Yang Chen Ruituo Wu, Yipeng Liu, *Senior Member* Ce Zhu, *Fellow IEEE*



Figure 1: In this work, we propose a novel universal solution for inverse problems based on implicit neural representation (INR) - HOIN. Compared with traditional INR methods such as WIRE [33], INCODE [22], and SIREN [37], HOIN significantly improves the model's ability to perceive high-frequency information, effectively characterizes signal details, and achieves the best performance in a series of classic inverse tasks such as image denoise, super-resolution, and inpainting.

ABSTRACT

Implicit neural representations (INR) suffer from worsening spectral bias, which results in overly smooth solutions to the inverse problem. To deal with this problem, we propose a universal framework for processing inverse problems called **High-Order Implicit Neural Representations (HOIN)**. By refining the traditional cascade structure to foster high-order interactions among features, HOIN enhances the model's expressive power and mitigates spectral bias through its neural tangent kernel's (NTK) strong diagonal properties, accelerating and optimizing inverse problem resolution. By analyzing the model's expression space, high-order derivatives, and the NTK matrix, we theoretically validate the feasibility of HOIN. HOIN realizes 1 to 3 dB improvements in most inverse problems, establishing a new state-of-the-art recovery quality and training efficiency, thus providing a new general paradigm for INR and paving the way for it to solve the inverse problem.

KEYWORDS

Implicit Neural Representation, Inverse Problem, High-Order Feature Interaction.

1 INTRODUCTION

Convolutional Neural Networks (CNNs) effectively learn signals but struggle with high-frequency signals due to high impedance, i.e. spectral bias, which has become a significant challenge in signal processing. But deep image prior (DIP) [42] takes advantage of spectral bias, successfully tackling image restoration tasks such as denoising, super-resolution, and other visual inverse challenges [5, 12, 17]. DIP benefits from its independence from vast external data sets. However, its need for many training parameters and extended time frames limit practical use [1, 18, 19].

Implicit Neural Representation (INR) [37] refines signal modeling by integrating coordinate inputs with neural networks. By its structural advantages [34], this method efficiently addresses inverse problems with reduced parameter count and processing time [22, 33]. However, traditional INR methods can lead to worsening spectral bias, which results in overly smoothed solutions that omit vital high-frequency details [33]. New strategies have been introduced to solve this problem, such as adding an encoding layer that elevates coordinate inputs to higher-dimensional spaces [14, 26, 28, 36, 39] and utilizing periodic [22, 25, 37] or non-periodic activation functions [30, 33]. These modifications aim to fine-tune frequency responses automatically, mitigating the issue of spectral bias and enhancing the DIP process, thus getting more detailed outcomes.

However, the existing solutions still have challenges. On the one hand, They tend to be tailored for specific tasks [33], needing more versatility for the broad spectrum of inverse problems. On the other hand, experimental evidence suggests that while these solutions may reduce spectral bias, they often fail to restore highfrequency details completely [22, 25, 33]. During our experiments, as shown in Figure 3, we observed that hash coding [15, 27], despite its efficiency in eliminating spectral bias, inadvertently blends high-frequency noise with the signal when applied to inverse problems. This issue is incredibly challenging in tasks such as image denoising and deblurring. Therefore, there is an evident need for a

All the authors are with the School of Information and Communication Engineering, University of Electronic Science and Technology of China (UESTC), Chengdu, 611731, China. (e-mail: yangchen2023@std.uestc.edu.cn, wrt786842305@gmail.com, eczhu@uestc.edu.cn, yipengliu@uestc.edu.cn).

solution universally applicable to all types of inverse problems and appropriately addresses spectral bias.

Incorporating high-order interaction structures [4, 44] into neural networks has dramatically expanded the hypothesis space, rapidly enhancing the ability to learn specific signal characteristics. This advancement is notably present in networks like Transformers [43] and Polynomial Neural Networks (PNNs) [9, 21, 46], which integrate multiplicative interactions and successfully address highfrequency signal processing challenges. Inspired by recent advancements, we diverge from traditional reliance on coding layers and activation functions, introducing an MLP block focused on higherorder feature interactions to present High-Order Implicit Neural Representations for Inverse Problems (HOIN), a novel, generalized approach for tackling inverse problems through INRs. Our research shows that HOIN enhances the translational invariance and eigenvalue distribution in the Neural Tangent Kernels (NTK) [8, 20] linked to INRs, thereby expanding the functional space of the model. HOIN significantly improves its capacity to mitigate spectral bias, excels at modeling high-frequency signals, and effectively minimizes noise interference.

To conclude, our contributions can be summarized as follows:

- We propose a new high-order interaction block to mitigate the worsening spectral bias in INR.
- We propose a universal inverse problem-handling framework, the HOIN, that can apply INR to any inverse problem.
- We analyze the expression ability, higher-order derivatives, and NTK matrices of the higher-order blocks and theoretically prove the higher-order blocks' effectiveness.
- HOIN maintains the state-of-the-art (SOTA) performance in various models that use INR to solve inverse problems and representation tasks.

2 BACKGROUND

2.1 Inverse Problem

Solving inverse problems, which aim to reconstruct original signals from measurements, is crucial in critical applications like image restoration and sound source localization. Traditional methods for these problems often rely on existing knowledge, proposing solutions that meet certain conditions or combine an understanding of the target's structure with sparsity assumptions [2, 6, 32, 40]. However, These approaches face challenges in more complex situations. Deep learning methods introduce innovative solutions like deep image priors (DIP) [42] and implicit neural representations (INR) [37]. Leveraging spectral deviation priors, INR can rapidly address and derive solutions for inverse problems from a single sample. This approach demands fewer parameters than methods based on convolutional networks and markedly abbreviates the training duration.

2.2 Implicit Neural Representation Details

INR [37] uses coordinate grids to approximate continuous signals, showing advantages in rendering, computational imaging, medical imaging, and virtual reality over traditional methods [7, 24, 26]. Recently, INR's approach to solving inverse problems has gained notable attention [33].

In an inverse problem, suppose coordinate inputs $\mathbf{x} \in \mathbb{R}^{D_i}$ corresponding to the clean signals $S(\mathbf{x}) : \mathbb{R}^{D_i} \mapsto \mathbb{R}^{D_o}$ and the noise signal $N(\mathbf{x}) : \mathbb{R}^{D_i} \mapsto \mathbb{R}^{D_o}$. For image, we have the coordinate input (x_i, x_j) , and the corresponding image $S(\mathbf{x}) \in \mathbb{R}^{3 \times H \times W}$. The noise signal can be modeled as

$$N\left(\mathbf{x}\right) = S\left(\mathbf{x}\right) + \mathbf{n},\tag{1}$$

where **n** is assumed to be *i.i.d.* Gaussian Noise drawn from $\mathcal{N}(0, \sigma^2 \mathbf{I})$ with **I** being the identity matrix.

INR parameterizes the clean signal $S(\mathbf{x})$ via a network $F_{\theta}(\mathbf{x}) : \mathbb{R}^{D_i} \mapsto \mathbb{R}^{D_o}$ and is optimized to fit the noisy signal $N(\mathbf{x})$, formulated as:

$$\theta^* = \underset{\theta}{\arg\min} \mathcal{L}\left(N\left(\mathbf{x}\right); F_{\theta}(\mathbf{x})\right), \quad S^*\left(\mathbf{x}\right) = F_{\theta^*}(\mathbf{x}). \tag{2}$$

Such parameterization allows lower-frequency contents to be fitted before the higher-frequency ones, exhibiting high impedance to signal noises or degradations. In practice, θ is usually learned using an MLP, and the overall network architecture of INR is as follows:

$$z_{0} = \gamma(\mathbf{x}),$$

$$z_{l} = \varphi(\mathbf{C}_{l} z_{l-1})$$

$$= \varphi(\mathbf{W}_{l} z_{l-1} + \mathbf{b}_{l}), l = 1, 2, \dots, L - 1,$$

$$F_{\theta}(\mathbf{x}) = \mathbf{W}_{L} z_{L-1} + \mathbf{b}_{L},$$
(3)

where \mathbf{z}^{l} denotes the output of layer $l, \theta = \{W^{l}, \mathbf{b}^{l} \mid l = 1, 2, ..., L - 1\}$, *L* is the number of layers, φ is the nonlinear activation function, $\gamma(\cdot)$ is the coding layer. C_{l} is linear function with respect to \mathbf{z}_{l-1} .

2.3 Motivation

In traditional approaches, INR can tackle inverse problems but is constrained by worsening spectral bias. This worsening spectral bias typically leads to excessively smooth solutions that lack crucial high-frequency details. Methods such as nonlinear activation functions and high-dimensional encoding have been implemented to mitigate this issue, but their effectiveness across a wide range of inverse problems is limited. Acknowledging these challenges, our goal is to conduct an in-depth analysis of the root causes of spectral bias and devise a solution strategy that is more universally applicable.

3 HOIN: HIGH-ORDER IMPLICIT NEURAL REPRESENTATIONS

3.1 Overview

In this Section, we introduce High-Order Implicit Neural Representations for Inverse Problems (HOIN). As illustrated in Figure 2, the HOIN framework is through the coding layer and the highorder interaction block stages. 1) The coding layer transforms the coordinate input of signals (e.g. audio, image, video, etc.) \mathbf{x} into a high-dimensional space (Section 3.2); 2) By utilizing various activation functions, the high-order interaction block facilitates complex interactions among features within this expanded space (Section 3.3). During training, we find the peak performance point and stop the fitting process there.



Figure 2: Overview of HOIN. We select the corresponding encoding layer based on the type of inverse problem, mapping the coordinate input x into a higher dimensional space $\gamma(x)$. Then, the low-frequency and high-frequency information in the signal is captured through a High-Order Block structure. During training, we find the peak performance point, stop the fitting process there, and find the solution $F_{\theta}(x)$. \odot denotes Hadamard product, \boxplus is addition, φ is the nonlinear activation function.

3.2 Encoding Layer

Using the encoding layer first in HOIN aims to alleviate spectral bias by mapping the signal coordinates to high-dimensional space, enhancing the model's ability to capture details [26, 39]. INR alleviates spectral bias by mapping the signal coordinates to high-dimensional space, enhancing the model's ability to capture details [26, 39]. In addressing inverse problems, the deployment of coding layers has become crucial. Essential coding methods include positional coding (Pos. Enc), Fourier features (FFN), and hash table mapping (Instant-NGP). In our proposed HOIN framework, as shown in Figure 2, we adopt the following specific encoding strategies based on different types of inverse problems:

• Hash Table [27]:

$$\gamma(\mathbf{x}) = \left(\bigoplus_{d=1}^{D_i} x_d \pi_d\right) \mod T,\tag{4}$$

• Position Encoding [26]:

$$\gamma(\mathbf{x}) = \left[\cos\left(2\pi\sigma^{j/m}\mathbf{x}\right), \sin\left(2\pi\sigma^{j/m}\mathbf{x}\right)\right]^{\mathrm{T}} j = 0, \dots, m, \quad (5)$$

• Fourier Features [39]:

$$\gamma(\mathbf{x}) = [\cos(2\pi \mathbf{B}\mathbf{x}), \sin(2\pi \mathbf{B}\mathbf{x})]^{\mathrm{T}}, \qquad (6)$$

where \oplus denotes the bit-wise XOR operation and π_d are unique large prime numbers. *T* is the size of the hash table. each entry in $\mathbf{B} \in \mathbb{R}^{m \times D_i}$ is sampled from $\mathcal{N}(0, \sigma^2)$, *m* is the mapping size, and σ is chosen for each task and dataset with a hyperparameter sweep.

3.3 High-Order Interaction Block

3.3.1 **Rethinking Plain Block and Residual Block**. To address the worsening spectral bias in inverse problems, past enhancements have mainly concentrated on refining the coding layer and activation function, overlooking the crucial role of the MLP architecture within the INR. This oversight leaves the cascade architecture of the MLP unexamined, which is instrumental in the root cause of the worsening spectral bias [25, 47].

For the classic INR model, the Plain MLP Block is

• Plain Block [29]:

$$\mathbf{z}_l = \varphi(\mathbf{C}_l \mathbf{z}_{l-1}). \tag{7}$$

3

The cascade effect observed in plain blocks is the primary cause of spectral bias in INR. This problem presents itself in two significant ways: Firstly, with an increase in the number of block layers, the vanishing gradient issue becomes more pronounced, making the training process more challenging [29]. Secondly, using the ReLU activation function can lead to the loss of high-order signal derivatives, further intensifying spectral bias [16].

Residual blocks featuring residual connections have been a method to improve gradient flow to deeper layers. The expression is as follows

• Residual Block [16]:

$$\mathbf{z}_{l} = \varphi((\mathbf{I} + \mathbf{C}_{l}) \, \mathbf{z}_{l-1}). \tag{8}$$

However, the residual block continues to face challenges with worsening spectral bias and learning high-frequency information [3]. It has yet to improve the efficiency of processing inverse problems significantly.

3.3.2 **HO Block**. Inspired by high-order interactions [10, 13, 31] in neural networks, we introduce a novel element into the MLP architecture of INR, which is called the **High-Order (HO) Block**. This addition aims to facilitate complex feature interactions at a higher level than traditional methods. This structure can be expressed in the following form:

$$\mathbf{z}_{l} = \varphi((\mathbf{J} + \mathbf{C}_{l}\mathbf{z}_{l-1}) \odot \mathbf{z}_{l-1}), \tag{9}$$

where \odot denotes Hadamard product, **J** is the all-one matrix

We incorporate second-order interaction within the plain block by multiplying the previous layer's outputs with those of the current layer and then summing them up. This approach evolves into creating HO blocks through a hierarchical linkage, empowering the model to facilitate $2^{(L-1)/2}$ -order feature interactions. Such an augmentation in high-order interactions diminishes the model's reliance on low-frequency learning. With increasing model depth, high-order blocks are designed to avoid the issue of gradient vanishing, enabling effective fitting of both high and low frequencies in the initial stages of training. This capacity allows for a swift alignment with the objective function of the real signal $G(\mathbf{x})$.



Figure 3: (a) Comparison of learning speeds at different frequencies. The target image is transformed into 10 frequency bands through the Fourier transform (x-axis, 0 represents the lowest frequency band), and we compare the learned components with the proper amplitude. On the color chart scale, 1 represents a perfect approximation. HO block can effectively alleviate spectral bias. Hash encoding does not exhibit spectral bias. (b) PSNR learning curves for different blocks. HO block maintains the highest PSNR.

The HOIN framework tailors its approach to various inverse problems by selecting suitable encoding layers and activation functions to meet the specific demands of each task. An overly aggressive correction for spectral bias and the rapid acceleration of high-frequency learning might inadvertently blend noise with the signal, often detrimentally affecting the task. We introduced HO blocks to SIREN [37], Pos.Enc [26], and FFN [39], creating **HO-SIREN**, **HO-Pos.Enc**, and **HO-FFN**, respectively. For particular inverse problem scenarios, we evaluate these models to identify the most effective one for deployment.

4 THEORETICAL ANALYSIS OF HOIN

In this section, we perform a theoretical analysis of HO blocks. In Sections 4.1, 4.2, and 4.3, we analyze the expressive ability, highorder derivatives, and NTK properties of various blocks.

4.1 Expression Ability Exploration

In INR frameworks, the dimension of a network's functional space is a crucial metric for assessing the network's capacity for expression [4]. The architecture of the network is denoted by $\mathcal{D} = \{D_1, ..., D_l\}$, where D_l indicates the number of neurons in the *l*-th layer. Any given activation function block can be decomposed into a series of polynomial functions with leading degree *r* through Taylor approximation. This process helps understand how activation functions and network configurations influence INR models' functional capacity and expressiveness.

For the network architecture D_l with an activation function of leading degree r, we represent the leading functional space of the neural network as $\mathcal{F}_{\mathcal{D},r}$. The leading functional variants of plain Block, residual block, and HO block can be defined as the *Zariski closure* [23] of their leading functional space, i.e. $\mathcal{V}_{\mathcal{D},r}^P, \mathcal{V}_{\mathcal{D},r}^R$ and $\mathcal{V}_{\mathcal{D},r}^{HO}$ (similar to the ones presented in [4, 23]). Using these definitions, we have:

THEOREM 4.1. For an activation function with leading degree $r \ge 1$ and network architecture $\mathcal{D} = \{D_1, ..., D_l\}$, the leading functional variety of Plain Block, $\mathcal{V}_{\mathcal{D},r}^P$, HO Block, $\mathcal{V}_{\mathcal{D},r}^{HO}$, and Residual Block, $\mathcal{V}_{\mathcal{D}r}^{R}$, satisfy:

$$\mathcal{V}_{\mathcal{D},r}^{HO} = \mathcal{V}_{\mathcal{D},2r}^{R} = \mathcal{V}_{\mathcal{D},2r}^{P}.$$
 (10)

Proof. See Supplementary

Theorem 4.1 posits that within an identical network architecture, HO block with a leading degree of r and plain and residual block with a leading degree of 2r exhibit the same variation in homogeneous polynomial functions. This implies that HO block can support a diversity of leading functions, specifically $(2r)^{l-1}$ homogeneous polynomials, in contrast to neural networks of the same structure and activation function, which are limited to r^{l-1} . Consequently, the HO block possesses a more expansive expression space, representing a broader range of frequency signal components.

Besides, we give the frequency decay rates for different blocks as follows:

THEOREM 4.2. For a single-layer MLP, the number of neurons is d. For frequency k, the frequency decay rate of the Plain and Residual Block is k^{-d} , while on the contrary, the frequency decay rate of the HO Block is $k^{-d/2}$.

Proof. See Supplementary

Theorem 4.2 suggests that HO blocks can capture more highfrequency information than plain blocks, facilitating a faster resolution of inverse problems.

We perform an image representation experiment to verify the phenomenon of spectral bias for different models. We employ an MLP model with three hidden layers, utilizing a ReLU activation function and positional encoding, to train on authentic natural images from the DIV2K [41] dataset. Our evaluation centers on the model's capacity to learn information across different frequency bands. Following the methodology outlined in [35], we partition the image spectrum into ten frequency bands and monitor the model's learning progress on these frequency bands during the training process. The experimental results are depicted in Figure 3, where darker red colors indicate weaker learning abilities of the model for those frequency bands. The findings suggest that the plain block struggles to learn high-frequency information in the image, even



Figure 4: (a) Visualization of NTK and corresponding eigenvalues in different models. (b) Draw the corresponding feature values. Because the maximum eigenvalue is much larger than the minimum eigenvalue, all eigenvalues are processed by logarithmic functions for visualization. HO blocks significantly enhance the eigenvalues on the diagonal of the NTK matrix, thus enhancing the ability of the INR to capture high-frequency information. Plain, residual, and high-order blocks are abbreviated as P, R, and HO.

with positional encoding. Conversely, the HO block captures highfrequency features early in training, showcasing fast learning rates and excellent representation ability.

Additionally, upon introducing mainstream hash encoding representations, we observe its capability to learn balanced high-frequency and low-frequency information, exhibiting no spectral bias. Hash coding, a lattice-based interpolation method, learns low and highfrequency information. However, in inverse tasks, this leads to the blending of signal noise and high-frequency details, causing noise to be fitted early in the training, which is undesirable for inverse tasks. Specific experiments on this are conducted in the supplementary material.

4.2 Derivative Analysis

In signal processing, first and second derivatives are pivotal for encapsulating rich high-frequency information. Spectral bias often arises because high-order derivatives gravitate towards zero during the learning process with plain and residual blocks, leading to a substantial loss of high-frequency details. The HO block is ingeniously designed to mitigate this issue. Specifically, the first derivative of the HO block ∇z_I is

$$\nabla \mathbf{z}_l = \mathbf{C}_l \odot \mathbf{z}_{l-1} + \mathbf{C}_l \mathbf{z}_{l-1} + \mathbf{J}. \tag{11}$$

The second derivative of the HO Block $\Delta \mathbf{z}_l$ is

$$\Delta \mathbf{z}_l = 2\mathbf{C}_l. \tag{12}$$

Furthermore, we compute the first-order derivatives of the plain and residual blocks, with the findings presented in Table 1. According to Table 1, the second derivative of plain blocks equals zero, indicating a significant loss of signal detail during processing. In contrast, the second derivative of the HO block remains constant, a property that significantly enhances its ability to capture detailed signal information. This characteristic of the HO block is instrumental in efficiently accelerating the resolution of inverse problems by preserving and leveraging high-frequency details often lost in traditional processing blocks. Table 1: First and second-order derivatives of different blocks. Plain, residual, and high-order blocks are abbreviated as P, R, and HO. O is an all-zero matrix.

Block	$\nabla \mathbf{z}_l$	$\Delta \mathbf{z}_l$
Р	C_l	0
R	$C_l + I$	0
HO	$\mathbf{C}_l \odot \mathbf{z}_{l-1} + \mathbf{C}_l \mathbf{z}_{l-1} + \mathbf{J}$	$2C_l$

4.3 Neural Tangent Kernel Perspective

Our proposed HOIN method can effectively capture the high-frequency components of the signal. However, it is difficult to study this characteristic of spectral bias theoretically. The function constructed by the neural network is implicit, and its dependence on low-frequency component learning cannot be directly analyzed. Recently, some researchers have studied the learning process of neural networks through kernel function approximation [20]. The neural tangent kernel theory uses a first-order Taylor expansion of the model parameters θ , that is:

$$F_{\theta}(\mathbf{x}) \approx F_{\theta_0}(\mathbf{x}) + (\theta - \theta_0)^{\top} \nabla_{\theta} F_{\theta_0}(\mathbf{x}).$$
(13)

When the width of the layer in $F_{\theta}(\mathbf{x})$ is close to infinity, and the learning rate of the optimizer is close to 0, $F_{\theta}(\mathbf{x})$ can converge to the kernel regression solution of the neural tangent kernel during the training process, i.e. the kernel function

$$\mathcal{K}_{\text{NTK}}\left(\mathbf{x}, \mathbf{x}'\right) = \mathbb{E}_{\theta \sim \mathcal{N}} \left\langle \frac{F_{\theta}(\mathbf{x})}{\partial \theta}, \frac{F_{\theta}(\mathbf{x}')}{\partial \theta} \right\rangle.$$
(14)

By analyzing the eigenvalue distribution of the NTK kernel function, we can deeply understand the learning behavior of the neural network [38, 47]. When the larger eigenvalues of the kernel function are mainly concentrated in the diagonal area, the kernel function exhibits better translation invariance [25]. This structural property enables the model to learn signals more efficiently during training.

Methods		8×			4×			2×			1×		
memous	#Param↓	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	LPIPS ↓	#Param↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
InstantNGP [27]	0.233	54.081	0.992	36.928	0.938	0.065	1.588	41.864	0.960	0.012 •	36.112	0.856	0.116
WIRE [33]	0.437	54.392 🔵	0.996 🔵	38.809	0.947	0.036	0.889	42.335	0.963	0.012	35.037	0.895	0.111
INCODE [22]	0.207	50.697	0.989	39.236	0.918	0.057	1.029	42.771 🛛	0.967	0.013	36.723	0.882	0.108
SIREN [37]	0.199	51.651	0.995	35.633	0.924	0.055	0.791 🗕	41.478	0.965	0.015	33.505	0.872	0.155
Pos. Enc [26]	0.204	31.016	0.901	31.322	0.872	0.136	0.805	36.789	0.939	0.046	32.954	0.884	0.159
FFN [39]	0.329	46.031	0.986	38.609	0.960	0.047	1.314	41.801	0.973	0.013	34.206	0.916	0.122
						Ours							
HO-SIREN	0.199 🗕	59.199 😐	0.997 😐	41.060	0.981	0.031 •	0.794	42.852 •	0.987	0.015	37.696	0.958	0.131
HO-Pos. Enc	0.206 •	44.638	0.991	40.034 •	0.978 •	0.024	0.805 •	41.974	0.986 •	0.017	37.095 •	0.954 •	0.103
HO-FFN	0.329	54.637	0.997	45.393 😐	0.991 😐	0.008 😐	1.317	46.845 😐	0.995 😐	0.004 😐	39.203 😐	0.967 😐	0.097 😐

Table 2: Results of image representation of different downsampling factors. Best 3 scores in each metric are marked with gold ●, silver ● and bronze ●.

In addition, when the feature value is larger, the model has a more vital high-frequency learning ability. This means that the model can respond more sensitively and learn high-frequency components in the signal.

We analyze the NTKs for various models, including activation functions like ReLU and SIREN (no encoding layers) and models that utilize encoding layers such as Pos. Enc and FFN coupled with a ReLU activation function. We visualize the NTK matrices for these models configured with Plain and high-order blocks. Unless otherwise specified, we all use three hidden layer networks to generate the NTK kernel matrix in the rest of this article.

As shown in Figure 4(a), the kernel function of the plain block exhibits poor diagonal properties, leading to challenges in learning both low and high frequencies. Conversely, the HO block kernel matrix showcases significant diagonal eigenvalues, facilitating effective learning of low-frequency signals while concurrently capturing high-frequency signals. Furthermore, the HO block exhibits excellent diagonal properties and large feature values for SIREN, Pos. Enc, and FFN. This attribute is a crucial factor contributing to the successful characterization of high-frequency signals by these INRs. The model's diagonal width is further reduced upon integrating the high-order structure, and the eigenvalues are augmented. This advancement enhances the model's capacity to capture high-frequency information beyond the original model, facilitating nearly simultaneous learning of high-frequency and low-frequency information. Figure 4(b) showcases the shift in the eigenvalue distribution, revealing a significant increase in the number of eigenvalues exceeding 10¹ when utilizing HO blocks. This observation underlines the enhanced capability of the HOIN framework to learn high-frequency information effectively.

5 EXPERIMENTS

In this Section, we conduct an extensive experimental evaluation of HOIN. Our experimental setup is detailed in Section 5.1. In Sections 5.2, 5.3, 5.4, 5.5, and 5.6, we explore the application of the HOIN framework to specific image inversion tasks, including image representation, denoising, super-resolution, CT reconstruction, and image completion. Additional ablation studies and visualizations

are provided in the supplementary experiments section for further insights into the effectiveness and operational mechanisms of HOIN.

5.1 Experimental details and setup

For our experimental benchmarks, we select four models: WIRE [33], SIREN [37], Pos. Enc [26], and FFN [39]. To these models, we integrate the HO block into SIREN, Pos. Enc, and FFN to identify the most effective configurations, collectively termed HOIN. We also add InstantNGP [27] and INCODE [22]. The experimental scope includes tasks such as image representation (5000 epochs), image denoising (2000 epochs), image super-resolution (2000 epochs), CT image reconstruction (5000 epochs), and image completion (1000 epochs). The data for these tasks comprise randomly selected 10 images with dimensions of $1644 \times 2040 \times 3$ from the DIV2K [41] dataset. We set the evaluation metrics of the Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) [45]. Further experimental details are available in the supplementary materials.

5.2 Image Represention

Image representation can be viewed as a particular inverse problem. Its performance intuitively presents the model's improvement in spectral bias. In the experiment, we subject the images to various degrees of downsampling, namely $1\times$, $2\times$, $4\times$, and $8\times$, to cater to the diverse requirements of image representation at different downsampling rates. For the experiments involving $1\times$ and $2\times$ downsampling rates, the hidden layers are configured with 512 neurons each. Conversely, in the experiments with $4\times$ and $8\times$ downsampling, the number of neurons per hidden layer is set to 256. To maintain a fair comparison across all models, we adjust the parameter count of the InstantNGP model to align with the order of magnitude of the parameters in the other models. The outcomes of these experiments are presented in Table 2.

Table 2 demonstrates that integrating the High-Order structure markedly enhances our model's capacity for high-frequency representation, yielding PSNR results compared to the baseline model. Notably, the HO-FFN model records the highest PSNR, registering approximately 8.1dB greater than the original FFN model. Figure 5 illustrates the error distribution during the reconstruction of 2×



Figure 5: Visualization of Image Representation. Here, we demonstrate the representation errors of different models. The brighter areas indicate higher representation errors. HO-FFN accurately reconstructs all the detailed information of the image.

Table 3: Image denoising results under different Gaussian noise σ .

Methods		10			25			50	50	
methods	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	
InstantNGP [27]	29.483	0.773	0.211	22.786	0.540	0.398	17.754	0.382	0.567	
WIRE [33]	30.802	0.846	0.185	26.411	0.725	0.323	22.864	0.589	0.451	
INCODE [22]	30.478	0.811	0.196	25.113	0.644	0.357	21.841	0.509	0.489	
SIREN [37]	31.140	0.849	0.169	26.748 🛛	0.734 鱼	0.309	23.407	0.625	0.426	
Pos. Enc [26]	27.271	0.736	0.281	25.756	0.685	0.306 😐	23.459	0.643	0.414	
FFN [39]	31.131	0.852	0.180	26.335	0.716	0.317	22.805	0.583	0.453	
				Ours						
HO-SIREN	32.338	0.896	0.144	27.043	0.766	0.300	23.451 •	0.642 •	0.425 •	
HO-Pos. Enc	32.452 😐	0.910 😐	0.123 😐	27.561 😐	0.801 😐	0.264 😐	23.858 😐	0.677 😐	0.390 😐	
HO-FFN	32.057 鱼	0.882 鱼	0.152 •	26.596	0.726	0.310	22.990	0.574	0.447	

downsampled images, revealing the HO-FFN model's near-complete reconstruction of high-frequency details, including edges and textures. For models like Pos. Enc, introducing HO structures can also effectively enhance their ability to represent high-frequency information.

5.3 Image Denoise

In the image denoising experiment, to each image to assess the denoising abilities of the models, we add Gaussian noise with three noise levels, including $\sigma = 10$, $\sigma = 25$, and $\sigma = 50$. Each model is set up with a hidden layer containing 256 neurons. The results of these experiments are detailed in the following Table 3.

Table 3 reveals that the HO-Pos.Enc model, benefiting from the moderate acceleration provided by the HOIN framework in learning high-frequency information, exhibits superior performance across all denoising experiments. Networks utilizing the ReLU activation function have effectively learned low-frequency information, whereas the HOIN framework has demonstrated a significant advantage in acquiring high-frequency details. Furthermore, in line with previous analyses, excessive acceleration in high-frequency information learning by models such as InstantNGP and HO-FFN can result in the undesirable amalgamation of high-frequency noise with details. This conflation can detrimentally affect the outcome of denoising tasks. Detailed discussions and visualizations related to the denoising experiments are thoroughly presented in the supplementary. Table 4: Results of Image Super-Resolution.

Methods	×2		×4		×	6	×8	
methods	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
InstantNGP [27]	19.74	0.400	16.42	0.203	16.44	0.212	16.35	0.246
WIRE [33]	31.50	0.846	29.09 ●	0.786	26.77 🛛	0.717	24.44	0.686
INCODE [22]	31.94	0.853	28.97	0.812	26.47	0.759	24.95 😐	0.694
SIREN [37]	31.61	0.851	28.26	0.803	26.23	0.736	24.18	0.715
Pos. Enc [26]	30.39	0.805	24.26	0.745	24.36	0.718	23.28	0.713
FFN [39]	31.38	0.856	27.93	0.795	26.16	0.783 \bullet	24.49	0.728 🛛
Ours								
HO-SIREN	33.03	0.898	29.61 😐	0.854 😐	27.53 😐	0.815 😐	25.69 😐	0.771 😐
HO-Pos. Enc	32.47 •	0.876 鱼	28.91	0.824 •	26.43	0.762	24.78	0.720
HO-FFN	33.10 😐	0.898 😐	29.30	0.839	27.30	0.798	25.44	0.759

5.4 Image Super-Resolution

In our image super-resolution experiment, we initially downsample the original images by 2, 4, 6, and 8 factors. These downsampled images are used in the training phase to leverage the inherent interpolation capabilities of INR. Subsequently, in the testing phase, we aim to restore them to their original dimensions. The comprehensive results of these experiments are meticulously documented in Table 4, showcasing the effectiveness of our approach in enhancing image super-resolution.

As shown in Table 4, the HOIN framework markedly enhances the performance across all evaluated models. Notably, HO-SIREN exhibits outstanding PSNR and SSIM metrics across most superresolution tasks. In contrast, due to its intrinsic methodology, the InstantNGP model, which relies on hash table indexes for reconstruction, proves less adept for pixel-aligned super-resolution tasks. Additional visualization results and detailed analyses are available in the supplementary.

5.5 CT Reconstruction

Our CT image reconstruction experiment utilizes 10 CT lung images from the publicly accessible lung nodule analysis dataset on Kaggle [11]. To assess the efficacy of our model in CT reconstruction tasks, these images are downsampled to a resolution of 256×256 . The experiment involves measuring reconstruction at four angles: 50, 100, 200, and 300. The findings of these experiments are comprehensively detailed in Table 5.



PSNR: 28.97 SSIM: 0.71

PSNR: 29.93 SSIM: 0.72

SSIM: 0.96 PSNR: 30.93 PSNR: 32.45 SSIM: 0.95

Figure 6: Visualization of Computed Tomography Reconstruction. Here, we demonstrate various methods for CT-based reconstruction of 256 × 256 images at 100 angles. HO-FFN maintains the best reconstruction results.

PSNR: 34.15

Methods	50		10	00	20	00	300		
	PSNR ↑	SSIM ↑							
InstantNGP [27]	17.56	0.569	18.65	0.662	20.59	0.743	22.21	0.795	
WIRE [33]	21.93	0.648	26.28	0.799	29.01	0.814	29.20	0.818	
INCODE [22]	22.76	0.674	26.63	0.701	31.16	0.819	32.22 ●	0.861	
SIREN [37]	22.96	0.714	26.96	0.745	27.97	0.822	30.32	0.847	
Pos. Enc [26]	22.72	0.734	23.78	0.784	24.20	0.809	24.30	0.801	
FFN [39]	26.03	0.779 鱼	30.17 •	0.898	31.79 🜒	0.925	32.03	0.936	
-			01	ırs					
HO-SIREN	28.02 😐	0.866	32.12 😐	0.932	34.41	0.963 😐	34.82	0.968 😐	
HO-Pos. Enc	26.94	0.906 😐	28.39	0.933 😐	28.79	0.944 😐	29.40	0.949 鱼	
HO-FFN	26.85 鱼	0.769	30.82	0.912 •	34.90 😐	0.962	34.83 😐	0.961	

Table 5: CT reconstruction results from different angles.

CT reconstruction involves creating computational images from sensor measurements, with sparse CT reconstruction tackling the challenge of producing accurate images from limited measurement data. This challenge is primarily due to the difficulty in reconstructing images with scarce data. HO block substantially improves the quality of the reconstruction results by efficiently capturing highfrequency components throughout the reconstruction process. As detailed in Table 5, the HOIN model outperforms others in all measurement scenarios, showcasing its superior performance. Figure 6 illustrates that the HO-SIREN model is particularly adept at reconstructing images' texture and contour details. In comparison, the SIREN model, much like the WIRE and INCODE models, is prone to artifacts, whereas InstantNGP struggles with significant pixel loss issues.

5.6 Image Inpainting

In the image inpainting experiment, we select an image of a Celtic spiral knot with a resolution of 572×582×3. The mask applied in this experiment is generated randomly, obscuring approximately 10% of the image's pixel area. The architectural configuration for this experiment is aligned with that used in the image representation task, with the findings presented in Figure 7.

Compared to existing SOTA methods based on INR, the HO-FFN model demonstrates considerable superiority in image inpainting tasks, particularly in accurately rendering details. To corroborate



SSIM: 0.94

PSNR: 18.66 SSIM: 0.58 PSNR: 19.47 SSIM: 0.67

Figure 7: Visualization of Image Inpainting. Here, we only use 10% of the original image's pixels for reconstruction. HO-SIREN effectively reconstructs detail levels such as texture edges.

the efficacy of our approach in image completion tasks, additional relevant experiments are included in the supplementary for further examination.

CONCLUSION 6

In this paper, we propose High-Order Implicit Neural Representations for Inverse Problems (HOIN), an innovative framework for addressing inverse problems. By integrating high-order interaction blocks into INR, HOIN substantially enlarges the functional

space of INR to enhance the model's capacity to capture highfrequency information. The NTK matrix associated with HOIN features notable diagonal and translational invariance, offering robust theoretical backing to mitigate spectral bias. Unlike alternative approaches, HOIN is adept at diminishing noise interference and swiftly and efficiently resolving inverse problems. Through comprehensive experiments, HOIN has been shown to outperform other models utilizing INR for inverse problem-solving and has also excelled in representation tasks.

REFERENCES

- Metin Ersin Arican, Ozgur Kara, Gustav Bredell, and Ender Konukoglu. 2022. Isnas-dip: Image-specific neural architecture search for deep image prior. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 1960–1968.
- [2] Richard G Baraniuk, Volkan Cevher, Marco F Duarte, and Chinmay Hegde. 2010. Model-based compressive sensing. *IEEE Transactions on information theory* 56, 4 (2010), 1982–2001.
- [3] Yuval Belfer, Amnon Geifman, Meirav Galun, and Ronen Basri. 2021. Spectral analysis of the neural tangent kernel for deep residual networks. arXiv preprint arXiv:2104.03093 (2021).
- [4] Jie Bu and Anuj Karpatne. 2021. Quadratic residual networks: A new class of neural networks for solving forward and inverse problems in physics involving pdes. In Proceedings of the 2021 SIAM International Conference on Data Mining (SDM). SIAM, 675–683.
- [5] Prithvijit Chakrabarty and Subhransu Maji. 2019. The spectral bias of the deep image prior. arXiv preprint arXiv:1912.08905 (2019).
- [6] Antonin Chambolle. 2004. An algorithm for total variation minimization and applications. Journal of Mathematical imaging and vision 20 (2004), 89–97.
- [7] Yinbo Chen, Sifei Liu, and Xiaolong Wang. 2021. Learning continuous image representation with local implicit image function. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 8628–8638.
- [8] Moulik Choraria, Leello Tadesse Dadi, Grigorios Chrysos, Julien Mairal, and Volkan Cevher. 2022. The spectral bias of polynomial neural networks. arXiv preprint arXiv:2202.13473 (2022).
- [9] Grigorios G Chrysos, Stylianos Moschoglou, Giorgos Bouritsas, Jiankang Deng, Yannis Panagakis, and Stefanos Zafeiriou. 2021. Deep polynomial neural networks. *IEEE transactions on pattern analysis and machine intelligence* 44, 8 (2021), 4021–4034.
- [10] Grigorios G Chrysos, Bohan Wang, Jiankang Deng, and Volkan Cevher. 2023. Regularization of polynomial networks for image recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 16123– 16132.
- [11] Kenneth Clark, Bruce Vendt, Kirk Smith, John Freymann, Justin Kirby, Paul Koppel, Stephen Moore, Stanley Phillips, David Maffitt, Michael Pringle, et al. 2013. The Cancer Imaging Archive (TCIA): maintaining and operating a public information repository. *Journal of digital imaging* 26 (2013), 1045–1057.
- [12] Mohammad Zalbagi Darestani and Reinhard Heckel. 2021. Accelerated MRI with un-trained neural networks. *IEEE Transactions on Computational Imaging* 7 (2021), 724–733.
- [13] Feng-Lei Fan, Mengzhou Li, Fei Wang, Rongjie Lai, and Ge Wang. 2023. On expressivity and trainability of quadratic networks. *IEEE Transactions on Neural Networks and Learning Systems* (2023).
- [14] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. 2020. Multiplicative filter networks. In International Conference on Learning Representations.
- [15] Sharath Girish, Abhinav Shrivastava, and Kamal Gupta. 2023. Shacira: Scalable hash-grid compression for implicit neural representations. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 17513–17524.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition. 770–778.
- [17] Reinhard Heckel and Paul Hand. 2018. Deep decoder: Concise image representations from untrained non-convolutional networks. arXiv preprint arXiv:1810.03982 (2018).
- [18] Reinhard Heckel and Mahdi Soltanolkotabi. 2019. Denoising and regularization via exploiting the structural bias of convolutional generators. arXiv preprint arXiv:1910.14634 (2019).
- [19] Kary Ho, Andrew Gilbert, Hailin Jin, and John Collomosse. 2021. Neural architecture search for deep image prior. *Computers & graphics* 98 (2021), 188–196.
- [20] Arthur Jacot, Franck Gabriel, and Clément Hongler. 2018. Neural tangent kernel: Convergence and generalization in neural networks. Advances in neural information processing systems 31 (2018).
- [21] Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF

conference on computer vision and pattern recognition. 4401-4410.

- [22] Amirhossein Kazerouni, Reza Azad, Alireza Hosseini, Dorit Merhof, and Ulas Bagci. 2024. INCODE: Implicit Neural Conditioning with Prior Knowledge Embeddings. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. 1298–1307.
- [23] Joe Kileel, Matthew Trager, and Joan Bruna. 2019. On the expressive power of deep polynomial neural networks. Advances in neural information processing systems 32 (2019).
- [24] Ålexandr Kuznetsov. 2021. NeuMIP: Multi-resolution neural materials. ACM Transactions on Graphics (TOG) 40, 4 (2021).
- [25] Zhen Liu, Hao Zhu, Qi Zhang, Jingde Fu, Weibing Deng, Zhan Ma, Yanwen Guo, and Xun Cao. 2023. FINER: Flexible spectral-bias tuning in Implicit NEural Representation by Variable-periodic Activation Functions. arXiv preprint arXiv:2312.02434 (2023).
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- [27] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant neural graphics primitives with a multiresolution hash encoding. ACM transactions on graphics (TOG) 41, 4 (2022), 1–15.
- [28] Nithin Raghavan, Yan Xiao, Kai-En Lin, Tiancheng Sun, Sai Bi, Zexiang Xu, Tzu-Mao Li, and Ravi Ramamoorthi. 2023. Neural Free-Viewpoint Relighting for Glossy Indirect Illumination. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, e14885.
- [29] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. 2019. On the spectral bias of neural networks. In *International conference on machine learning*. PMLR, 5301– 5310.
- [30] Sameera Ramasinghe and Simon Lucey. 2022. Beyond periodicity: Towards a unifying framework for activations in coordinate-mlps. In *European Conference* on Computer Vision. Springer, 142–158.
- [31] Yongming Rao, Wenliang Zhao, Yansong Tang, Jie Zhou, Ser Nam Lim, and Jiwen Lu. 2022. Hornet: Efficient high-order spatial interactions with recursive gated convolutions. Advances in Neural Information Processing Systems 35 (2022), 10353-10366.
- [32] Yaniv Romano, Michael Elad, and Peyman Milanfar. 2017. The little engine that could: Regularization by denoising (RED). SIAM Journal on Imaging Sciences 10, 4 (2017), 1804–1844.
- [33] Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G Baraniuk. 2023. Wire: Wavelet implicit neural representations. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 18507–18516.
- [34] Nimrod Shabtay, Eli Schwartz, and Raja Giryes. 2022. Pip: Positional-encoding image prior. arXiv preprint arXiv:2211.14298 (2022).
- [35] Zenglin Shi, Pascal Mettes, Subhransu Maji, and Cees GM Snoek. 2022. On measuring and controlling the spectral bias of the deep image prior. *International Journal of Computer Vision* 130, 4 (2022), 885–908.
- [36] Rajhans Singh, Ankita Shukla, and Pavan Turaga. 2023. Polynomial implicit neural representations for large diverse datasets. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2041–2051.
- [37] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. 2020. Implicit neural representations with periodic activation functions. Advances in neural information processing systems 33 (2020), 7462–7473.
- [38] Alex J Smola and Bernhard Schölkopf. 1998. On a kernel-based method for pattern recognition, regression, approximation, and operator inversion. *Algorithmica* 22 (1998), 211–231.
- [39] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. 2020. Fourier features let networks learn high frequency functions in low dimensional domains. Advances in neural information processing systems 33 (2020), 7537–7547.
- [40] Robert Tibshirani. 1996. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Series B: Statistical Methodology 58, 1 (1996), 267– 288.
- [41] Radu Timofte, Shuhang Gu, Jiqing Wu, and Luc Van Gool. 2018. Ntire 2018 challenge on single image super-resolution: Methods and results. In Proceedings of the IEEE conference on computer vision and pattern recognition workshops. 852–863.
- [42] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. 2018. Deep image prior. In Proceedings of the IEEE conference on computer vision and pattern recognition. 9446–9454.
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems 30 (2017).
- [44] Yan Wang, Lingxi Xie, Chenxi Liu, Siyuan Qiao, Ya Zhang, Wenjun Zhang, Qi Tian, and Alan Yuille. 2017. Sort: Second-order response transform for visual recognition. In Proceedings of the IEEE International Conference on Computer Vision. 1359–1368.

- [45] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. 2003. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference* on Signals, Systems & Computers, 2003, Vol. 2. Ieee, 1398–1402.
- [46] Zirui Xu, Fuxun Yu, Jinjun Xiong, and Xiang Chen. 2022. Quadralib: A performant quadratic neural network library for architecture optimization and design exploration. *Proceedings of Machine Learning and Systems* 4 (2022), 503–514.
- [47] Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. 2022. A structured dictionary perspective on implicit neural representations. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 19228–19238.

The supplementary material provides detailed theoretical analyses, proofs (see section A), and numerous additional experiments (see section B and C).

A THEORETICAL ANALYSIS

In this section, we detail the theory and proofs regarding the expression spaces and high-order derivatives of HOIN, as referenced in the main paper.

A.1 Expression Ability Exploration

In INR frameworks, the dimension of a network's functional space is a crucial metric for assessing the network's capacity for expression [4]. The architecture of the network is denoted by $\mathcal{D} = \{D_1, ..., D_l\}$, where D_l indicates the number of neurons in the *l*-th layer. Any given activation function block can be decomposed into a series of polynomial functions with leading degree *r* through Taylor approximation. This process helps understand how activation functions and network configurations influence INR models' functional capacity and expressiveness.

For the network architecture D_l with an activation function of leading degree r, we represent the leading functional space of the neural network as $\mathcal{F}_{\mathcal{D},r}$. The leading functional variants of plain Block, residual Block, and HO block can be defined as the *Zariski closure* [23] of their leading functional space, i.e., $\mathcal{V}_{\mathcal{D},r}^P, \mathcal{V}_{\mathcal{D},r}^R$ and $\mathcal{V}_{\mathcal{D},r}^{HO}$ (similar to the ones presented in [4, 23]) as follows

Definition A.1. Suppose there is a neural network $\mathcal{D} = \{D_1, ..., D_l\}$ whose leading space satisfies the following condition:

$$\mathcal{F}_{\mathcal{D},r} = \operatorname{Sym}_{r^{l-1}} \left(\mathbb{R}^{D_0} \right)^{D_l}.$$
 (15)

For a filling functional variety, its leading functional variety satisfies:

$$\mathcal{V}_{\boldsymbol{d},r} = \overline{\mathcal{F}_{\mathcal{D},r}} = \overline{\operatorname{Sym}_{r^{l-1}}\left(\mathbb{R}^{D_0}\right)^{D_l}}.$$
 (16)

Thus, the function space or the varieties of the network do not have to completely occupy the ambient space of homogeneous polynomials. Instead, we only need to consider the space of homogeneous polynomials whose leading degrees are contained as being adequately filled.

PROPOSITION A.2. For a single-layer network $\mathcal{D} = (D_l)$ utilizing a linearly activated (r = 1) High-Order (HO) Block, the network has a filling functional space of degree 2. That is, its leading functional space satisfies the following criteria:

$$\mathcal{F}_{\mathcal{D},1}^{HO} = \operatorname{Sym}_2(\mathbb{R}^{D_l})^{D_l}.$$
(17)

PROOF. We can relate the linear HO Block to a quadratic polynomial regression. Consider a HO Block:

$$z_{l} = (\mathbf{J} + \mathbf{C}_{l} z_{l-1}) \odot z_{l-1}$$

= $z_{l-1} + (\mathbf{C}_{l} z_{l-1}) \odot z_{l-1},$ (18)

Where \mathbf{z}_l represents a primary linear term with $\operatorname{Sym}_1(\mathbb{R}^{D_l})^{D_l}$, and $(C_l \mathbf{z}_{l-1}) \odot \mathbf{z}_{l-1}$ is a quadratic term with $\operatorname{Sym}_2(\mathbb{R}^{D_l})^{D_l}$. This quadratic term outlines the primary functional space of the HO Block. By Definition A.1, a single-layer HO Block encompasses a filling functional space of degree 2.

THEOREM A.3. For an activation function with leading degree $r \ge 1$ and network architecture $\mathcal{D} = \{D_1, ..., D_l\}$, the leading functional variety of Plain Block, $\mathcal{V}_{\mathcal{D},r}^P$, HO Block, $\mathcal{V}_{\mathcal{D},r}^{HO}$, and Residual Block, $\mathcal{V}_{\mathcal{D},r}^R$, satisfy:

$$\mathcal{V}_{\mathcal{D},r}^{HO} = \mathcal{V}_{\mathcal{D},2r}^{R} = \mathcal{V}_{\mathcal{D},2r}^{P}.$$
(19)

PROOF. This can be proven by discussing the equivalence of functional space for every Block using Proposition A.2. For the *i*-th layer in the HO Block, i = 1, 2, ..., l, before applying nonlinear activation, it has $\mathcal{V}_{(D_l),1}^{HO} = \operatorname{sym}_2(\mathbb{R}^{D_l})^{D_i} = \mathcal{V}_{(D_l),2}^P = \mathcal{V}_{(D_l),2}^R$ (since a single-layer Block with polynomial activation of degree 2 has a filling functional space of degree 2). This proves the case for r = 1. For nonlinear activations of leading degree r, applying the activation function to the space $\mathcal{V}_{(D_l),1}^{HO}$, we obtain: $\mathcal{V}_{(D_l),r}^{HO} = \left(\mathcal{V}_{(D_l),2}^P\right)^{\otimes r} = \mathcal{V}_{(D_l),2r}^P$, where \otimes denotes Kronecker product. Since the relation applies to each layer, thus we have $\mathcal{V}_{D,r}^{HO} = \mathcal{V}_{D,2r}^P = \mathcal{V}_{D,2r}^R$.

Theorem A.3 posits that within an identical network architecture, HO block with a leading degree of r and plain and residual block with a leading degree of 2r exhibit the same variation in homogeneous polynomial functions. This implies that the HO block can support a diversity of leading functions, specifically $(2r)^{l-1}$ homogeneous polynomials, in contrast to neural networks of the same structure and activation function, which are limited to r^{l-1} . Consequently, the HO block possesses a more expansive expression space, representing a broader range of frequency signal components.

Besides, we give the frequency decay rates for different blocks as follows:

THEOREM A.4. For a single-layer MLP, the number of neurons is d. For frequency k, the frequency decay rate of the Plain and Residual Block is k^{-d} , while on the contrary, the frequency decay rate of the HO Block is $k^{-d/2}$.

PROOF. For Plain and Residual Block, the frequency attenuation rate and proof are shown in [3]. For HO Block, the frequency attenuation rate and proof are shown in [8].

A.2 Derivative Analysis

The first and second derivatives are pivotal for encapsulating rich high-frequency information in signal processing. Spectral bias often arises because high-order derivatives gravitate towards zero during the learning process with plain and residual blocks, leading to a substantial loss of high-frequency details. The HO block is ingeniously designed to mitigate this issue.



Figure 8: Results of the audio representation. (a) Reconstruction error. (b) Reconstructed PSNR. HO-SIREN excels by minimizing reconstruction errors and demonstrating rapid convergence.

Consider a HO Block in 18, the first derivative of the HO block $\nabla \mathbf{z}_l$ is as follows

$$\nabla \mathbf{z}_{l} = \frac{\partial \mathbf{z}_{l}}{\partial \mathbf{z}_{l-1}} = \partial \mathbf{z}_{l-1} \left((\mathbf{J} + \mathbf{C}_{l} \mathbf{z}_{l-1}) \odot \mathbf{z}_{l-1} \right) = \left(\partial \mathbf{z}_{l-1} \left(\mathbf{J} + \mathbf{C}_{l} \mathbf{z}_{l-1} \right) \right) \odot \mathbf{z}_{l-1} + \left(\mathbf{J} + \mathbf{C}_{l} \mathbf{z}_{l-1} \right) \odot \left(\partial \mathbf{z}_{l-1} \left(\mathbf{z}_{l-1} \right) \right) = \mathbf{C}_{l} \odot \mathbf{z}_{l-1} + \mathbf{C}_{l} \mathbf{z}_{l-1} + \mathbf{J}.$$
(20)

The second derivative of the HO Block $\Delta \mathbf{z}_l$ is

$$\Delta \mathbf{z}_{l} = \nabla (\nabla \mathbf{z}_{l})$$

$$= \frac{\partial (\nabla \mathbf{z}_{l})}{\partial \mathbf{z}_{l-1}}$$

$$= \partial \mathbf{z}_{l-1} (\mathbf{C}_{l} \odot \mathbf{z}_{l-1} + \mathbf{C}_{l} \mathbf{z}_{l-1} + \mathbf{J})$$

$$= \partial \mathbf{z}_{l-1} (\mathbf{C}_{l} \odot \mathbf{z}_{l-1}) + \partial \mathbf{z}_{l-1} (\mathbf{C}_{l} \mathbf{z}_{l-1} + \mathbf{J})$$

$$= (\partial \mathbf{z}_{l-1} (\mathbf{C}_{l})) \odot \mathbf{z}_{l-1} + \mathbf{C}_{l} \odot (\partial \mathbf{z}_{l-1} (\mathbf{z}_{l-1})) + \mathbf{C}_{l}$$

$$= \mathbf{C}_{l} + \mathbf{C}_{l}$$

$$= 2\mathbf{C}_{l}.$$
(21)

The second derivative of the HO block remains constant, a property that significantly enhances its ability to capture detailed signal information. This characteristic of the HO block is instrumental in efficiently accelerating the resolution of inverse problems by preserving and leveraging high-frequency details often lost in traditional processing blocks.

B ADDITIONAL EXPERIMENTS AND DETAILS

In this section, we broaden the scope of our experiments to provide a more thorough comparison between our method, HOIN, and the current state-of-the-art (SOTA) method. We show that the inherent simplicity of HOIN leads to enhanced performance, particularly in terms of expressiveness and the ability to tackle inverse problems, compared to the corresponding SOTA method. These results underscore our approach's effectiveness in extending the INR network's capabilities and enhancing its applicability across various domains. We now include additional visualizations that distinctly highlight the advantages of our method.

B.1 Experimental details

All implementations utilize MLP networks with three hidden layers. Our experiments use PyTorch on an Nvidia RTX 3080 Ti GPU with 12GB of RAM. We employ the Adam optimizer, complemented by a learning rate scheduler that decreases the learning rate by 0.1 upon completion of each epoch. Details on specific datasets and architectures are provided for the corresponding tasks.

B.2 Signal Representation

B.2.1 Audio.

Data: We use the first 7 seconds of Bach's Cello Suite No. 1: Prelude [22], with a sampling rate of 44100 Hz as our example for the audio representation task.

Architecture: All models use three hidden layers with 256 neurons per hidden layer. We set the first layer $w_0 = 10000$ for SIREN, HO-SIREN, INCODE. Each model is trained for a total of 1000 epochs. **Analysis:** We present the audio representation visualization results and their corresponding error plots in Figure 8. These visualizations are crucial for illustrating the strengths of our approach. Regarding sound playback quality, SIREN tends to introduce a distinct squeak-like sound accompanying the main audio. With INCODE, certain moments experience annoying noise, as the error chart indicates. However, HO-SIREN significantly reduces noise interference and outperforms the other methods.

B.2.2 Image.

Data: In the main paper and supplementary material experiments, we select one of the larger nature images on the Internet with a size of $3 \times 4844 \times 3219$.

Architecture: For all models except InstantNGP, we use three hidden layers, each with 512 neurons. We set the first layer's frequency parameter $w_0 = 30$ for SIREN, HO-SIREN, and INCODE. For the Wire, we set the scaling parameter $s_0 = 20$. Each model is trained for a total of 1000 epochs.

Analysis: Figure 9 displays the visualization results of large-size images. In terms of reconstruction quality, InstantNGP incorrectly represents some colors. Because of its inherent spectral bias, SIREN struggles with reconstructing high-frequency details, such as the top of the house, with distinct light and dark variations. HO-SIREN

Wu and Wang, et al.

InstantNGP



0







SIREN

PSNR: 31.72

0.1



Figure 9: Image representation results. The second row is an error map. The darker the red color, the higher the error. For PSNR, red is the best, and blue is the second best. HO-FFN is the best representation result.

and HO-FFN excel in capturing high-frequency details and consistently maintain the best overall quality, 3 dB better than the current top-performing WIRE (SOTA) model.

B.2.3 Occupancy Volume.

Data: We use the Lucy and Thai Statue datasets from the Stanford 3D Scanning Repository and follow the WIRE strategy [33]. We create an occupancy volume through point sampling on a 512×512×512 grid, assigning values of 1 to voxels within the object and 0 to voxels outside.

Architecture: Our network and training configuration is similar to the image representation task, with the difference that INR now maps 3D coordinates to signed distance function (SDF) values. Each model is trained for a total of 100 epochs.

Analysis: The results in Table 6 demonstrate the effectiveness of

Table 6: Best 3 scores in each metric are marked with gold •, silver • and bronze •.

Methods	Thai	Lucy
INCODE	0.9879	0.9951 😐
WIRE	0.9903 鱼	0.9718
SIREN	0.9758	0.9885
Pos.Enc	0.9872	0.9927
	Ours	
HO-SIREN	0.9935 😐	0.9948
HO-Pos.Enc	0.9918	0.9945 ●

HO-SIREN as a formidable option for occupancy representation



Figure 10: The result of image denoising. HO-Pos. enc maintains the best PSNR and SSIM.



Figure 11: The result of image super-resolution. The second line is the error map. HO-SIREN maintains the best PSNR and SSIM and can accurately characterize textures, edges, and other detailed information.

tasks. HO-SIREN significantly improves representation by effectively utilizing the HO Block to capture complex interactions between features. This capability is especially evident in enhancing high-frequency information while maintaining excellent capture of low-frequency details. Our approach achieves higher "Intersection over Union" (IOU) values, significantly enhancing object detail and scene complexity and rendering more accurately than existing methods.

B.3 Inverse Problems

B.3.1 Image denoising.

Data: In the experiments in the main paper and supplementary materials, we employ an image from DIV2K dataset [41], downsampled by a factor of 1/4 from $1152 \times 2040 \times 3$ to $288 \times 510 \times 3$. we add Gaussian noise with three noise levels, including $\sigma = 10$, $\sigma = 25$, and $\sigma = 50$.

Architecture: The setup for the denoising experiment closely mirrors that of the image characterization experiment, with the modification that the neurons in each model were adjusted to 256.



Figure 12: Reconstruction results for different projects. HO-FFN always maintains the best reconstruction results, keeping the highest PSNR and SSIM.

Throughout the training process, we monitored the Peak signal-tonoise ratio (PSNR) of both the noisy and clean images, considering the peak PSNR of the clean image as the final result of the reconstruction. Each model is trained for a total of 2000 epochs. **Analysis:** Experimental results for the three noise scales are presented in the main paper. Here, we visualize the experimental results of σ = 25. As shown in Figure 10, HO-Pos.Enc substantially enhances the fidelity of noisy images, achieving a 9.9 dB improvement in PSNR and a 0.22 increase in Structural Similarity Index (SSIM). Compared to the INCODE and SIREN methods, HO-Pos.Enc more effectively reduces noise artifacts while delicately preserving image

Methods	20	%	40	%	60	%	80%	
	PSNR ↑	SSIM ↑						
InstantNGP	20.719	0.734	23.400	0.829	25.263	0.871	26.747	0.896
WIRE	21.118	0.710	24.015	0.827	26.073	0.885 •	26.677	0.895
INCODE	22.097 •	0.781	24.796	0.859	25.922	0.877	27.279 •	0.903 •
SIREN	21.318	0.719	24.044	0.828	25.495	0.865	26.380	0.885
FFN	21.862	0.773	25.127	0.868	27.617	0.912	29.754	0.938
				Ours				
HO-SIREN	22.119	0.777 •	24.961 •	0.866 •	26.376 •	0.881	27.279	0.894
HO-FFN	22.357 😐	0.802 😐	25.962 😐	0.896 😐	28.764 😐	0.936 😐	31.573 😐	0.958 😐

Table 7: The result of image inpainting.

details. Furthermore, our method surpasses the Pos. Enc. method in terms of SSIM by 0.02.

B.3.2 Image super resolution.

Data: We adopt an image from the DIV2K dataset [41] and down-sampled the image with the size of $1356 \times 2040 \times 3$ by factors of 1/2, 1/4, 1/6, and 1/8.

Architecture: We maintain the same architectural and training settings as the image representation task. By employing a downsampled image during training, we exploit the interpolation capabilities of INRs to reconstruct an image of its original size in the test. Each model is trained for a total of 500 epochs.

Analysis: The experimental results for four upsampling factors are shown in the main paper. Here, we visualize the experimental results for one map with an upsampling factor of 4. The application of INRs as interpolators holds significant promise in super-resolution, leveraging inherent biases within INRs that can be utilized to enhance performance in such tasks. As depicted in Figure 11, HO-SIREN and HO-FFN consistently achieve superior PSNR and SSIM values across various super-resolution scales, surpassing competing methods. Additionally, HO-SIREN excels in reconstructing detailed elements like high-quality and transparent textures, avoiding the background artifacts commonly associated with SIREN.

B.3.3 CT reconstruction.

Data: Our CT image reconstruction experiment utilizes 10 CT lung images from the publicly accessible lung nodule analysis dataset on Kaggle [11]. To assess the efficacy of our model in CT reconstruction tasks, these images are downsampled to a resolution of 256×256 . The experiment measures reconstruction at four angles and projects: 50, 100, 200, and 300.

Architecture: We maintain the same architectural and training settings as the image representation task. We generate a sinogram according to the projection level using the radon transform. The model predicts a reconstructed CT image. Subsequently, we calculate the radon transform for the generated output and compute the loss function between these sinograms to guide the model toward generating CT images with reduced artifacts. Each model is trained for a total of 5000 epochs.

Analysis: CT reconstruction involves creating computational images from sensor measurements. In sparse CT reconstruction, the challenge is generating accurate images using only a limited subset of the available measurements, complicating the imaging process. As shown in Figure 12, HOIN addresses this challenge by effectively integrating higher-order interactions between features using the HO Block. The HO-FFN leverages 100 measurements to produce a sharp reconstruction with crisp details, achieving a notable improvement of 5.12 dB over the standard FFN, thus distinguishing itself in performance. In contrast, SIREN, similar to WIRE and INCODE, exhibited artifacts. This underscores the robustness of HO-FFN in managing noisy and undersampled inverse problems, demonstrating its potential as a promising solution for constrained image reconstruction, where it effectively balances image fidelity with noise reduction.

B.3.4 Inpainting.

Data: We utilize Celtic spiral knots image with a $572 \times 582 \times 3$ resolution. The sampling masks are generated randomly, with an average of 20%, 40%, 60%, and 80% of pixels being sampled.

Architecture: We use the same structure as the image representation. Each model is trained for a total of 500 epochs.

Analysis: Image inpainting poses a significant challenge, as the task requires the model to predict entire pixel values based on only a small portion of trained pixel data. The experimental results are shown in Table 7. The high capacity of INRs offers a unique advantage in addressing this inverse problem. The strong prior embedded within the INR function space facilitates applications such as repairs from finite observations, where the model leverages its learned representation to predict and fill in missing values. As seen in other tasks, HO-FFN excels in capturing complex features, particularly edges, which allows it to outperform other methods that often yield ambiguous results.

C THEORETICAL EXPERIMENTAL VERIFICATION

C.1 Convergence rate comparison

C.1.1 Experimental settings.

We use the band correspondence metric in [35] to check the inputoutput correspondence across multiple bands in the frequency domain Multiple bands For this metric, let $\{\theta^{(1)}, \ldots, \theta^{(T)}\}$ denote the trajectory of *T* steps of gradient descent in the parameter space and let $\{F_{\theta^{(1)}}, \ldots, F_{\theta^{(T)}}\}$ denote the corresponding trajectory in the output space. We propose to analyze the Fourier spectrum of the output images $F_{\theta^{(t)},t=1,\ldots,T}$ to show the convergence dynamics of different frequency components of the target image. The Fourier



Figure 13: Frequency-band correspondence metric. The left image shows an example of correspondence map H, which is computed according to Eq. (22). We divide the correspondence map into N subgroups corresponding to N non-overlapping frequency bands. Since the correspondence map is symmetrical around the center, we group it uniformly according to the distance between its elements and center, as illustrated by the right image when N = 5. Different colors represent different subgroups. We compute the mean correspondence for each band to transform the 2D map into the 1D one.

spectrum of the output image $F_{\theta^{(t)}}$ is obtained by the Fourier transform \mathscr{F} , denoted as $\mathscr{F}\{F_{\theta^{(t)}}\}$ for step *t*. We similarly compute the Fourier transform for the target image *G*, denoted as $\mathscr{F}\{G_0\}$. We then compute an element-wise correspondence between both transforms as:

$$H_{\theta^{(t)}} = \frac{\mathscr{F}\{F_{\theta^{(t)}}\}}{\mathscr{F}\{G_0\}}.$$
(22)

Intuitively, $H_{\theta(t)}$ denotes to what extent any deep image prior at step *t* corresponds with image G_0 in the frequency domain; the closer the values are to 1, the higher the correspondence. As we are interested in the spectral bias of the deep image prior, we divide the correspondence map into *N* subgroups corresponding to *N* non-overlapping frequency bands. Since the correspondence map is symmetrical around the center, we group it uniformly according to the distance between its elements and its center, as illustrated in Figure 13. To transform the 2D map to the 1D one, we compute the mean correspondence for each band, denoted as $\bar{H}_{\theta(t)}^{(n)}$, with $n=1,\ldots,N$. The value of $\bar{H}_{\theta(t)}^{(n)}$ indicates the convergence dynamics of different frequency components of a target image.

C.1.2 Experimental results.

This section analyzes various models' spectral bias and convergence speed, including InstantNGP, INCODE, SIREN, and Pos. Enc, FFN, HO-SIREN, HO-Pos. Enc, and HO-FFN. We use the configuration in section C.1.1. We conducted an image representation experiment where, as described in the main paper, the image is divided into 10 frequency bands from low to high. The metric is the ratio of the learned frequency band values to the true image for each epoch.

As depicted in Figure 14, the darker the red color, the less information is learned in that frequency band. Models like SIREN, Pos. Enc. and FFN struggle to learn high-frequency information in the early stages of training. However, introducing the HO Block significantly enhances the model's perception of high-frequency information. Among these, HO-SIREN and HO-FFN consistently achieve the best results. But for InstantNGP, the approach involves dividing the image into countless grids and simultaneously fitting the values at these grid points, effectively learning low-frequency and high-frequency information simultaneously. This method contributes to its rapid fitting capabilities.

We also visualize the results across different epochs. After just two epochs, HO-SIREN and HO-FFN could capture the outline and color features of the tiger, performing even better than InstantNGP. Throughout subsequent iterations, HO-FFN consistently maintained strong visual reconstruction results.

C.2 Spectral bias in inverse tasks

In this section, we address the application of mitigating spectral bias in inverse problems. Properly reducing spectral bias to enhance the perception of high-frequency information can accelerate the resolution of inverse problems. However, excessive acceleration might lead to premature coupling of high-frequency noise with the signal's high-frequency information, complicating the resolution of inverse problems. Using image denoising as an example, we visualize the PSNR learning curves for noisy and clean images in Figure 15. HO-Pos. Enc and HO-SIREN effectively speed up the learning process compared to Pos. Enc and SIREN are reflected in both high noisy PSNR and clean PSNR. Conversely, InstantNGP and HO-FFN, due to their overly aggressive mitigation of spectral bias, experience coupling of noise and signal at high frequencies, which is detrimental for image denoising tasks.

For the representation task, where excessive mitigation of spectral bias is not a concern, HO-FFN emerges as the top-performing model. In denoising tasks, HO-Pos.Enc strikes the best balance and proves to be the most effective model. For tasks involving superresolution, CT reconstruction, and inpainting, both HO-SIREN and HO-FFN stand out as the best models.



Figure 14: Comparison of learning speeds at different frequencies. The target image is transformed into 10 frequency bands through the Fourier transform (x-axis, 0 represents the lowest frequency band), and we compare the learned components with the proper amplitude. On the color chart scale, 1 represents a perfect approximation. HO block can effectively alleviate spectral bias.



Figure 15: The PSNR learning curves for noisy and clean images are in Figure 10. HO-Pos. Enc and HO-SIREN effectively speed up the learning process compared to Pos. Enc and SIREN are reflected in both high noisy PSNR and clean PSNR.

Wu and Wang, et al.



Figure 16: The results across different epochs. After just two epochs, HO-SIREN and HO-FFN can capture the outline and color features of the tiger, performing even better than InstantNGP. Throughout subsequent iterations, HO-FFN consistently maintains strong visual reconstruction results.