

Sharp-NeRF: Grid-based Fast Deblurring Neural Radiance Fields using Sharpness Prior

Byeonghyeon Lee^{1*} Howoong Lee^{2,3*} Usman Ali^{2†} Eunbyung Park^{1,2†}

¹Department of Artificial Intelligence, Sungkyunkwan University

²Department of Electrical and Computer Engineering, Sungkyunkwan University

³Hanwha Vision

Abstract

Neural Radiance Fields (NeRF) has shown its remarkable performance in neural rendering-based novel view synthesis. However, NeRF suffers from severe visual quality degradation when the input images have been captured under imperfect conditions, such as poor illumination, defocus blurring and lens aberrations. Especially, defocus blur is quite common in the images when they are normally captured using cameras. Although few recent studies have proposed to render sharp images of considerably high-quality, yet they still face many key challenges. In particular, those methods have employed a Multi-Layer Perceptron (MLP) based NeRF which requires tremendous computational time. To overcome these shortcomings, this paper proposes a novel technique Sharp-NeRF—a grid-based NeRF that renders clean and sharp images from the input blurry images within a half an hour training. To do so, we used several grid-based kernels to accurately model the sharpness/blurriness of the scene. The sharpness level of the pixels is computed to learn the spatially varying blur kernels. We have conducted experiments on the benchmarks consisting of blurry images and have evaluated full-reference and non-reference metrics. The qualitative and quantitative results have revealed that our approach renders the sharp novel views with vivid colors and fine details, and it has considerably faster training time than the previous works. Our code is available at <https://github.com/benhenryL/SharpNeRF>.

1. Introduction

In the field of computer vision and graphics, with numerous applications in visual effects, e-commerce, AR/VR, and robotics, it is essential to model and reconstruct 3D scenes as representations that can synthesize high-quality images.

*Equal contributions

†Corresponding authors




Figure 1. Comparison in terms of training time and image quality on the real defocus dataset. Left: Evaluated under full-reference metric (PSNR). Right: Evaluated under no-reference metric (Niqe).

Recently, Neural Radiance Fields (NeRF) [21] has been proposed, and it has shown excellent performance in 3D scene representation and rendering photorealistic images for novel views. While successful, many existing NeRF methods assume that the training images are obtained in perfect condition, which bypasses vital issues that would arise in many practical settings. For example, training images captured with commodity hand-held devices inevitably contain a considerable amount of blur, e.g., motion or defocus blur [19, 39]. This violates 3D view consistency between different view images during training and deteriorates the rendered image quality.

A few approaches have been recently suggested to address the issue [16, 20, 28]. For example, Deblur-NeRF [20] proposed a deblurring neural radiance fields that renders clean images at novel views from blurry input images. It proposed to generate clean images first and apply convolution operation with the learnable blur kernels. At the time of inference, blur kernels are not utilized, and hence, it enables NeRF to produce sharp images. DP-NeRF [16] improved Deblur-NeRF by imposing 3D consistency to blur kernels and using depth information. Although promising results have been shown, they suffer from very long training and inference times due to the large number of neural network propagations to render images, please see Fig. 1.

In this work, we suggest Sharp-NeRF, a fast NeRF-based framework for rendering sharp images with blurred training images. Inspired by the recent approaches that incorporate additional data structures, such as voxel grids [3, 4, 9, 10], trees [41, 53] and hashes [24, 42], we leverage a decomposed-grid representation for neural fields architecture [4] to accelerate the training time and efficiently represent the colors and densities.

In addition, we propose *discrete learnable kernels* (without neural networks) for blurring convolution operations as opposed to the existing methods that utilize an extra MLP to generate the blur kernels. Using a neural network as a kernel-generating function introduces additional computational costs, and the bias induced by the neural network may prevent it from learning complex and sophisticated kernels that are necessary to handle various real-world scenarios. Furthermore, it has been known that the size of neural networks, e.g., width or depth, determines the complexity of the functions they can represent. Therefore, more optimal kernels require larger neural networks, which results in adding more computational burden during training.

As blurriness varies spatially across the scene, the ideal approach would involve retaining blur kernels for every individual pixel in every image. However, this becomes impractical due to the current prevalence of high-resolution images and multiple-view images. We introduce a sharpness prior to the proposed framework that enables us to maintain only a few number of discrete kernels. We propose to measure the per-pixel sharpness level and assign a blur kernel to a group of pixels whose sharpness level is similar. This strategy ensures that the blur kernels learn the spatially varying sharpness consistently and reliably with a marginal increase in learnable parameters. The overall architecture of our proposed method is shown in Fig. 2.

To optimize the training time further, we propose random patch sampling during training. As the blurriness is attributed to the intermingling of neighboring pixels, the existing deblurring methods render multiple neighboring pixels (usually 5 or 6) to render one pixel. It increases the training time proportionally to the number of rays per pixel. We develop a training strategy that renders a patch instead of a pixel and apply the convolution operation over the patch with a learned discrete kernel to produce the blurred patch. This strategy significantly reduces the training time while maintaining the rendered image quality.

We have tested the proposed approach on a challenging real-world dataset consisting of blurry training images. The experimental results have shown that Sharp-NeRF improves the training speed by a large margin and performs on par with the existing deblurring methods regarding visual quality. We also demonstrated that the proposed method outperforms the prior arts in more representative quantitative metrics, such as *Brisque* and *Niqe*, widely used in measur-

ing visual quality.

To sum up, our contributions are the following:

- We proposed novel learnable grid-based kernels to obtain sharp output from neural radiance fields. These kernels are optimized directly without requiring additional networks.
- For better and more reliable kernel optimization, a sharpness prior has been incorporated that explicitly measures the sharpness of pixels.
- We proposed to adopt random patch sampling to reduce the computational complexity of rendering.
- Novel view images rendered by our method are sharp, contain vivid colors and fine details, and have been achieved through the training time of a half an hour.

2. Related Works

2.1. Neural Radiance Fields

Recently, the Neural Radiance Fields (NeRF) has emerged as a powerful approach for synthesizing high-fidelity 3D scenes from 2D images. NeRF leverages deep neural networks to represent volumetric scene properties and enables realistic rendering from novel viewpoints. A radiance field is a continuous function f that maps a 3D location $x \in \mathbb{R}^3$ and a viewing direction $d \in \mathbb{S}^2$ to estimate density $\sigma \in [0, \infty)$ and a color value $c \in [0, 1]^3$ of the given point. A multi-layer perceptron (MLP) has been used by [21] to parameterize this function, where the weights of MLP are optimized to reconstruct a set of input images of a certain scene:

$$f_{\theta} : (\gamma(x), \gamma(d)) \rightarrow (c, \sigma) . \quad (1)$$

Here, θ denotes the network weights, and γ is a predefined positional encoding [40] applied to x and d . Given the volume density and color of points, volume rendering [7] is used to produce the images at novel views [21].

With the success of NeRF, many follow-up studies have been conducted that improve the one or multiple modules of NeRF for improved quality of rendering. Several approaches, e.g., grid-based F2-NeRF [42] and TensorRF [4], aim to improve the rendering speed [10, 12–14, 24, 31, 38, 53]. Few others, e.g., FreeNeRF [51], PixelNeRF [54] and MVSNeRF [5], try to predict a continuous neural scene representation conditioned on one or a few input images and SNeRG [13], MeRF [32], and BakedSDF [52] bake model to achieve real-time view synthesis.

Learning-based deblurring networks for NeRF. Some studies in the literature have explored a learning-based approach to directly train the neural networks to deblur the NeRF-based rendered images. These methods leverage a

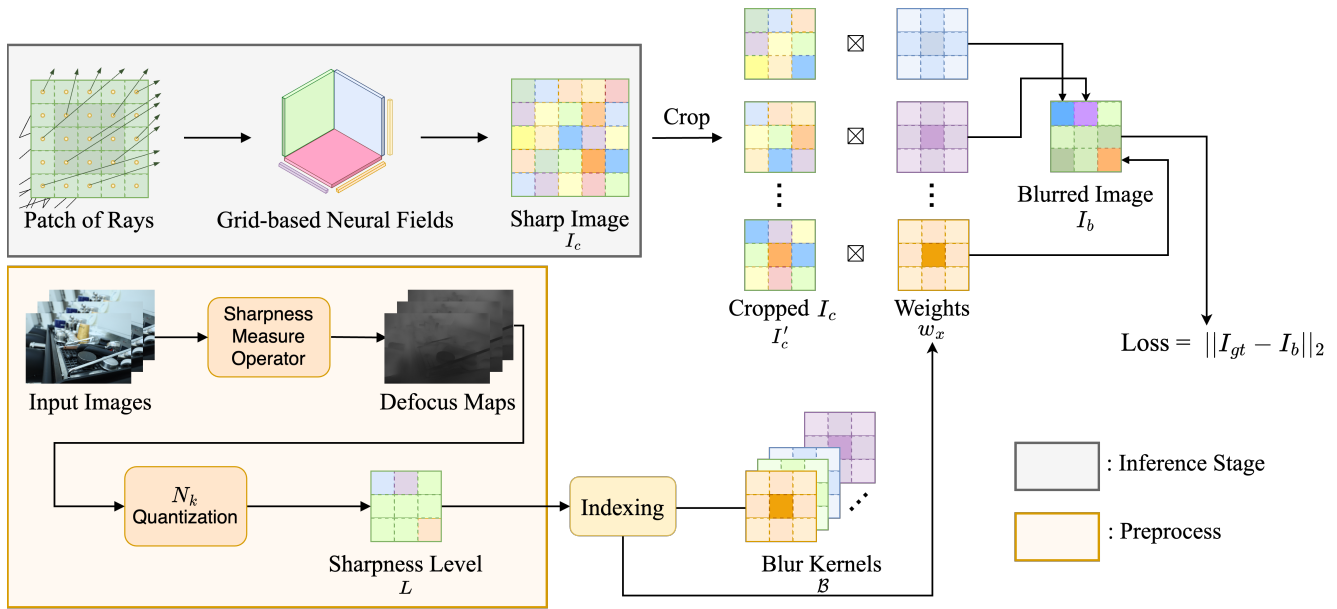


Figure 2. The overall architecture of Sharp-NeRF. \otimes stands for weighted sum. First, it computes defocus map of each training view using sharpness measure operator. Then the defocus map is quantized into N_k values which is used as per-pixel sharpness level map L . This is a preprocess and is not required to compute L during process. During training, it takes ray patches as inputs to backbone neural fields model [4] and render sharp and clean image I_c . I_c is then cropped into several small patches I'_c with stride 1 where size of each patch is $K \times K$. Preprocessed L of the input ray patch is also given as input data and used to indexing blur kernels B to obtain per-pixel weight w_x . Subsequently, w_x and I'_c are weighted sum to render blurred image I_b . Note that at the time of inference, only modules in gray box are valid which means blur kernel is no longer required and only I_c is used as a final rendered outcome.

pair of blurred and clear image datasets to learn the mapping between the blurred rendering and the corresponding sharp image. However, many of these tasks use MLP-based architectures. The Deblur-NeRF [20] and DP-NeRF [16] take a long time to train a MLP-based neural fields and kernel. The Hybrid Neural Rendering Model [6] has a grid kernel, but it is fixed and applies only to the camera motion blur [43]. In the defocus blurred scene, point spread function (PSF) is spatially varying and cannot be applied straightforward. DoF-NeRF [48] has a limitation that all-in-focus images are required for training. Although PDRF [28] is similar to our model in that it used grid-based representation, it still uses two MLPs; first to predict the scene density and the second to build blur kernels. In contrast, we propose a well-curated approach Sharp-NeRF—a grid-based model that can render sharp novel images from blurry input images within half an hour of training.

2.2. Image Blind Deblurring

It is common to observe that some parts of the pictures are blurred when we normally capture them using optical imaging devices. Among a number of factors that cause this blurriness are lens defocusing, camera shake, and fast motion [1, 35]. The out-of-focus or defocus blur occurs in the imaging process when the image plane is away from

the ideal reference plane. Mathematically, the deterioration of an image due to defocus blur is typically represented by $B = I \otimes K + \epsilon$, where I is the latent sharp image, B is the observed blurry image, K is the blur kernel, \otimes is the convolution operator, and ϵ is the additive white Gaussian noise that often appears in natural images. The earliest canonical method that deconvolves an image is [34] where a known PSF—the blur kernel—is used to iteratively minimize an energy function to get a maximum likelihood approximation of the original image. Usually, both the sharp image I and blurring kernel K are unknown. Recovering the latent sharp image given only a single blurry image B is known as image blind deblurring [46]. This is a highly ill-posed and long standing problem in the image and vision community. In this process, first a blurring kernel is estimated which is then used to derive the deblurred image. A variety of methods have been proposed in the literature to tackle the blurring problem. The traditional approaches rely on natural image priors and often formulate deblurring as an optimization problem [19, 27, 49, 57]. On the other hand, most of the deep learning based methods directly map the blurry image with the latent sharp image by employing convolutional neural networks (CNN) [26, 33, 56]. However, most of the learning-based approaches disregard the underlying imaging process, and hence they perform less effectively in

the case of severe blur.

2.3. Sharpness Measurement

Computing the sharpness level of image pixels is a crucial step for many imaging applications, including the object detection and shape from focus (SFF) system [2,47,50]. In the case of partially out-of-focus image, focused areas typically exhibit greater intensity variations than defocused areas. Any operator can be used as a sharpness measure as long as it can distinguish between sharp (focused) and blurred (defocused) pixels in a partially blurred (defocused) image. The sharpness measurements or levels of image pixels are obtained by convolving the image with the sharpness measure. A detailed analysis of sharpness measures for focus computation has been provided in [11,29]. Among different categories of sharpness measures, derivative-based measures appear to be an interesting option because of their inherent ability to capture the rate of change. A well-known and commonly used focus measure operator is modified laplacian (ML) [25]. Inspired by the effectiveness of sharpness measures, we propose to leverage a learning-based sharpness prior in our deblurring approach for NeRF. Our analysis has shown that enforcing the sharpness prior is effective in rendering better deblurred images.

3. Method

3.1. Preliminary: Tensorial Radiance Fields

We adopted Tensorial Radiance Fields (TensorRF) [4] as our backbone grid-based neural radiance fields. It encodes and reconstructs the scene using two 3D grids G_σ and G_c that estimate density and color respectively, which are further decomposed to several low-rank 2D and 1D tensors as following:

$$G_\sigma = \sum_{i=0}^{N_x} v_i^x \circ M_i^{y,z} + \sum_{i=0}^{N_y} v_i^y \circ M_i^{x,z} + \sum_{i=0}^{N_z} v_i^z \circ M_i^{x,y}, \quad (2)$$

where N_x is the number of the decomposed tensors, v^x is 1D vector along the x-axis, $M^{y,z}$ is 2D matrix along the y- and z-axis. The grid for color estimation G_c is defined identically to G_σ . TensorRF takes input coordinates $x \in \mathbb{R}^3$ and viewing directions $d \in \mathbb{S}^2$, and produces 4D vectors comprising of density and RGB values, such as,

$$\sigma(r) = \zeta(G_\sigma(x)), \quad (3)$$

$$c(r) = f_\theta(G_c(x), d), \quad (4)$$

where θ denotes the parameters of shallow MLP for color rendering and ζ denotes softplus activation function [8]. Then these $\sigma(r)$ and $C(r)$ are finally used to compute pixel

intensity using volume rendering [7]:

$$C(r) = \sum_{q=1}^Q T_q (1 - \exp(-\sigma(r)_q \Delta_q)) c(r)_q, \quad (5)$$

$$T_q = \exp(-\sum_{p=1}^{q-1} \sigma(r)_p \Delta_p), \quad (6)$$

where $C(r)$ is the approximated intensity of ray r , q is sampled points, and Δ_p is the distance between pixels.

3.2. Patch Sampling

Random ray sampling has been widely used in NeRF literature [18,21,45]. This approach selects rays arbitrarily across the scene, where it is highly unlikely that it will sample the rays belonging to a local region simultaneously. However, as mentioned earlier, defocus blur occurs due to the intermingling of neighboring pixels which happens due to the shallow depth of field of the camera. In order to model such interference of neighboring pixels, peripheral region of the interesting pixel must be rendered as well. Therefore, instead of conventional random ray sampling, we adopted random patch sampling that samples a group of rays within a close location, namely a patch. However, rendering the neighboring pixels increases the training time, and this increase is inevitable in deblurring neural fields literature. Naively reducing the number of rays to minimize the training time rather degrades the reconstruction quality. As the neural fields lack the notion of 3D consistency, we adopted the patch sampling approach to model the blurring phenomenon by learning the affinity among the neighboring pixels.

Fig. 3 suggests that $N \times K \times K$ pixels should be rendered to produce blurred intensity values of N pixels, where N is the number of target pixels and $K \times K$ is the size of blur kernel. Interestingly, under random patch sampling, target pixels can share their neighboring pixels within a patch so that the number of neighboring pixels is reduced. Such that, only $(P' + K + 1)^2$ pixels need to be rendered to produce the same number of target pixels, where P' is the width and height of region to compute intensity. The number of sampled rays is no longer dependent on N but P'^2 , so random patch sampling can effectively reduce the training time by shrinking the amount of ray rendering. This approach saves a lot of time, in contrast to NeRF that fetches thousands of rays per training iteration where $N \gg P'^2$. With random patch sampling, we have reduced the number of required neighboring rays per ray from K^2 to $\frac{(P'+K+1)^2}{P'^2}$.

3.3. Sharpness Prior

We propose to leverage a sharpness prior to accurately model the deblurring phenomenon for sharp NeRF. By applying this sharpness prior, sharpness level of image pixels

is determined. This sharpness information plays a key role in learning the suitable blur kernels for those pixels. We used the defocus map estimation network (DMENet) [17], which is the first end-to-end CNN framework to directly estimate spatially varying defocus map from a given defocus image. We estimate defocus map of every training view and quantize it uniformly into N_k values. Additionally, the sharpness prior is computed only once as part of the preprocessing, taking less than a couple of minutes, and it does not affect the subsequent training and testing times. Quantized defocus map is used as a sharpness level map L which divides pixels into N_k groups based on their degree of sharpness. Here, it is worth mentioning that our proposed method is not specific to applying DMENet [17], but is generic. That is, any operator from the literature can be utilized as a sharpness prior as long as it can reliably measure the sharpness level of image pixels.

The original NeRF and many other derived works [18, 21, 45] treated the image pixels and their corresponding rays independently. This approach considers that the colors/intensities of the pixels/points have association only along the ray, i.e., along the principal axis. This approach ignores any transversal relation among neighboring points or rays during volume rendering. Therefore, these works cannot mitigate the defocus effects and result in blurry rendered pixels. In contrast, we propose to exploit the transversal relation among neighboring rays as well. Specifically, we compute the sharpness level of points in their neighborhood. Furthermore, we leverage depth information to improve the extraction of regions with sharp focus. Initially, we segment the scene based on depth information and assess the level of focus within each segment. The calculation of depth information is accomplished through the application of the MIDAS [30].

3.4. Grid-based Blur Kernel

We propose multiple trainable grid-based blur kernels $\mathcal{B} \in \mathbb{R}^{N_{img} \times N_k \times K \times K \times C}$ where N_{img} is the number of training views, N_k is the number of blur kernels, K is kernel size, and C is the number of channels. We will omit channel C for brevity from now on. Our grid-based blur kernels can be directly optimized to model the spatially-varying blurriness of the given scene, in contrast to the previous works [16, 20, 28] which employed MLP to generate blur kernels. Learnable grid-based kernels allow to obtain blur kernels w_x of each pixel through simple indexing:

$$w_x = F(\mathcal{B}, L, x), \quad (7)$$

where F is a mapping function that maps each pixel to corresponding blur kernel $w_x \in \mathbb{R}^{P' \times P' \times K \times K}$ based on sharpness level L . Then we crop the clean image I_c along to height and width with stride 1 and get cropped image $I'_c \in \mathbb{R}^{P' \times P' \times K \times K}$. Using the obtained w_x and I'_c , blurred

image can be convolved from clean image as following:

$$I_b[h, w] = \sum_{i=0}^K \sum_{j=0}^K I'_c[h, w, i, j] w_x[h, w, i, j], \quad (8)$$

where I_b is a blurred image with the shape of $P' \times P'$, $[\cdot]$ is an indexing operation, and h, w represents index along the height and width of image and blur kernels respectively. Both blur kernels \mathcal{B} and neural fields are optimized via loss \mathcal{L}_{recon} in Eq. (9) which is an MSE loss within I_b and ground truth image I_{gt} :

$$\mathcal{L}_{recon} = \|I_b - I_{gt}\|_2. \quad (9)$$

Using learnable grid-based blur kernels, we can easily improve expressivity of the blur kernels by simply increasing the size of grid which does not require extra cost to produce kernels and successfully rendered clean image even from harshly defocus region.

Instead of optimizing per pixel blur kernels, we train N_k blur kernels which is identical to the number of the unique groups of L and pixels in a same group share a blur kernel, for effective optimization. It is worth to note that applying appropriate blur kernels to appropriate regions of the images plays pivotal role to estimate the scene blurriness because different regions in various depth of fields have varying degrees of blurriness. As pixels are blurred because of the interference of the neighboring pixels, the value of blur kernels tend to be spread widely to estimate the interruption of wide range of peripheral pixels if given pixels are severely blurry, while they tend to focus on the central part of neighboring pixels if given pixels are sharp. Fig. 4 shows different form of blur kernels depending on the blurriness of their interesting regions. Therefore, each blur kernel can estimate blurriness correctly only if it is assigned to right region where the degree of blurriness is similar within pixels. Table 3 shows the quantitative results on comparing building L with sharpness prior and other approaches without sharpness prior.

4. Experiments

4.1. Implementation Details

We used TensorRF [4] as the baseline and build a grid-based learnable kernel. Grid-based blur kernels are jointly optimized with the grid-based neural fields, starting from the first training iteration, as the grid-based neural fields converge rapidly. We set the size of each blur kernel to 11, learning rate of the kernel to 0.05 and initialize with 2D Gaussian distribution. We use 400 for N_k and skip convolution for pixels in 100 groups of each view with high sharpness value. For random patch sampling, we sample 28 patches with the patch size $P = 22$, so the number of




Figure 3. Left: random ray sampling. Right: random patch sampling. Blue pixels are $P' \times P'$ interesting pixels to be rendered and skyblue pixels are required neighboring pixels for blur convolution.




Figure 4. Visualization of blur kernels. From left to right, the values of kernels are widely spread which implies that leftmost kernel is responsible for sharp region and rightmost kernel is responsible for blurry region.

the interesting pixels P'^2 is 12^2 . We reduce the grid resolution from 640^3 to 480^3 and raise the number of grid component from 96 to 132. We use the learnable camera response function proposed by [20]. All the experiments were conducted on a single NVIDIA A100 GPU and please refer TensorRF [4] for more details.

4.2. Evaluation Metrics

Usually during the image acquisition and processing, distortions can cause an image’s quality to decrease. Examples of distortion include noise, blurring, ringing, and compression artifacts. Efforts have been made to develop impartial metrics of quality. A useful quality metric corresponds favorably with the subjective assessment of quality made by a human observer in many situations.

Full-reference quality metrics: If a distortion-free image is available, it can be used as a reference to measure the quality of other restored/reconstructed images. In such cases, full-reference quality metrics can be used to compare the target image with the reference image. We have used peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [44].

No-reference quality metrics: If distortion-free reference image is not available, an alternative is to employ a no-reference image quality metric. These metrics compute the image quality based on image statistics. When compared to full-reference metrics, all no-reference quality metrics typically perform better in terms of agreement with a subjective human quality score. We have used Blind/referenceless

image spatial quality evaluator (Brisque) [22], and Natural image quality evaluator (Niqe) [23]. All of these quality metrics have been defined in the supplementary material.

4.3. Dataset

We conducted experiments on the scenes which contain real and synthetic defocus blurred images of the Deblur-NeRF dataset [20]. There are ten scenes in the real defocus category and five in synthetic defocus category of this dataset. They utilized the in-built skills of Blender to produce depth-of-field images in order to reduce defocus blur. Realistic defocus blur effects were obtained by establishing the aperture size and arbitrarily choosing a focal plane between the nearest and farthest depths. They select a large aperture to capture the defocus image and use COLMAP [36, 37] to compute the camera pose of the blur and reference image in the real scene.

4.4. Results

In this section, we present full reports of experiments with both qualitative and quantitative results. We compare the proposed method to existing deblurring neural fields models, Deblur-NeRF [4], DP-NeRF [16], and PDRF [28]. We also compare ours with TensorRF [4] trained with deblurred image using Restormer [55]. Restormer is a modern Transformer-based model that showed noticeable image restoration performance. Table 1 shows the quantitative results on real defocus dataset evaluated using full-reference quality metrics PSNR and SSIM. We also evaluated our model using no-reference quality metrics, Brisque and Niqe as shown in Table 2. Sharp-NeRF achieves comparable results to other MLP-based models consistently under various image quality evaluation metrics and even attains state-of-the-art performance under Niqe metric while reducing training time to less than a half hour which is the shortest training time among all deblurring neural fields models. Due to the page limit, results on synthetic data and other experiments will be covered in supplementary materials.

4.5. Ablation Studies

Sharpness Prior. Firstly, we conducted an ablation study on the selection of focus measure operator. We compared not only the hand-crafted operators like sum modified laplacian (SML) [25] and Tenengrad [15], but also deep neural network-based model DMENet [17]. The results are presented in Table 3. The better values of PSNR and SSIM for DMENet [17] indicate that this sharpness prior is more effective in accurately measuring the sharpness levels of image pixels.

Sharpness Level. Next, we conducted experiments on building sharpness level map L . It can be observed in Table 4 that constructing L using sharpness measurement is a major factor that considerably affects the rendering quality.

Table 1. Quantitative results on real defocus dataset under Full-reference metrics.

	Cake		Caps		Cisco		Coral		Cupcake				
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑			
NeRF [21]	24.42	0.72	22.73	0.63	20.72	0.72	19.81	0.56	21.88	0.68			
TensoRF [4]	24.08	0.70	21.55	0.53	20.56	0.70	19.27	0.54	21.16	0.65			
Restormer [55]+TensoRF	23.96	0.68	21.54	0.51	20.28	0.69	19.15	0.52	21.37	0.65			
Deblur-NeRF [20]	26.27	0.78	23.87	0.71	20.83	0.72	19.85	0.59	22.26	0.72			
DP-NeRF [16]	26.16	0.77	23.95	0.71	20.73	0.72	20.11	0.61	22.80	0.74			
PDRF-5 [28]	27.03	0.79	24.29	0.72	20.67	0.72	19.65	0.58	22.94	0.74			
PDRF-10 [28]	27.06	0.80	24.06	0.71	20.68	0.72	19.61	0.58	22.95	0.74			
Ours	26.23	0.77	23.98	0.70	20.88	0.72	20.07	0.59	22.75	0.73			
	Cups		Daisy		Sausage		Seal		Tools		Average		Time
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	Hours ↓
NeRF [21]	25.02	0.75	22.74	0.62	17.79	0.48	22.79	0.62	26.08	0.85	22.40	0.66	2.30
TensoRF [4]	23.56	0.70	22.74	0.63	17.04	0.44	21.19	0.55	21.75	0.74	21.29	0.62	0.39
Restormer [55]+TensoRF	23.03	0.68	22.53	0.62	16.59	0.42	20.52	0.52	23.78	0.79	21.28	0.61	0.39
Deblur-NeRF [20]	26.21	0.79	23.52	0.68	18.01	0.49	26.04	0.77	27.81	0.89	23.46	0.71	10.35
DP-NeRF [16]	26.75	0.81	23.79	0.69	18.35	0.54	25.95	0.77	28.07	0.89	23.67	0.72	20.20
PDRF-5 [28]	26.37	0.80	24.30	0.73	18.82	0.55	26.12	0.78	28.00	0.89	23.82	0.73	1.00
PDRF-10 [28]	26.39	0.80	24.49	0.74	18.94	0.56	26.36	0.79	28.00	0.89	23.85	0.73	1.67
Ours	25.34	0.77	23.66	0.70	18.77	0.54	25.82	0.77	27.98	0.90	23.55	0.72	0.48

Table 2. Quantitative results on real defocus dataset under No-reference metrics

	Cake		Caps		Cisco		Coral		Cupcake			
	Brisque↓	Niqe↓	Brisque↓	Niqe↓	Brisque↓	Niqe↓	Brisque↓	Niqe↓	Brisque↓	Niqe↓		
TensoRF [4]	30.27	3.11	33.90	3.08	35.03	3.83	36.15	3.45	38.24	3.82		
Deblur-NeRF [20]	25.78	3.23	28.8	3.33	31.89	3.86	39.18	4.43	33.68	3.51		
DP-NeRF [16]	24.15	3.24	26.52	3.34	30.22	3.68	36.77	4.37	32.43	3.46		
PDRF-5 [28]	25.00	3.60	37.37	3.58	30.66	3.97	35.75	4.13	32.17	3.71		
PDRF-10 [28]	21.93	3.50	30.46	3.40	30.39	3.84	38.07	3.99	31.52	3.59		
Ours	27.54	3.26	31.66	3.22	28.56	3.63	29.43	3.59	31.05	3.49		
	Cups		Daisy		Sausage		Seal		Tools		Average	
	Brisque↓	Niqe↓	Brisque↓	Niqe↓	Brisque↓	Niqe↓	Brisque↓	Niqe↓	Brisque↓	Niqe↓	Brisque↓	Niqe↓
TensoRF [4]	37.41	3.58	36.79	4.07	36.86	3.56	34.72	3.18	39.65	3.49	35.90	3.52
Deblur-NeRF [20]	32.65	3.40	29.10	4.11	28.63	3.37	32.09	3.46	31.44	3.14	31.32	3.58
DP-NeRF [16]	31.00	3.33	30.01	4.09	30.17	3.38	31.82	3.29	29.86	3.05	30.29	3.52
PDRF-5 [28]	33.75	3.60	29.18	4.37	29.38	3.60	32.78	3.45	35.73	3.44	32.18	3.75
PDRF-10 [28]	29.09	3.46	28.40	4.24	26.67	3.49	29.31	3.50	34.28	3.29	30.01	3.63
Ours	33.70	3.29	26.35	3.85	28.51	3.45	32.63	3.19	33.10	3.13	30.25	3.41

Table 3. Ablation study on selection of the sharpness prior. Average values for the real defocus testset have been mentioned.

Methods	PSNR ↑	SSIM ↑
SML [25]	22.56	0.67
Tenengrad [15]	22.49	0.66
DMENet [17]	23.55	0.72

Table 4. Ablation study on constructing sharpness level map. Randomly assigning (Random) and allocating pixels into the groups based on their position (Local) show degraded performance.

Methods	PSNR ↑	SSIM ↑
Random	22.09	0.66
Local	21.43	0.63
Ours	23.55	0.72

Table 5. Comparison on kernel generation time.

Methods	Time ↓
MLP-based kernels	2.20 ms
Ours	0.18 ms

Kernel Generation Time. Finally, we investigated the time taken by the MLP-based and grid-based methods in generating the blur kernels. To do so, we computed the time taken by the recent MLP-based blur kernels of PDRF [28] and by our method. These times are computed for real defocus dataset, and their average values are mentioned in Table 5. The considerably reduced time of our method indicates its superiority in terms of computational efficiency.




Figure 5. Qualitative results on real defocus dataset. Our proposed method renders sharp images which have vivid colors and fine details.

5. Limitations & Future Works

Although our approach, involving learnable grid-based blur kernels, has produced appreciably good quality results, yet it has limitation. For example, our method produces degraded results for motion blurred images. This is because our approach relies on the sharpness measurement of the input images, and we have incorporated a sharpness prior only for the defocus blur, and not for the motion blur. In other words, when the input images have motion blur, our sharpness prior fails in accurately estimating the sharpness of the pixels, which ultimately leads to the poor quality of rendered images. However, it is worth mentioning that this limitation is due to the design methodology of our approach. For instance, instead of defocus blur, our approach can be adjusted and applied to motion blurred scenes by incorporating a sharpness prior designed for the motion blurred images. We anticipate that estimation of the camera motion blur can be achieved by computing the orientation and magnitude of blurriness in the pixels using higher order derivatives, and we will explore it in our future work. Further, it seems interesting that whether few other priors, or group of priors, can be incorporated in our method to produce good quality results for other image degrading situations, like haze, raining, and under water imaging. In future, we will learn the blur kernels in 3D and we will cover all types of blur. Further, we will seek any gain that can be achieved

by learning the kernels through joint optimization with the rest of the network.

6. Conclusion

This paper proposes first fully grid-based deblurring neural fields—Sharp-NeRF. In the literature, it has been observed that grid-based neural fields suffer from quality degradation when rendering under imperfect conditions such as defocused images [28]. To overcome this shortcoming, in this work, we proposed novel learnable grid-based blur kernels. This approach can produce blur kernels with simple indexing and sharpness measurement to promote blur kernel optimization. Coupled with grid-based kernel and random patch sampling, we could achieve state-of-the-art training time, taking only a half an hour to train.

7. Acknowledgements

This work was supported in parts by the Institute of Information and Communication Technology Planning Evaluation (IITP) grant (IITP-2019-0-00421), Creative Challenge Research Program (2021R1I1A1A01052521) through the National Research Foundation (NRF) of Korea, and the NRF grant (RS-2023-00245342) funded by the Ministry of Science and ICT (MSIT) of Korea.

References

- [1] Abdullah Abuolaim, Mahmoud Afifi, and Michael S Brown. Improving single-image defocus deblurring: How dual-pixel images help through multi-task learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1231–1239, 2022. 3
- [2] Usman Ali and Muhammad Tariq Mahmood. Robust focus volume regularization in shape from focus. *IEEE Transactions on Image Processing*, 30:7215–7227, 2021. 4
- [3] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. *CVPR*, 2023. 2
- [4] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, pages 333–350. Springer, 2022. 2, 3, 4, 5, 6, 7
- [5] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. *arXiv preprint arXiv:2103.15595*, 2021. 2
- [6] Peng Dai, Yinda Zhang, Xin Yu, Xiaoyang Lyu, and Xiaojuan Qi. Hybrid neural rendering for large-scale scenes with motion blur. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 3
- [7] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. *SIGGRAPH Comput. Graph.*, 22(4):65–74, jun 1988. 2, 4
- [8] Charles Dugas, Yoshua Bengio, François Bélisle, Claude Nadeau, and René Garcia. Incorporating second-order functional knowledge for better option pricing. *Advances in neural information processing systems*, 13, 2000. 4
- [9] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *CVPR*, 2023. 2
- [10] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 2
- [11] Guillermo Gallego, Mathias Gehrig, and Davide Scaramuzza. Focus is all you need: Loss functions for event-based vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12280–12289, 2019. 4
- [12] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021. 2
- [13] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. *ICCV*, 2021. 2
- [14] Tao Hu, Shu Liu, Yilun Chen, Tiancheng Shen, and Jiaya Jia. Efficientnerf efficient neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12902–12911, 2022. 2
- [15] Eric Krotkov. Focusing. *International Journal of Computer Vision*, 1(3):223–237, 1988. 6, 7
- [16] Dogyoon Lee, Minhyeok Lee, Chajin Shin, and Sangyoun Lee. Deblurred neural radiance field with physical scene priors. *arXiv preprint arXiv:2211.12046*, 2022. 1, 3, 5, 6, 7
- [17] Junyong Lee, Sungkil Lee, Sunghyun Cho, and Seungyong Lee. Deep defocus map estimation using domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5, 6, 7
- [18] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5741–5751, 2021. 4, 5
- [19] Yu-Qi Liu, Xin Du, Hui-Liang Shen, and Shu-Jie Chen. Estimating generalized gaussian blur kernels for out-of-focus image deblurring. *IEEE Transactions on circuits and systems for video technology*, 31(3):829–843, 2020. 1, 3
- [20] Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12861–12870, 2022. 1, 3, 5, 6, 7
- [21] 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, 2, 4, 5, 7
- [22] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12):4695–4708, 2012. 6
- [23] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012. 6
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 2
- [25] Shree K Nayar and Yasuo Nakagawa. Shape from focus. *IEEE Transactions on Pattern analysis and machine intelligence*, 16(8):824–831, 1994. 4, 6, 7
- [26] Thekke Madam Nimisha, Akash Kumar Singh, and Ambasamudram N Rajagopalan. Blur-invariant deep learning for blind-deblurring. In *Proceedings of the IEEE international conference on computer vision*, pages 4752–4760, 2017. 3
- [27] Jinshan Pan, Deqing Sun, Hanspeter Pfister, and Ming-Hsuan Yang. Blind image deblurring using dark channel prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1628–1636, 2016. 3
- [28] Cheng Peng and Rama Chellappa. Pdrf: Progressively deblurring radiance field for fast and robust scene reconstruction from blurry images, 2022. 1, 3, 5, 6, 7, 8
- [29] Said Pertuz, Domenec Puig, and Miguel Angel Garcia. Analysis of focus measure operators for shape-from-focus. *Pattern Recognition*, 46(5):1415–1432, 2013. 4
- [30] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *ICCV*, 2021. 5

- [31] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14335–14345, 2021. 2
- [32] Christian Reiser, Richard Szeliski, Dor Verbin, Pratul P. Srinivasan, Ben Mildenhall, Andreas Geiger, Jonathan T. Barron, and Peter Hedman. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *SIGGRAPH*, 2023. 2
- [33] Dongwei Ren, Kai Zhang, Qilong Wang, Qinghua Hu, and Wangmeng Zuo. Neural blind deconvolution using deep priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3341–3350, 2020. 3
- [34] William Hadley Richardson. Bayesian-based iterative method of image restoration. *JoSA*, 62(1):55–59, 1972. 3
- [35] Lingyan Ruan, Bin Chen, Jizhou Li, and Miuling Lam. Learning to deblur using light field generated and real defocus images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16304–16313, 2022. 3
- [36] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [37] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 6
- [38] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 2
- [39] Supasorn Suwajanakorn, Carlos Hernandez, and Steven M Seitz. Depth from focus with your mobile phone. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3497–3506, 2015. 1
- [40] 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
- [41] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yan-shun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenoctrees for dynamic radiance field rendering in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13524–13534, June 2022. 2
- [42] Peng Wang, Yuan Liu, Zhaoxi Chen, Lingjie Liu, Ziwei Liu, Taku Komura, Christian Theobalt, and Wenping Wang. F2-nerf: Fast neural radiance field training with free camera trajectories. *CVPR*, 2023. 2
- [43] Peng Wang, Lingzhe Zhao, Ruijie Ma, and Peidong Liu. Bad-nerf: Bundle adjusted deblur neural radiance fields. *arXiv preprint arXiv:2211.12853*, 2022. 3
- [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [45] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 4, 5
- [46] Oliver Whyte, Josef Sivic, Andrew Zisserman, and Jean Ponce. Non-uniform deblurring for shaken images. *International journal of computer vision*, 98:168–186, 2012. 3
- [47] Changyeon Won and Hae-Gon Jeon. Learning depth from focus in the wild. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part I*, pages 1–18. Springer, 2022. 4
- [48] Zijin Wu, Xingyi Li, Juewen Peng, Hao Lu, Zhiguo Cao, and Weicai Zhong. Dof-nerf: Depth-of-field meets neural radiance fields. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1718–1729, 2022. 3
- [49] Li Xu, Shicheng Zheng, and Jiaya Jia. Unnatural 10 sparse representation for natural image deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1107–1114, 2013. 3
- [50] Fengting Yang, Xiaolei Huang, and Zihan Zhou. Deep depth from focus with differential focus volume. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12642–12651, 2022. 4
- [51] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [52] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P. Srinivasan, Richard Szeliski, Jonathan T. Barron, and Ben Mildenhall. Bakedsd: Meshing neural sdf for real-time view synthesis. *arXiv*, 2023. 2
- [53] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenoctrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. 2
- [54] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. 2
- [55] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 2022. 6, 7
- [56] Hongguang Zhang, Yuchao Dai, Hongdong Li, and Piotr Koniusz. Deep stacked hierarchical multi-patch network for image deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5978–5986, 2019. 3
- [57] Meina Zhang, Yingying Fang, Guoxi Ni, and Tiejiong Zeng. Pixel screening based intermediate correction for blind deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5892–5900, 2022. 3