widely used natural dataset [34] which contains 16 images with 512×512 resolution is adopted.

5.1.1 FINER behaviors under different initializations

To compare the performance of FINER under different initializations better, the pre-set parameter ω_0 is set as 1 here. According to the analysis in Sec. 4, the supported frequency set of FINER increases with the initialization range of \vec{b} . Different curves in Fig. 1(b) reflect this behavior. Fig. 5 visualizes the learned images after 5000 training epochs. In the 3rd column, because the initialization range of \vec{b} follows the standard method [32], FINER has a small supported frequency set, as a result, high-frequency components in the image are lost. With the increase of the initialization range,

Table 1. Quantitative comparisons on image fitting. We color code each cell as **best**, **second best**, and **third best**.

Metrics	PEMLP	Gauss	SIREN	WIRE	FINER
PSNR \uparrow	29.60	35.39	38.52	31.31	40.76
SSIM \uparrow	0.8484	0.9455	0.9724	0.8738	0.9790
LPIPS \downarrow	1.21e-1	1.91e-2	5.52e-3	6.45e-2	2.56e-3

the supported frequency set of FINER increases and more high-frequency details are provided from the 3rd column to the last one in Fig. 5. Additionally, as analysed by Eqn. 6, the supported frequency set of SIREN is a subset of the one in FINER, thus FINER provides more clear details than SIREN when same initializations are applied, *i.e.*, the 2nd column vs the 3rd column in Fig. 5.

According to [40], the first layer of SIREN or FINER plays the role of frequency encoding. We visualize 8 neurons outputs from total 256 neurons in first layer from SIREN and FINER for comparison, where 4 neurons in the 1st row have the smallest frequencies and the last 4 neurons have the largest frequencies (see Fig. 6). It is observed that different neurons in SIREN encode similar frequencies, resulting in a waste of neurons. On the contrary, different neurons in FINER focus on different frequencies, therefore, better representational ability is achieved in FINER.

5.1.2 Comparisons with the State-of-the-arts

We compare FINER with four classical INRs, *i.e.*, the Fourier feature embedding (PEMLP) [23], INR with periodic activation functions (SIREN) [32], Gaussian activation functions (Gauss) [28] and wavelet activation functions (WIRE) [30]. For a fair comparison, all INRs are set with a same network configuration (3 hidden layers with 256 neurons per layer) and trained with the same Adam optimizer [16] and L_2 loss function between the network out-



Figure 7. Qualitative comparisons between the FINER and baselines on fitting images.




Figure 8. Qualitative comparisons on representing the signed distance field of Armadillo.

put and the ground truth, other parameters are set according to the open-source codes released by authors. Tab. 1 compares FINER with others quantitatively. FINER outperforms other method in all three metrics. Fig. 7 demonstrates the details of different methods. Among all 5 methods, FINER provides more clear results such as the texts ‘67’ and ‘SBS’ in the cyan and green boxes, respectively. On the contrary, these texts are over-smoothed in the results of SIREN and PEMLP. Although WIRE and Gauss could also provide clear texts here, unwelcome serious artefacts also appear in the smooth backgrounds, such as white and red billboards.

5.2. 3D Shape Representation

Signed distance field (SDF) is one of the most commonly used implicit surface representations in the computer graphics [13]. As the name implies, SDF characterizes the distance from the given 3D point to the closest surface using a continuous function, and the sign of the distance is used to denote whether the point is inside (negative) or outside (positive) the surface. Recently, representing the SDF using INR is drawing more and more attention [18]. Given a 3D point \vec{x} , INR learns a 3D mapping function $f : \mathbb{R}^3 \rightarrow \mathbb{R}^1$ to output the signed distance field values s . We apply the FINER to this task and compare to four classical INRs mentioned above. In the experiment, 5 shapes from public dataset [17, 24] are used for evaluation. For a fair com-

Table 2. Quantitative comparisons on representing signed distance field.

	Methods	Armadillo	Dragon	Lucy	Thai Statue	BeardedMan	Avg.
Chamfer ↓	PEMLP	3.559e-6	2.081e-6	2.224e-6	5.284e-6	4.058e-6	3.441e-6
	Gauss	1.778e-5	7.427e-6	5.494e-6	1.618e-5	1.620e-5	1.262e-5
	SIREN	3.505e-6	2.759e-6	2.493e-6	4.481e-6	3.952e-6	3.438e-6
	WIRE	3.346e-6	2.101e-6	2.238e-6	3.979e-6	4.597e-6	3.252e-6
	FINER	3.348e-6	2.364e-6	2.119e-6	3.580e-6	4.023e-6	3.087e-6
IOU ↑	PEMLP	9.870e-1	9.750e-1	9.760e-1	9.526e-1	9.939e-1	9.769e-1
	Gauss	9.768e-1	9.679e-1	9.601e-1	9.481e-1	9.932e-1	9.692e-1
	SIREN	9.895e-1	9.666e-1	9.721e-1	9.514e-1	9.948e-1	9.749e-1
	WIRE	9.893e-1	9.723e-1	9.707e-1	9.565e-1	9.911e-1	9.760e-1
	FINER	9.899e-1	9.725e-1	9.756e-1	9.625e-1	9.943e-1	9.790e-1

parison, all methods use a same network configuration, *i.e.*, 3 layers with 256 neurons per layer, additionally, the same coarse-to-fine loss function is used according to [19]. In the training stage, 10k points are randomly sampled in each iteration and is repeated 200k iterations. In the testing stage, a 512^3 grid is extracted for evaluation and visualization.

Tab. 2 provides quantitative comparisons between the proposed FINER and four baselines. Because FINER provides more freedom for tuning the supported frequency set, best results are achieved among all methods. Fig. 8 compares the reconstructed details visually on the Armadillo rendered using Marching Cubes [21]. Two representative regions, *i.e.*, the low-frequency smooth pectoral and the high-frequency rough shank, are zoomed-in for comparisons. SIREN represents the smooth pectoral well but loses

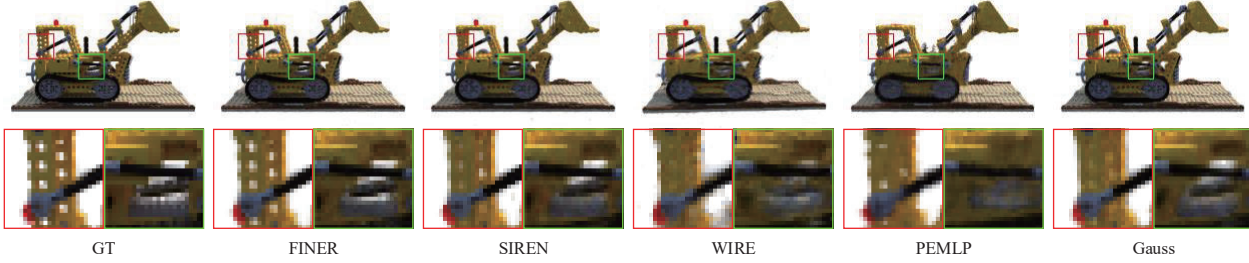


Figure 9. Qualitative comparisons between the FINER and baselines on NeRF.

the details of the shank, WIRE overcomes the limitation of SIREN for representing high-frequency shank at the cost of rough pectoral. For PEMLP, because the pre-defined frequency may not match the frequency distribution in the SDF of Armadillo, both the smooth pectoral and rough shank are not well represented. Gauss could not provide stable representation for all shapes (Tab. 2) and there are obvious artefacts in the reconstructed SDF such as the noise outside the shape in Fig. 8. Compared with these baselines, FINER could provide consistent performance for reconstructing both the low- and high-frequency components.

5.3. Neural Radiance Fields Optimization

Novel view synthesis, which aims at rendering realistic images at uncaptured poses from a set of images captured at different positions, is one of the key problems in both communities of computer vision and graphics. Recently, representing scenes as neural radiance fields (NeRF) [23] using INR dominates this task due to the advantages of realism and scalability for embedding different rendering processes [35]. Given a 3D point \vec{x} and a 2D observed direction \vec{d} , NeRF focuses on learning a 5D mapping function $f: \mathbb{R}^5 \rightarrow \mathbb{R}^4$ with 5D (\vec{x}, \vec{d}) to its 3D color c and 1D opacity σ . For any pixel p in novel view images, its ray function in 3D space is firstly calculated using the in/extrinsic matrices of camera [9], then N points are sampled along the ray within a predefined depth range, furthermore the direction and position coordinates of these N points are fed into the INR for querying the radiance value (c, σ) , finally the color $C(p)$ of the pixel is calculated using the differentiable volume rendering technique [22, 23],

$$C(p) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) c_i \quad (10)$$

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right),$$

where δ_i represents the distance between the neighbors in the sampled N points.

We evaluate the FINER for this task and compare to four classical methods mentioned above. To better verify the advantages of FINER for representing high-frequency compo-

Table 3. Quantitative comparisons on novel view synthesis.

	Methods	Chair	Drums	Ficus	Hotdog	Lego	Materials	Mic	Ship
PSNR \uparrow	PEMLP	31.32	20.18	24.49	30.59	25.90	25.16	26.38	21.46
	Gauss	32.68	23.16	26.10	32.17	28.29	26.19	33.59	22.28
	SIREN	33.31	24.89	27.26	32.85	29.60	27.13	33.28	22.25
	WIRE	29.31	22.22	25.91	30.11	25.76	25.05	32.35	21.15
	FINER	33.90	24.90	28.70	33.05	30.04	27.05	33.96	22.47
SSIM \uparrow	PEMLP	0.960	0.814	0.914	0.945	0.904	0.909	0.960	0.754
	Gauss	0.967	0.883	0.933	0.956	0.932	0.916	0.980	0.782
	SIREN	0.971	0.912	0.947	0.960	0.948	0.932	0.979	0.788
	WIRE	0.938	0.858	0.931	0.938	0.886	0.898	0.978	0.734
	FINER	0.973	0.911	0.958	0.959	0.951	0.928	0.981	0.792
LPIPS \downarrow	PEMLP	0.026	0.185	0.056	0.037	0.070	0.050	0.057	0.190
	Gauss	0.019	0.082	0.055	0.032	0.040	0.037	0.013	0.136
	SIREN	0.017	0.051	0.033	0.032	0.030	0.028	0.014	0.116
	WIRE	0.035	0.106	0.045	0.048	0.075	0.058	0.014	0.172
	FINER	0.015	0.052	0.024	0.033	0.025	0.032	0.010	0.108

nents, we follow the experimental setting of WIRE that only 25 images are used for training instead of commonly used 100 images, additionally the same ω_0 is set in both FINER and SIREN. Tab. 3 and Fig. 9 provide quantitative and qualitative comparisons of FINER against different methods on the Blender dataset [23]. FINER achieves the best performance in almost all 8 scenes. Fig. 9 demonstrates the advantage of FINER for representing high-frequency components visually. For example, the holes (red boxes) and the highlights (green boxes) in the frame of the truck are over-smoothed in the reconstructions of all baselines, only FINER could provide clear details here.

6. Conclusion

We have proposed and verified the FINER which uses variable-periodic functions for activating the INR. We have pointed out that current INRs suffer from a limited supported frequency set due to the under-utilization of the definition domain of activation functions. The proposed FINER overcomes this problem by introducing variable-periodic activation function and initializing the bias vector to different ranges, where different sub-functions with different frequencies along the definition domain will be selected for activation. As a result, the supported frequency set of FINER could be flexibly tuned, and the capacity of INR could be significantly enlarged. We have demonstrated the advantages of FINER over other INRs in image fitting, 3D shape representation and neural rendering.

References

- [1] Jinshuai Bai, Gui-Rong Liu, Ashish Gupta, Laith Alzubaidi, Xi-Qiao Feng, and YuanTong Gu. Physics-informed radial basis network (pirbn): A local approximating neural network for solving nonlinear partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 415: 116290, 2023. 5
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2
- [3] Hao Chen, Bo He, Hanyu Wang, Yixuan Ren, Ser Nam Lim, and Abhinav Shrivastava. Nerv: Neural representations for videos. In *Advances in Neural Information Processing Systems*, pages 21557–21568, 2021. 2
- [4] Shin-Fang Chng, Sameera Ramasinghe, Jamie Sherrah, and Simon Lucey. Gaussian activated neural radiance fields for high fidelity reconstruction and pose estimation. In *European Conference on Computer Vision*, pages 264–280. Springer, 2022. 3
- [5] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021. 2
- [6] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2020. 2
- [7] Ruohan Gao, Zilin Si, Yen-Yu Chang, Samuel Clarke, Jeanette Bohg, Li Fei-Fei, Wenzhen Yuan, and Jiajun Wu. Objectfolder 2.0: A multisensory object dataset for sim2real transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10598–10608, 2022. 2
- [8] Zhongkai Hao, Songming Liu, Yichi Zhang, Chengyang Ying, Yao Feng, Hang Su, and Jun Zhu. Physics-informed machine learning: A survey on problems, methods and applications. *arXiv preprint arXiv:2211.08064*, 2022. 2
- [9] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 8
- [10] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18398–18408, 2022. 2
- [11] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, and Qing Wang. Inverting the imaging process by learning an implicit camera model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21456–21465, 2023. 2
- [12] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018. 5
- [13] M.W. Jones, J.A. Baerentzen, and M. Sramek. 3d distance fields: a survey of techniques and applications. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):581–599, 2006. 7
- [14] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021. 1, 2
- [15] Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. Layered neural atlases for consistent video editing. *ACM Transactions on Graphics (TOG)*, 40(6):1–12, 2021. 2
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. 6
- [17] Stanford Computer Graphics Laboratory. The stanford 3d scanning repository, 2014. <http://graphics.stanford.edu/data/3Dscanrep/>. 7
- [18] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H. Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8456–8465, 2023. 7
- [19] David B Lindell, Dave Van Veen, Jeong Joon Park, and Gordon Wetzstein. Bacon: Band-limited coordinate networks for multiscale scene representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021. 2, 7
- [20] Renhao Liu, Yu Sun, Jiabei Zhu, Lei Tian, and Ulugbek S Kamilov. Recovery of continuous 3d refractive index maps from discrete intensity-only measurements using neural fields. *Nature Machine Intelligence*, 4(9):781–791, 2022. 2
- [21] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. 7
- [22] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 8
- [23] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. 2, 6, 8
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 2, 7
- [25] Nithin Raghavan, Yan Xiao, Kai-En Lin, Tiancheng Sun, Sai Bi, Zexiang Xu, Tzu-Mao Li, and Ravi Ramamoorthi. Neural free-viewpoint relighting for glossy indirect illumination. In *Computer Graphics Forum*, page e14885. Wiley Online Library, 2023. 2
- [26] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019. 2

- [27] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019. [2](#)
- [28] Sameera Ramasinghe and Simon Lucey. Beyond periodicity: Towards a unifying framework for activations in coordinate-mlps. In *European Conference on Computer Vision*, pages 142–158. Springer, 2022. [2](#), [3](#), [6](#)
- [29] Vishwanath Saragadam, Jasper Tan, Guha Balakrishnan, Richard G Baraniuk, and Ashok Veeraraghavan. Miner: Multiscale implicit neural representation. In *European Conference on Computer Vision*, pages 318–333. Springer, 2022. [2](#)
- [30] Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G Baraniuk. Wire: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18507–18516, 2023. [2](#), [3](#), [6](#)
- [31] Rajhans Singh, Ankita Shukla, and Pavan Turaga. Polynomial implicit neural representations for large diverse datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2041–2051, 2023. [2](#)
- [32] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020. [1](#), [2](#), [3](#), [4](#), [6](#)
- [33] Yannick Strümpfer, Janis Postels, Ren Yang, Luc Van Gool, and Federico Tombari. Implicit neural representations for image compression. In *European Conference on Computer Vision*, pages 74–91. Springer, 2022. [2](#)
- [34] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. [2](#), [5](#), [6](#)
- [35] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, W Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *Computer Graphics Forum*, pages 703–735. Wiley Online Library, 2022. [1](#), [8](#)
- [36] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34:27171–27183, 2021. [2](#)
- [37] Zhijie Wu, Yuhe Jin, and Kwang Moo Yi. Neural fourier filter bank. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14153–14163, 2023. [2](#)
- [38] Shaowen Xie, Hao Zhu, Zhen Liu, Qi Zhang, You Zhou, Xun Cao, and Zhan Ma. DINER: Disorder-invariant implicit neural representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2023. [2](#)
- [39] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Computer Graphics Forum*, 41(2):641–676, 2022. [1](#)
- [40] Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. A structured dictionary perspective on implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19228–19238, 2022. [2](#), [3](#), [6](#)
- [41] Hao Zhu, Zhen Liu, You Zhou, Zhan Ma, and Xun Cao. DNF: Diffractive neural field for lensless microscopic imaging. *Optics Express*, 30(11):18168–18178, 2022. [2](#)
- [42] Junyu Zhu, Hao Zhu, Qi Zhang, Fang Zhu, Zhan Ma, and Xun Cao. Pyramid nerf: Frequency guided fast radiance field optimization. *International Journal of Computer Vision*, pages 1–16, 2023. [2](#)