MIRE: Matched Implicit Neural Representations

Supplementary Material

7. Additional experiments

7.1. Image super resolution

As the learning process of an INR enables the discovery of a continuous functional mapping from low-dimensional spatial coordinates to high-dimensional signal space, one gains the ability to query the learned representation as needed. Consequently, unlike grid-based image representations, this representation becomes decoupled from the spatial resolution of the image, allowing for image super-resolution through interpolation of the learned representation.

For the purpose of demonstrating image superresolution, we selected two images with different levels of detail complexity. The first, referred to as the "boy image" from the Set 14 Dataset [1], is a dynamic image that contains sharp frequency contents throughout the image, which could potentially challenge the detail retrieval capabilities of super-resolution. The second image, a woman, where sharp frequency details are only concentrated primarily around facial features, offering a localized test for the super-resolution process (See Fig. 8). Both images were downsampled by a factor of two and then used to train the INR. The trained INR, unconstrained by spatial resolution, was then used to infer high-resolution images.

The performance of image super-resolution demonstrates the generalizability of an INR. As evident from the results, MIRE consistently exhibits not only the best PSNR and SSIM values across both cases but also the most visually coherent result, irrespective of the complexity of the image. Among the INRs, WIRE shows a notable deficiency in generalizing to a large set of unseen coordinates during testing, particularly evident in complex images. For example, in the case of the "boy image," WIRE struggles to generalize the down-scaled image's implicit mapping. When attempting super-resolution, WIRE introduces a significant number of random values for unseen coordinates, highlighting its limitations in handling images with varying complexity levels. Conversely, in the "women image," WIRE demonstrates a more reasonable representation. A simillar situation can be seen from both GAUSS and MFN. Therefore, indicating a dependency of WIRE, GAUSS, and MFN's super-resolution capabilities on the detailedness of an image. These methods tend to excel with images where details are localized, such as the "women image," but fail with highly detailed images like the "boy image".

A closer examination of SIREN's super-resolution images reveals a learned low-passed representation, missing high-frequency components in the decoded high-resolution image. In contrast, *MIRE* exhibits unparalleled performance in all scenarios tested, confirming the robustness and reliability of its learned representation. The performance of *MIRE* is consistent across different types of images. These findings further corroborate our hypothesis: an INR equipped with an appropriately sequenced activation framework learns a superior representation than a fixed sequence of activations that depends on the given signal.

7.2. Edge detection

INRs are characterized by their inherent ability to encode signals in an implicit manner, where a key attribute of a generalized INR lies in its ability to undertake tasks which are typically reserved for explicit representations. With the explict signal representations, especially with images, the spatial filters like Sobel, Gaussian, Canny, etc., enable the extraction of edges. This process is essential for activities like identifying the boundaries of objects, detecting text, or recognizing facial features, all of which rely significantly on precise edge detection [40]. Even though it's possible to obtain edge maps from an INR through its decoded outputs, INRs provide the benefit of representations that are learned and differentiable. Therefore, an INR must not only precisely capture the signal but also facilitate the retrieval of its gradient data. For an INR to efficiently function as an edge detector, an INR requires to encode pixel-level relationships within its weights and biases.

The capabilities of extracting edge maps using INRs become evident when applied to an image which has clear visual edges, like the Monarch picture from the Set 5 Dataset [1]. To identify the possibility of edge detection capabilities of INRs, initially, the image representation has been learned through training the INR. Thereafter, the gradient operator is utilized between this learned model and the coordinates used during training. Figure 9 illustrates the RGB image of the Monarch alongside its edge maps derived from INRs.

MIRE stands out by providing the cleanest and most well-defined edge map, identifying necessary edges with minimal false edge identifications. Among other INRs, while MFN can identify edges, its edge map often contains excessive texture content, potentially leading to false edge detection. GAUSS on the other hand, produced smooth edge maps, which can diminish the true edge signal and result in a less reliable representation. In contrast, *MIRE* preserves intricate patterns, offering a more accurate portrayal of the butterfly's natural markings. SIREN also demonstrates its effectiveness as an edge detector, leveraging its built-in capabilities for identifying edges. However, even in the absence of such inherent advantages, *MIRE* excels in detecting edges, highlighting its proficiency in capturing com-



Figure 8. **Image super resolution capabilities of** *MIRE*. The top row showcases super resolution results for the "boy image," where it contains a rapidly varing dynamic content throughout the image, while the bottom row depicts a woman's picture with localized frequencies around the face region. *MIRE* emerges as the sole INR consistently delivering the highest PSNR and SSIM scores for both cases, showcasing its superiority in learned implicit neural representation.

plex patterns and delivering precise representations. Therefore, *MIRE* showcases its effectiveness as a reliable edge detector compared to existing state-of-the-art INRs.

7.3. High frequency encoding capabilities

A straightforward method to evaluate the high-frequency encoding capabilities of INRs are by assessing their effectiveness in representing sharp frequencies. This can be achieved by examining a simple image comprised solely of sharp transitions, where the abrupt color changes in the spatial domain correspond to high-frequency components in the frequency domain.

For this experiment, we selected the image displayed on the left side of Fig. 10. INRs were then trained on this image. Once each INR had learned the implicit representation, it was decoded by inputting the corresponding coordinates. Subsequently, we obtained normalized error plots comparing the ground truth with the decoded representations.

As illustrated in Fig. 10, *MIRE* exhibits the smallest error margin between the decoded image and the ground truth, in comparison to other models such as SIREN and WIRE, which demonstrate noticeable errors even in areas of uniform color. This outcome distinctly highlights *MIRE*'s superior capability in accurately encoding high-frequency components. Such precision in representation validates the underlying hypothesis of *MIRE* in frequency domain perspective as well: the sequence of activations within an INR should adapt based on the characteristics of the input signal.

8. Additional results

8.1. Image representation

Unlike traditional INRs, which use the same activation function throughout the network with suboptimal parameters, that may not be ideal for a given signal, *MIRE* offers a framework that allows the network to adapt its internal configuration to better match the signal, and produces an INR which is matched to the given signal. This added flexibility enables *MIRE* to consistently outperform existing INR methods. This is evident from our thorough evaluation of *MIRE* on the Kodak image dataset, as shown in Fig. 3. Furthermore, we have also provided decoded representations of the Kodak dataset, which are presented in Fig. 11. Moreover, the average PSNR and SSIM metrics for *MIRE* and baseline models across the dataset are summarized in Tab. 1.

8.2. Image inpainting

A key difference between traditional signal representation mechanisms and INRs is that INRs attempt to establish a continuous implicit functional relationship between normalized coordinates and signal values. The generalizability of this relationship depends heavily on the type of activation function used. *MIRE*, by exploring its dictionary to find activation atoms that best match the given signal and task, consistently outperforms all INR baselines, regardless of the specific task.

This is clearly demonstrated by the thorough evaluation we conducted on the image inpainting task using the entire Kodak dataset. The PSNR variations across all meth-



Figure 9. MIRE's edge detection capabilities: Edge maps obtained through applying the gradient operator on the learned representation.



Figure 10. Ability to encode high frequencies in INRs. The results present normalized error plots obtained after training each INR on the ground truth image. It is evident that *MIRE* exhibits the lowest error among the INRs when encoding signals into their weights and biases. Notably, it excels in encoding high frequencies exceptionally well compared to the other INRs as *MIRE* is tailored to the given signal. Conversely, WIRE and SIREN demonstrate errors not only in encoding high-frequency content but also in constant color areas.

ods are shown in Fig. 4. As shown, MIRE consistently achieves the highest accuracy metrics for every image in the dataset. Sample inpainting outcomes are displayed in Fig. 12. The first column represents the ground truth image, while the second column shows the text-masked image. The text mask used contains different fonts and sizes, with overlapping parts, and this same mask was applied to the entire dataset. This presents a significant challenge for INRs to recover the original signal. However, as seen, MIRE recovers the original image with the highest accuracy. Further, MIRE does not showcase a tendency towards overfitting to the training data, nor does it tend to produce low-pass signal representations lacking fast-varying components. The average PSNR across the dataset is presented in Tab. 2. An additional image inpainting experiment has been conducted with a different text mask, and the results obtained are shown in Fig. 13. In this experiment, a randomly generated text mask with varying font sizes is used. As evident from the results, MIRE is the only architecture that achieves the highest PSNR and SSIM values, indicating the best recovery compared to existing INRs.

8.3. Occupancy fields

The complete decoded occupancy fields, along with the ground truth, are shown in Fig. 14. *MIRE* consistently de-

livers the occupancy field that is closest to the ground truth. This capability is attributed to the ability of *MIRE* to find the optimal activation sequence for each given signal.

8.4. Neural radiance fields

In addition to the novel views provided in Sec. 4.4, novel views from different viewing angles and positions are provided in Fig. 15, and Fig. 16 for Chair and Hotdog datasets respectively.

9. Explanations

9.1. Pseudo code of MIRE

For enhanced understanding of *MIRE*, we provide its pseudo code in Algorithm 1. It's important to highlight that *MIRE* offers the flexibility to both expand and condense its dictionary. This means elements can be added to or removed from *MIRE*'s dictionary as needed. For example, if we remove all activation atoms except for the Gabor Wavelet, *MIRE* transforms into WIRE, but with layer wise training.



Figure 11. Additional image representation results from the Kodak dataset: As shown by the results, *MIRE* is the architecture that consistently delivers the cleanest image representation. Moreover, it not only excels in producing clean output but also achieves the highest accuracy metrics.

Method	PSNR (dB)	SSIM
MIRE	36.55	0.93
WIRE	32.30	0.90
SIREN	32.19	0.88
GAUSS	31.34	0.84
MFN	30.15	0.85
INCODE	30.53	0.85
FINER	34.32	0.92

Table 1. Average metrics for image representation task across the entire Kodak dataset

Algorithm 1 Pseudo Code of MIRE

- 1: **Input**: Explicit signal representation
- 2: Output: Implicit Neural Representation
- 3: Initialization:
- 4: Dictionary \leftarrow Dictionary of activation functions
- 5: $N \leftarrow$ number of hidden layers
- 6: Model ← MLP (single hidden layer, matching I/O dimensions)
- 7: $\forall j \in \{1, \dots, N\}$, Model.activation_j \leftarrow None
- 8: Minimum $\text{Loss}_j \leftarrow \infty$, Best Activation $_j \leftarrow$ None 9: for $j \in \{1, ..., N\}$ do

```
10:
         for activation function in Dictionary do
              Model.activation<sub>i</sub> \leftarrow activation function
11:
              for epoch \in \{1, \ldots, \text{num\_epochs}\} do
12:
                   Train the Model
13:
                   Find Loss_i
14 \cdot
                   if Loss_i < Minimum Loss_i then
15:
                       Minimum Loss_i \leftarrow Loss_i
16:
17:
                       Best Activation<sub>i</sub> \leftarrow activation function
                   end if
18:
              end for
19:
         end for
20:
         Model.add a hidden layer
21:
22:
         Model.activation<sub>i</sub> \leftarrow Best Activation<sub>i</sub>
23: end for
```

9.2. Spatial domain variation

In Fig. 17, the seven activation atoms used in *MIRE* are depicted, showcasing their spatial domain variations. These activations include Sinc, Raised Cosine, Root Raised Cosine, Prolate Spheroidal Wave Function (PSWF), Gabor Wavelet, Sinusoid, and Gaussian. The diversity of these filters enables the *MIRE* framework to capture a wide range of features, from sharp edges to smooth transitions, periodic patterns, and localized details. This variety ensures the network can represent different structures effectively, making the method highly adaptable.

9.3. How does MIRE differ from baselines?

The *MIRE*'s dictionary consists of seven activation functions, three of which, sinusoidal, Gaussian, and Gabor wavelet, have been previously introduced in other studies. This naturally raises the question of how *MIRE* differs from these baselines. The key distinction lies in the initialization of activation function parameters. While earlier studies used these same activation functions, their parameters (such as α and β in the sinusoidal function, as noted in Sec. 3.2 of the main paper) were typically fine-tuned via exhaustive



Figure 12. Additional image inpainting outcomes for the Kodak dataset: *MIRE*'s performance can be attributed to its ability to dynamically tailor the sequence of activations to the specific signal and task.

INR	PSNR (dB)	SSIM
MIRE	32.34	0.90
WIRE	29.82	0.85
SIREN	28.44	0.79
GAUSS	26.41	0.69
MFN	26.02	0.70
INCODE	30.39	0.87
FINER	30.78	0.89

Table 2. Average metrics for image inpainting task across the entire Kodak dataset

grid searches for each specific INR application. In contrast, *MIRE* initializes the parameters of each activation function either randomly or using their base configurations (for instance, for a sinusoid, $\alpha = 1$ and $\beta = 0$). As demonstrated in Sec. 4.5 of the main paper, when baseline models (which use a single activation function throughout the network) are initialized with these configurations, they often fail to achieve comparable performance. Therefore, unlike these baselines, *MIRE* does not rely on any pre-optimized activation function parameters, offering a more flexible and robust approach.

9.4. Robustness of *MIRE* to random initializations

This section directly relates to Sec. 4.6, which provides a detailed discussion of additional parameters. Furthermore, we provide a mathematical inclined explanation of why *MIRE* is robust to random initializations. The corresponding results are shown in Tab. 3. We explain the robustness as follows:

Larger solution space: Carefully initialization is necessary for conventional INRs such as SIREN, limiting the expressivity of INRs. Choosing an activation function adaptively for each layer can overcome this limitation, as the optimization process will automatically select the most suitable activation function from the available dictionary for each stage, regardless of the initial conditions.

9.5. Total and effective number of epochs and learning curves

Compared to other INRs, MIRE's total number of epochs is determined by the size of the dictionary. As MIRE can accommodate any dictionary size, let us assume the dictionary contains k activations. For simplicity, consider a model with three hidden layers. If each activation is trained for x_1, x_2 , and x_3 epochs in the first, second, and third layers respectively, the total number of search epochs is given by $k \times x_1 + k \times x_2 + k \times x_3$. In our representation experiments, we set $x_1 = 100$, $x_2 = 100$, and $x_3 = 200$. The training plot for the third image from the Kodak dataset is shown in Fig. 18. For the image representation task, we used seven different activation functions, resulting in a total of 2800 epochs. To ensure a fair comparison, all other baselines were trained for $k \times (x_1 + x_2 + x_3)$ epochs. The convergence plot for the baselines is displayed in Fig. 19. When it comes to MIRE, with a dictionary size of k atoms and n layers, let us assume that in the i^{th} layer, MIRE is trained for x_i epochs. From the full set of $\sum_{i=1}^n x_i \cdot k$ epochs, MIRE extracts a signal-specific activation sequence using only $\sum_{i=1}^{n} x_i$ effective epochs, as one activation is only needed per layer, without any need for an exhaustive





Ground Truth



 35.90 dB
 33.68 dB

 0.946
 0.944

Figure 13. **Image inpainting capabilities of** *MIRE***.** The top row presents the ground truth and the masked image, where the text is random with varying font sizes. It is evident that *MIRE* achieves the highest PSNR and SSIM values while producing the most visually coherent inpainting outcome among the considered INRs.

grid search over activation function parameters.

10. Ablation studies

10.1. Effect of network hyperparameters

A study has been conducted to investigate the impact of varying the number of hidden nodes and layers, and learn-

ing rate on performance. The results related to changes in the number of hidden nodes are shown in left plot of Fig. 20. It is evident that *MIRE* surpasses all existing INRs regardless of the number of hidden neurons employed. This superiority is attributed to *MIRE*'s ability to self-optimize based on the input signal, ensuring that the optimal sequence of activations yields superior results compared to



Figure 14. **Complete decoded occupancy fields**: It is evident from these representations that most baselines showcase a low-pass decoded representation, primarily due to the inability of their activation functions to effectively capture the rapidly varying components across the INR. In contrast, *MIRE* delivers the occupancy field closest to the ground truth, achieving the highest IoU metric.

Distribution	MIRE	WIRE	SIREN	GAUSS
U(0,1)	39.63 ± 0.53	21.11 ± 1.01	17.48 ± 2.54	16.92 ± 3.16
U(-1,1)	39.51 ± 0.67	20.74 ± 1.51	15.66 ± 2.61	18.58 \pm 0.41
U(-10, 10)	$\textbf{40.07} \pm \textbf{0.66}$	37.35 ± 1.02	23.96 ± 6.81	24.01 ± 8.45
U(-100, 100)	$\textbf{36.32} \pm \textbf{1.83}$	24.05 ± 6.80	35.70 ± 2.84	22.45 ± 1.96
$\mathcal{N}(0,1)$	$\textbf{39.97} \pm \textbf{0.78}$	22.41 ± 4.03	15.93 ± 2.58	17.06 ± 3.17
$\mathcal{N}(0, 10)$	$\textbf{38.92} \pm \textbf{1.26}$	32.86 ± 4.79	28.87 ± 4.14	27.46 ± 4.78

Table 3. The PSNR variation of existing INRs when activation functions are drawn from different probability distributions.

pre-optimized INRs.

The middle plot of Fig. 20 depicts the effect of the number of layers on the PSNR. As can be seen, *MIRE* demonstrates exceptionally competitive PSNR metrics with merely one or two hidden layers. This performance stands out in contrast to other INRs, which begin with activation parameters already tailored for convergence. However, *MIRE* starts with random parameter initializations, highlighting its efficiency in self-optimization to reduce the loss between explicit and implicit representations. When the model configuration expands to include three hidden layers, *MIRE*'s ability to adapt and tailor its approach for the specific signal allows it to outperform all existing INRs, showcasing its robust optimization capabilities.

Lastly, the right side plot of Fig. 20 illustrates *MIRE*'s performance with log scale learning rate. As can be seen, *MIRE* shows strong performance, maintaining a high PSNR around 40 dB at a low learning rate of 10^{-3} . Even when the learning rate is increased by an order of magnitude to 10^{-2} , *MIRE* manages to sustain a relatively high PSNR. However, as the learning rate further increases to 10^{-1} , there is a noticeable decrease in PSNR, although it still performs better than every baseline except MFN. On the other hand, WIRE shows significant sensitivity to the learning rate. It reaches its peak PSNR at a learning rate of 10^{-2} , but its performance drastically drops as the learning rate increases further. This demonstrates that WIRE's optimal learning rate range is narrower, and its performance quickly deteri-



Figure 15. *MIRE*'s novel view synthesis capabilities on the Chair dataset: A closer inspection of the novel views generated by *MIRE* reveals that its synthesized images capture finer details, such as the intricate carvings on the top of the chair and its legs, to a much greater extent than the baselines. Additionally, MIRE successfully preserves the gold pattern on the chair's cushion without excessively smoothing it, unlike the baselines, which tend to produce low-pass representations. This ability to retain both fine details and texture while avoiding over-smoothing highlights *MIRE*'s superior performance in generating realistic novel views.

orates outside this range. SIREN also demonstrates a decline in performance, but its behavior is steadier compared to WIRE. As the learning rate increases, SIREN's PSNR decreases consistently, never reaching the peak values observed with *MIRE* or even WIRE. It can be concluded that, *MIRE* stands out for its ability to maintain high PSNR values across a wider range of learning rates, highlighting its robustness.

10.2. Effect of restricting the dictionary to one atom

A further ablation was conducted by restricting the dictionary to contain only a single atom (i.e., k = 1). The observed results are shown in Tab. 4.

Task	MIRE	Sinusoid	Gaussian	Sinc
Rep.	39.04	32.73	37.70	37.95
Inp.	37.57	34.82	36.09	35.56

Table 4. Performance (dB) for representation and inpainting tasks when restricting the dictionary to one atom.

10.3. Effect of residual connections

The effect of residual connections has also been studied. The results are shown in Tab. 5

Task	With Residual	Without Residual
Rep.	37.46	39.04
Inp.	35.75	37.57

Table 5. Performance (dB) for representation and inpainting tasks with and without residuals.

10.4. Effect of weight initialization

In most recent INR literature, it has been suggested that INRs equipped with space-frequency compact activation functions are generally less sensitive to the weight initialization mechanism, unlike those using sinusoidal activations [24, 25]. All the presented experimental results in this study is obtained through Pytorch's default weight initialization scheme. In this weight initialization process, the standard deviation (stdv) is defined as the reciprocal of the square root of the number of input units $n_{\rm in}$. Specifically, stdv = $\frac{1}{\sqrt{n_{\rm in}}}$, where $n_{\rm in}$ represents the number of features or input dimensions in the linear layer. The weights of the layer are then initialized by drawing each element of the weight matrix W from a uniform distribution in the range [-stdv, stdv], i.e.,



Figure 16. *MIRE*'s novel view synthesis capabilities on the Hotdog dataset: Upon closer inspection of the novel views generated by *MIRE*, it is evident that *MIRE* produces more realistic and detailed representations compared to the baselines. Specifically, *MIRE* captures the fine texture of the hotdog bun and the subtle lighting effects on the plate and condiments, which are much closer to the ground truth. In contrast, the baseline methods, such as WIRE and SIREN, tend to over-smooth the details, losing the sharpness in the condiments and the realistic shadows around the plate. *MIRE* not only preserves the structure and texture of the hotdog but also handles the complex lighting conditions more effectively, resulting in the most visually accurate and high-fidelity novel views.



Figure 17. Spatial domain activation function variation. *MIRE* benefits from all types of activations, as each has unique spatial and frequency characteristics.



Figure 18. **PSNR variation of** *MIRE* with epochs: *MIRE*'s training procedure is sequential. Initially, it optimizes a single hidden-layer INR, changing the activation function of the first layer every 100 epochs. Once the matched activation for the first hidden layer is found, a second hidden layer is added. The network is then trained while changing the activation function of the second layer every 100 epochs. After determining the matched activation for the second layer, a third hidden layer is added, and the network is retrained, changing the activation function of the third hidden layer every 200 epochs.



Figure 19. PSNR variation of the baselines for the third Kodak image

$$W_{ij} \sim \mathcal{U}\left(-rac{1}{\sqrt{n_{\mathrm{in}}}}, rac{1}{\sqrt{n_{\mathrm{in}}}}
ight)$$

This initialization strategy helps in keeping the weights small enough to ensure stable learning by preventing large gradients during the early stages of training. However, as different weight initializations often lead to different training dynamics of the network, we checked *MIRE*'s performance and the obtained activation sequences for different weight initialization schemes. For this experiment we utilized the Parrot image in the main paper. The obtained resuls are shown in Tab. 6.

As shown in Tab. 6, the sequence of activations changes when the weight initialization mechanism of the network is altered. This can be attributed to the fact that *MIRE* optimizes the network starting from the given weight initialization scheme, attempting to navigate the loss landscape toward a local minimum. When the distribution of initial weights is significantly changed, it affects the optimization trajectory. The network's parameters are adjusted based on gradients calculated from this starting point, and the training dynamics follow different paths depending on the initial conditions. Since the loss landscape may contain multiple local minima, the optimization path taken by the network can vary. This explains why different activation function sequences are optimal for different weight initializations—each initialization leads the network to explore a different part of the loss landscape, and the network adapts



Figure 20. Effect of the MLP architecture's hyperparameters on MIRE's performance

Weight Initialization Mechanism	Sequence	PSNR (dB)
Pytorch Default	Sinc, Gaussian, Gabor	39.04
Xavier Normal	PSWF, Gaussian, Gaussian	37.55
Orthogonal	RRC, Gaussian, Gaussian	37.15
SIREN-like ($\alpha \neq 30$)	PSWF, PSWF, Gaussian	36.57
SIREN-like ($\alpha = 30$)	PSWF, Sinc, PSWF	36.73

Table 6. Comparison of weight initialization mechanisms

its activations accordingly to minimize the loss. Therefore, due to the variability in the loss landscape and the dependence on the starting point W_0 , the training process may converge to different local minima, resulting in variations in the matched activation sequences for each weight initialization scheme. Therefore, it can be concluded that PyTorch's default weight initialization is the most effective for *MIRE*.

10.5. Effect of positional encoding

Positional embedding schemes are often utilized in INRs as a method of coordinate transformation. This technique functions similar to a frequency modulation mechanism, enabling the embedding of high frequencies within the signal. In our study, we assessed the performance of *MIRE* when incorporating this coordinate transformation.

The left side of Fig. 21 displays the decoded representations obtained from *MIRE* with positional embedding, while the right side shows those from *MIRE* without the positional embedding. Given that *MIRE* adopts a sequential training approach as outlined in Sec. 9.1, it allows for the observation of the decoded image at each layer. These images, referred to as 'Layer 1', 'Layer 2', and 'Layer 3' in Fig. 21, demonstrate the outcomes at each stage.

The comparison for the selected image reveals that the integration of positional embedding within *MIRE* results in improved performance for both the first and second layers, in contrast to the standard *MIRE* without any positional embedding. However, upon determining the best activations for the initial two layers, the introduction of the third layer shows that the standard *MIRE*, devoid of positional embed-

ding, optimizes the MLP in a manner that it outperforms the *MIRE* with positional embedding for this image.

11. Experimental setup

Our numerical implementation utilized the PyTorch framework and the Adam optimizer with a base learning rate of 0.001, which decays by 0.1 after each training epoch. Presented results were obtained using an MLP with three hidden layers, each containing 300 hidden nodes, and activation parameters drawn from either a uniform distribution with support between 0 and 1 or each activation's base configuration. For performance evaluation, we used Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) for images, and Intersection over Union (IoU) for occupancy fields.



Figure 21. Effect of positional embedding on *MIRE*: The first column displays images obtained under *MIRE*'s sequential training with positional embedding, while the second column depicts results without using positional embedding. It is evident that for the first two layers, *MIRE* with positional embedding scheme yields better results compared to the *MIRE* without positional embedding. However, upon introducing the third layer, *MIRE* without positional embedding surpasses the *MIRE* with positional embedding for this image.