

StegaNeRF: Embedding Invisible Information within Neural Radiance Fields

Chenxin Li^{1*}, Brandon Y. Feng^{2*}, Zhiwen Fan^{3*}, Panwang Pan⁴, Zhangyang Wang³
¹Chinese University of Hong Kong, ²University of Maryland,
³University of Texas at Austin, ⁴ByteDance

Abstract

Recent advancements in neural rendering have paved the way for a future marked by the widespread distribution of visual data through the sharing of Neural Radiance Field (NeRF) model weights. However, while established techniques exist for embedding ownership or copyright information within conventional visual data such as images and videos, the challenges posed by the emerging NeRF format have remained unaddressed. In this paper, we introduce StegaNeRF, an innovative approach for steganographic information embedding within NeRF renderings. We have meticulously developed an optimization framework that enables precise retrieval of hidden information from images generated by NeRF, while ensuring the original visual quality of the rendered images to remain intact. Through rigorous experimentation, we assess the efficacy of our methodology across various potential deployment scenarios. Furthermore, we delve into the insights gleaned from our analysis. StegaNeRF represents an initial foray into the intriguing realm of infusing NeRF renderings with customizable, imperceptible, and recoverable information, all while minimizing any discernible impact on the rendered images. For more details, please visit our project page: <https://xgqnet.github.io/StegaNeRF/>

1. Introduction

Implicit neural representation (INR) is an emerging concept where the network describes the data through its weights [14, 24, 40, 42, 45, 58, 65]. After training, the INR weights can then be used for content distribution, streaming, and even downstream inference tasks, all without sending or storing the original data. Arguably the most prominent INR is Neural Radiance Fields (NeRF) [42], where a network learns a continuous function mapping spatial coordinates to density and color. Due to its lightweight size and superb quality, NeRF has immense potential for 3D content representation in future vision and graphics applications.

*Equal contribution

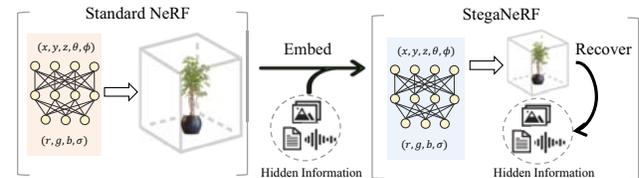


Figure 1. We introduce the new problem of NeRF steganography: hiding information in NeRF renderings. Our proposed framework, StegaNeRF, can embed and recover customized hidden information while preserving the original NeRF rendering quality.

While there is a plethora of work dedicated towards better quality [4, 63, 72, 91], faster rendering [43, 56, 61, 71, 83], and sparse view reconstruction [10, 11, 44, 57, 80, 81, 85, 88], in this paper, we look beyond the horizon and explore a new question: *Can we achieve steganography with NeRF?*

Established digital steganography method [12] focus on embedding hidden messages in 2D images. The recent growth of deep learning and social media platforms further gives rise to many practical use cases of image steganography. As countless images and videos are shared online and even used to train deep learning models, 2D steganography methods [2, 3] allow users and data providers to protect copyright, embed ownership, and prevent content abuse.

Now, with the ongoing advances in 3D representations powered by NeRF, we envision a future where people frequently share their captured 3D content online just as they are currently sharing 2D images and videos online. Moreover, we are curious to explore the following research questions: ❶ Injecting information into 2D images for copyright or ownership identification is common, but can we preserve such information when people share and render 3D scenes through NeRF? ❷ NeRF can represent large-scale real-world scenes with training images taken by different people, but can we preserve these multiple source identities in the NeRF renderings to reflect the collaborative efforts required to reconstruct these 3D scenes? ❸ Common image steganography methods embed either a hidden image or a message string into a given image, but can we allow different modalities of the hidden signal in NeRF steganography?

Driven by these questions, we formulate a framework to embed customizable, imperceptible, and recoverable information in NeRF renderings without sacrificing the visual quality. Fig. 1 presents an overview of our proposed framework, dubbed StegaNeRF. Unlike the traditional image steganography that embeds hidden signals only into a specific source image, we wish to recover the same intended hidden signal from NeRF rendered at arbitrary viewpoints.

Despite many established works on 2D steganography and hidden watermarks for image and video [2, 3, 9, 25, 36, 64], naively applying 2D steganography on the NeRF training images is not practical, since the embedded information easily gets lost in the actual NeRF renderings. In contrast, our framework enables reliable extraction of the hidden signal from NeRF renderings. During NeRF training, we jointly optimize a decoder network to extract hidden information from the NeRF renderings. To minimize the negative impact on the visual quality of NeRF, we identify weights with low impact on rendering and introduce a gradient masking strategy to steer the hidden steganographic information towards those low-importance weights. Extensive experimental results validate StegaNeRF balances between the rendering quality of novel views and the high-fidelity transmission of the concealed information.

StegaNeRF presents the first exploration of hiding information in NeRF models for ownership identification, and our contributions can be summarized as follows:

- We introduce the new problem of NeRF steganography and present the first effort to embed customizable, imperceptible, and recoverable information in NeRF.
- We propose an adaptive gradient masking strategy to steer the injected hidden information towards the less important NeRF weights, balancing the objectives of steganography and rendering quality.
- We empirically validate the proposed framework on a diverse set of 3D scenes with different camera layouts and scene characteristics, obtaining high recovery accuracy without sacrificing rendering quality.
- We explore various scenarios applying StegaNeRF for ownership identification, with the additional support of multiple identities and multi-modal signals.

2. Related Work

Neural Radiance Fields Following the success of NeRF [42], numerous techniques have been incorporated including ray re-parameterizations [4, 5], explicit spatial data structures [16, 23, 34, 43, 61, 84], caching and distillation [19, 22, 54, 69], ray-based representations [1, 15, 59], geometric primitives [33, 35], large-scale scenes [39, 63, 79, 89], generalizable architectures [60, 68, 73], and dynamic

settings [18, 32, 46, 51, 78]. Unlike prior works, this paper explores the uncharted problem of embedding information in NeRF renderings, with critical implications for copyright protection and ownership preservation. As early-stage NeRF-based products already become available [37, 49] and more are expected, we believe now is the right time to open up the exploration of NeRF steganography.

Image Steganography Steganography hides intended signals as invisible watermarks (*e.g.*, hyperlinks, images) within the cover media called carriers (*e.g.*, images, video and audio) [12, 26]. Classical methods focus on seeking and altering the least significant bits (LSBs) [17, 47, 62, 76] and transforming domain techniques [7, 8, 52, 70]. Prior research also uses deep neural networks for steganography [2, 3, 21, 66]. Among them, DeepStega [2] conceals a hidden image within an equal-sized carrier image. Subsequent works [25, 36] use invertible networks to improve the performance of deep image steganography. Another line of work conceals information in other carrier media like audio [13, 20, 82] and video [38, 55, 75]. The above advances all play a critical part in the era when traditional media formats like images and videos are dominant. However, as MLP-based neural representations of 3D scenes are gaining momentum to become a major format of visual data, extending steganography to NeRF is bound to become an important problem in the upcoming future.

Lifting Steganography to 3D Prior to NeRF, meshes are commonly used to represent 3D shapes. Early pioneers apply steganography 3D meshes [6, 50] for copyright protection when meshes are exchanged and edited. More recent work [77] has also explored embedding multi-plane images within a JPEG-compressed image. Their problem can be regarded as a special case of 2D steganography, hiding multiple 2D images inside a single 2D image. In contrast, we try to hide a natural image into a 3D scene representation (NeRF), fundamentally differing from these prior methods where 2D images act as the carrier of hidden information.

3. Method

The goal of StegaNeRF is to inject customized (steganographic) information into the NeRF weights with imperceptible visual changes when rendering. Given a NeRF with weights θ_0 and the information I to be injected, when we render with θ on any viewpoint, we hope the injected I can be recovered by a decoder F_ψ with learnable weights ψ .

A seemingly obvious solution is to use prior image steganography methods to 1) inject I on training images, 2) train NeRF on those images with embedded information, 3) apply their provided F_ψ to extract I from NeRF-rendered images. This approach has been successfully applied for GAN fingerprinting [86]. However, it fails in the NeRF context (see Fig. 3), where the subtle changes induced by prior

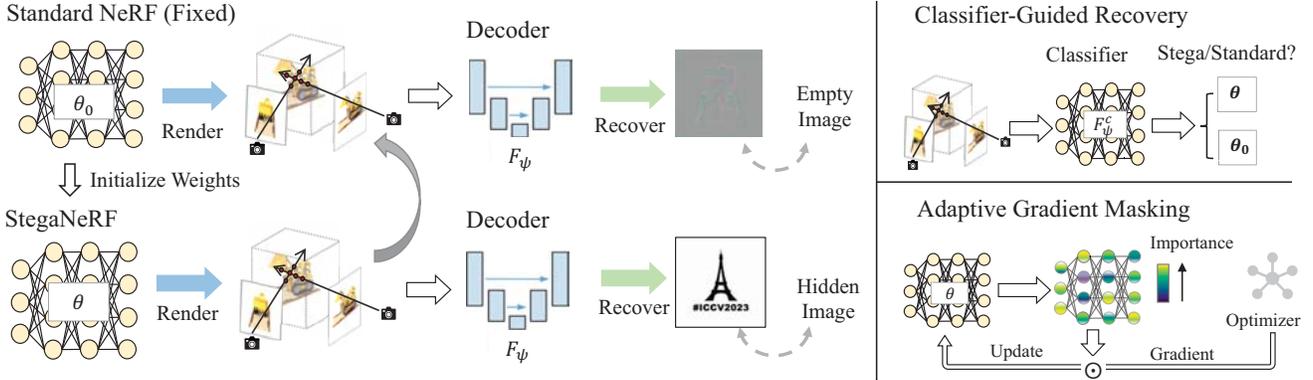


Figure 2. StegaNeRF training overview. At the first stage, we optimize θ_0 with standard NeRF training. At the second stage, we initialize θ with θ_0 and optimize for the steganography objectives. We train the decoder F_ψ to recover hidden information from StegaNeRF renderings and no hidden information from original NeRF renderings. We introduce *Classifier Guided Recovery* to improve the accuracy of recovered information, and *Adaptive Gradient Masking* to balance between steganography ability and rendering visual quality.

steganography methods easily get smoothed out, inhibiting the decoder from identifying the subtle patterns necessary for information recovery.

Such shortcomings of off-the-shelf steganographic methods are not surprising since they are developed for the traditional setting where 2D images are the ultimate form of visual information. In contrast, our problem setting involves the new concept of INR as the underlying representation, and 2D images are just the final output rendered by NeRF.

3.1. Two-Stage Optimization

Recognizing the formulation difference due to the emergence of implicit representation with NeRF, we move away from the traditional 2D image-based pipeline that trains an encoder network to inject subtle changes to the given 2D images. Instead, we incorporate the steganography objective into the gradient-based learning process of NeRF.

We re-design the training as a two-stage optimization. The first stage is the original NeRF training procedure, involving the standard photometric loss between the rendered and ground truth pixels to guarantee the visual quality. After finishing the weights update θ_0 at the first stage, we dedicate the second stage to obtain the final NeRF weights θ containing steganographic information. We introduce several techniques to achieve robust information recovery with imperceptible impact on the rendered images. The training workflow is depicted in Fig. 2 and Alg. 1.

3.2. Information Recovery

Let P denote the camera pose at which we render an image from a NeRF network. We want to recover I from the rendered image $\theta(P)$. Importantly, we also want to avoid false positives on images $\theta_0(P)$ rendered by the original network without steganography ability, even if the images rendered by θ_0 and θ look visually identical. Therefore, we

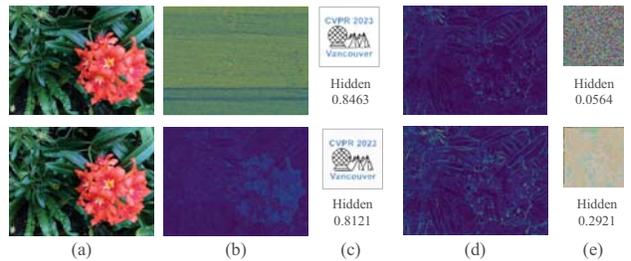


Figure 3. Results of applying prior 2D steganography methods, LSB [9] (top) and DeepStega [2] (bottom). From left to right, we show (a) training image after applying 2D steganography, (b) residual error of (a) over ground truth, (c) hidden image recovered from (a), (d) residual error of NeRF rendering at the pose of (a), and (e) hidden image recovered from (d). Prior 2D steganography methods fail in the NeRF context since the hidden information injected in training images mostly disappear in NeRF renderings.

optimize θ to minimize the following contrastive loss terms:

$$\mathcal{L}_{dec}^+ = |F_\psi(\theta(P)) - I|, \quad \mathcal{L}_{dec}^- = |F_\psi(\theta_0(P)) - \emptyset|, \quad (1)$$

where \emptyset is an empty image with all zeros. Effectively, \mathcal{L}_{dec}^+ serves as a regularization term, facilitating the decoder to recover the embedded image pattern based on renderings obtained from a StegaNeRF. In contrast, \mathcal{L}_{dec}^- prevents the decoder from falsely producing any plausible image pattern when given renderings from a Standard NeRF that is without hidden signals.

Classifier-Guided Recovery The decoder F_ψ is easily implemented as a U-Net to decode I as the form of 2D images, but accurately recovering all the details in I might be challenging. Therefore, we additionally train a classifier

Algorithm 1 Train StegaNeRF on a single scene

Data: Training images $\{Y_i\}$ with poses $\{P_i\}$, hidden information I , learning rate $\eta = [\eta_0, \eta_1]$

Output: Steganographic NeRF θ and decoder F_ψ

Initialize NeRF θ_0 and decoder network F_ψ

Optimize θ_0 on $\{Y_i, P_i\}$ with standard NeRF training

Compute mask m for θ_0 as in Eq. (3)

for each training iteration t **do**

 Randomly sample a training pose P_i and render $\theta(P_i)$

 Compute stega. losses $\mathcal{L}_{dec}^c, \mathcal{L}_{dec}^+, \mathcal{L}_{dec}^-$ as Eq. (1), (2)

 Compute \mathcal{L}_{rgb} and combine total loss \mathcal{L} as in Eq. (4)

 Update θ with $\eta_0 \cdot (\frac{\partial \mathcal{L}}{\partial \theta} \odot m)$ and F_ψ with $\eta_1 \cdot \frac{\partial \mathcal{L}}{\partial \psi}$

end for

network F_ψ^c to solve the easier task of classifying whether the given NeRF rendering contains hidden information. F_ψ^c is optimized by the following cross-entropy loss [30, 67]:

$$\mathcal{L}_{dec}^c = -\log(F_\psi^c(\theta(P))) - \log(1 - F_\psi^c(\theta_0(P))). \quad (2)$$

We then use its prediction to guide the process of decoding pixel-wise information by adding the classification output as input to F_ψ , such that $F_\psi(x) = F_\psi(x, F_\psi^c(x))$.

Although the above discussion focuses on hiding images, our framework can be easily extended to embed other modalities like strings, text, or even audio, all of which may be represented as 1D vectors. We can simply modify the architecture of F_ψ to have a 1D prediction branch.

3.3. Preserving Perceptual Identity

Since we want to hide information without affecting the visual perception of the rendered output, an intuitive regularization is to penalize how much θ deviates from θ_0 . However, we find that naively summing the deviation penalty across all weights makes it difficult for the NeRF network to adjust its weights for the steganographic objective. Instead, motivated by the fact that INR weights are not equally important and exhibit strong sparsity [28, 87], we propose an adaptive gradient masking strategy to encode the hidden information on specific weights’ groups.

Formally, given the initial set of weights $\theta_0 \in \mathbb{R}^N$, we compute a soft importance mask $m \in \mathbb{R}^N$. Specifically, for each weight $w_i \in \theta_0$, we compute:

$$m_i = \frac{|w_i|^{-\alpha}}{\sum_i^N |w_i|^{-\alpha}}, \quad (3)$$

where the magnitude $|\cdot|$ of each model weight w_i serves as its individual importance indicator, and $\alpha > 0$ is a scaling factor of power controlling the relative distribution of importance across the weights. We mask the gradient as $\frac{\partial \mathcal{L}}{\partial \theta} \odot m$ when optimizing θ based on the total loss \mathcal{L} in the second stage, where \odot is a Hadamard product. Effectively,

Algorithm 2 Typical usage scenario of StegaNeRF

- 1: Alice captures some images of a 3D scene
 - 2: Alice trains a StegaNeRF to hide a personalized image
 - 3: Alice shares the model θ online for other people to enjoy and explore the 3D scene themselves
 - 4: Bob grabs the model θ and reposts it with his own account without crediting Alice or asking for permission
 - 5: Alice sees Bob’s post, deploys the decoder F_ψ , and verifies the owner of θ is Alice, not Bob
 - 6: Bob takes down the post or gets banned for copyright infringement
-

more significant weights are “masked out” to minimize the impact of steganographic learning on the rendered visual quality.

We retain the standard photometric error in steganography learning to maintain NeRF rendering fidelity: $\mathcal{L}_{rgb} = |\theta(P) - \theta_0(P)|$. The overall training loss at the second stage can be formulated as follows:

$$\mathcal{L} = \lambda_0 \mathcal{L}_{dec}^c + \lambda_1 \mathcal{L}_{dec}^+ + \lambda_2 \mathcal{L}_{dec}^- + \lambda_3 \mathcal{L}_{rgb}. \quad (4)$$

4. Experiments

4.1. Implementation Details.

Dataset. We use common datasets LLFF [41] and NeRF-Synthetic [42], with forward scenes $\{flower, fern, fortress, room\}$ from LLFF and 360° scenes $\{lego, drums, chair\}$ from NeRF-Synthetic. We further experiment on the *Brandenburg Gate* scene from NeRF-W dataset [39], with over 800 views of in-the-wild collected online.

Training. On LLFF and NeRF-Synthetic, we adopt Plenoxels [83] as the NeRF backbone architecture for efficiency. On NeRF-W, we use the available PyTorch implementation [53]. The first stage of training is performed according to the standard recipes of those implementations. We then perform the second stage of steganography training for 55 epochs unless otherwise noted. On LLFF and NeRF-W, we downsize the training images by 4 times following common practice, and we use the original size on NeRF-Synthetic.

For hyper-parameters in Eq. (4), we set the weight of NeRF reconstruction error $\lambda_3 = 1$ for all experiments. We set $\lambda_0 = 0.01, \lambda_1 = 0.5, \lambda_2 = 0.5$ for all scenes on LLFF dataset, and $\lambda_0 = 0.1, \lambda_1 = 1, \lambda_2 = 1$ for the scenes in NeRF-Synthetic, and $\lambda_0 = 0.05, \lambda_1 = 1, \lambda_2 = 1$ for NeRF-W. For Eq. (3), we set $\alpha = 3$ for all scenes. We run experiments on one NVIDIA A100 GPU.

Evaluation. We assume a typical authentication scenario shown in Alg. 2. We consider the recovery quality of the embedded information, including the metrics of classification accuracy (Acc.) of the classifier, and the structural sim-

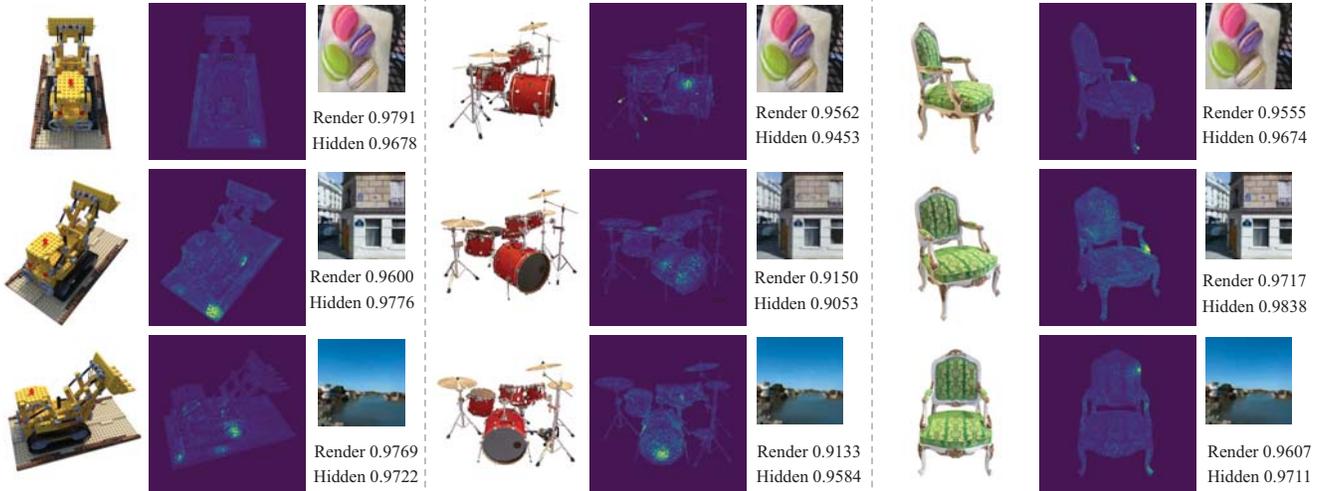


Figure 4. Results on the NeRF-Synthetic dataset. Within each column, we show the SteganeRF rendering, residual error from the initial NeRF rendering, and the recovered hidden image. We show the SSIM for the SteganeRF renderings and the recovered hidden images.

Table 1. Quantitative results of SteganeRF rendering and hidden information recovery. *Standard NeRF* is the initial NeRF θ_0 with standard training (upper-bound rendering performance). Prior 2D steganography fails after NeRF training while SteganeRF successfully embeds and recovers hidden information with minimal impact on the rendering quality. Results are averaged over the selected LLFF and NeRF-Synthetic scenes.

Method	NeRF Rendering			Hidden Recovery	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Acc. (%) \uparrow	SSIM \uparrow
Standard NeRF	28.15	0.8575	0.1377	50.00	N/A
LSB [9]	28.13	0.8569	0.1452	N/A	0.0132
DeepStega [2]	26.94	0.8432	0.1641	N/A	0.2098
SteganeRF (Ours)	28.13	0.8563	0.1460	100.0	0.9730
Standard NeRF	31.13	0.9606	0.0310	50.00	N/A
LSB [9]	31.12	0.9604	0.0310	N/A	0.0830
DeepStega [2]	31.13	0.9606	0.0313	N/A	0.2440
SteganeRF (Ours)	30.96	0.9583	0.0290	99.72	0.9677

ilarity (SSIM) [74] of the hidden image recovered by the decoder. We evaluate the final NeRF renderings with PSNR, SSIM and LPIPS [90]. All the metrics are computed on the test set and averaged over all the scenes and embedded images.

4.2. Case I: Embedding in a Single Scene.

We first explore the preliminary case of ownership identification on a specific NeRF scene. We select random images from ImageNet [27] as the hidden information to be injected to NeRF renderings.

Failure of 2D Baseline Due to the lack of prior study on NeRF steganography, we consider a baseline from 2D image steganography by training NeRF from scratch with

Table 2. Quantitative results of NeRF rendering and hidden information recovery. We consider two conditions, embedding each scene with a common hidden image (*One-for-All*) or a scene-specific hidden image (*One-for-One*). We report the quality difference compared to the single-scene settings as $\Delta_{SSIM}^{(10^{-2})}$. Results are averaged over the selected LLFF scenes.

Setting	NeRF Rendering				Hidden Recovery		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	$\Delta_{SSIM}^{(10^{-2})} \uparrow$	Acc. (%) \uparrow	SSIM \uparrow	$\Delta_{SSIM}^{(10^{-2})} \downarrow$
<i>One-for-All</i>	24.99	0.8013	0.1786	-0.10	100.00	0.9860	+0.19
<i>One-for-One</i>	24.99	0.8016	0.1779	-0.07	100.00	0.9122	-7.19
<i>One-for-All</i>	27.90	0.8513	0.1236	+0.01	100.00	0.9844	-0.76
<i>One-for-One</i>	27.90	0.8515	0.1195	+0.03	100.00	0.9448	-4.45
<i>One-for-All</i>	30.27	0.8498	0.1289	+0.02	100.00	0.9430	-5.42
<i>One-for-One</i>	30.12	0.8480	0.1302	-0.16	100.00	0.9102	-8.67

the watermarked images. We implement two off-the-shelf steganography methods including a traditional machine learning approach called Least Significant Bit (LSB [9]), and a deep learning pipeline as DeepStega [2]. An ideal case is that the embedded information can be recovered from the synthesized novel views, indicating the successful NeRF steganography. However, as can be seen in Fig. 3, the embedded information containing the hidden images cannot be recovered from the NeRF renderings. By analyzing the residual maps of training views (between GT training views and the watermarked) and novel views (between GT novel views and the actual rendering), we observe the subtle residuals to recover the hidden information are smoothed out in NeRF renderings, and similar failures occur on other datasets as well. Therefore, for the rest of our study, we mainly focus on analyzing our new framework that performs direct steganography on 3D NeRF.

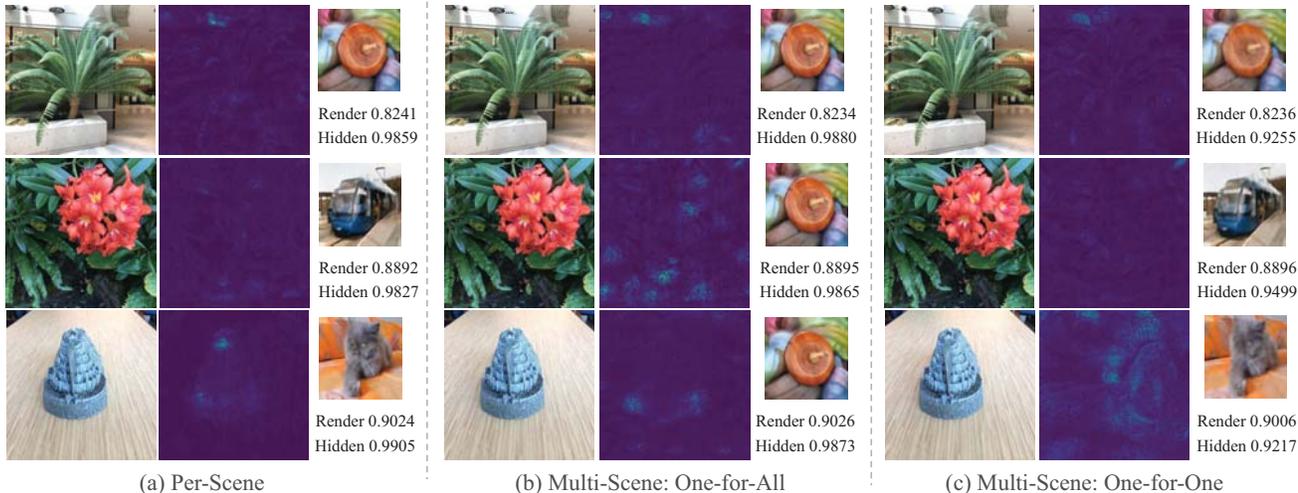


Figure 5. Results on three multi-scene settings. Within each column, we show the StegaNeRF rendering, residual error from the initial NeRF rendering, and recovered hidden image. We show the SSIM for the StegaNeRF renderings and the recovered hidden images.

Results Tab. 1 contains quantitative results on LLFF and NeRF-Synthetic scenes. While NeRF trained by 2D steganography methods hardly recovers the embedded information, StegaNeRF accurately recovers the hidden image with minimal impact on the rendering quality measured by PSNR. Fig. 4 provides qualitative results of StegaNeRF on three embedded images on the NeRF-Synthetic scenes. An interesting observation is the common regions where hidden information (high residual error) emerges in the renderings, *e.g.*, the right rear side of *lego* and left handrail of *chair*. We also notice that, within each scene, these regions are persistent across multiple viewpoints.

4.3. Case II: Embedding in Multiple Scenes.

Settings We extend our steganography scheme to embed information within multiple scenes at once. Specially, we use three LLFF scenes $\{flower, fern, fortress\}$ to test the two sub-settings of multi-scene embedding with: (1) *One-for-All*, a common hidden image and (2) *One-for-One* scene-specific hidden images. All scenes share the same decoder F_ψ and classifier F_ψ^c . The difference between the two sub-settings is the number of hidden images that F_ψ and F_ψ^c need to identify and recover. We sample one scene for every training epoch, and due to the increased data amount, we increase training epochs until convergence.

Results Tab. 2 provides quantitative results on the multi-scene setting. The performance drop compared to single-scene training is sometimes noticeable, but it is not surprising due to the inherent requirement of per-scene fitting for our NeRF framework. Fig. 5 presents visual comparisons of multi-scene steganography against single-scene scheme.

4.4. Case III: Embedding Multiple Identities.

Settings Constructing NeRF of large-scale cultural landmarks is a promising application, and community contributions are crucial to form the training images. Since every NeRF prediction is indebted to some particular training images, it would be meaningful to somehow identify the contributing users in the rendered output. Specifically, we present a proof-of-concept based on the following scenario: Given a collection of user-uploaded training images, our task is to apply StegaNeRF to ensure the final NeRF renderings hide subtle information about the relevant identities whose uploaded images help generate the current rendering.

To simulate this multi-user scenario in the public NeRF-W dataset, we randomly select M anchor views with different viewing positions, and then find the K nearest neighbour views to each anchor to form their respective clusters. We set $M = 3, K = 20$ in experiments. We assume a common contributor identity for each cluster, and we want the NeRF rendering to contain customized hidden information about those M contributors when we render within the spatial regions spanned by their own cluster. Thus, our classifier network F_ϕ^c is modified to output M -class cluster predictions and another class for identity outside of the M clusters. Since the decoder should extract no information for views outside of those M clusters to prevent false positives, we also compute \mathcal{L}_{dec}^- (1) and \mathcal{L}_{dec}^c (2) for those poses.

Results We employ the same network backbone as NeRF-W [66] to handle in-the-wild images with different time and lighting effects. Fig. 6 presents qualitative results of embedding multiple identities in a collaborative large-scale NeRF.



Figure 6. Qualitative results of NeRF steganography on NeRF-W dataset. Within each block we show the StegaNeRF rendering, the initial NeRF rendering, the residual error between StegaNeRF and initial, and the recovered hidden image. At the bottom right, we provide the SSIM computed with their respective ground truth. “Unmarked Identity” denotes the remaining training views not related to any of the considered users, and the decoder does not recover meaningful information from renderings at those poses.

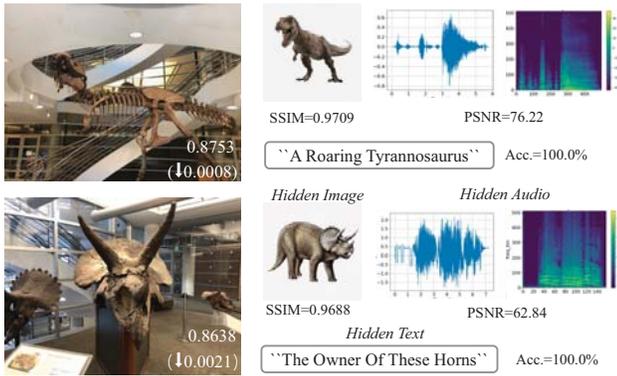


Figure 7. Multi-modal hidden information. We show StegaNeRF rendering, recovered hidden images, audio (waveform and spectrum), and text. Metrics for hidden information in each modality are labeled respectively. For the StegaNeRF rendering, we report both SSIM and the relative change from initial NeRF renderings.

4.5. Case IV: Embedding Multi-Modal Information.

Settings We further show the potential of StegaNeRF in embedding multi-modal hidden information, such as images, audio, and text. We modify the decoder network to build a modal-specific decoder for each modality.

Results Fig. 7 shows recovered multi-modal embedded signals in the *trex* and *horns* scenes from LLFF. Evidently, the StegaNeRF framework can easily extend to embed multi-modal information with high recovery performance without scarifying the rendering quality.

4.6. Ablation Studies

The effect of removing each component of StegaNeRF is presented in Tab. 3. *Standard NeRF* uses the initial NeRF

Table 3. Ablation study of different components of StegaNeRF. Results are averaged on the selected LLFF scenes.

Method	NeRF Rendering			Hidden Recovery	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Acc. (%) \uparrow	SSIM \uparrow
StegaNeRF	28.13	0.8568	0.1430	100.0	0.9486
No Classifier	26.85	0.8077	0.2417	N/A	0.5629
No Classifier Guided	27.12	0.8239	0.2073	100.0	0.6716
No Gradient Masking	27.86	0.8375	0.1710	100.0	0.8822
No Soft Masking	28.05	0.8558	0.1526	94.44	0.8751
Standard NeRF	28.15	0.8575	0.1377	50.00	N/A

rendering. *No Classifier* removes the classifier F_{ψ}^c , while *No Classifier Guided* retains the classification task (hence impact on NeRF rendering) but does not condition the decoder on classifier output. *No Gradient Masking* removes the proposed masking strategy (Sec. 3.3), and *No Soft Masking* uses the binary mask [29, 31] with a threshold of 0.5. Our StegaNeRF makes a good balance between the rendering quality and the decoding accuracy. Fig. 8 shows the visual impact of removing each component. It appears that when any component is removed, the performance drops accordingly, revealing the effectiveness of our design.

Fig. 9 shows the impact of varying gradient masking strategies. We can see a different degree of affected by ratio variation on hidden recovery and rendering quality. Fig. 10 reports the performance of StegaNeRF against the common perturbations including JPEG compression and Gaussian noise. The lines show the average accuracies across selected scenes and the shaded regions indicate the range of 0.5 standard deviation. It appears that the hidden recovery of StegaNeRF is robust to various JPEG compression ratios and Gaussian blur degradation.

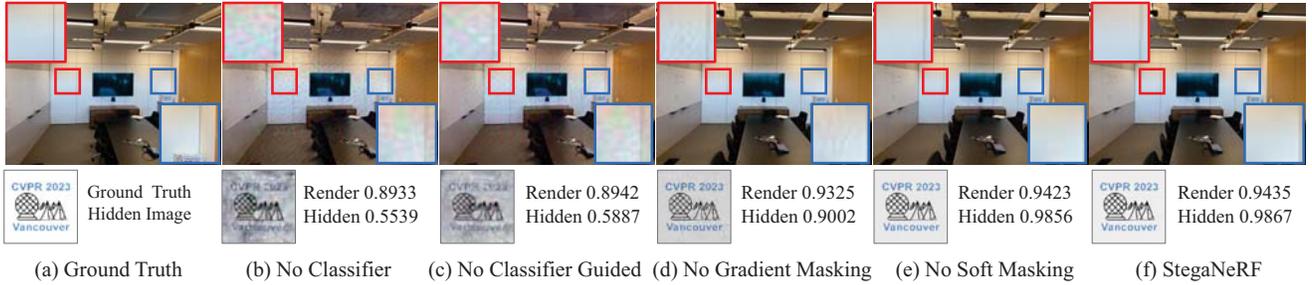


Figure 8. Impact on visual quality when changing different components of the proposed StegaNeRF framework as in Tab. 3. We show the SSIM for the renderings and the recovered hidden images.

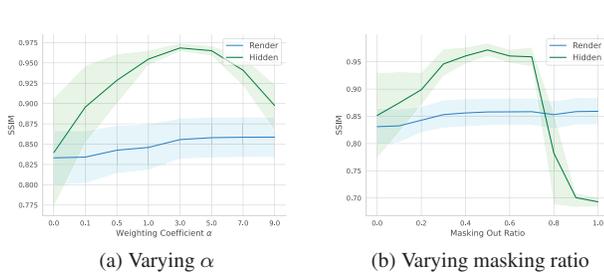


Figure 9. Analysis of gradient masking. (a) Varying α from Eq. (3). (b) Varying portions of masked out weights from the gradient updates at steganography learning. We provide the SSIM of rendered views (blue) and recovered hidden images (green).

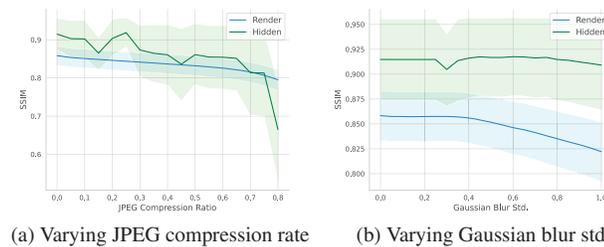
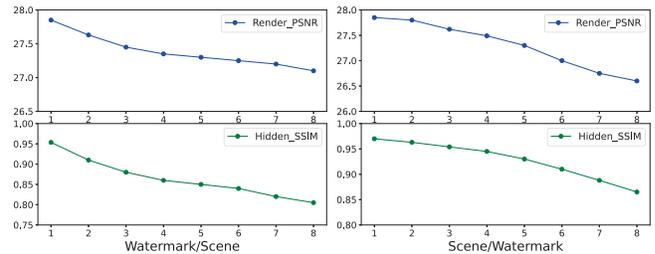


Figure 10. Analysis of robustness over (a) JPEG compression and (b) Gaussian blur. We provide the SSIM of rendered views (blue) and recovered hidden images (green).

5. More Discussions

How useful is steganography for NeRF? Although NeRF-based 3D content has yet to become mainstream, we believe it will play a major future role not only for social platforms, but also for 3D vision research and applications. As people upload their personal NeRF models online for viewing purposes, NeRF steganography for ownership identification is apparently an important feature. Moreover, future 3D vision research will likely demand large-scale NeRF datasets of real-world images, and in this context, NeRF steganography can be a crucial tool for responsible



(a) Varying #watermark per scene (b) Varying #scene per watermark

Figure 11. Results using a universal decoder to embed (a) multiple images for the *flower* scene and (b) a common image into different scenes from the LLFF dataset. The successful application under these two setups suggests that StegaNeRF can generalize across different scenes, user identities, and hidden images.

Table 4. The further empirical study on StegaNeRF. Results are averaged on the selected LLFF scenes.

Method	NeRF Rendering			Hidden Recovery	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Acc. (%) \uparrow	SSIM \uparrow
StegaNeRF	28.13	0.8568	0.1430	100.0	0.9486
No Empty Sampling	28.12	0.8557	0.1466	100.0	0.9364
NeRF Frozen	28.15	0.8575	0.1377	50.00	0.3975
Standard NeRF	28.15	0.8575	0.1377	50.00	N/A

and ethical uses of training data and deployed models.

Hidden or Memorized? StegaNeRF follows the standard setup of NeRF which prescribes per-scene training. Therefore, it may be reasonable to wonder whether the decoder is trained to “memorize” the hidden image without it really being injected into NeRF. To address this concern, we perform an experiment where the *NeRF weights are frozen* after the first stage and only update the decoder of StegaNeRF. If the decoder actually memorizes the image by itself, then we should expect good recovery performance under this setup. However, as shown in Tab. 4, with *NeRF Frozen*, we observe a significant performance degradation. Therefore, we can verify that the decoder alone cannot trivially memorize

the hidden image, and the information is actually hidden in the NeRF renderings. Furthermore, we extend StegaNeRF to generalize across different scenes and hidden images while sharing the same universal decoder. Fig. 11a shows the rendering and decoding quality of embedding the different numbers of hidden images to a scene (*flower*). Fig. 11b shows the performance of embedding a common image into a varying number of scenes from the LLFF dataset. While the performance slightly decreases in the generalized setting, it still remains at a comparable level.

Why not just share the rendered images instead of NeRF weights? Directly sharing the NeRF model gives end users the freedom to render at continuous viewpoints, which fundamentally differs from the traditional setting centered around discrete 2D image frames. It would not be realistic to assume people will stick with sharing 2D images instead of 3D weights, and therefore, it is necessary for the research community to see beyond 2D steganography and explore 3D NeRF steganography.

Robustness: As shown in Fig. 12, in addition to using a resolution of 128×128 for hidden images in most of our experiments, we also tested resolutions ranging from 96×96 to 224×224 . We found that increasing the size of embedded images had no noticeable impact on the performance of StegaNeRF, with only a 0.6% drop in rendering PSNR and a 1.2% reduction in hidden SSIM. We also demonstrated the robustness of StegaNeRF to different sampling approaches during rendering, such as avoiding sampling in empty areas [43, 48]. Instead of considering all the points on a ray during rendering, we disregard the contribution of points to the integration along the entire ray if their density falls below a threshold, like 0.001. As shown in Tab. 4, replacing the conventional rendering strategy with *No Empty Sampling* has minimal effect on the performance of StegaNeRF, revealing our robustness to different rendering fashions.

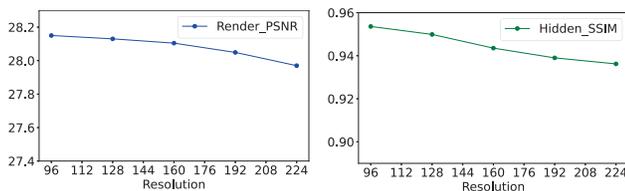


Figure 12. Ablation on the resolutions of embedded images by the average results over the selected scenes from the LLFF dataset.

Limitations: The current implementation is limited by the computational efficiency of NeRF, since the steganographic training stage requires performing gradient descent updates through the NeRF model. Another limitation is that, the recovery accuracy of the injected hidden information varies across different 3D scenes. More systematic analysis is required to understand the underlying factors that make some scenes easier to hide information than others.

6. Conclusion

With the progression of NeRF-based content, now serving as a pivotal element in diverse 3D applications, it is incumbent upon the research community to proactively tackle the imminent requirement to embed information within NeRF models as they are exchanged and disseminated across online platforms. This study introduces StegaNeRF, a pioneering framework for the automated embedding of steganographic information within NeRF renderings. We outline the essential components essential for achieving this functionality, unveiling potential challenges and revealing insights that can be instrumental for subsequent developers. By shedding light on the prospect of ownership identification within the realm of NeRF, this paper underscores the need for increased attention and concerted efforts toward addressing related challenges and avenues.

References

- [1] Benjamin Attal, Jia-Bin Huang, Michael Zollhoefer, Johannes Kopf, and Changil Kim. Learning Neural Light Fields With Ray-Space Embedding Networks. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19787–19797, 2022. 2
- [2] Shumeet Baluja. Hiding Images in Plain Sight: Deep Steganography. *Advances in Neural Information Processing Systems*, 30, 2017. 1, 2, 3, 5
- [3] Shumeet Baluja. Hiding Images Within Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(7):1685–1697, 2019. 1, 2
- [4] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 1, 2
- [5] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5460–5469, 2022. 2
- [6] Oliver Benedens. Geometry-Based Watermarking of 3D Models. *IEEE Computer Graphics and Applications*, 19:46–55, 1999. 2
- [7] Ning Bi, Qiyu Sun, Daren Huang, Zhihua Yang, and Jiwu Huang. Robust Image Watermarking Based on Multiband Wavelets and Empirical Mode Decomposition. *IEEE Transactions on Image Processing*, 16(8):1956–1966, 2007. 2
- [8] Rajarathnam Chandramouli, Mehdi Kharrazi, and Nasir Memon. Image Steganography and Steganalysis: Concepts and Practice. In *International Workshop on Digital Watermarking*, pages 35–49. Springer, Springer, 2003. 2
- [9] Chin-Chen Chang, Ju-Yuan Hsiao, and Chi-Shiang Chan. Finding Optimal Least-Significant-Bit Substitution in Image Hiding by Dynamic Programming Strategy. *Pattern Recognition*, 36(7):1583–1595, 2003. 2, 3, 5

- [10] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021. 1
- [11] Di Chen, Yu Liu, Lianghua Huang, Bin Wang, and Pan Pan. GeoAug: Data Augmentation for Few-Shot NeRF With Geometry Constraints. In *European Conference on Computer Vision*, pages 322–337. Springer, Springer, 2022. 1
- [12] Ingemar Cox, Matthew Miller, Jeffrey Bloom, Jessica Fridrich, and Ton Kalker. *Digital Watermarking and Steganography*. Morgan Kaufmann, 2007. 1, 2
- [13] Wenxue Cui, Shaohui Liu, Feng Jiang, Yongliang Liu, and Debin Zhao. Multi-Stage Residual Hiding for Image-Into-Audio Steganography. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2832–2836, 2020. 2
- [14] Zhiwen Fan, Yifan Jiang, Peihao Wang, Xinyu Gong, Dejia Xu, and Zhangyang Wang. Unified implicit neural stylization. In *European Conference on Computer Vision*, pages 636–654. Springer, 2022. 1
- [15] Brandon Yushan Feng and Amitabh Varshney. SIGNET: Efficient Neural Representation for Light Fields. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14204–14213, 2021. 2
- [16] 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
- [17] Jessica Fridrich, Miroslav Goljan, and Rui Du. Detecting LSB Steganography in Color, and Gray-Scale Images. *IEEE Multimedia*, 8(4):22–28, 2001. 2
- [18] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2021. 2
- [19] 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
- [20] Margarita Geleta, Cristina Puntì, Kevin McGuinness, Jordi Pons, Cristian Canton, and Xavier Giro-i Nieto. PixInWav: Residual Steganography for Hiding Pixels in Audio. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2485–2489, 2022. 2
- [21] Jamie Hayes and George Danezis. Generating Steganographic Images Via Adversarial Training. *Advances in Neural Information Processing Systems*, 30, 2017. 2
- [22] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021. 2
- [23] 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
- [24] Xudong Huang, Wei Li, Jie Hu, Hanting Chen, and Yunhe Wang. Refsr-nerf: Towards high fidelity and super resolution view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8244–8253, 2023. 1
- [25] Junpeng Jing, Xin Deng, Mai Xu, Jianyi Wang, and Zhenyu Guan. HiNet: Deep Image Hiding by Invertible Network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4733–4742, 2021. 2
- [26] Gary C Kessler and Chet Hosmer. An Overview of Steganography. *Advances in Computers*, 83:51–107, 2011. 2
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification With Deep Convolutional Neural Networks. *Communications of the ACM*, 60(6):84–90, 2017. 5
- [28] Jaeho Lee, Jihoon Tack, Namhoon Lee, and Jinwoo Shin. Meta-Learning Sparse Implicit Neural Representations. *Advances in Neural Information Processing Systems*, 34:11769–11780, 2021. 4
- [29] Chenxin Li, Mingbao Lin, Zhiyuan Ding, Nie Lin, Yihong Zhuang, Yue Huang, Xinghao Ding, and Liujuan Cao. Knowledge condensation distillation. In *European Conference on Computer Vision*, pages 19–35. Springer, 2022. 7
- [30] Chenxin Li, Wenao Ma, Liyan Sun, Xinghao Ding, Yue Huang, Guisheng Wang, and Yizhou Yu. Hierarchical deep network with uncertainty-aware semi-supervised learning for vessel segmentation. *Neural Computing and Applications*, pages 1–14, 2022. 4
- [31] Chenxin Li, Yunlong Zhang, Zhehan Liang, Wenao Ma, Yue Huang, and Xinghao Ding. Consistent posterior distributions under vessel-mixing: a regularization for cross-domain retinal artery/vein classification. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 61–65. IEEE, 2021. 7
- [32] Tianye Li, Mira Slavcheva, Michael Zollhofer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022. 2
- [33] Zhi-Hao Lin, Wei-Chiu Ma, Hao-Yu Hsu, Y. Wang, and Shenlong Wang. NeurMiPs: Neural Mixture of Planar Experts for View Synthesis. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15681–15691, 2022. 2
- [34] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural Sparse Voxel Fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. 2
- [35] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason M. Saragih. Mixture of Volumetric Primitives for Efficient Neural Rendering. *ACM Transactions on Graphics (TOG)*, 40:1–13, 2021. 2

- [36] Shao-Ping Lu, Rong Wang, Tao Zhong, and Paul L Rosin. Large-Capacity Image Steganography Based on Invertible Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10816–10825, 2021. 2
- [37] Luma AI. <https://lumalabs.ai/>. 2
- [38] Xiyang Luo, Yinxiao Li, Huiwen Chang, Ce Liu, Peyman Milanfar, and Feng Yang. Dvmark: A Deep Multi-scale Framework for Video Watermarking. *ArXiv Preprint ArXiv:2104.12734*, 2021. 2
- [39] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 2, 4
- [40] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy Networks: Learning 3d Reconstruction in Function Space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 1
- [41] 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. 4
- [42] 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, 4
- [43] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant Neural Graphics Primitives With a Multi-resolution Hash Encoding. *ACM Transactions on Graphics (TOG)*, 41:1 – 15, 2022. 1, 2, 9
- [44] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing Neural Radiance Fields for View Synthesis From Sparse Inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 1
- [45] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning Continuous Signed Distance Functions for Shape Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 1
- [46] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 2
- [47] Tomáš Pevný, Tomáš Filler, and Patrick Bas. Using High-Dimensional Image Models to Perform Highly Undetectable Steganography. In *International Workshop on Information Hiding*, pages 161–177. Springer, Springer, 2010. 2
- [48] Martin Píala and Ronald Clark. Terminerf: Ray termination prediction for efficient neural rendering. In *2021 International Conference on 3D Vision (3DV)*, pages 1106–1114. IEEE, 2021. 9
- [49] Polycam. <https://poly.cam/>. 2
- [50] Emil Praun, Hugues Hoppe, and Adam Finkelstein. Robust Mesh Watermarking. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pages 49–56, 1999. 2
- [51] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 2
- [52] Yinlong Qian, Jing Dong, Wei Wang, and Tieniu Tan. Deep Learning for Steganalysis Via Convolutional Neural Networks. In *Media Watermarking, Security, and Forensics 2015*, volume 9409, pages 171–180. SPIE, SPIE, 2015. 2
- [53] Chen Quei-An. Nerf-pl: A Pytorch-Lightning Implementation of NeRF. https://github.com/kwea123/nerf_pl/. 4
- [54] 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
- [55] Mennatallah M Sadek, Amal S Khalifa, and Mostafa GM Mostafa. Video Steganography: A Comprehensive Review. *Multimedia Tools and Applications*, 74(17):7063–7094, 2015. 2
- [56] Katja Schwarz, Axel Sauer, Michael Niemeyer, Yiyi Liao, and Andreas Geiger. Voxgraf: Fast 3d-Aware Image Synthesis With Sparse Voxel Grids. *ArXiv Preprint ArXiv:2206.07695*, 2022. 1
- [57] Qihong Shen, Xingyi Yang, and Xinchao Wang. Anything-3d: Towards single-view anything reconstruction in the wild. *arXiv preprint arXiv:2304.10261*, 2023. 1
- [58] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit Neural Representations With Periodic Activation Functions. *Advances in Neural Information Processing Systems*, 33, 2020. 1
- [59] Vincent Sitzmann, Semon Rezhchikov, William T. Freeman, Joshua B. Tenenbaum, and Frédo Durand. Light Field Networks: Neural Scene Representations With Single-Evaluation Rendering. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [60] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Generalizable patch-based neural rendering. In *European Conference on Computer Vision*, pages 156–174. Springer, 2022. 2
- [61] 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. 1, 2
- [62] Abdelfatah A Tamimi, Ayman M Abdalla, and Omaima Al-Allaf. Hiding an Image Inside Another Image Using Variable-Rate Steganography. *International Journal of Advanced Computer Science and Applications (IJACSA)*, 4(10), 2013. 2

- [63] Matthew Tancik, Vincent Casser, Xichen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-Nerf: Scalable Large Scene Neural View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 1, 2
- [64] Matthew Tancik, Ben Mildenhall, and Ren Ng. Stegastamp: Invisible Hyperlinks in Physical Photographs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2117–2126, 2020. 2
- [65] Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. *Advances in Neural Information Processing Systems*, 33, 2020. 1
- [66] Weixuan Tang, Shunquan Tan, Bin Li, and Jiwu Huang. Automatic Steganographic Distortion Learning Using a Generative Adversarial Network. *IEEE Signal Processing Letters*, 24(10):1547–1551, 2017. 2, 6
- [67] Jeroen JG Van Merriënboer, Liesbeth Kester, and Fred Paas. Teaching complex rather than simple tasks: Balancing intrinsic and germane load to enhance transfer of learning. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 20(3):343–352, 2006. 4
- [68] Mukund Varma, Peihao Wang, Xuxi Chen, Tianlong Chen, Subhashini Venugopalan, and Zhangyang Wang. Is attention all that nerf needs? In *The Eleventh International Conference on Learning Representations*, 2022. 2
- [69] Huan Wang, Jian Ren, Zeng Huang, Kyle Olszewski, Menglei Chai, Yun Fu, and S. Tulyakov. R2L: Distilling Neural Radiance Field to Neural Light Field for Efficient Novel View Synthesis. In *European Conference on Computer Vision*. Springer, Springer, 2022. 2
- [70] Huaqing Wang and Shuozhong Wang. Cyber Warfare: Steganography vs. Steganalysis. *Communications of the ACM*, 47(10):76–82, 2004. 2
- [71] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun 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*, pages 13524–13534, 2022. 1
- [72] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning Neural Implicit Surfaces by Volume Rendering for Multi-View Reconstruction. *Advances in Neural Information Processing Systems*, 34, 2021. 1
- [73] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning Multi-View Image-Based Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021. 2
- [74] 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. 5
- [75] Xinyu Weng, Yongzhi Li, Lu Chi, and Yadong Mu. High-Capacity Convolutional Video Steganography With Temporal Residual Modeling. In *Proceedings of the 2019 International Conference on Multimedia Retrieval*, pages 87–95, 2019. 2
- [76] Raymond B Wolfgang and Edward J Delp. A Watermark for Digital Images. In *Proceedings of 3rd IEEE International Conference on Image Processing*, volume 3, pages 219–222. IEEE, IEEE, 1996. 2
- [77] Yue Wu, Guotao Meng, and Qifeng Chen. Embedding Novel Views in a Single JPEG Image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14519–14527, 2021. 2
- [78] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9421–9431, 2021. 2
- [79] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. BungeeNeRF: Progressive Neural Radiance Field for Extreme Multi-Scale Scene Rendering. In *The European Conference on Computer Vision (ECCV)*, volume 2, 2022. 2
- [80] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Humphrey Shi, and Zhangyang Wang. Sinnerf: Training neural radiance fields on complex scenes from a single image. In *European Conference on Computer Vision*, pages 736–753. Springer, 2022. 1
- [81] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4479–4489, 2023. 1
- [82] Hyukryul Yang, Hao Ouyang, Vladlen Koltun, and Qifeng Chen. Hiding Video in Audio Via Reversible Generative Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1100–1109, 2019. 2
- [83] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance Fields Without Neural Networks. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5491–5500, 2022. 1, 4
- [84] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for Real-Time Rendering of Neural Radiance Fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. 2
- [85] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. Pixelnerf: Neural Radiance Fields From One or Few Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1
- [86] Ning Yu, Vladislav Skripniuk, Sahar Abdelnabi, and Mario Fritz. Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14448–14457, 2021. 2

- [87] Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. A Structured Dictionary Perspective on Implicit Neural Representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19228–19238, 2022. [4](#)
- [88] Jian Zhang, Yuanqing Zhang, Huan Fu, Xiaowei Zhou, Bowen Cai, Jinchu Huang, Rongfei Jia, Binqiang Zhao, and Xing Tang. Ray Priors Through Reprojection: Improving Neural Radiance Fields for Novel View Extrapolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18376–18386, 2022. [1](#)
- [89] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and Improving Neural Radiance Fields. *ArXiv Preprint ArXiv:2010.07492*, 2020. [2](#)
- [90] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features As a Perceptual Metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. [5](#)
- [91] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul E. Debevec, William T. Freeman, and Jonathan T. Barron. NeRF-Factor: Neural Factorization of Shape and Reflectance Under an Unknown Illumination. *ACM Trans. Graph*, 40:237:1–237:18, 2021. [1](#)