

SCINeRF: Neural Radiance Fields from a Snapshot Compressive Image

Yunhao Li^{1,2} Xiaodong Wang^{1,2} Ping Wang^{1,2} Xin Yuan² Peidong Liu^{2†}
¹Zhejiang University ²Westlake University
 {liyunhao, wangxiaodong, wangping, xyuan, liupeidong}@westlake.edu.cn



Figure 1. Given a single snapshot compressed image, our method is able to recover the underlying 3D scene representation. Leveraging the strong novel-view image synthesis capabilities of NeRF, we can render multi-view consistent images in high quality from **the single measurement**.

Abstract

In this paper, we explore the potential of Snapshot Compressive Imaging (SCI) technique for recovering the underlying 3D scene representation from a single temporal compressed image. SCI is a cost-effective method that enables the recording of high-dimensional data, such as hyperspectral or temporal information, into a single image using low-cost 2D imaging sensors. To achieve this, a series of specially designed 2D masks are usually employed, which not only reduces storage requirements but also offers potential privacy protection. Inspired by this, to take one step further, our approach builds upon the powerful 3D scene representation capabilities of neural radiance fields (NeRF). Specifically, we formulate the physical imaging process of SCI as part of the training of NeRF, allowing us to exploit its impressive performance in capturing complex scene structures. To assess the effectiveness of our method, we conduct extensive evaluations using both synthetic data and real data captured by our SCI system. Extensive experimental results demonstrate that our proposed approach surpasses the state-of-the-art methods in terms of image reconstruction and novel view image synthesis. Moreover, our method also exhibits the ability to restore high frame-rate multi-view consistent images by leveraging SCI and the rendering capabilities of NeRF. The code is available at <https://github.com/WU-CVGL/SCINeRF>.

1. Introduction

Conventional high-speed imaging systems often face challenges such as high hardware cost and storage requirement. Drawing inspiration from Compressed Sensing (CS) [5, 8], video Snapshot Compressive Imaging (SCI) [50] system has emerged to address these limitations. A conventional video SCI system usually contains a hardware encoder and a software decoder. The hardware encoder employs a series of specially designed 2D masks to modulate the incoming photons across exposure time into a single compressed image. It enriches the low-cost cameras the ability to capture high-speed scenes, which further reduces the storage requirement. The whole encoding process can also be achieved via software implementation on pre-captured images, which can reduce the storage/transmission requirement and offer additional privacy protection. On the other hand, the software decoder restores the high frame-rate images using the compressed measurement and corresponding binary masks.

In recent years, several image reconstruction algorithms have been proposed for SCI, which range from the model-based methods [18, 20, 49] to deep-learning based methods [6, 7, 21, 30, 38, 39]. These algorithms can reconstruct encoded images or video frames with relatively high quality. However, these methods usually do not consider the underlying 3D scene structure to ensure multi-view consistency and can only recover images corresponding to the applied encoding masks. To address these challenges, we propose SCINeRF which recovers the underlying 3D scene representation from a single compressed image. High frame-rate

[†]Corresponding author.

multi-view consistency images can then be rendered from the learned 3D representation, as shown in Fig. 1.

Our method exploits neural radiance fields (NeRF) [26] to represent the 3D scene implicitly. Different from prior explicit 3D representation techniques, NeRF takes pixel location and its corresponding camera pose as input, and predicts the pixel intensity by volume rendering. This characteristic makes it well suited for modeling the pixel-wise modulation process of an SCI system. Since it is impossible to recover the camera poses from a single compressed image via COLMAP [32], we conduct a joint optimization on both the camera poses and NeRF, via minimizing the difference between the synthesized compressed image and real measurement from the encoder. Since SCINeRF performs test-time optimization, it does not suffer from the generalization performance degradation as prior end-to-end deep learning based methods. With the help of our proposed SCINeRF, we can capture the scenes within a single exposure time (can be shorter than 20ms or even 10ms) from a fast-moving camera, then recover the underlying 3D scene representations.

To better evaluate the performance of our method, we setup a **real hardware platform** to collect real snapshot compressed images. We also conduct quantitative evaluations using synthetic images generated via Blender. Experimental results demonstrate that SCINeRF achieves superior performance over previous state-of-the-art methods in terms of image restoration and novel view image synthesis. In summary, our contributions are list as follows.

- We present a novel method to restore 3D aware multi-view images from a single snapshot compressed image, under the framework of NeRF.
- We experimentally validate that our approach is able to synthesize high quality novel view images, surpassing existing state-of-the-art SCI image/video reconstruction methods.
- SCINeRF also presents an alternative approach for both efficient and privacy protection transmission between the edge devices and cloud infrastructure for practical NeRF deployment.

2. Related Work

We review two main areas of the prior works: snapshot compressive imaging and NeRF, which are the most related components in our work.

Snapshot Compressive Imaging. Early SCI image decoding/reconstruction methods focus on regularized optimization-based approach [18, 20, 47, 49]. It estimates the input compressed images from SCI encoding system by solving an optimization problem iteratively. In SCI, different regularizers and priors have been used, including sparsity [47] and total variation (TV) [49]. When solving the

optimization problems, instead of using conventional gradient descent algorithm, most of the existing methods employ alternating direction method of multipliers (ADMM) [4], which leads to good results and easier to adapt to different systems. The decompress SCI (DeSCI) [20] and GAP-TV [49] are state-of-the-art optimization-based approaches. However, this type of methods suffer from long running time and low flexibility to high resolution images.

With the rise of deep learning and neural networks, most of the recent SCI decoding/reconstruction methods are based on deep learning. Various kinds of network architectures have been adopted into SCI decoder, including U-net [31], and generative adversarial networks (GAN) [10]. These deep learning-based methods takes thousands or even millions of synthetic SCI measurements and masks as training dataset (since obtaining large amount of real SCI measurement was time consuming and extremely difficult), and optimize the network using various kinds of loss functions including mean square error (MSE), feature loss [15] and GAN loss [24]. Qiao et al. [30] built an End-to-End CNN (E2E-CNN) with reconstruction loss to retrieve compressed images. Cheng et al. [7] employed bidirectional recurrent neural networks (BIRNAT), which includes a bi-directional residual neural network (RNN) [12] to reconstruct the temporal frames in a sequential manner. RevSCI [6] reduced the time and memory complexity during large-scale video SCI training by designing a multi-group reversible 3D CNN architecture. ADMM-Net [21] modeled the decoding process as a tensor recovery problem from random linear measurements and interprets the ADMM process into a deep neural network. MetaSCI [42] adopted a meta modulated CNN to make the reconstruction network more adaptable to large scale data and novel masks. Plug-and-play fast and flexible denoising CNN (PnP-FFDNet) [51] combined deep denoising network with ADMM, leading to the fast and flexible reconstruction. Later, Yuan et al. [52] developed fast deep video denoising network (FastDVDNet), which improves the performance of PnP-FFDNet by introducing the the most recent deep denoising network. Wang et al. introduced spatial-temporal Transformer (STFormer) [39] and EfficientSCI [38] to exploit the spatial and temporal correlation within image decoding process utilizing Transformer [37] architecture. These deep learning-based methods can reconstruct images with relatively high quality. However, since these methods contain pre-trained models from synthetic datasets, they lack generalization abilities and perform poor on real datasets. Moreover, existing deep learning-based methods can only reconstruct images corresponding to the masks, making it difficult to restore high frame-rate images and videos. On the contrary, our method performs test-time optimization with no generalization problems, and it can restore high frame-rate multi-view images from estimated 3D scenes.

NeRF and Its Variants. Mildenhall et al. [26, 35] proposed NeRF, an epochal novel scene representation method. NeRF estimates the scene appearance and geometry continuously and its performance surpasses state-of-the-art scene representation schemes by producing better novel-view images. To make NeRF robust to more complicated real-world environments, researchers have proposed many dedicated NeRF variants recently. Some researches extend NeRF to make it compatible for large scale scene reconstruction and representation [11, 34, 36, 45]. Others focus on developing NeRF which can deal with non-rigid object reconstruction [1, 9, 28, 29]. There are also many NeRF-based frameworks for high dynamic range (HDR) image modeling [13], scene editing [43, 44, 46, 53] and scene appearance decomposition [2, 3].

We will mainly review NeRF-based methods close to our work in detail in this section. The original NeRF requires the accurate camera poses as input. Since many existing image datasets lack camera pose information, many NeRF-based methods, including the original NeRF framework, estimate the camera poses from input images. They employ widely used structure-from-motion (SfM) software COLMAP [32] to get camera poses. But in many cases, the camera poses estimated from SfM are unavailable or inaccurate. To address this issue, many frameworks focus on NeRF optimization without camera poses. Wang et al. propose NeRF- [42], which jointly estimates the scene and camera parameters at each training iteration. Jeong [14] develops a NeRF-based camera self-calibration scheme. iMAP [33], introduced by Sucar et al., integrates NeRF into Simultaneous Localization and Mapping (SLAM) to optimize the scene and camera trajectory together. Meng et al. [23] implement GAN to solve the problem of optimizing NeRF without camera poses. Other researches focus on refining inaccurate camera poses during the NeRF optimization process. Bundle adjusting neural radiance fields (BARF) [19] proposed by Lin et al. employs a simple coarse-to-fine registration strategy, which applies a smooth mask on the positional encoding serving as a dynamic low-pass filter. This strategy allows BARF to optimize the scene representation with camera poses. Following the idea of BARF, Wang et al. [40] propose bundle-adjusted deblur NeRF (BAD-NeRF), which optimizes the camera trajectories and NeRF network together by linearly interpolating camera poses to imitate the image formation model of motion blur, so that it can estimate the clear scenes from blurred images.

3. Method

Our method takes a single compressed image and encoding masks as input, and recovers the underlying 3D scene representation as well as camera poses. High frame-rate images can then be rendered from the learned 3D scene rep-

resentation. To achieve this, we exploit NeRF [26] as the underlying 3D scene representation due to its impressive representation capability. We follow the real image formation process to synthesize snapshot compressed image from NeRF. By maximizing the photometric consistency between the synthesized image and the actual measurements, we optimize both NeRF and the camera poses. Fig. 2 presents the overview of our method. We detail each component as follows.

3.1. Background on NeRF

Given a set of input multi-view images (together with both the camera intrinsic and extrinsic parameters), NeRF [26] transfers the pixels of the input images into rays. It then samples points along each ray, and takes 5D vectors (i.e., the 3D position of sampled point and the 2D viewing directions) as input. The volume density σ and view-dependent RGB color \mathbf{c} of each sampled 3D point are then estimated by a Multi-layer Perceptron (MLP). The reason that NeRF predicts color from both position and viewing direction is to better deal with the specular reflection of the scene. After obtaining the volume density and color of each sampled point along the ray, it employs a conventional volumetric rendering technique to integrate the density and color to synthesize the corresponding pixel intensity C of the image. The whole process can be formally defined via following equation:

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \quad (1)$$

where t_n and t_f are near and far bounds for volumetric rendering respectively, $\mathbf{r}(t)$ is the sampled 3D point along the ray \mathbf{r} at the distance t from the camera center, $\sigma(\mathbf{r}(t))$ represents the predicted density of the sampled point $\mathbf{r}(t)$ by the MLP, $T(t)$ denotes the accumulated transmittance along the ray from t_n to t , and is defined as $\exp(-\int_{t_n}^t T(t)\sigma(\mathbf{r}(s))ds)$, \mathbf{d} is the viewing direction in the world coordinate frame, and $\mathbf{c}(\mathbf{r}(t), \mathbf{d})$ is the predicted color of the sampled point $\mathbf{r}(t)$ by the MLP.

The photometric loss, i.e. the mean squared error (MSE) between the rendered pixel intensity and the real captured intensity, is usually used to train the networks:

$$\mathcal{L} = \sum_{\mathbf{r} \in \mathcal{R}} \left\| \hat{C}(\mathbf{r}) - C(\mathbf{r}) \right\|^2, \quad (2)$$

where both $C(\mathbf{r})$ and $\hat{C}(\mathbf{r})$ denote the rendered and real captured pixel intensities for ray \mathbf{r} respectively, and \mathcal{R} denotes the set of sampled rays.

3.2. Image Formation Model of Video SCI

The formation process of a video SCI system is similar to that of a blurry image. The difference is that the cap-

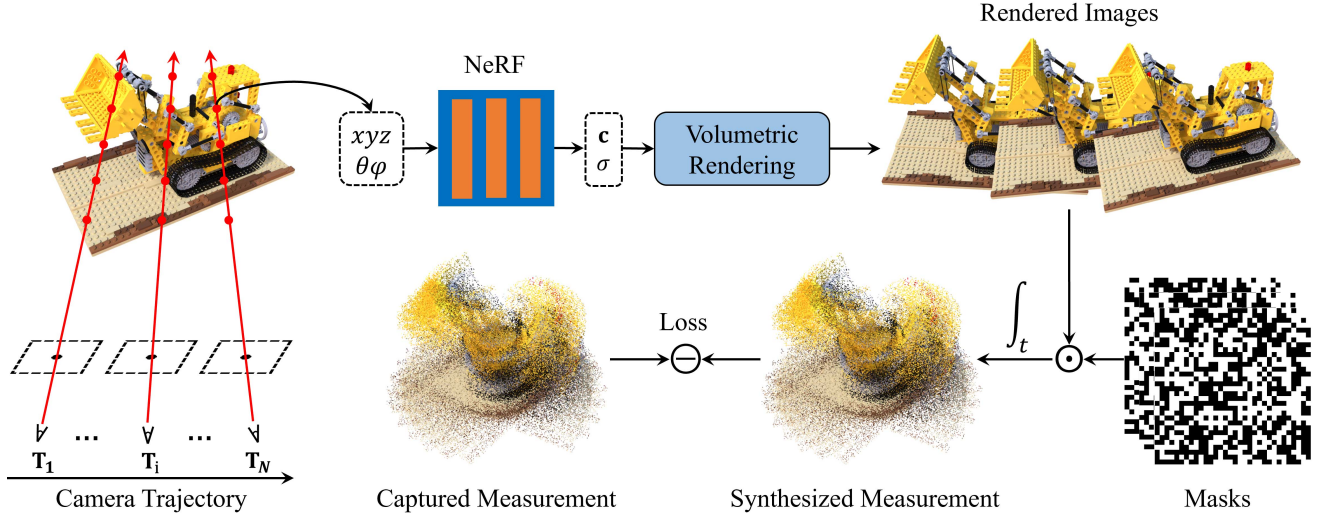


Figure 2. **Overview of the proposed SCINeRF.** Our method takes a single snapshot compressed image and corresponding masks as input, and recovers the underlying 3D scene representation as well as the camera motion trajectory within a single exposure time.

tured image of a video SCI system is modulated by N binary masks $\{\mathbf{M}_i\}_{i=1}^N \in \mathbb{R}^{H \times W}$ across the exposure time, where both H and W are image height and width respectively. For practical hardware implementation, those masks are achieved by displaying different 2D patterns on the Digital Micro-mirror Device (DMD) and a spatial light modulator, e.g., liquid crystal. The image sensor then accumulates the modulated photons across exposure time to a compressed/coded image. The number of masks or different patterns on the DMD within exposure time determine the number of coded frames, i.e., the temporal compression ratio (CR). Due to the mask modulation, the image restoration problem is not ill-posed anymore, i.e., the N virtual images can be recovered from a single compressed image alone. The whole imaging process can be described formally as:

$$\mathbf{Y} = \sum_{i=1}^N \mathbf{X}_i \odot \mathbf{M}_i + \mathbf{Z}, \quad (3)$$

where $\mathbf{Y}, \mathbf{X}_i \in \mathbb{R}^{H \times W}$ are the captured compressed image and the i^{th} virtual image within exposure time respectively, N is the temporal CR, \odot denotes element-wise multiplication, and $\mathbf{Z} \in \mathbb{R}^{H \times W}$ is the measurement noise. The individual pixel value in the binary mask is randomly generated. For N masks across the exposure time, the probability of assigning 1 to the same pixel location is fixed and defined as overlapping ratio. The optimal overlapping ratio is selected via ablation study.

Given the NeRF representation and corresponding camera poses, we are able to render \mathbf{X}_i to synthesize the compressed image \mathbf{Y} . We can see that \mathbf{Y} is differentiable with respect to NeRF and the poses, which lay the foundation for

our optimization step.

3.3. Proposed Framework

Based on Eq. (3), we can model the SCI measurement (i.e. \mathbf{Y}) as the compressed multi-view images. However, we cannot estimate the scene from SCI measurement directly because of lack of camera poses. Since the SCI measurement contains compressed frames and it is in the form of 2D image, it is unrealistic to use the conventional SfM method to estimate the camera pose for each frame, as most of the NeRF-based methods.

To deal with the unknown poses problem, we follow the idea of prior methods [19, 40, 41], which start from the initialized inaccurate poses and optimize the parameters of NeRF together with poses simultaneously. Since the compressed multi-view images are taken within a relatively-short exposure time, we assume that the camera trajectory during the imaging process is linear and obtain poses using linear interpolation. For more complex motions, we can exploit higher-order spline [17] or directly optimize individual poses without the loss of generality. The virtual camera pose can thus be represented as:

$$\mathbf{T}_i = \mathbf{T}_1 \exp\left(\frac{i}{N} \log(\mathbf{T}_1^{-1} \mathbf{T}_N)\right), \quad (4)$$

where \mathbf{T}_i is the pose of i^{th} frame, N denotes the number of compressed frames, $\mathbf{T}_1 \in \mathbf{SE}(3)$ and $\mathbf{T}_N \in \mathbf{SE}(3)$ are the poses of the first and last frame respectively. They represent start and end pose of the camera trajectory. Both \mathbf{T}_1 and \mathbf{T}_N are initialized to be nearly identity matrix with small random perturbations to the translation components. The intrinsic parameters, including the camera focal length

and principle points, come from camera intrinsic calibration, which is a well-studied problem with abundant methods available.

Given the image formation model and the camera pose representations, we are now able to jointly optimize both the NeRF parameters and camera poses by minimizing following loss function:

$$\mathcal{L} = \sum_{\mathbf{r} \in \mathcal{R}} \left\| \hat{\mathbf{Y}}(\mathbf{r}) - \sum_{i=1}^N \mathbf{M}(\mathbf{r}, i) \odot \mathbf{C}(\mathbf{r}, i) \right\|^2, \quad (5)$$

where \mathcal{R} denotes the set of sampled rays \mathbf{r} , $\hat{\mathbf{Y}}(\mathbf{r})$ is pixel value of the real captured image corresponding to \mathbf{r} , $\mathbf{M}(\mathbf{r}, i)$ is the binary mask value of the i^{th} mask corresponding to ray \mathbf{r} , \odot represents element-wise multiplication, and $\mathbf{C}(\mathbf{r}, i)$ denotes the rendered pixel value of the i^{th} frame corresponding to \mathbf{r} .

4. Experiments

We validate our SCINeRF on both synthetic and real datasets captured by our system, and it is evaluated against state-of-the-art (SOTA) SCI image restoration methods. The experimental results demonstrate that SCINeRF delivers higher performance compared to existing works in terms of restoration quality.

4.1. Experimental Setup

Synthetic datasets. We generate synthetic datasets from commonly-used multi-view datasets for image-based rendering and novel-view synthesis, including *Airplants* in LLFF [25] and *Hotdog* in NeRF Synthetic360 [26]. We generate additional synthetic datasets from the scenes provided in DeblurNeRF [22] with Blender, to better simulate the capturing process along a trajectory of real camera. There are four virtual scenes in the datasets synthesized by Blender, including *Cozy2room*, *Tanabata*, *Factory*, and *Vender*. For these four scenes, we generate images from 5 trajectories and compute the average image quality metrics. The compression ratio of benchmark datasets is 8. To verify the adaptability to different image resolutions of our SCINeRF, we use different spatial resolutions in synthetic datasets: For LLFF, we use 512×512 , for NeRF Synthetic360 we use 400×400 . For Blender-synthesized datasets, we use 600×400 resolution.

Real-world datasets. For real dataset, we collect the SCI measurement from a real video SCI setup. The setup consists of an iRAYPLE A5402MU90 camera and a FLDISCOVERY F4110 DMD. The compression ratio of the real dataset is 8, with the resolution of SCI measurement being 1024×768 . Fig. 3 shows the experimental setup we used to collect real dataset.




Figure 3. **Experimental setup for real dataset collection.** This SCI imaging system contains a CCD camera to record snapshot measurement, primary and rely lens, and a DMD to modulate input frames.

Baseline methods and evaluation metrics. Since our SCINeRF can render high-quality images from the estimated scene, we compare our method against SOTA SCI image restoration methods, including GAP-TV [49], PnP-FFDNet [51], PnP-FastDVDNet [52], and EfficientSCI [38]. For fair comparisons, we fine-tuned EfficientSCI [38] with the masks used in our datasets. Additionally, we also compare the performance on novel view image synthesis of SCINeRF against that of vanilla NeRF with restored images from prior SOTA methods as input. We noticed that due to the lack of high-quality details and multi-view consistency, we cannot estimate camera poses via SfM using the images reconstructed from those SCI image reconstruction methods. So when feeding the reconstructed images into NeRF, we use camera poses estimated from ground truth images instead, making the comparison unfair to our SCINeRF (but providing reasonable results). Widely used metrics are employed for quantitative evaluations, such as the structural similarity index (SSIM), peak signal-to-noise-ratio (PSNR), and learned perceptual image patch similarity (LPIPS) [54].

Implementation details. We implement our method with PyTorch [48]. We exploit the vanilla NeRF from [26] to represent the 3D scene. For higher performance, more recent representation can be exploited [27]. We use two separate Adam optimizers [16] with learning rate for scene model decreasing from 5×10^{-4} to 5×10^{-5} exponentially, and the learning rate for pose optimization decreasing from 1×10^{-3} to 1×10^{-5} exponentially. We train our model for 100-200K iterations, with 5000 rays as batch size.

Results The experimental results on the synthetic dataset provides empirical evidence on the efficacy of our SCINeRF in estimating and representing high-quality 3D scenes from a single SCI measurement, as shown in Fig. 4 and Table 1. It is noteworthy that, in certain scenes, the SSIM metric of our method does not exceed that of EfficientSCI. This observation can be attributed to the fact that our approach does not directly recover compressed images from the input

	Airplants			Hotdog			Cozy2room			Tanabata			Factory			Vender		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
GAP-TV [49]	22.85	.4057	.4986	22.35	.7663	.3179	21.77	.4321	.6031	20.42	.4264	.6250	24.05	.5666	.5149	20.00	.3678	.6882
PnP-FFDNet [51]	27.79	.9117	.1817	29.00	.9765	.0511	28.98	.8916	.0984	29.17	.9032	.1197	31.75	.8977	.1142	28.70	.9235	.1315
PnP-FastDVDNet [52]	28.18	.9092	.1757	29.93	.9728	.0522	30.19	.9132	.0793	29.73	.9333	.0980	32.53	.9165	.1055	29.68	.9395	.1043
EfficientSCI [38]	30.13	.9425	.1129	30.75	.9568	.0461	31.47	.9327	.0476	32.30	.9587	.0600	32.87	.9259	.0709	33.17	.9401	.0456
ours	30.69	.9335	.0728	31.35	.9878	.0310	33.23	.9492	.0445	33.61	.9638	.0374	36.60	.9638	.0221	36.40	.9840	.0298

Table 1. **Quantitative SCI image reconstruction comparisons on the synthetic datasets** The results are computed from the rendered images from estimated scenes via SCINeRF and recovered images from state-of-the-art SCI image restoration methods. The experimental results demonstrate that our SCINeRF can render images with higher quality than those from existing methods.

	Airplants			Hotdog			Cozy2room			Tanabata			Factory			Vender		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
NeRF+GAP-TV	23.72	.4684	.4195	23.80	.7672	.2847	21.99	.5027	.5212	20.91	.4122	.5531	25.48	.6853	.4401	21.68	.4365	.5536
NeRF+PnP-FFDNet	26.72	.8831	.2249	29.14	.9721	.0618	30.15	.9030	.0919	28.00	.9025	.1400	31.96	.8850	.1823	30.02	.9343	.1317
NeRF+PnP-FastDVDNet	26.91	.8766	.2061	29.31	.9729	.0548	31.17	.9170	.0797	30.79	.9362	.1000	32.56	.9026	.1532	31.30	.9451	.1151
NeRF+EfficientSCI	28.62	.9140	.1595	29.82	.9766	.0527	31.79	.9357	.0641	31.35	.9535	.0769	32.72	.8942	.0676	32.77	.9431	.0648
ours	30.61	.9384	.0764	30.59	.9814	.0442	33.12	.9477	.0357	32.84	.9574	.0297	36.33	.9602	.0426	34.75	.9744	.0257

Table 2. **Quantitative comparisons on novel-view synthesis** The results are the rendered novel-view images from SCINeRF and novel-view images from vanilla NeRF with reconstructed images from existing state-of-the-art methods. Since we cannot estimate accurate poses from the SOTA outputs, we use ground truth images to estimate camera poses instead. The experimental results show that our SCINeRF outperforms existing approaches.

SCI measurement; instead, it utilizes NeRF to render images based on the underlying scene representation. While NeRF successfully preserves the majority of scene details, the rendering process may introduce a marginal loss of image information. Nevertheless, our method exhibits superior performance in terms of both PSNR and LPIPS metrics. We notice that our SCINeRF performs significantly higher than existing methods in scenes with rich textures and characters, where existing methods fail to retrieve these details as shown in Fig. 4.

Table 2 presents the quantitative comparisons on novel-view image synthesis. Even in the unfair situation (i.e. the other methods exploit camera poses computed by COLMAP [32] with ground truth images), our method still outperforms existing methods.

To evaluate the performance of SCINeRF on real datasets, we also conduct qualitative comparisons against SOTA methods. Fig. 5 illustrates the experimental results, depicting the outcomes for real datasets. Notably, existing SOTA techniques exhibit limitations when confronted with high-frequency details and characters, resulting in noticeable deficiencies in the output images. In contrast, our SCINeRF surpasses these methods on real datasets by effectively recovering scenes with fine details, thereby achieving superior performance.

4.2. Additional Study

Mask overlapping rate. We evaluate the effect of different mask overlapping rates when modulating the input frames during SCI image formation process. When capturing the SCI measurement, the camera moves along a trajectory and a series of fixed masks modulate the input multi-view images to form the compressed image. It is possible that there is spatial overlap between masks. We define the mask over-

Mask Overlapping Rate	Synthetic Dataset		
	PSNR↑	SSIM↑	LPIPS↓
0.125	31.93	.9562	.0520
0.25	33.61	.9599	.0420
0.5	32.92	.9494	.0465
0.75	30.08	.8917	.1398

Table 3. **Additional studies on the effect of mask overlapping rate.** The results demonstrate that image quality first increase then decrease when the mask overlapping rate increase. The best mask overlapping rate is 0.25.

lapping rate as:

$$\mathbf{OR}(x, y) = \frac{\sum_{i=1}^N \mathbf{M}_i(x, y)}{N}, \quad (6)$$

where \mathbf{OR} denotes mask overlapping rate, \mathbf{M}_i is the i -th mask and (x, y) indicates pixel coordinate, and N is total number of compressed images. From Eq. (6), we know that when mask overlapping rate becomes low, the mask will become sparse in spatial domain, leading to lower sampling rate but higher information loss. However, when overlapping rate becomes high, the measurement values will contain pixels values from more frames. It may lead to pixel value ambiguous and deterioration on the estimated scenes from NeRF. We evaluate the effect of different mask overlapping rates via synthetic datasets. We tested different overlapping rates, ranging from 0.125 to 0.75. As shown in Table 3, when mask overlapping rates increase from 0.125 to 0.25, the image quality becomes higher, indicating the benefit to sample more pixels. However, the image quality decreases rapidly when the overlapping rate continues to increase. It demonstrates that sampling too many pixels would also lead to performance degradation due to the




Figure 4. **Qualitative evaluations of our method against SOTA SCI image restoration methods on the synthetic dataset.** Top to bottom shows the results for different scenes, including *Cozy2room*, *Tanabata*, *Factory* and *Vender*. The experimental results demonstrate that our method achieves superior performance on image restoration from a single compressed image (the far-left column).

increased ambiguity, which can be better perceived for an extreme case. For example, when the overlapping rate becomes 1, the formation model is the same as that of a blurry image and it is severe ill-posed to recover sharp images from a single blurry image. Empirically, we choose the overlapping rate not to exceed 0.25 for all our experiments.

High compression ratio We study the performance of our SCINeRF under different compression ratios, i.e., the number of compressed images N within one SCI measurement. High compression ratio is a challenging task for SCI image reconstruction algorithms because the higher compression

ratio leads to more information loss, making the image reconstruction and recovery more difficult. In this paper, we test our SCINeRF under the compression ratios of 8, 16, 24 and 32, by running SCINeRF on synthetic *Cozy2room*, *Tanabata*, *Factory*, *Vender* datasets at different compression ratios. Table 4 shows the mean metrics of our SCINeRF on four scenes under different compression ratios. It is noteworthy that the image quality of state-of-the-art methods decrease rapidly when compression ratios increase. On the other hand, the image quality deterioration of SCINeRF is small, and our SCINeRF still keeps high image quality even under very-high compression ratios.




Figure 5. **Qualitative evaluations of our method against SOTA SCI image restoration methods on the real dataset captured by our system in Fig. 3.** Top to bottom shows the results for different scenes. Since the compressed ground truth images in real datasets are unavailable, we capture separate scene images after capture the snapshot compressed image used for reference. For qualitative evaluation purpose, we render images from the learned 3D scene representations by SCINeRF. The results demonstrate that our SCINeRF surpasses existing image restoration methods on real datasets.

	CR=16		CR=24		CR=32	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
GAP-TV	19.50	.4975	18.63	.4700	18.25	.4425
PnP-FFDNet	24.35	.7950	22.21	.7325	23.55	.6975
PnP-FastDVDNet	26.59	.8600	25.85	.8475	24.56	.8100
EfficientSCI	31.60	.9093	29.10	.8717	28.43	.8443
Ours	34.88	.9600	34.37	.9532	33.95	.9481

Table 4. **Additional studies on the performance of methods with different compression ratios.** We test the performance of different methods under the compression ratios (CR) of 16, 24 and 32. The image quality metrics of our SCINeRF decrease slightly when the compression ratio increase.

5. Conclusion

In this paper, we present SCINeRF, a novel approach for 3D scene representation learning from a single snapshot compressed image. SCINeRF exploits neural radiance fields as its underlying scene representation due to its impressive representation capability. Physical image formation process of an SCI image is exploited to formulate

the training objective for jointly NeRF training and camera poses optimization. Different from previous works, our method considers the underlying 3D scene structure to ensure multi-view consistency among recovered images. To validate the effectiveness of SCINeRF, we conduct thorough evaluations against existing state-of-the-art techniques for SCI image recovery with both synthetic and real datasets. Extensive experimental results demonstrate the superior performance of our method in comparison to existing methods, and the necessity to consider the underlying 3D scene structure for SCI decoding.

Acknowledgements. This work was supported in part by NSFC under Grants 62202389 and 62271414, in part by a grant from the Westlake University-Muyuan Joint Research Institute, and in part by the Westlake Education Foundation, Science Fund for Distinguished Young Scholars of Zhejiang Province (LR23F010001), Research Center for Industries of the Future (RCIF) at Westlake University and the Key Project of Westlake Institute for Optoelectronics (Grant No. 2023GD007).

References

- [1] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignerf: Fully controllable neural 3d portraits. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 20364–20373, 2022. [3](#)
- [2] Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition. *arXiv preprint arXiv:2008.03824*, 2020. [3](#)
- [3] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerf: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. [3](#)
- [4] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011. [2](#)
- [5] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509, 2006. [1](#)
- [6] Ziheng Cheng, Bo Chen, Guanliang Liu, Hao Zhang, Ruiying Lu, Zhengjue Wang, and Xin Yuan. Memory-efficient network for large-scale video compressive sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16246–16255, 2021. [1](#), [2](#)
- [7] Ziheng Cheng, Ruiying Lu, Zhengjue Wang, Hao Zhang, Bo Chen, Ziyi Meng, and Xin Yuan. Birnat: Bidirectional recurrent neural networks with adversarial training for video snapshot compressive imaging. In *European Conference on Computer Vision*, pages 258–275. Springer, 2020. [1](#), [2](#)
- [8] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006. [1](#)
- [9] 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. [3](#)
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. [2](#)
- [11] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985*, 2021. [3](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#)
- [13] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18398–18408, June 2022. [3](#)
- [14] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5846–5854, 2021. [3](#)
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 694–711. Springer, 2016. [2](#)
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [17] Moyang Li, Peng Wang, Lingzhe Zhao, Bangyan Liao, and Peidong Liu. USB-NeRF: Unrolling shutter bundle adjusted neural radiance fields. In *arXiv preprint arXiv:2310.02687*, 2023. [4](#)
- [18] Xuejun Liao, Hui Li, and Lawrence Carin. Generalized alternating projection for weighted-2,1 minimization with applications to model-based compressive sensing. *SIAM Journal on Imaging Sciences*, 7(2):797–823, 2014. [1](#), [2](#)
- [19] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5741–5751, 2021. [3](#), [4](#)
- [20] Yang Liu, Xin Yuan, Jinli Suo, David J Brady, and Qionghai Dai. Rank minimization for snapshot compressive imaging. *IEEE transactions on pattern analysis and machine intelligence*, 41(12):2990–3006, 2018. [1](#), [2](#)
- [21] Jiawei Ma, Xiao-Yang Liu, Zheng Shou, and Xin Yuan. Deep tensor admm-net for snapshot compressive imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10223–10232, 2019. [1](#), [2](#)
- [22] Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12861–12870, 2022. [5](#)
- [23] Quan Meng, Anpei Chen, Haimin Luo, Minye Wu, Hao Su, Lan Xu, Xuming He, and Jingyi Yu. Gnerf: Gan-based neural radiance field without posed camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6351–6361, 2021. [3](#)
- [24] Xin Miao, Xin Yuan, Yunchen Pu, and Vassilis Athitsos. 1-net: Reconstruct hyperspectral images from a snapshot measurement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4059–4069, 2019. [2](#)
- [25] 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](#)
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view syn-

- thesis. *Communications of the ACM*, 65(1):99–106, 2021. **2, 3, 5**
- [27] 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, July 2022. **5**
- [28] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021. **3**
- [29] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. **3**
- [30] Mu Qiao, Ziyi Meng, Jiawei Ma, and Xin Yuan. Deep learning for video compressive sensing. *Appl Photonics*, 5(3), 2020. **1, 2**
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. **2**
- [32] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. **2, 3, 6**
- [33] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229–6238, 2021. **3**
- [34] Matthew Tancik, Vincent Casser, Xintan 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. **3**
- [35] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. **3**
- [36] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12922–12931, 2022. **3**
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **2**
- [38] Lishun Wang, Miao Cao, and Xin Yuan. Efficientsci: Densely connected network with space-time factorization for large-scale video snapshot compressive imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18477–18486, 2023. **1, 2, 5, 6, 7, 8**
- [39] Lishun Wang, Miao Cao, Yong Zhong, and Xin Yuan. Spatial-temporal transformer for video snapshot compressive imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. **1, 2**
- [40] Peng Wang, Lingzhe Zhao, Ruijie Ma, and Peidong Liu. Bad-nerf: Bundle adjusted deblur neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4170–4179, June 2023. **3, 4**
- [41] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. **4**
- [42] Zhengjue Wang, Hao Zhang, Ziheng Cheng, Bo Chen, and Xin Yuan. Metasci: Scalable and adaptive reconstruction for video compressive sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2083–2092, 2021. **2, 3**
- [43] Silvan Weder, Guillermo Garcia-Hernando, Aron Monszpart, Marc Pollefeys, Gabriel J Brostow, Michael Firman, and Sara Vicente. Removing objects from neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16528–16538, 2023. **3**
- [44] Fangyin Wei, Thomas Funkhouser, and Szymon Rusinkiewicz. Clutter detection and removal in 3d scenes with view-consistent inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18131–18141, 2023. **3**
- [45] Yuanbo Xiangli, Lining 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 *European conference on computer vision*, pages 106–122. Springer, 2022. **3**
- [46] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13779–13788, 2021. **3**
- [47] Peihao Yang, Linghe Kong, Xiao-Yang Liu, Xin Yuan, and Guihai Chen. Shearlet enhanced snapshot compressive imaging. *IEEE Transactions on Image Processing*, 29:6466–6481, 2020. **2**
- [48] Lin Yen-Chen. Nerf-pytorch. <https://github.com/yenchenlin/nerf-pytorch/>, 2020. **5**
- [49] Xin Yuan. Generalized alternating projection based total variation minimization for compressive sensing. In *2016 IEEE International conference on image processing (ICIP)*, pages 2539–2543. IEEE, 2016. **1, 2, 5, 6, 7, 8**
- [50] Xin Yuan, David J Brady, and Aggelos K Katsaggelos. Snapshot compressive imaging: Theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 38(2):65–88, 2021. **1**

- [51] Xin Yuan, Yang Liu, Jinli Suo, and Qionghai Dai. Plug-and-play algorithms for large-scale snapshot compressive imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1447–1457, 2020. [2](#), [5](#), [6](#), [7](#), [8](#)
- [52] Xin Yuan, Yang Liu, Jinli Suo, Fredo Durand, and Qionghai Dai. Plug-and-play algorithms for video snapshot compressive imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):7093–7111, 2021. [2](#), [5](#), [6](#)
- [53] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. [3](#)
- [54] 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](#)