

# Waterf: Robust Watermarks in Radiance Fields for Protection of Copyrights

Youngdong Jang<sup>1</sup> Dong In Lee<sup>1</sup> MinHyuk Jang<sup>1</sup> Jong Wook Kim<sup>1</sup>  
 Feng Yang<sup>2</sup> Sangpil Kim<sup>1\*</sup>

<sup>1</sup>Korea University <sup>2</sup>Google Research

## Abstract

The advances in the Neural Radiance Fields (NeRF) research offer extensive applications in diverse domains, but protecting their copyrights has not yet been researched in depth. Recently, NeRF watermarking has been considered one of the pivotal solutions for safely deploying NeRF-based 3D representations. However, existing methods are designed to apply only to implicit or explicit NeRF representations. In this work, we introduce an innovative watermarking method that can be employed in both representations of NeRF. This is achieved by fine-tuning NeRF to embed binary messages in the rendering process. In detail, we propose utilizing the discrete wavelet transform in the NeRF space for watermarking. Furthermore, we adopt a deferred back-propagation technique and introduce a combination with the patch-wise loss to improve rendering quality and bit accuracy with minimum trade-offs. We evaluate our method in three different aspects: capacity, invisibility, and robustness of the embedded watermarks in the 2D-rendered images. Our method achieves state-of-the-art performance with faster training speed over the compared state-of-the-art methods. Project page: <https://kuai-lab.github.io/cvpr2024waterf/>

## 1. Introduction

Digital watermarking plays a pivotal role in reinforcing the copyright of digital assets, e.g., text, image, audio, video, and 3D content. Digital assets are easily misused without permission from the creators and owners of the digital assets. One way to protect the copyright is by encoding invisible watermarks in the digital assets.

Neural Radiance Fields (NeRF) [35] have emerged into the spotlight in 3D content creation and modeling since NeRF can represent 3D objects or scenes in a compact way. With increasing interest in 3D content for the Metaverse, virtual reality, and augmented reality, NeRF has become important in digital watermarking research.

A straightforward approach for protecting rendered 3D

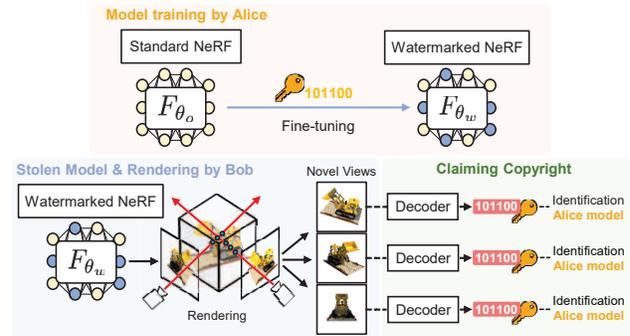


Figure 1. The NeRF model can be fine-tuned to embed the watermark into the images of the novel view. The figure above shows the model owner Alice and her NeRF model is stolen by Bob. Alice can identify the copyright by using her watermark decoder. Our method embeds the watermark into all of the rendered images.

images from the NeRF model is using the post-generation watermarking method that embeds the watermark into the 2D rendered image from NeRF with existing watermarking methods. However, this method only protects copyright for the rendered images but not the copyright of the NeRF model. A fundamental solution to protecting the model and rendered image is fine-tuning the NeRF model to encode the watermark into the NeRF model itself. The watermark will be encoded into the rendered image from the fine-tuned NeRF every time during the rendering process.

NeRF has two primary representations: implicit and explicit representations. The implicit representation [1, 35, 37] uses a multi-layer perceptron (MLP) to represent a 3D scene and the explicit representation [3, 10, 58] uses a traditional 3D structure, such as voxels. Prior methods [22, 30] were limited to applying to only one type of NeRF representation. However, in order to overcome this limitation, we explore a new question: *How can we embed the watermark into two distinct NeRF representations(implicit and explicit)?*

In this paper, we propose an innovative watermarking method that can be applied to both types of NeRF representations. Our method integrates the watermarking process into the rendering process without changing the model architecture. It adjusts the pre-trained NeRF model so that all

\*Corresponding author.

of the rendered images have same watermark embedded. We pre-train the decoder with the conventional deep learning method HiDDeN [59] to extract the watermark from the rendered image. Our method has several advantages. First, since it is a fine-tuning process, we can deal with the implicit and the explicit NeRF representations. Second, we do not need any extra processes for concealing watermarks in the rendered images. Therefore, less computational cost is required. Third, our method can protect the model and the rendered images simultaneously.

Invisibility and robustness are very important aspects in the digital watermarking domain. Concealing watermarks in the rendered image without degrading the original image quality and ensuring identification of the watermark is challenging. For instance, an embedded watermark in the image is better to be as invisible as possible, and also watermark should be identified after various distortion attacks. We intend to hide the watermark message in the low-level subband in the frequency domain by carefully designing the loss function that evaluates the loss value in the frequency domain. To enhance the image quality, we consider the local structure of the object by introducing patch-wise loss. We subdivided the rendered image and cached the gradient in the grid to encode the watermark locally in the image.

Our extensive experiments demonstrate that our method successfully encodes watermarks by fine-tuning the pre-trained NeRF models and identifying the watermark message using a simplified decoder from HiDDeN [59]. Moreover, we show that our method is robust from diverse watermark attacks. We evaluate our method by measuring the capacity and robustness under diverse watermark attacks. The training time of our method is approximately six times faster than the CopyRNeRF [30]. Since our method can be applied to both implicit [35] and explicit [3] NeRF, our method can be used in more general applications than the other watermarking methods. Our method outperforms the other state-of-the-art watermarking methods within all metrics. In summary, our contributions are as follows:

- Our method can be applied to implicit and explicit NeRF models, unlike other existing watermarking methods.
- We propose a novel watermarking method for NeRF that fine-tunes the NeRF model by minimizing the loss function which is evaluated in the frequency domain.
- We propose a patch-wise loss to improve rendering quality and bit accuracy and enable encoding the watermark locally in the image, reducing the color artifacts.
- The proposed watermarking method achieves state-of-the-art performance, and we show that our methods are robust in diverse watermark attacks.

## 2. Related Work

**Neural Radiance Fields.** Due to the success of Neural Radiance Field (NeRF) [35], a highly photo-realistic

view synthesis of complex scenes has been achieved. Recently, NeRF has been used in various research and applications, including faster inference [26, 52, 54], 3D reconstruction [32, 39, 53], image processing [19, 31, 47], dynamic scenes [24, 42] and generative models [25, 33, 40]. Therefore, NeRF become the dominant 3D representation and is widely used. Several commercial products [36, 45] utilized NeRF 3D representations. Thus, the management of the copyright of NeRF-based 3D representations has been emerging as a crucial aspect. As all of those applications are based on implicit, explicit, or both representations, we explore a versatile watermarking method that can be trained on both representations.

**Digital watermarking.** Watermarking methods have been evolving over the decades [2, 43]. There are two main ways to recover watermarks [23]: multi-bit watermark and zero-bit watermark. The multi-bit watermark allows for the encoding and decoding of multi-bit messages from media. The zero-bit watermark can be used for ownership authentication. In this paper, we focus on the zero-bit watermark, which involves encoding and decoding multi-bit messages. Recently, with the great success of deep learning watermark methods like HiDDeN [59], considerable research has been conducted to embed watermarks into carriers, especially into 2D images [8, 27, 28], video [7, 29] and 3D [9, 49, 51].

On the NeRF field, the watermarks do not survive during the volume rendering. To overcome this issue, several methods [22, 30] have emerged. StegaNeRF [22] designed a steganography model on explicit NeRF that hides natural images in 3D scene representation. CopyRNeRF [30] secured the copyright of images rendered from an implicit 3D model by embedding a watermark into the rendered color representation. However, these methods have two big limitations. First, they could only be applied to one of the implicit or explicit representations, while both representations of NeRF greatly impact the advancement of NeRF. Second, these watermarks were not robust enough. Since StegaNeRF required the original image for message extraction, this method was not robust to distortion-like cropping, which made a lot of difference between the original and rendered images and CopyRNeRF did not resist cropping. To address these issues, we design a robust watermarking NeRF model that not only can be utilized on both implicit and explicit representation but is also robust to diverse attacks.

**Frequency Domain.** It is known that embedding a watermark into the spatial domain is vulnerable to attacks such as cropping and compression because it directly modifies the image pixels [46]. For this reason, many studies have been conducted based on frequency domains [44, 57]. The following are the frequency transformation methods.

- Discrete Fourier Transform (DFT) expresses images in terms of phase and amplitude. When a watermark is em-

bedded into the frequency domain converted by DFT, the magnitude is invariant, making it robust against geometric attacks such as rotation and scaling [14, 20, 41].

- Discrete Cosine Transform (DCT) decomposes the energy of image data into a sum of cosine functions by representing signals in the frequency domain. It is useful for image compression because most of the energy of the image is concentrated in the top-left corner coefficient. Therefore, DCT is extensively used in watermarking methods that energy compaction. [5, 17, 38]
- Discrete Wavelet Transform (DWT) decomposes the image into four subbands, LL(Low-Low), LH(Low-High), HL(High-Low), and HH(High-High). The LL band has the most energy and contains low-frequency information. Additionally, the LL band can be recursively decomposed into n-levels. DWT is used in several watermarking methods [6, 50] and shows significant differences in performance depending on the level [11, 21, 48]. In this paper, we use DWT and the LL subband for fine-tuning NeRF to embed the watermark.

### 3. Preliminaries: Representation of NeRF

In this paper, we focus on both implicit and explicit representations of a NeRF model. The implicit NeRF [35] represents a scene using a multi-layer perceptrons (MLP)  $\Phi$  whose input is 3D location  $\mathbf{x} = (x, y, z) \in \mathbb{R}^3$  and 2D viewing direction  $\mathbf{d} \in \mathbb{R}^2$  and output is color  $\mathbf{c} \in \mathbb{R}^3$ , volume density  $\sigma \in \mathbb{R}^+$ . These MLP-based radiance fields can be written as:

$$\sigma, \mathbf{c} = \Phi(\mathbf{x}, \mathbf{d}) \quad (1)$$

The explicit NeRF, such as TensoRF [3], models a scene as a 4D tensor, which represents a 3D voxel grid. These grid-based radiance fields can be written as:

$$\sigma, \mathbf{c} = \mathcal{T}_\sigma(\mathbf{x}), \mathcal{S}(\mathcal{T}_c(\mathbf{x}), \mathbf{d}), \quad (2)$$

where  $\mathcal{T}_\sigma \in \mathbb{R}^{I \times J \times K}$  is a geometry tensor and  $\mathcal{T}_c^{I \times J \times K \times C}$  is an appearance tensor for channel  $C$  and the voxel grid resolution  $I, J, K$ .  $\mathcal{S}$  is a decoding function including MLP and spherical harmonics function.

Both representations are followed by a volumetric rendering equation, which is used to synthesize novel views:

$$\begin{aligned} C(r) &= \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt, \\ T(t) &= \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right), \end{aligned} \quad (3)$$

where  $\mathbf{r}(t)$  is camera ray with near and far bounds  $t_n$  and  $t_f$ .  $C(r)$  is an expected color of the ray  $\mathbf{r}(t)$ .

Considering the NeRF representation type like the above, how can we embed a watermark in the various types, not just one type? Since these types of NeRF share a common property, rendering images through the same volume rendering function, we use the rendered images to embed watermarks into the models.

## 4. Method

We propose a process of fine-tuning the NeRF that does not involve alterations to the model’s architecture to embed the watermark message. Our method aims to embed the watermark into the weights  $\theta$  of the NeRF model in the frequency domain of rendered images. Our method stands apart from traditional watermarking methods, which focusing on training encoders and decoders. The difference lies in the fine-tuning process, which embeds the watermark without using an encoder. There are two phases: (1) Pre-training the watermark decoder  $D$ , (2) Fine-tuning the NeRF model  $F_\theta$  to embed the message. Our method is illustrated in Fig. 2 and described in detail below.

### 4.1. Pre-training the watermark decoder

We select HiDDeN [59] architecture as our watermark decoder. HiDDeN comprises two convolution networks for data hiding: a watermark encoder  $E$  and a watermark decoder  $D$ . For robustness, it includes a noise layer  $N$ . However, in this training phase, where we focus solely on the decoder’s performance, we have excluded an adversarial loss responsible for improving visual quality. After training the HiDDeN model, watermark encoder  $E$  was not utilized in the second phase.

The encoder  $E$  takes in a cover image  $I_o \in \mathbb{R}^{H \times W \times 3}$  and a binary message  $M \in \{0, 1\}^L$  with length  $L$ , as input. Then  $E$  embeds  $M$  into  $I_o$  and produces an encoded image  $I_w$ . In order to make the decoder resistant to various distortions such as rotation and JPEG compression,  $I_w$  is transformed with a noise layer  $N$ . The decoder  $D$ , made of several convolution layers, receives  $I_w$  as input and extracts a message  $M'$ .

$$M' = D(N(I_w)) \quad (4)$$

We utilize a sigmoid function to set the range of the extracted message  $M'$  between  $[0, 1]$ . The message loss is calculated with Binary Cross Entropy (BCE) between  $M_L$  and sigmoid  $sg(M'_L)$ .

$$\begin{aligned} \mathcal{L}_{\text{message}} &= - \sum_{i=1}^L M_i \cdot \log sg(M'_i) \\ &\quad + (1 - M_i) \cdot \log(1 - sg(M'_i)) \end{aligned} \quad (5)$$

The decoder is trained to detect watermarks in images that have passed through the trained encoder. However, we do not use the encoder in the second stage. We find that when the decoder receives a vanilla-rendered image, there is a bias between the extracted message bits. Thus, after training the decoder, we conduct PCA whitening to a linear decoder layer to remove the bias without reducing the extraction ability.

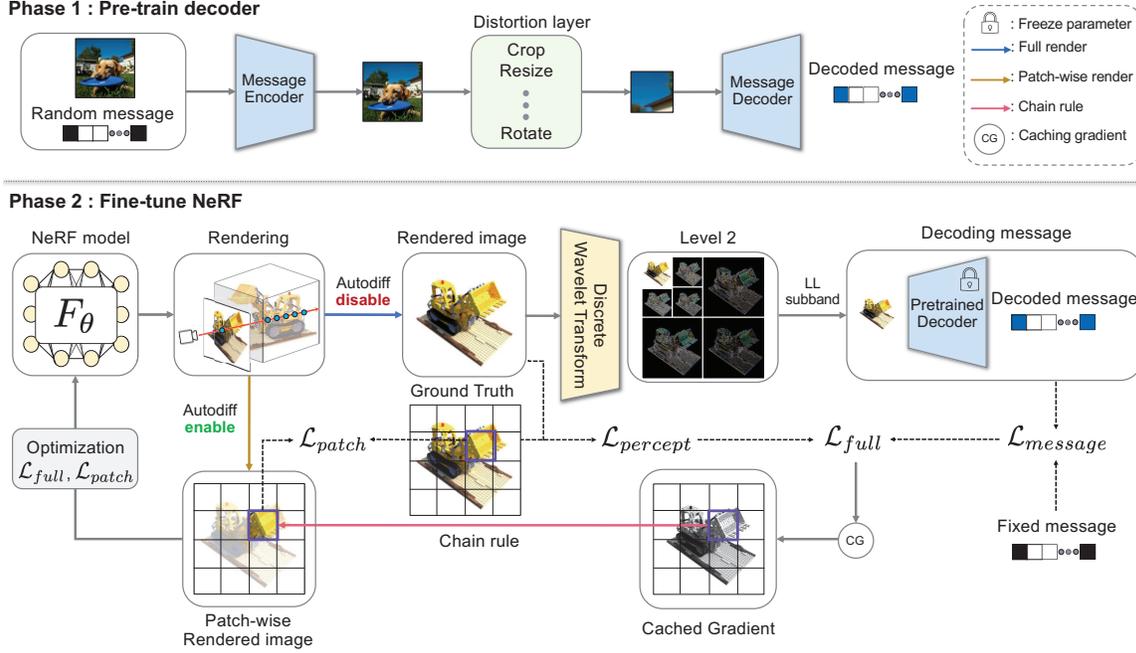


Figure 2. **WaterRF overview.** Phase 1: We train the encoder and the decoder to extract messages. After phase 1, we do not use the encoder. Phase 2: We fine-tune the NeRF to embed the messages into the rendered images. (a) We disable auto-differentiation and render a full-resolution image to save memory. (b) We use DWT for the rendered images and choose the LL subband as the input of the pre-trained decoder. (c) We enable auto-differentiation and render the images patch by patch. Then the NeRF is optimized using Eq 7 and Eq 8.

## 4.2. Embedding and extracting watermark on DWT

Recently, a fine-tuning watermarking method for NeRF [22] in the spatial domain has emerged. Although the fine-tuning method of embedding a message in the spatial domain shows incomparable invisibility and message extraction ability, it is vulnerable to attacks that distort the spatial domain, such as cropping. Directly applying spatial domain techniques from the latent diffusion model [8] does not allow for effective adjustment of NeRF’s weights.

To tackle these problems, we propose a fine-tuning method in a frequency domain instead of a spatial domain. Various watermarking techniques for images use the frequency domain have seen development and improvement over the years. We find that DWT is the appropriate domain for encoding the message into the weights of a NeRF model.

The NeRF model renders diverse views of the 3D model given corresponding camera parameters. We transform the pixels of the rendered images, denoted by  $X = (x_c, y_c) \in \mathbb{R}^{H \times W \times 3}$ , into wavelet forms, with  $c$  representing the channel. The DWT is defined as [12]:

$$W_\varphi(j_0, m, n) = \frac{1}{\sqrt{MN}} \sum_{x_c=0}^{M-1} \sum_{y_c=0}^{N-1} f(x_c, y_c) \varphi_{j_0, m, n}(x_c, y_c),$$

$$W_\psi^i(j, m, n) = \frac{1}{\sqrt{MN}} \sum_{x_c=0}^{M-1} \sum_{y_c=0}^{N-1} f(x_c, y_c) \psi_{j, m, n}^i(x_c, y_c)$$

(6)

where  $\varphi(x, y)$  is a scaling function and  $\psi(x, y)$  is a wavelet function.  $W_\varphi(j_0, m, n)$  is called by an LL subband, which is an approximation of the image at scale  $j_0$ .  $W_\psi^i$  represents LH, HL, HH subbands, where  $i = \{H, V, D\}$  and each denotes horizontal, vertical and diagonal coefficients.

Previous studies [13, 15, 18] have selected  $LH$ ,  $HL$ , and  $HH$  subbands for embedding watermarks because the  $LL$  subband contains significant information about the image. However, we choose the LL subband as an input of our decoder  $D$  and get the extract message with  $M' = D(W_\varphi)$ . We experimentally discover that embedding the watermark in the LL subband is more robust and effective than other subbands for the HiDDeN decoder.

The DWT is characterized by its subbands being computed across different levels; therefore, choosing an optimal level for our purpose is necessary. The 1-level separates the images to 4 subbands ( $LL_1, LH_1, HL_1, HH_1$ ), then 2-level separate the  $LL_1$  subband into 4 subbands ( $LL_2, LH_2, HL_2, HH_2$ ). We experimentally discover that selecting a level too high decreases visual quality due to excessive adjustment of crucial visual elements. Thus, we use the 2-level DWT because it best maintains a good balance between bit accuracy and reconstruction quality.

## 4.3. Fine-tuning the NeRF model

NeRF has two different representations, implicit and explicit representations (see Sec. 3). Previous studies [22, 30] focused on applying their methods to just one form of NeRF

representation. In contrast, we propose a method applicable to both NeRF representations: implicit and explicit. Our approach involves fine-tuning the NeRF to ensure that novel view images include an embedded message.

Our method starts with preparing each pre-trained decoder for different lengths of message bits and also preparing an initial NeRF model  $F_{\theta_0}$ . Then we establish fixed binary messages  $M = (m_1, \dots, m_L) \in \{0, 1\}^L$ .

Since rendering with NeRF requires a lot of memory, we turn off the auto-differentiation at first. The rendered images  $X \in \mathbb{R}^{H \times W \times 3}$  at full resolution are rendered by  $F_{\theta_0}$ . Following this, the full resolution images  $X$  are transformed into seven groups of wavelet subbands:  $\{LL_2, LH_2, HL_2, HH_2, LL_1, LH_1, HL_1, HH_1\}$ . Our pre-trained decoder  $D$  takes the  $LL_2$  subband as an input and extracts a message  $M' = D(W_\varphi(j_0, m, n))$  from it. The message loss is calculated by comparing  $M$  and  $M'$ :  $L_m = BCE(M, sg(M'))$ . We choose the Waston-VGG [4] to calculate the full image perceptual loss between the image without watermark  $X_0$  and the image with a watermark  $X$ :  $L_i(X_0, X)$ . Full resolution loss function is a sum of the full image perceptual loss and the message loss.

$$\mathcal{L}_{full} = \lambda_i \mathcal{L}_{percept} + \lambda_m \mathcal{L}_{message} \quad (7)$$

We cache the gradient calculated by Eq. 7 prior to the patch-wise rendering process to update the parameters. Then, we turn on the auto-differentiation so that the NeRF parameters enable patch-wise rendering. Our patch loss  $L_{patch}$  is calculated by combining mean absolute error (MAE) across rendered pixel colors, total variation (TV) regularization, and SSIM loss to maintain the balance between bit-accuracy and visual quality.

$$\mathcal{L}_{patch} = \lambda_{MAE} \mathcal{L}_{MAE} + \lambda_{TV} \mathcal{L}_{TV} + \lambda_{SSIM} \mathcal{L}_{SSIM} \quad (8)$$

Then we optimize from the initial NeRF model  $F_{\theta_0}$  to watermarked  $F_\theta$  using Eq. 7 and Eq. 8.

#### 4.4. Deferred back-propagation with patch loss

In conventional NeRF models, rendering at full resolution is inefficient regarding memory usage. To solve the memory-inefficient problem, pixels are often sampled and rendered randomly to compute the L2 loss for training. However, the main limitation of random pixel-wise rendering is in its incompatibility with using CNN-based perceptual loss and our watermark decoder for extracting message bits, both of which require a full-resolution image as input.

Our method implements a full-resolution rendering derived from our pre-trained decoder during the fine-tuning process. Each iteration presents a considerable challenge due to the extensive memory consumption. In order to relieve the memory constraints, we adopt the deferred back-propagation [55], specifically designed for memory-efficient, patch-wise optimization within NeRF.

Our method proposes a novel patch loss to be applied to the deferred back-propagation as detailed in Eq. 8. When the model renders the full-resolution image with auto-differentiation disabled, the image and message loss are computed by the rendered image, and the derived gradients are cached (see Sec. 4.3 for more details). Then, the model renders the full-resolution image with image patches that are re-rendered with auto-differentiation enabled and back-propagated using the corresponding gradients from the cached set.

Our innovation is not solely relying on cached gradients for back-propagation during re-rendering. Instead, we calculate the loss between the re-rendered patch and the ground truth (GT) patch and carry out an additional back-propagation process. The patch loss consists of a sum of L1 loss (MAE), a Total Variation loss, and an SSIM loss, as defined in Eq. 8. By utilizing back-propagation, once for the combined loss from the full-resolution image and the message loss and once for the patch loss, we are able to enhance both image quality and message bit accuracy at once.

## 5. Experiments

### 5.1. Dataset

We evaluate our method using datasets commonly used by NeRF: the Blender dataset [35], which consists of synthetic bounded scenes, and the LLFF dataset [34], which consists of real unbounded, forward-facing scenes. We follow the conventional literature in NeRF papers, which evaluate and compare using eight scenes each from the Blender dataset (including chair, drums, ficus, hotdog, lego, materials, mic) and the LLFF dataset (including fern, flower, fortress, horns, leaves, orchids, room, trex).

### 5.2. Implementation Details

Vanilla NeRF [35] and TensorRF [3] employ similar training methods, with some specifics as detailed below. The training process is carried out on an image-by-image basis, with a batch size set to one. Ray tracing is implemented along with the camera pose, and the Adam optimizer [16] is used for the learning process. The training is completed with epochs ranging from 5 to 10. Since our pre-trained decoder requires full-resolution images, NeRF renders full-resolution images during the fine-tuning process. However, NeRF requires a great amount of memory usage. We adopt the deferred back-propagation in ARF [55] to solve the memory issue and combine it with the patch-wise loss to improve rendering quality and bit accuracy.

The initial  $\lambda_i$ ,  $\lambda_m$  was set to 0.1 and 0.9 for the LLFF dataset and 0.05 and 0.95 for the Blender dataset, respectively. Additionally, within the patch loss equation (Eq. 8), the parameters  $\lambda_{MAE}$ ,  $\lambda_{TV}$ , and  $\lambda_{SSIM}$  were set to 0.1, 0.06, and 0.02.

	Bit Accuracy(%) $\uparrow$				
	4 bits	8 bits	16 bits	32 bits	48 bits
HiDDeN [59]+NeRF [35]	50.31	50.25	50.19	50.11	50.04
CopyRNeRF [30]	<b>100.00</b>	<b>100.00</b>	91.16	78.08	60.06
Ours (w/ NeRF [35])	<b>100.00</b>	<b>100.00</b>	94.24	86.81	70.43
Ours (w/ TensoRF [3])	<b>100.00</b>	<b>100.00</b>	<b>95.67</b>	<b>88.58</b>	<b>85.82</b>

Table 1. Bit accuracies with different message lengths are compared with the state-of-the-art method. (HiDDeN [59]+NeRF [35] is pre-embedded before training). The results are the average of LLFF and Blender datasets.

### 5.3. Evaluation

We evaluate our method with three key essential aspects in watermarking techniques: 1) *capacity*, measuring the length of the messages that can be embedded into the image; 2) *invisibility*, which involves evaluating the difficulty in detecting embedded watermarks in images by peoples; 3) *robustness*, evaluating the robustness of our watermarking method under various image distortions such as Gaussian Noise, Rotation, Scaling, Gaussian Blur, Crop, Brightness, JPEG compression and combination of these distortions. We employ Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [56] for the evaluation metrics to assess invisibility while utilizing bit accuracy to evaluate capacity and robustness. We conduct experiments on the invisibility and robustness of the message length of  $M_L \in \{4, 8, 16, 32, 48\}$ .

**Baseline** To the best of our knowledge, CopyRNeRF [30] stands out alone in its ability to embed bits with the implicit NeRF model, which can be converted into various modality data such as messages or images. The CopyRNeRF performs state-of-the-art on bit-accuracy and visual quality. As a result, we conduct our evaluations utilizing the CopyRNeRF. We select the following relevant comparative models: HiDDeN [59]+NeRF [35], which employ the classical 2D watermarking method HiDDeN [59] on training images prior to training the NeRF model.

To demonstrate the applicability of our method, we apply our method on two distinct NeRF models: Vanilla NeRF [35], symbolizing implicit representation NeRF, and TensoRF [3], representing explicit representation NeRF based on voxel grids. We select the Vanilla NeRF [35] since it is the first NeRF model to be introduced and can be considered the foundational and ancestral model for all subsequent NeRFs. TensoRF [3] is chosen because it is an influential NeRF model with significantly improved performance and training speed.

### 5.4. Experimental results

**Capacity** As shown in Tab. 1, there is a clear trade-off between bit accuracy and message bit length. We conduct bit

	Bit Acc(%) $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
HiDDeN [59]+NeRF [35]	50.19	26.53	0.917	0.035
CopyRNeRF [30]	91.16	26.29	0.910	0.038
Ours (w/ NeRF [35])	94.24	28.81	<b>0.954</b>	<b>0.025</b>
Ours (w/ TensoRF [3])	<b>95.67</b>	<b>32.79</b>	0.948	0.033

Table 2. Bit accuracies and reconstruction qualities comparison with baselines.  $\uparrow$  ( $\downarrow$ ) means higher (lower) is better. We show the results on 16 bits. The results are evaluated in the same way as baselines. The best performances are highlighted in **bold**.

accuracy tests on messages of lengths 4, 8, 16, 32, and 48 bits. Models implemented with our method show a gradual decrease in bit accuracy as the length of the message increases. However, when compared to CopyRNeRF [30], the decline in bit accuracy is less dramatic and less steep. Notably, CopyRNeRF [30] shows a modest 60.6% bit accuracy for 48 bits, whereas Vanilla NeRF [35] and TensoRF [3] using our method demonstrates a respectable 85.82% and 70.43% accuracy, respectively. Additionally, to demonstrate the stability of our method across a range of messages, we fine-tune 500 different messages, each corresponding to unique, randomly generated bits. We rigorously evaluate the bit accuracy for all images rendered by Vanilla NeRF [35] and TensoRF [3]. As illustrated in Fig. 3, the plot shows a minor descending trend that corresponds with the increasing number of unique messages, and it shows that our method is robust with using many different messages, which is a key value to identify unique NeRF model.

**Invisibility** Our method demonstrates a trade-off between bit accuracy and the quality of reconstructed data, as detailed in Tab. 2. This trade-off is analyzed using several metrics, including bit accuracy and visual quality metrics. With TensoRF [3], our method achieves the best bit accuracy and PSNR results. Conversely, when applied to Vanilla

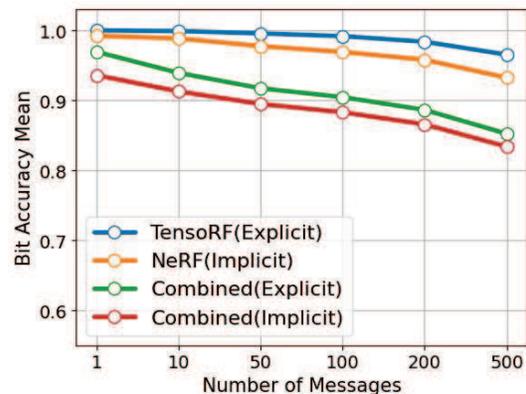


Figure 3. **Identification results.** We average the bit accuracy for each 500 messages. Our method not only works efficiently with just one key but also works well with arbitrarily generated messages. We show the results on  $M_L = 16$  bits.

	No Distortion	Gaussian Noise ( $\nu = 0.1$ )	Rotation ( $\pm\pi/6$ )	Scaling (25%)	Bit Accuracy(%) $\uparrow$				Combined (Crop, Brightness, JPEG)
					Gaussian Blur (deviation = 0.1)	Crop (40%)	Brightness (2.0)	JPEG Compression (10% quality)	
CopyRNeRF [30]	91.16	90.04	88.13	89.33	90.06	-	-	-	-
Ours (w/ NeRF [35])	94.24	94.06	85.02	91.35	94.12	83.48	84.14	86.88	73.64
Ours (w/ TensorRF [3])	<b>95.67</b>	<b>95.36</b>	<b>93.13</b>	<b>93.29</b>	<b>95.25</b>	<b>95.40</b>	<b>90.91</b>	<b>86.99</b>	<b>84.12</b>

Table 3. Robustness under diverse attacks compared with the state-of-the-art method. The bit accuracy results are the average of LLFF and Blender datasets. We show the results on  $M_L = 16$  bits. The best performances are highlighted in **bold**.

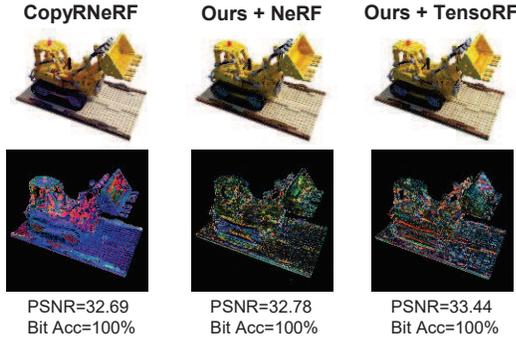


Figure 4. **Qualitative comparisons** We show the differences ( $\times 10$ ) between the rendered images and the ground truth. Our method achieves higher PSNR and bit accuracy than CopyRNeRF.

NeRF [35], it yields the highest SSIM and the lowest LPIPS (noting that higher SSIM and lower LPIPS are desirable). Although the SSIM and LPIPS scores for TensorRF [3] using our method are not the highest, they are comparable to those achieved with Vanilla NeRF [35]. In summary, models trained with our method attain superior bit accuracy, PSNR, SSIM, and the lowest LPIPS, indicating a well-balanced trade-off between bit accuracy and reconstruction quality.

Additionally, to verify the invisibility of our Result, We compute the differences between ground truth and our rendered results. As illustrated in Fig. 4, the results of applying our method to both Vanilla NeRF [35] and TensorRF [3] achieve a better balance between reconstruction quality and bit accuracy compared to CopyRNeRF [30]. This indicates that our method can embed watermarks with minimal impact on the reconstruction quality.

**Robustness** We evaluate the robustness of the watermarked image to different attacks by applying Gaussian noise, Rotation, Scaling, Gaussian Blur, Crop, Brightness, JPEG compression and combination of these distortions. As shown in Tab. 3, We experiment on both Vanilla NeRF [35] and TensorRF [3] when applied with our method, resulting in robustness against various attacks compared to CopyRNeRF [30].

## 5.5. Ablation study

**Patch Loss and Frequency Domain** This section delves into the effectiveness of integrating patch loss and utilizing the frequency domain. Fig. 5 presents a comparative analysis of our method, applied to Vanilla NeRF [35], both with

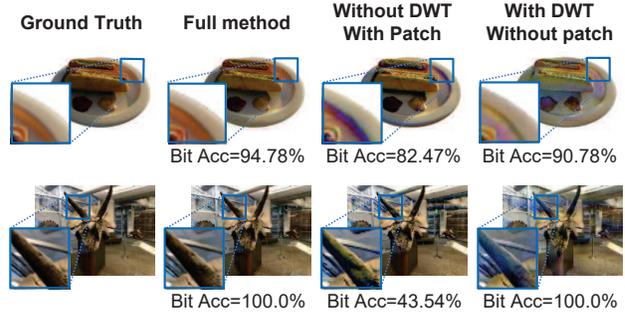


Figure 5. **Reconstruction quality comparisons** We evaluate our full method, our method without patch loss and our method without frequency domain about 16 bits with NeRF [35].

	Bit Acc(%) $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
DFT	71.06	19.04	0.822	0.1719
DCT	43.75	<b>35.33</b>	<b>0.967</b>	<b>0.0205</b>
No frequency	72.68	32.80	0.944	0.0392
DWT-Level 1	91.16	33.04	0.947	0.0332
DWT-Level 2	<b>95.67</b>	32.79	0.948	0.0336
DWT-Level 3	90.96	32.45	0.950	0.0331
DWT-Level 4	85.02	31.82	0.952	0.0329

Table 4. Bit accuracies and reconstruction qualities compared with spatial and frequency domains (DFT, DCT) and DWT levels 1, 2, 3, and 4. We focus on low frequency as our assumption. We show the results on 16 bits. The results are evaluated in the same way as baselines. The best performances are highlighted in **bold**.

and without the incorporation of patch loss and use in the frequency domain. As shown in Fig. 5, the absence of DWT significantly decreases bit accuracy. Without the patch-wise loss, there is a decrease in reconstruction quality.

**DWT Level** During our experiments with DWT, we evaluate performance across commonly used levels: 1, 2, 3, and 4. As shown in Tab. 4, it is indicated that the results from levels 1, 3, and 4 do not perform as well as the outcomes achieved at level 2, both in terms of bit accuracy and reconstruction quality balance. Additionally, we conduct applying DFT, DCT, and no frequency instead of DWT. As a result, our findings show that the DWT at level 2 is the most effective. Thus, we use the DWT at level 2 based on the results from the ablation experiments.

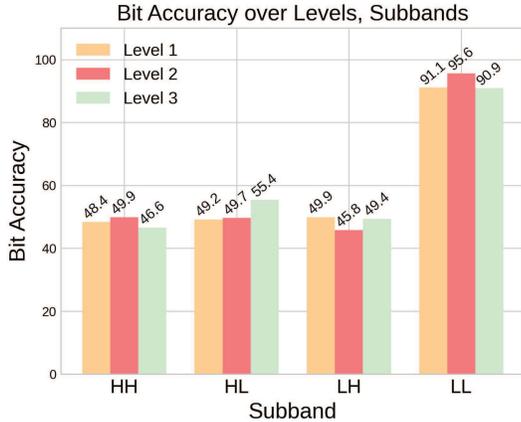


Figure 6. **Bit accuracy comparison between DWT Subbands** We evaluate the bit accuracy on level 1,2,3 and all subbands. We show the results on  $M_L = 16$  bits.

**DWT Subband** The choice of DWT subband is as critical as the level of DWT itself. Typically, subbands such as LH, HL, and HH, which are composed of high-frequency components less perceptible to the people, are employed for watermark embedding in conventional watermarking techniques. Conversely, in our experimental approach, we opt for the LL subband representing the image’s coarse approximation and containing essential low-frequency details. This decision is based on the watermark decoder’s pre-training process to extract watermarks specifically from the image’s spatial domain, which is visually similar to the LL subband and its demonstrated robustness to JPEG compression. In Fig. 6, our empirical investigations confirm that the LL subband yields the highest bit accuracy.

**Comparison of Training time.** Even though our proposed method necessitates the use of a pre-trained model’s weights as a starting point for fine-tuning NeRF, the time requirement for training is significantly less compared to that of CopyRNeRF [30]. As depicted in Fig. 7, it is evident that the model utilizing our techniques substantially outperforms CopyRNeRF [30] regarding training efficiency. Our method, when applied to an explicit NeRF architecture, such as TensorRF [3], achieves a remarkable speed increase, performing up to six times faster than CopyRNeRF [30]. Similarly, implementing our techniques within the Implicit NeRF framework, like Vanilla NeRF [35], yields a significant speed improvement, with our model operating six times faster than CopyRNeRF [30] but it also surpasses the bit accuracy of the compared model.

## 6. Conclusion

We propose a neural 3D watermarking method for the NeRF model. Our method trains a 2D watermark decoder and NeRF model separately. Therefore, our pipeline only requires training the decoder once and re-use it on differ-

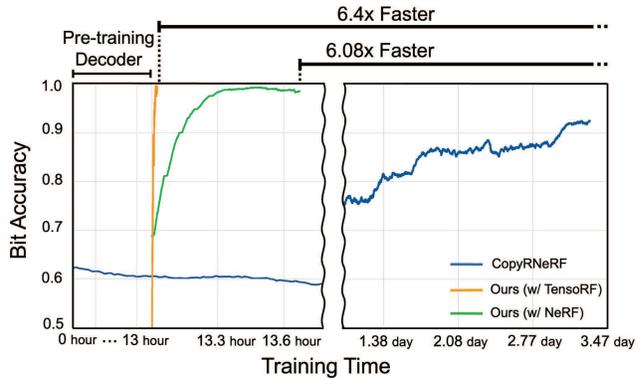


Figure 7. **Analysis of train time** for reconstruction from blender dataset at 160x160 resolution. Note the introduction of a scale break within the graph to accommodate the significant disparities in training durations, with time represented in minutes on the left axis and in days on the right axis.

ent NeRF watermark models. We adopt conventional watermark techniques in image watermarking, which transfer the image from the spatial domain to the frequency domain to encode the watermark into the image efficiently. We find that DWT and patch-wise loss enhance the overall reconstruction quality while performing high message bit accuracy. Through extensive experiments, we demonstrate that our method outperforms the state-of-the-art method and is six times faster than the previous work.

**Limitations.** Our method shows outstanding performance in the digital watermark of NeRF, but training a decoder that extracts watermark from rendered images is time-consuming, taking approximately 12 hours on a single RTX 3090. Our method only encodes a unique message per model, requiring fine-tuning for message insertion every time. We observe that the bit accuracy drops when we apply numerous messages(see Fig. 3). Although previous studies, including us, report robustness against various distortions of images, there is no consideration about cases where the model is attacked. Future work may explore watermarks resilient to cases of model attacks.

**Acknowledgments.** This research was supported by the Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture(Project : 4D Content Generation and Copyright Protection with Artificial Intelligence, R2022020068, 80%), the Korea Research Institute for Defense Technology Planning and Advancement (KRIT) Grant funded by the Korean Government(Grant KRIT-CT-23-021, 10%), the National Research Foundation of Korea grant (NRF-2022R1F1A1074334, 9%), and the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT)(No. 2019-0-00079, Artificial Intelligence Graduate School Program(Korea University), 1%).

## References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, pages 5855–5864, 2021. 1
- [2] Mahbuba Begum and Mohammad Shorif Uddin. Digital image watermarking techniques: a review. *Information*, 11(2): 110, 2020. 2
- [3] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *ECCV*, pages 333–350. Springer, 2022. 1, 2, 3, 5, 6, 7, 8
- [4] Steffen Czolbe, Oswin Krause, Ingemar Cox, and Christian Igel. A loss function for generative neural networks based on watson’s perceptual model. *Advances in Neural Information Processing Systems*, 33:2051–2061, 2020. 5
- [5] Chinmayee Das, Swetalina Panigrahi, Vijay K Sharma, and KK Mahapatra. A novel blind robust image watermarking in dct domain using inter-block coefficient correlation. *AEU-International Journal of Electronics and Communications*, 68(3):244–253, 2014. 3
- [6] Xin Deng, Chao Gao, and Mai Xu. Pirnet: Privacy-preserving image restoration network via wavelet lifting. In *ICCV*, pages 22368–22377, 2023. 3
- [7] Elhameh Farri and Peyman Ayubi. A robust digital video watermarking based on ct-svd domain and chaotic dna sequences for copyright protection. *Journal of Ambient Intelligence and Humanized Computing*, 14(10):13113–13137, 2023. 2
- [8] Pierre Fernandez, Guillaume Couairon, Hervé Jégou, Matthijs Douze, and Teddy Furon. The stable signature: Rooting watermarks in latent diffusion models. *ICCV*, 2023. 2, 4
- [9] Felipe ABS Ferreira and Juliano B Lima. A robust 3d point cloud watermarking method based on the graph fourier transform. *Multimedia Tools and Applications*, 79(3):1921–1950, 2020. 2
- [10] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, pages 5501–5510, 2022. 1
- [11] Youcai Gao, Jinwei Wang, and Yun-Qing Shi. Dynamic multi-watermarking and detecting in dwt domain. *Journal of Real-Time Image Processing*, 16:565–576, 2019. 3
- [12] Rafael C Gonzalez. *Digital image processing*. Pearson education india, 2009. 4
- [13] A Hernández-Joaquín, G Melendez-Melendez, and R Cumplido. A secure dwt-based dual watermarking scheme for image authentication and copyright protection. *Multimedia Tools and Applications*, 82(27):42739–42761, 2023. 4
- [14] Ningombam Jimson and K Hemachandran. Dft based coefficient exchange digital image watermarking. In *2018 Second International Conference on Intelligent Computing and Control Systems (ICICCS)*, pages 567–571. IEEE, 2018. 3
- [15] Priyank Khare and Vinay Kumar Srivastava. A reliable and secure image watermarking algorithm using homomorphic transform in dwt domain. *Multidimensional Systems and Signal Processing*, 32:131–160, 2021. 4
- [16] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015. 5
- [17] Hung-Jui Ko, Cheng-Ta Huang, Gwoboa Horng, and WANG Shih-Jeng. Robust and blind image watermarking in dct domain using inter-block coefficient correlation. *Information Sciences*, 517:128–147, 2020. 3
- [18] Sanjay Kumar and Binod Kumar Singh. Dwt based color image watermarking using maximum entropy. *Multimedia Tools and Applications*, 80:15487–15510, 2021. 4
- [19] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas J Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic neural fields: A semantic object-aware neural scene representation. In *CVPR*, pages 12871–12881, 2022. 2
- [20] Guezouli Larbi, Hachani Roumaissa, and Herri Sara. Embedding watermark in the magnitude matrix of the dft of image. In *Proceedings of the 2018 International Conference on Computing and Pattern Recognition*, pages 106–110, 2018. 3
- [21] Yong-Seok Lee, Seo Young-Ho, and Kim Dong-Wook. Blind image watermarking based on adaptive data spreading in n-level dwt subbands. *Security and communication Networks*, 2019, 2019. 3
- [22] Chenxin Li, Brandon Y Feng, Zhiwen Fan, Panwang Pan, and Zhangyang Wang. Steganerf: Embedding invisible information within neural radiance fields. In *ICCV*, pages 441–453, 2023. 1, 2, 4
- [23] Yue Li, Hongxia Wang, and Mauro Barni. A survey of deep neural network watermarking techniques. *Neurocomputing*, 461:171–193, 2021. 2
- [24] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. Dynibar: Neural dynamic image-based rendering. In *CVPR*, pages 4273–4284, 2023. 2
- [25] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, pages 300–309, 2023. 2
- [26] Haotong Lin, Sida Peng, Zhen Xu, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Efficient neural radiance fields for interactive free-viewpoint video. In *SIGGRAPH Asia Conference Proceedings*, 2022. 2
- [27] Xiyang Luo, Ruohan Zhan, Huiwen Chang, Feng Yang, and Peyman Milanfar. Distortion agnostic deep watermarking. In *CVPR*, pages 13548–13557, 2020. 2
- [28] Xiyang Luo, Michael Goebel, Elnaz Barshan, and Feng Yang. Leca: A learned approach for efficient cover-agnostic watermarking. In *Media Watermarking, Security, and Forensics*, 2022. 2
- [29] Xiyang Luo, Yinxiao Li, Huiwen Chang, Ce Liu, Peyman Milanfar, and Feng Yang. Dvmark: a deep multiscale framework for video watermarking. *IEEE Transactions on Image Processing*, 2023. 2
- [30] Ziyuan Luo, Qing Guo, Ka Chun Cheung, Simon See, and Renjie Wan. Copyrnerf: Protecting the copyright of neural

- radiance fields. In *ICCV*, pages 22401–22411, 2023. [1](#), [2](#), [4](#), [6](#), [7](#), [8](#)
- [31] 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 *CVPR*, pages 12861–12870, 2022. [2](#)
- [32] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, pages 7210–7219, 2021. [2](#)
- [33] Luke Melas-Kyriazi, Iro Laina, Christian Ruppert, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *CVPR*, pages 8446–8455, 2023. [2](#)
- [34] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019. [5](#)
- [35] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [36] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. [2](#)
- [37] Thu Nguyen-Phuoc, Feng Liu, and Lei Xiao. Snerf: stylized neural implicit representations for 3d scenes. *ACM Trans. Graph.*, 41(4), 2022. [1](#)
- [38] Shabir A Parah, Javaid A Sheikh, Nazir A Loan, and Ghulam M Bhat. Robust and blind watermarking technique in dct domain using inter-block coefficient differencing. *Digital Signal Processing*, 53:11–24, 2016. [3](#)
- [39] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, pages 5865–5874, 2021. [2](#)
- [40] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*, 2023. [2](#)
- [41] SN Prajwalasimha, S Chethan Suputhra, and C Mohan. Performance analysis of combined discrete fourier transformation (dft) and successive division based image watermarking scheme. *International Journal of Recent Technology and Engineering*, 8(3):34–39, 2019. [3](#)
- [42] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, pages 10318–10327, 2021. [2](#)
- [43] Deepti Shukla, Nirupama Tiwari, and Deepika Dubey. Survey on digital watermarking techniques. *International Journal of Signal Processing, Image Processing and Pattern Recognition*, 9(1):239–244, 2016. [2](#)
- [44] S Suthaharan and S Sathanathan. Transform domain technique: robust watermarking for digital images. In *Proceedings of the IEEE SoutheastCon 2000: Preparing for The New Millennium' (Cat. No. 00CH37105)*, pages 407–412. IEEE, 2000. [2](#)
- [45] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Justin Kerr, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, David McAllister, and Angjoo Kanazawa. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. [2](#)
- [46] Ron G Van Schyndel, Andrew Z Tirkel, and Charles F Osborne. A digital watermark. In *Proceedings of 1st international conference on image processing*, pages 86–90. IEEE, 1994. [2](#)
- [47] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *CVPR*, pages 3835–3844, 2022. [2](#)
- [48] Jinwei Wang, Shiguo Lian, and Yun-Qing Shi. Hybrid multiplicative multi-watermarking in dwt domain. *Multidimensional Systems and Signal Processing*, 28:617–636, 2017. [3](#)
- [49] Yinghui Wang, Jing Liu, Yajie Yang, Douli Ma, and Ruijiao Liu. 3d model watermarking algorithm robust to geometric attacks. *IET Image Processing*, 11(10):822–832, 2017. [2](#)
- [50] Muyu Xu, Fangneng Zhan, Jiahui Zhang, Yingchen Yu, Xiaojin Zhang, Christian Theobalt, Ling Shao, and Shijian Lu. Wavenerf: Wavelet-based generalizable neural radiance fields. In *ICCV*, pages 18195–18204, 2023. [3](#)
- [51] Innfarn Yoo, Huiwen Chang, Xiyang Luo, Ondrej Stava, Ce Liu, Peyman Milanfar, and Feng Yang. Deep 3d-to-2d watermarking: embedding messages in 3d meshes and extracting them from 2d renderings. In *CVPR*, pages 10031–10040, 2022. [2](#)
- [52] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *ICCV*, pages 5752–5761, 2021. [2](#)
- [53] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, pages 4578–4587, 2021. [2](#)
- [54] Heng Yu, Joel J. Julian, Zoltan A Milacski, Koichiro Niinuma, and László A Jeni. Dylon: Making light field networks dynamic. In *CVPR*, pages 12397–12406, 2023. [2](#)
- [55] Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields. In *ECCV*, pages 717–733. Springer, 2022. [5](#)
- [56] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. [6](#)
- [57] Qiumei Zheng, Nan Liu, Baoqin Cao, Fenghua Wang, and Yanan Yang. Zero-watermarking algorithm in transform domain based on rgb channel and voting strategy. *Journal of Information Processing Systems*, 16(6):1391–1406, 2020. [2](#)
- [58] Haoyi Zhu. X-nerf: Explicit neural radiance field for multi-scene 360deg insufficient rgb-d views. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5766–5775, 2023. [1](#)

- [59] Jiren Zhu, Russell Kaplan, Justin Johnson, and Li Fei-Fei. Hidden: Hiding data with deep networks. In *ECCV*, pages 657–672, 2018. [2](#), [3](#), [6](#)