MIRE: Matched Implicit Neural Representations

Dhananjaya Jayasundara[1], Heng Zhao[2], Demetrio Labate[3], and Vishal M. Patel[1]
[1]Johns Hopkins University, [2]The Rockefeller University, [3]University of Houston
{vjayasu1,vpatel36}@jhu.edu, {hzhao25,dlabate2}@central.uh.edu

Abstract

Implicit Neural Representations (INRs) are continuous function learners for conventional digital signal representations. With the aid of positional embeddings and/or exhaustively fine-tuned activation functions, INRs have surpassed many limitations of traditional discrete representations. However, existing works only find a continuous representation for the digital signal by solely using a single, fixed activation function throughout the INR, and it has not yet been explored to match the INR to the given signal. As current INRs are not matched to the signal being represented by the INR, we hypothesize that this approach could restrict the representation power and generalization capabilities of INRs, limiting their broader applicability. A way to match the INR to the signal being represented is through matching the activation of each layer in the sense of minimizing the mean squared loss. In this paper, we introduce MIRE, a method to find the highly matched activation function for each layer in INR through dictionary learning. To showcase the effectiveness of the proposed method, we utilize a dictionary that includes seven activation atoms: Raised Cosines (RC), Root Raised Cosines (RRC), Prolate Spheroidal Wave Function (PSWF), Sinc, Gabor Wavelet, Gaussian, and Sinusoidal. Experimental results demonstrate that MIRE not only significantly improves INR performance across various tasks, such as image representation, image inpainting, 3D shape representation, novel view synthesis, super-resolution, and reliable edge detection, but also eliminates the need for the previously required exhaustive search for activation parameters, which had to be conducted even before INR training could begin.

1. Introduction

Implicit Neural Representations (INRs), also known as coordinate-based neural networks, operate by learning a continuous (implicit) functional representation when provided with the coordinates of an explicit signal representation. In general, an INR is structured as a multilayer perceptron (MLP) with several fully connected layers, where the explicit signal's coordinates serve as the input. Through the learning process of the MLP, the explicit representation is encoded into the weights and biases of the neural network. A distinctive feature of INRs is their versatility in handling different types of signals, from two-dimensional images through three-dimensional shapes and beyond. For example, in the context of images, an INR utilizes the coordinates from a two-dimensional grid to produce the corresponding color values at those coordinates, effectively learning a continuous representation for the image.

INRs stand in contrast to traditional discrete signal representation techniques, offering a more flexible and potentially more efficient means of representing complex signals [8]. Once the conversion of an explicit signal representation to an implicit representation through an INR is completed, a continuous functional relationship between the signal's coordinates and its values is established. This learned continuous implicit functional relationship, facilitated by INRs, serves as a robust representation mechanism for the underlying signal, allowing it to perform operations like precise querying of the learned representation and differentiation, etc. In contrast, discrete representations of signals encounter limitations in operations such as querying, which are constrained by quantized interpolations, and also differentiation may not yield desired outputs due to the discrete nature. Therefore, the inherent capabilities of INRs offer significant advantages in accurately representing and manipulating signals compared to discrete representations. In addition, while the memory requirement for conventional representations increases exponentially with the signal resolution, INRs are not tied to the resolution, making this approach highly memory-efficient [8, 34].

Despite the potential advantages of using INRs for the applications mentioned above, their performance critically depends on the architecture of the MLP, particularly the choice of activation function. Traditional activation functions, such as ReLU, Sigmoid, and Tanh, which are commonly used in deep learning models, have shown very poor

Project Page: https://dsgrad.github.io/MIRE

performance in INRs [27, 36]. This inefficiency is primarily due to their inability to effectively pass the high-frequency components of signals through the network [38]. As a solution, a fixed coordinate transformation has been proposed prior the training [36], commonly referred to as positional embedding, which embeds high-frequency content into the input coordinates of an INR. While positional embeddings can enhance representation, [27] found that they suffer from limited representational capacity and struggle to generalize effectively. To mitigate these issues, they introduced sinusoidal activations, with a specific frequency and a carefully designed MLP weight initialization, bypassing the need for positional embeddings. Nevertheless, the reliance on exact weight initialization and frequency tuning for sinusoidal activations presents a significant limitation, despite their generalization strength. The reliance on a very specific weight initialization has been minimized after the introduction of Gaussians and Gabor Wavelets as activation functions, leveraging their strong space-frequency localization [24, 25]. However, these non-linear activations still require specific parameter tuning for each signal and INR application, often necessitating exhaustive grid searches [25]. In practice, a common workaround is to reuse previously identified activation parameters when converting explicit signals to implicit representations, but this limits the ability of INRs to adapt and understand the unique characteristics of each signal, ultimately hindering performance. Moreover, to the best of our knowledge, prior research has focused only on using a single activation function to improve INR capabilities. This leads to the following questions: 1). Could using multiple activation functions adaptively enhance both the representation power and generalization of INRs?, 2). How can we ease the burden of the time-consuming and extensive grid search process to determine activation parameters even before training any INR?

To address these existing issues within INRs, we present "Matched Implicit Neural Representations" (*MIRE*), a novel approach that designs an INR based on the signal that is fed to it. To achieve this, we propose a method that draws upon principles from dictionary learning-based methods, dynamically adjusting the activation function for each layer and resulting a matched INR network for the given signal. To dynamically adjust, and find the matching activation, we utilize a method that closely aligns with matching pursuit-algorithm [19]. For the experiments, we utilize a dictionary that includes four new activation "atoms"—Raised Cosines [2], Root Raised Cosines [12], Prolate Spheroidal Wave Functions [15, 16, 29–31], and Sinc functions [26]—chosen for their strong space-frequency localization, a feature commonly leveraged in signal and image processing [5, 22, 33]. In addition, we incorporate three widely used activations from the INR literature: Sinusoids [27], Gabor Wavelets [9, 25], and Gaussians [24]. For performance evaluations of *MIRE*, we present several applications, including image representation, image inpainting, super-resolution, occupancy field representation, novel view synthesis, edge detection, and high-frequency encoding capabilities. Our thorough evaluation of *MIRE* shows that it surpasses state-of-the-art INR solutions by a clear margin. Furthermore, our comprehensive ablations reveal that *MIRE* helps to relax the need for specific activation function parameter determination, previously deemed necessary prior to training any INR.

2. Related works

Activation Functions Neural network activation functions, also referred to as transfer functions [3], determine the output of each neuron based on the weighted sum of the inputs they receive from the previous layer. These functions are typically non-linear and aid neural networks in capturing non-trivial functional relationships with a reduced number of nodes [35]. Unlike network's weights and biases, which are updated based on training data, activation functions are typically chosen beforehand and remain unchanged throughout the training process [17]. However, data-dependent activations, i.e., trainable activations, were recently proposed using the classical sigmoid function [3]. Since then several other trainable activations have been proposed [7, 39]. Several studies have also explored the connections between deep neural networks and activation functions from a frequency perspective [4, 37], providing additional insight to understand their behavior and impact on neural network dynamics.

Compactly Supported and Band-limited Signals have the property of having non-zero values only within a bounded set of their domain or Fourier transform domain, respectively. This property is often desirable in signal processing, communications, and other fields as it comes with efficient approximation, transmission, and recovery properties [23, 28]. Since it is mathematically impossible to have both compact support and band-limitedness at the same time, the next best property is to have some form of *space-frequency concentration*, that is, compact support with rapid frequency decay or band-limited with rapid space decay or rapid decay in both domains. With the advances in INRs, it has been demonstrated that when an activation function has good space-frequency concentration, it not only significantly enhances INR performance but also tends to reduce the need for specific INR weight initialization [24, 25].

Implicit Neural Representations (INRs) have recently garnered attention from the computer vision research community, mainly due to their simplistic network architecture and the performance improvements observed in various vi-

sion tasks compared to traditional parameter-heavy vision models [6, 25, 27]. The emergence of INRs has begun mainly after the introduction of neural radiation fields with ReLU activations [20], which has led to multiple follow-up studies [10, 21], and the use of Sinusoidal activation as an alternative to conventional ReLU activations [27]. Thereafter, [24] has shown the existence of a broader class of activations that are suitable for INRs. A more recent work, [25] has proposed Gabor Wavelets, which are not only compactly supported but also benefit from exponential damping, as a non-linearity for INRs, and showed improved INR performance compared to previous INR models.

3. Methodology

3.1. Formulation of an INR

Consider an INR denoted as F_θ, where θ represents the neural network parameters. F_θ takes coordinates from a K-dimensional space, denoted as $\mathbb{R}^K$, and maps them to a M-dimensional signal, denoted as $\mathbb{R}^M$. Therefore, this mapping can be expressed as:

$$F_\theta : \mathbb{R}^K \rightarrow \mathbb{R}^M.$$

If $W^{(i)}$ and $b^{(i)}$ are the weight and bias matrices of the i^{th} layer, the input to the $(i+1)^{\text{th}}$ layer is given by $\sigma^{(i)}(W^{(i)}x^{(i)} + b^{(i)})$, where $\sigma^{(i)}$, and $x^{(i)}$ represent the activation function, and the input to the i^{th} layer respectively. The representation capacity or the learning dynamics of an INR is governed by the activation function σ [24, 25, 27, 36]. Most studies have used a single activation type for the entire network, i.e., $\sigma^{(i)} = \sigma$ for all i. Through extensive experiments, we show that this approach often leads to suboptimal learning outcomes for INRs since constraining the INR to a single activation function throughout the network potentially limits the representation power of the learned model. Consequently, the model struggles when attempting to generalize to unseen or untrained coordinates, undermining the intended purpose and functionality of INRs. Therefore, this limits the INR's adaptability and effectiveness, whereas robustness and generalization are essential aspects when moving from one representation to another.

3.2. Dictionary of activations for *MIRE*

We adopt a dictionary comprising of seven activation functions. This includes two functions with rapid decay in both the spatial and Fourier domains (Complex Gabor Wavelets and Gaussian) and five band-limited functions (Sinc, Raised Cosine, Root Raised Cosine, Sinusoid, and PSWF).

1. **Sinc Function:** The sinc function is the Fourier transform of a rectangular pulse in the Fourier domain (in digital communication literature, this is also referred to as the Nyquist pulse [26]). It is defined as, $\text{sinc}(\alpha x) = \frac{\sin(\alpha x)}{\alpha x}$, and it decays as $\frac{1}{\alpha x}$, where α is a parameter.
2. **Raised Cosine (RC):** It is another band-limited function whose decay in the space domain is of order x^{-2}, hence faster than the Sinc. Defined by the parameters α, β, and γ, the general functional form of a raised cosine is given by: $\frac{\text{sinc}(\alpha x)\cos(\beta x)}{1-|\gamma|x^2}$.
3. **Root Raised Cosine (RRC):** This is a modified version of the raised cosine obtained by taking the square root of the frequency response of the raised cosine pulse. This modification improves the decay of the signal. With α, β, γ, a, and b as the parameters, the general functional form of a root-raised cosine is given by, $\frac{a\sin(\alpha x)+b\cos(\beta x)}{1-|\gamma|x^2}$.
4. **Prolate Spheroidal Wave Function (PSWF):** These are the solutions to the Helmholtz equation in prolate spheroidal coordinates. The Helmholtz equation in prolate spheroidal coordinates can be transformed to the following ordinary differential equation, where m, n, and c are parameters, and $R_{mn}(c,x)$ are the PSWFs: $(x^2-1)\frac{d^2R_{mn}(c,x)}{dx^2} + 2x\frac{dR_{mn}(c,x)}{dx} - \beta R_{mn}(c,x) = 0$ where $\beta = \left[\lambda_{mn}(c) - c^2x^2 + \frac{m^2}{x^2-1}\right]$. This differential equation arises in the context of bandlimited signals when the signal that has the highest possible energy concentration within a given interval [11]. Although finding a closed-form solution is difficult, a discretized approximation for PSWFs is taken in this study. To define it as an activation atom, the natural cubic spline approximation has been used.
5. **Gabor Wavelet:** Gabor Wavelets involve a Gaussian-modulated cosine or sine wave. They offer rapid decay in both spatial and frequency domains, and have already been employed for INRs showing better performance compared to sinusoidal activations. With α, γ as the parameters, the general functional form of a complex Gabor Wavelet is given by $e^{j\alpha x - |\gamma|x^2}$.
6. **Gaussian:** Similar to Complex Gabor Wavelet, Gaussian functions also offer rapid decay in both spatial and frequency domains, and have already been utilized in INRs. The general functional form of a Gaussian is given by, $e^{-|\gamma|x^2}$ where γ is a parameter.
7. **Sinusoid:** Sinusoidal functions are band-limited, and have been used in INRs. The general functional form of a sinusoid is given by $\sin(\alpha x + \beta)$, where α, β are parameters.

The distinct spatial characteristics of each activation function can be observed from the provided figure in supplementary material, where the variation of the activation function value with spatial distance is depicted.

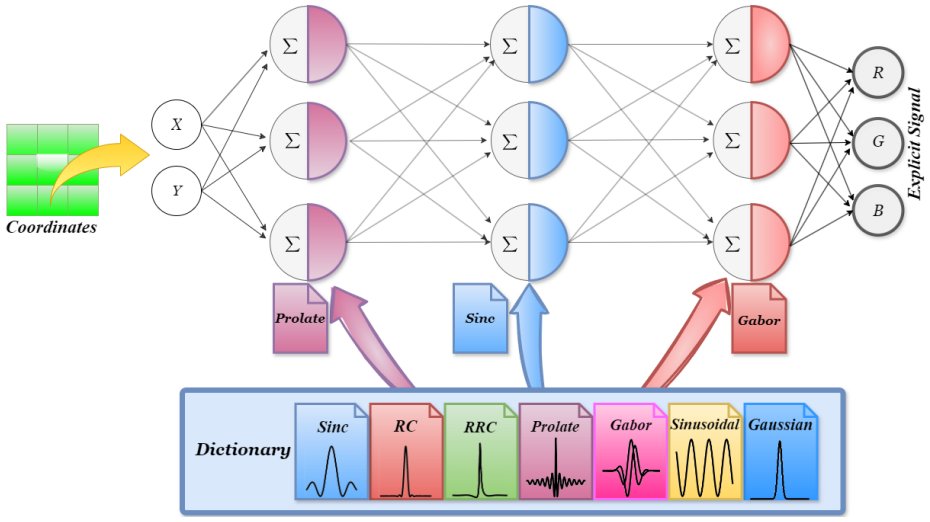

Figure 1. **Illustration of *MIRE*.**: Existing INRs lack the capacity to fully adapt to a given signal and capture its intrinsic characteristics, as they often rely on predetermined activation function parameters that are not specifically tailored to the signal being converted to an implicit representation. To dynamically create an INR that truly understands the given explicit signal, MIRE leverages a dictionary of activation atoms to identify the most suitable activation sequence for each signal.

3.3. Premise of *MIRE*

MIRE begins by constructing a dictionary of predefined activation functions, as described in Sec. 3.2, each equipped with trainable parameters. These activation atoms have their parameters randomly initialized, drawn either from a uniform or normal distribution. *MIRE* is then initialized as a single-layer MLP, with input and output dimensions customized to match the explicit signal representation. The algorithm iterates through each activation function in the dictionary, and applies it as the non-linearity for the hidden layer over a specified number of training epochs, where it allows the activation function parameters to be optimized. During this process, the algorithm tracks the performance of each activation function by calculating the mean square loss between two representations, saving the parameters of the activation that achieves the lowest loss. After all activation atoms have been evaluated, the algorithm selects the activation function that produces the minimum loss for the first layer.

Once the most suitable activation function for the first-hidden layer is identified based on the minimum loss criterion, it is fixed as the non-linearity for the first-hidden layer, along with its associated parameters. *MIRE* then proceeds to add a second-hidden layer. It again starts the MLP training process afresh, with one key difference: the first-hidden layer's activation function and its optimized parameters, which minimized the loss when using a single hidden layer, are retained. The algorithm then tests each activation function from the dictionary as the non-linearity for the second-hidden layer, again over a predetermined number of epochs. Similarly, the performance of each activation is recorded. At the end of this sweep, based on the performance, the activation that gives the minimum loss is selected, and fixed as the second layer's non-linearity. This process continues for all the hidden layers. Figure 1 illustrates the process of selecting an optimal activation sequence for image representation tasks by *MIRE* through the activation function dictionary. For a better understanding of the *MIRE*'s training process, please refer the pseudo-code provided (See the supplementary material) along with Sec. 3.3.

3.4. Activation function parameter initialization

The performance of an INR with parametric activation functions heavily depends on activation parameter initialization, where poor initialization often degrades performance across tasks. Previous studies have used extensive grid searches to determine the activation parameters [25] for each application, but the effectiveness of this method is highly dependent on the diversity of the signal and may perform poorly with signals that differ from those used during parameter determination. In contrast, *MIRE* either randomly initializes the activation function parameters or initializes them with their base configurations, allowing the network to learn the optimal parameters for each application and signal during optimization. This adaptability enables *MIRE* to optimize based on the specific characteristics of the signal. Experimental results demonstrate that *MIRE* outperforms existing INRs, including WIRE [25], SIREN [27], GAUSS [24], and MFN [9], in both performance and generalization in various tasks.

4. Experimental results

4.1. Image representation

As mentioned earlier, a direct application of *MIRE* is learning an implicit representation of an image, commonly referred to as image representation. To clearly illustrate how *MIRE* functions, we select an image that exhibits high spatial variation and has a broad frequency range (shown on the left of the top row in Fig. 2). Here, the network is provided with the normalized coordinates of the signal, without any positional embedding, and *MIRE* is trained to predict the corresponding RGB values. Then, for a more comprehensive evaluation, the representation capacity of *MIRE* is evaluated across the Kodak [14] data set. The resulting PSNR for each image, along with the baseline results, is shown in Fig. 3. For the average PSNR and the decoded representations of each method, refer the supplementary material.

As shown in the top row of Fig. 2, *MIRE* achieves the highest PSNR and SSIM values, indicating the lowest distortion and best preservation of structural information, texture, and contrast compared to existing INRs. For this experiment, we used all the activations defined in Sec. 3.2. As detailed in Sec. 3.3, *MIRE* begins with a single hidden layer and searches the dictionary to determine which activation produces the highest PSNR (or the lowest loss). This process is carried out for 100 epochs for each activation atom in the dictionary. Upon identifying the activation that is mostly matched to the image according to the loss criterion, Sinc in this case, it locks this activation for the first hidden layer (bottom left, Fig. 2). Subsequently, *MIRE* adds the second hidden layer and resumes training the entire network afresh while keeping the first layer's identified best activation frozen, adjusting only the activation of the second-hidden layer at every 100 epochs. At the end of this phase, it determines the activation that provides the highest PSNR for the second hidden layer, which in this instance is Gaussian (bottom middle, Fig. 2). Following this, *MIRE* introduces the third hidden layer and begins training the network afresh now while keeping both first and second layers' activations frozen, this time modifying the activation of the third layer at every 200 epochs. Upon completion of training, *MIRE* identifies the activation for the third hidden layer that results in the highest PSNR, which, for this stage, is the Gabor Wavelet (Bottom right, Fig. 2). Therefore, the matched sequence of activation atoms for the Parrot image, as identified by *MIRE*, is Sinc, Gaussian, and Gabor Wavelet. Note that although *MIRE* determines the most suitable activation for each layer based on loss, PSNR plots are used for illustrative purposes.

Considering the bottom row of Fig. 2, we can conclude that, once the network has the matched sequence of activation functions determined by *MIRE*, INRs start showcasing a faster convergence. In the case of *MIRE*, when the first two layers' activation functions are determined, it only needs at most 200 epochs to obtain a minimum loss between implicit and explicit representations. Therefore, showing a much faster convergence rate compared to the current state-of-the-art INRs. As can be evidenced from both Fig. 2 and Fig. 3, an INR achieves the highest accuracy metrics when its activation functions are customized for a specific signal, rather than using a pre-optimized, uniform activation sequence throughout the INR. This thorough evaluation confirms our hypothesis that designing an INR that is matched to the signal that is fed to it significantly enhances the INR's performance, even when the activation parameters are randomly initialized.

4.2. Image inpainting

Unlike explicit discretized signal representations, an INR learns a continuous implicit representation of a given signal through the MLP training process. Therefore, once the corresponding explicit representation is encoded into the weights and biases of an INR, one should be able to query the model as desired. The inpainting task serves as a good measure of INRs to assess whether the model is overfitted, as the primary purpose of adopting a new representation is to generalize it through learned continuous mapping. To demonstrate the functionality of *MIRE* for inpainting, we selected an image with intricate details, shown in the top left of Fig. 5. The adjacent image on the right shows the same image with a text mask applied. Additionally, we evaluated *MIRE*'s inpainting performance on the Kodak dataset to provide a more comprehensive assessment. The PSNR results for each image, along with baseline comparisons, are shown in Fig. 4. For the inpainted images, and average performance metrics on image inpainting on the Kodak dataset please refer the supplentary matrial.

For this experiment, the newly introduced activations i.e., RC, RRC, Sinc, PSWF, and Gabor Wavelet have been used to showcase the effectiveness of these activations. The bottom row of Fig. 5 showcases the PSNR performance observed in each layer when following the procedure in Sec. 3.3. It should be noted that in the case of image inpainting, the loss calculation for deciding the activation is based on the partial image data. The results clearly demonstrate that *MIRE* delivers the cleanest and most visually coherent image inpainting outcomes compared to all existing INRs. Beyond producing the most visually coherent images, *MIRE* also achieves the highest PSNR and SSIM values for the inpainting tasks.

4.3. Occupancy field representation

As INRs offer a continuous functional mapping from low-dimensional coordinate space to signal space, they can be used to effectively represent three-dimensional signed distance fields. In this scenario, the mapping extends from the

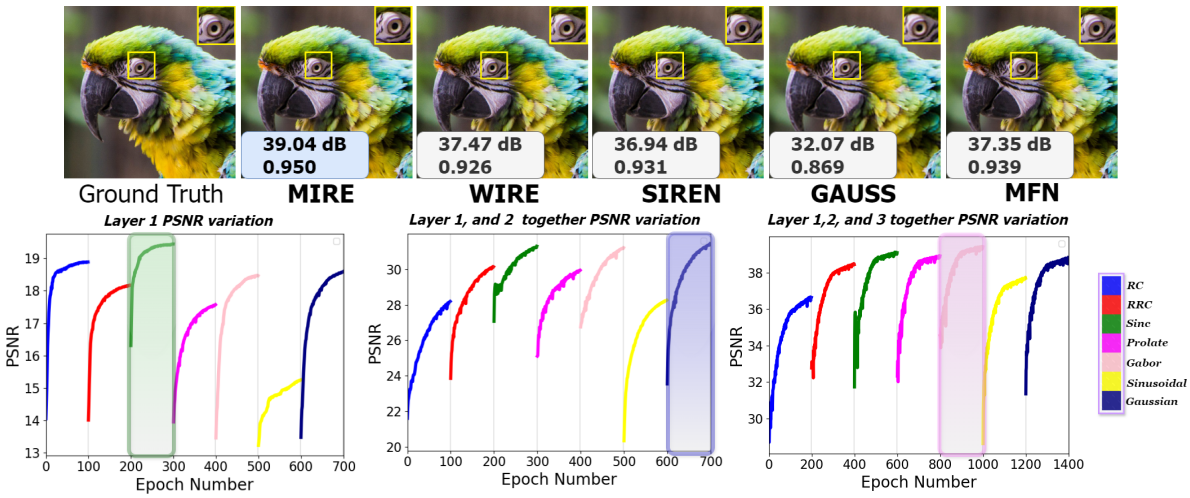

Figure 2. **Image representation capacity of *MIRE***: The top row depicts the image reconstruction using various types of INRs. *MIRE* stands out as the INR that achieves the highest PSNR and SSIM metrics, indicating minimal distortion and maximum preservation of structural information. The bottom row illustrates how *MIRE* achieves these results through sequential training. The designed INR for this signal by *MIRE* has the following activation sequence: Sinc, Gaussian, and Gabor Wavelet under randomly initialized activation parameters.

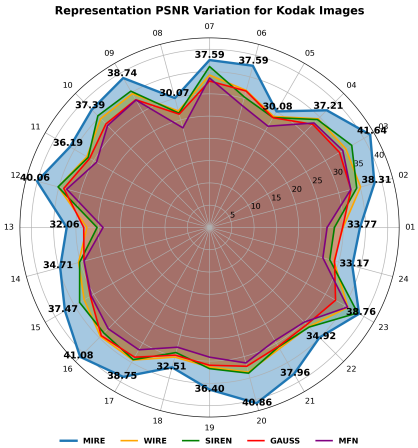

Figure 3. **Image representation capabilities of *MIRE* on the Kodak dataset**

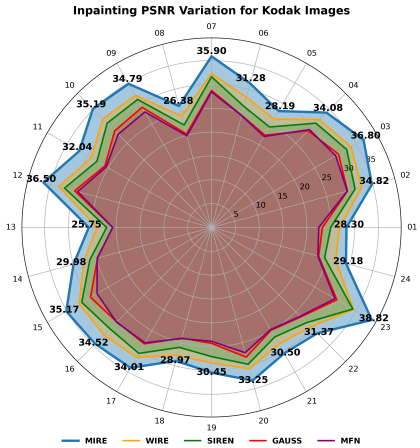

Figure 4. **Image inpainting capabilities of *MIRE* on the Kodak dataset**

three-dimensional space to a one-dimensional space, where the signal space is represented by binary values: either 1 or 0. Here, 1 denotes that the signal lies within the specified region, while 0 indicates its absence in the given region. For this experiment, two datasets, Thai Statue and Stanford Lucy, were obtained from Stanford 3D datasets [32]. The sampling procedure followed the method described in [25], using a $512 \times 512 \times 512$ grid. Voxels inside the volume were assigned a value of 1, while those outside the volume were assigned a value of 0. A part of sampled volumes for Stanford Lucy and the Thai Statue are displayed in the first column of the 1st and 2nd rows, respectively, in Fig. 6.

Figure 6 shows the decoded representations for each INR along with the ground truth. As can be clearly seen *MIRE* achieves the highest Intersection over Union (IoU) metric, demonstrating the greatest representation capacity among all existing INRs regardless of the occupancy field. A closer examination of decoded statues reveal that *MIRE* precisely encodes intricate high-frequency details. In contrast, INRs like WIRE[1] and SIREN tend to converge toward a low-pass representation, highlighting the challenge of encoding rapidly varying, detailed features in these models. These findings clearly indicate that not only for images but for any signal, when the matched sequence of activations is identified, an INR can accurately learn the implicit representation. In these experiments, *MIRE* determined the matched sequence of activations for Stanford Lucy as Sinc, RRC, and PSWFs for the first, second, and third layers, respectively. For the Thai Statue, the activations were RC, RRC,

[1]* Reproduced result with 300 hidden neurons

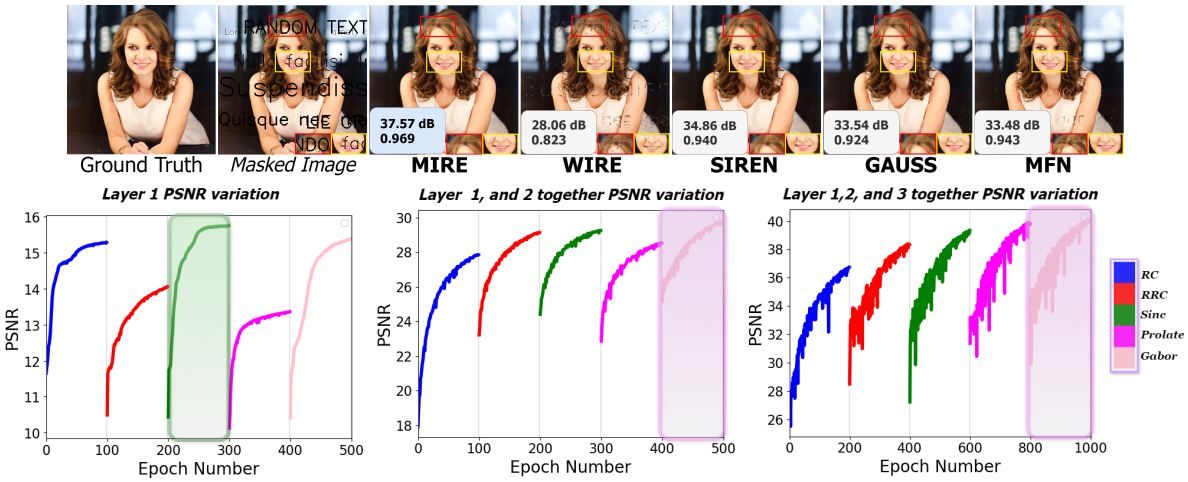

Figure 5. **Image inpainting capabilities of *MIRE*.** The top row displays the recovered images using various types of INRs. *MIRE* stands out as the INR that not only achieves the highest PSNR and SSIM metrics but also yields the most visually coherent inpainting outcome. The higher metrics indicate that *MIRE* has restored the image with minimal distortion and maximum preservation of structural information. The bottom row illustrates how *MIRE* achieves this result through sequential training. By tailoring activations to the specific image, the corresponding sequence consists of Sinc, Gabor Wavelet, and Gabor Wavelet.

and Gabor Wavelet for the respective layers. The complete occupancy fields corresponding to Fig. 6 is shown in Occupancy Fields section in supplementary material.

4.4. Neural radiance fields

INRs have gained popularity in the computer vision community, largely due to the impact of NeRFs [20]. In which, a 3 dimensional scene is encoded in an INR by inputting the viewer's spatial coordinates (x, y, z) and viewing angles (θ, ϕ) into the network with the aid of collection of images captured around the scene. The INR is tasked with predicting the color and density at those locations. When the INR is trained, the INR can generate unseen perspectives from new spatial positions and viewing angles which are not present in the training data. For this experiment, we utilized a vanilla NeRF architecture with Chair, and Hotdog datasets. The top and bottom rows of Fig. 7 show novel views generated from the trained INR models on the Chair and Hotdog datasets, respectively. Additional novel views are provided in the supplementary material.

4.5. Effect of activation parameter initialization

As outlined in Sec. 3.4, the performance of conventional INRs heavily depends on the initialization of activation function parameters. In contrast, *MIRE* designs an INR through matching the activation sequence for a specific task without needing a much precise initialization. To substantiate this claim, we have sourced activation function parameters from both uniform and normal distributions. The primary reason for selecting these distributions is to understand how INRs perform when parameters are derived from distributions that are either spread evenly across a range or centered around a mean value. A uniform distribution over $[a, b]$ is denoted as $U(a, b)$, and a normal distribution with mean μ and standard deviation σ as $\mathcal{N}(\mu, \sigma)$. The results in Tab. 3 in supplentary show the average PSNR (in dB) from five trials on the Parrot image in Fig. 2, with variability expressed as the standard deviation next to the $\pm$ symbol. The bold number indicates the highest PSNR, and the following number represents the lowest standard deviation. As illustrated in Tab. 3 in supplementary, *MIRE* emerges as the only INR which is capable of delivering consistent PSNR across various distributions while exhibiting minimal variation around the mean. *MIRE* not only maintains PSNR consistency but also records the highest PSNR values. In contrast, WIRE demonstrates commendable performance exclusively under the $U(-10, 10)$ distribution, suggesting its activation parameters require initialization within a narrow range (-10 to 10) for the tested parrot image. Similarly, SIREN shows enhanced performance when its activation parameters are selected from $U(-100, 100)$ distribution. These observations underline the dependency of existing INRs on specific initial conditions for their activation parameters to guide the network towards convergence.

4.6. Additional experiments and ablation Studies

Comprehensive experiments on image super-resolution, edge detection, and high-frequency encoding are presented in supplementary material, where we also include comparisons with recently released state-of-the-art INR methods, namely INCODE [13] and FINER [18]. Details of the experimental setup and ablation studies are provided, evalu-

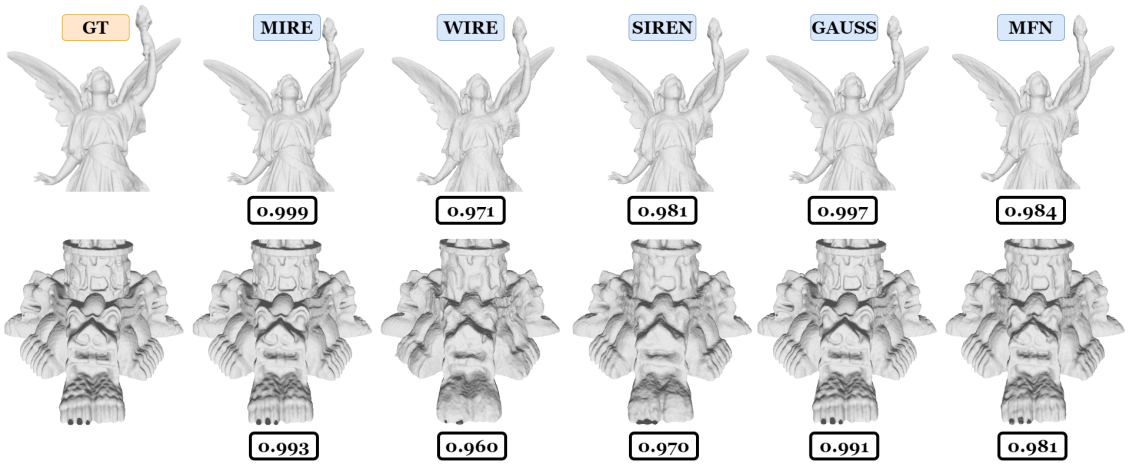

Figure 6. **Occupancy field representation capacity of *MIRE***: The image illustrates the reconstruction capabilities of various INRs for occupancy volumes. *MIRE* stands out as the INR that not only achieves the highest IoU metric but also the INR which preserves the highest amount of fine details in its weights and biases.

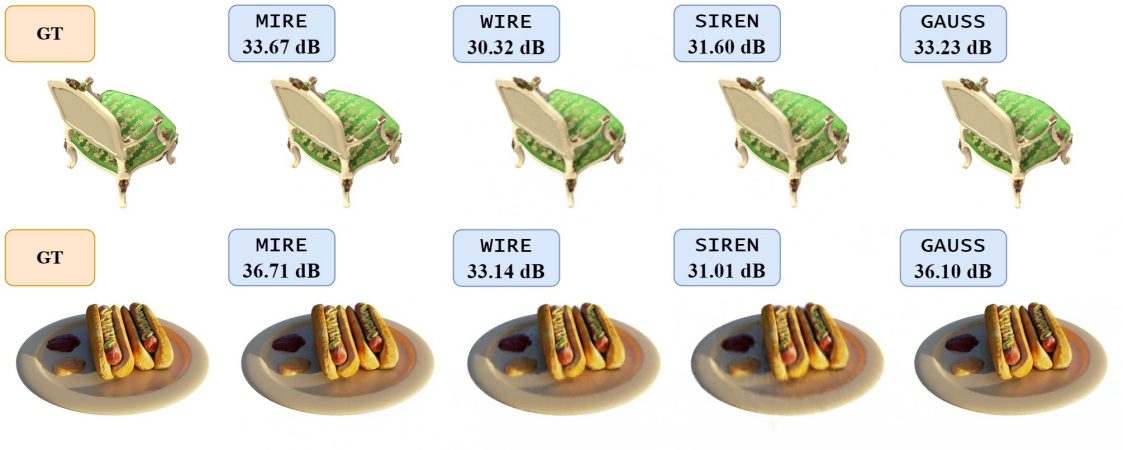

Figure 7. ***MIRE*'s novel view synthesis capabilities**: *MIRE* consistently achieves the highest performance metrics and captures more intricate details than the baselines. For the chair dataset, *MIRE* accurately produced the fine textures, carvings, and lighting effects, closely matching the ground truth. Similarly, for the hotdog dataset, *MIRE* preserves the texture, shadows, and reflections with greater fidelity.

ating MIRE's performance with respect to hidden neurons, layers, learning rates, weight initialization, and positional encoding. Additionally, we provide explanations of training curves, variations in activations within the spatial domain, and how *MIRE* differs from baseline methods.

5. Conclusion

To learn a continuous functional representation of a digital signal using an Implicit Neural Representation (INR), one typically needs to conduct an exhaustive grid search to optimize activation function parameters. However, the lack of insight into the parameter space complicates finding an optimal configuration, leading to the common practice of reusing previously discovered parameters. This workaround restricts the INR's ability to fully adapt to each unique signal, ultimately limiting its expressive power and generalizability. In this work, we introduce *Matched Implicit Neural Representations (MIRE)*, a dynamic method that adapts INRs to each input signal using a dictionary learning approach, where the dictionary consists of non-linear activation atoms. By dynamically aligning activation functions to signal characteristics, *MIRE* unlocks greater potential for generalization across diverse data types. Our method, demonstrated with seven distinct non-linearities (extendable or customizable), significantly enhances INR performance across various tasks in representation and generalization for all data modalities. Additionally, it relaxes the constraints of grid searches and parameter reuse, providing a more efficient and adaptive solution.

6. Acknowledgment

This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number 140D0423C0076. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

[1] Awesome-Super-Resolution/dataset.md at master · ChaofWang/Awesome-Super-Resolution — github.com. https://github.com/ChaofWang/Awesome-Super-Resolution/blob/master/dataset.md. [Accessed 01-10-2024]. 1

[2] Nader Sheikholeslami Alagha and Peter Kabal. Generalized raised-cosine filters. *IEEE transactions on Communications*, 47(7):989–997, 1999. 2

[3] Andrea Apicella, Francesco Donnarumma, Francesco Isgrò, and Roberto Prevete. A survey on modern trainable activation functions. *Neural Networks*, 138:14–32, 2021. 2

[4] Nuri Benbarka, Timon Höfer, Andreas Zell, et al. Seeing implicit neural representations as fourier series. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2041–2050, 2022. 2

[5] Boualem Boashash. *Time-frequency signal analysis and processing: a comprehensive reference*. Academic press, 2015. 2

[6] Pablo Cervantes, Yusuke Sekikawa, Ikuro Sato, and Koichi Shinoda. Implicit neural representations for variable length human motion generation. In *European Conference on Computer Vision*, pages 356–372. Springer, 2022. 3

[7] Shiv Ram Dubey, Satish Kumar Singh, and Bidyut Baran Chaudhuri. Activation functions in deep learning: A comprehensive survey and benchmark. *Neurocomputing*, 2022. 2

[8] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021. 1

[9] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2020. 2, 4

[10] Kyle Gao, Yina Gao, Hongjie He, Dening Lu, Linlin Xu, and Jonathan Li. Nerf: Neural radiance field in 3d vision, a comprehensive review. *arXiv preprint arXiv:2210.00379*, 2022. 3

[11] Maria C Gonzalez. *Engineering Applications of Prolate Spheroidal Wave Functions and Sequences and Legendre Polynomials: Filtering and Beamforming in 1D and 2D*. University of California, Davis, 2018. 3

[12] Michael Joost. Theory of root-raised cosine filter. *Research and Development*, 47829, 2010. 2

[13] Amirhossein Kazerouni, Reza Azad, Alireza Hosseini, Dorit Merhof, and Ulas Bagci. Incode: Implicit neural conditioning with prior knowledge embeddings. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1298–1307, 2024. 7

[14] Kodak. Kodak lossless true color image suite. https://r0k.us/graphics/kodak/. Accessed: October 1, 2024. 5

[15] Henry J Landau and Henry O Pollak. Prolate spheroidal wave functions, fourier analysis and uncertainty—ii. *Bell System Technical Journal*, 40(1):65–84, 1961. 2

[16] Henry J Landau and Henry O Pollak. Prolate spheroidal wave functions, fourier analysis and uncertainty—iii: the dimension of the space of essentially time-and band-limited signals. *Bell System Technical Journal*, 41(4):1295–1336, 1962. 2

[17] Johannes Lederer. Activation functions in artificial neural networks: A systematic overview. *arXiv preprint arXiv:2101.09957*, 2021. 2

[18] Zhen Liu, Hao Zhu, Qi Zhang, Jingde Fu, Weibing Deng, Zhan Ma, Yanwen Guo, and Xun Cao. Finer: Flexible spectral-bias tuning in implicit neural representation by variable-periodic activation functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2713–2722, 2024. 7

[19] Stéphane G Mallat and Zhifeng Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing*, 41(12):3397–3415, 1993. 2

[20] 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. *Communications of the ACM*, 65(1):99–106, 2021. 3, 7

[21] Amirali Molaei, Amirhossein Aminimehr, Armin Tavakoli, Amirhossein Kazerouni, Bobby Azad, Reza Azad, and Dorit Merhof. Implicit neural representation in medical imaging: A comparative survey. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2381–2391, 2023. 3

[22] Arogyaswami J Paulraj and Constantinos B Papadias. Space-time processing for wireless communications. *IEEE signal processing magazine*, 14(6):49–83, 1997. 2

[23] John G Proakis. *Digital communications*. McGraw-Hill, Higher Education, 2008. 2

[24] 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, 4, 8

[25] 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, 4, 6, 8

[26] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. 2, 3

[27] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. 2, 3, 4

[28] Bernard Sklar. *Digital communications: fundamentals and applications*. Pearson, 2021. 2

[29] David Slepian. Prolate spheroidal wave functions, fourier analysis and uncertainty—iv: extensions to many dimensions; generalized prolate spheroidal functions. *Bell System Technical Journal*, 43(6):3009–3057, 1964. 2

[30] David Slepian. Prolate spheroidal wave functions, fourier analysis, and uncertainty—v: The discrete case. *Bell System Technical Journal*, 57(5):1371–1430, 1978.

[31] David Slepian and Henry O Pollak. Prolate spheroidal wave functions, fourier analysis and uncertainty—i. *Bell System Technical Journal*, 40(1):43–63, 1961. 2

[32] Stanford University Computer Graphics Laboratory. The stanford 3d scanning repository. https://graphics.stanford.edu/data/3Dscanrep/. Accessed: 2024-09-30. 6

[33] LJubisa Stankovic, Srdjan Stankovic, and Igor Djurovic. Space/spatial-frequency analysis based filtering. *IEEE Transactions on Signal Processing*, 48(8):2343–2352, 2000. 2

[34] Yannick Strümpler, 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. 1

[35] Tomasz Szandała. Review and comparison of commonly used activation functions for deep neural networks. *Bio-inspired neurocomputing*, pages 203–224, 2021. 2

[36] 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, 3

[37] Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019. 2

[38] 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

[39] Brosnan Yuen, Minh Tu Hoang, Xiaodai Dong, and Tao Lu. Universal activation function for machine learning. *Scientific reports*, 11(1):18757, 2021. 2

[40] Djemel Ziou, Salvatore Tabbone, et al. Edge detection techniques-an overview. *Pattern Recognition and Image Analysis C/C of Raspoznavaniye Obrazov I Analiz Izobrazhenii*, 8:537–559, 1998. 1