

# Transientangelo: Few-Viewpoint Surface Reconstruction Using Single-Photon Lidar

Weihan Luo<sup>1</sup>  
 weihan.luo@mail.utoronto.ca

Anagh Malik<sup>1,2</sup>  
 anagh@cs.toronto.edu

David B. Lindell<sup>1,2</sup>  
 lindell@cs.toronto.edu

<sup>1</sup>University of Toronto    <sup>2</sup>Vector Institute

<https://weihan1.github.io/transientangelo/>

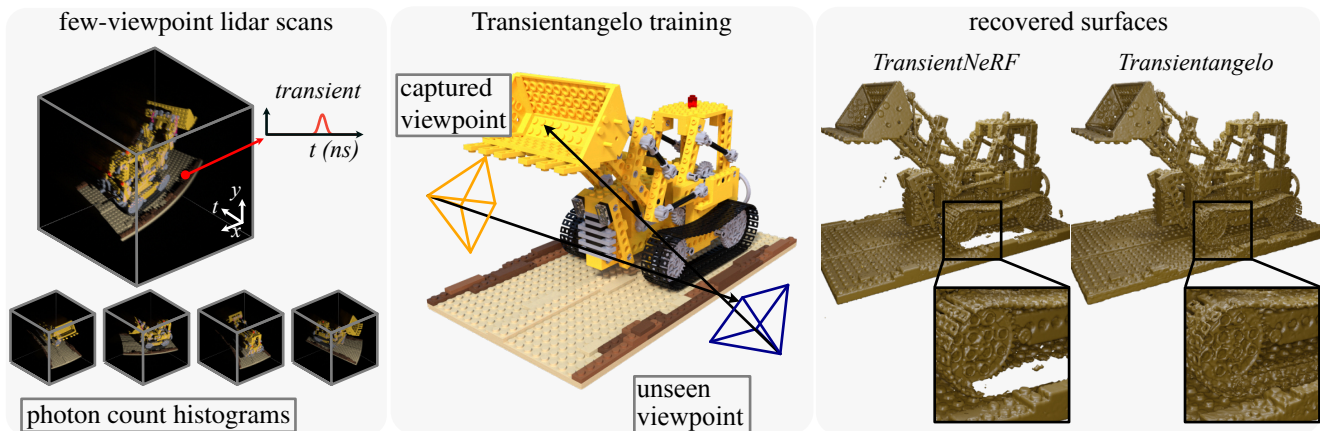


Figure 1. Transientangelo takes as input raw lidar scans from sparse viewpoints. These scans are used to optimize a scene representation based on a signed distance function, which is further regularized to constrain the geometry from both captured and unseen viewpoints. The method recovers higher-fidelity surfaces than previous methods in the sparse-view and low-photon regime (i.e., from tens to hundreds of measured photons per pixel). The above scene was trained using simulated single-photon lidar data from five viewpoints, with an average of 150 photons per pixel over the occupied regions of the scene.

## Abstract

We consider the problem of few-viewpoint 3D surface reconstruction using raw measurements from a lidar system. Lidar captures 3D scene geometry by emitting pulses of light to a target and recording the speed-of-light time delay of the reflected light. However, conventional lidar systems do not output the raw, captured waveforms of backscattered light; instead, they preprocess these data into a 3D point cloud. Since this procedure typically does not accurately model the noise statistics of the system, exploit spatial priors, or incorporate information about downstream tasks, it ultimately discards useful information that is encoded in raw measurements of backscattered light. Here, we propose to leverage raw measurements captured with a single-photon lidar system from multiple viewpoints to optimize a neural surface representation of a scene. The measurements consist of time-resolved photon count histograms, or transients, which capture information about backscattered light at picosecond time scales. Additionally, we develop new

regularization strategies that improve robustness to photon noise, enabling accurate surface reconstruction with as few as 10 photons per pixel. Our method outperforms other techniques for few-viewpoint 3D reconstruction based on depth maps, point clouds, or conventional lidar as demonstrated in simulation and with captured data.

## 1. Introduction

Within the last few years, significant progress has been made on the problem of multi-view 3D surface reconstruction. Recent methods combine efficient neural surface representations with volumetric rendering models, enabling high-fidelity scene reconstruction from multi-view images [23, 47]. Overall, achieving accurate scene reconstruction is now relatively straightforward when tens to hundreds of input images are available. However, a significant remaining challenge is achieving robust reconstruction when the number of input viewpoints is reduced to only a few (e.g., less than ten). In

this case, scene reconstruction from conventional images alone is ill-posed [18] and requires strong priors to solve robustly [54].

Instead of using conventional images, we investigate this challenge using lidar measurements captured from a few (2–5) different viewpoints. Lidar measures the distance to each scene point by illuminating it with a pulse of light and measuring the speed-of-light time delay of the back-reflected pulse. These measurements provide additional constraints to the multi-view reconstruction problem. However, conventional lidar systems do not output the raw, captured waveforms of backscattered light; instead, they post-process these data into a 3D point cloud. Since this procedure typically does not accurately model the noise statistics of the system, exploit spatial priors, or incorporate information about downstream tasks, it ultimately discards useful information that is encoded in raw measurements of backscattered light. We are interested in using the raw measurements captured by a single-photon lidar system—the time-resolved photon counts histograms, or *transients*, which contain information about the precise arrival times of backscattered light particles at picosecond time scales.

Previous methods for multi-viewpoint 3D reconstruction—whether based on conventional images or lidar—have several drawbacks. First, image-based methods are usually trained with tens or hundreds of images, making acquisition difficult [23, 31, 34, 50]. While several recent works attempt to tackle the few-view reconstruction problem [28, 35, 49, 62], they typically fail for the case where the baseline between input images becomes large, and the problem becomes increasingly ill-posed. Recent methods that leverage large-scale diffusion models [8, 54] alleviate this limitation to some extent, but are also prone to hallucinating scene geometry. Other techniques improve few-view 3D reconstruction by introducing additional geometric constraints from point clouds [44] or monocular depth estimators [62]. In this vein, our work is closest to TransientNeRF [28], which shows that using raw lidar data instead of depth maps or point clouds leads to significant improvements in reconstructed scene appearance in the few-view setting. However, TransientNeRF is based on a volumetric representation and is designed for appearance modeling rather than surface reconstruction. Further, TransientNeRF uses long acquisition times ( $\approx 20$  minutes per view) to capture per-pixel photon count histograms with several thousand photon arrivals each.

To address these limitations we propose Transientangelo, a method for few-viewpoint scene reconstruction using transient supervision and a scene representation based on signed distance functions (SDFs). Inspired by Neuralangelo [23], we use an efficient neural representation with a hashing-based feature grid to parameterize an SDF and radiance field. The representation is then trained using supervision from photon count histograms. To further constrain the recon-

structed surfaces, we regularize the scene geometry from unseen viewpoints and use supervision from the captured reflectivity (i.e., the time-integrated lidar measurements).

Overall, we improve few-viewpoint 3D surface reconstruction from raw lidar measurements compared to previous work (see Figure 1), while also demonstrating significant improvements in reconstructed geometry compared to image-based or depth-based approaches. Our approach also improves robustness to lidar measurements captured with relatively few photon counts (e.g., 10–300), which is especially important when operating at fast acquisition speeds or when imaging targets at long range. In summary, we make the following technical contributions.

- We introduce Transientangelo: a method to perform surface reconstruction using few-viewpoint single-photon lidar.
- We propose regularization techniques that improve the reconstructed geometry and provide improved robustness in the low-flux regime where only tens to hundreds of photons are collected per pixel.
- We introduce a multi-viewpoint simulated and captured transient dataset with a range of photon count levels, which we use to demonstrate state-of-the-art performance for few-viewpoint surface reconstruction.

## 2. Related Work

Transientangelo brings together the areas of active single-photon imaging, neural surface reconstruction, and few-viewpoint 3D reconstruction.

**Active single-photon imaging.** We are specifically concerned with lidar systems based on single-photon avalanche diodes (SPADs), which output precise timestamps corresponding to the arrival times of individual captured photons. For SPAD-based lidar, a transient is measured by repeatedly illuminating a point with pulses of light and accumulating the individual photon arrival times into a time-resolved histogram. Such captured histograms of a scene can be used for direct scene reconstruction [13, 17, 24, 28], non-line of sight scene reconstruction [3, 5, 38, 43, 55], seeing through participating media [25], visualizations of light propagation [9, 27, 48] or to uncover statistical properties of captured photons [41, 42]. Our work differs from these in that we specifically aim to recover high-fidelity surfaces from multi-viewpoint SPAD data. We are motivated by previous work [28] that demonstrates the strong geometric information encoded in transients and their promise for multi-viewpoint surface reconstruction.

While many previous techniques on 3D imaging with SPADs use high-powered lasers or long acquisition times to capture hundreds or thousands of photons [24, 25, 28, 38],

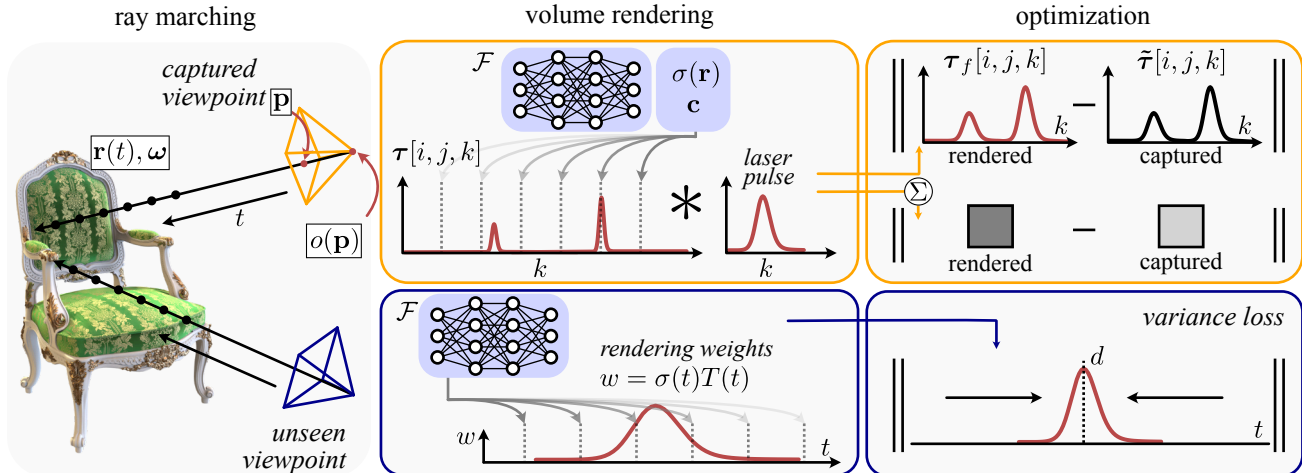


Figure 2. Transientangelo training procedure. For a pixel  $\mathbf{p}$ , we cast out a ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{c}\omega(\mathbf{p})$ . **Captured viewpoint:** 3D ray coordinates are passed through the neural surface representation  $\mathcal{F}$  to retrieve a radiance and an SDF value, which gets converted to density  $\sigma$  (see Equation 3). These values are then binned into a transient, which, after a convolution with the laser pulse, gives the final rendered transient  $\tau_f[i, j, k]$ . The network is supervised with an L1 loss between the rendered and captured transient and an L1 loss between the integrals of the rendered and captured transients (reflectivity loss). **Unseen viewpoint:** 3D ray coordinates are passed through the neural representation to retrieve rendering weights, which are used to calculate the variance of the weights around the depth  $d$ . The network is trained to minimize this variance, resulting in thinner surfaces and removing spurious zero-level sets.

our work is more closely related to single-photon imaging techniques that operate in the photon-starved regime. Such methods exploit spatial or temporal correlations in the photon arrivals to recover depth or reflectivity from only a few captured photons at each pixel [20,41]. Instead, our approach exploits cross-viewpoint information and reflectivity-based regularization for scene reconstruction from as few as 10 collected photons per pixel.

**Neural surface reconstruction.** Neural surface representations typically use either occupancy [36,37,39] or a signed distance function (SDF) [23,50,51,59,60] to model scene geometry. The former uses a neural network to classify points in 3D space that fall inside an object [37]. The latter learns a network that predicts the SDF of a 3D point (i.e., its signed distance to the closest surface). The surface is extracted from the SDF by finding the zero-level set. Our work, inspired by Neuralangelo [23], builds on a hashgrid-based SDF representation [34]. However, unlike in previous work, we use transient measurements as supervision.

Recent works [7,12,32,40,57] have also explored training neural surface representations using data from emerging sensors. Concurrent work [33] reconstructed 3D shapes of objects using transient measurements from a low-cost SPAD sensor. Their setup differs from ours as they utilize hundreds of training poses, while we investigate 3D surface reconstruction from 2 to 5 viewpoints. To the best of our knowledge, our work is the first to tackle the problem of few-viewpoint surface reconstruction using transient measurements.

**Few-viewpoint neural reconstruction.** Prior work has addressed few-viewpoint reconstruction using geometric supervision (e.g., based on depth maps, point clouds, or transients) [4,28,45,49,62], or by introducing other strategies to regularize reconstruction from multi-view images [15,19,35,53,56,58,61,62]. Our work is closest to TransientNeRF [28] and RegNeRF [35]. In contrast to TransientNeRF—which reconstructs a volumetric scene model from few-viewpoint transients—we focus on surface reconstruction using an SDF-based representation and improve robustness to noisy (low photon) measurements. Similar to RegNeRF, we regularize views rendered from novel viewpoints; however, we use a different depth-based regularization strategy and exploit transient data to achieve improved performance in the few-view setting.

### 3. Method

In this section, we provide a mathematical description of single-photon lidar [28], which we then relate to the rendering model used to optimize the surface representation. An overview of the method is shown in Figure 2.

#### 3.1. Measurement Model

Suppose that light from a laser pulse is reflected from a scene and focused onto a sensor at point  $\mathbf{p} \in \mathbb{R}^2$ . Also, assume that laser and sensor are co-axial, such that light propagates to the scene and back along the same path. Then, we can model the forward path of light along a ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{c}\omega(\mathbf{p})$ , where  $\mathbf{o} \in \mathbb{R}^3$  represents the camera origin,  $\mathbf{c}$

is the speed of light, and  $\omega$  is a unit vector describing the direction of the ray.

Light incident along  $\mathbf{r}$  is integrated at a sensor pixel  $(i, j)$  over time bins  $n$ , resulting in a transient measurement  $\tilde{\tau}[i, j, n]$ . The number of photons collected in each time bin follows a Poisson distribution [28, 41]:

$$\tilde{\tau}[i, j, n] \sim \text{Poisson}(N\eta\lambda[i, j, n] + B),$$

where  $B = N(\eta A[i, j] + D)$ . (1)

Here,  $\lambda$  is the expected number of incident photons and depends on the reflectivity of the point imaged by pixel  $(i, j)$ . The value  $N$  denotes the number of emitted laser pulses per pixel,  $\eta$  is the detection efficiency of the sensor, and  $B$  is the number of background detections, which depends on the average ambient photon rate at pixel coordinates  $(i, j)$ ,  $A[i, j]$ . Finally, the value  $D$  is the dark count—i.e., the number of false detections produced by the sensor.

For a Poisson distribution, the signal-to-noise ratio (SNR) is defined as the ratio between the mean and its standard deviation. Hence, in our case, the SNR of  $\tilde{\tau}$  is given by:

$$\frac{N\eta\lambda[i, j, n] + B}{\sqrt{N\eta\lambda[i, j, n] + B}}. \quad (2)$$

Decreasing the number of laser pulses reduces the acquisition time, the number of collected photons, and the SNR of the measurements.

### 3.2. Surface-based Transient Rendering

We render transients by combining a neural representation that parameterizes the scene appearance and surface geometry with a time-resolved volume rendering equation.

**Neural surface representation.** We use a neural surface representation  $\mathcal{F}$  to represent the scene appearance and surface geometry. The representation consists of a multi-resolution hash-grid, a base multi-layer perceptron (MLP)  $\mathcal{F}_b$ , and a colour MLP  $\mathcal{F}_c$  [23, 34]. The MLP  $\mathcal{F}_b$  maps a point along a ray  $\mathbf{r}(t) \in \mathbb{R}^3$  to a geometric latent feature vector  $\mathbf{l}_{\text{geo}} \in \mathbb{R}^{16}$  and an SDF value  $f \in \mathbb{R}$ . The surface normal is given as  $\mathbf{n} = \frac{\nabla f(\mathbf{r})}{\|\nabla f(\mathbf{r})\|}$ , where the gradient  $\nabla f(\mathbf{r})$  is computed via finite differences. We then recover the radiance at  $\mathbf{r}(t)$  as  $\mathbf{c} = \mathcal{F}_c(\mathbf{n}, \gamma(\omega), \mathbf{l}_{\text{geo}})$ , where  $\gamma$  is a positional encoding function [46].

To make the representation compatible with time-resolved volume rendering, we follow NeuS [50] and convert the predicted SDF to a density field  $\sigma(\mathbf{r})$  as

$$\sigma(\mathbf{r}) = \max\left(\frac{-\frac{d}{dt}\Phi_s(f(\mathbf{r}))}{\Phi_s(f(\mathbf{r}))}, 0\right), \quad (3)$$

where  $\Phi_s(x) = (1 + e^{-sx})^{-1}$  is the sigmoid function, and  $s$  is a learnable parameter controlling the support of the density field around the zero level set [52].

**Time-resolved volume rendering.** Following Transient-NeRF [28] we use the predicted density and radiance values from this representation to render transients  $\tau$  at each pixel  $(i, j)$  and time bin  $n$  using the time-resolved volume rendering equation:

$$\tau[i, j, n] = \int_{\mathcal{P}_{i,j}} \int_{\mathcal{T}_n} (tc)^{-2} T(t)^2 \sigma(\mathbf{r}) \mathbf{c}(\mathbf{r}, \omega) dt d\mathbf{p}, \quad (4)$$

where  $T(t) = \exp\left(-\int_{t_0}^t \sigma(\mathbf{r}) ds\right)$ .

The value  $T(t)$  is the transmittance from a distance  $t_0$  along the ray to  $t$ , and this term is squared to reflect the two-way propagation of light [1]. The value  $(tc)^{-2}$  accounts for the inverse-square falloff of intensity,  $\mathcal{T}_n$  is the transient bin width, and  $\mathcal{P}_{i,j}$  models the pixel footprint. Note that in practice, we approximate the integral using quadrature by sampling at locations  $t_i$  along the ray [29]. We approximate the integral over  $\mathcal{P}_{i,j}$  by explicitly sampling multiple rays within the area of each pixel. Finally, we convolve  $\tau$  with the calibrated temporal impulse response of the lidar system to yield the final rendered transient  $\tau_f$  [28].

### 3.3. Optimization

**Data term.** We supervise the rendered transients to be consistent with the measurements  $\tilde{\tau}$  by minimizing the following loss function.

$$\mathcal{L}_\tau = \frac{1}{|\mathcal{R}_{\text{train}}|} \sum_{\mathcal{R}_{\text{train}}} \|\tilde{\tau} - \tau_f\|_1, \quad (5)$$

where  $\mathcal{R}_{\text{train}}$  is the set of rays corresponding to the captured transient measurements used for training.

**Weight variance regularization.** To recover thinner and smoother surfaces and aid the model in generalizing to viewpoints far from the training distribution, we regularize the scene density from sampled unseen viewpoints by penalizing the variance of the rendering weights along each ray.

To sample unseen viewpoints, we follow RegNeRF [35] and sample camera origins on a sphere around the scene estimated from the training views (see supplement for details). We then regularize the density from these unseen views using RawNeRF’s weight-variance regularizer [30]. We find that replacing their rendered depth estimate with the argmax depth [28] recovers more accurate surfaces, as shown in the supplement.

More specifically, let  $t_i$  be the samples along the ray used in the quadrature approximation of Eq. 4. Then, the values  $w_i = \sigma(t_i)T(t_i)$  represent alpha compositing weights applied to the radiance values along a ray. We use  $w_i$  to compute the weight-variance regularizer, which we apply to rays  $\mathcal{R}'_{\text{train}}$  sampled from the set of unseen views:



$$\mathcal{L}_{\text{weight\_var}} = \frac{1}{|\mathcal{R}'_{\text{train}}|} \sum_{\mathcal{R}'_{\text{train}}} \sum_i w_i \frac{(t_i - d)^3 + (t_{i-1} - d)^3}{3(t_i - t_{i-1})}. \quad (6)$$

The value  $d = \arg \max_{t^*} T(t^*)\sigma(t^*)$  represents the depth, or the distance along the ray with the maximum probability of ray termination [28].

**Reflectivity loss.** To improve performance in the low photon count regime, we introduce a reflectivity loss, which provides additional supervision using the integrated transients. Since the integrated transients have a higher signal-to-noise ratio than the individual photon count histogram bins, this provides a complementary source of supervision. We implement this loss as

$$\mathcal{L}_{\text{ref}} = \frac{1}{|\mathcal{R}_{\text{train}}|} \sum_{\mathcal{R}_{\text{train}}} \left\| \sum_n \tau[i, j, n] - \sum_n \tilde{\tau}[i, j, n] \right\|_1, \quad (7)$$

where the integrated transients are computed by summing over  $n$ .

In summary, our training loss can be expressed as

$$\mathcal{L} = \mathcal{L}_{\tau} + \lambda_{\text{ref}} \mathcal{L}_{\text{ref}} + \lambda_{\text{sc}} \mathcal{L}_{\text{sc}} + \lambda_{\text{eik}} \mathcal{L}_{\text{eik}} + \lambda_{\text{weight\_var}} \mathcal{L}_{\text{weight\_var}} + \lambda_{\text{sparse}} \mathcal{L}_{\text{sparse}}. \quad (8)$$

Here, the space carving loss,  $\mathcal{L}_{\text{sc}}$ , penalizes densities at locations where the transient is below a background level  $B$ :

$$\mathcal{L}_{\text{sc}} = \sum_{\substack{i,j,n \\ \tilde{\tau}[i,j,n] < B}} \int_{\mathcal{P}_{i,j}} \int_{\mathcal{T}_n} T(t)\sigma(\mathbf{r}) dt d\mathbf{p}. \quad (9)$$

The Eikonal regularization [10] is given as

$$\mathcal{L}_{\text{eik}} = \frac{1}{|\mathbf{r} \in \mathcal{R}_{\text{train}}|} \sum_{\mathbf{r} \in \mathcal{R}_{\text{train}}} (\|\nabla f(\mathbf{r})\| - 1)^2, \quad (10)$$

and encourages the SDF to have a gradient magnitude equal to one. Finally, the sparsity loss is

$$\mathcal{L}_{\text{sparse}} = \frac{1}{|\mathbf{r} \in \mathcal{R}_{\text{train}}|} \sum_{\mathbf{r} \in \mathcal{R}_{\text{train}}} \exp(-\alpha|f(\mathbf{r})|), \quad (11)$$

where  $\alpha$  is a scalar hyperparameter. This loss helps to prevent floating artifacts in the reconstruction.

### 3.4. Implementation Details

We implement our method in Pytorch-Lightning [6, 11]. We use the NerfAcc [22] version of InstantNGP [34] for efficient ray marching, and we build on top of TransientNeRF’s [28] time-resolved rendering.

**Training and optimization.** We initialize the SDF to approximate a sphere [2]. We follow Neuralangelo’s coarse-to-fine optimization strategy by initializing the step-size in the gradient calculation  $\epsilon$  as the coarsest hash-grid voxel size and exponentially decreasing it during training. For the hash-grid, we start at four hash resolutions and unmask two new hash resolutions every 5000 steps. We optimize the network using the AdamW [26] optimizer, with  $\beta_1 = 0.9$  and  $\beta_2 = 0.99$  and employ a two-step learning rate scheduler by initializing the learning rate to  $1 \times 10^{-5}$  and linearly increasing it until it reaches  $1 \times 10^{-3}$  by the end of the 5000 steps. Then, the learning rate is exponentially decreased over the remaining steps, with a total of 250K steps. All experiments were run on a single Nvidia RTX A6000 GPU with 48GB of memory. We provide additional implementation details, including hyperparameter settings, in the supplement.

## 4. Multiview Transient Dataset

To test the robustness of different methods on low-SNR transient data, we adapt the simulated and captured datasets introduced by TransientNeRF [28]. For each captured and simulated scene, we synthesize transients with an average of 300, 150, 50, and 10 photons per occupied (non-background) pixel. The paragraphs below outline how these datasets are created.

**Simulated dataset.** To synthesize transients at different photon counts in simulation, we match the flux from the time-resolved Mitsuba renderer [28] to the target number of photons per pixel. We then follow Malik et al. [28] and set the background flux to 0.001 photons per 2850 captured scene photons, which is set to match the statistics of the captured dataset. The flux is then Poisson sampled to simulate photon counts.

**Captured dataset.** We create low-flux captured datasets using the raw photon arrival timestamps from TransientNeRF. Specifically, we subsample the single photon arrivals using thinning [21] to achieve a set number of detected photons. We then bin the photon arrivals into a transient histogram. Additional dataset details can be found in the supplement. Both the simulated and captured variable photon count datasets will be released upon publication.

## 5. Results

We evaluate our method for transient and intensity image synthesis from novel viewpoints, depth rendering, and mesh recovery. We also benchmark the performance of our method’s robustness to low-SNR data. Our supplementary contains a more detailed set of results and ablation studies.

**Baselines.** We compare the rendered intensity, depth maps, and meshes against state-of-the-art methods in surface re-

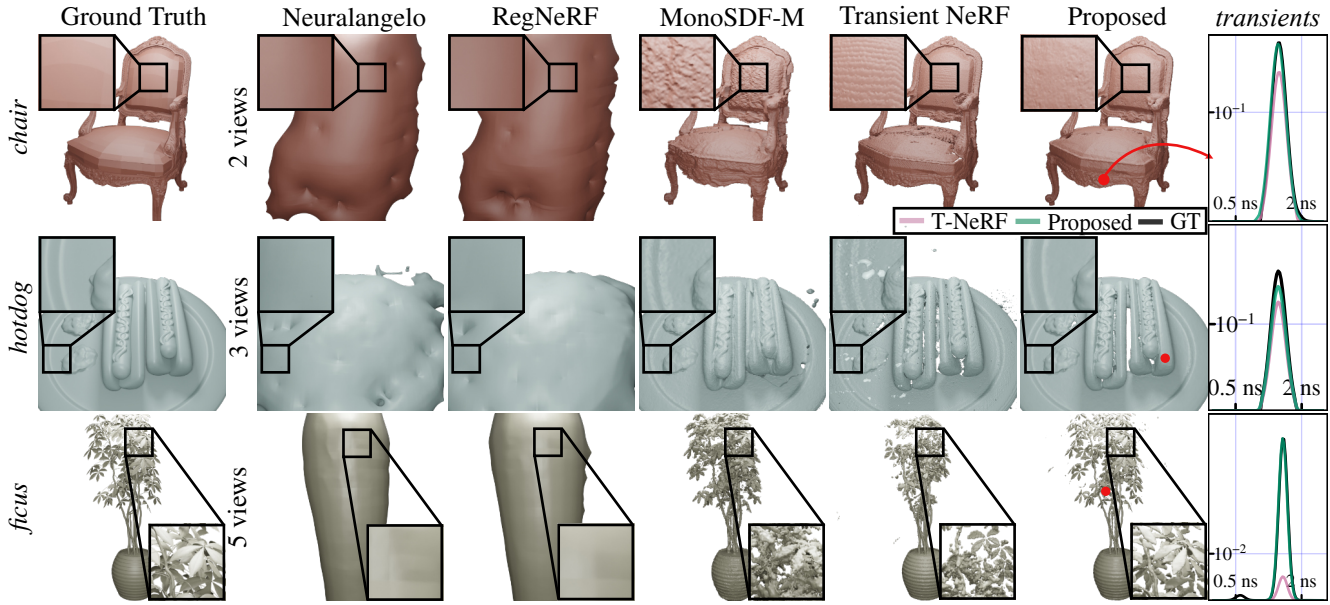


Figure 3. Results on the simulated dataset using an average of 6000 photons per occupied (non-background) pixel. We show the recovered meshes from the baselines and the proposed method. For the transient-based methods, we also show rendered transients for the indicated pixels. Our method recovers smoother meshes with fewer missing parts. We also recover transients that better match the ground truth.

Table 1. Simulated results assessing image quality, depth accuracy, and mesh quality. The left column indicates average photons per pixel.

Method	Chamfer Distance ↓			PSNR (dB) ↑			LPIPS ↓			L1 (depth) ↓			Transient IOU ↑			
	2 views	3 views	5 views	2 views	3 views	5 views	2 views	3 views	5 views	2 views	3 views	5 views	2 views	3 views	5 views	
6000	Neuralangelo [23]	4.56	7.99	7.17	18.57	19.44	22.03	0.388	0.442	0.434	0.162	0.122	0.090	-	-	-
	RegNeRF [35]	4.94	3.48	4.38	18.32	19.89	21.67	0.376	0.412	0.418	0.264	0.142	0.100	-	-	-
	MonoSDF w/ mask [62]	3.78	1.67	0.18	19.76	20.74	24.65	0.240	0.244	0.132	0.045	0.026	0.015	-	-	-
	TransientNeRF [28]	0.27	0.44	0.31	21.38	23.48	28.39	0.172	0.151	0.115	0.015	0.011	0.013	0.31	0.40	0.55
	Proposed	0.07	0.06	0.05	25.28	26.69	29.49	0.176	0.155	0.150	0.006	0.006	0.005	0.54	0.58	0.68
300	TransientNeRF	0.26	0.33	0.30	21.34	22.79	25.88	0.201	0.152	0.121	0.012	0.012	0.015	0.32	0.41	0.58
	Proposed	0.07	0.06	0.06	25.70	26.65	28.50	0.196	0.162	0.149	0.007	0.007	0.005	0.56	0.60	0.70
150	TransientNeRF	0.25	0.65	0.27	21.04	22.18	25.35	0.203	0.161	0.117	0.015	0.013	0.010	0.31	0.39	0.56
	Proposed	0.09	0.07	0.29	25.56	26.52	28.19	0.193	0.165	0.151	0.007	0.006	0.005	0.56	0.60	0.70
50	TransientNeRF	0.53	0.66	0.50	20.16	21.57	24.01	0.232	0.184	0.133	0.021	0.017	0.014	0.29	0.40	0.55
	Proposed	0.15	0.08	0.06	24.73	25.76	26.55	0.218	0.163	0.167	0.012	0.005	0.007	0.53	0.59	0.63
10	TransientNeRF	2.86	0.31	5.50	18.94	22.85	16.73	0.265	0.218	0.209	0.052	0.018	0.092	0.27	0.43	0.19
	Proposed	0.22	0.28	0.18	22.76	23.98	25.51	0.254	0.216	0.208	0.016	0.015	0.014	0.43	0.47	0.54

construction and few-view surface reconstruction: Neuralangelo [23], RegNeRF [35], and MonoSDF [62]. To ensure a fair comparison, we implement these baselines using the same Neuralangelo [23] backbone architecture, and we integrate their proposed loss functions. All three methods are supervised using integrated transient measurements. MonoSDF is also supervised using depth maps obtained by applying a log-matched filter [16] to the transient measurements, and the input normal maps are estimated from the depth. We also compare our method to TransientNeRF [28].

**Evaluation criteria.** To measure the mesh reconstruction quality, we use Chamfer Distance. For image reconstruction evaluation, we use PSNR and LPIPS [63]. We also evaluate the quality of the rendered transients using Tran-

sient IOU [28], defined as the intersection over the union of the rendered transient and the ground truth transient (see supplement for details).

## 5.1. Simulated Results

Figure 3 shows the recovered meshes and rendered transients for specific pixels from the proposed method compared to baselines. We show results for measurements with an average of 6000 photons per occupied pixel (results for other photon levels are included in the supplement). Due to a lack of any depth-based priors, Neuralangelo [23] and RegNeRF [35] both fail to recover accurate scene geometry. MonoSDF [62] recovers noisier surfaces than the proposed approach, especially when trained with fewer input viewpoints. Compared to TransientNeRF [28], the proposed

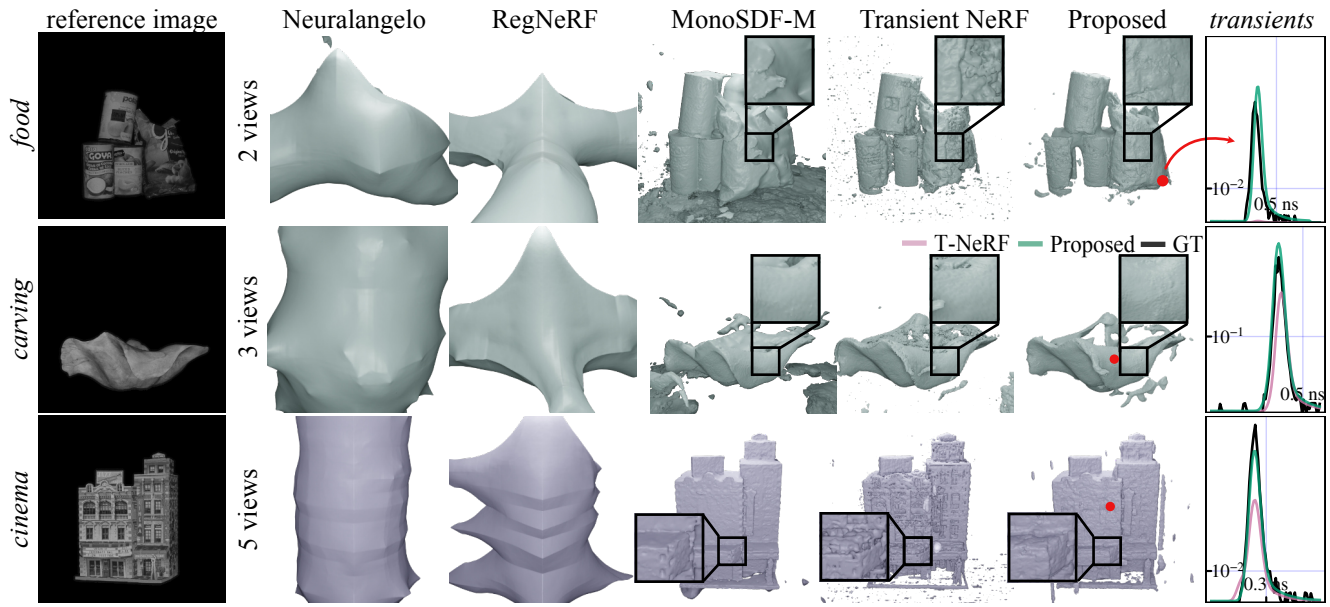


Figure 4. Results on the captured dataset. We show the recovered meshes from the baselines and the proposed method. Due to the lack of ground-truth, we include the closest captured image for reference of the scene. As can be seen, Neuralangelo recovers smoother meshes with fewer missing parts. We also recover transients that better match the ground truth.

Table 2. Captured results assessing image quality, depth accuracy, and mesh quality. The left column indicates average photons per pixel.

Method	PSNR (dB) $\uparrow$			LPIPS $\downarrow$			L1 (depth) $\downarrow$			Transient IOU $\uparrow$			
	2 views	3 views	5 views	2 views	3 views	5 views	2 views	3 views	5 views	2 views	3 views	5 views	
1500	Neuralangelo [23]	17.79	19.28	21.07	0.380	0.355	0.300	0.076	0.071	0.061	-	-	-
	RegNeRF [35]	18.50	19.59	20.29	0.375	0.320	0.330	0.085	0.073	0.076	-	-	-
	MonoSDF w/ mask [44]	17.63	21.16	27.25	0.353	0.248	0.157	0.033	0.026	0.021	-	-	-
	TransientNeRF [28]	22.11	21.83	22.72	0.271	0.212	0.172	0.005	0.006	0.010	0.34	0.41	0.50
	Proposed	21.31	23.95	25.22	0.241	0.170	0.145	0.007	0.006	0.006	0.40	0.50	0.58
300	TransientNeRF	20.14	21.99	20.53	0.270	0.210	0.177	0.012	0.007	0.014	0.26	0.39	0.36
	Proposed	21.60	23.53	23.79	0.223	0.170	0.151	0.005	0.005	0.012	0.42	0.49	0.55
150	TransientNeRF	19.13	21.06	19.95	0.277	0.206	0.196	0.016	0.015	0.018	0.21	0.33	0.28
	Proposed	21.23	23.56	25.33	0.222	0.171	0.146	0.007	0.005	0.004	0.41	0.50	0.60
50	TransientNeRF	18.80	19.63	19.31	0.274	0.194	0.190	0.020	0.020	0.018	0.19	0.26	0.24
	Proposed	19.14	22.33	25.34	0.280	0.180	0.148	0.017	0.011	0.007	0.34	0.44	0.59
10	TransientNeRF	16.40	16.36	15.40	0.246	0.259	0.272	0.037	0.041	0.090	0.08	0.08	0.02
	Proposed	22.40	23.05	25.47	0.242	0.208	0.173	0.008	0.008	0.007	0.35	0.36	0.45

method recovers smoother surfaces with fewer missing parts and fewer floating artifacts, especially in the case of training with two or three input viewpoints. The results highlight the benefit of training a neural surface representation with transient supervision.

Table 1 shows the quantitative performance of the proposed method against the baselines averaged across all scenes in the two, three, and five training viewpoint settings. We benchmark performance for varying measurement SNRs by training using transients with lower photon counts. To compute the metrics, we first scale the rendered transients so that the average photon count per occupied pixel (across all viewpoints) is the same across all photon levels. The

recovered meshes from the proposed method achieve significantly lower Chamfer Distance compared to the other methods and roughly  $5\times$  lower than the second-best method, TransientNeRF. The recovered depth maps are also more accurate than the baseline methods, and the image metrics are comparable. Finally, the transients rendered from our method are more similar to the ground truth based on the Transient IOU metric.

## 5.2. Captured Results

Figure 4 shows the recovered meshes for our method and the baselines on the *food*, *carving*, *cinema* scenes for two, three, and five training viewpoints respectively (see




Figure 5. Novel view synthesis for varying photon levels. We show the rendered novel view on the *cinema* scene, trained on five viewpoints with an average of 50, 300 photons. We also show peak-time visualizations [27], which show the full transient in a single visualization. Hue encodes the time of peak intensity, brightness is modulated by the maximum intensity, and each band corresponds to an isochrone, or wavefront of equal path length. We show transient plots (right) for the pixel indicated on the ground-truth image (blue dot).

Table 3. Ablation study on the *Lego* scene with 6000ppp. We present results without the sparsity loss (w/o  $\mathcal{L}_{\text{sparse}}$ ), without the space carving loss (w/o  $\mathcal{L}_{\text{sc}}$ ), without the weight variance regularizer (w/o  $\mathcal{L}_{\text{weight\_var}}$ ) and without the reflectivity loss (w/o  $\mathcal{L}_{\text{ref}}$ ).

Method	Chamfer Distance ↓			PSNR (dB) ↑			Transient IOU ↑		
	2 views	3 views	5 views	2 views	3 views	5 views	2 views	3 views	5 views
Ours w/o $\mathcal{L}_{\text{sc}}$	0.07	0.04	0.11	25.63	26.32	26.22	0.53	0.57	0.53
Ours w/o $\mathcal{L}_{\text{weight\_var}}$	0.28	0.37	1.04	21.58	22.66	23.81	0.48	0.53	0.60
Ours w/o $\mathcal{L}_{\text{sparse}}$	0.02	0.03	0.08	25.43	26.87	26.58	0.53	0.58	0.63
Ours w/o $\mathcal{L}_{\text{ref}}^*$	5.50	2.91	0.53	22.38	24.92	26.22	0.49	0.55	0.65
Ours	0.02	0.02	0.01	25.72	26.66	27.47	0.53	0.58	0.63

\*For this ablation we set the Eikonal loss to 0 and the weight variance loss to  $10^{-5}$ .

the supplement for qualitative results for all scenes). As in TransientNeRF [28], we mask out the background of the transients for the captured dataset. Similarly to the simulated dataset, the meshes recovered from the proposed method are more accurate than the baseline methods.

We report metrics averaged across six captured scenes in Table 2. Due to a lack of ground-truth meshes, we are not able to report Chamfer distances, and the depth L1 metric is calculated by comparing depth maps estimated from log-matched filtering of transients from held-out views. Our quantitative results show an improvement over TransientNeRF across all image and depth metrics. Overall, our method remains robust across various photon levels, whereas TransientNeRF’s image rendering quality deteriorates rapidly as the mean photon count decreases.

Figure 5 shows the results of training our method and TransientNeRF [28] on transients with an average of 50 and 300 photons per occupied pixel. At these photon levels, using TransientNeRF results in rendered transients that underestimate the measured intensities compared to our method.

### 5.3. Ablation Study

Table 3 evaluates the effect of ablating the regularization terms on the simulated *Lego* scene. The results indicate that all the added regularizers are crucial for performance; however,  $\mathcal{L}_{\text{ref}}$  and  $\mathcal{L}_{\text{weight\_var}}$  are the most essential. When ablating the reflectivity loss (row 4), we deviated from the default settings of  $\mathcal{L}_{\text{eik}}$  and  $\mathcal{L}_{\text{weight\_var}}$ , as indicated in the table, to improve convergence (see details in the supplement).

## 6. Discussion

We show that using a neural surface representation with transient supervision and additional regularization recovers high-fidelity surfaces in the few viewpoint setting with as few as 10 measured photons per pixel. Our work improves the practicality of transient-based reconstruction by reducing the required light levels, which could be especially useful for 3D imaging of photo-sensitive materials or long-range (and low-SNR) remote sensing.

In future work, we hope to address certain limitations of our method. For example, computational efficiency could potentially be improved by developing new time-resolved rendering frameworks based on Gaussian splatting [14]. Additionally, while the main focus of our work is scene reconstruction using the direct component of light, transients carry much richer scene information, including higher order bounces of light, future work could exploit indirect light transport to infer scene properties such as geometry, reflectance, and material characteristics.

## 7. Acknowledgments

DBL acknowledges support of the Canada Foundation for Innovation, the Ontario Research Fund, and NSERC under the RGPIN, RTI, and Alliance programs.



## References

- [1] Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. Törf: Time-of-flight radiance fields for dynamic scene view synthesis. *Proc. NeurIPS*, 34, 2021. 4
- [2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proc. CVPR*, 2020. 5
- [3] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proc. CVPR*, 2019. 2
- [4] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proc. CVPR*, 2022. 3
- [5] Daniele Faccio, Andreas Velten, and Gordon Wetzstein. Non-line-of-sight imaging. *Nat. Rev. Phys.*, 2(6):318–327, 2020. 2
- [6] William A Falcon. Pytorch lightning. *GitHub*, 3, 2019. 5
- [7] Yuki Fujimura, Takahiro Kushida, Takuya Funatomi, and Yasuhiro Mukaigawa. Nlos-neus: Non-line-of-sight neural implicit surface. *arXiv preprint arXiv:2303.12280*, 2023. 3
- [8] Ruiqi Gao\*, Aleksander Holynski\*, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul P. Srinivasan, Jonathan T. Barron, and Ben Poole\*. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv*, 2024. 2
- [9] Genevieve Garipey, Nikola Krstajić, Robert Henderson, Chunyong Li, Robert R Thomson, Gerald S Buller, Barmak Heshmat, Ramesh Raskar, Jonathan Leach, and Daniele Faccio. Single-photon sensitive light-in-flight imaging. *Nat. Commun.*, 6(1):6021, 2015. 2
- [10] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. 5
- [11] Yuan-Chen Guo. Instant neural surface reconstruction, 2022. <https://github.com/bennyguo/instant-nsr-pl>. 5
- [12] Sacha Jungerman and Mohit Gupta. Radiance fields from photons, 2024. 3
- [13] Sacha Jungerman, Atul Ingle, Yin Li, and Mohit Gupta. 3d scene inference from transient histograms. In *Proc. ECCV*, 2022. 2
- [14] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4), July 2023. 8
- [15] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *Proc. CVPR*, 2022. 3
- [16] Ahmed Kirmani, Dheera Venkatraman, Dongeek Shin, Andrea Colaço, Franco NC Wong, Jeffrey H Shapiro, and Vivek K Goyal. First-photon imaging. *Science*, 343(6166):58–61, 2014. 6
- [17] Tzofi Klinghoffer, Xiaoyu Xiang, Siddharth Somasundaram, Yuchen Fan, Christian Richardt, Ramesh Raskar, and Rakesh Ranjan. PlatoNeRF: 3D reconstruction in Plato’s cave via single-view two-bounce lidar. In *Proc. CVPR*, 2023. 2
- [18] Kiriakos Kutulakos and Steven Seitz. A theory of shape by space carving. *Int. J. of Comp. Vis.*, 38:199–218, 01 2000. 2
- [19] Min-Seop Kwak, Jiuhn Song, and Seungryong Kim. Geconerf: Few-shot neural radiance fields via geometric consistency. *arXiv preprint arXiv:2301.10941*, 2023. 3
- [20] Jongho Lee, Atul Ingle, Jenu V Chacko, Kevin W Eliceiri, and Mohit Gupta. Caspi: collaborative photon processing for active single-photon imaging. *Nat. Commun.*, 14(1):3158, 2023. 3
- [21] PA W Lewis and Gerald S Shedler. Simulation of nonhomogeneous poisson processes by thinning. *Naval research logistics quarterly*, 26(3):403–413, 1979. 5
- [22] Ruilong Li, Matthew Tancik, and Angjoo Kanazawa. Nerfacc: A general nerf acceleration toolbox. *arXiv preprint arXiv:2210.04847*, 2022. 5
- [23] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proc. CVPR*, 2023. 1, 2, 3, 4, 6, 7
- [24] David B Lindell, Matthew O’Toole, and Gordon Wetzstein. Towards transient imaging at interactive rates with single-photon detectors. In *Proc. ICCP*, 2018. 2
- [25] David B Lindell and Gordon Wetzstein. Three-dimensional imaging through scattering media based on confocal diffuse tomography. *Nat. Commun.*, 11(1):4517, 2020. 2
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [27] Anagh Malik, Noah Juravsky, Ryan Po, Gordon Wetzstein, Kiriakos N. Kutulakos, and David B. Lindell. Flying with photons: Rendering novel views of propagating light. In *Proc. ECCV*, 2024. 2, 8
- [28] Anagh Malik, Parsa Mirdehghan, Sotiris Nousias, Kiriakos N. Kutulakos, and David B. Lindell. Transient neural radiance fields for lidar view synthesis and 3d reconstruction. *Proc. NeurIPS*, 2023. 2, 3, 4, 5, 6, 7, 8
- [29] Nelson Max. Optical models for direct volume rendering. *IEEE Trans. Vis. Comput. Graph.*, 1(2):99–108, 1995. 4
- [30] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron. NeRF in the dark: High dynamic range view synthesis from noisy raw images. In *Proc. CVPR*, 2022. 4
- [31] 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. 2
- [32] Parsa Mirdehghan, Maxx Wu, Wenzheng Chen, David B Lindell, and Kiriakos N Kutulakos. Turbosl: Dense accurate and fast 3d by neural inverse structured light. *Proc. CVPR*, 2023. 3
- [33] Fangzhou Mu, Carter Sifferman, Sacha Jungerman, Yiquan Li, Mark Han, Michael Gleicher, Mohit Gupta, and Yin Li. Towards 3D vision with low-cost single-photon cameras. In *Proc. CVPR*, 2024. 3

- [34] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):1–15, 2022. 2, 3, 4, 5
- [35] 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 *Proc. CVPR*, 2022. 2, 3, 4, 6, 7
- [36] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proc. CVPR*, 2020. 3
- [37] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proc. ICCV*, 2021. 3
- [38] Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 2018. 2
- [39] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Proc. ECCV*, 2020. 3
- [40] Mohamad Qadri, Kevin Zhang, Akshay Hinduja, Michael Kaess, Adithya Pediredla, and Christopher A Metzler. Aoneus: A neural rendering framework for acoustic-optical sensor fusion. *arXiv preprint arXiv:2402.03309*, 2024. 3
- [41] Joshua Rapp and Vivek K Goyal. A few photons among many: Unmixing signal and noise for photon-efficient active imaging. *IEEE Trans. Comp. Imaging*, 3(3):445–459, 2017. 2, 3, 4
- [42] Joshua Rapp, Yanting Ma, Robin MA Dawson, and Vivek K Goyal. High-flux single-photon lidar. *Optica*, 8(1):30–39, 2021. 2
- [43] Joshua Rapp, Charles Saunders, Julián Tachella, John Murray-Bruce, Yoann Altmann, Jean-Yves Tournier, Stephen McLaughlin, Robin MA Dawson, Franco NC Wong, and Vivek K Goyal. Seeing around corners with edge-resolved transient imaging. *Nat. Commun.*, 11(1):5929, 2020. 2
- [44] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *Proc. CVPR*, 2022. 2, 7
- [45] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proc. CVPR*, 2022. 3
- [46] 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. *Proc. NeurIPS*, 33, 2020. 4
- [47] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Yifan Wang, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, Tomas Simon, Christian Theobalt, Matthias Niessner, Jonathan T. Barron, Gordon Wetzstein, Michael Zollhoefer, and Vladislav Golyanik. Advances in neural rendering, 2022. 1
- [48] Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mounqi Bawendi, Diego Gutierrez, and Ramesh Raskar. Femtophotography: capturing and visualizing the propagation of light. *ACM Trans. Graph.*, 32(4):1–8, 2013. 2
- [49] Guangcong Wang, Zhaoxi Chen, Chen Change Loy, and Ziwei Liu. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. In *Proc. ICCV*, 2023. 2, 3
- [50] 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. *arXiv preprint arXiv:2106.10689*, 2021. 2, 3, 4
- [51] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proc. ICCV*, 2023. 3
- [52] Zian Wang, Tianchang Shen, Merlin Nimier-David, Nicholas Sharp, Jun Gao, Alexander Keller, Sanja Fidler, Thomas Müller, and Zan Gojcic. Adaptive shells for efficient neural radiance field rendering. *arXiv preprint arXiv:2311.10091*, 2023. 4
- [53] Haoyu Wu, Alexandros Graikos, and Dimitris Samaras. S-volsdf: Sparse multi-view stereo regularization of neural implicit surfaces. In *Proc. ICCV*, 2023. 3
- [54] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with diffusion priors. *arXiv preprint arXiv:2312.02981*, 2023. 2
- [55] Shumian Xin, Sotiris Nousias, Kiriakos N Kutulakos, Aswin C Sankaranarayanan, Srinivasa G Narasimhan, and Ioannis Gkioulekas. A theory of fermat paths for non-line-of-sight shape reconstruction. In *Proc. CVPR*, 2019. 2
- [56] 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 *Proc. ECCV*, 2022. 3
- [57] Dongyu Yan, Xiaoyang Lyu, Jieqi Shi, and Yi Lin. Efficient implicit neural reconstruction using lidar. In *Proc. ICRA*. IEEE, 2023. 3
- [58] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proc. CVPR*, 2023. 3
- [59] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Proc. NeurIPS*, 34, 2021. 3
- [60] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Proc. NeurIPS*, 33, 2020. 3

- [61] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proc. CVPR*, 2021. [3](#)
- [62] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Proc. NeurIPS*, 35, 2022. [2](#), [3](#), [6](#)
- [63] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*, 2018. [6](#)