

# In Search of a Data Transformation That Accelerates Neural Field Training

Junwon Seo\* Sangyoon Lee\* Kwang In Kim Jaeho Lee  
 Pohang University of Science and Technology (POSTECH)

{junwon.seo, sangyoon.lee, kimkin, jaeho.lee}@postech.ac.kr

## Abstract

*Neural field is an emerging paradigm in data representation that trains a neural network to approximate the given signal. A key obstacle that prevents its widespread adoption is the encoding speed—generating neural fields requires an overfitting of a neural network, which can take a significant number of SGD steps to reach the desired fidelity level. In this paper, we delve into the impacts of data transformations on the speed of neural field training, specifically focusing on how permuting pixel locations affect the convergence speed of SGD. Counterintuitively, we find that randomly permuting the pixel locations can considerably accelerate the training. To explain this phenomenon, we examine the neural field training through the lens of PSNR curves, loss landscapes, and error patterns. Our analyses suggest that the random pixel permutations remove the easy-to-fit patterns, which facilitate easy optimization in the early stage but hinder capturing fine details of the signal.<sup>1</sup>*

## 1. Introduction

Neural field is a form of data representation that parameterizes each target signal as a neural network that maps spatiotemporal coordinates to the signal values [33]. For example, a colored image can be represented by a model that maps (X,Y) pixel coordinates to the corresponding (R,G,B) values. This parameterization enjoys many advantages in faithfully and efficiently representing high-dimensional signals with fine detail, and thus is being widely used for modeling signals of various modalities, such as image [4], video [10], 3D scene [19], or spherical data [8].

A key obstacle that prevents the widespread adoption of neural fields is their *training cost*. To represent each datum as a neural field, one must train a neural network using many SGD iterations. For instance, NeRF requires at least 12 hours of training time on GPUs to represent a single 3D scene [19]. Representing a set of data thus requires a considerable amount of computation and time, making it very

difficult to develop practical applications of neural fields.

Many prior works view the “optimization bias” as a major cause behind the long training time of neural fields [29]. In particular, the spectral bias of the SGD-based training is known to bias the neural network to prioritize fitting the low-frequency components of the target signal and leave high-frequency components for the late stage of training [22]. It has been observed that such a tendency greatly hinders the neural fields from expressing natural data (*e.g.*, images) with fine, high-frequency details [29, 37].

The prevailing strategy to mitigate such bias is to introduce a useful *prior* (or inductive bias) that can help neutralize the negative effects of the bias. One popular approach is to develop new network components that bring a favorable *architectural prior*, such as the Fourier features [19], sinusoidal activation [27], or spatial encoding [18, 20]. Other works also attempt to introduce the prior in the form of *initial parameters* that have been meta-learned from a large set of signals [26, 30]. Such meta-learned initializations tend to have rich high-frequency spectra, which can help fitting natural signals within a small number of SGD steps [30].

In this paper, we approach the problem from a different angle. In particular, we ask the following question:

*“Can we **exploit** the optimization bias of SGD, instead of fighting against it?”*

Precisely, we ask whether we can *transform* the datum in a way that the optimization bias acts favorably in fitting the transformed data with a neural field. If there exists such a transformation, and if it admits an efficiently computable inverse, we may be able to use the following strategy to reduce the neural field training cost: we train a neural field that expresses the transformed signal, and the original signal can be recovered from the model by applying an inverse transformation to the signal generated by the model (Fig. 1).

**Contribution.** As a first step to answer this question, we conduct an extensive empirical study on how applying simple data transformations affect the computational cost of training the corresponding neural field. In particular, we compare the number of SGD steps needed to fit the transformed data to a certain fidelity level (*e.g.*, PSNR 50 for images), with the steps to fit the original data. We try to-

\*equal contribution

<sup>1</sup>code: <https://github.com/eff1-lab/DT4Neural-Field>

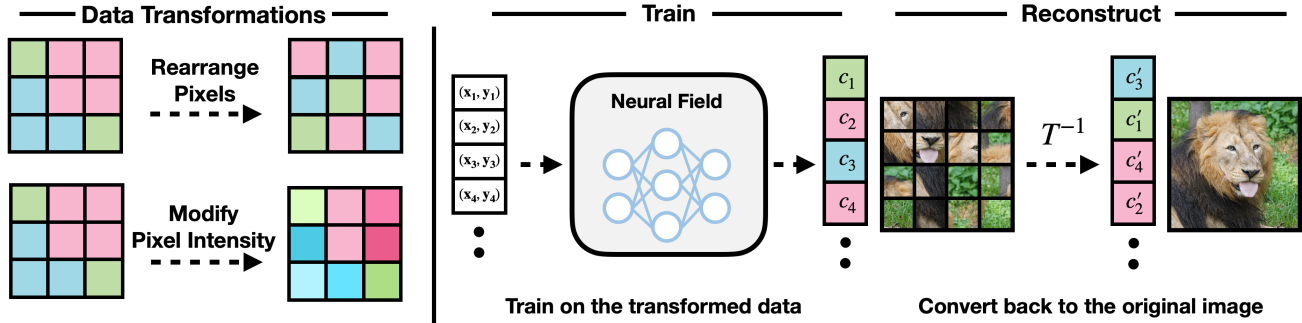


Figure 1. **Overall pipeline.** We consider a three-step procedure to train neural fields. (1) Apply a data transformation to the target datum. (2) Train a neural field to fit the transformed data. (3) Reconstruct the original data by generating the transformed datum from the neural field, and then applying the inverse of the data transformation. By selecting the data transformation carefully, we can significantly reduce the computational cost required to train a neural field that achieves the desired quality of approximation.

tal seven different data transformations, from the one that permutes the location of the pixels to the one that scales the intensity values of each pixel.

From the experiment, we make a surprising observation: we find that the *random pixel permutation* (abbr. RPP) provides a consistent acceleration on a range of datasets and various neural field architectures with sophisticated encoding schemes. The RPP is in fact the only data transformation that provided acceleration on every experimental setups that we considered. We find that the original data requires, on average,  $\sim 30\%$  greater number of SGD steps to achieve the similar fidelity level than the RPP data. As RPP tends to bias the data toward high-frequency, it is quite counter-intuitive that they can be fit faster than the original signals.

*Why does the random pixel permutation accelerate neural field training?* To explain this phenomenon, we articulate the following “blessings of no pattern” hypothesis:

Original data often have smooth, representative patterns that facilitate easy optimization, particularly in the early phase of learning. However, this smoothness quickly transforms into an obstacle when aiming for a sufficiently high level of fidelity. Random pixel permutation removes these easy-to-fit patterns, which accelerates optimization in the long run. (★)

To corroborate our hypothesis ★, we take an in-depth look into the training dynamics of the original and the RPP data. We find that, indeed the training speed on the RPP data is slower than on the original data during the early training phase (Sec. 4.1). However, after training for a sufficient number of epochs, the neural field trained on RPP data finds a “linear loss highway” on which the optimization is very easy; the optimization on original data fails to find one (Sec. 4.2). Furthermore, we find that the errors in the RPP images are more evenly distributed over pixels and lack visually distinguishable structures; errors in original images tend to have periodic or axis-aligned patterns, which might

be the artifacts of the encoding scheme (Sec. 4.3).

To sum up, our contribution can be summarized as:

1. Through systematic study, we find that simple data transformations can dramatically change the training speed of neural field ( $\times 0.1$ – $\times 20$ ), even on state-of-the-art neural fields architectures with sophisticated encodings [20].
2. We discover that random pixel permutations (RPP) provide consistent acceleration over fitting the original image, providing  $\times 1.08$ – $\times 1.50$  speedups.
3. We conduct an in-depth analysis which sheds light on how RPP speeds up the training by removing the easy-to-fit patterns that slow down the training eventually.

Despite the limitation of RPP that it may be difficult to be applied for applications where generalization is critical, our study has at least two potential impact areas. First, the RPP itself may be useful in the applications of neural field where a strong training fidelity is the core performance criterion, such as data compression [24]. In fact, one of the core challenges in neural-field-based data compression is the slow encoding speed (see, *e.g.*, COOL-CHIC [12]). Second, our study provides a concrete starting point in developing a new data transformation that strikes the balance between the training speed and the generalizability.

## 2. General framework

The general framework of finding data transformations that accelerate neural field training can be formalized as an optimization problem, as we describe in this section.

### 2.1. Formalisms

Suppose that we want to use the neural field parameterization to approximate some *signal*  $\mathbf{x} \in \mathcal{X}$ . Here, the *signal space*  $\mathcal{X}$  is the space of all signals that have the same data type. For example, we can let  $\mathcal{X}$  be the set of all  $256 \times 256$  RGB images; in such case, we may have  $\mathcal{X} \subseteq \mathbb{R}^{256 \times 256 \times 3}$ .

The *data transformation*  $T : \mathcal{X} \rightarrow \mathcal{X}$  maps an element in the signal space to another element. We define the *trans-*

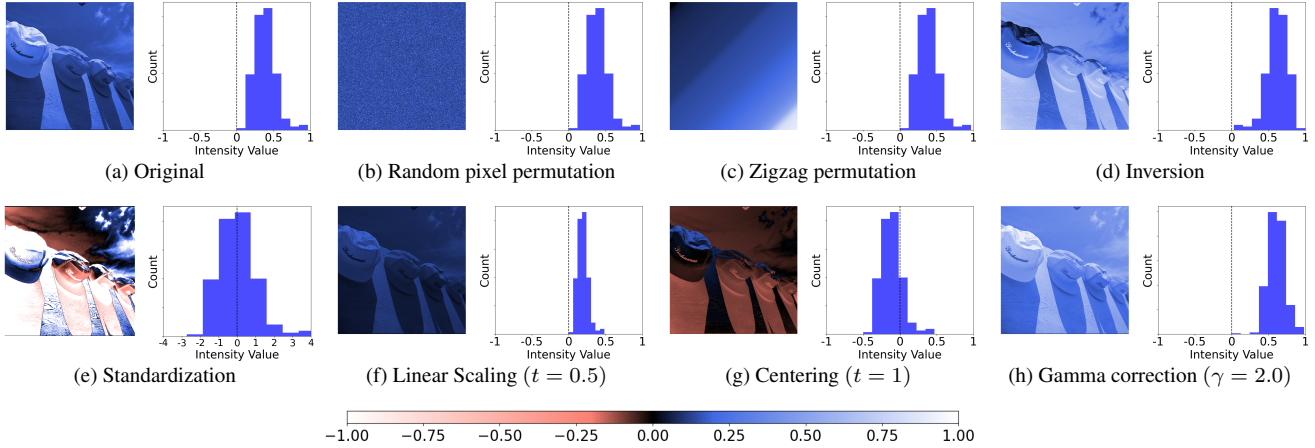


Figure 2. **Data transformations considered.** In each subfigure, we visualize the data transformations by illustrating how the transformed image (left) and the intensity histogram (right) looks like on a Kodak image.

formation space  $\mathcal{T}$  as a set of data transformations on  $\mathcal{X}$  that has an inverse.

Our goal is to find a data transformation  $T$  that minimizes the training cost of fitting the transformed signal  $T(\mathbf{x})$ . To formalize this, we define the *training cost*

$$\text{cost} : \mathcal{X} \times \mathcal{T} \rightarrow \mathbb{R}_+. \quad (1)$$

Here, the function  $\text{cost}(\mathbf{x}, T)$  measures the computational burden required to train a neural field for  $T(\mathbf{x})$ , whenever its outcome is inverted back via  $T^{-1}$ , approximates the original signal  $\mathbf{x}$  with the desired level of precision. The burden may be measured in various ways, *e.g.*, the number of SGD steps, FLOPs, or wall-clock time. We note that the training cost of  $\mathbf{x}$  using  $T(\cdot)$  need not be identical to the training cost of fitting  $T(\mathbf{x})$  with the same precision, *i.e.*,

$$\text{cost}(\mathbf{x}, T) \neq \text{cost}(T(\mathbf{x}), \text{Id}), \quad (2)$$

especially when the transformation involves scaling of the signal. Also, the training cost may heavily depend on the choice of neural field architectures, hardware type, and the batch size—using a larger batch size tends to reduce the number of steps until convergence [25], but it also requires training with a larger memory.

Given these tools, we work to consider the optimization minimize  $\text{cost}(\mathbf{x}, T)$ , subject to  $T \in \mathcal{T}^*$ , (3)

where  $\mathcal{T}^* \subseteq \mathcal{T}$  is a subset of the transformation space that satisfies some desired properties. The set  $\mathcal{T}^*$  can be configured in many different ways, to account for the expected usages of the trained neural fields. This problem can also be extended to a version where we have a probability distribution of the signals  $\mathbf{x}$ , and find a single transformation  $T$  that minimizes the expected cost.

**Scope of this paper.** While the most general framework is to solve the optimization (3) directly, the problem is intractable unless we have an expressive, well-paramaterized

transformation space  $\mathcal{T}^*$  that admits an effective optimization method. In this paper, we focus on the *proof-of-concept* that there exists some choice of  $T$  such that

$$\text{cost}(\mathbf{x}, T) < \text{cost}(\mathbf{x}, \text{Id}), \quad (4)$$

(Sec. 3), and deepening our understanding on when and why such  $T(\cdot)$  can accelerate the training (Sec. 4).

## 2.2. Example cases

Here are some examples desired properties of data transformation, and how they relate to practical applications. We provide a more in-depth discussions in the ??.

**Efficient invertibility.** The inverse  $T^{-1}$  should be able to be computed efficiently, *e.g.*, by performing a linear operation. Ideally, one may wish to be able to combine the inverse transformation into the neural field by modifying the parameters of the trained neural field directly. For instance, if  $T(\mathbf{x}) = -\mathbf{x}$ , we can incorporate the inverse in the neural field by negating the final layer weights of the neural field. This property is useful whenever we expect many repeated inferences of the neural field, *e.g.*, for real-time interaction.

**Retains interpolatability.** We want our transform  $T$  to admit a principled way to sample the value of interpolated coordinates from the neural field that approximates the transformed signal. In some neural field applications, *e.g.*, super-resolution [4], this interpolating capability plays an essential role. On the other hand, some applications like data compression do not involve any interpolation [24].

## 3. Data transformations vs. training speed

In this section, we compare the training cost of the neural field training on various data transformation schemes. As it turns out, there indeed exists a nice data transformation that satisfies ineq. 4: the *random pixel permutation* (RPP).

Architecture	Dataset	Random pixel perm.	Zigzag perm.	Inversion	Standardization	Linear Scaling		Centering		Gamma Correction	
						$t = 0.5$	$t = 2.0$	$t = 1.0$	$t = 2.0$	$\gamma = 0.5$	$\gamma = 2.0$
SIREN	Kodak	1.26×	17.90×	0.72×	2.20×	0.79×	1.39×	0.76×	0.97×	0.36×	0.80×
	DIV2K	1.30×	21.74×	0.70×	2.36×	0.82×	1.45×	0.97×	1.22×	0.30×	0.90×
	CLIC	1.08×	17.07×	0.68×	2.18×	0.76×	1.36×	0.87×	1.14×	0.24×	0.82×
Instant-NGP	Kodak	1.50×	0.86×	0.21×	0.45×	1.15×	0.75×	1.24×	0.80×	0.29×	0.38×
	DIV2K	1.32×	0.80×	0.20×	0.61×	1.12×	0.64×	1.39×	0.97×	0.14×	0.33×
	CLIC	1.35×	0.78×	0.07×	0.53×	1.03×	0.68×	0.99×	0.66×	0.12×	0.32×

Table 1. **Acceleration factors of data transformations.** We compare the average acceleration factor (eq. 5) of seven data transformations under six different combinations of datasets and models. We take average of cost ratios, instead of ratios of averages, to avoid the statistic being driven by a small number of outliers that require very long training time. The shaded figures denote the speedup.

### 3.1. Experimental setup

We focus on the task of 2D image regression. This choice reduces the computational burden of fitting each datum, so that we can make experimental validations on multiple datasets and models with extensive hyperparameter tuning.

For the training cost, we use *the number of SGD steps* until we reach the training PSNR over 50dB (a near-perfect reconstruction for 8-bit images). The number of SGD steps is an informative indicator of the total FLOPs and runtime for any fixed choice of hardware, model architecture, and data. Importantly, we tune the learning rate for each (transformed) image, making it close to the scenario where all settings have been optimized to minimize the runtime.

For each choice of an image and the data transformation, we measure the *acceleration factor*, *i.e.*,

$$\text{acc}(\mathbf{x}, T) \triangleq \frac{\text{cost}(\mathbf{x}, \text{Id})}{\text{cost}(\mathbf{x}, T)}. \quad (5)$$

Other configurations are as follows.

**Datasets.** We use three different image datasets:

- *Kodak*. Consists of 24 different images from the Kodak lossless true color image suite [11].
- *DIV2K*. Consists of 100 images in the validation split of the DIV2K image super-resolution challenge (HR) [1].
- *CLIC*. Consists of 100 images from the validation split of the dataset for the Challenge on Learned Image Compression 2020 (CLIC) [31]. There are total 102 images in the validation split, but we removed two images that have a sidelength shorter than 512 pixels.

We pre-process each image in the datasets as follows. Following BACON [14], we resize each image into  $512 \times 512$  pixels by first center-cropping the image to a square with sidelengths equal to the shorter sidelength of the original image, and then resizing with the Lanczos algorithm. Next, we convert the image to grayscale. Finally, we apply sRGB to Linear RGB operation, as in Instant-NGP [20].

The resulting images have  $512 \times 512$  pixels (thus total  $2^{18}$  pixels) with intensities lying in the interval  $[0, 1]$ .

**Models.** We use two different neural field architectures.

- *SIREN*. A classical architecture that uses multi-layered perceptrons with sinusoidal activation functions [27]. We

configure the model to have three hidden layers with 512 neurons in each layer and output dimension 1. We use the default frequency scaling hyperparameter  $\omega_0 = 30$ .

- *Instant-NGP*. A more recently proposed neural field with multi-resolution hash encodings [20]. We use the default setup for the image regression, with the maximum hash table size and 16-bit parameters.

**Training.** We use full-batch gradient descent, *i.e.*, the batch size  $2^{18}$ . This batch size, in general, minimizes the number of SGD steps required to fit the target image. We note that we sample directly from the coordinate grid, instead of sampling from interpolated coordinates (as in, *e.g.*, [20]).

**Hyperparameters.** We tune the learning rate using the grid search. For SIREN, we tune the learning rate in the range  $\{2^{-8}, \dots, 2^{-16}\}$ . For Instant-NGP, we tune the learning rate in the range  $\{2^{-4}, \dots, 2^{-15}\}$ .

**Data transformations.** We consider total seven elementary data transformations (Fig. 2): Two transformations that only change the location of the pixels (random pixel permutation, zigzag permutation), and five transformations that change the intensity values of each pixels (inversion, standardization, linear scaling, centering, and gamma correction).

- *Random pixel permutation* (RPP). We randomly permute the location of each pixels. This transformation generally increases the frequency spectrum of the image.
- *Zigzag permutation*. We sort all the pixels in the ascending order of their intensities. Then, we place the sorted pixels in the zigzag order, starting from the upper left corner. This produces a low-frequency image.
- *Inversion*. We invert the intensity of each pixel, *i.e.*, perform  $z \mapsto 1 - z$  on each intensity values.
- *Standardization*. We measure the mean  $\mu$  and standard deviation  $\sigma$  for the intensities of all pixels in the image. Then we standardize the intensities via  $z \mapsto \frac{z - \mu}{\sigma}$ .
- *Linear scaling*. We scale the intensities of each pixel by  $t$  without any centering, *i.e.*  $z \mapsto tz$ .
- *Centering*. We center the interval  $[0, 1]$  at zero and scale the intensities by  $t$ , *i.e.*,  $z \mapsto t(z - 1/2)$ .
- *Gamma correction*. We nonlinearly scale the intensity of each pixels as  $z \mapsto z^{1/\gamma}$ . Choosing  $\gamma > 1$  makes images brighter, and  $\gamma < 1$  makes darker.

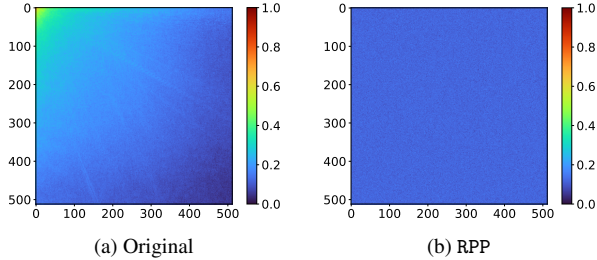


Figure 3. **Frequency spectra of original vs. RPP.** We compare the average DCT coefficients of the original and RPP Kodak images. Upper left region denotes the low-frequency, and the lower right region denotes the high-frequency.

### 3.2. Results

We report the average acceleration factors in Tab. 1. From the table, we make several observations:

- RPP provides a consistent acceleration over the original image; see Sec. 4 for an in-depth analysis.
- The zigzag permutation works tremendously well on SIREN—speeding up training by over  $\times 17$ —but slows down training Instant-NGP. We hypothesize that this is due to the axis-aligned inductive bias imposed by the spatial grid encoding; see Sec. 4.3 for more discussion.
- Inversion consistently slows down the training. This fact, ironically, implies that inversion might have been an effective accelerator if all “natural images” looked like the color-inverted images.
- Scaling up the intensities tends to have opposite effects on different architectures. Scaling up speeds up training in SIRENs, but slows down in Instant-NGPs.
- Gamma correction slows down the training, regardless of taking the power of 2 or  $1/2$ . We suspect that this is due to the precision errors from taking powers. In other words, excessive scaling operations may not be helpful in accelerating the training.

### 4. A closer look at the random permutation

We now focus on a specific type of data transformation: the random pixel permutation (RPP). The RPP relocates each pixel to a new random location, *i.e.*,

$$(\text{coords}_i, \text{values}_i) \mapsto (\text{coords}_{\pi(i)}, \text{values}_i) \quad (6)$$

for some random permutation  $\pi(\cdot)$ . The marginal distributions of the input coordinates or the output values remain the same. As we have seen in Tab. 1, the RPP transformation accelerates the fitting of the corresponding neural field, which is very unexpected and difficult to explain.

**Why is this strange?** From the perspective of the *spectral bias*, the RPP operation should have slowed down the training speed, instead of accelerating it. As demonstrated by Rahaman et al. [22], typical neural networks first rapidly fit the low-frequency components of the target signal, and fit

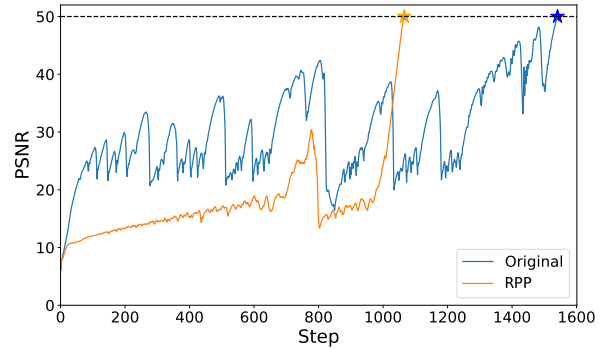


Figure 4. **PSNR curves for a single Kodak image.** The original image (orange) excels during the early stage of training but the RPP image quickly reaches PSNR 50 in the middle of the training.

	Target PSNR		Ratio
	30dB	50dB	
Original	105.8	1371.4	0.08
Random pixel perm.	851.8	1100.5	0.79

Table 2. **The number of steps to reach PSNR 30dB and 50dB.** We measure the average steps to reach moderate and high PSNR levels over Kodak images, and report the average ratio.

the high-frequency components much later. The RPP transformation, in this sense, should have been detrimental to the training. In fact, RPP images tend to have more higher-frequency components than the original images. In Fig. 3, we provide the average discrete cosine transform (DCT) coefficients of the original and RPP images; for visualization, we raise the coefficients to the power of 0.03, similar to [35]. We observe that the original image is more biased toward low-frequency than RPP. However, we also observe that RPP images fit faster than the original images.

In the remainder of this section, we identify three different aspects that the training dynamics of RPP images are critically different from that of original images. These aspects, when put together, support our “blessings of no pattern” hypothesis (★) that RPP accelerates training by removing easy-to-fit patterns that can be harmful to reaching a high level of fidelity. In particular, we show that

- RPP is slow to reach moderate PSNR levels, but is fast to reach high PSNR (Sec. 4.1).
- there is a linear path that connects the moderate-to-high PSNR points in the RPP loss landscape (Sec. 4.2).
- model trained on RPP images tends to have less structured error, while vanilla training leads to error patterns that reflect the underlying encoding schemes (Sec. 4.3).

#### 4.1. The PSNR curve: Slower to approximate well, but faster to approximate “very well”

First, we observe that RPP images tend to achieve moderate PSNRs (*e.g.*, 30dB) much later than the original images, but they quickly reach high PSNRs (*e.g.*, 50dB) afterwards.

In Fig. 4, we provide an example plot of fitting a SIREN

on a natural image from the Kodak dataset. From the figure, we can immediately make the following observations:

- *RPPs are slow starters.* The training PSNR of the original image quickly arrives at  $\sim 30$ dB the early stage of training, even before taking 300 SGD iterations. On the other hand, the PSNR of the RPP image remains very small, under 15 dB at the similar number of steps.
- *Explosive surge in later stage.* After staying at a low PSNR level (under 20 dB) for a while, the PSNR of RPP images bursts explosively. The PSNR jumps from  $\sim 15$ dB to 50dB in less than 100 SGD steps. On the other hand, the PSNR curve of the original image oscillates wildly at 30–35dB level for a long time, and gradually reaches the 50dB at around 1600 steps.

These observations are consistent over the choice of the target image; we provide a comprehensive collection of figures in the supplementary materials.

For a more quantitative comparison, we compare the average number of SGD steps to arrive the PSNR of 30dB and 50dB for the original and RPP images (Tab. 2). Again, we use the images from the Kodak dataset and fit with SIRENs.

We observe that, for original images, the number of steps to reach 30dB is around 100 steps, while the number of SGD steps to reach the 50dB is over 1300 steps. That is, it takes almost  $13\times$  more steps to reach the high PSNR than for moderate PSNRs. On the other hand, for RPP images, it takes over 800 steps to reach the moderate PSNR level, while it takes only 300 more steps to reach the high PSNR.

#### 4.2. Linear paths in the loss landscape: RPP images find a “linear expressway”

From the previous observation, it seems likely that, for RPP images, there may exist an *expressway* in the neural field parameter space that connects the moderate-PSNR parameter with the high-PSNR parameter; on this expressway, the SGD may not encounter too many hills to circumvent, and may mostly follow a smooth linear path.

To validate this intuition, we visualize the loss landscape of the SIREN trained on Kodak images. In particular, we capture two different phases of training.

1. *Early phase.* We plot the linear path from the initialization to 30dB point, *i.e.*, the parameter that first achieves PSNR 30dB during the training (Fig. 5)
2. *Late phase.* We plot the linear path from the 30dB point to the PSNR 50dB point (Fig. 6).

To generate such visualization, we project the parameter space into a two-dimensional space, as in [13]. Unlike [13], we fix one axis to a directional vector between two parameter points that we want to capture (*e.g.*, 30dB point and 50dB point); the other axis has been decided randomly.<sup>2</sup>

During the early phase (Fig. 5), we observe that the loss landscape of RPP is quite hostile; unlike in original images,

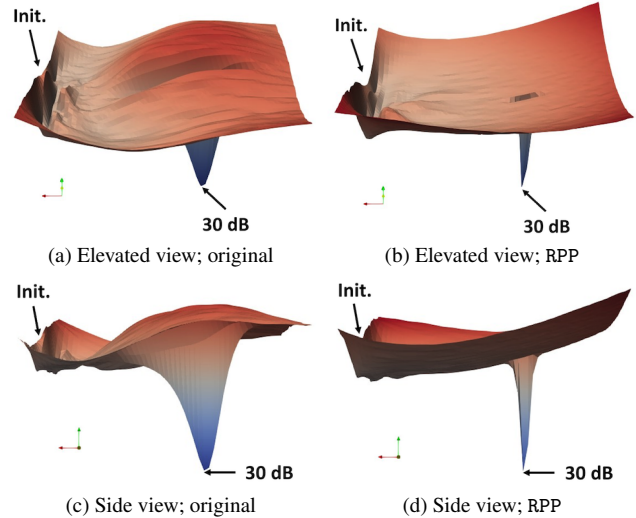


Figure 5. **SIREN loss landscape: from initial point to 30dB point.** For original and RPP versions of a Kodak image, we plot the loss landscape between the initial point and the parameter that achieves PSNR 30. For RPP, the minima is much narrower and lacks a clear pathway toward it, unlike in the original image.

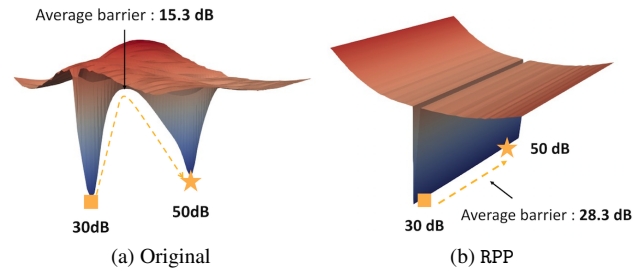


Figure 6. **SIREN loss landscape: from 30dB point to 50dB point.** For original and RPP versions of a Kodak image, we plot the loss landscape between the parameters that first achieve PSNR 30 (■) and PSNR 50 (★), respectively. The loss barriers between two minima are dramatically different; 15.3dB for the original image, and 28.3dB for the RPP (averaged over all Kodak images).

the minima that 30dB point belongs to is quite narrow and there is no clear pathway toward the 30dB point.

During the late phase (Fig. 6), for the RPP image, once the parameter arrives at the minima, there exists a linear path that connects the 30dB point to the 50dB point. If we measure the loss barrier, *i.e.* the maximum loss on the linear path between 30dB and 50dB points, it is as low as 28.3dB for the RPP versions of Kodak images. For original images, the loss barrier is quite high (15.3dB).

#### 4.3. Patterns in the error: The errors are less structured in RPP images

Another interesting property of RPP is that the errors that the neural field makes are more evenly distributed among the pixels, lacking a clearly distinguishable pattern. It turns out that the *whiteness* of the RPP error can help generat-

<sup>2</sup>A similar method has been considered by Goodfellow et al. [7].

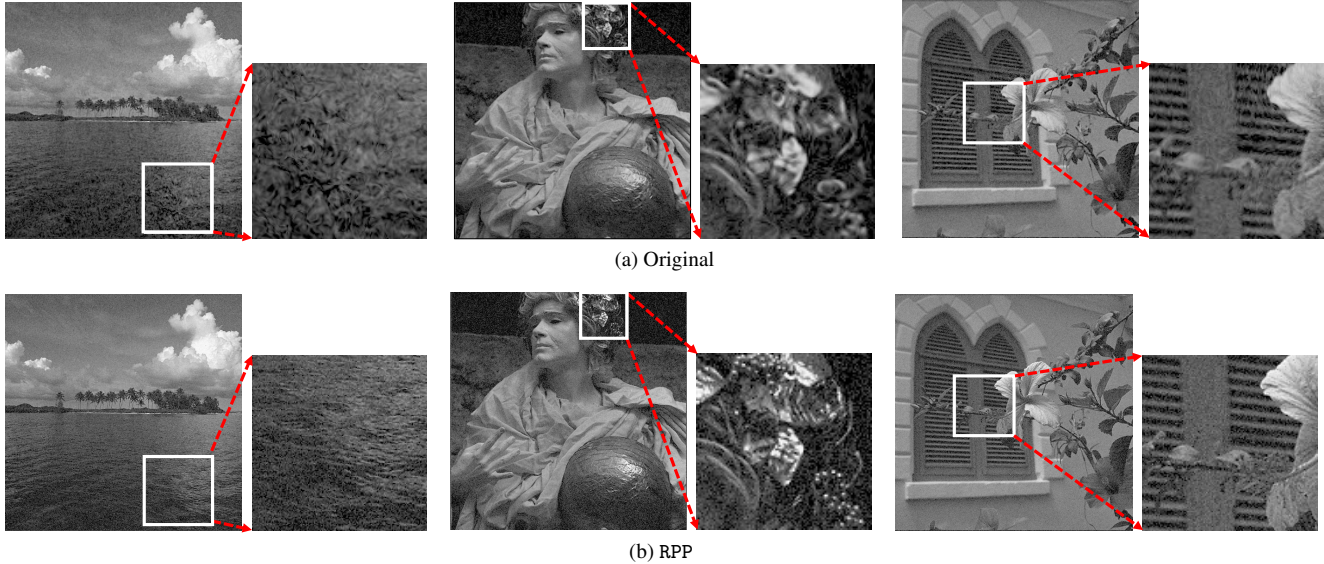


Figure 7. **Example reconstructions at a target PSNR value of 30 (SIREN).** Random pixel permutation enables the network to prioritize the restoration of high-frequency, fine details, whereas the original representation tends to steer the network toward restoring low-frequency components initially. While this behavior may pose challenges during the early stages of training, it ultimately facilitates achieving visually more plausible reconstructions much earlier than the original representation.

	Target PSNR			
	20	30	40	50
Original	10.51e-04	23.68e-06	19.46e-08	5.90e-10
Random pixel perm.	5.68e-04	5.89e-06	7.59e-08	5.81e-10
Average ratio	2.03×	5.20×	5.80×	1.55×

Table 3. **Loss variance over pixels.** We measure the average pixel loss variance in the SIREN trained on Kodak dataset, measured at various levels of PSNR. The “average ratio” denotes the mean of  $\text{Var}(\text{original})/\text{Var}(\text{RPP})$ . RPP has a lower loss variance at all PSNRs.

ing more visually sharp and satisfactory images than neural fields trained on original images.

In Tab. 3, we measure the pixel loss variance in both original and RPP images trained to various target PSNR levels. From the table, we observe that the loss variance is up to  $5.80\times$  larger in the neural fields trained on the original images than on RPP images. This implies that, for original images, some pixels in the original image are fit much faster than other pixels. Putting that differently, the neural field trained on the unpermuted images tends to prioritize learning a specific type of patterns first, while the neural field trained on RPP images does less so.

Exactly what pattern is RPP avoiding to fit? In Figs. 7 and 8, we compare the images generated by the neural fields trained on original and RPP images on SIREN and Instant-NGP, respectively. In Fig. 7, we observe that SIRENs trained on original images have thicker wavy patterns, which is likely to be a consequence of the sinusoidal encoding structure of SIREN. The models trained on RPP images, on the other hand, have much sharper texture with fine-grained noise. A similar phenomenon is observed in

the case of Instant-NGP (Fig. 8). Here, we observe that the models trained on original images attain axis-aligned artifacts; black (or white) horizontal (or vertical) blocks appear in the rendered images. On the other hand, the error is much less structured in models trained on RPP images.

## 5. Related work

**Optimization bias of neural networks.** It has been well known, both theoretically and empirically, that SGD-based optimization algorithms are biased toward learning “simpler solutions,” for various notions of simplicity. A prominent example is the *spectral bias*, a tendency to prioritize learning lower-frequency components of the target function [3, 22]. Other types of biases have also been studied in the literature: SGD tends to prefer learning smaller-norm solutions [9, 21] and low-rank solutions [32, 34]. These works mostly focus on explaining how such simplicity biases are advantageous in learning a solution that can generalize well (see, e.g., [23]). Our work, in contrast, describes how such bias can be *disadvantageous* in terms of the training speed, instead of the generalization performance.

**Random labels and memorization.** The question of how neural nets behave when trained on random labels has been actively studied since its connection to the generalization capability of deep learning has been established [38]. Arpit et al. [2] finds that neural networks tend to learn simple patterns first and then *memorize* the unexplainable samples—such as random label and data—in the later phase. A more recent study argues that training on randomly labeled data can actually be advantageous [17]; pre-training on random



Figure 8. **Example reconstructions at a target PSNR value of 20 (Instant-NGP).** Similarly to the SIREN case, we observe that Instant-NGPs trained on the RPP images have much less visually distinguishable artifacts. The models trained on original images tend to have axis-aligned blocks of error, which is likely to have originated from the spatial grid encoding of the model.

data helps fitting the training data, when eventually fine-tuning on the dataset with correct labels. Different from these works, our work discovers how random label can be beneficial in the neural field context, without any consideration on further fine-tuning.

**Faster training of neural fields.** Other than the works described in the introduction, there have been many other attempts to accelerate the neural field training by designing better encoding schemes. Positional encodings [5, 19, 27, 29] and spatial grids [6, 16, 18, 20] are the most popular options, and some works also use tree-like structures [28, 36]. Our work is complementary to this line of works, and argues that transforming the data in the original data space can provide auxiliary training speed boosts. We also note that a concurrent work also aims to boost the training speed by

modifying the given data [15]. In particular, [15] finds that separating the visual signals into many sub-segments can facilitate the convergence of SIRENs and proposes a meta-learning algorithm to speed up training. Our work considers a more general class of data transformations, and focus on understanding how transformations accelerate training.

## 6. Conclusion

Learning neural fields poses a unique challenge, as the primary objective is to *overfit* to the given data—an aim contrary to the typical goals in other problem domains where learning attempts to achieve *regularization* to prevent overfitting. In our study of this distinctive problem, we discovered that certain data transformations can expedite the network’s acquisition of fine details more rapidly than when learning from the original images. Specifically, we demonstrated that random pixel permutations, which explicitly make learning low-frequency details more challenging, guide the network away from fixating on low-frequency components. This resulted in significantly enhanced PSNR within a given computational budget and, furthermore, yielded visually more plausible reconstructions at the same PSNR level. Our insights are substantiated by thorough empirical evaluations.

**Limitation & future direction.** A notable limitation of our approach is that its direct applicability may be constrained when learning neural networks for other problem domains such as classification and regression, where regularization is crucial. As a future work, we aim to conduct more comprehensive studies to understand how data transformations affect both trainability and generalization of neural fields (and other neural networks). Such exploration may offer new insights on how to modify the given workload to jointly optimize the training cost and the model performance.

Another, more narrowly scoped future direction will be to gain a more concrete, theoretical understanding on the effect of random pixel permutations on the training dynamics of neural fields. Although our empirical analyses provide some insights into how RPP can help accelerate the training, we are yet to fully understand why fitting the easy patterns slows down the later phase of training. Deeper understandings on this matter may lead to impactful lessons applicable to any deep learning domain that involves long training.

**Acknowledgments.** This work was partly supported by the National Research Foundation of Korea (NRF) grant (RS-2023-00213710, Neural Network Optimization with Minimal Optimization Costs), and partly by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant (No.2022- 0-00713, Meta-learning applicable to real-world problems, No.2019-0-01906, Artificial Intelligence Graduate School Program (POSTECH)), all funded by the Korea government (MSIT).



## References

- [1] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017. 4
- [2] Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *ICML*, 2017. 7
- [3] Ronen Basri, David Jacobs, Yoni Kasten, and Shira Kritchman. The convergence rate of neural networks for learned functions of different frequencies. In *NeurIPS*, 2019. 7
- [4] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *CVPR*, 2021. 1, 3
- [5] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *ICLR*, 2021. 8
- [6] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 8
- [7] Ian J. Goodfellow, Oriol Vinyals, and Andrew M. Saxe. Qualitatively characterizing neural network optimization problems. In *ICLR*, 2015. 6
- [8] Daniele Grattarola and Pierre Vandergheynst. Generalised implicit neural representations. *NeurIPS*, 2022. 1
- [9] Suriya Gunasekar, Blake Woodworth, Behnam Neyshabur Srinadh Bhojanapalli, and Nathan Srebro. Implicit regularization in matrix factorization. In *NeurIPS*, 2017. 7
- [10] Subin Kim, Sihyun Yu, Jaeho Lee, and Jinwoo Shin. Scalable neural video representations with learnable positional features. In *NeurIPS*, 2022. 1
- [11] E. Kodak. Kodak dataset, 1999. 4
- [12] Théo Ladune, Pierrick Philippe, Félix Henry, Gordon Clare, and Thomas Leguay. COOL-CHIC: Coordinate-based low complexity hierarchical image codec. *arXiv preprint 2212.05458*, 2022. 2
- [13] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *NeurIPS*, 2018. 6
- [14] David B. Lindell, Dave Van Veen, Jeong Joon Park, and Gordon Wetzstein. BACON: Band-limited coordinate networks for multiscale scene representation. In *CVPR*, 2022. 4
- [15] Ke Liu, Feng Liu, Haishuai Wang, Ning Ma, Jiajun Bu, and Bo Han. Partition speeds up learning implicit neural representations based on exponential-increase hypothesis. In *ICCV*, pages 5474–5483, 2023. 8
- [16] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *NeurIPS*. Curran Associates, Inc., 2020. 8
- [17] Hartmut Maennel, Ibrahim M. Alabdulmohsin, Ilya O. Tolstikhin, Robert Baldock, Olivier Bousquet, Sylvain Gelly, and Daniel Keysers. What do neural networks learn when trained with random labels? In *NeurIPS*, 2020. 7
- [18] Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Adaptive coordinate networks for neural scene representation. *ACM Transactions on Graphics*, 40(4), 2021. 1, 8
- [19] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 8
- [20] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):102:1–102:15, 2022. 1, 2, 4, 8
- [21] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *ICLR*, 2015. 7
- [22] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *ICML*, 2019. 1, 5, 7
- [23] Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. In *NeurIPS*, 2020. 7
- [24] Jonathan Schwarz, Jihoon Tack, Yee Whye Teh, Jaeho Lee, and Jinwoo Shin. Modality-agnostic variational compression of implicit neural representations. In *ICML*, 2023. 2, 3
- [25] Christopher J. Shallue, Jaehoon Lee, Joseph Antognini, Jascha Sohl-Dickstein, Roy Frostig, and George E. Dahl. Measuring the effects of data parallelism on neural network training. *Journal of Machine Learning Research*, 20(112): 1–49, 2019. 3
- [26] Vincent Sitzmann, Eric Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. MetaSDF: Meta-learning signed distance functions. In *NeurIPS*, 2020. 1
- [27] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *NeurIPS*, 2020. 1, 4, 8
- [28] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In *CVPR*, 2021. 8
- [29] 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. In *NeurIPS*, 2020. 1, 8
- [30] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P. Srinivasan, Jonathan T. Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. In *CVPR*, 2021. 1
- [31] George Toderici, Wenzhe Shi, Radu Timofte, Lucas Theis, Johannes Ballé, Eirikur Agustsson, Nick Johnston, and Fabian Mentzer. Workshop and challenge on learned image compression (CLIC2020), 2020. 4
- [32] Guillermo Valle-Pérez, Chico Q. Camargo, and Ard A. Louis. Deep learning generalizes because the parameter-function map is biased towards simple functions. In *ICLR*, 2019. 7

- [33] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Comput. Graph. Forum*, 2022. [1](#)
- [34] Greg Yang and Hadi Salman. A fine-grained spectral perspective on neural networks. *arXiv preprint 1907.10599*, 2019. [7](#)
- [35] L. P. Yaroslavsky. Compression, restoration, resampling, ‘compressive sensing’: Fast transforms in digital imaging. *Journal of Optics*, 2015. [5](#)
- [36] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. [8](#)
- [37] Gizen Yüce, Guillermo Ortiz-Jiménez, Beril Besbiar, and Pascal Frossard. A structured dictionary perspective on implicit neural representations. In *CVPR*, 2022. [1](#)
- [38] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017. [7](#)